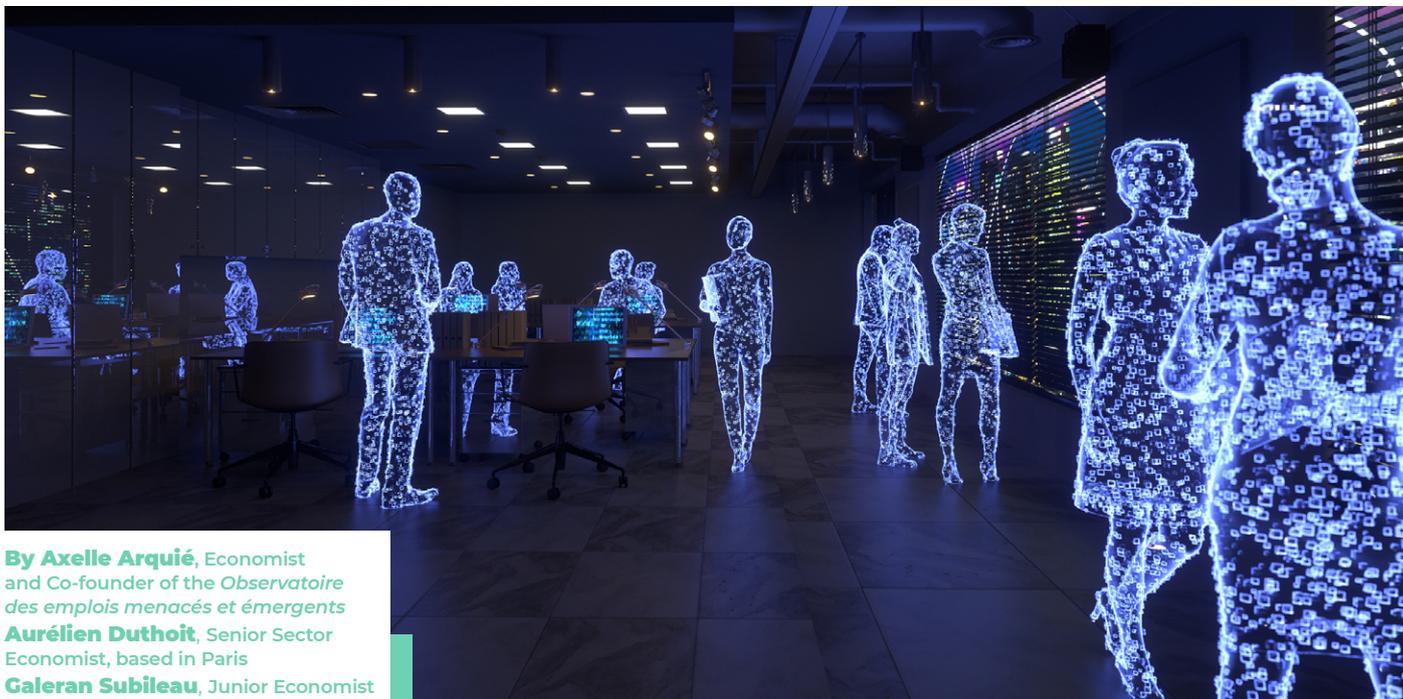


FOCUS



By **Axelle Arquié**, Economist and Co-founder of the *Observatoire des emplois menacés et émergents*
Aurélien Duthoit, Senior Sector Economist, based in Paris
Galeran Subileau, Junior Economist

The Next Automation Frontier: A Scenario Map of AI Labour Exposure

EXECUTIVE SUMMARY

Three years after the release of ChatGPT, AI seems to be everywhere in corporate discourse, and yet still almost nowhere in aggregate labour-market statistics. The first effects are however beginning to appear, mainly at the margins, among the younger, entry-level workers of the most exposed occupations and industries. This apparent paradox may simply reflect timing, with the current phase of AI adoption still largely focused on tools that assist work without reorganizing it, and the current evidence capturing early signals rather than the main shock itself. We believe the rise of agentic AI systems, that is, systems capable of planning, coordinating and executing multi-step workflows, could accelerate and amplify the disruptive impact of AI on labour markets by moving from assistance on isolated tasks to the execution of entire workflows.

To move beyond these early signals, the Observatoire des Emplois Menacés et Émergents (OEM) developed a task-based framework to measure the technical exposure of occupations to successive phases of AI development. While many existing studies rely either on expert judgment or one-off assessments produced by large language models, the OEM approach breaks complex tasks into elementary actions, applies explicit and reproducible scoring rules and projects exposures along successive technological phases rather than a single point in time. The result is a granular, reproducible and prospective way to assess how advances in AI could reshape occupations, industries and labour markets.

Together with the OEM, Coface contributed to extending this framework by developing a weighting method for tasks, refining further both the prospective automation scenarios and the notation rules elaborated by the OEM, and by broadening the empirical scope of the analysis. This framework is deliberately gross and supply-side: it measures technical exposure to automation, not net job destruction. It abstracts from demand dynamics, from potential new task creation and from frictions that may slow deployment.





Our results point to a different automation pattern from earlier technological waves such as industrial robotics of the 1990s and 2000s and computer software. Rather than concentrating on routine, middle-skill work, the automation potential of AI in its agentic form is highest in occupations made up largely of cognitive, non-routine tasks often performed by more highly educated and better-paid workers. This is especially visible in engineering & computational occupations (29% of task content at risk), legal & financial, creative & content occupations (27%), and management & administrative roles (24%). At the national level, automation-exposed task content ranges from around 12% in Turkey to close to 20% in the UK, with countries such as France (16%), Germany (17%) and the United States (17%) lying closer to the middle of the distribution. Because the occupations most exposed are among the most central to income formation, tax revenues and value creation in advanced economies, the consequences are likely to extend well beyond employment alone. AI could create a significant redistribution of income away from labour and toward capital. If employment losses in exposed occupations are not fully offset by new job creation, wage income would come under pressure while productivity gains would accrue first to companies deploying AI. Part of these gains may then be captured further upstream by the technology companies controlling the core infrastructure of the AI value chain. This raises, for many countries with public finances heavily dependent on labour taxation, a double risk of erosion of the domestic wage tax base and partial leakage of capital income toward foreign countries where AI profits are concentrated.

The implications also extend beyond employment and include a rethinking of education systems as AI begins to test the labour-market value of qualifications and diplomas and shift the premium toward judgment, adaptability and AI-complementary skills. They also include stronger supply-chain and geopolitical dependencies as AI becomes a critical input to production while its key assets (semiconductors, data centers, models) remain highly concentrated across a small number of companies and countries, thereby increasing exposure to external shocks, regulatory frictions and strategic vulnerabilities.

While these results should be read with caution since the translation from the gross technical exposure of tasks to the net employment effects is neither immediate nor mechanical, the underlying questions would remain even under a more gradual or limited disruption scenario. The framework should therefore be understood not as a forecast, but as a structured map of where labour-market, macroeconomic and strategic pressures are most likely to emerge as agentic AI diffuses.

SOME OF THE THINGS YOU'LL LEARN...

13% of occupations
will have more than
30% of their tasks
automatable by agentic
AI in coming years

p.6

14 to 19%, depending
on countries,
of the workforce task
content could potentially
be automated
in coming years

p.8

Jobs are more
at risk in Luxembourg,
UK, Switzerland,
Netherlands, Norway
and Denmark

p.10

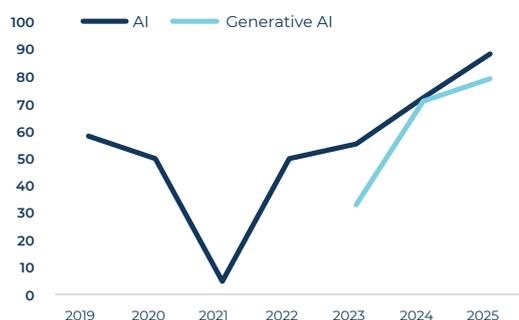
20-25% of task
content is at risk
among occupations
held by the highest
10% paid of workers

p.11

Three years on since the launch of ChatGPT, AI is everywhere but in aggregate labour statistics

Generative artificial intelligence has fast-forwarded the use of AI in companies. The public release of ChatGPT in late 2022 marked an inflection point in the corporate adoption of artificial intelligence (AI). The share of companies reporting the use of AI in at least one business function rose from 55% to 88% between 2022 and 2025, while the adoption of generative AI alone increased from 33% to 79% (Chart 1).

Chart 1 - Organizations using AI in business functions
(% of respondents)



Sources: The state of AI in 2025, McKinsey, November 2025, Coface, OEM. The survey was conducted over a sample of 1,993 global participants.

Evidence of an impact at the aggregate level, however, has not yet materialized: the clearest effects appear concentrated among younger workers in highly exposed occupations and firms¹. In the United States, recent studies find weaker entry-level hiring in AI-exposed firms and occupations, with employment declines concentrated among workers at the start of their careers rather than among senior staff^{2,3}. Early wage evidence points in the same direction: pay has come under pressure for junior and mid-level roles in exposed firms, while senior compensation has remained stable or continued to rise⁴.

In France, a similar pattern can be observed⁵. Employment adjustments in AI-exposed sectors, such as information-technology and information services, publishing, and head offices and management consulting, are concentrated among young workers (15–29 years, excluding apprentices), with noticeable declines in entry-level hiring and early-career positions since late 2023, even as value added remains correctly oriented on an upward trend.

The labour market effects documented so far likely reflect only the first and most limited phase of AI deployment. The real disruption is unlikely to come from large language models (LLMs) used in isolation, which leave production processes largely unchanged, but from agentic systems capable of autonomous, multi-step action across entire workflows. Agentic artificial intelligence, however, is still only beginning to be deployed at scale⁶. What current studies may be capturing, in other words, is less the main shock than its early signal.

To anticipate that possibility, the Observatoire des Emplois Menacés et Émergents (OEM) develops a prospective task-based framework for measuring

the technical exposure of occupations to successive phases of AI development. This approach extends the standard occupation-by-task method by making it more granular, reproducible and forward-looking.

Complex tasks are decomposed into elementary actions, each defined by a verb, a direct object, and a contextual element, such as "summarize / text / confidential" or "welcome / customer / restaurant". Each elementary action is then scored according to semantic rules designed to capture AI capabilities and limitations (see Box 1). Tasks are scored based on the elementary actions they contain; occupations, in turn, receive an aggregate score based on the tasks they require.

Together with the OEM, Coface contributed to extending this framework by developing a weighting method for tasks based on their importance, relevance and frequency, refining further both the prospective automation scenarios and the notation rules elaborated by the OEM, and by broadening the empirical scope of the analysis, first to twelve individual countries using granular data enriched with sociodemographic indicators such as income, education and sector, and then to around thirty additional European countries using slightly less granular but standardized European data.

This framework is deliberately gross and supply-side: it measures technical exposure to automation, not net job destruction. It abstracts from demand dynamics, from the organizational, economic and regulatory frictions that may slow deployment and from potential new task creation. The usual offsetting argument should, however, be taken with caution: unlike past technologies, AI has no fixed ceiling, so the new tasks it creates may themselves be automatable in turn.

Applied to 923 occupations and twelve countries, the framework suggests that the next automation wave may differ sharply from previous ones: exposure is concentrated not in routine middle-skill work, but in many cognitive, high-skill and highly paid occupations, while physically anchored and highly variable forms of work remain comparatively protected.

The impact of AI on labour markets is likely to be (very) different from past technological disruptions

AI's capacity to automate cognitive, non-routine tasks is a game changer. Prior waves of automation targeted mainly repetitive tasks: the acceleration of industrial robotisation from the 1980s onwards displaced manual labour in factories, while rule-based software and RPA (Robotic Process Automation) in the 2000s automated routine cognitive tasks in offices. In both cases, machines could only react to explicitly programmed instructions: every situation had a predetermined response. Only tasks with a sufficient degree of routine could effectively be automated.

The shift to neural AI changed this fundamentally. Today's models rely on algorithms that learn autonomously from data (machine learning), enabling them to handle a form of bounded uncertainty. Generative AI adds a mastery of natural language, which, combined with a capacity for improvisation rooted in machine learning, has suddenly broadened the scope of what can be automated, including tasks that are non-repetitive.

1 - Martha Gimbel, Molly Kinder, Joshua Kendall and Maddie Lee, Evaluating the Impact of AI on the Labor Market: Current State of Affairs, The Budget Lab at Yale, October 1, 2025
 2 - Amir Hosseini and John Lichtinger, Generative AI as Seniority-Biased Technological Change, SSRN, July 24, 2025
 3 - Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence, November 2025, revised in February 2026.
 4 - Xavier Giné et al., Quasi-Experimental Evidence From Occupational Exposure, SSRN, December 12, 2025. The paper reports lower starting wages for junior and mid-level roles in more exposed firms, while senior wages remain stable or increase.
 5 - Insee, *Note de Conjoncture*, March 2026
 6 - McKinsey's The State of AI in 2025: Agents, Innovation, and Transformation.

BOX 1 - THE OEM METHODOLOGY

**1 - FOUNDATIONS: TASK DECOMPOSITION
INTO GENERIC ACTIONS (SEE ANNEX)**

The methodology developed by the OEM (*Observatoire des Emplois Menacés et Émergents*) extends the reference academic approach for assessing the impact of technological change on occupations, adding three dimensions absent from most existing studies: granularity, reproducibility and temporality.

This method, initiated by Autor, Levy & Murnane⁷, consists of decomposing each occupation into tasks, then evaluating the automation potential of each one. The overall score for an occupation is derived from the aggregation of the scores of its constituent tasks. The international O*NET database, the most widely used by researchers, details tasks for 923 occupations whose automation potential is then rated. Most studies on the impact of AI base this assessment either on evaluations by "domain experts" or on a rating by a LLM, both conducted at a single point in time.

These methods suffer from three limitations. First, lack of granularity: O*NET tasks group together complex and heterogeneous actions, so that evaluating them globally conceals critical nuances and may introduce some bias. Second, lack of reproducibility and auditability: the way an expert or LLM arrives at a rating remains a non-auditable black box: experts may revise their assessments, and LLMs are highly context-dependent, producing different scores from marginally different prompts. Third, absence of dynamics: occupations are assessed at a single point in time, with no forward-looking view of technological trajectories. An occupation with low exposure today under isolated LLMs may become highly vulnerable in two to three years with mature agentic AI, a blind spot that prevents strategic anticipation.

2 - THE THREE-LEVEL SCORING ARCHITECTURE

Each of these limitations is addressed while retaining the fundamental principle of task decomposition. O*NET tasks are broken down into simple generic elementary actions structured as a triplet (verb, direct object, contextual element: for instance, "summarize / text / confidential") allowing each component of a complex task to be evaluated separately and revealing nuances invisible to conventional approaches. A fully reproducible rating system is then applied: unlike expert opinion or LLM judgement, the same occupation evaluated twice produces identical scores. Finally, prospective scores are produced for each technological phase rather than at a single point in time.

**A FULLY REPRODUCIBLE RATING SYSTEM IS THEN
APPLIED ACROSS THREE LEVELS:****Level 1: Verb**

Each verb is first classified using the semantic categories from WordNet (an English lexical database, developed by Princeton, that organizes words into semantic categories and hierarchical synonym sets, such as *verb.cognition*, *verb.communication*, *verb.creation*), as well as hyperonym hierarchies. Verbs associated with simple cognitive operations, such as calculate, extract, summarize, translate, receive positive base scores reflecting strong AI performance on routine informational tasks. By contrast, verbs implying higher-order reasoning, such as infer, design, solve, receive more moderate scores, except in the SuperHuman scenario where causal reasoning would be achieved, because they involve multi-step logical chains AI models still handle unreliably.

Level 2: Object

The direct object determines the domain and nature of what is being acted upon (see matrix in Chart 1). Object categories apply either a boost or a penalty to the verb base score. For example, objects falling in the "data", "text" or "numeric" categories receive strong positive boosts, while objects in "sensitive", "strategic" categories receive heavy penalties reflecting AI limitations. Those penalties decrease as AI capabilities are supposed to be advancing across the scenarios studies.

Level 3: Contextual element

The contextual qualifier aims at capturing the situational context that can radically change the automation potential of an otherwise identical action. For instance, the contextual element "digital" boosts the triplet, while "confidential" triggers a penalty across all phases, reflecting both regulatory constraints and the unreliability of AI in high-stakes sensitive processing.

**3 - SEMANTIC RULES CAPTURING
AI LIMITATIONS: BOOST AND PENALTIES**

Penalties are applied when an action requires capabilities that AI systems do not reliably possess in each scenario considered (depending on the assumptions made). Rules apply boosts or penalties to whole categories (verb, object, contextual element, or pairs of each of these categories). They are designed to encode theoretically grounded AI limitations rooted in assumptions for each of the technology scenarios described below.

In addition to general categories from WordNet, semantic categories have been constructed manually to capture AI limitations and capabilities that standard lexical resources are ill-equipped to reflect. A selection of purely illustrative examples is provided below.

Example 1: Error-handling and exception management

This group is designed to capture one of the most fundamental limit to AI automation currently: the difference between detecting statistically likely errors in structured data, which AI handles relatively well, and reasoning about novel, context-dependent exceptions in open-ended situations. The underlying limitation is that error management in professional settings is rarely a pattern-matching problem; it requires understanding why something went wrong, not merely that it deviates from a norm.

Accordingly, actions in this group attract a significant penalty in Phase 1, where isolated LLMs have no access to external data and no capacity for autonomous cross-checking. The penalty attenuates in Phase 2, as multi-agent architectures introduce cross-validation mechanisms that partially mimic error-detection routines, but it never disappears, because the root limitation (absence of causal understanding) persists across most phases modeled. Furthermore, multi-agent architectures introduce an additional risk: errors generated at one node can propagate and amplify across the system, making compounding failures a structural concern rather than a marginal one.

**Example 2: Planning and improvisation
(orchestrate, improvise)**

These two groups capture related but conceptually distinct limitations, both rooted in the same underlying deficiency: AI models trained by next-token prediction do not plan; they predict.

The orchestrate group, covering actions that imply coordinating multiple agents, sub-tasks, or processes toward a compound goal, is heavily penalized in Phase 1, precisely because this coordination capacity is what isolated LLMs categorically lack. The penalty is strongly attenuated in Phase 2: is the most structurally significant phase-transition signal in the entire system: it encodes the defining architectural innovation of agentic AI, which is precisely its (expected) ability to decompose goals, delegate sub-tasks to specialized agents, and orchestrate their outputs. An action that is essentially impossible for a standalone LLM becomes a relative strength of a well-designed multi-agent architecture.

The improvise group captures a different and more persistent limitation: the inability to generate genuinely novel responses to situations absent from training data. Where orchestration is a planning problem amenable to architectural solutions, improvisation is a generalization problem that probabilistic models structurally struggle with. Accordingly, this penalty attenuates from Phase 1 to Phase 2 as agentic systems gain flexibility.

Example 3: Relational objects

Interpersonal mediation, conflict management, employee wellbeing carry a strong penalty in Phase 1, reflecting the absence of genuine empathic understanding in language models. This penalty attenuates substantially in Phase 2, capturing the fact that agentic systems can handle scripted relational exchanges with increasing competence. A residual penalty remains, however, encoding the view that open-ended empathic interaction, not following a script but genuinely adapting to emotional context, remains beyond reliable AI reach.

Example 4: "Care" objects

Patient care, therapeutic interaction, clinical management receive a near-total penalty across all phases. This is the only category where phase progression is moderated: the combination of physical embodiment requirements, irreversible real-world consequences and thus responsibility and regulatory issues, creates a ceiling that architectural improvements alone cannot raise.

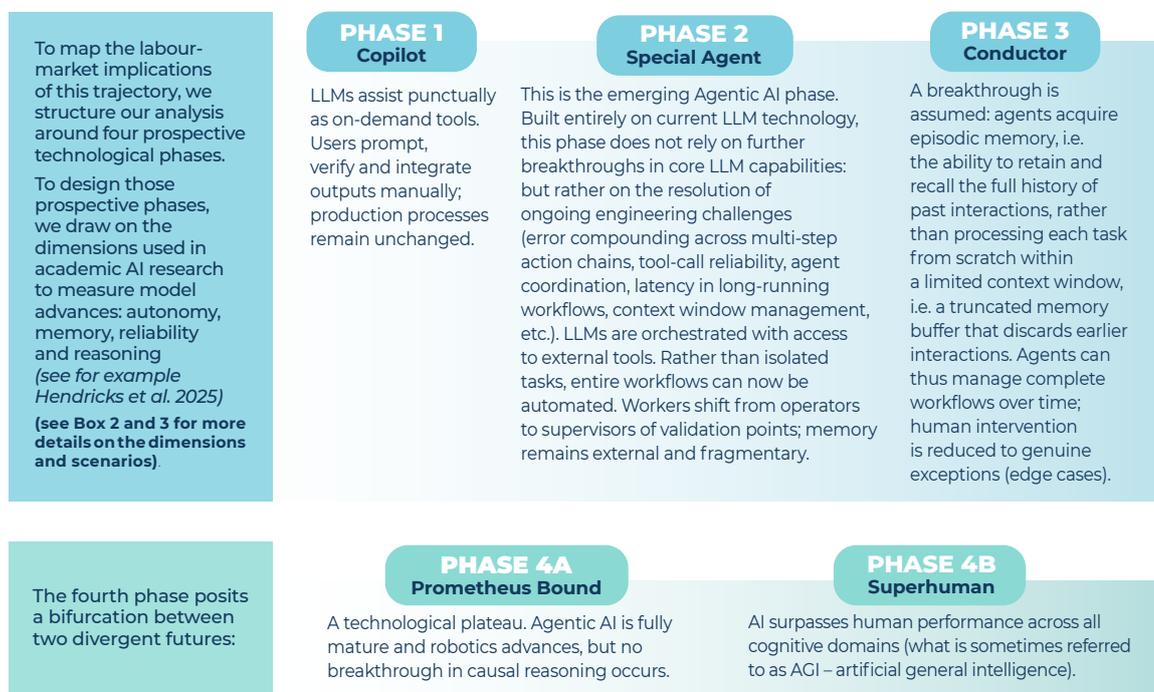
A model based on machine learning techniques does not need to be told the rules: it infers them from data, provided past cases are sufficiently abundant and structured. Yet this improvisation remains bounded. Large Language Models (LLM) operate primarily through statistical pattern-matching. Tasks requiring genuine causal reasoning, adaptation to truly novel situations absent from training data (“out-of-distribution problems”), or sustained coordination across evolving and ambiguous contexts remain beyond their reliable reach - for now. It is this residual frontier, and its gradual erosion across successive phases of agentic AI development, that the following analysis maps.

Artificial intelligence also stands out from past technologies because of the speed at which new AI capabilities are reaching the market. This acceleration reflects several reinforcing factors:

- **Digital economics:** once AI infrastructure and AI models are brought to the market, AI services can be

replicated and duplicated at near-zero marginal cost for technology providers, allowing them to expand aggressively and improve service quality at software speed rather than through slower capital replacement cycles more typical of hardware innovations (robotics, in particular).

- **Low-friction interfaces:** chat-based tools for non-technical employees and application programming interfaces (APIs) for developers reduce adoption barriers and allow for greater experimentation and faster integration into business processes.
- **Industrial-scale distribution:** AI technologies are largely funded by some of the world’s most profitable companies engaged in fierce competition for market dominance. Adoption is also facilitated by those companies’ ability to incorporate AI into their existing services (cloud services from Microsoft, Alphabet, Amazon or Oracle, CRM software from Salesforce, ERPs from SAP, developer toolchains like GitHub, etc.).



While the agentic AI considered for these scenarios (phases 2 and 3) does not require any true disruptive innovation, **agentic AI systems face major limitations that undermine their deployment in real-world environments.**

The first stems from the fundamental opacity of these systems: LLM decisions are not based on readable logical rules, but on a matrix of billions of interconnected numerical coefficients, making robust technical auditability extremely difficult with current methods. This lack of mechanistic interpretability makes it legally complex to delegate critical tasks.

There are also serious engineering bottlenecks: within agentic architectures, errors propagate in cascades through agent chains, real-time monitoring of inter-agent flows remains difficult, resource contention causes sharp performance degradations, and communication delays accumulate desynchronizations that are detrimental to operational responsiveness.

These systems also suffer from limited causal reasoning capabilities, making them fragile when confronted with novel situations that require solving out-of-distribution problems. In this technological paradigm, errors and hallucinations are likely to remain a major issue.

The economic calculus further reinforces these technical reservations. Multi-agent architectures multiply model inference costs, generating potentially prohibitive inference costs, to which hidden costs must be added (maintenance, human supervision, cybersecurity, model updates) and which could render automation unprofitable as soon as the architecture grows in complexity.

Mapping the impact of agentic AI in the coming years: the risk of a major automation shock

We assess the impact of AI automation potential from occupations to occupational families, jobs, industries through to countries and regions. We examine exposure to AI automation across several levels of aggregation: occupations, occupational families, industries, national labour markets and finally local labour markets. Unless otherwise stated, the following results originate from our task-based automation framework covering 923 occupations under the Special Agent scenario.

Drawing on our analysis of occupations by task, we suggest the following matrix as a simplified but robust heuristic to assess the labour-market disruption potential of agentic AI (Chart 2). It classifies tasks by the substrate they operate on (matter, people, or data) and the substrate in which they deliver value, using the resulting 3x3 mapping to rank the likely automation potential of agentic AI.

We find that exposure to automation is lowest when work operates from matter to matter, and highest when it operates from data to data. **In the first case, value is created through direct intervention in the physical world** (building, repairing, maintaining, handling) where automation remains limited by the need for dexterity, mobility and adaptation to heterogeneous, unstable and often unpredictable environments where large language models, as they are today, are facing clear limitations. **In the second, value is produced entirely within the world of information**, where agentic AI already encounters its most favourable terrain: digital inputs are machine-readable, outputs are easily standardized and tasks such as processing, transcribing, drafting or restructuring can be executed at scale with relatively little friction. **Tasks involving people occupy an intermediate position. They are often more exposed than physical work because they rely heavily on language and information exchange, yet less exposed than purely data-based work because they also depend on trust, care, persuasion, authority and other relational dimensions that remain harder to codify and delegate reliably.** The overall pattern is one of declining resistance as work moves away from matter and toward information.

Starting at the occupation level, agentic AI puts one in eight occupations on course for deep transformation. Existing studies have focused on thresholds identifying risks of job destruction, with scores between 30% and 70% interpreted as medium employment risk, and above 70% as high risk of displacement⁸. Revisiting the risk of automation. Economics Letters). No established threshold exists for structural transformation short of displacement. We propose 30% of tasks susceptible to automation as a threshold for significant structural transformation, not destruction, of an occupation. We therefore consider that prospective scores higher than 30% correspond to occupations likely to be profoundly transformed by the deployment of Agentic AI systems. Among the 923 occupations studied, 120 (13%) show an exposure to automation above 30% to the second automation phase "Special Agent" (Chart 3).

While highly automation-exposed occupations cluster in sectors with a high cognitive and informational intensity, the depth of impact varies significantly. "Architecture & Engineering", "Computer & Math", "Office & Administrative Support" and "Business and Finance" have the highest count of exposed occupations. Yet the highest average automation-exposure among occupations scoring higher than 30% is found in "Sales" and "Office & Administrative Support", surpassing "Architecture & Engineering", "Computer & Math".

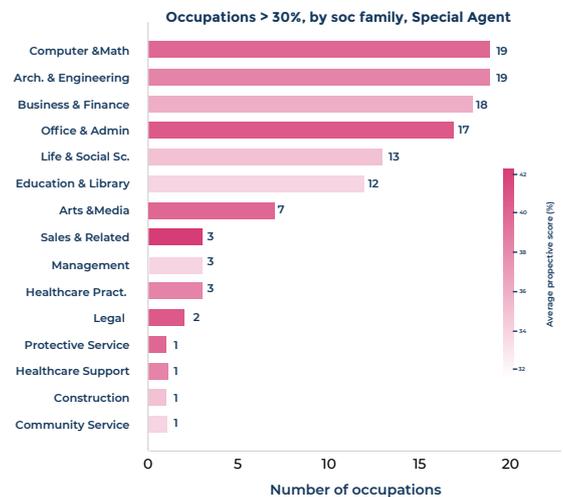
Chart 2 - Mapping agentic AI automation potential across task substrates (%)

Value / Output substrate \ Work / Input substrate	Matter	Person	Data
Matter	"Handling" Lowest Construction, maintenance, repair, craft work	"Assisting" Low Bringing comfort, reassurance security, authority, Physical assistance to people	"Sensing" Average to high Visual quality insurance, inventory counting, medical imaging...
Person	"Caring/Serving" Low Nursing care, childcare, hospital-ity service, skilled trades...	"Influencing" Medium Negotiating, convincing, leading, motivating...	"Transcribing" High Call center trans-cription, meeting notes, form filling, claims intake...
Data	"Controlling" Average to high Physical automation: vehicles, robots, production lines...	"Explaining" High Customer replies, explainers, tutoring content...	"Processing" Highest Information drafting, summarizing, classifying, restructuring...

Sources: OEM, Coface

How to read: the matrix classifies tasks according to the type of input they operate on (matter, people or data) and the type of output through which they create value. Tasks running from matter to matter - "Handling", for example construction, repair or craft work - involve direct intervention in the physical world and therefore remain the least exposed to agentic AI. At the other end of the matrix, tasks running from data to data ("Processing"), such as drafting, summarizing or restructuring information, are the most exposed because both their inputs and outputs are already digital and codifiable. Intermediate cases, especially those involving people, tend to fall in between: they are often more exposed than physical work because they rely heavily on language and information exchange ("Explaining", "Transcribing") but less exposed than purely data-based work because they also depend on trust, judgment, persuasion or care ("Caring", "Serving", "Influencing").

Chart 3 - Occupations with ≥ 30% of tasks automatable by occupational family, Special Agent Scenario



Sources: OEM, Coface

How to read: Each bar shows the number of occupations within a SOC family whose task exposure to automation exceeds 30%. For instance, in the "Legal" family, two occupations cross this threshold. Bar color reflects the family-level average share of tasks exposed to automation: darker shades indicate higher average exposure.

Some characteristics of the underlying task structure of these occupations stand out as vulnerable to agentic AI. The first dimension of tasks operates on digital raw material (documents, emails, forms, databases, compliance files) whose inputs are already encoded in formats that AI agents can directly manipulate (see Chart 2). The second involves information transformation and manipulation tasks: drafting, summarizing, translating, transcribing.

The third covers tasks that require human judgment, but which in fact can be documented by sufficiently structured and abundant data: compliance checking, fraud detection, underwriting, regulatory reporting. Machine learning does not need explicit rules: it infers decision patterns from historical cases, provided the training data is rich enough. AI performs reliably well when the variability of situations is real but bounded: what we call low **horizontal complexity**. The underlying logic is statistical: the higher the diversity of cases an occupation must handle, the lower the probability that any given real-world instance falls within the coverage of available training data, and the greater the risk of out-of-distribution failure. Where horizontal complexity is low, past cases are more likely to be informative about future ones.

The fourth vulnerability factor is **vertical complexity**. AI systems acquire formalized technical expertise, i.e. domain knowledge at scale, meaning that depth of specialization no longer confers protection. Among the 120 most automation-exposed occupations, 56 belong to high-skill occupational families (Computer & Mathematical, Architecture & Engineering, Life & Social Sciences, Legal and Healthcare Practitioners). **Expertise and formal knowledge alone no longer shield an occupation.**

Three distinct profiles emerge from this data. The first comprises some very specific occupations in sales, especially telemarketing tasks, and administrative roles. Here, Agentic AI transcends simple rule-based RPA to handle complex, semi-structured workflows and telephone client interactions with a higher autonomy.

The second includes knowledge and creative production occupations where AI agents can automate a significant share of tasks: this includes both analytical work (statistics, data science, business analysis), content production (design, animation, translation) and document-intensive legal and financial work. AI Agents demonstrate strong capabilities in generating final outputs (whether drafting contracts, analyzing regulatory frameworks, or producing graphic designs) challenging the assumption that complex cognitive and creative tasks are immune to automation.

The third profile covers engineering and quantitative science occupations where coding, modelling and simulations tasks are heavily exposed.

One in seven occupations remains shielded from agentic AI automation. Among the 923 occupations studied, 136 (15%) show an exposure below 5% to the "Special Agent" phase, suggesting they will remain largely untouched by this wave of automation (Chart 4). The pattern is clear: low-exposure roles are overwhelmingly concentrated in physically and sensorially anchored work. Production, Construction & Extraction and Installation & Repair alone account for 84 of the 136 occupations, followed by Transportation, Food Preparation, Healthcare Practitioners and Farming & Forestry.

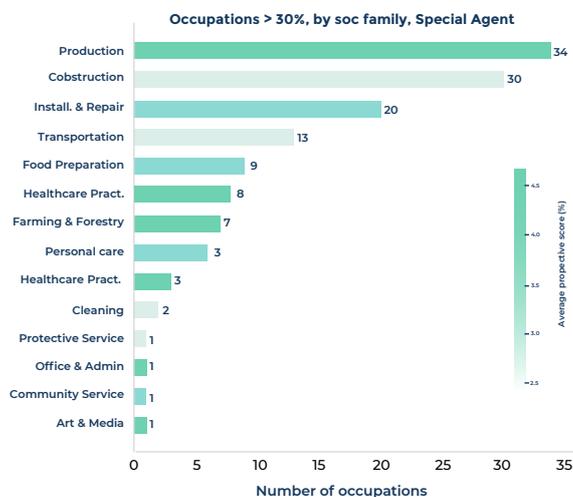
The primary shield for these occupations is **embodied complexity**. Tasks relying on precise physical manipulation, spatial navigation, or real-time sensory feedback, such as those performed by metal pourers (0.1%) or painting and coating workers (0.25%), obviously remain beyond the reach of agentic AI and basic robotics.

Low automation scores for Transportation Security Screeners (Protective Service) and Athletes (Arts & Media) reflect the physical nature of their core tasks. Similarly, Mental Health Social Workers rely on a therapeutic presence that resists digitization. Within 'Office Support,' roles like Mail Clerks are protected by their focus on physical handling rather than data processing. Ultimately, it is the bodily and relational nature of work that determines resistance to agentic automation.

A second protective dimension is **horizontal complexity** and judgment under genuine uncertainty. However, it does not explain the lowest-exposure occupations, which are almost exclusively physical. It instead accounts for why a large set of cognitively intensive occupations cluster in the intermediate zone between 5% and 30%, remaining partially shielded despite being information-rich.

Community and social service occupations illustrate this protective dimension: mental health counselors (8.2%), child and family social workers (9.1%), or healthcare social workers (9.4%) score well below the high-exposure threshold. Their work involves continuous adaptation to contextually unique human situations: relational attunement and improvised judgment needed cannot be reduced to pattern-matching over historical data. It is precisely the irreducible variability of the cases a worker faces, rather than the cognitive depth of the work or formalized knowledge required (**vertical complexity**), that provides protection here.

Chart 4 - Occupations with < 5% of tasks automatable by occupational family, Special Agent Scenario



Sources: OEM, Coface

What share of tasks is exposed and among how many employees?

While this occupation-level view highlights the magnitude of the AI automation wave across a wide range of jobs and industries, it does not yet indicate the volume of task content actually exposed to automation across the workforce. To move from the automation-exposure of occupations to the exposure of tasks at workforce scale, we weight each occupation's exposure by its employment in national economies. We then aggregate the 923 occupations into eight employment-weighted occupational families⁹ to help identify broad patterns. To limit country-specific composition effects, we report the median value across the twelve countries we analyse.

9 - The details of the occupational families can be found in the appendix

Management, creative, legal, financial, engineering, and IT occupations, whose tasks largely fall into the “Explaining”, “Processing” and “Influencing” categories of our mapping, all show more than 25% of tasks at risk. By contrast, in-person service occupations (e.g. cleaning and food preparation) and skilled trades and industrial production occupations (manufacturing, transport, installation, and maintenance) remain below the 10% threshold (Chart 5) - their tasks mostly fall into the “Handling” and “Caring & Serving” categories of our mapping.

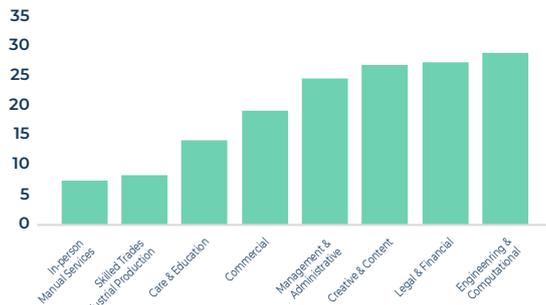
Because industries employ different mixes of occupations, occupational exposure naturally translates into sectoral exposure. Using sector-level data for France, we find that sectors employing a high share of cognitively intensive occupations stand out clearly. Information and communication, finance and insurance, and public administration all show automation-exposure rates close to 25% (Chart 6). The results also suggest that the traditional divide between industry and services is too simplistic. Although service activities are, on average, more exposed, dispersion across sectors is substantial. Several large service sectors including accommodation and food services (8%), health and social work (13%), and transport (15%) display lower automation-exposure than manufacturing as a whole (16%). More detailed data also indicate that cognitively intensive industries with a strong R&D component, such as pharmaceuticals (18%)

and electronics (20%), are particularly exposed¹⁰. Overall, these patterns support our view that a task-based approach provides a more granular and informative picture of future AI exposure than broad sectoral or occupational classifications.

Sectoral exposure in turn aggregates to the level of national labour markets, depending on each country’s economic structure. Reflecting differences in employment composition and economic structure, exposure varies markedly across countries (Chart 7). The Netherlands and the UK consistently rank among the most automation-exposed, followed by Poland, Denmark, Australia and Canada. Germany, Sweden, the United States, France and Finland appear comparatively less exposed, while Japan stands out as a clear outlier.

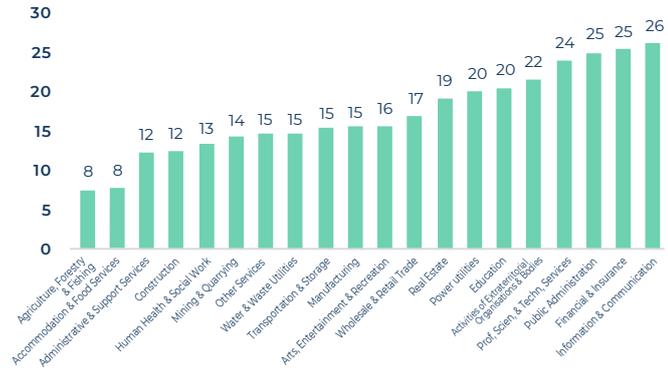
Our granular approach also shows that just three occupational families account for more than 60% of all task content at risk across the workforce. The following chart compares two measures: the average exposure of each occupational family to AI risk (that is, the share of task content at risk within that family, as in Chart 4), and the average contribution of each family to total task contents at risk across the workforce at the macroeconomic level, reflecting both exposure and employment volume.

Chart 5 - Task content at risk across the workforce by occupational family, Special Agent (France, %)



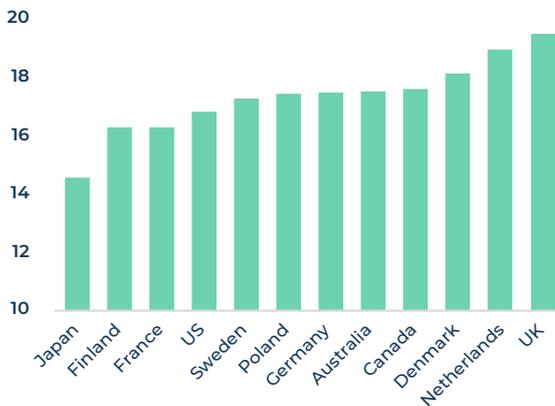
Sources: OEM, Coface

Chart 6 - Task content at risk across the workforce by sector, Special Agent Scenario by sector (France, %)



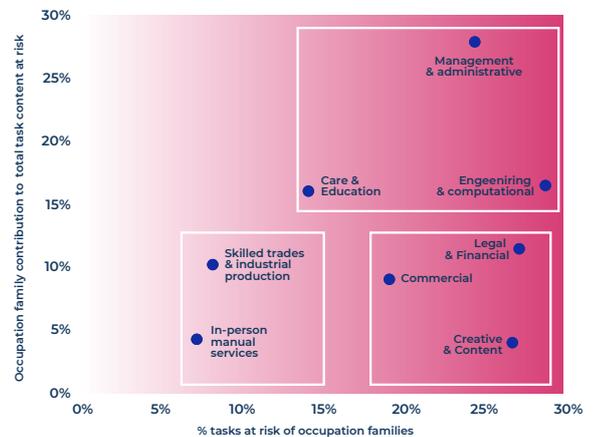
Sources: OEM, Coface

Chart 7 - Task content at risk across the workforce by country, Special Agent Scenario (%)



Sources: OEM, Coface

Chart 8 - Automation exposure of occupational families and their contribution to total task content at risk, Special Agent scenario (%)



Sources: OEM, Coface.

How to read: in our Special Agent scenario, 24% of all tasks belonging to the “Management and Administrative” occupational family are at risk. These tasks represent 28% of all tasks at risk across the economy, when employment volumes are considered.

10 - More detailed figures for the 85 sectors of the French industry classification can be found in the appendix. Data for Sweden and Japan used for the preparation of this report, while less granular, display very similar patterns.

Our analysis points to three broad categories:

- “Skilled Trades & Industrial Production” and “In-person Manual Services” combine relatively low exposure to AI with a limited contribution to total task content at risk. They are therefore expected to play a secondary role in the current AI wave.
- “Creative and Content” and “Commercial” occupations are much more exposed to AI, but because they represent a relatively small share of total employment across countries, their contribution to total task content at risk remains modest. “Legal and Financial” occupations broadly fall into this category as well, although they are closer to the boundary.
- By contrast, “Engineering & Computational,” and “Management and Administrative” combine both high exposure to AI and a large contribution to potential task automation. As a result, they are the most likely to appear in aggregate employment figures as emerging agentic AI is adopted by companies. “Care & Education” displays a more balanced profile, combining a relatively high contribution to total task content at risk with only mid-range average exposure within the family. This reflects the coexistence of relatively exposed “Processing,” “Transcribing” and “Explaining” tasks (documentation, preparation, coordination, reporting...) with much less exposed “Caring / serving” and “Assisting” tasks at the core of teaching and care. This is consistent with the lower scores observed in the most physically and relationally intensive parts of the family, notably healthcare support and personal care occupations.

Differences in AI exposure between countries cannot be explained by any single factor.

Looking at the contribution of occupational families to total task content at risk across the workforce at a country level this time, we find results to be generally homogenous with the same top three occupational families accounting for 55% to 65% of total task content at risk across all countries. Differences between countries are therefore better explained by a combination of factors, rather than by a single factor alone. The table below summarizes how each country differs from the cross-country average as regards the contribution of occupational families to task content at risk across the workforce (Chart 9).

Tech-heavy Nordic welfare states: Denmark, Finland and Sweden share a far greater contribution of Care & Education occupations, consistent with the weight of welfare-state services in the region, alongside a comparatively lower contribution from highly exposed administrative and management categories. Their strong positions in high technology sectors, reflected in engineering and computational occupations, increase their exposure to AI.

The “headquarters trap” of the UK and the Netherlands:

both countries concentrate more employment in highly exposed, cognitively intensive occupational families - most notably in finance, IT, legal and media jobs. This is consistent with their role as European hubs and headquarters locations for a large number of multinational companies and business services activities. They also have the lowest contributions from the “Skilled Trades and Industrial Production” occupations encompassing a large share of employment from the industrial and construction sectors.

Anglo business-services economies:

Canada, the United States and Australia display a comparatively greater contribution from management, administrative, legal, and financial occupations - much like the UK. Unlike the UK however, the contribution of their “Engineering and Computational” occupations is far more limited with a sharp divide between their large and growing IT sectors but comparatively narrow engineering base.

Das Land der Ingenieure: Deutschland.

Germany forms a category of its own because of its record high contribution from “Engineering & Computational” occupations, reflecting the country’s prominent position in industrial R&D in Europe.

Outliers:

Japan, France and Poland. While their aggregate national scores do not stand out (15%, 16% and 17% of task content at risk across the workforce, respectively), those countries display more original patterns in terms of occupational family contributions and are unique in their own way. Japan’s very distinct profile is largely explained by the relatively higher weight of occupational families such as office and administrative support and sales, alongside a broader

Chart 9 - Occupational family contributions to task content at risk across the workforce, Special Agent scenario

Occupation families	Care & education	Commercial	Creative and content	Engineering & computational	In-person manual services	Legal & financial	Management & administrative	Skilled trades & industrial production
Countries								
Average contribution across countries	16%	9%	4%	17%	4%	12%	28%	10%
Denmark	Way above average	Below average	Close to average	Way above average	Below average	Close to average	Way below average	Below average
Finland	Way above average	Close to average	Close to average	Way above average	Close to average	Below average	Way below average	Close to average
Sweden	Way above average	Below average	Close to average	Above average	Close to average	Close to average	Way below average	Close to average
Netherlands	Close to average	Close to average	Above average	Above average	Close to average	Way above average	Way below average	Way below average
United Kingdom	Close to average	Below average	Above average	Above average	Close to average	Above average	Above average	Way below average
Canada	Close to average	Close to average	Above average	Below average	Close to average	Way above average	Close to average	Close to average
United States	Below average	Above average	Below average	Way below average	Above average	Above average	Above average	Close to average
Australia	Above average	Way below average	Close to average	Way below average	Close to average	Close to average	Way above average	Below average
Germany	Way below average	Below average	Close to average	Way above average	Close to average	Below average	Close to average	Close to average
France	Way below average	Close to average	Close to average	Below average	Close to average	Close to average	Above average	Way above average
Poland	Below average	Below average	Close to average	Below average	Below average	Way above average	Above average	Above average
Japan	Way below average	Way above average	Below average	Way below average	Above average	Way below average	Way above average	Above average

Sources: OEM, Coface

How to read: Across the twelve countries analyzed, care and education jobs account on average for 16% of total tasks at risk in our Special Agent scenario. Denmark, Finland, and Sweden show a much higher contribution from care and education jobs, while Germany, France, and Japan show a much lower one

mix that includes production- and personal-service-related jobs. Its fragmented retail sector and the high reliance on human-centric service, often cited as drags on productivity, serve here as protective barriers against AI automation. France's and Poland's profiles are closer to international averages, but they still display distinctive features. In particular, they show a lower contribution from "Care & Education" occupations, consistent with the relatively smaller share of these jobs in total employment compared with the European average. By contrast, "Skilled Trades & Industrial Production" occupations account for a significantly larger share of task content at risk across the workforce - in the case of France, despite the country's shrinking manufacturing base. This unusual pattern likely reflects the outsized role of construction in both the French and Polish economy, as well as the continued importance of small-scale trades and craft-based activities.

Because occupational family exposure is broadly similar across our eight European countries, the results can be extrapolated to the wider European Union. The relationship between occupational exposure and national exposure also makes it possible to apply the framework beyond the countries covered by detailed microdata. Using the average profile of our eight European countries as a proxy for countries where occupational data are less detailed, we can extend the analysis to the entire European Union, including large economies such as Italy and Spain, as well as to major non-EU economies such as Switzerland, Norway and Turkey.

We implement this extension in two steps. First, we estimate AI exposure at a detailed occupational level (ISCO 4-digit) using data from the eight countries in our baseline sample and aggregate the results to broader ISCO 2-digit occupational groups. Second, for countries lacking comparable microdata, we apply these group-level exposure rates to employment shares by ISCO 2-digit category from Eurostat, computing national exposure as the employment-weighted average across occupational groups. The broader European sample largely confirms the patterns identified in our initial twelve-country comparison (**Chart 10**).

Norway (17% of task content at risk across the workforce) and Iceland (16%) can be added to the tech-heavy Nordic welfare states cluster. Their profiles remain broadly consistent with the Nordic pattern, combining

a strong contribution from welfare-state-related occupations with a lower contribution from clerical, sales and industrial categories. Estonia (16%) also leans in this direction, reflecting its stronger specialization in information and communication technologies (ICT).

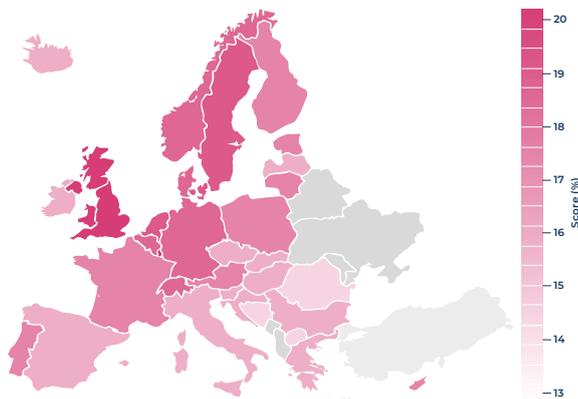
The "headquarters trap" identified for the **UK and the Netherlands** also extends to **Ireland (17%) and Luxembourg (21%, the highest score of all European countries)**. Both countries display a comparatively greater contribution from highly exposed corporate occupations, notably in finance, legal and ICT alongside a lower contribution from skilled trades and industrial production occupations. **Switzerland (18%)** also largely fits into this cluster.

A broader **Central European cluster** can also be identified around Austria, Czechia, Slovakia (16% for all three countries) as well as Slovenia, Hungary and Croatia (15% for all three countries). These countries share many similarities with Germany, and tend to combine a relatively stronger contribution from engineering, technical associate and administrative support occupations with a generally higher contribution from skilled trades and industrial production than in the "headquarters trap" cluster. Their employment structures remain more closely tied to industrial ecosystems than to headquarters-type activities.

Finally, a looser **Southern cluster** can be defined around Portugal and Italy (16%), Spain, Greece, Cyprus and Bulgaria (15%), Romania (13%) and Turkey (12%). Compared with northern and northwestern Europe, these countries generally display a lower contribution from highly exposed corporate and ICT-intensive occupations, and a greater contribution from sales, personal services, construction, transport and, in some cases, agriculture-related occupations. This is consistent with a larger role for tourism, fragmented retail, local services and labour-intensive activities in the structure of employment.

Overall, GDP per capita appears to be the most useful single heuristic for approximating AI exposure across countries (**Chart 11**). It acts as a proxy for economic structure and, in particular, for the share of cognitively intensive occupations in the economy - especially the "Processing", "Transcribing" and "Explaining" categories of our mapping. Country-level differences still reflect a broader combination of structural factors, but GDP per capita captures an important part of that variation.

Chart 10 - Task content at risk across the workforce, by country, Special Agent scenario (%)



Sources: OEM, Coface

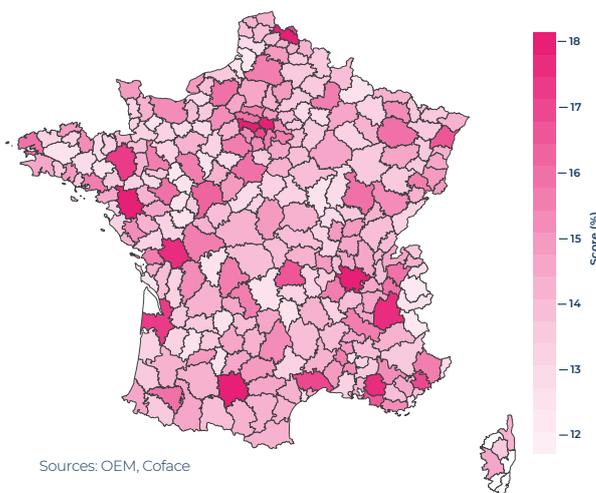
Chart 11 - 2025 GDP per capita (chain-linked 2020 euros) and % task content at risk across the workforce by country, Special Agent Scenario



Sources: Eurostat, OEM, Coface

Finally, national averages mask substantial internal variation. The same mechanisms also operate within countries, as regions specialize in different economic activities. Differences in employment composition and economic structure also play out at the regional level, raising the prospect of a geographic concentration of AI exposure (Chart 12). Applying our framework to French data at the level of employment areas, defined by INSEE as “commuting zones”, i.e. areas within which most people both live and work, we find substantial dispersion both across regions and within them. Large, urban economic hubs stand out in particular, including Paris, Lyon, Grenoble, Lille, Rennes, Nantes, Bordeaux, and Toulouse.

Chart 12 - Task content at risk across the workforce, by French commuting zones, Special Agent scenario (%)



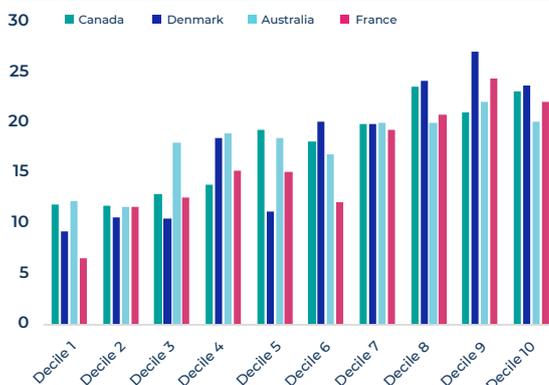
The concentration of AI automation-exposure could add strain on local systems including employment services and ripple through work-adjacent markets, including residential and commercial real estate. It also suggests possible policy responses must take into account the specifics of the local labour markets. This dynamic could also give rise to a territorial rebalancing of economic activity. Metropolitan areas have historically concentrated high-value-added, knowledge-intensive employment. Yet it is precisely this cognitive capital that AI is now targeting. Conversely, territories often perceived as less economically dynamic, built around agriculture (such as the Médoc), artisanal production or tourism (such as Corsica), possess a structural buffer. Their value added resides in physical gesture and embodied presence, two dimensions that AI cannot yet virtualize, at least as long as robotics remains insufficiently developed to replicate them at scale.

A massive impact on domestic labour markets would have macroeconomic consequences

Agentic AI could break the historical pattern: this time, disruption is likely to run up the income distribution. The occupations identified as highly exposed in the Special Agent phase are not peripheral to labour markets: they are high-volume, well-remunerated roles in finance, engineering, law, business services and office administration. Their gradual erosion through non-replacement, accelerating as agentic AI matures, could represent a structural contraction of cognitive, well-remunerated employment, not a transitory shock. This dynamic marks a break with past automation waves: unlike robots and rule-based software, which largely spared highly educated workers, agentic AI

disproportionately targets occupations requiring tertiary education and commanding high wages (Chart 13). Where previous waves hollowed out the middle of the income distribution - a phenomenon known as “polarization”¹¹ - the current wave threatens to compress it from the top.

Chart 13 - Task content at risk across the workforce, by country and income decile, Special Agent scenario (%)



Sources: OEM, Coface

How to read: each decile groups workers by their place in the national labour income distribution, from the 1st decile (the lowest-paid 10%) to the 10th decile (the highest-paid 10%). For example, the 10th decile measures exposure among the highest-paid 10% of workers and shows that 20% to 25% of the tasks of their occupations are exposed to automation in our Special Agent scenario.

AI could tilt income distribution toward capital and away from labour, a structural trend at play since the 1980s in many advanced economies. Unlike most previous waves of technological change, and similarly to the spread of software since the 1980s, many AI applications are highly scalable, allowing companies to expand output and profits without a proportional increase in employment or wages. If job destruction or non-replacement is not fully compensated, whether in the short term or over time, a larger share of the productivity gains induced by AI may accrue to the companies deploying it rather than to workers.

This mechanism may be amplified by the strong concentration of capital in the AI value chain, from advanced semiconductor manufacturing to data centers and foundation models. In other words, not only could the capital share of income increase, but income within that share could also become more concentrated among a limited number of companies. Because labour income is largely domestic but capital income crosses borders, a greater reliance on capital could also shift part of AI-related productivity gains abroad. Payments for cloud services, software licences and intellectual property may allow foreign, predominantly U.S.-based technology companies to capture a significant share of the gains generated by domestic adoption of AI.

For countries whose redistributive models rely heavily on the taxation of labour, this creates a double fiscal challenge. In the short term, should well-remunerated occupations become scarcer and displaced workers find themselves unemployed, a fiscal scissor effect emerges: the state and social security administrations simultaneously lose tax revenues and face increased expenditure, as they must finance unemployment benefits for those laid off, at least until displaced workers potentially transition into the “new occupations” expected to emerge from technological change.

More structurally, a shift from high-value-added, heavily taxed employment toward a service- and craft-based economy - personal services, tourism, construction and artisanal trades - all of which are extremely socially

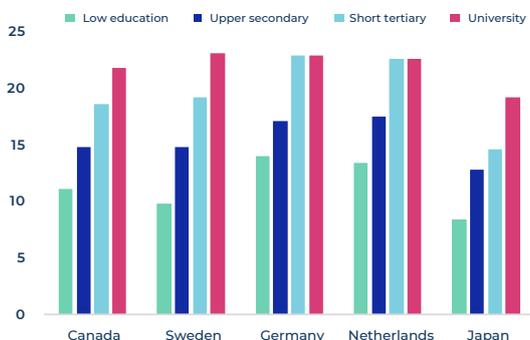
¹¹ - Maarten Goos and Alan Manning, “Lousy and Lovely Jobs: The Rising Polarization of Work in Britain,” Review of Economics and Statistics 89, no. 1 (February 2007): 118-133.

valuable but fiscally less productive given their lower wage levels, would require a **fundamental rethinking of redistributive models** in countries where financing is still based primarily on labour taxation. In such a configuration, increasing the taxation of capital to compensate may prove difficult if part of the taxable base increasingly accrues to foreign technology firms rather than remaining within domestic tax jurisdictions. In the current context - fiscal, but not only - this represents a significant macroeconomic risk for most countries, and the financing of the social protection system could be threatened in the short to medium term.

Beyond employment: diplomas, dependencies and the new fault lines of AI-led growth

AI could test the labour-market value of diplomas by impacting the most qualified occupations. Significant job destructions due to the growing prevalence of AI among occupations once occupied by highly educated individuals could also challenge the assumption that longer studies translate into greater job safety and higher wages. Employment data broken down by education level, which are available for Canada, Sweden, Germany, the Netherlands and Japan, that is countries with different education systems and economic structure, underline a clear correlation between education level and exposure to AI (**Chart 14**).

Chart 14 - Task content at risk across the workforce by country and education level, Special Agent scenario (%)



Sources: OEM, Coface.

Germany and Netherlands incorporate short tertiary and university degrees into a same category.

Prominent personalities from the technology industry including Sam Altman (Open AI), Demis Hassabis (Google DeepMind), Jensen Huang (Nvidia) or Elon Musk (xAI) have questioned whether higher education as it is today will retain the same labour-market value in an AI-intensive economy. As AI systems increasingly perform knowledge retrieval and routine cognitive tasks, they argue that the signalling power of degrees may decline, with employers placing greater weight on skills, adaptability and the ability to effectively use AI tools than on traditional academic credentials. In this view, higher education would move away from the transmission of knowledge toward the cultivation of judgment, critical thinking and learning capacity - competencies that remain complementary to, rather than substitutable by, AI.

Many academics, including Erik Brynjolfsson (Stanford/MIT), Daron Acemoglu (MIT) and Fei-Fei Li (Stanford), emphasize a closely related point but draw a different implication: as AI automates execution, human workers will increasingly be tasked with supervision, interpretation and decision-making, raising demand for analytical training and technical literacy. This would reinforce the role of higher education in skill formation,

provided it adapts its content and pedagogy, including by selectively integrating AI as a learning tool.

AI could strengthen supply-chain and geopolitical dependencies as it becomes a core input to production.

This same concentration of critical technological assets into a few hands also carries strategic implications as regards supply chain resilience among firms and strategic autonomy among countries. By shifting production away from distributed labour to more concentrated forms of capital, an AI-intensive economy could increase the risk of systemic failures, whether triggered by geopolitical shocks (export controls, sanctions, regulatory fragmentation) or supply chain constraints (semiconductors, energy, computing capacity) disruptions, and amplify operational fragilities (cyberattacks, cloud service disruptions). Even as AI raises productivity across large parts of the economy, it might as well create new "single points of failure" through which shocks could propagate.

How far could disruption ultimately extend if AI capabilities continue to advance?

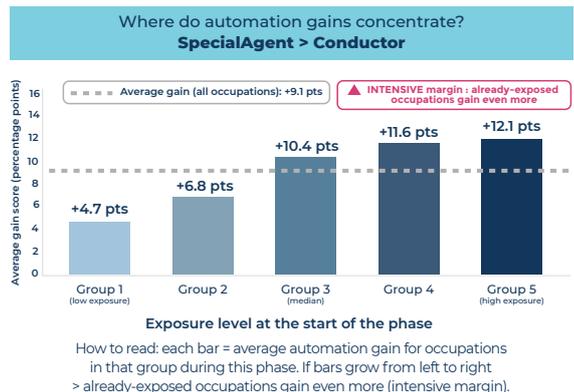
All the results discussed so far refer to our near-term "Special Agent" scenario, centred on the emerging agentic phase of AI; the following section asks how much further disruption could extend if capabilities continue to advance beyond that threshold. To address this question, we turn to the three more speculative phases that lie well beyond current capabilities, beginning with the Conductor phase, where episodic memory enables AI systems to store dated interaction episodes, thereby unlocking sophisticated planning and coordination abilities.

The Conductor phase marks a significant acceleration.

With an automation score of 26.3%, this phase is the first in which the majority of occupations cross a 30% automation exposure threshold: 375 out of 923 (40.6%). While the distribution becomes notably more symmetric, reflecting a broad diffusion of automation potential across the occupational spectrum rather than concentration on a few occupations at the top, this phase is still largely driven by already-exposed occupations becoming disproportionately more at risk (the "intensive margin").

The Conductor phase reshuffles the automation exposure hierarchy toward coordination and orchestration roles.

This pattern is theoretically coherent: a Conductor-level AI can sustain a plan, allocate resources across sub-agents, and manage inter-task dependencies, i.e., functions that previously defined the comparative advantage of human managers. What unifies the top gaining occupations is precisely what episodic memory unlocks: the ability to maintain a persistent, updateable internal representation of a situation across time, rather than processing each interaction as a stateless event.



Sources: OEM, Coface

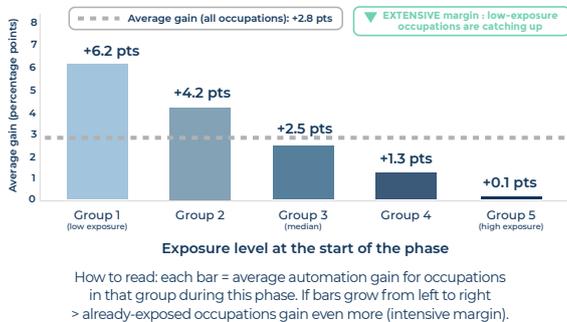
The Prometheus Bound phase represents a consolidation rather than a rupture. The mean score rises more modestly to 29.1%, the smallest incremental gain of any transition. Crucially, however, the distribution compresses. The 30% threshold is now crossed by 44.6% of occupations (412 out of 923).

This transition is characterized above all by breadth rather than depth: it is not the already-exposed occupations that become even more affected, but a set of relatively less-impacted occupations that become more heavily exposed (the “extensive margin”). Low-exposure occupations gain 6.2 points on average while those already heavily exposed gain a near-negligible 0.1 points. This pattern reflects the core assumption embedded in this scenario, namely that physical robotics reaches sufficient maturity to automate embodied, manual, and spatially-located tasks. The occupations gaining most in this transition are predominantly trades and physical installation work.

The share of occupations below the 5% threshold falls to zero: no occupations under this extreme scenario remains structurally shielded. The most protected occupations remain anchored in physical, relational or embodied work. But unlike in earlier phases, horizontal complexity no longer provides a shield: when AI systems face no cognitive ceiling, the irreducible variability of human situations ceases to be a barrier to automation.

The chart brings these scenario results together by tracking the share of occupations with at least 30% of tasks exposed to automation at each stage of AI development. Taken together, these scenarios show that advances in AI capabilities do not simply deepen exposure in the occupations already most at risk, but that they also broaden it.

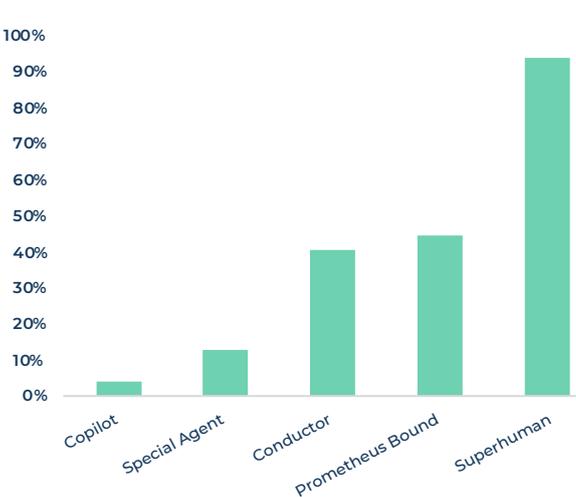
**Where do automation gains concentrate?
 Conductor > PrometheusBound**



Sources: OEM, Coface

The Superhuman phase offers a speculative but instructive outer bound. This most exploratory phase marks a qualitative break with previous scenarios. It captures what becomes technically automatable if AI systems reach a level of general capability comparable to or exceeding that of humans across all cognitive domains, a condition that remains highly speculative and uncertain. With a mean of 49.5%, almost all occupations cross the 30% automation threshold (93.7%). The Conductor-to-Superhuman transition is driven by occupations that were least exposed at the start of this phase gaining nearly twice as much as those already heavily automated (the “extensive margin”): Group 1 (low exposure) gains +27.6 points against only +15.8 for Group 5 (high exposure). Far from compounding existing inequalities, this final leap dissolves them: automation reaches into the last sheltered occupations.

Chart 15 - Occupations with ≥ 30% of tasks automatable along scenarios (%)



Sources: OEM, Coface

**Where do automation gains concentrate?
 Conductor > Superhuman**



Sources: OEM, Coface

BOX 2 - TECHNICAL DIMENSIONS

THE PHASES ARE ARTICULATED AROUND ADVANCES RELATING TO:

Autonomy: what complexity, what task duration can a system carry out autonomously, i.e. without human supervision?

For example, moving from the simple calculation of an individual bonus to managing the complete monthly payroll cycle, including anomalies and adjustments.

Memory: what is the model's context window size, i.e. the elements the model is able to refer to coherently in order to respond to user requests?

For example, moving from managing an isolated monthly monitoring report for a real estate programme, to tracking the evolution of a project over 6 months while retaining the context of previous decisions, through to steering a multi-year real estate programme with full memory of technical and budgetary choices.

Reliability: what is the error rate, and in particular the propensity to hallucinate, i.e. to fabricate entirely false content?

For example, moving from extracting invoice amounts with 15-20% errors requiring systematic human verification, to 5-8% errors thanks to finance-specialised models, through to reaching error rates below those of humans, requiring virtually no further human validation.

Reasoning: how many reasoning steps can a model process?

Is it capable of reasoning in terms of causality, or does it only perform pattern matching, meaning it applies statistics without genuine understanding? Is it capable of handling novel problems it has never encountered before, including problems absent from training data? For example, moving from extracting and summarising contractual clauses, to detecting inconsistencies between clauses, through to analysing the full legal impact of a cascading amendment on warranties, liabilities, etc.

These various dimensions are interdependent: better memory can, for example, enable improvements in a model's reasoning capabilities. The "autonomy" dimension implicitly encompasses these various dimensions from which it results: more tasks will be delegated autonomously to a system that has good memory, a low error rate, and genuine reasoning capabilities.

BOX 3 - MORE DETAILS ON THE TECHNOLOGICAL PHASES STUDIED

THE TABLE BELOW SUMMARIZES THE TECHNOLOGICAL PHASES USING THE CORE TECHNICAL DIMENSIONS UNDERLYING OUR SCENARIOS.

Technical Dimension	Phase 1 Copilot	Phase 2 Special Agent	Phase 3 Conductor	Phase 4A Prometheus Bound	Phase 4B Super human
Defining features	Language models are called upon as simple, one-off tools by users. They accelerate task execution, but production processes themselves remain unchanged.	Transition from "chatbot" to "agent." Multiple models are orchestrated and given access to tools, allowing them to execute sequences of actions.	Main engineering bottlenecks are assumed to be resolved, and memory becomes intrinsic and continuous rather than external and fragmentary.	Technological plateau: agenticAI is fully mature and robotics advances, but no breakthrough in causal reasoning occurs.	Technological revolution: AI surpasses human performance across all cognitive domains.
Architecture / technical setup	A single LLM is used as a one-off supplement.	Multiple existing language models are orchestrated in sequence or in parallel. Models are given access to tools such as search, APIs, file opening, email or voice calls.	Technical framework remains relatively close to the current one, namely LLMs, but with genuine advances in memory and much more robust multi-agent operation.	Mature agentic systems plus progress in robotics.	Systems reach a level of general capability comparable to or exceeding that of humans across cognitive domains.
Autonomy	Near-zero autonomy: humans create prompts and use the responses formulated by the models to carry out their work.	Autonomy progresses but remains limited: the system can plan and execute chains of sub-tasks, but critical steps still require validation.	Autonomy increases considerably: human intervention points are significantly reduced. AI can manage certain cases end-to-end.	Very high autonomy within the limits of absent causal reasoning. Robotics extends automation to more embodied tasks.	Maximum autonomy: cognitive ceilings disappear.
Memory	Very limited to a fairly short context window. Long-term temporal tracking is difficult to automate.	Memory remains external, not integrated into the models. The orchestrator retrieves relevant elements and injects them into context.	Memory becomes intrinsic and continuous. Agents retain the complete history of their interactions and learn cumulatively from past errors.	Mature memory, but still without a breakthrough in causal reasoning.	No meaningful cognitive memory constraint remains.
Reliability	Hallucinations are frequent and reduce model reliability, making precision-demanding tasks harder to automate.	Reliability evolves in a paradoxical way: specialized models improve performance, but orchestration introduces cascading error risks.	Reliability improves: cross-agent verification moderates the propagation of cascading errors.	Reliability continues to improve, especially for embodied execution through robotics, but systems remain bounded by the plateau in reasoning.	Reliability reaches or exceeds human levels across cognitive tasks.
Reasoning	Limited: for tasks requiring more than one reasoning step, capabilities collapse.	Reasoning progresses over chains of 3-4 logical steps thanks to orchestration, but remains fragile beyond that.	Reasoning improves incrementally, without any major breakthrough: longer reasoning chains are handled more reliably, but models still operate through pattern-matching rather than genuine causal understanding.	No breakthrough in causal reasoning. Systems remain highly capable but cognitively bounded.	AI reaches or exceeds human-level reasoning across domains.
Human role	The worker remains the operator: they prompt, verify, correct and send.	The worker becomes a flow supervisor, validating critical steps and correcting agent deviations.	The worker becomes a complex-case manager, intervening mainly for ambiguous, novel or exceptional cases.	Human intervention is concentrated in the hardest embodied, strategic or genuinely novel edge cases that remain outside the system's plateau.	Human oversight is no longer technically required for most cognitive tasks.
Persistent limitations and bottlenecks	Short context window, hallucinations, weak multi-step reasoning.	Cascading errors, fragile coordination, external memory, and limited reasoning beyond short chains.	System remains probabilistic in its operation and does not "understand" causal meaning the way a human does.	Plateau in causal reasoning; remaining constraints are no longer mainly organizational but technological.	This phase is explicitly speculative and represents an outer bound rather than a near-term scenario.
Expected impact on work	AI serves as a one-off supplement.	Workers shift from operators to supervisors of validation points.	AI becomes a virtual collaborator capable of managing certain cases end-to-end; management of edge cases remains human.	Consolidation of mature agentic automation and wider diffusion into physical work through robotics.	Automation potential extends across almost all occupations, including the last previously sheltered cognitive tasks.

PHASE 1 corresponds to the current use of LLMs as punctual assistants embedded in otherwise unchanged production processes.

A Payroll Administrator may ask ChatGPT to draft an email explaining a deduction, but still verifies, corrects and sends it.

A Management Controller can obtain plausible explanations for a budget variance, but complex multi-step financial reasoning remains unreliable.

A Real Estate Programmes Specialist cannot delegate the management of a multi-month project to a system that "forgets" prior decisions.

PHASE 2 marks the shift from chatbot to agent. Instead of scripting each step, the user defines an objective and the agent plans and executes the sequence of actions. An agent may extract data from a payroll system, analyse it and send a summary email.

But this gain in autonomy comes with new risks: as multiple models interact, errors may propagate across the chain, which is why human-in-the-loop validation remains central.

PHASE 3 assumes that the main engineering bottlenecks of agentic AI have been resolved and that episodic memory becomes intrinsic rather than external. AI can then manage complete case files over time, while the human role shifts from continuous supervision to the treatment of exceptions. The key limit is that even at this stage, the system remains probabilistic and does not achieve genuine causal understanding.

PHASE 4A AND 4B represent the final bifurcation of the framework:

either a technological plateau in which agentic AI matures without a breakthrough in causal reasoning, or a superhuman scenario in which AI exceeds human performance across cognitive domains.

BOX 4 - GLOSSARY OF TECHNICAL TERMS

Agentic AI:

An architecture that combines multiple language models orchestrated together and equipped with concrete tools (internet search, file access, email). Unlike a simple chatbot, agentic AI can break down a complex task into sub-tasks, delegate them to specialised agents, and cross-check their results to improve the overall reliability of the system.

Agent coordination

(also: multi-agent orchestration):

In architectures where multiple specialized sub-agents collaborate on a task, coordination refers to the mechanisms ensuring consistent task allocation, information sharing, and conflict resolution between agents. Failures include redundant actions, contradictory outputs, and deadlocks.

Context window:

The quantity of information a language model can process simultaneously when responding to a query. A short context window limits the model to a few pages of text, while an extended window allows it to process the equivalent of entire books. For managing long-running projects, a large context window is necessary to maintain coherence.

Context window management

(also: context length, long-context handling):

The challenge of maintaining coherent, relevant information within the finite input window of an LLM across extended tasks. Strategies include summarization, retrieval-augmented memory, and sliding window techniques, each introducing tradeoffs between fidelity and efficiency.

Episodic memory:

A system's capacity to intrinsically retain the full history of its past interactions and experiences. Unlike current external memory systems, which retrieve filtered excerpts, true episodic memory would allow a system to continuously recall all previous exchanges and to learn cumulatively from its mistakes.

Error compounding

(also: cascading errors, error propagation):

In multi-step agentic pipelines, errors made at one step are passed as inputs to subsequent steps, where they are amplified rather than corrected. The further down the chain, the greater the divergence from the intended outcome.

Latency in long-running workflows

(also: inference latency, end-to-end latency):

The cumulative time delay introduced by sequential LLM inference calls and tool invocations across a multi-step agentic chain. While single-call latency has improved substantially, compounded latency over long workflows remains a practical bottleneck for real-time or interactive applications.

LLM (Large Language Model):

A large-scale language model trained on billions of words to predict the next word in a sequence. These models, such as ChatGPT, operate through statistical pattern-matching: they identify regularities in their training data to produce fluent, coherent text, without necessarily "understanding" in the human sense.

Machine learning:

A subfield of artificial intelligence that enables a computer system to learn without being explicitly programmed for each task. The system can learn either from training data (for example, by analysing thousands of transactions to detect fraud) or through interaction with its environment, receiving rewards or penalties based on its actions (for example, learning to play chess through trial and error).

Orchestration:

A mechanism that coordinates multiple language models in sequence or in parallel to accomplish a complex task. The orchestrator breaks down a global request into sub-tasks, assigns them to the appropriate specialised agents, synchronises their contributions, and combines their outputs to produce a coherent and reliable response.

Pattern-matching (or pattern recognition):

The operating method of language models, which consists of identifying statistical regularities in data to generate responses. The model finds the most probable next word based on billions of learned parameters, without necessarily reasoning in terms of causality or conceptual understanding as a human would.

RPA (Robotic Process Automation):

A software technology that automates repetitive, rule-based tasks by replicating the interactions a human user would perform on a computer: clicking, copying, pasting, filling in forms, extracting data from documents. Unlike machine learning, RPA does not learn from data: it follows explicitly programmed sequences of instructions. RPA was the dominant automation paradigm of the 2010s.

Tool-call reliability

(also: tool use, function calling):

The ability of an LLM-based agent to correctly invoke external tools (APIs, search engines, code interpreters, file systems) with well-formed parameters and appropriate timing. Current failure modes include hallucinated tool names, malformed arguments, and incorrect sequencing.

Training data:

A set of examples used to teach an AI model how to perform a task. The model analyses this data to identify statistical patterns and adjust its internal parameters, known as "weights". For example, to train a language model, billions of texts are provided (books, articles, websites): the model learns to predict the next word in a sentence and to generate coherent text. This is referred to as "self-supervised" training.

Appendix

Extracting Elementary Actions from Each Occupation Using the O*NET Database

Each O*NET task is automatically decomposed into simple elementary actions according to the following structure: verb + direct object + contextual element.

Concrete example for the occupation "Procurement Manager":
O*NET task: "Locate suppliers of materials, equipment or supplies, and interview them in order to determine product availability and terms of sale."

Extracted elementary actions:

- locate, suppliers
- conduct, interview
- determine, availability, product
- determine, terms, commercial

Why does this granular analysis matter?

Because each elementary action has a distinct propensity for AI automation. Rating the overall task as a whole does not allow for sufficient analytical precision or adequate tracking. The granular approach increases the transparency and accuracy of the method.

Rating Each Elementary Action

A set of rules is used to rate each elementary action. The rules apply either to verbs, direct objects, or contextual elements in isolation, or to the verb,direct object verb,direct object pair, or to the full verb,direct object, contextual element triplet.

Distinct semantic groups have been constructed in order to capture technological advances across technical dimensions, including models' capacity for strategic decision-making, improvisation, and so on (**see Box 1**).

Aggregation into an Overall Occupation Score

Each O*NET task receives a score corresponding to the average of its elementary actions, and each occupation receives an aggregated score corresponding to the weighted average of the scores of each task. The weighting draws on information from the O*NET database and combines the frequency with which each task is performed and its importance, as assessed by a panel of incumbent workers.

Mapping AI exposure in different countries

Occupational classifications differ across countries, but concordance tables allow them to be mapped to a common framework. This enables us to use U.S.-specific information, such as detailed occupation task content. To perform cross-country comparisons, identify national patterns and incorporate sociodemographic variables available only in some datasets, we convert occupation-level data from national statistical sources into the U.S. Standard Occupational Classification (SOC) using well-documented crosswalks provided by the International Labour Organization (ILO) and other official sources.

All the countries included in the analysis are members of the ILO and therefore aim to align their labour statistics with the International Standard Classification of Occupations (ISCO).

From 923 distinct occupations to 8 occupational families

For the sake of clarity and analysis, we have grouped the 923 occupations into eight occupational families, drawing on the two digit occupational families used by the U.S. Standard Occupational Classification (U.S. SOC). Our grouping is based on common patterns in the task content of occupations and, consequently, in exposure to AI. The reconciliation between the two classification is as follows:

US SOC code	US SOC occupation family	Coface occupation family
11	Management occupations	Management & administrative
13	Business and financial operations occupations	Legal & financial
15	Computer and mathematical occupations	Engineering & computational
17	Architecture and engineering occupations	Engineering & computational
19	Life, physical, and social science occupations	Engineering & computational
21	Community and social service occupations	Care & education
23	Legal occupations	Legal & financial
25	Educational instruction and library occupations	Care & education
27	Arts, design, entertainment, sports, and media occupations	Creative & content
29	Healthcare practitioners and technical occupations	Care & education
31	Healthcare support occupations	Care & education
33	Protective service occupations	In-person manual services
35	Food preparation and serving related occupations	In-person manual services
37	Building and grounds cleaning and maintenance occupations	In-person manual services
39	Personal care and service occupations	Care & education
41	Sales and related occupations	Commercial
43	Office and administrative support occupations	Management & administrative
45	Farming, fishing, and forestry occupations	In-person manual services
47	Construction and extraction occupations	Skilled trades & industrial production
49	Installation, maintenance, and repair occupations	Skilled trades & industrial production
51	Production occupations	Skilled trades & industrial production
53	Transportation and material moving occupations	Skilled trades & industrial production

Task content at risk across the workforce by sector: the case of France

The following table details the sector impact of our Special Agent scenario for the French economy ranked by ascending order:

NACE label	NACE code	% employment weighted tasks at risk
Services to buildings and landscape activities	81	6.5
Fishing and aquaculture	3	6.8
Crop and animal production, hunting and related service activities	1	7.3
Food and beverage service activities	56	7.5
Accommodation	55	8.8
Other personal service activities	96	10.2
Forestry and logging	2	10.8
Employment activities	78	11.2
Manufacture of leather and related products	15	11.6
Mining of metal ores	7	11.7
Specialised construction activities	43	11.9
Social work activities without accommodation	88	12.0
Manufacture of wood and of products of wood and cork	16	12.8
Manufacture of food products	10	12.9
Sports activities and amusement and recreation activities	93	13.2
Veterinary activities	75	13.4
Manufacture of textiles	13	13.4
Manufacture of furniture	31	13.5
Residential care activities	87	13.5
Manufacture of other non-metallic mineral products	23	13.5
Manufacture of fabricated metal products	25	13.6
Waste collection, treatment and disposal activities	38	13.8
Land transport and transport via pipelines	49	13.9
Other mining and quarrying	8	13.9
s	24	14.0
Remediation activities and other waste management services	39	14.3
Manufacture of paper and paper products	17	14.4
Wholesale and retail trade and repair of motor vehicles and motorcycles.	45	14.7
Construction of buildings	41	14.7
Manufacture of rubber and plastic products	22	14.9
Human health activities	86	14.9
Civil engineering	42	14.9
Sewerage	37	15.3
Manufacture of wearing apparel	14	15.4
Other manufacturing	32	15.6
Warehousing and support activities for transportation	52	15.6
Libraries, archives, museums and other cultural activities	91	15.6
Manufacture of motor vehicles, trailers and semi-trailers	29	15.7
Air transport	51	15.7
Security and investigation activities	80	16.0
Manufacture of beverages	11	16.1
Retail trade, except of motor vehicles and motorcycles	47	16.1

NACE label	NACE code	% employment weighted tasks at risk
Water transport	50	16.3
Repair and installation of machinery and equipment	33	16.4
Manufacture of coke and refined petroleum products	19	16.7
Water collection, treatment and supply	36	17.1
Printing and reproduction of recorded media	18	17.2
Manufacture of electrical equipment	27	17.4
Rental and leasing activities	77	17.4
Manufacture of chemicals and chemical products	20	17.5
Gambling and betting activities	92	17.5
Manufacture of tobacco products	12	17.7
Repair of computers and personal and household goods	95	17.8
Manufacture of basic pharmaceutical products and pharmaceutical preparation	21	17.9
Manufacture of machinery and equipment n.e.c.	28	17.9
Activities of membership organisations	94	18.2
Wholesale trade, except of motor vehicles and motorcycles	46	19.0
Extraction of crude petroleum and natural gas	6	19.1
Real estate activities	68	19.1
Creative, arts and entertainment activities	90	19.6
Electricity, gas, steam and air conditioning supply	35	19.9
Travel agency, tour operator reservation service and related activities	79	20.0
Manufacture of other transport equipment	30	20.2
Education	85	20.4
Manufacture of computer, electronic and optical products	26	20.5
Postal and courier activities	53	21.2
Activities of extraterritorial organisations and bodies	99	21.5
Motion picture, video and television programme production, etc.	59	21.7
Scientific research and development	72	22.3
Mining support service activities	9	22.4
Office administrative, office support and other business support activities	82	22.4
Other professional, scientific and technical activities	74	23.0
Architectural and engineering activities; technical testing and analysis	71	23.4
Advertising and market research	73	23.5
Activities of head offices; management consultancy activities	70	23.8
Activities auxiliary to financial services and insurance activities	66	24.6
Public administration and defence; compulsory social security	84	24.8
Telecommunications	61	25.2
Financial service activities	64	25.4
Insurance, reinsurance and pension funding	65	25.5
Programming and broadcasting activities	60	26.1
Information service activities	63	26.5
Computer programming, consultancy and related activities	62	26.8
Publishing activities	58	26.8
Legal and accounting activities	69	27.0

DISCLAIMER

This document reflects the opinion of Coface's Economic Research Department at the time of writing and based on the information available. The information, analyses and opinions contained herein have been prepared on the basis of multiple sources considered reliable and serious; however, Coface does not guarantee the accuracy, completeness or reality of the data contained in this guide. The information, analyses and opinions are provided for information purposes only and are intended to supplement the information otherwise available to the reader. Coface publishes this guide in good faith and on the basis of commercially reasonable efforts as regards the accuracy, completeness, and reality of the data. Coface shall not be liable for any damage (direct or indirect) or loss of any kind suffered by the reader as a result of the reader's use of the information, analyses and opinions. The reader is therefore solely responsible for the decisions and consequences of the decisions he or she makes on the basis of this guide. This handbook and the analyses and opinions expressed herein are the exclusive property of Coface; the reader is authorised to consult or reproduce them for internal use only, provided that they are clearly marked with the name «Coface», that this paragraph is reproduced and that the data is not altered or modified. Any use, extraction, reproduction for public or commercial use is prohibited without Coface's prior consent. The reader is invited to refer to the legal notices on Coface's website: <https://www.coface.com/Home/General-informations/Legal-Notice>

COFACE SA

1, place Coste et Bellonte
92270 Bois-Colombes
France
www.coface.com