



# **Workplace Risk Tradeoffs in the Era of AI: Policies for Success**

Matteo **Mio** and Stefano **Serra**

May 12th; 2025

VOYCE

TECHNOLOGY COMMISSION

# Contents

- 1 Introduction ..... 1**
- 2 What is Professional Risk? ..... 2**
  - 2.1 Income Security ..... 3
  - 2.2 Social Security and Downside Protection ..... 3
  - 2.3 Legal and Fiscal Risk ..... 4
  - 2.4 Autonomy and Agency ..... 4
- 3 Analysis ..... 6**
  - 3.1 Large Companies ..... 6
  - 3.2 Small and Medium and Innovative Companies ..... 14
  - 3.3 Entrepreneurship ..... 21
- 4 From “Risk Framework” to “Policy scoreboard” ..... 25**
  - 4.1 What the scoreboard is ..... 25
  - 4.2 Why individual rights aren’t enough in workplaces ..... 27
  - 4.3 Why the EU level matters ..... 28
- 5 What Does Good Look Like? ..... 30**
  - 5.1 Income Insecurity: Transitions that are fast, fair, and do not scar earnings ..... 30
  - 5.2 Social Security and Downside Protection: Coverage that follows workers across status and shocks ..... 31
  - 5.3 Legal and Fiscal Risk: accountability that is workable, predictable, and contestable ..... 32
  - 5.4 Autonomy and Agency: productivity without pervasive surveillance or psychological harm ..... 32
- 6 Current EU Policy Landscape ..... 34**
  - 6.1 Income security ..... 34
  - 6.2 Social security and downside protection ..... 34

6.3	Legal and Fiscal risks .....	35
6.4	Autonomy and agency .....	35
6.5	Where do these tools stop? .....	35
<b>7</b>	<b>Risk Patterns Across Firm Types .....</b>	<b>36</b>
7.1	Large Firms: .....	37
7.2	SMEs: .....	38
7.3	Synthesis: .....	40
<b>8</b>	<b>From Diagnosis to Intervention .....</b>	<b>41</b>
8.1	EU Channels of Interventions .....	42
8.2	Income Security Levers .....	44
8.3	Social Protection and Downside-Protection Levers .....	46
8.4	Legal and Fiscal Risk Levers .....	48
8.5	Cross-cutting Design Principles: How to Combine EU Levers .....	51
<b>9</b>	<b>From Intervention Menu to Prioritization .....</b>	<b>53</b>
9.1	Adoption and Competitiveness vs Protection and Oversight .....	54
9.2	Harmonization vs Subsidiarity and National Diversity .....	56
9.3	Simplicity and Proportionality vs Robustness and Accountability .....	57
9.4	Transparency and Contestability vs Confidentiality and Administrative Feasibility .....	58
9.5	Prevention vs Compensation .....	59
9.6	Uniform Obligations vs Differentiated Pathways .....	61
<b>10</b>	<b>Prioritised Recommendations .....</b>	<b>62</b>
10.1	Policy Package at a Glance .....	64
10.2	R1 - Professional AI Proficiency Certificate .....	66
10.3	R2 - AI Transition & Reskilling Guarantee .....	79
10.4	R3 - SME AI Capability & Compliance Hub .....	89
10.5	R4 - AI Augmentation Incentive Standard .....	99

10.6 R5 - AI Entrepreneurship Bridge .....	105
10.7 Policy Direction - An EU Wide Movement to Implement Real Change .....	110
<b>11 Conclusion .....</b>	<b>112</b>
<b>Bibliography .....</b>	<b>115</b>

# 1 Introduction

The recent exponential integration of Generative AI tools in the Workplace has been a transformative event for most companies, businesses, and more generally the whole workforce. Mckinsey reports 71% of the businesses targeted by their studies on AI in the workplace globally were using GenAI tools in at least one function by the end of 2024 (McKinsey & Company, 2025a). Job markets all around the globe have seen volatility both on the supply and the demand side, with many entry-level positions becoming unnecessary as AI and automation allow companies to optimize and reduce the need for more workers. Meanwhile, on the supply side applying to jobs becomes easier and faster as applicants use GenAI powered platforms to generate resumes, cover letters and even to directly apply to positions (World Economic Forum, 2025).

As the AI industry grows and adapts, its long-term impact becomes inevitable. More particularly, its possible role in replacing some jobs is crucial, as it allows companies to cut costs and enhances others, allowing for higher individual productivity in many tasks. Workforce sentiment about the future of jobs is shown to be pessimistic by reports, with workers reporting fears of fewer opportunities due to AI replacement (Pew Research Center, 2025). Simultaneously, thanks to the advent of Large Language Models, Low-Code Development tools and the possibility to automate tasks through AI Agents, the opportunity for entrepreneurship becomes more accessible and less expensive (Gartner, 2025; JPMorgan, 2025). This report will take on the task of exploring the impacts of these new technologies on the traditional tradeoffs between working as an employee in a company structure and being an entrepreneur.

To do so, it will analyse different European case studies and expert insights through the lens of Professional Risk, identify the impact of AI on its subcategories and finally explore the possibility for European Union (EU) policymakers to seize this moment of change as an opportunity to give European workers an edge to optimize their means of value creation in an AI-augmented environment.

## 2 What is Professional Risk?

The notion of risk has always been a factor in human decisions. From ancient times, when the risk of death was the price of survival in a primitive natural state, to modern capitalist societies where financial risk dominates the direction of our markets, risk has remained a driving force of economic crisis and opportunity. Through our centuries of history, it has also defined our career choices. More often than for genuine interest or curiosity, humans choose certain professional paths in direct consequence to their propensity to risk encountering financial and social instability.

Traditionally, it's common to perceive entrepreneurial careers as more dangerous, although potentially more profitable, while corporate careers, contributing to someone else's business as their employee, has always been seen as more stable and attractive to ensure a certain, basic level of wellbeing. Think of how common and stereotypical it has become to hear of humans dreaming of becoming artists, self-made millionaires or sportsmen once they are young and then settling for an office job once entering the workforce. This is the result of the high levels of risk associated with entrepreneurial paths, in which one has to create value for themselves often starting from scratch, and the positive feeling of financial security linked to being an employee, creating value for someone else that can provide a stable and sure salary as well as standardized working conditions.

Now, for the purpose of this research piece, it is important to define professional risk more precisely. The following framework will be the main tool used in this piece in order to analyse the impact of the qualitative and quantitative research results as well as the impact of the policies that will be assessed in the Policy part. It builds on two major sources: The *OECD Job Quality Framework* (OECD, 2015) and the *UNECE Quality of Employment Framework*. Both sources represent some of the most complete definitions of a qualitative workplace and are therefore used to map the required topics for this framework. On top of using both sources to build the foundation of this analytical tool, a focus on professional risk has been added, to adapt the whole

structure to assess not just the general quality of a job position but more in detail the risks and opportunities that surround it, from the perspective of an individual.

The framework is represented by four main sub-dimensions.

## **2.1 Income Security**

This part includes all factors that may affect the stability of current and future income of a worker. It is highly influenced by the human wish to have their financial situation under control.

In case of the entrepreneur/employee tradeoff, the most relevant factors are:

- **Earning predictability and consistency**;
- **Time-to-security**, or the time it can take from the first day of work to reach a stable income that can persist in the future;
- **Displacement risk**, or the chances of being forced to change to an inferior role or being laid off completely; and
- **Skill and knowledge obsolescence risk**, or the likelihood that the basis by which a person occupies a certain role, creates a certain value, becomes obsolete and therefore puts their position at risk.

## **2.2 Social Security and Downside Protection**

Here we group all tools that can be used to protect workers in times of crisis and shocks or to guarantee some level of income and wellbeing aside from salaries earned during working hours.

They can take different forms, such as:

- **Safety nets** like contribution and access to unemployment insurance, severance pay, paid maternity/paternity leave.
  - In this section, we also find corporate contributions to institutional benefits such as health insurance, pension funds, subsidized childcare and transport;

- **Personal financial cushions**, such as savings or individual contributions to pension funds and investments; and
- **Asset ownership and protection**. This part describes the risk of working in an environment in which your own assets can be either used or leveraged in the business. This can include personal funds, financial assets and liabilities, office space, and others.

## 2.3 Legal and Fiscal Risk

This part is quite self explanatory: any position holds a certain set of responsibilities, which are linked to different legal requirements and fiscal compliance needs that the worker or their company must comply with. The larger the set of responsibilities, the higher the risk of legal troubles and therefore the more effort and resources are needed to make sure that everything runs smoothly. It is important to note that, although many of such responsibilities are similar for both employees and entrepreneurs, employees who belong to a large corporate structure can benefit from tailored departments that handle legal and fiscal procedures for them. These can include:

- **Tax compliance;**
- **Labor law constraints;**
- **Regulatory barriers and industry-specific regulations;** and
- **Fines, legal costs and operational burden.**

## 2.4 Autonomy and Agency

It can be defined as the psychological and social part of professional risk. Its mission is to enable the inclusion of more abstract factors, such as stress, needs for social inclusion, and perception of social utility into any analysis of professional risks and satisfaction of the labour force. We often tend to think of professional risks and benefits as merely tangible concepts, such as income, safety nets, and career opportunities, but in reality it is as important for working professionals to

feel not only included in a certain social environment but also stress-free, autonomous in their tasks and decisions and potentially to believe that the value they contribute to producing has a positive impact on themselves and the world around them. The different key factors to map here are:

- **Decisional autonomy and authority** can on one hand be a positive factor, as freedom to decide and take initiative can increase the sense of purpose and the motivation to work, and on the other hand be a negative as it increases responsibility and agency over potential risks of managing people or capital autonomously;
- **Career path predictability** which resonates with the consistency of a certain career path over time and the ability to expect a similar journey through a company or a field as older professionals who can be taken as example. This allows recent graduates to figure out which entry-level jobs can lead to which long-term career for example, or it can be useful for senior professionals to understand which position to aim for in a company;
- **Work-life balance autonomy** allows one to get a sense of an individual's capacity to develop himself outside of the workplace and therefore is an important signal of stability and possibilities for sustainable growth in that certain working environment; and
- **Social integration and sense of belonging** as the main social factor of this definition, to map the differences in social satisfaction between more dynamic and human centric environments and others which may be less socially rewarding.

This definition of professional risk will be the fundamental tool used in this report in order to assess the impact of the integration of AI in different work environments on the multiple factors identified in each category.

The methodology used will be such that:

- Each case study is analysed to identify direct impacts on any of these factors "all else equal", such that the risk of bias or hidden involvement of any other influence is minimized.
- Once such impact is isolated and possibly quantifiable, the goal is to answer the following questions:

- Does it increase/decrease professional risk?
- Is it likely to be permanent or just temporary?
- Is it impactful to a large enough scale that political intervention is required?

## **3 Analysis**

### **3.1 Large Companies**

To begin analysing the AI integrated professional landscape, this report will tackle large companies (250+ employees), their propensity to integrate AI into their operations, and the impact it can have on different roles in the company. First of all, some context is needed to introduce this category. As the amount of European businesses using AI in their operations continues to rise, the proportion of large businesses integrating such technologies appears significantly higher than that of SMBs (Eurostat, 2025a). The amount of European companies concerned by this data is very likely to constantly increase over the next few years, due to industry trends and integration of AI in commonly used softwares such as CRMs and outreach systems, which are undergoing relevant changes due to the possibility of reaching higher efficiency thanks to AI augmented features. These examples only represent a small part of a wider movement towards AI augmented software in the workplace.

This comes with many benefits. With increased AI adoption reportedly set to contribute up to €416 billion in additional GDP annually to the European economy by 2030 (McKinsey & Company, 2025b), European business leaders project a 15 percentage point rise in productivity and a 13 percentage point rise in revenues over the next 18 months due to generative AI (Accenture, 2025). Its direct impact on each employee is especially positive thanks to the potential to automate activities that currently consume 60% to 70% of employee time (McKinsey & Company, 2025b), especially useful for skilled workers who might be wasting their time doing repetitive tasks in which their skills are not fully exploited.

However, all of the positive recent data supporting AI integration in large companies and its benefits can not overshadow the fears of experts and members of the workforce about dangerous influences and volatility on the demand of jobs in this context, as well as their worries on the potential issues with AI agency over the workers' tasks and its likelihood of committing mistakes. Accenture reports that approximately 60% of European workers are worried about job displacement and the accuracy of AI outputs, while Mckinsey finds that the scale of automation potential thanks to AI contributes to worker anxiety (Accenture, 2025; McKinsey & Company, 2025b).

The primary subcategory of professional risk concerned by this case is **Income Security**. A large company structure has for long been synonymous with high *Earning predictability and consistency* as well as low *Time-to-security*, with overall good scores on each factor of this first subcategory, traditionally offering workers positions with stable compensation and the chance to look at the future confidently.

However, this status quo might be shifting: the intervention of AI in the workplace seems to create major risks on multiple of these factors.

First of all, *Displacement risk* can increase due to possible takeover of human positions by cheaper, more efficient AI-powered systems. Such claims are backed by research showing GenAI tools already matching if not surpassing human capacities in various benchmarks (International AI Safety Report, 2025).

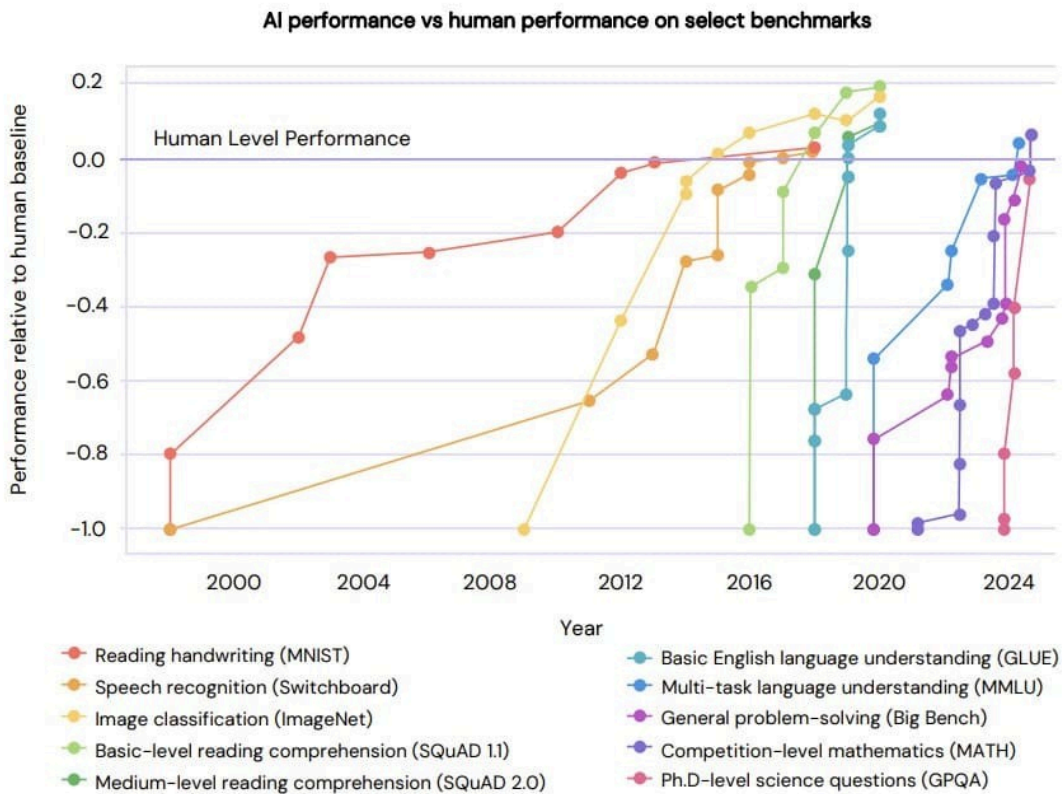


Figure 1: AI Performance Across Various Skills Over Time Compared to Human Ability  
 Source: *International AI safety report (International AI Safety Report, 2025)*

In practice, these tools are already able to outperform human doctors on medical diagnosis, increasing research output and boosting consultants' working speed and output quality (Centre for Future Generations, 2025). A survey by AI Impact reports that according to around 3000 researchers in the field, AI could perform most professional tasks that humans do before 2035. This information suggests rational incentives for companies, especially those operating at large scale and therefore able to afford and integrate such innovative technologies, to proactively explore the benefits of complementing their operation processes with AI tools, which may come at the disadvantage of human capital.

As mentioned previously, 60% of European workers are worried about job displacement (Accenture, 2025), which is not only caused by the potential of AI to be cheaper and less expensive than human workers, but also by *Skill and knowledge obsolescence risk*, a primary factor in large companies as updating their workers' skills to the AI-augmented workplace would require a very large amount of time and resources. Reports have calculated that 37%

of the European workforce lacks basic digital skills and 60% of EU firms identify a general lack of digital skills in the labor market, alighting fears that companies might find it more useful to fully replace some roles with AI instead of training their workers to complete the tools (Accenture, 2025). Although 73% of large businesses provide ICT training (Eurostat, 2025a), the scale at which European companies will need retraining services appears hard to fulfill, with reports estimating that 94 million European workers will require reskilling by 2030 (McKinsey Global Institute, 2024). Such fears already translate into reality. While a large part of the German Tiktok moderators went on strike in July 2025 amidst concerns of AI replacement in their positions, companies such as IBM have already paused the hiring process for thousands of back-office positions that they identified as “AI-replaceable”, while Duolingo is reportedly planning on “gradually stopping using contractors to do work that AI can handle” and “requiring teams to prove AI cannot perform a role before hiring humans”.<sup>11</sup>As PWC reported, one CEO in 4 is expecting to reduce human capital needs by 5% in 2026 thanks to AI integration (Centre for Future Generations, 2025). The Littler 2025 European Employer Survey Report (Littler, 2025) finds a basis to support this argument by polling qualitative data from more than 400 European employers and HR professionals. Amidst the highlights of this survey, the claim that more than 25% of respondents have either undertaken or is in the process of undertaking workforce reductions measures as a result of the integration of AI in the workplace, with another 25% planning to do so. Another phenomenon observed by these studies is tendencies for reduced hiring, which falls into the category of Earning predictability (unemployed members of the job market will likely have a more difficult time predicting their likelihood of finding a position in the short to middle term) and is a consequence of Skills and knowledge obsolescence (as the skills possessed by these individuals do not manage to keep their relevance over time through this AI development wave).

**Has your organisation taken any of the following actions as a result of the actual, or anticipated, increased use of AI in the workplace?**



Figure 2: *Little 2025 European Employer Survey Report (Little, 2025)*

Moreover, the role of skill obsolescence is increasingly important. With 37% (Accenture, 2025) of the European workforce already lacking basic digital skills to begin with, many worry about the decreasing relevance of their previous knowledge and expertise, as McKinsey finds that Generative AI has the technical potential to automate activities that currently consume 60% to 70% of employee time. The decisive factor in protecting these employees from being replaced here is reskilling and training, a process that appears too expensive and complicated to many large companies (and even more to SMBs). That is especially due to a lack of multidisciplinary options to train large groups of employees on their digital skills and their readiness to working alongside AI tools even though many of them have very different needs and require specific technical expertise in different fields which therefore complicate the training process (McKinsey & Company, 2025b).

*Career path predictability* is another factor to keep an eye on, as AI shapes company structures and traditional perceptions of career steps leaving space for a new, dynamic environment. The first evidence of such impact is the recent decline in entry-level job openings. Brynjolfsson et al. (2025) show that in AI-exposed occupations, employment for workers aged 22–25 falls by about 15–16 log points, while employment for more experienced workers in the same occupations remains stable. They argue that AI automates codified, entry-level tasks while complementing

tacit, experience-based tasks, so firms cut junior inflows rather than senior staff. This shows that the traditional progression from junior to mid-career rungs becomes less reliable in AI-exposed roles. At the same time, the ability of less experienced workers to transition between similar firms decreases, changing the status quo of horizontal options open to workers in their early-mid phase of their career, as they might be forced out of a position or wish to take a similar one at a different company, but now have less opportunities to do so. In the future, it is likely that such an impact will also influence more experienced workers, as AI integration goes further and AI capabilities cover more and more tasks.

As we move towards a job market in which the traditional junior to senior structure is transformed, young graduates cannot rely on entry-level jobs as much and experienced professionals worry about the future of their position and the increasing difficulties in transitioning to different companies or keeping up to date with new skill necessities, the workers' ability to predict their career path and approach their professional life with a sense of stability becomes a factor to worry about.

On the topic of Autonomy and Agency, GenAI tools not only increase worries of lack of human agency in the end result of work produced by any company, but also represent a risk in the psychological dimension of careers (European Parliamentary Research Service, 2025). AI augmented supervision and control over employees alongside increased reliance on technological tools and automations can decrease freedom in one's role and in some cases represent a threat in terms of privacy and satisfaction in a working environment. Excessive algorithmic surveillance can lead to decreases in the quality of a working environment, as demonstrated by Schlund & Zitek (2024): "algorithmic surveillance led participants to perceive they had less autonomy and to engage in more resistance behaviors, such as complaining more, performing worse, and intending to quit" compared to human oversight". As a study from the EPRS (European Parliamentary Research Service, 2025) has found, continuous monitoring alongside performance pressure from automated metrics and diminished human oversight can contribute to increased stress, burnout risk, and anxiety. This study also identifies regulatory gaps which

increase the need for tailored policies to tackle this issue.

### **A practical case study: Volkswagen and its “New Auto” Strategy. (Erdogan et al., 2024)**

Recent documentation on Volkswagen’s efforts to integrate AI in the workplace provides great insights on how the previous analysis translates in reality.

The German car manufacturer has entered a revolutionary phase of digitalization and AI integration to transform the company from a traditional vehicle manufacturer into a software-oriented provider of sustainable mobility.

Their main enabler is the VW Industrial Cloud, developed in partnership with Amazon Web Services (AWS). This cloud structure consolidates data from 124 factory sites and a network of over 1,500 suppliers into a single architecture, allowing for advanced data collection and processing which ends up translating into innovative industrial technologies such as data-driven production. For example, by leveraging IoT and machine learning, the cloud runs algorithms on data gathered from shop-floor sensors to drive predictive maintenance and optimize manufacturing processes. On the factory floor, this strategy has transformed VW’s production chain into a more efficient and reliable process.

- **Intelligent “CoBots”**: Collaborative robots powered by AI work alongside human operators now automate repetitive tasks and streamline assembly lines, resulting in higher productivity and reduced cycle times;
- **Error Reduction via Computer Vision**: AI systems are used for quality testing, where computer vision detects the “minutest defects” in components that might be missed by the human eye. It also ensures the accurate application of country-specific labels on vehicles to maintain regulatory compliance; and
- **Autonomous Logistics**: The company has developed fully autonomous mobile charging robots that can navigate confined spaces like underground garages to charge electric vehicles without human intervention.

In office spaces, a similar process has been ongoing with the same ambitions. Rather than replacing human roles, VW uses AI to augment human expertise, particularly in complex office-based functions:

- **Strategic Decision-Making:** AI models are used to identify patterns to support and empower employees in strategic decisions rather than supplanting their expertise;
- **Supply Chain Monitoring:** Using an algorithm called Prewave, VW scans public media and social networks in over 50 languages across 150 countries to identify sustainability risks, such as human rights abuses or environmental pollution, providing early warnings to supply chain managers; and
- **Financial and Marketing Operations:** AI-powered systems evaluate the probability of fraud in credit applications (via Actico) and optimize media-buying decisions in the German market through predictive analytics (via Blackwood Seven).

The impact of such strategies on the company's efficiency is predicted to be very positive, as the Industrial Cloud initiative aims for a 30% increase in productivity and a 30% decrease in factory costs, while the company anticipates potential savings of €1 billion in supply chain costs through AI-driven optimization.

However, the impact on the workforce is yet to be seen. In this particular case, the company truly appears to focus on creating an AI-augmented workforce which improves alongside AI tools instead of being replaced by them, however training issues and the comparative advantage of using AI tools instead of members of the workforce which skills might be obsolete in many fields still remain, and eventually companies will be confronted with situations in which it makes no financial sense to open new positions or keep current ones in sectors in which AI can do a better job for cheaper. To prepare for such situations, members of the workforce need to be exposed to training and reskilling as well as protected by regulations aiming to minimize human replacement by AI in fields at risk and to incentivize a smooth transition to an AI-augmented

workforce in which individuals from all backgrounds and demographics can have a place.

### 3.2 Small and Medium and Innovative Companies

Depending upon the size of the company, the impact of AI on its operations changes. Smaller companies run into similar risks (skill obsolescence, displacement risk) and opportunities (higher efficiency and cheaper operations) while often disposing of less resources to face such risks and to take advantage of such opportunities. However, the impact of each AI tool on a single SMB can be higher in proportion to the company’s size and ambitions compared to the impact it can have on a large company that would need to integrate many more tools to optimize, meaning potentially higher growth opportunities for SMBs once they manage to augment their employees’ work with AI. As of today, evidence shows small and medium businesses are using AI tools significantly less than large companies (Eurostat, 2025b).

**Enterprises using AI technologies by size class, EU, 2024 and 2025**  
(% of enterprises)

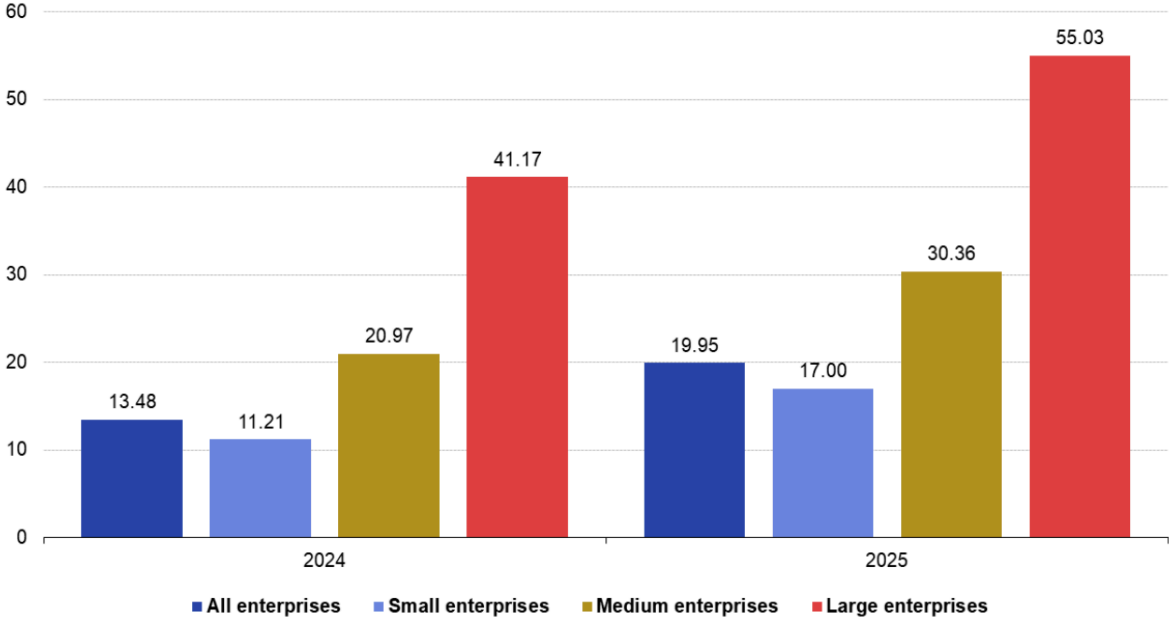


Figure 3: Eurostat “Use of artificial intelligence in enterprises” (Eurostat, 2025b)

This phenomenon is consistent with the digitalization gap shown by Eurostat (Eurostat, 2025a), as their research shows that only 58% of SMBs reach basic digital intensity vs 91% for large firms. The logical consequence is that SMBs start from a weaker digital baseline when they now

face the added requirement of AI-specific skills. A recent OECD report identifies skill shortages and financial constraints as major barriers to AI adoption by SMEs, more so than for large firms (OECD, 2025a).

As explained in §3.1, an AI augmented workplace stresses the need for training and upskilling in order to adapt each worker's skillset to these new technologies, especially in tasks such as prompting, supervising automated tasks, and the setting or monitoring of AI agents. For large companies, higher capital and resources allow for easier planning and organization of such programs, however for SMBs this might not be the case. While large firms already show higher AI adoption rates and can amortize the fixed costs of internal academies, tailored AI leadership and multi-year training programmes, SMBs often lack budget, HR capacity and information to identify their AI skills needs and find suitable training, leading to a widening capability gap.

### **The Risks:**

The main risk for SMBs in this context would be not being able to catch up quickly enough. SMBs already face efficiency gaps in operations like supply chains, marketing, and admin. As large firms use AI for predictive analytics, automation, and real-time decisions, this gap could become larger. Persistent adoption lags could undermine SMB market share and reinforce the winner-takes-all dynamics of AI. As highlighted by Babina et al (2024), AI productivity gains are often disproportionately distributed to favor large, fast growing companies, partly because smaller firms face delays in adopting and integrating these technologies.. Their research suggests that businesses who fail to catch the "AI wave" will drift further behind, especially relevant in the case of SMBs, as large companies in leading positions in AI integration are already showing up to 5% relative revenue growth compared to competitors with less resources. The paper highlights the following reasons:

- **Fixed Cost Spreading:** large companies can more easily spread the high fixed costs of AI investment over higher revenues;

- **Data Productivity:** larger firms are able to deploy AI more productively because they possess or are able to purchase larger datasets, which are essential for training and refining models; and
- **Management Expertise:** larger organizations can count on larger and more skilled management to navigate the complex organizational changes required for AI integration.

In case of loss of market share due to lower AI adaptation, workers in SMBs would have increased professional risk in many categories, adding up to the risk increases in common with those explained in the large companies section, as:

- **Earning predictability and consistency**, as instability could hit their firms, potentially because of lost customers to larger AI-ready competitors. This is especially true for micro/small enterprises (less than 10 or 50 employees), as any volatility in costs and incomes can affect each employee more than for larger companies; and
- **Displacement risk**, as a consequence of potential downsizing. However, so far research shows that AI in SMEs is used more to ease constraints and fill skill gaps than to shed workers. In fact, the OECD17 found that among SMEs using generative AI, 39% say it helps compensate for skill gaps, and roughly one third report reduced staff workload, while 83% report no change in overall staff needs;
- **Career path predictability**, as a consequence of shocks in the structure of the company itself as well as potentially less opportunities to change firms, as seen in Part 1 for large companies. Careers in SMEs can be impacted in more specific ways, as in small firms with few rungs, any decrease in junior hiring or elimination of a role can collapse an entire internal ladder (as there often are no lateral departments to absorb workers), making paths more fragile than in large multinationals; and
- **Skill Obsolescence** remains a common risk with the previous part of the analysis. As discussed in §3.1, resources available for reskilling programs are crucial at this stage, and SMEs could suffer from lack of resources in this case, as well as needing to spend more on each employee's training, simply as large companies can create tailored training programs to apply to large number of employees working in similar positions, therefore benefitting from

lower average costs per employee, which does not apply to smaller companies which may need to provide tailored training only applying to a few workers.

When looking at Small and Medium Businesses, a disclaimer needs to be made. Very often, the reason for which they might not be using Generative AI, is that it simply is not suited to what they do (OECD, 2025b). This is especially true for sectors such as agriculture, manufacturing, mining, and construction, as they rely on physical abilities outside the (current) capabilities of generative AI, but it is also reported as a barrier in some service sectors, such as finance and professional, scientific and technical activities. Therefore, when looking at data on AI adoption we always need to remember that not all businesses actually need such tools, even though it might not seem so when looking at how quickly this industry is growing.

However, this is not an issue for Innovative Companies (Startups and Scaleups). They tend to be made of generally younger members, therefore more digitally literate than the average in other companies, and often build digital, technological products or services which positions them more clearly on the edge of the AI wave. This translates into generally higher adoption of AI18 in their processes and in their products.

Most of these risks are common to both large companies and SMEs, but what really stands out in the case of SMEs and Innovative Companies are the opportunities. In the next paragraphs, we will explore how the many opportunities for growth and for risk reduction AI can bring to individuals working in small teams and simple company structures, in which optimizing solid operations with limited resources is often crucial.

In the case of AI being an opportunity for economic growth in companies, its benefits in large companies are likely to transform into higher revenue and shareholder value (Accenture, 2025), with low benefits for the employees itself due to the company's vertical structure. However, in small and medium enterprises, more simple and horizontal company structures can lead to such benefits impacting each worker's professional life at a larger scale.

The main opportunities that can arise are those that allow SMEs to face capability gaps, one of their most recurrent problems. Due to more simple company structures, reduced workforce compared to large firms and lack of resources to train their workers in a complete, continuous way, these companies constantly face lower readiness to execute certain tasks and to optimize their operations compared to their bigger competitors.

Especially in matters of digitalization, SMEs often find themselves not having the right people and skills to undertake certain processes, as their resources and human capital are reduced and less polyvalent compared to large firms (Eurostat, 2025a). Eurostat shows that while 73 % of EU SMEs have at least a basic level of digital intensity (Eurostat, 2025a), only 6% of them reach a very high level, compared with 41% of large firms, showing a broad capability gap in advanced digital tools. This issue is also endogenous to the topic of AI integration itself, as an Eurostat survey found that among firms which had considered AI but decided not to adopt it, lack of relevant expertise was the most cited reason (Eurostat, 2025b). Cedefop finds that managers and staff in many SMEs “frequently lack the digital and strategic skills needed to integrate new technologies effectively” listing skills shortages as one of the main barriers to SME digitalization (Cedefop, 2025). This can be the case in areas such as software development, internal automation, digital organization, and integration of digital tools such as CRMs or Cloud storage spaces.


Other than for digital matters, capability gaps have high impact on SMEs in many other fields, with the legal and fiscal compliance ones standing out.

Legal capability gaps are very common, given that SMEs rarely have a proper legal department of their own to count on and have to rely on external firms and cabinets to ensure compliance. This translates into higher legal risks and costs (European Commission, 2015), as they incur higher per-unit legal burdens because they cannot spread specific compliance costs across large unit volumes and often outsource to consultants for EU law transposition and enforcement. The EU’s lack of legal harmonization between national and supranational

frameworks contributes to this issue, creating grey areas on which national interpretation of EU laws makes it even more difficult to comply, forcing SMEs to incur even higher risks and expenditure on external advice (European Commission, 2015). This issue also impacts cross-border growth of SMEs, as per a recent Eurobarometer Survey, 33% of SMEs cite understanding regulatory differences as the top cross-border barrier, with 30% struggling with access to information on rules and requirements (Eurobarometer, 2025).

This same issue is relevant for fiscal matters. An European Commission (referred to as Commission from here onwards) Report finds that in EU countries SMEs can face exponentially higher tax compliance costs, the ratio between total tax-related compliance costs and paid taxes, of 30.1% for micro enterprises, while large companies have a cost to tax revenue ratio of only 5.9% (European Commission, 2022).

**Figure 28: Estimated average TETCC to tax revenue ratio (R3) by size and sector**



*N=2 201. The figures indicate the estimated average ratio for the 28 countries in each category. Confidence intervals (95%) are indicated by the black brackets. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1. Source: VVA / KPMG (2021)*

Figure 4: European Commission “Tax compliance costs for SMEs: An update and a complement Final Report” (European Commission, 2022)

The Commission’s estimated tax compliance costs EU-wide burden of €204bn per year is reported to weigh disproportionately more on SMEs due to lack of internal compliance systems

and difficulties in minimizing costs (European Commission, 2022). Taxation issues being an important barrier to expansion for EU SMEs, compliance risks increase with cross-border operations just as legal risks, considering different tax regimes and collection mechanisms between EU countries.

Going back to the professional risk framework, we can see that many of its points are touched by the need for filling capability gaps. All of the **Legal and Fiscal Risk** is of course concerned, alongside **Income Security** which is clearly affected by the potential costs of capability gaps and their impact on potential income, as they can be the cause of suboptimal efficiency.

Solving capability gaps in digitalization can help SMEs reach higher efficiency thanks to digital tools and therefore potentially more revenue and lower costs, while on the legal and fiscal sides it would decrease the need for outsourcing and harmonize the company's legal and fiscal operations, leading to less legal risks and compliance costs.

AI integration in these companies, if done well, can be an important tool to help close these gaps. On the digital side, AI tools can help advanced skills towards non-technical users of digital platforms (Eurostat, 2025b) thanks to automations and applications of human-like conversational input technologies that can then achieve technical tasks, such as vibe coding. The integration of AI assistants in tools such as CRMs or Cloud managers can also allow non-technical workers to fully navigate such platforms without needing to be trained on very advanced usage cases for these tools. This can also be achieved, in the future, by agentic technologies that we can already find in products such as Claude Code or Cowork and more controversial ones like ClawdBot. These AI agents can access many of the digital platforms that an employee might struggle with and complete even advanced tasks just by following the employee's input through a conversational tool. The need for humans in the loop frameworks and training to use such advanced and unsupervised tools is of course an issue. An OECD survey points towards GenAI helping closing skill gaps as the main reason for AI driven performance increases among SMEs, as among SMEs that have seen positive effects from using such tools 46% indicated

that GenAI is particularly beneficial where a lack of skills is the main constraint for employee performance.

### 3.3 Entrepreneurship

Entrepreneurship is probably the clearest professional category in terms of risk perception. Traditionally seen as very risky because of the lack of safety nets and the high responsibility of starting from scratch without relying on a solid company structure, it features high levels of risk in terms of **Earning Predictability, Time-to Security, Safety Nets, Asset Ownership, Legal and Fiscal Risk and Career path predictability**, as well as many other factors dictated by its uncertain nature and unpredictability of outcomes.

Its risks in the current AI professional revolution are similar to those of SMEs, as they both have less resources to compete with large companies and can therefore risk of losing opportunities to stay alive in markets in which productivity and profitability gaps may widen due to the comparative advantage that large companies can achieve thanks to the economies of scale of AI integration.

However, the introduction of GenAI tools for entrepreneurs, whether they want to create a tech startup, a digital agency or even a more traditional, non-digital company in fields such as manufacturing, trade or education, can decrease risk in many factors if implemented well. Logically, decreasing risks results in creating opportunity.

The first factor to look at is **Income security**. The uncertainty of cost and revenue structures are probably one of the main barriers to entrepreneurship. Heavy starting costs combined with barriers to entry in many markets as well as the need for professional recognition even before reaching the first few sources of revenue all contribute to the traditional perception of entrepreneurship being a high-risk career. AI can make early-stage entrepreneurship less financially volatile by different mechanisms such as shrinking fixed costs, improving test and validation

speed as well as increasing odds of product-market fit (Cai et al., 2025).

The first main mechanism is allowing for lower cost to test and iterate:

- No-code tools using GenAI to simulate more technical tasks allow entrepreneurs to create prototypes of digital products, marketing campaigns and basic automations that previously would have required a full team, cutting needs for burning cash upfront and lowering operational times.
- This reduces Time-to-security because entrepreneurs can test multiple ideas quickly instead of spending months on a single one to then have to pivot later on.

Secondly, it allows for better odds of revenue per attempt:

- Very small teams can run sophisticated testing strategies, audience research and outreach campaigns faster and easier than ever thanks to AI agents and marketing tools which allow for easier access to information and easier creation of essential datasets to figure out efficient strategies that traditionally require more time and resources to be reached.
- That increases the chances that at least one of the strategies tested reaches potential for bringing in vital revenue to early-stage funding of an entrepreneurial project, which lowers the burden on the entrepreneurs' financial responsibility and risk.

GenAI tools also decrease the need for outsourcing of tasks throughout the whole entrepreneurial journey, decreasing the costs of hiring external consultants or freelancers to do specific tasks that can now be done by AI agents. These tasks include, for example, recording financial statements, organizing and managing data through CRM tools, and tracking numbers for relevant insights on outreach performance and efficiency of advertisement.

The second factor to consider is **Social Security and Downside Protection**. Entrepreneurs are still exposed significantly, as they often have their own assets on the line and can not benefit from wealthy company structures to provide financial cushions for them. They remain more responsible for contributing to savings, private insurance and large parts of their assets are at risk of being involved in the business to cover additional costs.

However, AI integration can contribute in lowering such risks (Cai et al., 2025):

- Lower operating costs (due to usage of AI tools in infrastructure, administrative work and staff management instead of more expensive options such as external partners in the early stages) mean a larger share of early revenue can go into personal buffers rather than covering costs.
- AI integration can also decrease the difficulty of starting a micro-business which, requiring less time in early stages, allows people to reach a solid proof-of-concept and a plan on the side of working a full-time job, maintaining a more stable position before fully jumping into a personal project.
- Reduced needs for assets (as offices become less necessary and overall outsourcing of inventory and fabrication facilities becomes easier thanks to digitalization) means lower risks are taken by the entrepreneur, as the downside in case of failure becomes less heavy.

**Legal Risks** are also impacted, as AI integration can lower the costs and resources needed to reach basic levels of compliance through:

- Easier access to information, as LLMs and specialized platforms can resume large legal documents, extracting essential information without the need of hiring a third party just to go through a certain regulation framework; and
- Specific B2B services using AI to draft legal documentation such as contracts, statutes and compliance forms. Such platforms can also help fill tax forms and other procedures lowering the need for external legal spend on simple to more difficult procedures.

AI can make a larger legal toolbox available to entrepreneurs in a more immediate, less expensive way than what it would cost them usually, lowering the unit cost of being compliant (International Monetary Fund, 2025). However, it is important to highlight that these advancements come with uncertainties about data protection and accuracy risks, increasing the need for human in the box frameworks and responsible usage of new tools.

Lastly, AI's impact on **Autonomy and Agency** needs to be considered (Agarwal & others, 2025). In the case of AI augmented human work, entrepreneurs can hold decision authority

while outsourcing research work for analysis and forecasting to AI tools, allowing them to maintain their decisional autonomy and agency over their work. Allowing AI solutions to take over repetitive tasks such as customer support and basic operations can allow them to focus on more creative, idea-related tasks which can reduce stress and increase professional satisfaction as they dispose of more time and resources to focus on developing more interesting parts of their project which are likely those that made them want to create it in the first place.

However, this requires them to be conscious of the potential risks of over-relying on AI. The line between AI-augmented work and cognitive offload (Forbes, 2025) to AI is very thin and educational frameworks on the topic are needed to ensure that entrepreneurs are ready to integrate innovative tools without risking negative downsides. For this purpose, frameworks such as the AI Literacy Framework for Primary and Secondary Education created by EU institutions alongside the OECD are fundamentals, especially to train the next generations of professionals, but efforts must be made to create more precise frameworks targeting existing professionals. Even if it meant reducing the scope of such a project to counsel on responsible usage of short-term AI implementation, it would still be crucial as the way AI gets adopted in the next few years can dictate the habits of the decades to come.

We will now move onto the policy proposal section of the report, where we will translate the above-reported professional risk framework, income security, social security & downside protection, legal & fiscal risk, and autonomy & agency, into a practical EU policy scoreboard for assessing where AI is improving working lives and where it is instead amplifying downside risks. Our earlier analysis shows that AI is reshaping the traditional employee-entrepreneur tradeoff, lowering barriers to entrepreneurship and also raising productivity, while creating new forms of volatility in job demand, career ladders, and workplace governance.

We have seen that risk exposure is uneven across firm types. Within large companies, pressures cluster around income security through displacement and skill obsolescence, calling for reskilling at scale. Additionally, there are early signals of reduced entry-level opportunities among AI-exposed roles. Meanwhile, SMEs face the critical challenge of a widening capability gap: where lower AI adoption and weaker digital baselines interact with disproportionately high legal and fiscal compliance burdens. On top of these factors, there remains fragmented implementation and interpretation across MS.

Against this landscape, we consider the EU's policy task to be to reduce downside risk while ensuring that adoption isn't halted. Our proposal aims to (i) strengthen transition and reskilling support, (ii) close portability and coverage gaps in social protection, (iii) set enforceable safeguards for algorithmic management and workplace surveillance, and (iv) lower compliance barriers for SMEs through harmonized guidance, tools, and enforcement coordination. The proposal will be supported by a monitoring framework with clear KPIs that are aligned to the four risk dimensions.

## **4 From “Risk Framework” to “Policy scoreboard”**

### **4.1 What the scoreboard is**

The **Professional Risk Scoreboard** is a practical way to evaluate the AI-in-the-workforce policies through a simple question: do they reduce or reallocate downside risks faced by workers

and firms? Our goal is to avoid treating “AI impact” as a one-dimensional question of jobs gained or lost, and the scoreboard operationalizes the four risk dimensions from earlier which together shape the employee-entrepreneur tradeoff:

**1. Income Security:** the way that AI affects earnings stability, displacement risk, and the speed and quality of transitions into new tasks and roles;

**2. Social Security and Downside Protection:** whether protections (unemployment, sickness, pensions) remain continuous and portable across job changes across employment statuses. A key reference point in this regard is the EU’s own emphasis on access to social protection for workers and the self-employed;

**3. Legal and Fiscal Risk:** whether accountability, compliance expectations, and redress routes are clear enough to be workable, especially when new AI obligations interact with fragmented national implementation; and

**4. Autonomy and Agency:** whether AI-enabled management practices expand productivity while simultaneously respecting worker autonomy. This point is especially focused on monitoring, scheduling, performance evaluation, and contestability. We have a concrete template for organizational duties rather than “rights on paper” alone through the Platform Work Directive’s approach to transparency and human oversight in algorithmic management.

It is important to point out that through the scoreboard we aren’t trying to “rank” MS. Its purpose is more diagnostic and decision-oriented: making trade-offs visible, helping to identify where risks are concentrated (by sector, firm size, or country), and forcing policy proposals to specify which risks they are trying to reduce, which ones they may increase, and how these eventual side-effects are mitigated. For example, a policy may accelerate SME adoption by simplifying compliance and improving legal-risk outcomes and competitiveness, but it would still require guardrails to ensure that cost-cutting doesn’t translate into intensified monitoring or unmanaged psychosocial strain. Through this framing, we set up our next step, identifying why workplace AI can’t be governed through individual rights alone, and the reasons why

implementation capacity, enforcement, and collective/organizational duties should be kept in mind.

## **4.2 Why individual rights aren't enough in workplaces**

In principle, individuals are already protected around data use and automated decision-making scenarios thanks to EU law. However, in practice, workplaces are a structurally “non-ideal” environment for rights that rely only on individual action, and AI makes that gap bigger.

First, we have a power imbalance between employer and worker that limits the effectiveness of “notice-and-consent” models. Under the GDPR, consent isn't a valid legal basis in situations where we have a clear imbalance between the individual and the controller, and EU guidance explicitly highlights the employer-employee relationship as a typical example (European Commission, 2016; GDPR, 2016). The EDPB's consent guidelines reinforce that consent is invalid where people have no genuine choice or would face disadvantage from refusing or withdrawing, conditions which can easily arise when monitoring and analytics are embedded into routine management (EDPB, 2020).

Second, AI-driven decisions amplify information asymmetries and transaction costs in ways that are especially difficult for individual rights to overcome. Even where a worker formally has a right to contest a decision, being able to exercise that right is only possible if the decision is understood, under what basis, with what consequences, and through which channels. These are evaluations that often occur under time pressure and with limited access to internal documentation. The EU-level guidance endorsed by the EDPB on automated decision-making and profiling stresses that safeguards depend on effective transparency and meaningful human involvement, not merely the existence of rights in the abstract (EDPB, 2018).

Third, workplace AI often produces collective effects which are poorly addressed by case-by-case individual remedies. Algorithmic scheduling, monitoring policies, performance scoring, or AI-supported restructuring are able to reshape work intensity, allocation of opportunities, and

promotions pathways across teams. Exactly for that reason, EU labour law has long relied on organizational and collective mechanisms (such as information and consultation frameworks) rather than leaving governance to individualized disputes after harm occurs (Directive 2002/14/EC).

Finally, many of the most salient risks are health and safety risks, including psychosocial strain and intensified work. It is clear that a prevention model here is more appropriate than after-the-fact redress. Employers have duties to assess risks, adopt protective measures, and enable worker participation. This approach maps naturally onto systematic AI-enabled management practices (Directive 89/391/EEC).

For the above reasons, the policy approach in this report will treat individual rights as a necessary backdrop, but prioritizes organizational duties, worker voice, enforceable processes, and implementation capacity as the primary tools for governing workplace AI.

### **4.3 Why the EU level matters**

EU-level action is of the utmost importance because workplace AI is being increasingly shaped by cross-border markets and shared compliance infrastructures, while the most important protections which are crucial for workers are still being implemented through national labour and social systems. In light of this, we get a predictable policy failure. Rules and standards that may look coherent on paper can actually become fragmented in practice. This increases compliance costs and creates uneven protection across MS, especially for firms that operate across borders or procure AI-enabled HR and management tools internationally. The EU has a comparative advantage, reducing this fragmentation by setting a common governance spine for AI risk management and accountability, anchored in internal-market harmonization and fundamental-rights protection.

A second EU-level rationale is one that targets capacity and convergence. While adoption is rising, it is an uneven phenomena across firm sizes and sectors, meaning that the same regulatory

and organizational demands can produce different burdens and outcomes. At the same time, the Digital Decade governance framework explicitly treats SME digitalization as a strategic objective, with a focus on a convergence agenda rather than laissez-fair divergence (Decision (EU) 2022/2481). In practice, Eurostat’s Digital Intensity reporting illustrates how far many SMEs remain from those ambitions, reinforcing that “compliance-by-paper” without capability support risks widening gaps rather than managing risks (Eurostat, 2024).

Third, EU action matters because workplace AI risks aren’t just individual but also organizational and systemic. As a result, implementation architecture (like guidance, templates, enforcement capacity, and cross-border consistency) is at least as important as formal rights. For this purpose, the EU has existing building blocks which can be scaled beyond their immediate scope. For example, the AI Act already provides a risk-based compliance logic for systems that can affect fundamental rights, including employment-related contexts (AI Act, 2024). Meanwhile, the Platform Work Directive offers a more concrete template for organizational duties in relation to algorithmic management, which can be used to inform broader good practice beyond platform work.

Finally, EU-level action is crucial to connect the standards of delivery of the policies. Even when MS retain primary competence over labour-market institutions, the EU has the ability to make policies implementable by: (i) producing common compliance toolkits and “minimum documentation sets” (especially for SMEs), (ii) coordinating enforcement approaches to avoid diverging interpretations, and (iii) scaling transition and skills investments through existing funding instruments and reporting cycles (Regulation (EU) 2021/1057). These levers are exactly how downside risks can be reduced without freezing adoption, one of the main objectives which follows from the report’s risk-scoreboard framing.

## 5 What Does Good Look Like?

In order to use the four risk dimensions as a policy scoreboard, we need to define what “improvement” means in each category. Through this we aren’t trying to prescribe a single EU labour-market model, but to set directional targets compatible with national systems while still being comparable across countries, sectors, and firm sizes. In practical terms, “good” performance means that AI adoption can proceed without shifting hidden costs onto workers (through volatility, gaps in protection, unaccountable decisions, or intensified work), and without creating a compliance environment which is only navigable by large firms.

### 5.1 Income Insecurity: Transitions that are fast, fair, and do not scar earnings

A “good score” on this metric means that AI-driven task change doesn’t translate into prolonged unemployment, repeated involuntary transitions, or persistent wage penalties. Instead, labour markets should be able to support job-to-job mobility and fast re-employment speed, using training and matching systems meant to make transitions routine rather than exceptional.

Concretely, this metric has three main elements:

**1. Shorter and less unequal transition times.** A well-performing system limits the duration of unemployment spells following displacement, and prevents long-term unemployment from concentrating in specific groups or regions (Eurostat, 2025c).

**2. Lower wage scarring after transition.** Re-employment shouldn’t systematically mean a step down in earnings or job quality; policies should be judged on the quality of reattachment to work, not just its existence.

**3. Observable reallocation capacity.** Because AI is often a “task reshuffling” technology, our scoreboard should track labour-market flows and transition probabilities as a core indicator family, using existing statistical infrastructure where possible (Eurostat, 2026a).

These identified elements are compatible with the evidence base that AI's labour market effects are frequently mediated through changing tasks and work organization rather than immediate headline job losses, making transition capacity central to welfare outcome conversations (OECD, 2023).

## **5.2 Social Security and Downside Protection: Coverage that follows workers across status and shocks**

Here, a “good score” means that people don't face large, discontinuous jumps in risk when they change jobs, working hours, or employment status. The EU's own policy benchmark emphasizes access to social protection for workers and the self-employed, our scoreboard aims to translate that benchmark into measurable continuity outcomes (Council Recommendation 2019/C 387/01).

Concretely, the target state is:

**1. Coverage continuity during transitions.** Switching jobs or status shouldn't create “benefit cliffs”.

**2. Portability and transparency.** Entitlements should be trackable and transferable, that way workers can make informed choices and firms are able to operate with predictable obligations.

**3. Measurable gaps shrink over time.** We aim to track social-protection and income-security outcomes using harmonized EU data sources, with explicit identification of what can be measured now and what still needs improved data. EU-SILC provides the methodological backbone for comparable indicators on income, poverty, and living conditions that can proxy the effectiveness of downside protection (Eurostat, 2026b).

### **5.3 Legal and Fiscal Risk: accountability that is workable, predictable, and contestable**

We can define a “good score” on the legal and fiscal risk as workplace AI that is deployed within a governance environment where obligations are clear to follow, proportionate enough to be feasible for SMEs, and credible enough to provide real redress where systems can cause harm. This target state has two simultaneous requirements:

**1. Clarity and predictability for firms.** Compliance obligations should be clear, stable, and feasible to implement without the reliance on costly case-by-case legal advice. In practical terms, firms should be able to follow a standard minimum documentation and risk-management process, especially in common workplace AI use cases.

**2. Contestability and accountability for workers.** Decisions that may affect access to employment, working conditions, or advancement must be reviewable through accessible channels, with meaningful safeguards beyond formal rights.

The AI Act provides support for the “predictable and accountable” target state. Policies should be judged by whether they make that governance operable in workplaces, not only by whether new rules exist (AI Act, 2024).

In parallel, GDPR-era guidance on automated decision-making and profiling sets out a logic whereby safeguards are dependent on effective transparency and meaningful human involvement. This factor is helpful to translate “contestability” into a measurable policy objective (EDPB, 2018).

### **5.4 Autonomy and Agency: productivity without pervasive surveillance or psychological harm**

A “good score” here means that AI-enabled management doesn’t silently trade productivity for deteriorating job quality. The target state isn’t “no monitoring”, but instead a bounded

and justified monitoring, predictable working time and scheduling, and active management of psychological risks.

Concretely, the target state includes:

**1. Prevention-based risk management for psychological harm.** Occupational safety and health law already reflect a prevention model. The scoreboard will treat algorithmic management as a potential workplace risk factor that must be assessed and mitigated (Directive 89/391/EEC).

**2. Working-time quality and predictability.** Wherever algorithmic scheduling affects volatility and intensity, “good performance” means that compliance and outcomes are aligned with EU working-time standards and predictability expectations (Directive 2003/88/EC; Directive (EU) 2019/1152).

**3. Measurable job-quality outcomes.** The scoreboard will lean on existing pan-European survey infrastructure to measure work intensity, discretion, working time quality, and related job-quality dimensions. Eurofound’s European Working Conditions Survey (EWCS) provides a long-running basis for comparable monitoring of working conditions across Europe (Eurofound, 2024).

**4. Establishment-level readiness and practice.** EU-OSHA’s European Survey of Enterprises on New and Emerging Risks (ESENER) adds a complementary employer/establishment lens on the way that workplaces manage safety and health risks in practice, including psychological risks and worker participation (EU-OSHA, 2024).

With these “good score” definitions in place, we can move onto mapping what the EU already has, hard law, guidance, funding, and enforcement infrastructure, against each target state, and to identify what implementation gaps are present which prevent policy from improving the scoreboard.

## **6 Current EU Policy Landscape**

Building upon the professional risk framework, this section will look at which existing instruments at an EU level already address each dimension and where gaps in implementation, coverage and capabilities remain in practice. It will first look at each dimension separately and later analyse highlight where rules “stop” in practice due to fragmentation, grey areas or capacity bottlenecks, especially for SMEs.

### **6.1 Income security**

- Coordination of social security via Regulations 883/2004 and 987/2009 allows aggregation of contribution periods across all MS, helps protect earnings stability when workers move or are displaced.
- Reskilling and upskilling programs are already addressed by regulations such as ESF+ (Regulation 2021/1057) and RRF (Regulation 2021/241) They also cover active labour market policies that can cushion job or role changes linked to automation.
- The European Globalization Adjustment Fund (EGF, Regulation 2021/691) can support workers affected by major restructuring, including technology-driven mass redundancies.

### **6.2 Social security and downside protection**

- Coordination of unemployment, sickness and pension benefits via EU social security regulations ensures that mobile workers do not lose social security rights when moving to another country or changing status.
- The Work–Life Balance Directive (2019/1158) and other directives create the basic minimum standards around the topics of leave and care that help smooth transitions and labor shocks.
- The previously mentioned ESF+ regulation can finance income support programs and complementary safety nets during training and job transitions.

### **6.3 Legal and Fiscal risks**

- The AI Act has brought a necessary new framework for regulating the usage and integration of AI technologies. This also applies to professional environments, therefore adding a level of legal standards that was not existing before.
- Tools to facilitate legal and fiscal compliance are particularly complex to put in place due to the fragmentation of policies and requirements between different states.

### **6.4 Autonomy and agency**

- Work-life balance issues are addressed by tools such as the Working Time Directive (2003/88/EC) and the Transparent and Predictable Working Conditions Directive (2019/1152), although there are not yet any similar directives which can relate to the special characteristics of an AI augmented workplace.
- The OSH Framework Directive (89/391/EEC) addresses health and safety conditions including psychological sides of the topic. In the case of our topic, it can relate to the psychological risks and benefits of using AI, such as for example the negative effects of AI-driven monitoring or on the other hand the increased creativity due to the automation of purely technical tasks.
- The Platform Work Directive (EU 2024/2831) established a base amount of individual and collective rights which can often be useful to guarantee agency and autonomy in the future workplace. It backs explanation rights and human review safeguards against fully automated decisions, including a number of advanced algorithmic management (AM) safeguards.

### **6.5 Where do these tools stop?**

Such tools can only address their related issues until a certain point. Reasons for their limited reach and their inefficiency in some cases include:

- Fragmentation, as national transposition creates grey areas, especially for cross-border digital services.
- Coverage gaps for non-standard workers, as self-employed and micro-entrepreneurs may often fall between regimes on social protection and have weaker autonomy protections.
- Rules often are technology-neutral. Labour inspectorates and work councils struggle to translate existing OSH or working conditions directives into concrete limits on AI surveillance without specific guidance. As mentioned in the case of the Working Time Directive (2003/88/EC) and the Transparent and Predictable Working Conditions Directive (2019/1152), such tools still lack methods to be adapted to the future technology driven issues in an AI-augmented workplace.

## **7 Risk Patterns Across Firm Types**

The above-defined four-dimensional risk framework provides a common language to evaluate how AI affects professional life. However, these risks don't materialize in the same way across all firms. The effects of AI aren't only dependent on the technology itself, but also on firm size, organizational capacity, internal labour-market structure, and the ability to absorb regulatory and adjustment costs.

This distinction across firm types matters for policy as a uniform approach risks solving the wrong problem for the wrong actors. On the one hand, large firms are generally better placed to adopt AI, reorganize workflows, and manage compliance internally. But that very capacity makes them more likely to generate large-scale transition pressures, including displacement, role design, and changes to career ladders. On the other hand, smaller firms are less exposed to immediate large-scale restructuring, yet face a different risk configuration: weaker digital capabilities, lower AI adoption readiness, and proportionately heavier legal and compliance burdens.

For the above reasons, this chapter aims to translate the general risk framework into a firm-type diagnosis. Its purpose isn't to repeat the landscape analysis, it is trying to identify which risk dimensions become most salient in different organizational settings, and as a result the policy levers which are most important. Below we will argue that large firms are primarily a managed-transition problem, while SMEs are primarily capability and compliance problems. Making this distinction is essential for the recommendations which will follow, as it displays why “one-size-fits-all” policy design would either under-protect workers in large organizations, or overburden smaller firms that are already struggling to adopt and comply.

## **7.1 Large Firms:**

In large firms, the main AI-related pressures don't come from a lack of adoption capacity, but from the scale and speed at which adoption can reshape internal labour demand. Large organizations are typically better positioned than SMEs to invest in AI tools, redesign workflows, and absorb compliance costs. However, this greater capacity doesn't eliminate professional risk. Instead, it works to shift the center of gravity of risk toward income security, internal labour-market restructuring and in some cases autonomy and agency.

The first and most visible risk is that of income security through displacement or task compression. In large firms, AI is often introduced not as a marginal efficiency tool, but as part of a wider process of optimization. This implies that it can alter staffing needs across entire functions rather than individual tasks in isolation. As such, the policy concern is not only job loss in a narrow sense, but the broader problem of how quickly workers can be reallocated, retrained, or reabsorbed into roles of comparable quality. This is a policy concern because, wherever firms automate selectively but at scale, the main welfare effect can come less from immediate unemployment than from repeated role redesign, stalled progression, or downward pressure on entry pathways.

A second large-firm risk is the weakening of traditional career ladders, especially for junior and routine-intermediary roles. Wherever AI systems substitute for portions of analytical, administrative, or support work which previously functioned as training grounds for younger workers, firms may preserve output while reducing the number of entry-level opportunities through which skills and institutional knowledge are normally built.

A third risk is that of the quality of reorganization inside firms. Large firms are more likely to have the managerial and technical capacity needed to integrate AI into performance management, workflow allocation, and monitoring systems. While this may create gains in coordination and productivity, it can also intensify work, narrow discretion, and increase the opacity of managerial decisions in the case that safeguards are weak. In other words, large-firm risks aren't just about "jobs disappearing", they are also about whether AI-supported reorganization leads to a more adaptive workplace or to a more tightly monitored and less contestable one.

Taking the above together, we can see that the large-firm policy problem is best understood as a question of managed transition. Large firms generally have more capacity than SMEs to comply, invest, and reorganize, meaning that the main policy priority is not basic adoption support but ensuring that reorganization is accompanied by internal transition planning, reskilling, worker information and consultation, and safeguards around AI-enabled management practices. Shifting to the scoreboard, large firms are the context in which income security and autonomy and agency risks become especially salient, even when legal-compliance capacity is relatively stronger.

## **7.2 SMEs:**

Moving onto SMEs, the main AI-related risk pattern differs from that of large firms. The main issue here is not that of large-scale internal restructuring, but more whether smaller firms have the capability, confidence, and support needed to adopt AI productively and lawfully in the first place. Within this context, professional risk is shaped less by immediate mass displacement

and more by interaction between lower adoption readiness, weaker organizational slack, and disproportionately heavy legal and compliance burdens.

The first risk we look at is not simply one of “lower adoption”, but greater vulnerability to failed or partial adoption. In smaller firms, even modest investments in AI are difficult to absorb if they require complementary changes in workflow design, staff training, procurement, data practices, or legal review. These dynamics which may appear manageable for large organizations may become disruptive in a small one because the same adjustment must be handled by fewer people, less internal expertise, and less room for experimentation. From a policy perspective, the focus is shifted away from AI uptake as such and toward the question of whether SMEs have the institutional support needed to adopt productively rather than superficially.

A second SME-specific risk is that of fixed-cost exposure. SMEs are less able to spread the costs of compliance, due diligence, training, or advisory support across a large revenue base or dedicated support functions. Even when formal legal obligations are identical across firm types, their practical burden isn't. In effect, SMEs face a more fragile cost-benefit calculation. Uncertainty over what counts as adequate compliance, documentation, or governance may be enough to delay adoption, narrow the range of tools considered, or push firms toward informal workarounds. So, we reach a problem that isn't purely based on regulation itself, but the absence of usable, standardized, low-friction pathways for complying with it.

Third, SMEs are more exposed to implementation fragmentation. Wherever guidance is abstract, enforcement expectations vary, or support ecosystems differ across MS, smaller firms are the ones least able to navigate the resulting uncertainty. Differently from large organizations, they often cannot rely on in-house legal, HR, or digital teams to translate broad regulatory principles into everyday practice. So, implementation quality is a central part of the SME risk profile. Therefore, the policy issue isn't whether SMEs should be exempted from governance requirements, but whether governance is intelligible, proportionate, and replicable for firms with limited internal capacity. Existing EU support models, such as the European Digital Innovation

Hubs (EDIH), are useful here because they treat technical adoption, organizational change, and skills support as interconnected issues rather than separate ones (European Commission, 2025a).

For the above reason, the SME challenge in this report is best understood as one of constrained absorption under asymmetric burdens. Smaller firms aren't just more likely to struggle with adoption capacity, they are also more likely to experience legal uncertainty, compliance costs, and adjustment demands as mutually reinforcing pressures. Looking at the scoreboard, it means that legal and fiscal risks, income insecurity linked to weak competitiveness, and elements of social vulnerability during transition may compound one another more easily in SMEs than in large firms. The policy implication is that SME-oriented measures must prioritize standardization, practical guidance, and capability support, rather than assuming that formal obligations alone will generate convergence.

### **7.3 Synthesis:**

The comparisons made in this chapter are a demonstration of the ways in which the four dimensions of professional risk concentrate differently across firm types. In large firms, the central challenge is typically one of managed transition. By contrast, in SMEs, the dominant challenge is one of constrained absorption capacity: the core risks arising from limited organizational slack, weaker implementation capacity, and the disproportionate burden of navigating adoption, compliance, and adjustment with fewer internal resources.

This distinction helps us show that AI-related professional risk is not reducible to a single policy problem. A framework designed primarily for large firms would assume that organizations already have the internal capacity to interpret rules, absorb compliance costs, and manage workforce transitions. Meanwhile, a framework designed primarily for SMEs would focus too narrowly on reducing burden, while underestimating the scale of transition and governance issues that emerge in larger organizations. In practice, both of these risks are real, but they are distributed differently. As such, the policy challenge isn't only to protect workers from the

downside effects of AI, but to do so in ways which are proportional to the firms' organizational capacities and risk profiles.

For the purpose of this report, it leads to a clear conclusion: effective EU policy must be differentiated in design, even if it remains common in principle. Larger firms will require stronger emphasis on transition planning, worker information and consultation, and safeguards for AI-supported management. On the other hand, SMEs will require a stronger emphasis on capability support, standardized compliance pathways, and low-friction implementation tools. Our next chapter will build on this diagnosis by identifying the policy levers through which the EU can respond to each of the risk patterns without freezing adoption or deepening divergence across firm types.

## **8 From Diagnosis to Intervention**

The differences uncovered in the previous chapter imply that the policy question be framed as simply as whether the EU needs “more regulation” or “more innovation support”. It is more important to ask through which channels the EU can reduce the four dimensions of professional risk without freezing adoption or widening divergence across firms. Here, the central difficulty becomes the gap between broad regulatory principles, uneven implementation capacity, and firms' very different ability to absorb legal, organizational, and adjustment costs. There are existing EU frameworks which already provide parts of the necessary architecture, but they have not yet been systematically adapted to the practical realities of workplace AI.

This chapter will map the EU's intervention menu. Its purpose is to identify the main channels through which professional risk can be reduced across the four dimensions of the scoreboard: implementation tools, guidance and coordination, enforcement capacity, and funding/support instruments. Our goal isn't to prioritize individual proposals yet, but to clarify which types of policy lever are available, what problems they are best suited to address, and how they relate to the differentiated risk patterns identified in the previous chapter. This will provide the foundation

for the trade-off analysis and recommendations which will follow.

## 8.1 EU Channels of Interventions

The EU's capacity to shape the impact of AI in the workforce isn't dependent on a single lever. It operates through a combination of regulatory frameworks, implementation tools, institutional coordination, and funding instruments. Each of these address different parts of the professional-risk problems. It is important to make this distinction, as the four dimensions of the scoreboard don't all respond to the same type of intervention. For example, some risks can arise from gaps in legal clarity, others from weak implementation, others from limited organizational capacity, and others from the absence of effective transition support. As such, the task of policy design isn't only to identify the right objectives, but also to match each objective to the intervention channel most capable of delivering it.

The first channel is **implementation tools**. These are the mechanisms which translate broad rules and policy commitments into repeatable organizational practice: standard templates, model notices, minimum documentation sets, risk-assessment formats, reporting routines, and operational checklists. These are important due to the fact that many workplace-AI problems aren't caused by the total absence of rules, but more by the difficulty of turning abstract obligations into usable procedures. Implementation tools prove especially relevant wherever firms (particularly SMEs) need to know not simply what the law requires in principle, but how to comply in day-to-day workplace settings.

The second channel is **guidance and coordination**. Here, the EU acts to reduce interpretive uncertainty and limiting fragmentation across MS. It includes common guidance, examples of compliant practice, cross-border coordination, and the development of shared expectations around governance standards. This is a valuable channel for policy design, as it lowers the cost of legal ambiguity. Where firms face uncertainty over how rules apply, guidance can improve predictability even before new legislation is introduced. It is particularly important in workplace

AI, where legal obligations often intersect with labour law, data protection, product governance, and occupational safety and health.

The third channel is **enforcement capacity**. Rights and duties have limited practical capacity if institutions cannot verify, interpret, or enforce them in the workplace. In this context, enforcement includes labour inspectorates, data protection authorities, courts, complaint channels, and the broader administrative capacity needed to audit or contest AI-enabled decisions. This channel is crucial, as workplace AI often creates harms which are difficult for individuals to challenge on their own; without credible enforcement, formal protections remain weak in practice. Especially for this reason, enforcement capacity should be understood not as a secondary add-on, but as a core part of how the EU reduces legal risk and autonomy-related harms.

The fourth channel is **funding and support instruments**. Some of the most important barriers identified in earlier chapters are not primarily legal, but organizational and distributive. Firms may lack absorptive capacity, workers may need retraining, and transitions may require public support to avoid avoidable income and protection losses. Instruments such as the ESF+ are relevant here as they provide a way to finance reskilling, employment support, and social inclusion measures, while the Digital Decade and EDIH illustrate how the EU can support digital capability-building and practical adoption support. In other words, funding matters not only because it subsidizes change, but because it makes policy goals implementable for actors who would otherwise struggle to meet them.

Together, these four channels provide the basic infrastructure for the rest of the chapter. The following sections use them to map how the EU can intervene across each dimension of professional risk. It is important to recognize how no single channel is sufficient on its own: rules without implementation remain abstract, guidance without enforcement remains weak, and funding without governance may improve adoption without reducing risk. A credible EU response therefore depends on combining these channels in ways that reflect both the risk

dimension at stake and the organizational realities of different firms.

## 8.2 Income Security Levers

For the income security dimension, the EU's main levers aren't concentrated in a single instrument. Instead, they are spread across funding channels for labour-market participation, targeted adjustment funds for displaced workers, skills-coordination architecture, and a set of procedural labour-law frameworks that shape how restructuring is managed inside firms.

The first concrete lever is the **European Social Fund Plus (ESF+)**, which is the EU's main financing instrument for employment, labour mobility, education, and social inclusion (Regulation (EU) 2021/1057). For income security purposes, it is particularly useful because it can support reskilling, upskilling, job-search assistance, and broader labour-market adaptation capacity in MS. In other words, ESF+ is not a compensation tool but more of a transition-capacity tool: helping to determine whether workers can move into new employment quickly enough to avoid prolonged income loss when technology reshapes labour demand. In this context, it is especially relevant where AI-related adjustment takes the form of task change and employability pressure rather than immediate mass dismissal.

The second concrete lever we look at is the **European Globalization Adjustment Fund for Displaced Workers (EGF)**, which is the EU's clearest instrument for supporting workers after major restructuring events (Regulation (EU) 2021/691). Under current regulation, MS may apply for EGF support for measures that target displaced workers and self-employed persons, including cases involving at least 200 affected persons. The regulation also explicitly contemplates situations that involve SMEs in the same region or sector. As such, the EGF is particularly relevant where AI-related restructuring is concentrated enough to create a visible displacement shock, but where public support can still help workers re-enter employment through a coordinated package of personalized services.

A third set of levers is that of the EU's skills-governance architecture: the **European Skills Agenda**, the **Pact for Skills**, and the **Union of Skills** strategy (European Commission, 2020a; 2020b; 2025b). These instruments aren't meant to replace income support, but are crucial as they aim to reduce the likelihood that technological change turns into durable earnings loss. The European Skills Agenda aims to provide a strategic framework for sustainable competitiveness, social fairness, and resilience (European Commission, 2020a). Meanwhile, the Pact for Skills provides a more practical coordination mechanism, meant to encourage stakeholders to translate commitments into more concrete upskilling and reskilling acts (European Commission, 2020b). Lastly, the Union of Skills, presented in March 2025, links skills development to labour shortages, competitiveness, and high-quality jobs (European Commission, 2025b). Putting these together, we can understand them as anticipatory transition levers: trying to improve income security before disruptions cause unemployment or wage reductions.

The fourth relevant channel is that of procedural labour law, especially **Directive 2002/14/EC** on informing and consulting employees and **Directive 98/59/EC** on collective redundancies (Directive 2002/14/EC; Directive 98/59/EC). These aren't exactly income-support measures, but they matter because they structure the way in which workforce adjustment is handled within firms. The information-and-consultation framework is a good procedural footwork for workers whenever organizational change is being planned (Directive 2002/14/EC). Meanwhile, the collective redundancies framework requires timely notification and consultation, giving public authorities a role in responding to projected large-scale redundancies (Directive 98/59/EC). Regarding income security, they bring value thanks to their ability to make restructuring more visible, more contestable, and more manageable. These things increase the chances that transition support, redeployment, or retraining can occur before income losses become entrenched.

Putting these instruments together, we can see that the EU already governs income security through a combination of funding, skills coordination, and procedural restructuring law. As such, the main limitation which comes up is not the complete absence of tools, but the fact that these tools are only partially integrated into coherent responses to AI-driven adjustment. Practically

speaking, the strongest income-security strategy likely combines ESF+ employability support, EGF-type adjustment support for concentrated shocks, skills coordination for anticipatory reskilling, and the procedural frameworks that can push firms to make a managed transition rather than purely ex-post displacement.

### **8.3 Social Protection and Downside-Protection Levers**

The EU's toolkit for social protection is built less around a single distributive instrument than around a combination of coverage benchmarks, cross-border coordination rules, administrative implementation tools, and operational support for fair mobility. Specifically in the context of AI, this toolkit is relevant because downside risk arises not only from job loss, but from the fact that workers move frequently across jobs, contractual forms, or self-employment, and may therefore encounter coverage gaps, benefit interruptions, or status penalties (Council Recommendation 2019/C 387/01).

The most obvious EU benchmark on this issue is the **Council Recommendation** on access to social protection for workers and the self-employed (Council Recommendation 2019/C 387/01). It is important as it explicitly treats workers and the self-employed together, framing reform around the principles of formal coverage, effective coverage, adequacy, and transparency (Council Recommendation 2019/C 387/01). In practice, it makes the EU's main normative lever for addressing the kinds of transitions that matter in an AI-shaped labour market: movement across labour statuses, more fragmented careers, and non-standard forms of work that don't fit neatly into systems built around full-time open-ended employment. Unfortunately, the Recommendation is limited, as it is soft law. It simply sets direction and benchmarks, but implementation still depends heavily on MS.

The second lever is the EU's **social security coordination framework**, centered on **Regulation (EC) No 883/2004** (Regulation (EC) No 883/2004) and its implementing **Regulation (EC) No 987/2009** (Regulation (EC) No 987/2009). These don't create a common European

welfare state, instead, they coordinate national social security systems so that people can move across MS without automatically losing rights or contributions. For the purpose of our report, we identify their value in their ability to make portability and continuity across borders a binding part of the EU toolkit, rather than a purely national matter. In parallel, by design these regulations are strongest in cases of cross-border mobility; they don't by themselves solve all domestic protection gaps that arise when workers switch status within a single national system.

A third lever is that of **administrative and digital implementation**, where the most relevant current initiative is the **European Social Security Pass (ESSPASS)** (European Commission, 2025c). It is designed to simplify the exercise of social security rights in another European country by digitalising entitlement documents and enabling real-time cross-border verification. As of February 2026, the EC has launched a public consultation to prepare a proposal for ESSPASS as part of the **Fair Labour Mobility Package** due in autumn 2026 (European Commission, 2026a). From this we see how important an example ESSPASS is for the main idea of this chapter. Some EU levers aren't new benefit schemes, but implementation infrastructures that can reduce friction, delays, and proof-of-status problems that make portability difficult in practice.

Fourth, we have the **European Labour Authority (ELA)** (Regulation (EU) 2019/1149). It was established to ensure fair labour mobility across EU MS and to support the coordination of social security systems. Within the context of downside protection, it is relevant in an operational sense rather than a redistributive one: cooperation between national authorities is facilitated, which is an important factor whenever portability, entitlement verification, or cross-border enforcement becomes administratively complex.

Finally, **ESF+** remains relevant as a supporting lever, despite it not being a social-security coordination regime. However, it does support social inclusion, institutional capacity, and labour-market adaptation measures meant to help MS close some of the practical gaps identified by the 2019 Recommendation mentioned above (Council Recommendation 2019/C 387/01).

Taken as a whole, these levers are a showing of the EU's existing architecture for reducing downside protection risk: a soft-law benchmark for broad access (Council Recommendation 2019/C 387/01); binding cross-border coordination rules (Regulation (EC) No 883/2004; Regulation (EC) No 987/2009); operational support through the ELA (Regulation (EU) 2019/1149); and an emerging digital portability tool through ESSPASS (European Commission, 2025c). We can see that the main gap is not the total absence of EU involvement, but the fact that the current architecture seems to be stronger on cross-border coordination than on domestic continuity across changing forms of work. Exactly for this reason, social protection will be treated in this report as a question of continuity and portability, not only of benefit generosity.

## 8.4 Legal and Fiscal Risk Levers

The EU's toolkit for this dimension is more developed than it may first appear. In our report, the risk dimension isn't limited to formal liability after harm occurs; it also takes note of the cost of legal uncertainty, the administrative burden of compliance, and the availability of credible redress when AI affects employment access, working conditions, or managerial decisions. As a result, the most relevant EU levers won't just be substantive rules, but also the rules which make those rules intelligible, standardized, and enforceable in practice.

The central legal authority we look at is the **AI Act**. It creates a harmonized, risk-based framework for AI across the internal market, and explicitly treats AI systems used in employment, worker management, and access to self-employment as part of the high-risk architecture. Practically speaking, it gives the EU a clear lever to reduce legal uncertainty in workplace AI: the issue isn't whether there is a framework, but whether firms can understand how to apply it in specific workplace contexts. It is also relevant as it combines protection with innovation-support measures directed at SMEs and start-ups, meaning it is already designed as more than a pure prohibition regime.

Secondly, there are the levers which concern **implementation support under the AI Act**. The **European AI Office** (European Commission, 2024) and the broader **AI Act governance** (European Commission, 2025d) matter because legal risk is often driven by uncertainty over interpretation rather than by the bare existence of obligations. The AI Office plays a central role in implementation and oversight, while the EC has also moved to support implementation through guidelines, an **AI Act Service Desk**, and a **Single Information Platform** (European Commission, 2025e; 2025f). This lever is especially significant for SMEs, as we have explained in earlier chapters.

A third level is **standardization**, which is often underestimated in policy discussions but is central to reducing compliance cost. Under the AI Act, harmonized EU standards are intended to support implementation and provide legal certainty, especially where firms need to demonstrate conformity in high-risk contexts (European Commission, 2026b). Standardization has policy value as it can translate abstract legal obligations into more predictable technical and organizational requirements. Standards help convert “what the law requires” into “what a firm actually has to document, test, and maintain.” This is one of the clearest channels through which the EU can lower the fixed cost of compliance for firms while ensuring that safeguards aren’t weakened.

A fourth lever is the combination of **regulatory sandboxes and simplification measures**. Regulatory sandboxes are provided for in the AI Act, meant as a support instrument for innovation and testing, and the EC has been developing the implementing rules for the way that such sandboxes should operate (European Commission, 2025g). Additionally, the **EC’s January 2026 Digital Omnibus on AI Regulation Proposal** explicitly frames simplification as a way to ensure that implementation remains timely, smooth, and proportionate (European Commission, 2026c). This is important, as fiscal risk is often less about taxation than about the cost burden of compliance, documentation, and external advisory dependence, especially for smaller firms. So, sandboxes and simplification are important not because they remove regulation, but because they make compliance more proportionate and testable before costs escalate.

Alongside the AI Act, the **GDPR** remains a major part of the legal-risk architecture whenever workplace AI relies on personal data (GDPR, 2016). **Article 22** and the associated **EDBP–endorsed guidance on automated decision-making and profiling** are particularly relevant as they address cases where individuals are affected by decisions based solely on automated processing (EDPB, 2018). For this report, we are less interested in the abstract privacy doctrine and more in the way they shape contestability, meaningful human involvement, and the informational conditions under which workers can challenge AI-mediated decisions (EDPB, 2018). Essentially, legal risk in the workplace is governed through an overlap between the AI Act’s ex ante system governance and the GDPR’s data-protection and decision-rights logic.

Furthermore, another component of the EU architecture is **redress and liability**. We refer to the recent development of the **revised Product Liability Directive**, which updates EU product-liability rules in a way that reflects software and digital products clearly (Directive (EU) 2024/2853). This directive strengthens the broader accountability environment for harms linked to defective digital products, including AI-enabled products, and improves legal certainty for both claimants and businesses. At the same time, the EC’s 2025 **Work Programme** confirms that the separate **AI Liability Directive proposal** was withdrawn (European Commission, 2025h). The consequence for our report is the following: the EU’s current liability architecture is stronger on product liability and system governance than on a standalone civil-liability regime tailored specifically to AI. It is for this reason that guidance, documentation, enforcement coordination, and contestability mechanisms are so important.

Taking these levers together, we see that the EU already has a substantial legal-risk architecture: a harmonized ex ante framework through the AI Act, implementation and interpretive support through the AI Office and EC guidance, compliance support through standards and sandboxes, worker-facing rights through the GDPR and EDPB guidance, and a broader redress layer through product liability. So, the main weakness which arises is not the absence of legal tools, but more the fact that firms (especially SMEs) still face a gap between formal obligations and a usable compliance pathway. This gap is exactly where recommendations on templates,

documentation sets, and coordinated enforcement can add the most value.

## 8.5 Cross-cutting Design Principles: How to Combine EU Levers

The previous sections showcased how the EU already has a meaningful set of levers across all four dimensions of professional risk. The key policy question which arises is therefore not only which instruments exist, but how they should be combined. Throughout the report, we identify five cross-cutting principles which matter the most: **proportionality**, **operability**, **enforceability**, **worker voice** and **contestability**, and measurability. These aren't separate criteria in themselves, they determine whether the existing toolkit can actually reduce risk without freezing adoption or widening divergence across firms.

The first principle is that of **proportionality**, which in our context means that obligations should be common in principle but differentiated in intensity according to risk and organizational capacity. The AI Act is explicitly built on a risk-based logic rather than a one-size-fits-all model, while the OSH Framework Directive explicitly states that worker-protection measures should avoid imposing administrative, financial, and legal constraints that would hold back the creation and development of SMEs. Putting these frameworks together, they suggest that good policy should not exempt smaller firms from governance altogether, but should ensure that compliance pathways remain proportionate to their internal capacity. This principle is especially important in the scope of this report, which previously showed how large firms and SMEs face different dominant pressures even when they operate under the same broad legal environment.

The second principle is **operability**. A policy framework is weak if firms are able to identify the legal objective but not the practical steps required to meet it. In the context of workplace AI, operability means that obligations should be capable of being translated into usable procedure, standard documentation, clear internal responsibilities, and repeatable implementation routines. It matters since many of the problems identified in earlier sections are not caused by the total absence of law, but by the gap between broad legal principles and workable compliance

pathways. As such, a usable framework is one in which implementation does not depend entirely on bespoke consultancy, informal interpretation, or firm-by-firm trial and error. In practice, operability is the principle that connects the legal architecture mapped in 5.4 to our later recommendations for templates, documentation sets, and standardized support pathways.

The third principle is **enforceability**. Formal rights and duties matter only if institutions can verify, interpret, and act on them in real workplaces. Here, the EU's existing labour and OSH framework is especially important. **Directive 2002/14/EC** establishes a general framework for informing and consulting employees, while the **OSH Framework Directive** requires information, dialogue, and balanced participation in relation to workplace risks (Directive 89/391/EEC). These instruments point to a broader lesson: EU workplace governance is strongest when it does not rely solely on individual complaints after harm occurs, but also creates institutional channels through which change can be anticipated, discussed, and monitored. In the AI context, enforceability therefore depends not only on courts or regulators, but also on the presence of credible organizational processes through which workers, managers, and authorities can identify and respond to risks before they become entrenched.

The fourth principle is **worker voice and contestability**. This principle is related to enforceability but distinct from it. Enforceability asks whether rules can be applied; contestability asks whether people affected by AI-enabled decisions can understand, question, and challenge those decisions in practice. This matters because many workplace harms linked to AI are mediated through opaque management systems rather than direct legal commands. The most concrete EU model here is the **Platform Work Directive**, which addresses situations where decisions are taken or supported by the automated processing of personal data and responds to problems of opacity, lack of explanation, and weak human contact for affected workers (Directive (EU) 2024/2831). Combined with older information-and-consultation and OSH participation logics (Directive 2002/14/EC; Directive 89/391/EEC), this suggests that good policy should treat contestability not as a narrow procedural add-on, but as a core condition for maintaining worker autonomy and agency in AI-enabled workplaces.

The fifth and last principle is **measurability and policy learning**. A credible EU strategy shouldn't just deploy levers, it should also make it easier to observe whether they are improving outcomes over time. This is a crucial distinction to make, as the four professional-risk dimensions are inherently multi-dimensional and may move in different directions at once. For the purpose of this report, we will tie workplace-AI governance to a small number of measurable indicators, in a way that makes it possible for policy to be adjusted when adoption improves but risk doesn't, or when worker protection improves at the cost of excluding smaller firms from implementation. Essentially, measurability is the principle that turns our four-risk framework from a conceptual lens into a usable policy scoreboard.

The previous sections displayed the substantial set of levers already possessed by the EU across the four dimensions of professional risk. However, the policy challenge isn't simply to accumulate instruments, but also to determine how to combine, sequence, and prioritize these measures in such a way that reduces risks without imposing unnecessary burdens, weakens adoption, or produces new forms of inequality across firms and workers.

For this reason, our next chapter shifts focus from the menu of available interventions to the trade-offs involved in using them. Some of the above-mentioned levers strengthen protection but raise implementation costs; others support adoption but leave gaps in accountability or worker voice if used alone. So, prioritization is not only technical but also strategic. It requires a decision which efficiently combines instruments that are the most capable of improving the professional-risk scoreboard while maintaining proportionality, operability, and political feasibility.

## **9 From Intervention Menu to Prioritization**

In the previous chapter, we demonstrated the existing set of meaningful levers available to the EU across the four dimensions of professional risk. Our policy challenge, however, is not simply to identify these tools one by one. We need to determine how they should be combined,

sequenced, and prioritized in a way that improves the overall professional-risk scoreboard without freezing adoption or deepening divergence across firms.

As noted earlier, this is a necessary step as the EU's policy levers are not automatically complementary. Measures that strengthen accountability could lead to a rise in implementation costs. Others that simplify compliance could reduce burden, but run the risk of weakening robustness if simplification becomes mere dilution. Funding and transition support may reduce short-term disruption, but may not by themselves address opacity, psychological risk, or contestability in AI-enabled management. So, our policy problem is inherently multi-objective: different instruments improve different dimensions of risk, and some combinations can solve one problem while worsening another.

For the above reason, this chapter will examine the main trade-offs which arise when the EU's existing levers are used together. We are not trying to eliminate these trade-offs with this chapter, but to clarify which tensions are unavoidable, which can be mitigated through design, and which criteria should guide prioritization. The following sections will assess the most important tensions in our policy space between adoption and protection, harmonization and subsidiarity, simplicity and accountability, transparency and feasibility, prevention and compensation, and uniformity and differentiated pathways.

## **9.1 Adoption and Competitiveness vs Protection and Oversight**

One of the central tensions in EU AI policy is that adoption is itself a policy objective, not just a market outcome. The **Digital Decade Policy Programme 2030** explicitly treats the digital transformation of businesses as a strategic priority (Decision (EU) 2022/2481). Meanwhile, the **AI Act** is not only framed around fundamental-rights protection, but also around improving the functioning of the internal market and supporting innovation (AI Act, 2024). So, it is clear that the EU is operating under a dual mandate: being expected to accelerate diffusion of AI and to

ensure that such diffusion doesn't occur through the unmanaged externalization of risk onto workers, job quality, or smaller firms.

The tension arises because measures that improve protection and oversight often raise the visible cost of adoption. Requirements that concern risk management, documentation, worker information, or human oversight can increase the time, expertise, and organizational effort needed to deploy AI lawfully. As previously mentioned, the main issue is not simply “innovation versus regulation”, but the unequal distribution of compliance burdens across firm sizes.

At the same time, if a strategy prioritizes speed of adoption but fails to have adequate governance it can create a different distortion. Firms may boost productivity through opaque decision systems, intensified monitoring, or poorly managed restructuring, while the resulting legal, social, and psychological costs are borne elsewhere. It is because of this reason that simply “lightening” protections won't solve the trade-off: lower short-term burdens can support faster take-up, but also permit forms of adoption that undermine autonomy, contestability, and worker protection.

So, the question which arises is not whether protection should be sacrificed for competitiveness, but what form of protection is least likely to suppress productive adoption while also constraining harmful deployment. We find that the adoption-protection trade-off is best managed through proportionality, implementation support, and sequencing. Proportionality ensures that stricter obligations apply where worker-facing risks are higher. Implementation support matters because adoption is more likely to stall when compliance depends on uncertainty, fragmentation, or expensive external advice rather than usable pathways. Lastly, sequencing matters because templates, guidance, and support infrastructures can work to lower friction before more demanding obligations go into force.

The practical implication is that EU policy should distinguish between productive compliance costs and unproductive compliance costs. Productive costs are ones that create accountability, predictability, and safer deployment. Meanwhile, unproductive costs arise from ambiguity,

duplication, fragmented interpretation, or lack of implementation support. So, the aim which arises is that of reducing the latter without dismantling the former.

## **9.2 Harmonization vs Subsidiarity and National Diversity**

The second major trade-off we identify concerns the relationship between EU-level harmonization and national diversity in labour-market and social-protection systems. AI-related risks are shaped by cross-border markets, common digital tools, and shared compliance expectations, while the institutions that manage employment relations, income support, and workplace adjustment are deeply national. As such, the result is a familiar dilemma. Too little EU action would leave costly fragmentation; too much standardization runs the risk of colliding with institutional diversity that is both legally and politically entrenched.

Looking at harmonization, the case for it is strongest wherever fragmentation directly generates legal or operational risk. A good example is that of the AI Act, as it establishes a harmonized framework for AI across the internal market. Wherever firms deploy similar systems across multiple MS, divergent national interpretations would increase compliance costs, reduce predictability, and weaken the legal clarity which the Act is supposed to provide. As such, harmonization is presented as especially valuable where the goal is to reduce uncertainty, duplication, and cross-border inconsistency.

However, not all risk dimensions lend themselves to the same degree of centralization. Social protection is a clear counterexample. The EU plays an important role through the Council Recommendation on access to social protection and through the coordination of social security systems under Regulations 883/2004 and 987/2009, but these instruments don't necessarily create a common European welfare state. They coordinate national systems and set common orientations while leaving benefit design, financing, and labour-market institutions largely to MS. This is a reflection of both subsidiarity and practical reality. Continuity and portability can

be improved at EU level, but can't be fully standardized without moving far beyond the EU's current settlement.

The most effective response is unlikely to come from choosing either full harmonization or full decentralization. A more resistant design is usually one with a layered approach: common principles, minimum standards, and shared governance where fragmentation is costly; coordination, funding, and implementation support where national diversity remains legitimate but capacity gaps are real, and room for MS to adapt those frameworks to their own labour and social systems.

### **9.3 Simplicity and Proportionality vs Robustness and Accountability**

A third trade-off concerns the relationship between simplicity and robustness in policy design. Many firms (especially SMEs) find it difficult to translate rules into usable, proportionate procedures. As such, there is a strong case for simplification. However, it does carry its own risk: if governance becomes too light, too generic, or too formalistic, it can reduce the burden only by hollowing out the very protections it is meant to preserve.

The case for simplicity is strong as workplace-AI governance can otherwise become consultant-dependent and administratively heavy. In practice, firms comply better when obligations can be expressed through standard processes, clear documentation routines, and repeatable implementation steps, rather than through ad hoc interpretation or costly external advice. It is also why harmonized standards under the AI Act are important, they are meant to improve legal certainty and make implementation more predictable.

The problem with simplification is that it can slide into substantive dilution. A governance system can look "simple" because it uses short checklists or minimal paperwork, but it can fall short in providing meaningful accountability if it doesn't require firms to document real risks, allocate clear responsibilities, or maintain credible safeguards. In such a case, simplification produces box-ticking compliance rather than genuine governance. This is an especially danger-

ous dynamic in workplace AI, where harms such as reduced discretion, opaque managerial systems, or psychosocial strain are often indirect and cumulative.

Following the arguments above, we believe the right goal is not minimal regulation, but low-friction robustness. This means attempting to reduce costs that do not improve protection while preserving what does matter, such as clear responsibility, risk assessment, worker information, and reviewability. Previously, we pointed out several EU levers that point towards this middle path. Harmonized standards can reduce uncertainty without removing substance, regulatory sandboxes that can help firms test compliant pathways before full deployment costs escalate, and the Digital Omnibus on AI Regulation Proposal which similarly frames simplification as making implementation timely, smooth, and proportionate, not stripping obligations away.

For the scope of this report, we should judge simplicity based on how well it supports accountable implementation, not by how much formal burdens it removes. Proportionality should then apply to the pathway into compliance; and robustness must still apply to the substance of governance.

## **9.4 Transparency and Contestability vs Confidentiality and Administrative Feasibility**

The fourth trade-off concerns transparency and contestability versus confidentiality and administrative feasibility. Many workplace harms that are associated with AI don't arise from openly stated managerial instructions, but from the systems which shape recruitment, scheduling, evaluation, or work allocation in ways that workers may struggle to understand or challenge. In such a setting, transparency isn't just informational; it is a condition of contestability. The reason for this is that workers can't realistically question an AI-mediated decision if they aren't even aware that automated processing was involved, what type of decision was affected, or how to seek human review.

Simultaneously, transparency doesn't automatically justify maximal disclosure. Firms can argue that requiring full disclosure of source code, model architecture, or every internal decision parameter would create problems of trade-secret exposure, system gaming, and administrative overload. Too much poorly targeted information can therefore undermine contestability almost as effectively as too little.

For that reason, it is important to make a policy distinction between full technical disclosure and meaningful process transparency. The objective here isn't to give every worker access to source code. Instead, it is to ensure that workers can understand the type of system being used, the kind of decision it supports or influences, the main criteria or data categories involved, the practical consequences for them, and the pathway for requesting review or explanation. This is an important distinction which is increasingly visible across EU practices. The Platform Work Directive strengthens transparency, fairness, human oversight, safety, and accountability in algorithmic-management procedures without the need for platforms to hand over code (Directive (EU) 2024/2831). EC materials on the AI Act also emphasize guidelines and a Code of Practice to support compliance with transparency obligations (European Commission, 2025e).

The implication here is that, if our goal is to balance transparency with feasibility, the best route will be to define a minimum package of contestability-relevant information. This means a notification that AI is being used, a plain-language description of the decision context, the role of automated processing, the categories of information relied upon, the responsible contact point, and the existence of a human review route. Like this, confidentiality is preserved where appropriate, and it is still made materially easier for workers to identify when AI matters and how to respond.

## **9.5 Prevention vs Compensation**

A fifth trade-off concerns timing: whether the EU should prioritize preventing professional risk before it materializes or to compensate for it after harm occurs. As we previously discussed, the

four dimensions of the scoreboard do not all lend themselves to the same dynamics, and therefore the same timing. Some of them are organizational and foreseeable (like opaque management systems, weak worker information, or psychological strain linked to AI-enabled systems) while others emerge through economic shocks and transition failures (such as displacement, prolonged unemployment, or disputed harmful outcomes).

As such, the case for prevention is strongest wherever risk is embedded in workplace design and can be anticipated *ex ante*. The AI Act once again proves to be a useful example, as it is built around risk management before deployment for high-risk systems. Likewise, the OSH Framework Directive is preventive: it requires employers to assess risks, adopt protective measures, and involve workers in managing workplace hazards. In our workplace-AI context it is of relevance because harms related to autonomy and agency are often predictable consequences of how systems are implemented and governed.

Additionally, prevention is important when the issue is procedural fairness and worker voice rather than direct income loss. The **information-and-consultation framework under Directive 2002/14/EC** matters as it gives workers a foothold before any organizational decisions are made. The Platform Work Directive follows a similar logic with algorithmic management, with safeguards around transparency, human contact, and oversight being intended to reduce the likelihood of opaque or unchallengeable outcomes (Directive (EU) 2024/2831). Furthermore, prevention isn't only more humane here, it is also often cheaper institutionally than relying on *ex post* disputes or litigations.

However, compensation and adjustment support cannot be completely ignored, as they remain important in situations where disruption can't be fully prevented. It is especially true for the income-security dimension. As such, it follows that the EU's role isn't to stop all labour-market change, but to reduce the damage when it inevitably occurs. Here, the ESF+ is important for employability, reskilling, and labour-market adaptation (Regulation (EU) 2021/1057), while the EGF can more directly be used to help workers affected by major restructuring shocks

(Regulation (EU) 2021/691) The core implication from the arguments above is that prevention should be used where risk is organizational, foreseeable, and cumulative, while compensation remains necessary where shocks, failures, or residual harms cannot be fully avoided. We can see that autonomy/agency and legal risk should lean more towards preventive governance, while income security and some elements of downside protection must require stronger compensatory and adjustment tools.

## **9.6 Uniform Obligations vs Differentiated Pathways**

The sixth and last trade-off we look at concerns the relationship between uniform obligations and differentiated implementation pathways. Clearly, this follows from the discussion previously carried out in Chapter 4, larger firms and SMEs don't have to deal with the same dominant AI-related risks, despite being under the same broad legal environment.

Like previously discussed, there are valid reasons to retain a common baseline of obligations. Uniformity is important as the report's framework assumes that workers shouldn't lose protection simply because they are employed in a smaller or a less digitally advanced firm. However, a fully differentiated regime comes with its risks. Worker protection could be weakened by allowing lower-capacity firms to evade core safeguards around transparency, contestability, or psychological risk, while also undermining legal clarity by replacing a common governance logic with a patchwork of exceptions.

A strong design will likely involve uniform obligations in substance, but offer differentiated pathways in implementation. Under such a model, the EU can maintain common expectations around risk assessment, worker information, contestability, and transition support. The way that such expectations are operationalized should vary according to organizational capacity and use case.

This trade-off also relates to the tension we discussed earlier between simplicity and robustness. Lowering core obligations for SMEs would go towards weakening protection and

further fragmenting governance. Identical implementation burdens would slow adoption, raise advisory dependence, and widen the divergence between firms that can and can't cope. With differentiated pathways we have a middle route: they preserve a common legal floor while simultaneously adapting the compliance process for different organizational realities. The EU should opt for a layered model where core protections remain common, and implementation pathways, support mechanisms, and documentation requirements are calibrated to risk exposure and capacity.

## 10 Prioritised Recommendations

### Policy Focus:

Before delving into the actual mix of policies which European institutions (as well as their MS) could and should implement, it is important to remember what these policies should achieve. They should focus on two different approaches: **Protecting and Empowering**.

Protecting the existing structures of the labor market which are today more at risk. This category mostly relates to employees in large firms which may fear job displacement, those whose skills may become obsolete, young graduates facing disruptions in the entry-level jobs landscape and senior professionals whose expertise might become less compatible with the new AI augmented workplace. It also relates to increasing worries about algorithmic management and control as well as the needs for ethical and responsible usage of AI in any company structure, to protect professional ecosystems against overreliance, irresponsible habits and excessive offloading in their usage of AI tools.

Empowering the fields and careers which may see increased opportunity when using AI as a reliable partner. This part can be crucial for SMEs, to fill skill gaps and quality gaps compared to larger competitors, to facilitate reaching good standards of legal and fiscal compliance as well as to speed up time to security with more efficient business operations which can also help cut unnecessary costs. The same principle goes for innovative companies and entrepreneurs, as

testing mechanisms become faster and more efficient and access to crucial information to figure out the potential of a certain product or service can become easier and faster.

### **A Note on the Nature of These Recommendations**

The proposals that follow are policy recommendations. They are not binding prescriptions, implementation blueprints, or model legislation. Their purpose is to make the strongest possible argument for a set of policy directions that, taken together, would meaningfully reduce the professional risk that AI integration is creating for European workers, and open up the opportunities it is generating to a larger share of the workforce than current frameworks are allowing to reach.

Each recommendation comes with a proposed structure, a suggested allocation of responsibilities, and a set of illustrative implementation details. These details such as which institution should lead, which instrument should be used, how a benefit should be calibrated, what a KPI set should look like, are our best attempt to concretely show how each policy could work in practice. They are not the only way these policies could be designed, and in many cases they are not even the most important part of the argument. They are just one possible version of how each recommendation's logic could be translated into action.

The role of turning any of this into actual policy belongs entirely to policymakers. As such, each policy needs to be evaluated by any policymaker who finds value in them in order to figure out the actual implementation set that merges the technical and legal sides of turning these tools into sustainable and efficient outcomes at an EU level.

## **10.1 Policy Package at a Glance**

Five problems. Five answers. The policies that follow address the full range of professional risks that AI is creating for European workers, from the graduate who can no longer find a first job, to the experienced worker displaced without support, to the SME drowning in compliance, to the entrepreneur who sees the opportunity but cannot afford the leap. Together, they form a coherent package that reduces downside risk without halting adoption.

Here is a glance of the policy package, after which the report dives into each recommendation separately.

POLICY	WHY	WHAT IT IS	WHO IT'S FOR	RISK DIMENSIONS IT TACKLES
<b>R1. Professional AI Proficiency Certificate (PAPC)</b>	Employment for workers aged 22-25 in AI-exposed occupations has fallen sharply, not because firms do not need modern human skills in the era of AI, but because they have no way to verify whether a young hire can actually provide them.	A modular, EU-anchored <b>credential</b> that certifies workers can operate with, oversee, and govern AI tools. It turns undefined employer demand into a legible <b>market-entry signal</b> and transforms AI Act Article 4 literacy obligations from a compliance headache into a documented, verifiable standard that both workers and firms can rely on.	Recent graduates entering AI-exposed occupations. Firms navigating AI Act literacy obligations.	<b>Income Security:</b> Directly tackles displacement risk and time-to-security at the point of entry into the labour market, by shifting graduates toward augmentation-oriented roles where employment has held up.  <b>Autonomy and Agency:</b> gives workers a defined and legible role in the AI governance chain, preserving career path predictability.
<b>R2. AI Transition Reskilling Guarantee</b>	AI is reorganising large firms faster than workers can adapt on their own. The problem is not only displacement. It is that skills become obsolete before anyone has a plan, restructuring happens without warning, and workers are left to absorb the full cost of adjustment alone. 94 million European workers will need reskilling by 2030, and the current policy landscape has no instrument that makes transition a managed process rather than a personal emergency.	A dual-obligation framework pairing two things. First, a funded <b>personal transition bundle</b> covering skills assessment, career counselling, accredited training, and temporary income support. Second, a mandatory firm-side <b>Workforce Transition Plan</b> that requires companies to anticipate and manage AI-driven restructuring proactively. Together, they ensure that adjustment is a social transition, not a silent displacement.	Employees in large firms facing role compression, task obsolescence, or redundancy due to AI adoption.	<b>Income Security:</b> Directly addresses displacement risk, skill and knowledge obsolescence, and time-to-security, and limits wage scarring during transitions.  <b>Social Security and Downside Protection:</b> the personal transition bundle provides coverage continuity and an income floor during the adjustment period.
<b>R3. SME AI Capability Compliance Hub</b>	SMEs are not held back by lack of motivation. They are held back by: a weaker digital baseline than large firms, compliance costs that are disproportionately high, and a fragmented regulatory environment that forces them to pay for external advisors just to stay legal. The result is a widening capability gap that threatens both SME competitiveness and the workers who depend on them.	A one-stop <b>support architecture</b> that brings together standardised compliance toolkits, help-desk guidance, capability-building assistance, and an outcome-linked voucher scheme. It turns lawful, productive AI adoption from an expensive and fragmented puzzle into a navigable pathway for firms that do not have the internal capacity to figure it out alone.	SMEs and their workers, facing disproportionately high compliance burdens and a widening AI capability gap relative to large firms.	<b>Legal and Fiscal Risk:</b> Reduces compliance uncertainty, lowers dependence on costly external advice, and provides proportionate, standardised pathways that are actually feasible for smaller firms to follow.  <b>Income Security:</b> by supporting firm-level AI adoption and operational resilience, the Hub protects the earnings stability of workers employed in SMEs.
<b>R4. AI Augmentation Incentive Standard (AAIS)</b>	Firms have a rational incentive to replace workers with AI rather than augment them, and right now nothing in the policy environment makes that choice visible, measurable, or consequential. One in four CEOs expects to cut headcount by 5% in 2026 through AI alone. The EU is currently watching this happen with no lever to pull before the damage is done. This is the only risk in the package that can be addressed before displacement occurs, making prevention the only viable approach.	A classification mechanism that makes firms' AI deployment strategies, whether augmentation or replacement, visible and fiscally consequential. It activates a package of real incentives for firms that demonstrably expand human capability: tax top-ups, procurement preferences, EIB access, and an Augmentation Employer certification. It is the only recommendation in the package that intervenes at the moment firms are designing their strategies, before any harm to workers has occurred.	Workers in large firms whose roles are at risk of replacement rather than enhancement. Firms willing to commit to a human-centred AI deployment model.	<b>Income Security (preventive):</b> Unlike the other recommendations, it acts before displacement occurs, shaping firm-level AI investment decisions toward strategies that sustain rather than eliminate human roles.  <b>Autonomy and Agency:</b> by incentivising augmentation over automation, the Standard preserves workers' decisional authority and professional space within AI-enhanced workflows.
<b>R5. AI Entrepreneurship Bridge</b>	AI is genuinely lowering the barriers to entrepreneurship, reducing fixed costs and enabling individuals to do what once required full teams. But most displaced workers cannot take that opportunity. The moment a worker leaves employment, social protection disappears. The rational choice becomes staying in a degraded role rather than betting personal assets on a new idea, leaving the most promising escape route from AI disruption structurally inaccessible to those who need it most.	A structured three-phase <b>transition programme</b> that sits on top of existing social protection entitlements. It provides an additive income and support buffer during the move from employment to self-employment, so that displaced workers can test, build, and validate a new professional path without putting their financial cushion, or their future, on the line.	Workers formally displaced by AI-driven restructuring, or whose roles have been significantly degraded. Those for whom entrepreneurship is the logical next step, but the social protection gap makes it the one they cannot afford to take.	<b>Social Security and Downside Protection:</b> Directly addresses the protection cliff at the transition moment, ensuring that switching status from employee to entrepreneur does not mean losing the safety net.  <b>Income Security:</b> the programme's phased structure supports time-to-security for the early entrepreneurial period, reducing the financial volatility that makes self-employment inaccessible.  <b>Legal and Fiscal Risk:</b> Bridge participants gain access to the R3 Hub's compliance toolkits and guidance infrastructure from day one.

Table 1: Policy package at a glance

## 10.2 R1 - Professional AI Proficiency Certificate

### What Is It?

Our first recommendation proposes the creation of a **Professional AI Proficiency Certificate (PAPC)**: a modular, EU-anchored credential designed to make the AI competency of young graduates legible and verifiable to employers at the point of labour market entry. Its central purpose is to close the verification gap that AI-driven task automation has opened between the demand for human oversight of AI systems and firms' inability to assess whether a new hire can reliably provide it.

We base this recommendation on a specific finding from our report: Brynjolfsson et al. confirm that generative AI disproportionately automates the codified, entry-level tasks that once formed the first rung of professional career ladders, while complementing the tacit and experience-based judgement of senior workers, creating a structural mismatch between where human value is being preserved and where professional life begins. Early-career workers aged 22 to 25 in the most AI-exposed occupations experienced a 16% relative employment decline since the widespread adoption of generative AI, while employment for experienced workers remained stable or grew. Crucially, this effect is concentrated in automation-oriented AI applications, with significantly more muted changes in roles where AI augments rather than substitutes. As such, the resulting policy problem isn't simply that entry-level jobs are disappearing. Firms don't stop needing human oversight of AI, they just don't know how to verify that a young hire is competent to provide it. Without a standardized signal, employers default to more experienced workers who have built AI competence through years of tacit exposure, compressing the entry points available to graduates and weakening the generational renewal of AI-exposed occupations.

The PAPC aims to fill that verification gap, turning the AI-readiness of young graduates into a legible, market-signalling credential.

The recommendation builds directly on three foundational instruments already within the EU policy architecture, rather than attempting to construct an entirely new institutional logic from scratch.

- First, **AI Act Article 4**, in force since February 2025, legally obliges all providers and deployers of AI systems to ensure staff have “a sufficient level of AI literacy”, but creates no standard mechanism for verifying this for new hires. Enforcement rules apply from August 2026 onwards, meaning the verification demand is already becoming operationally urgent.
- Second, **DigComp 3.0**, published in December 2025, now integrates AI competence transversally across all five digital competence areas and all 21 competences, and is explicitly designed as a cross-border policy reference point. Its revised proficiency levels and learning outcomes provide a ready-made competence spine on which a professional credential can be anchored.
- Third, the **EU-funded AI Skills Academy**, launching in 2026 under the AI Continent Action Plan, and the European Digital Skills Certificate feasibility process are already developing modular credential frameworks and a quality label for digital skills training. As such, the PAPC would not create a new certification system from scratch. Instead, it would be the entry-level, young-graduates-focused piece of an infrastructure that already exists in outline, targeting labour market insertion specifically, which none of the above instruments address directly.

The certification is structured around four modular tracks, each targeting a distinct professional function reshaped by AI adoption:

- **Track 1 - AI Workflow Operator**: Certifies the ability to design, automate, and quality-check professional workflows using generative AI tools: prompt engineering, task chaining, agent supervision. This directly targets the entry-level task automation roles that have displaced junior coordination and administrative work. Relevant for: operations, marketing, HR, finance, and consulting support roles.

- **Track 2- AI Output Auditor:** Certifies the ability to review, test, and validate AI-generated outputs: code, reports, or automated decisions, and to perform human-in-the-loop oversight. This is the key track for human-in-the-loop policy goals, ensuring that AI deployment in professional environments is matched by a generation of workers capable of scrutinizing its outputs. Relevant for: software teams deploying AI-generated code, legal and compliance teams reviewing AI-drafted documents, and data teams.
- **Track 3 - AI Systems Integrator (Non-Technical):** Certifies the ability to configure, adapt, and manage AI tools integrated into workplace software: CRMs, ERP systems and AI-augmented platforms, without requiring deep coding skills. This track bridges directly to the SME Hub recommendation, as it targets the AI integration functions that smaller firms most need to fill but currently cannot verify in new hires. Relevant for: SME operations, project management, and client-facing roles.
- **Track 4 - Responsible AI Practitioner:** Certifies understanding of the EU regulatory landscape: AI Act risk tiers, GDPR interaction, bias detection basics, and the ability to flag compliance issues in AI deployments. This makes the holder a junior compliance resource, directly reducing the per-unit regulatory burden on firms that would otherwise have to rely on external legal support for routine AI governance tasks. Relevant for: legal, HR, and management roles.

Each track will therefore be earned through the study of specific coursework, developed either by EU bodies or in accredited collaboration with academic institutions, on condition of passing a standardized examination. Earning all four tracks grants the holder the Full Professional AI Proficiency Certificate.

It is important to note that the certification is targeted rather than universal. It is not a general AI training initiative for all workers at all times. Rather, it is specifically positioned as a market-entry credential for young graduates entering AI-affected occupations, activated through structured coursework and examination rather than passive participation. Track 1 and Track 3 are designed to be accessible without technical prerequisites, while Track 2 and Track 4

presuppose a degree-level background in relevant fields, a reflection of our previous conclusions that proportionality in the design of new obligations matters for both accessibility and credibility.

This certification could blend very well with the current landscape of educational and regulatory frameworks in two specific ways. First, it would be **anchored to DigComp 3.0**: the competence descriptors for each track are explicitly mapped to DigComp 3.0 areas, which gives it cross-border legibility and avoids creating a parallel standard. Second, it is designed to be **recognized under AI Act Article 4**: firms that hire a certified holder for an AI-relevant role can use the certificate as partial documentation of their Article 4 AI literacy obligation for that position, giving employers a direct and immediate incentive to value the credential. This last design choice is particularly important. The EDSC feasibility study found only moderate readiness across MS for a new EU-level digital skills certificate, and stakeholders consistently preferred a quality label model over a standalone new credential. By attaching employer value directly to an existing legal compliance obligation, the PAPC avoids the adoption risk that has historically limited the reach of voluntary digital skills frameworks. Rather than relying on voluntary uptake alone, the recommendation embeds demand creation into the legal architecture itself.

### **What Risk Dimensions Does it Improve?**

This recommendation primarily targets the **income security** dimension of the professional risk framework. Its core function is to improve the labour market insertion quality of young graduates in AI-exposed occupations, directly addressing the structural mismatch we previously identified between automation-driven task compression and the weakened entry pathways available to early-career workers.

More specifically, the PAPC improves income security through three distinct mechanisms. First, it fights structural entry barriers by converting the vague employer demand for “AI-competent new hires” into a verifiable signal, lowering the information cost of hiring young

workers for AI-relevant roles and reducing the effective hiring premium that currently accrues to experienced workers who have built AI competence through tacit exposure alone. Secondly, it can improve time-to-security for graduates entering AI-affected occupations: rather than spending the early years of employment navigating unstructured and self-directed learning of AI tools, certified workers can demonstrate role-readiness at the point of hiring, compressing the period between graduation and stable employment. Third, by targeting Track 1 and Track 3 at roles which require operational AI proficiency rather than deep technical skills, the certificate creates viable entry points into augmentation-oriented roles, precisely the category where, as we previously showed, entry-level employment has not declined. This shifts graduates away from the most automation-exposed positions and toward roles where AI complements rather than substitutes their contribution, directly targeting the displacement mechanism we identified as central to this risk.

Additionally, the recommendation produces meaningful secondary effects on the **autonomy and agency** dimension. One of the core concerns voiced throughout the report was that AI-related reorganization can make professional progression harder to read, weaken traditional career ladders, and increase anxiety for workers unsure whether their current skills will remain viable or whether alternatives will be available to them. The PAPC addresses this directly, as it replaces the opacity of undefined “AI skills” expectations with a legible, structured, and internationally recognized competence architecture. Track 4 in particular strengthens agency for young workers in AI-affected workplaces: by certifying basic regulatory literacy, it equips holders to understand the AI systems they work with, identify when deployments create compliance concerns, and exercise their rights under the AI Act and GDPR more effectively. In that respect, the recommendation supports autonomy not by shielding workers from AI-augmented workplaces, but by making those workplaces more intelligible and less arbitrary from the worker’s perspective.

A third, more indirect contribution falls under the **legal and fiscal risk** dimension. As we previously argued, one of the most acute legal-risk pressures on smaller and mid-sized firms is the compliance burden created by Article 4’s literacy obligation, which is already in force yet

leaves firms with no standardized way to demonstrate compliance. The PAPC provides a ready-made, low-cost compliance documentation tool, reducing the friction of AI Act compliance and lowering legal-risk exposure for both employers and their newly hired graduates. This is particularly true for Track 4, which creates a generation of entry-level workers capable of actively supporting compliance functions rather than requiring firms to maintain costly external legal support for routine AI governance tasks.

### **Who is Responsible for What?**

The Professional AI Proficiency Certificate needs to be structured as a shared-responsibility framework distributing design, delivery, accreditation, and incentive functions across EU institutions, MS, higher education providers, employers, and graduates themselves. It follows from the broader logic of the report: credentials only function as genuine market signals if all relevant actors, issuers, verifiers, and users, operate within a shared and coherent framework.

At the EU level, the primary responsibility is to define the certification architecture, provide the regulatory anchor, and connect the credential to the existing EU policy and funding infrastructure. Concretely, the EU, through institutions such as the AI Office, the Joint Research Centre, and the Digital Europe Programme, should define the competence descriptors for each track in explicit alignment with DigComp 3.0, establish the formal link between PAPC certification and Article 4 compliance documentation, and could use the AI Skills Academy as the operational launch vehicle for the initiative. Funding for curriculum development and accreditation infrastructure should flow through the Digital Europe Programme and ESF+ channels already available for advanced digital skills initiatives. Practically speaking, the EU's role isn't to centralize skills delivery, but to create a common credentialing spine that reduces fragmentation, provides standard competence descriptors and accreditation templates, and supports implementation through funding and monitoring.

MS would be mainly responsible for the operational integration side. Higher education systems, vocational qualification frameworks, and graduate employment services remain largely national, so MS would have to be responsible for the translation of the credential into functioning insertion pathways. The duties should include the accreditation of institutions authorized to deliver each track, the integration of PAPC coursework into relevant degree programmes or final-year electives where appropriate, and ensuring that certification completion is visible within national public employment service databases so that certified graduates are identifiable at the point of labour market entry.

This responsibility of educational institutions transforms into an opportunity when observing how many universities and professional institutions are moving towards AI integration in most steps of their teaching process with high levels of curiosity and commitment. From large, secular institutions such as Sciences PO Paris offering programs merging traditional programs with technological innovation and AI to new institutions such as Forward College redefining learning approaches by including AI and future design into curriculums centered around politics or economics, the European educational landscape seems more ready than ever to seize this opportunity.

Higher education institutions, vocational training providers, and accredited private learning bodies can also carry the delivery responsibility, alongside guidance and directional usage of the EU's centralized frameworks for each track. They must develop and maintain track-specific curricula in alignment with the EU-defined competence framework, ensure examination integrity, and update content in response to changes in AI tools and regulatory requirements. It is important to clarify that this doesn't amount to a requirement for all providers to deliver all tracks. Instead, it aims to create an accreditation logic whereby any institution wishing to deliver PAPC tracks meets a common quality standard, a function consistent with the report's broader preference for creating governance architecture that is enforceable through measurable procedures rather than aspirational commitments alone. Institutions already delivering DigComp-aligned digital skills

programmes, including those operating under the AI Skills Academy umbrella, should be the natural first cohort of accredited providers.

Employers and firms also carry a dual responsibility. As primary users of the credential signal, they should incorporate PAPC track recognition into hiring criteria for AI-relevant roles and, where applicable, use certification as part of their Article 4 AI literacy documentation. Larger firms and AI Pact signatories should be the initial adoption target, as their participation would generate the employer-side credibility needed for the certificate to function as a genuine market signal rather than a formal qualification with limited labour market weight. Additionally, firms, especially through sectoral associations, should contribute to periodic track content reviews to ensure that what is being certified reflects actual workplace AI needs rather than academic representations of them.

Worker representatives and consultation structures must also play an important role. Earlier on, we argued that many of the workplace AI risks are difficult to govern through individualized action because of the information asymmetries workers face inside firms. In the context of this recommendation, worker voice is relevant not only as a matter of procedural fairness, but also as a practical input for ensuring that track content, especially for Tracks 2 and 4, reflects the balance between employer compliance interests and workers' own rights and interests under the AI Act and GDPR.

On top of these, part of the responsibility also lies with public employment services and graduate support bodies. Where formal certification exists on paper, it can still have limited value if graduates can't access it in a timely, affordable, or geographically accessible way. Public employment services should ensure that certified graduates' credentials are visible in national referral systems, while ESF+-funded graduate support programmes should cover at minimum partial track costs for recent graduates from lower-income backgrounds or structurally weaker regions.

Lastly, it's important to keep in mind that this is only one way of imagining the delivery of such a framework. Tradeoffs between easy access and reputation as well as accountability need to be taken into account to design the most optimal system to allow this certificate to be available to anybody but also to be difficult enough to achieve to transform into a solid, top tier credential.

Together, these responsibilities are distributed in a way that matches the report's broader policy framing. The PAPC is neither a purely public skills programme nor a purely private hiring signal. Rather, it is a coordinated credentialing architecture where each actor is responsible for a different part of making the market signal work in practice.

### **Pros, Cons, and Mitigation**

The main strength of this recommendation is the way **it addresses the specific structural failure we identified** for young workers in AI-exposed occupations. The entry-level employment decline documented by Brynjolfsson et al. is not merely a consequence of automation, it is amplified by employers' inability to identify which young graduates are competent to perform AI oversight functions. Rather than attempting to reverse the automation trend or relying only on ex-post support for displaced graduates, the PAPC repositions young workers as credibly competent participants in AI-augmented workflows, directly targeting the role categories where entry-level employment has held up.

A second advantage of the recommendation is how it creates **simultaneous value for three distinct actors** without requiring a zero-sum trade-off.

- **Graduates** gain a legible market-entry credential.
- **Employers** gain a verified, low-cost compliance documentation tool for their Article 4 obligations.

- **The EU** gains an operational mechanism that connects its AI literacy mandate, its skills policy architecture, DigComp 3.0, the AI Skills Academy, the EDSC process, and its employment policy goals into a single coherent instrument.

This multi-stakeholder value design is particularly important, as it substantially increases the likelihood of meaningful uptake, as opposed to the voluntary adoption dynamics which have historically limited the reach of EU digital skills frameworks.

A third strength is that the recommendation **builds on existing EU instruments** rather than requiring an entirely new architecture. The PAPC's institutional value lies in connecting DigComp 3.0, the AI Skills Academy, the EDSC process, and Article 4 into a more coherent labour market insertion tool for young graduates, which makes it more politically and administratively plausible than a novel independent credential would be.

The first major limitation which the recommendation could encounter is **employer recognition risk**. A credential only functions as a labour market signal if hiring managers use it. The EDSC feasibility study found limited readiness across MS for a new EU-level digital skills certificate, and concluded that substantial efforts would be required to generate real employer uptake. If too few employers reference the PAPC in their hiring criteria, the certificate risks becoming a well-designed credential with low market penetration. A **key mitigation strategy** is the Article 4 compliance linkage: by giving firms a direct regulatory incentive to reference the certificate in hiring and documentation, the design builds demand creation into the legal architecture rather than relying solely on voluntary uptake. A secondary mitigation is phased rollout via AI Pact signatories before scaling to the broader employer population.

Secondly, there is the limitation of **quality fragmentation across MS**. If accreditation standards differ significantly across countries, the same certificate could reflect meaningfully different competence levels, undermining its cross-border legibility value. This risk can be mitigated by establishing a common EU-level quality assurance mechanism modelled on the

EDSC quality label framework, with periodic external audits of accredited providers and public reporting on examination pass rates and employment outcomes by provider.

A third weakness is **curriculum obsolescence**. The pace at which AI tools evolve means that a curriculum designed today around generative AI applications and current AI Act requirements could be significantly outdated within two to three years, a particularly acute risk for Tracks 1 and 2, which are directly tied to specific tool categories. This risk can be mitigated by building mandatory curriculum review cycles into the governance structure, linking these to the DigComp update schedule, and involving employer and sectoral representatives directly in the review process so that track content remains anchored in live labour market demand rather than fixed educational standards.

Taken together, the limitations above don't undermine the core recommendation, but they do define the conditions for its credibility. To be effective, the PAPC must be employer-recognized rather than merely institutionally endorsed, quality-controlled rather than accreditation-light, dynamically maintained rather than set once, and entry-accessible rather than financially exclusionary. In that form, it remains consistent with the report's broader principles of proportionality, operability, contestability, and policy learning.

## **KPIs and Data Sources**

To ensure that the recommendation is more than a declaratory commitment to closing the verification gap, it should be linked to a focused monitoring framework. In line with the broader logic of the report, our aim is not to create a long list of indicators, but a manageable set of measures which show whether young graduates are entering AI-relevant roles faster, more securely, and with greater capacity to navigate AI-augmented workplaces over time. The distinction between participation and outcomes is especially important here, given the well-documented risk of credential frameworks that generate enrolment without generating employment.

The first KPI group should focus on **labour market insertion outcomes** for certificate holders. Most importantly, we can focus on indicators like the employment rate of PAPC holders in AI-relevant roles at six and twelve months after graduation, and the median starting wage in those roles, disaggregated by track, institution type, and MS. The main data sources here would be national graduate tracking surveys and **Eurostat Labour Force Survey** microdata.

A second KPI group would capture **employer recognition and uptake**, not just whether the credential exists, but whether it is being used. The key indicators include the share of job postings for AI-relevant roles referencing the PAPC or DigComp-aligned credentials, the share of firms using PAPC documentation as part of their Article 4 compliance records, and the share of AI Pact signatories formally recognising the credential in hiring practices. The relevant data sources are employer panels, job posting analytics, and AI Office monitoring data from Article 4 compliance reporting cycles.

Thirdly, a group of indicators should be dedicated to assessing **track completion quality and curriculum effectiveness**. The indicators should include pass rates per track and per accredited provider, the share of completers employed in roles corresponding to their certified track, and the share of tracks updated in response to identified skill gaps within the preceding review cycle. Such indicators would help distinguish between meaningful certification and low-impact participation. The relevant data sources include certification body administrative data and periodic employer surveys conducted alongside track review processes.

Fourthly, a KPI group should focus on **accreditation coverage and access equity**. The important indicators include the number of accredited providers per MS, the share of eligible graduates in low-income or structurally weaker regions with access to at least one subsidized track, and the uptake rate of certification among graduates from underrepresented groups. The main data sources include national accreditation databases and MS reporting under ESF+ implementation frameworks.

A fifth KPI group should aim to capture the **credential's contribution to AI Act compliance**, specifically whether firms are using it as a practical Article 4 documentation tool. Useful indicators here include the share of firms' Article 4 documentation referencing PAPC-certified staff, and the share of AI Office inspection findings citing AI literacy gaps in firms that had not hired PAPC-certified individuals. Through these indicators, we would be able to tell whether the certificate is generating real compliance value or remaining a formal qualification without regulatory traction. The main sources would be the AI Office's enforcement monitoring database and the AI Act Service Desk's compliance reporting infrastructure.

Finally, a more indirect set of indicators could be used to monitor **autonomy and agency outcomes**. They could include the share of PAPC holders reporting that the certification meaningfully shaped their job search strategy, and the share reporting greater confidence in navigating AI-augmented workplaces after certification. These softer indicators can be gathered through programme follow-up surveys and complemented, where useful, by broader European survey infrastructures such as the EWCS.

Taken as a whole, the **Professional AI Proficiency Certificate responds directly to one of the report's most specific findings**: that entry-level employment in AI-exposed occupations has declined not because firms no longer need human oversight of AI, but because the mechanisms for identifying and signalling that competence in new graduates don't yet exist.

By creating a modular, DigComp-anchored, Article 4-linked credential, the recommendation turns that signalling gap into a solvable governance problem. Its value lies not in attempting to resist automation, but in ensuring that young workers entering AI-affected occupations can credibly demonstrate the oversight, integration, and compliance competences that AI-augmented workplaces increasingly require. In this way, it primarily strengthens income security for early-career workers while also improving labour market legibility, regulatory compliance capacity, and the professional agency of a generation navigating the employment consequences

of AI adoption for the first time.

### 10.3 R2 - AI Transition & Reskilling Guarantee

#### What is it?

Our second recommendation proposes an EU-backed transition framework that aims to make reskilling and job-to-job mobility the default response to AI-driven task automation: The **AI Transition & Reskilling Guarantee**. We base this recommendation on a core finding in our report: the main labour-market risk created by AI isn't limited to immediate job destruction, it is also about weakened transition capacity, compressed career ladders, and faster obsolescence of existing skills. Therefore, the resulting policy problem isn't just whether jobs disappear, but also whether workers and institutions can adapt quickly enough when tasks, roles, and internal labour-market structures are reorganized by AI adoption.

As such, our recommendation proposes a dual framework which combines worker-side entitlements with firm-side obligations. Focusing on the worker side first, we propose that individuals which are materially affected by AI-related restructuring would gain access to a funded transition bundle. It would include a skills assessment, career counselling, access to accredited training, and job-matching support, with temporary income support where necessary to make participation realistic. The aim here is to turn the abstract goal of upskilling into a more concrete and usable transition right. As previously discussed, workers face uncertainty not only about whether their role is becoming vulnerable, but also about which skills are worth acquiring, which adjacent roles are realistic, and how to move into them without bearing the full financial and informational cost of adjustment alone. Our recommendation intends to reduce this uncertainty by turning the transition from an individualized burden into a structured pathway.

Secondly, looking at the firm side, our policy would require larger firms to prepare a **Workforce Transition Plan** before or alongside AI implementation. The logic for this requirement is preventive rather than punitive. Based on our previous findings, we argue that firms should

be expected to anticipate negative effects of AI adoption on workers rather than treating worker adjustment as a residual social cost. As such, a Transition Plan should set out which functions are affected, how tasks are expected to change, what skills gaps are likely to emerge, which workers could be redeployed internally, what retraining measures are planned, and how workers will be informed and consulted. It is important to clarify that this doesn't amount to a ban on restructuring or a requirement to preserve every existing role unchanged. Instead, it aims to create a procedural duty whereby AI deployment is linked to a credible transition strategy rather than unmanaged displacement alone. This function is consistent with the report's broader preference for preventing governance where risks are organizational and foreseeable.

The recommendation is fortified by the combination of these two elements. Without employer planning, worker entitlements would simply mean that transition support is reactive, only coming into place following the relevant adjustment decisions that have already been made. On the other side, employer obligations with no worker-side support would risk becoming too formalistic; a well-designed transition plan could fail if workers lacked access to funded training, guidance, and income continuity during the adjustment period. This recommendation aims to transform AI-related restructuring from a purely private reorganization decision into a managed social transition.

It is important to note that the guarantee portion of the recommendations should be targeted rather than universal. Worker-side entitlement isn't intended as a general training voucher for all workers at all times. Instead, it is meant to act as a transition right activated where AI-related restructuring creates a material risk of displacement, occupational downgrading, or role obsolescence. Likewise, planning obligations should be strongest for large firms and cases of substantial workforce impact, a reflection of our previous conclusions that AI-related risks are distributed unevenly across firm types. This way, the recommendation preserves proportionality, while still ensuring that the burden of adjustment isn't externalized onto workers.

More broadly, the recommendation is meant to move EU policy away from passive adaptation and toward structured transition management. There already exist EU instruments which provide parts of the necessary architecture, including funding for reskilling and labour-market adaptation, support for displaced workers, and procedural frameworks for restructuring and worker consultation. Nonetheless, we previously argued that these tools remain fragmented and aren't yet integrated into a coherent response to AI-drive labour-market change. The AI Transition & Reskilling Guarantee would therefore not create an entirely new institutional logic from scratch; rather it would work to connect existing transition, skills, and restructuring levers into an operational framework tailored to the specific pressures created by AI in the workplace.

### **What Risk Dimensions Does it Improve?**

This recommendation primarily aims to target the **income security** dimension of the professional-risk framework. In the context of the dimension which we previously discussed throughout the report, the recommendation tries to improve transition capacity, which subsequently protects income security at its core.

More specifically, the recommendation is meant to improve income security in four different ways. First, it fights **displacement risk** by shifting the policy focus from ex post reactions to ex ante transition management solutions. With firms being required to map affected roles, identify retraining needs, and consider internal redeployment options before restructuring occurs, workers are less likely to become victims of abrupt and unmanaged exits from employment. Secondly, it can improve **time-to-security** as workers would have access to a concrete transition bundle instead of having to navigate reskilling and job search on their own. This is especially important, as the quality of labour-market adjustment is dependent not only on whether workers eventually find another job, but also how quickly and smoothly they do so. Third, it would limit **wage scarring**, as it would increase the chance that re-employment occurs through retraining into adjacent or upgraded roles rather than through downward occupational moves. Last but not least, it directly addressed **skill and knowledge obsolescence risk**, which we previously identified

as being a central mechanism through which AI increases professional risk. This is particularly true wherever firms can automate faster than workers are able to re-skill independently.

Additionally, the recommendation indirectly improves the dimension of **social security and downside protection**. One of the core concerns previously voiced in the report was that workers should not have to face sharp increases in risk just because they move between roles, employment conditions, or labour-market statuses during a period of technological adjustment. The structured transition guarantee would contribute to this objective, as it would reduce the likelihood of workers falling into prolonged gaps between jobs without support or being forced to absorb the full cost of retraining themselves. By including temporary income support, coordinated access to employment services, and a bridge into new employment, the transition bundle can help smooth shocks that AI-related restructurings could otherwise impose on workers. Essentially, the recommendations wouldn't replace social protection systems, but would instead strengthen their practical effectiveness by making transitions more manageable and less discontinuous. This last point is especially important for the logic used across the report, where "good" downside protection is defined as the workers' ability to change roles or status without encountering severe benefit cliffs or abrupt drops in security.

The recommendation contributes also in the **autonomy and agency** dimension, particularly through career-path predictability and workers' sense of control over professional change. We showed that AI-related reorganization can weaken traditional junior-to-senior pathways, make professional progression harder to read, and increase anxiety for workers as they become unsure whether their current role will remain viable or whether there will be alternatives open to them. The transition guarantee would help mitigate this problem, as it would replace uncertainty with a more legible adjustment process. While skills assessments, counselling, and redeployment pathways don't fully eliminate labour-market disruption, they do help in making it more understandable and navigable from the worker's perspective. In parallel, employer transition plans would reduce the opacity of restructuring by requiring firms to identify expected workforce effects in advance and to communicate them in a more structured way. This is important, as

professional agency isn't just about freedom inside the workplace; it is also dependent on workers' ability to anticipate likely transitions, understand their options, and make decisions under conditions of reduced uncertainty. In that respect, the recommendation would support autonomy not by blocking change, but by making adaptation more intelligible and less arbitrary.

As a whole, the recommendation can be best understood as having primary effects on income security and strong secondary effects on downside protection and autonomy. The distribution of effects is consistent with our previous diagnosis of AI risk across firm types. In large firms especially, the dominant problem isn't the lack of adoption capacity but more the risk that internal restructuring proceeds faster than workers can adapt, which weakens earnings stability, transition prospects, and career pathways. The AI Transition & Reskilling Guarantee would address that problem directly by improving the institutional capacity for adaptation. Instead of attempting to freeze labour-market change, it would try to ensure that AI-driven adjustment takes place through a structured, supported, and more predictable transition.

### **Who is Responsible for What?**

The AI Transition & Reskilling Guarantee needs to be structured as a shared-responsibility framework involving the EU, MS, firms, and implementation bodies. It follows from the broader logic from the report: AI-related transition risk doesn't emerge at single levels of governance. Immediate effects of AI-driven restructuring are visible inside firms, but the capacity to support workers through retraining, redeployment, and income continuity falls on public institutions, funding arrangements, and implementations systems which go beyond the firm itself. As such, the recommendations are framed neither as a purely employer-side duty nor as a purely worker-facing entitlement. Instead, it should act as a coordinated governance model where each level is responsible for a different part of the transition process.

At the EU level, a common policy framework must be provided as well as the enabling architecture within which national systems and firms can operate. The EU should be the

body which defines the broad principles of the guarantee, including what counts as AI-related restructuring, what minimum elements should be included in a worker transition bundle, and what information should be contained in a firm-level Workforce Transition Plan. It also falls under the EU's responsibility to link the recommendation to the relevant levers which we identified above, especially ESF+ and the European Globalization Adjustment Fund. Practically speaking, the EU's role isn't to centralize labour-market adjustment, but to create a common governance spine that reduces fragmentation, provides standard templates and guidance, and supports implementation through funding and monitoring.

MS would be mainly responsible for the operational delivery side. Labour-market institutions, public employment services, training systems, and social-protection administration remain largely national, so MS would have to be responsible for the translation of the guarantee into functioning support pathways. The duties should include the organization or accreditation of the bodies that provide skills assessments, career counselling, training access, and job-matching services; linking the guarantee to existing unemployment and active labour-market policy systems; and ensuring that workers have access to support quickly enough for transitions to be preventive rather than merely reactive.

Firms should be the primary responsibility holders for anticipating and managing the workforce consequences of their own restructuring decisions. As previously discussed, the main AI-related issue in large firms is that of internal reorganization proceeding faster than workers can adapt, which produces displacement pressures, stalled career ladders, and skill obsolescence. As such, there firms should be required to identify the roles and tasks which are likely to be affected by major AI deployments, assess where retraining or redeployment is feasible, and prepare Workforce Transition Plans that link AI adoption to concrete adjustment measures.

Worker representatives and consultation structures must also play an important role. Earlier on, we argued that many of the workplace AI risks are difficult to govern through simple atomized individual action because of the information asymmetries and structural power imbal-

ances that workers face inside firms. Exactly for this reason, worker voice is relevant not only as a matter of procedural fairness, but also as a practical input for better transition planning. These groups would play a vital role in identifying which tasks are genuinely vulnerable, whether management's skill-mapping is realistic, whether retraining measures correspond to actual future roles, and whether the internal mobility pathways proposed are credible.

On top of these, part of the responsibility also lies with public employment services, accredited training providers, and related implementation bodies. Where formal entitlement is present on paper, it can still have limited value if workers aren't able to access timely, relevant, and high-quality support in practice. Firstly, the public employment services should be central in coordinating skills assessments, referrals, and external job matching. Secondly, training providers would take the responsibility of ensuring that the programmes funded under the guarantee are linked to real labour-market demand rather than generic participation targets. Lastly, the sectoral or social-partner bodies could help in ensuring that training provisions reflect actual changes in skill needs in the exposed occupations.

Lastly, workers themselves would also carry notable responsibilities. There would be expectations for them to participate in skills assessments, engage with retraining and redeployment opportunities, and actually make use of the transition support which is made available to them.

Together, these responsibilities are distributed in a way that matches the report's broader policy framing. The allocations described above are important as the AI-driven labour-market adjustment is neither a purely private restructuring issue nor a purely public welfare problem.

### **Pros, Cons, and Mitigation**

The main strength found in this recommendation is the way it addresses the report's core diagnosis of AI-related labour-market risk: AI adoption not being the central challenge, but instead whether restructuring happens faster than workers are able to adapt. Instead of halting technological change or relying only on ex-post compensation, it aims to build transition

capacity directly into AI-related restructuring, making redeployment, retraining, and supported mobility more likely than unmanaged displacement.

A second advantage of the recommendation is how it improves multiple risk dimensions at once. It mainly affects income security, through lower displacement risk, less skill obsolescence, and better transition outcomes. Simultaneously, however, it also supports downside protection by reducing gaps during restructuring as well as strengthening autonomy and agency by making career transitions more predictable and less arbitrary. This is an especially important factor, as we previously argued that AI-related risks tend to compound across dimensions instead of appearing in isolation.

A third strength is that the recommendation builds on existing EU instruments rather than requiring an entirely new architecture. ESF+, the EGF, procedural labour-law frameworks, and the wider skills agenda already provide many of the relevant levers. The recommendation's institutional value lies in connecting these tools into a more coherent transition framework for workplace AI, which makes it more politically and administratively plausible.

The first major limitation which the recommendation could encounter is cost. Meaningful transition guarantees would require funding for training, counselling, job matching, and in some cases temporary income support. Due to the scale of Europe's reskilling challenge which was identified earlier in the report, the costs could become high while outcomes remain uneven.

Secondly, there is the limitation of administrative burdens and compliance resistance. Firms could interpret transition plans as yet another reporting requirement on top of existing obligations. If too rigid, the measure could create more formal paperwork rather than meaningful governance. A key mitigation strategy could be differentiated implementation: stronger planning duties for large firms and major restructuring cases, combined with lighter, standardized pathways for smaller firms through templates, simplified documentation, and implementation support.

A third weakness is the risk of low-quality training. The reskilling policy could easily become participation-heavy but remain outcome-light if workers are channelled into generic courses with little connection to actual labour-market demand. This risk can be mitigated by making the guarantee outcome-oriented rather than participation-oriented. Training should be tied to skills assessments, adjacent roles, and demand in the labour market. Providers should be accredited, and the programmes themselves should be monitored through re-employment and wage outcomes instead of enrolment alone.

Taken together, the limitations above don't undermine the entire recommendation, but they do help in defining the conditions for its credibility. To be effective, the guarantee must be targeted rather than universal, differentiated rather than one-size-fits-all, outcome-oriented rather than participation-oriented, anticipatory rather than purely reactive, and enforceable through measurable procedures rather than aspirational commitments alone. In that form, it remains consistent with the report's broader principles of proportionality, operability, contestability, and policy learning.

### **KPIs and Data Sources**

To ensure that the recommendation is more than a simple declaratory commitment, it should be linked to a focused monitoring framework. In line with the broader logic of the report, our aim is not to create a long list of indicators, but a manageable set of measures which show whether AI-related transitions are becoming faster, fairer, and less damaging over time.

The first KPI group should be focused on transition **speed and re-employment**. Most importantly, we can focus on indicators like average time between AI-related displacement or restructuring exposure and re-employment, and the share of affected workers re-employed within six and twelve months. The main data sources here would be **Eurostat labour market transitions** and **national administrative employment records**.

A second KPI group would capture **transition quality**, not just speed. Having rapid re-employment is not enough if it leads to systematic downgrading. As such, the key indicators include post-transition wage change and the share of workers re-employed in roles of comparable or higher skill level. The relevant data sources are **EU-SILC** and, where possible, **national earnings and social-security records**.

Thirdly, a group of indicators should be dedicated to **assessing reskilling effectiveness**. The indicators should include the share of workers completing funded training who enter employment within a defined period, and the share moving into roles linked to the training received or into shortage occupations. Such indicators would help distinguish between meaningful reskilling and low-impact participation. The relevant data sources include **public employment service records, administrative training data, and national programme evaluation systems**.

Fourthly, a KPI group should be focused on **tracking firm behaviour and implementation quality**. The important indicators include the share of covered firms submitting Workforce Transition Plans, the share of affected workers redeployed internally, and the share of AI-related restructuring cases involving documented worker consultation before implementation. The main data sources include **firm reporting, labour inspectorates**, or equivalent authorities, and **administrative monitoring systems**.

A fifth KPI group should aim to capture the **continuity of support and downside protection**. Useful indicators here include the share of affected workers receiving support before or immediately upon exit, the share experiencing income interruption during transition, and the take-up rate of transition bundles among eligible workers. Through these indicators we would be able to tell whether the guarantee is reaching workers early enough and smoothly enough to prevent avoidable insecurity. The main sources would be **national administrative data, public employment services, and social-protection delivery systems**.

Finally, a more indirect set of indicators could be used to **monitor autonomy and agency success**. They could include the share of workers reporting that they received timely information

and clear guidance about their transition options, and the share reporting that they understood the reasons for restructuring and the pathways available to them. These softer indicators can be gathered through **programme follow-up surveys and complemented**, where useful, by broader **European survey infrastructures** such as the EWCS.

Taken as a whole, the AI Transition & Reskilling Guarantee responds directly to the report's core finding that AI is creating a managed-transition problem, especially in larger firms where organizational change can outpace workers' ability to adapt on their own. By combining worker-side entitlements with firm-side planning duties, the recommendation shifts the policy response from passive adaptation towards structured transition management. Its value lies not in trying to stop technological change, but in ensuring that AI-driven restructuring is matched by credible pathways into retraining, redeployment, and re-employment of acceptable quality. In this way, it primarily strengthens income security while also improving continuity of protection, career predictability, and the procedural quality of adjustment. If designed in a targeted, proportionate, and measurable way, it would provide a practical mechanism for making AI-related labour-market disruption more governable and less socially costly.

## **10.4 R3 - SME AI Capability & Compliance Hub**

### **What it is**

Our third recommendation is called the **SME AI Capability & Compliance Hub**, a practical support architecture meant to address the report's core diagnosis of the SME problem in workplace AI. As previously discussed, SMEs aren't mainly held back by an absence of incentives to adopt, but rather by a combination of weaker digital readiness, higher fixed compliance costs, and a greater exposure to fragmented implementation environments. As such, the relevant policy need isn't simply more rules or lighter rules, but a usable pathway which makes lawful adoption, organizational adjustment, and capability-building easier for firms with limited internal capacity.

In light of the above, the recommendation proposes an SME-facing hub meant to combine four elements into a single support framework: **compliance toolkits, help-desk support, capability-building assistance, and an outcome-linked voucher scheme**. The main goal is that of reducing the fixed cost of productive and compliant AI adoption for SME without weakening substantive safeguards. The logic is fully consistent with the report's broader argument that SME-oriented policy should have standardization, practical guidance, and low-friction implementation tools as a priority instead of assuming that formal obligations alone can generate convergence.

The first pillar of the hub comprises a set of compliance toolkits for common SME AI use cases. They would include practical materials like standard checklists, minimum documentation sets, basic procurement guidance, risk-assessment templates, and model internal procedures for the adoption of AI tools in ordinary business functions. The goal here is to translate broad legal and governance expectations into repeatable, low-friction procedures that SMEs can actually use. This pillar aims to tackle a central theme identified in the report: many workplace AI governance problems aren't caused by the absence of rules, but by the difficulty of turning abstract legal obligations into workable organizational practice, especially for firms that don't have dedicated legal, HR, or compliance teams.

Moving onto the second pillar, we have the help-desk and support function. Even if templates exist thanks to the first pillar, many SMEs will still find it difficult to determine which guidance applies to their use case, whether they need additional review, or how to sequence adoption steps in practice. With this in mind, a help-desk provides a first layer of support, offering clarity on which toolkits are relevant, answering standard implementation questions, and referring firms to more specialized assistance wherever necessary. In the fragmented context of legal uncertainty that SMEs have across MS, support is valuable not only because it lowers costs, but also because it lowers uncertainty and reduces reliance on expensive external advice for standard use cases.

The third pillar of our recommendation would be capability-building support, with the purpose of helping SMEs develop the internal capacity needed to adopt AI tools well instead of superficially. It would consist of subsidized staff training, practical AI literacy support, assistance with workflow adaptation, and guidance on the organizational changes required to use AI tools effectively.

The final pillar would be a targeted voucher scheme, which would help SMEs adopt compliant AI tools, access implementation support, and train staff on those tools where internal resources are missing. It is important to note that the voucher must not operate as a generic technology subsidy. Remaining consistent with the report's logic, support must be tied to concrete outcomes like the successful deployment of a compliant tool or the completion of staff training linked to actual use. This distinction is especially important, as we previously stated that AI policy should distinguish between reducing productive compliance costs and simply subsidising adoption without sufficient accountability. As such, the value of the vouchers doesn't lie in their ability to lower costs, but in the way they can steer SMEs toward compliant and capability-enhancing adoption pathways.

On a broader level, the hub would function as an implementation infrastructure, not merely a funding window. We have treated the SME problem as one of constrained absorption under asymmetric burdens, and this recommendation responds to this by lowering the cost of interpretation, coordination, and initial capability-building rather than by weakening the underlying legal floor.

Taking the pathways above together, the SME AI Capability & Compliance Hub would be able to convert a fragmented set of demands into a usable support pathway for smaller firms. By refusing to assume that SMEs are able to absorb the same implementation burdens as large organizations, it can provide them with the practical tools, advisory support, capacity-building, and targeted financial assistance that would lead to real outcomes. As such, its purpose arises not only as an accelerator for AI uptake, but to make SME adoption more productive, more

compliant, and less dependent on ad hoc external advice or informal workarounds.

### **What Risk Dimensions Does it Improve?**

This recommendation primarily improves the **legal and fiscal risk** dimensions of the professional-risk framework. This dimension was previously identified as not being limited to liability after the harm occurs; it also includes the cost of legal uncertainty, the administrative burden of compliance, and the operational difficulty of understanding and applying rules in practice. This is one of the clearest pressure points especially for SMEs: facing higher per-unit compliance costs, weaker in-house capacity, and greater difficulty translating broad rules into workable day-to-day practice. The hub would directly address these issues by reducing uncertainty, lowering dependence on external advice for standard cases, and providing practical compliance pathways using toolkits, help-desk support, and standardized guidance. In this sense, legal and fiscal risks are amortized not by weakening the safeguards, but by making them more easily navigable by SMEs.

A second major effect is that on **income security**. For SMEs, the relevant problem isn't primarily large-scale restructuring, but the risk of falling behind if firms can't adopt AI productively while larger competitors move ahead. With capability gaps and compliance frictions blocking adoption, the result can be weaker productivity, lower growth, reduced hiring, and greater long-term instability for workers. As such, by assisting SMEs in adopting AI more effectively and lawfully, the hub can support firm resilience and therefore also the income security of those working within them.

Additionally, the recommendation has a supporting effect on **autonomy and agency**. While it isn't the main target here, it is still relevant as smaller firms often adopt tools in a more improvised way, with fewer formal procedures and less internal oversight. The hub would encourage better tool selection, staff training, and basic governance steps, making AI use more structured and intelligible, reducing arbitrariness and preserving a clearer role for human oversight.

Lastly, its effect on **social security and downside protection** is the most limited and indirect. While the hub isn't a social-protection measure in itself, it does improve SME resilience and reduces avoidable business stress. This may also lower the likelihood that workers are exposed to abrupt shocks that then have to be absorbed through existing protection systems.

Taken together, the SME AI Capability & Compliance Hub is best understood as a recommendation with **primary effects on legal and fiscal risk, strong secondary effects on income security, and more limited supporting effects on autonomy, agency, and resilience**. This matches the report's broader diagnosis that SMEs are mainly a capability-and-compliance problem, and that effective policy should reduce the fixed cost of lawful AI adoption without lowering the common governance floor.

### **Who is Responsible for What?**

From the previous arguments in the report, it follows that the SME AI Capability & Compliance Hub should be organized as a layered governance instrument where the EU sets the common architecture, MS manages delivery, intermediary support bodies provide operational access, and SMEs implement the supported measures. Such an architecture reflects the report's broader conclusion that the main SME problem isn't the absence of formal rules, but the gap between formal obligations and a usable compliance pathway.

At the EU level, the main responsibility lies in defining the common structure of the hub and anchoring it in the Union's existing implementation architecture. This would include the production of the standardized toolkits, minimum documentation sets, model procedures, and common guidance that the report previously identified as especially significant for SMEs. The main legal and governance reference point is the wider AI Act, which can be supported by the Commission's implementation work through the AI Office, Commission guidelines, the AI Act Service Desk, and the emerging standardization and simplification agenda. In practice, the EU

wouldn't have to run the support service firm by firm, but simply provide the common templates, baseline expectations, and interpretative consistency that reduce fragmentation across MS.

Additionally, the responsibility for the funding architecture of the hub falls to the EU. As previously mentioned, the most relevant instrument here is ESF+, which we identified as the Union's main funding instrument for employment, skills, and adaptation capacity. Specifically for this recommendation, it would be the most plausible vehicle for co-financing subsidized training, advisory support, and parts of the voucher scheme. This last point is especially true wherever support is linked to staff capability-building and practical AI adoption. Broadly speaking, the recommendation is also aligned with the logic of the Digital Decade Policy Programme 2030, which treats SME digitalization as a strategic objective and provides the wider policy rationale for EU-level capability support.

The main role of MS would be operational delivery and institutional integration. Business support systems, administrative interfaces, and SME assistance structures are largely national, so MS face the responsibility of embedding the hub into their existing support ecosystems. That is to say that they would need to adapt the EU's common templates and support to domestic administrative practices where needed, creating a clear link between the hub and national SME and digitalization programmes.

In addition, an important delivery role falls upon intermediary support bodies, especially the EDIHs and equivalent national SME-support structures. We previously identified EDIHs as useful because they connect technical adoption, organizational change, and skills support rather than treating them as separate problems. In practice, they would be most suitable for the operational front-end of the hub: helping to host the help-desk function, distributing toolkits, providing first-line advisory support, and other functions.

SMEs themselves would bear the responsibility of using the hub in good faith and for implementing the minimum governance and capability steps which are attached to the support they stand to receive. This would include engaging with training and advisory support and adopting

the relevant documentation or procedural steps wherever the toolkit requires them. It stems from this responsibility that the recommendation is supportive, but not unconditional. Firms won't be asked to replicate the compliance architecture of larger organizations, but there would be an expectation to implement proportionate internal measures once support is made available.

Taken together, this allocation of responsibilities reflects the policy logic developed earlier in the report. The EU provides the legal and implementation spine through the AI Act ecosystem, common guidance, standardization, and ESF+-backed support; MS integrate and deliver the hub through national systems; EDIHs and similar intermediary bodies act as the practical access point for firms; and SMEs are responsible for using the support to build capability and implement proportionate governance steps. In this form, the hub is not simply a funding scheme or an information portal. It is a structured implementation instrument aimed at reducing the fixed cost of lawful AI adoption for smaller firms without lowering the underlying accountability framework.

### **Pros, Cons, and Mitigation**

The main strength of our recommendation is the way it responds directly to our previous diagnosis of the SME problem. We identified how smaller firms are held back less by a lack of interest in AI than by weaker internal capacity, higher fixed compliance costs, and fragmented implementation environments. The recommendation therefore presents itself as well targeted: rather than extending large-firm-style obligations to SMEs in a one-size-fits-all way, it offers support, standardized pathways, and low friction implementation tools that are useful for smaller firms.

Secondly, the recommendation can **reduce unproductive compliance costs** while maintaining safeguards intact. By offering templates, guidance, help-desk support, and standardized documentation, the hub can lower the cost of navigating rules while preserving the common

governance floor. Through these mechanisms it will support both adoption and accountability simultaneously.

The first major limitation is the **risk of the hub becoming a generic subsidy scheme** rather than a targeted governance tool. If support is too generic or vouchers are too broad, adoption may be subsidized without a clear improvement in compliance or capability. A way to dampen this risk would be to make the support outcome-linked: like vouchers being tied to concrete steps such as completing staff training linked to actual use or implementing minimum governance procedures.

A second limitation is the **risk of low take-up**, especially among the smallest or least digitally prepared firms. In the case that access is too complex or support is poorly communicated, it would be the firms most in need that wouldn't be able to use it. A mitigation strategy here could be a one-stop-shop design with simple onboarding and delivery through familiar intermediary bodies such as EDIHs, SME agencies, or chambers. Additionally, part of the support could be reserved for first-time adopters and lower-digital-intensity firms.

Thirdly, there is the possibility of **checkbox compliance**. While templates and simplified toolkits are useful, they can also encourage firms to treat governance as a box-ticking exercise. A possible mitigation here is to combine the standardization with minimum-consent requirements, and light monitoring. In such a way, simplification reduces friction without hollowing out accountability.

Last but not least, there is serious **risk of uneven implementation across MS**. In the case that national versions diverge too much, the hub would be reproducing the exact fragmentation it's meant to reduce. This can be mitigated using a layered model: common EU templates and baseline standards, combined with national flexibility only in delivery and institutional embedding.

Putting these limitations together, we can clarify the conditions that our recommendation needs for credibility. A well-functioning hub must be targeted rather than broad, easy to access

but not shallow, standardized in core design but adaptable in delivery, and tied to measurable outcomes rather than uptake alone. With these characteristics, it remains consistent with the report's wider principles of proportionality, operability, and policy learning.

## **KPIs and Data Sources**

As with our other recommendations, to ensure that the SME AI Capability & Compliance Hub acts as more than a support scheme in name only, it must be tied to a focused monitoring framework. Our aim shouldn't be to build a very large indicator set, but to track whether the hub actually reduces capability gaps, lowers compliance friction, and supports more productive and lawful AI adoption among smaller firms. As such, we will focus primarily on legal and fiscal risk and capability uptake, with secondary attention to business outcomes and worker-facing quality safeguards.

The first KPI group we define is the one that should **measure hub uptake and capacity-building**. The simplest indicators would be the number and share of SMEs using the hub, the number receiving vouchers or support services, and the number completing subsidized training or advisory pathways. These indicators are especially helpful as they would help us identify whether the hub actually reaches smaller firms, especially those with weaker internal capacity. The main data sources here would be **administrative record of the hub, EDIH or national delivery data, and programme-level reporting by intermediary bodies**.

Our second KPI group should focus on **compliance usability**. The recommendation is partly justified by the need to reduce unproductive compliance costs, so useful indicators include the share of supported SMEs adopting minimum governance or documentation steps, the time needed to complete common compliance processes, and user-reported clarity of compliance requirements after support. Additionally, it should measure the share of firms reporting reduced need for external advisory support in standard use cases, since it would directly show whether

the hub is actually lowering consultant dependence. The main data sources here would be **firm surveys, hub reports, and targeted programme evaluation**.

Thirdly, a group should capture **adoption and operational outcomes**. Here, the central question is whether support through the hub actually helps firms move from interest in AI to successful and compliant use in practice. The relevant indicators are the share of supported SMEs successfully deploying AI tools, the share of reporting productivity or efficiency improvements, and the share which continues to use supported tools after voucher-backed support ends. On top of this, it could be useful to track the rate of failed or abandoned adoption attempts, since one of the main concerns we previously identified was that SMEs often face fragile or partial adoption under constrained conditions. The measures would rely mainly on **programme follow-up surveys and self-reporting by supported firms**.

A fourth KPI group should address **implementation and quality targeting**. As previously discussed, one of the risks of the recommendation is that it may mainly benefit firms which are already relatively capable. As such, the framework should track the share of support going to micro and small firms, the share going to first-time adopters, and the share reaching lower-digital-intensity firms. These indicators would help show whether the hub is actually reducing asymmetries or simply reinforcing existing ones. The most relevant sources here are once again **administrative data from the hub and programme monitoring records**.

Finally, a lighter KPI group should focus on **worker-facing quality safeguards**. Despite this recommendation not being a workplace-governance measure explicitly, it is clear that adoption support shouldn't come at the expense of basic oversight and intelligibility. For this reason, it would be useful to track the share of supported firms that provide staff training before redeployment, the share adopting basic human-oversight or governance steps, and, where feasible, the share reporting that workers were informed or consulted before AI tool rollout in relevant use cases. Data here would have to come mainly from **firm self-reporting, hub monitoring, and occasional programme audits**.

Taken together, the KPI framework for this recommendation aims to remain focused on a small core of indicators, with the most relevant data sources being administrative records of the hub, EDIH and national delivery data, firm surveys, programme evaluations, and national or EU digitalization monitoring tools. This way, the recommendation remains aligned with the broader principle of the report that policy should be measurable and open to adjustment over time, without losing proportionality.

Overall, the SME AI Capability & Compliance Hub would respond to the report's core finding that smaller firms face a distinct combination of weaker adoption capacity, higher fixed compliance costs, and greater exposure to fragmented implementation environments. Its main values are found in the fact that it doesn't lower the governance floor, but instead reduces the cost and complexity of reaching it for firms with limited internal capacity. In this sense, the hub acts as a differentiated SME policy response that primarily improves legal and fiscal risk while also supporting competitiveness, income stability, and more proportionate implementation of workspace-AI governance.

## **10.5 R4 - AI Augmentation Incentive Standard**

### **What is it?**

The fourth recommendation closes what is a very important gap in the policy framework as it currently stands. Our report shows incentivising AI integration strategies that augment human work rather than replace it as an explicit policy goal, yet not one of the preceding three recommendations operationalizes it. R4 does exactly that, through an **AI Augmentation Incentive Standard (AAIS)**: a taxonomy style classification mechanism that makes firms' augmentation versus replacement strategies visible, measurable, and fiscally consequential.

The case for acting here is not speculative. Brynjolfsson et al. (2025), drawing on high-frequency payroll data covering millions of workers across the US economy, find that employment declines are concentrated in occupations where AI use is primarily automative, while occupa-

tions where AI use is primarily augmentative have continued to see employment grow. The divide is not between workers who are exposed to AI and those who are not; it is between firms that have chosen to deploy AI as a replacement tool and those that have deployed it as an extension of human capability. The Volkswagen case study illustrates what the second path looks like in practice: combining AI-enabled cobots with human operators, using AI to sharpen decision-making rather than remove decision-makers, VW targeted a 30% productivity gain and a 30% reduction in factory costs without needing displacement to achieve either. Against that backdrop, the fact that 25% of European employers have already undertaken or planned AI-driven workforce reductions is not evidence of an inevitable structural shift. It is evidence of a choice that current policy gives firms no fiscal reason to make differently.

The AAIS borrows its structural logic from the EU Green Taxonomy. Just as the Taxonomy creates a classification standard for sustainable investment that then unlocks a range of fiscal instruments and preferential financing, the AAIS creates a classification standard for AI deployment strategy that firms can use to demonstrate augmentation alignment and access a progressive set of policy benefits. The mechanism rests on five tools that are designed to work together.

**Annual disclosure standard.** Large firms with significant AI deployments would be required, within their existing annual reporting obligations, to disclose a set of augmentation KPIs across affected roles, specifically:

- Headcount evolution at 6 and 12 months post-deployment;
- The share of affected workers redeployed internally versus exited; and
- Training investment per affected worker.

The underlying logic is rewarding transparency. Firms can only access the benefits of being “augmentation focused” if that alignment is verifiable, and as of today corporate reporting gives workers, investors, and regulators no way to distinguish firms committed to growing with their workforce from those using AI investment as a restructuring vehicle in a non sustainable way.

- **Fiscal differentiation within the existing R&D tax credit framework.** Firms that demonstrate augmentation alignment through their annual disclosure qualify for a top-up rate on qualifying AI-related R&D expenditure. This is not a hiring mandate. It does not prevent firms from restructuring where restructuring is warranted. It simply prices augmentation as the preferred path for AI investment, in the same way the US Inflation Reduction Act's prevailing wage multiplier prices responsible labour practice as the preferred path for clean energy investment. The top-up is intentionally modest and scoped tightly to AI R&D expenses, so it works with the existing tax credit architecture rather than around it.
- **EU public procurement preference.** Augmentation-aligned firms receive preferential standing in EU and MS procurement bids involving AI services or digital transformation contracts. This matters because it gives the classification immediate commercial value that extends beyond the tax benefit, particularly for firms whose business model depends significantly on public sector contracts.
- **EIB/EIF preferential lending access.** Augmentation-aligned firms become eligible for European Investment Bank and European Investment Fund instruments at preferential rates for qualifying AI investment. For mid-sized firms where the R&D deduction differential is too narrow to shift behaviour on its own, access to cheaper capital is likely to be the more decisive lever.
- **Augmentation Employer certification.** Firms that meet the alignment threshold through their disclosed KPIs receive a certification from the relevant authority. This can serve as a market signal to workers, graduates, and investors, and its reputational value adds even more benefits over the financial incentives.

The AAIS does not ask for a new institutional architecture. The disclosure obligation sits naturally within CSRD reporting already in force for large firms. The fiscal top-up operates through existing R&D credit mechanisms at MS level. The procurement preference follows a model already embedded in EU procurement directives. The EIB access replicates the channel the Green Taxonomy already uses to direct preferential financing toward aligned investment.

What the AAIS adds is the connective tissue: a single augmentation alignment classification, grounded in annual workforce data, that activates each of these instruments at once and gives firms a concrete reason to build augmentation into their AI strategies from the start, not after the workforce consequences have already landed.

### **What Risk Dimensions Does it Improve?**

This is the only recommendation in this framework that operates **preventively on the income security dimension**. The second recommendation supports workers once a transition is already underway. The AAIS intervenes earlier, at the moment firms are designing their AI deployment strategies, and tries to make the augmentation path more attractive before displacement becomes the default. If it works as intended, it reduces the volume of transitions that R1 and R5 then have to absorb.

The other effect is on **autonomy and agency**. AI deployment focused on augmentation strategies, as both the VW case and Brynjolfsson's data show, tends to preserve space and decision authority for workers in affected roles rather than engineering those roles away. A policy framework that incentivizes augmentation over automation therefore also becomes a framework that protects the conditions under which professional autonomy remains meaningful, which sits at the core of the agency dimension in our risk framework.

### **Who is Responsible for What?**

The **Commission** owns the framework definition. Its job is to establish what augmentation alignment means in practice: the KPI set, the disclosure format, the threshold above which firms are required to report and the alignment criteria that determine access to each of the five tools. This sits comfortably within the Commission's existing competence over corporate reporting standards (through CSRD), state aid and R&D credit coordination, and EU procurement directives. The Commission is also responsible for coordinating with the EIB and EIF to formalize

augmentation-aligned status as a qualifying criterion for preferential lending instruments, following the same channel that the Green Taxonomy uses for sustainable investment.

**MS** are responsible for implementation within their own fiscal and institutional architectures. Concretely, this means designating the competent authority responsible for awarding the Augmentation Employer certification, integrating the top-up rate into their national R&D credit frameworks within the parameters set at EU level and including the procurement preference into their national public procurement practices. MS also carry responsibility for enforcement of the disclosure standard, in a consistent way with how they currently enforce CSRD obligations for large firms.

**Firms** are the active participants. Large firms above the defined threshold compile and submit their augmentation KPIs as part of their annual reporting cycle, apply for certification where they qualify, and design their AI deployment strategies with the incentive structure in mind. The obligation is disclosure-based, not outcome-based, which means that firms are not required to achieve a specific augmentation score, only to report honestly against the defined KPIs. The policy framework then prices the results accordingly.

**Social partners** play an advisory role at the framework design stage, particularly in defining what counts as meaningful internal redeployment versus a managed exit, and in setting the training investment threshold that can qualify as a real upskilling commitment. Their involvement at that stage matters because it is where the risk is the highest: without the right input from worker representatives, firms could structure redeployment or training programmes that satisfy the framework of the KPIs without reflecting any real commitment to augmentation.

### **Pros, Cons, and Mitigation**

The case for the AAIS rests on structural advantages worth stating clearly before addressing the risks.

**It acts where no current recommendation does.** None of R1 to R3 touches firm-level AI deployment strategy. The AAIS is the only mechanism in the framework that tries to shape how companies integrate AI before the workforce consequences have already landed, rather than managing those consequences afterwards.

**It works with existing infrastructure.** Because it relies on CSRD reporting, existing R&D credit frameworks, and procurement directives already in place, it does not require building a new enforcement architecture from scratch. The policy asks firms to do more with instruments they already interact with.

**It creates compounding incentives.** The five tools do not just stack on top of each other, but rather reinforce each other. A firm that pursues augmentation alignment unlocks a tax benefit, cheaper capital, procurement preference, and a public certification at once, making the aligned path progressively more attractive the more a firm commits to it.

There are however two main risks that need to be acknowledged to find the right mitigation strategies.

- **Classification risk.** The line between genuine augmentation and a restructuring exercise dressed up as one is not obvious. A firm could technically maintain headcount in affected roles while quietly shifting those workers to peripheral tasks unrelated to the AI deployment, satisfying the disclosure KPIs without any real commitment to augmenting human capability. This is the most serious design risk in the recommendation.
- **Mitigation:** The KPI set needs to be deliberately focused on outcomes rather than intentions. These are harder to dress up than metrics based on activities such as “number of AI training sessions delivered”. Social partner involvement at the framework design stage, as described in the “Who Is Responsible” section, is also critical here, as worker representatives are well positioned to flag cases where redeployment is nominal rather than genuine.
- **Disclosure burden.** Adding a structured augmentation disclosure requirement on top of CSRD and other reporting obligations risks creating a compliance exercise that large firms

manage mechanically and that becomes a source of operational friction rather than a genuine behavioural signal.

- **Mitigation:** The disclosure standard should be integrated into existing CSRD reporting formats wherever possible, rather than designed as a standalone instrument. The goal is to add specific fields to a report firms are already producing, not to create a new reporting cycle. Keeping the KPI set short and output-focused, as proposed here, is the most practical way to manage this.

## 10.6 R5 - AI Entrepreneurship Bridge

### What is it?

As this report has argued throughout the entrepreneurship section, AI is already doing something structurally important: it is lowering the fixed costs of starting a business, compressing the time it takes to reach a first proof of concept, and allowing very small teams to run operations that would have previously required a much larger workforce. The opportunity is real. The barrier that prevents most workers from acting on it is not ambition or capability; it is the social protection cliff that opens up the moment someone leaves employment. This recommendation is designed to remove that barrier for the specific group of workers for whom the opportunity is most urgent and most relevant, viz. those whose careers have already been disrupted by AI.

The target group covers two types of workers. The first is those who have been formally displaced, meaning actually laid off as a direct consequence of AI-driven restructuring in their firm. The second is those whose roles have been significantly degraded by AI integration, not laid off, but working in positions where the tasks have been hollowed out to the point where the skills no longer carry a meaningful future. Both groups are documented in this report's analysis. Both face the same structural problem: the transition to entrepreneurship is the most logical path forward for many of them, but the loss of social protection coverage at the moment of transition makes it a leap that most people are rationally unwilling to take.

The **AI Entrepreneurship Bridge** addresses this through a structured three-phase programme, tied to existing social security mechanisms, that provides an additive benefit to support workers through the transition rather than asking them to jump without a safety net.

**Phase 1** is a short preliminary window of roughly two to three months, and it does not count toward the programme clock. Its purpose is orientation. The worker gets access to R3's hub infrastructure, including reskilling courses and compliance toolkits, to help them figure out what kind of business is feasible given their background, what the AI-augmented landscape looks like in their area of expertise, and whether entrepreneurship is genuinely the right direction for them. This is not a commitment phase, and there is no business plan required to enter it.

**Phase 2** begins once the worker presents a basic business plan and some contextual information about their intended activity to the relevant national authority. This is not a complex document; it is a straightforward declaration of what the business is, what market it operates in, and what the first 12 months of activity are intended to look like. Approval to move into Phase 2 activates the full benefit, including the additive social protection coverage.

**Phase 3** is the implementation window, running for up to 12 months from Phase 2 activation. During this period the worker receives the additive benefit on top of whatever existing entitlements they already hold, adjusted based on their individual situation and national context. At the end of 12 months there is a reconsideration point. For the majority of participants, the programme ends there. For a small subset of profiles where the business shows genuine activity but needs a limited additional runway, an extension may be considered on an exceptional basis. The programme is not designed to be open-ended, and the reconsideration is a threshold, not an automatic continuation.

The benefit is deliberately designed as additive and context-dependent rather than as a flat new entitlement. A worker who has been formally laid off and is already entitled to unemployment insurance under their national system will have that entitlement extended and conditioned on entrepreneurship activity, with a top-up structured around the gap between the insurance

coverage and what a basic operating standard requires. A worker who left voluntarily before formal displacement but whose role degradation is demonstrable will access a lighter version of the same structure. The precise calibration is left to MS within parameters set at EU level, consistent with how most social security coordination already works.

### **What Risk Dimensions Does it Improve?**

This recommendation directly targets the **social security and downside protection** dimension. The social protection cliff is the single most concrete barrier to entrepreneurship identified in this report's entrepreneurship section, and R5 addresses it head on by ensuring that the transition moment, the period between leaving employment and reaching the first stable revenue, is no longer a period of total exposure.

It also meaningfully improves **income security**, specifically the time-to-security factor. As the related section argues, AI is already compressing the time it takes to test an idea and reach early revenue. What AI cannot do is bridge the gap between the day a worker leaves their job and the day that early revenue arrives. R5 does exactly that, lowering the financial volatility of the early entrepreneurial phase and making it possible for workers to take the decision to start without betting their entire personal financial cushion on it.

### **Who is Responsible for What?**

The **Commission** sets the framework parameters: eligibility criteria for both the displaced and degraded worker categories, the minimum structure of the three phases, and the benefit calibration guidelines that MS work within. It also coordinates with the European Labour Authority to ensure that the programme interacts correctly with existing social security coordination rules under Regulation 883/2004.

**MS** are responsible for implementing the programme within their national social security systems, designating the competent authority that reviews Phase 2 business plans, and calibrat-

ing the additive benefit based on national insurance levels and local labour market context. They also carry responsibility for the reconsideration process at the end of Phase 3 and for the exceptional extension decisions.

**Workers** engage with the programme voluntarily. They enter Phase 1, decide whether to proceed, present their business plan to move into Phase 2, and are responsible for demonstrating minimum activity throughout the implementation window. Loss of the benefit is triggered by clear inactivity or fraud, not by business failure.

### **Pros, Cons, and Mitigation**

The qualities of this proposal are deeply tied to the report's emphasis on AI creating clear tradeoffs between past and future careers.

- **It converts a structural barrier into a structural opportunity.** The social protection cliff is not a minor inconvenience; it is the single most cited reason why workers with genuine entrepreneurial potential do not act on it. Removing it for workers already affected by AI disruption is one of the highest-leverage interventions available to the policy framework.
- **It is supply-side in a framework that is otherwise mostly demand-side.** R1 through R4 are largely about protecting workers within employment or shaping how firms behave. R5 is about creating an alternative path for those for whom the employment track has genuinely closed, making it structurally complementary rather than redundant.
- **It builds on existing mechanisms rather than creating new ones.** Because it operates as an extension and conditioning of existing unemployment insurance and social security entitlements, it does not require a new system, only a new eligibility frame within instruments that already exist.

There are two risks worth acknowledging.

- **Moral hazard.** An additive benefit tied to entrepreneurship activity could incentivize workers to enter the programme without genuine entrepreneurial intent, simply to extend their social protection coverage beyond what standard unemployment insurance would provide.
  - **Mitigation:** The Phase 2 business plan requirement, combined with the minimum activity threshold at reconsideration, creates enough friction to deter purely opportunistic entry without making the programme inaccessible to workers who are genuinely trying. The benefit is also time-limited and non-renewable in most cases, which limits its attractiveness as a pure income extension mechanism.
- **The definition of “career degradation” is contestable.** Formal displacement is straightforward to verify. Significant role degradation is not, and a poorly defined eligibility criterion here could either exclude many of the workers the programme is designed to help or open the programme to workers whose situations are not genuinely linked to AI disruption.
  - **Mitigation:** The eligibility criteria should reference the occupational AI exposure framework already used in this report’s analysis, particularly the automation versus augmentation distinction documented by Brynjolfsson et al. (2025), to provide a defensible and evidence-based basis for which roles and sectors qualify. Precise definitional choices are best left to policymakers in consultation with social partners.

## **KPIs and Data Sources**

At the individual level:

- **Share of bridge recipients who report active business operations at the end of Phase 3,** to measure genuine entrepreneurial uptake rather than programme participation.
- **Time-to-first-revenue for bridge recipients versus a control group of displaced workers who did not access the programme,** to assess whether the bridge actually shortens the financial exposure window.

At policy level:

- **Number of workers transitioning through the bridge annually**, broken down by displacement type, sector, and MS, to track, reach and identify gaps in national implementation.

These can be tracked through Eurostat’s Labour Market Transitions dataset, national social security administrative data, and the European Working Conditions Survey. As with R4, the precise measurement methodology and reporting thresholds are questions for policymakers and the relevant national competent authorities, who will be better placed to calibrate these indicators to their specific implementation context.

## **10.7 Policy Direction - An EU Wide Movement to Implement Real**

### **Change**

The five recommendations in this framework are grounded in a belief that the impact of AI on professional life is not a force of nature. It is a set of choices, made by firms, by policymakers, and by the institutions that shape the environment in which they operate. R1 through R5 represent our best attempt to translate that belief into concrete, contemporary policy instruments. But there is a deeper question that sits behind all of them, and it would be intellectually dishonest to leave it unaddressed.

The question is this: to what extent can European institutions meaningfully govern the impact of AI on European workers if the tools, the infrastructure, and the data layers through which AI operates are built, owned, and controlled outside Europe? The answer, as things currently stand, is limited. Over 80% of Europe’s digital infrastructure and core technologies are imported. The AI systems being integrated into European workplaces at the pace documented throughout this report are built on computing infrastructure, foundation models, and data governance frameworks that European regulators can respond to but cannot shape from the inside. This is not an argument for protectionism or for slowing AI adoption, but rather a call to move towards real governance capacity.

Europe is, amongst all major global powers, the one whose foundational values are most closely aligned with the policy direction this report has tried to articulate. The professional risk framework we have used throughout, income security, social security and downside protection, legal and fiscal risk, autonomy and agency embraces those values. It represents a set of commitments about what working life should look like: that workers deserve stability, that the transition costs of technological change should not weigh entirely on individuals, that human autonomy and dignity are not negotiable even in environments of rapid technological disruption. These are often not universally applied concepts. However, they very often are in European environments, and they are the reason why this report's policy proposals look the way they do.

This is why the broader technological direction matters so much for the specific policies we propose.

The human-centred vision of AI at the core of this framework, going all in on AI where it genuinely augments human capability while **insisting that the human element remains the focal point of efficient and socially sustainable business.**

But that requires more than smart regulation. It requires that Europe has a meaningful presence in the technological spaces that will shape the careers of tomorrow. That means European cloud infrastructure. It means European foundation models trained on European data under European governance. It means, as the EuroStack initiative has argued, a full-stack approach to digital sovereignty in which Europe is not just a consumer of other people's technological choices but an active builder of its own. The policies proposed in R1 through R5 will be considerably more effective if the AI tools deployed in European workplaces are ones that European institutions have shaped, that European workers can trust, that are designed to reflect the values this report has tried to push forward.

The supply-side conditions for this kind of ecosystem do not yet fully exist. One of the most persistent structural reasons is also one of the most tractable: the fragmentation of the European legal and corporate landscape makes it genuinely difficult for European founders, investors, and

companies to operate at the scale required to build the kind of technology that can compete and lead globally.

Talent that could be building European alternatives leaves for jurisdictions where the company formation, the capital markets, and the regulatory environment are simpler and more hospitable. Capital that could be funding the next generation of European AI companies follows it. This is the problem the Commission's March 2026 EU Inc proposal is directly designed to address. Although more nuanced in practice, its ideal is full of potential: a single pan-European corporate legal framework, available to any founder across all 27 MS, offering incorporation in under 48 hours, a harmonized employee ownership framework, and a centralized judicial system so that the rules mean the same thing everywhere they apply. To achieve such an ideal, however, this regulation needs to break free of European habits of compromising many of the focal points of innovative lawmaking to satisfy the tradeoffs that come with it. The point is, such regulation needs to be carefully designed not to disrupt the MS' sovereignty, but such compromise should not come at the expense of the proposal's potential, which is of removing bureaucratic barriers and fragmentation in EU wide value creation. It is the most significant structural reform proposal for European entrepreneurship and value creation at scale in a generation, and this report's recommendation R5 is explicitly designed to connect disrupted workers to exactly the kind of entrepreneurial pathway that EU Inc makes more viable.

## **11 Conclusion**

As we have argued, AI isn't just disrupting the way that work is done. It is also changing the conditions for professional security, predictability, and value. As we have shown, the relevance of AI in the workplace isn't simply captured by headline narratives of jobs lost or productivity gained alone. Its significance lies in the way that it can reshape professional risk: altering the balance between stability and opportunity, between the protections traditionally associated with employment, and the flexibility traditionally associated with entrepreneurship. In some settings,

AI can lower barriers to entry, reduce start-up costs, and give smaller firms capabilities that were once reserved for much larger organisations. Meanwhile, in others, it can compress entry-level pathways, accelerate skill obsolescence, weaken the visibility of career ladders, and introduce new forms of opacity into management, evaluation, and workplace control. The policy question which arises, then, is not simply whether AI is good or bad for work in the abstract. Instead, it is whether Europe can adopt it in ways that would work to widen opportunities without normalising insecurity, and raise productivity without hollowing out the human foundations of professional life.

From that perspective, the central conclusion which arises in our report is clear. Europe should neither approach workplace AI through passive optimism nor through defensive resistance. That means that downside risk should be reduced wherever AI amplifies volatility, opacity, and asymmetry, while preserving and expanding the real gains it can bring in efficiency, innovation, and new forms of value creation. On top of this, it means recognizing that these aren't risks that fall evenly. Large firms and SMEs aren't facing the same AI transition. The former are primarily sites of managed-transition risk, where restructuring, task compression, and algorithmic management have the possibility of outpacing workers' ability to adapt; the latter are primarily sites of capability and compliance strain, with weaker digital readiness and heavier fixed costs that can turn adoption itself into a source of divergence. As such, we propose policy responses that are differentiated in design, but coherent in purpose. They aim to protect workers from unmanaged displacement, arbitrary transitions, and declining professional legibility, while also empowering firms, entrepreneurs, and new labour-market entrants to use AI in ways that can strengthen human capability rather than substitute for it by default.

At its core, this report also argues about the kind of economy that Europe wants to build. Our comparative advantage shouldn't be seen as lying in reluctance toward technological change, but in our capacity to embrace that change and embed it with a strong social and institutional vision of what progress is for. The humanist vision of AI at the core of this framework represents the wish to embrace innovation where it genuinely augments human capability, while standing

firm on the belief that the human element remains the essential point of efficient and socially sustainable business. Europe has the power and responsibility to be the global example for the human element in the era of AI.

Such a bet, *going all-in on what AI can genuinely do while insisting that human capability, dignity, and agency remain the north star, should not be seen as a retreat from technological ambition. **It is actually the most distinctively European form of it.***

The policies discussed in this paper can be the first steps to submerging AI innovation into deep European Values. Now the only barrier to progress is fear: fear of bold actions and fear of change. But change will happen anyway, so to shape it to its image, Europe must be Brave.

## Bibliography

Accenture. (2025, June). *Europe's AI Reckoning*. <https://www.accenture.com/gb-en/insights/data-ai/europes-ai-reckoning>

Agarwal, & others. (2025). *Leveraging Artificial Intelligence to Foster Entrepreneurial Mindsets and Enhance Well-Being: A Multi-Dimensional Framework*. <https://acr-journal.com/article/leveraging-artificial-intelligence-to-foster-entrepreneurial-mindsets-and-enhance-well-being-a-multi-dimensional-framework-969/>

Babina, T., & others. (2024). Artificial Intelligence, Firm Growth, and Product Innovation. *Journal of Financial Economics*. <https://www.sciencedirect.com/science/article/pii/S0304405X2300185X>

Brynjolfsson, E., & others. (2025). *Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence*. <https://digitaleconomy.stanford.edu/publication/canaries-in-the-coal-mine-six-facts-about-the-recent-employment-effects-of-artificial-intelligence/>

Cai, J. J., Gu, X., Sheng, L., Xia, M., Zhao, L., & Zhu, W. (2025). AI as "Co-founder": GenAI for Entrepreneurship. *Arxiv Preprint Arxiv:2512.06506*. <https://arxiv.org/abs/2512.06506>

Cedefop. (2025, ). *How Europe's SMEs Are Navigating the Digital Transition*. <https://www.cedefop.europa.eu/en/news/how-europes-smes-are-navigating-digital-transition>

Centre for Future Generations. (2025, August). *Preparing for AI Labour Shocks Should Be a Resilience Priority for Europe*. <https://cfg.eu/ai-labour-shocks/>

*Council Directive 89/391/EEC on the Introduction of Measures to Encourage Improvements in the Safety and Health of Workers at Work*. (1989). <https://eur-lex.europa.eu/eli/dir/1989/391/oj/eng>

*Council Directive 98/59/EC on the Approximation of the Laws of the Member States Relating to Collective Redundancies*. (1998). <https://eur-lex.europa.eu/eli/dir/1998/59/oj/eng>

*Council Recommendation of 8 November 2019 on Access to Social Protection for Workers and the Self-Employed (2019/C 387/01)*. (2019). [https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=oj:JOC\\_2019\\_387\\_R\\_0001](https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=oj:JOC_2019_387_R_0001)

*Decision (EU) 2022/2481 Establishing the Digital Decade Policy Programme 2030*. (2022). <https://eur-lex.europa.eu/eli/dec/2022/2481/oj/eng>

*Directive (EU) 2019/1152 on Transparent and Predictable Working Conditions in the European Union*. (2019). <https://eur-lex.europa.eu/eli/dir/2019/1152/oj/eng>

*Directive (EU) 2024/2831 on Improving Working Conditions in Platform Work*. (2024, ). <https://eur-lex.europa.eu/eli/dir/2024/2831/oj/eng>

*Directive (EU) 2024/2853 on Liability for Defective Products*. (2024). <https://eur-lex.europa.eu/eli/dir/2024/2853/oj/eng>

*Directive 2002/14/EC Establishing a General Framework for Informing and Consulting Employees*. (2002). <https://eur-lex.europa.eu/eli/dir/2002/14/oj/eng>

*Directive 2003/88/EC Concerning Certain Aspects of the Organisation of Working Time*. (2003). <https://eur-lex.europa.eu/eli/dir/2003/88/oj/eng>

EDPB. (2018). *Guidelines on Automated Individual Decision-Making and Profiling for the Purposes of Regulation 2016/679 (WP251rev.01)*. [https://www.edpb.europa.eu/our-work-tools/our-documents/guidelines/automated-decision-making-and-profiling\\_en](https://www.edpb.europa.eu/our-work-tools/our-documents/guidelines/automated-decision-making-and-profiling_en)

EDPB. (2020). *Guidelines 05/2020 on Consent under Regulation 2016/679 (Version 1.1)*. [https://www.edpb.europa.eu/sites/default/files/files/file1/edpb\\_guidelines\\_202005\\_consent\\_en.pdf](https://www.edpb.europa.eu/sites/default/files/files/file1/edpb_guidelines_202005_consent_en.pdf)

Erdogan, Hiziroglu, A., & Hiziroglu. (2024). *An In-Depth Case Study of Volkswagen's AI Integration*. <https://ceur-ws.org/Vol-3698/paper7.pdf>

EU-OSHA. (2024). *European Survey of Enterprises on New and Emerging Risks (ESENER)*. <https://osha.europa.eu/en/facts-and-figures/esener>

- Eurobarometer. (2025). *Eurobarometer Report No 559*. <https://futurium.ec.europa.eu/en/border-focal-point-network/news/latest-eurobarometer-survey-reveals-key-barriers-cross-border-growth-smes>
- Eurofound. (2024). *European Working Conditions Survey (EWCS)*. <https://www.eurofound.europa.eu/en/surveys-and-data/surveys/european-working-conditions-survey>
- European Commission. (2015). *Cost of the Cumulative Effects of Compliance with EU Law for SMEs*. <https://ec.europa.eu/docsroom/documents/16321/attachments/1/translations/en/renditions/native>
- European Commission. (2016). *When Is Consent Valid?*. [https://commission.europa.eu/law/law-topic/data-protection/rules-business-and-organisations/legal-grounds-processing-data/grounds-processing/when-consent-valid\\_en](https://commission.europa.eu/law/law-topic/data-protection/rules-business-and-organisations/legal-grounds-processing-data/grounds-processing/when-consent-valid_en)
- European Commission. (2020a). *Communication from the Commission to the European Parliament, the Council, the European Economic and Social Committee and the Committee of the Regions: European Skills Agenda for Sustainable Competitiveness, Social Fairness and Resilience* (No. COM(2020) 274 final). <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52020DC0274>
- European Commission. (2020b). *Pact for Skills*. [https://employment-social-affairs.ec.europa.eu/policies-and-activities/skills-and-qualifications/working-together/pact-skills\\_en](https://employment-social-affairs.ec.europa.eu/policies-and-activities/skills-and-qualifications/working-together/pact-skills_en)
- European Commission. (2022). *Tax Compliance Costs for SMEs: An Update and a Completion Final Report*. <https://taxation-customs.ec.europa.eu/system/files/2022-12/221208%20DG%20GROW%20report%20-%202022%20Tax%20Compliance%20Costs%20SMEs.pdf>
- European Commission. (2024). *European AI Office*. <https://digital-strategy.ec.europa.eu/en/policies/ai-office>

European Commission. (2025h). *2025 Commission Work Programme and Annexes*. [https://commission.europa.eu/publications/2025-commission-work-programme-and-annexes\\_en](https://commission.europa.eu/publications/2025-commission-work-programme-and-annexes_en)

European Commission. (2025f). *Artificial Intelligence*. [https://commission.europa.eu/topics/artificial-intelligence\\_en](https://commission.europa.eu/topics/artificial-intelligence_en)

European Commission. (2025g). *Commission Seeks Feedback on Draft Implementing Act to Establish AI Regulatory Sandboxes under the AI Act*. <https://digital-strategy.ec.europa.eu/en/consultations/commission-seeks-feedback-draft-implementing-act-establish-ai-regulatory-sandboxes-under-ai-act>

European Commission. (2025a). *European Digital Innovation Hubs*. <https://digital-strategy.ec.europa.eu/en/policies/edihs>

European Commission. (2025c). *European Social Security Pass*. [https://employment-social-affairs.ec.europa.eu/policies-and-activities/moving-working-europe/eu-social-security-coordination/digitalisation-social-security-coordination/european-social-security-pass\\_en](https://employment-social-affairs.ec.europa.eu/policies-and-activities/moving-working-europe/eu-social-security-coordination/digitalisation-social-security-coordination/european-social-security-pass_en)

European Commission. (2025d). *Governance and Enforcement of the AI Act*. <https://digital-strategy.ec.europa.eu/en/policies/ai-act-governance-and-enforcement>

European Commission. (2025e). *Supporting the implementation of the AI Act with clear guidelines*. Shaping Europe's digital future. <https://digital-strategy.ec.europa.eu/en/news/supporting-implementation-ai-act-clear-guidelines>

European Commission. (2025b). *Union of Skills*. [https://employment-social-affairs.ec.europa.eu/policies-and-activities/skills-and-qualifications/union-skills\\_en](https://employment-social-affairs.ec.europa.eu/policies-and-activities/skills-and-qualifications/union-skills_en)

European Commission. (2026c). *Digital Omnibus on AI Regulation Proposal*. <https://digital-strategy.ec.europa.eu/en/library/digital-omnibus-ai-regulation-proposal>

European Commission. (2026a). *Public Consultation on European Social Security Pass (ESS-PASS) Launched*. [https://employment-social-affairs.ec.europa.eu/news/public-consultation-european-social-security-pass-esspass-launched-2026-02-03\\_en](https://employment-social-affairs.ec.europa.eu/news/public-consultation-european-social-security-pass-esspass-launched-2026-02-03_en)

- European Commission. (2026b). *Standardisation of the AI Act*. <https://digital-strategy.ec.europa.eu/en/policies/ai-act-standardisation>
- European Commission, & OECD. (2025). *AI Literacy Framework for Primary and Secondary Education*. <https://ailiteracyframework.org/>
- European Parliamentary Research Service. (2025). *Digitalisation, Artificial Intelligence and Algorithmic Management in the Workplace: Shaping the Future of Work*. [https://www.europarl.europa.eu/RegData/etudes/STUD/2025/774670/EPRS\\_STU\(2025\)774670\\_EN.pdf](https://www.europarl.europa.eu/RegData/etudes/STUD/2025/774670/EPRS_STU(2025)774670_EN.pdf)
- Eurostat. (2024). *How Digitalised Have the EU's Enterprises Become?*. <https://ec.europa.eu/eurostat/web/products-eurostat-news/w/ddn-20240829-1>
- Eurostat. (2025a). *Digitalisation in Europe – 2025 Edition*. <https://ec.europa.eu/eurostat/web/interactive-publications/digitalisation-2025>
- Eurostat. (2025c). *Unemployment Statistics and Beyond*. [https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Unemployment\\_statistics\\_and\\_beyond](https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Unemployment_statistics_and_beyond)
- Eurostat. (2025b). *Use of Artificial Intelligence in Enterprises*. <https://ec.europa.eu/eurostat/statistics-explained/SEPDF/cache/106920.pdf>
- Eurostat. (2026b). *EU statistics on income and living conditions (EU-SILC) methodology - introduction*.
- Eurostat. (2026a). *Labour Market Transitions*. <https://ec.europa.eu/eurostat/web/experimental-statistics/labour-market-transitions>
- Forbes. (2025). *Are You At Risk Of Acute Agency Decay Amid AI?*. <https://www.forbes.com/sites/corneliawalther/2025/03/29/are-you-at-risk-of-acute-agency-decay-amid-ai/>
- Gartner. (2025). *Gartner Predicts Agentic AI Will Autonomously Resolve 80 Percent of Common Customer Service Issues Without Human Intervention by 2029*. <https://www.gartner.com/en/newsroom/press-releases/2025-03-05-gartner->

predicts-agentic-ai-will-autonomously-resolve-80-percent-of-common-customer-service-issues-without-human-intervention-by-20290

GDPR. (2016). *Regulation (EU) 2016/679 (General Data Protection Regulation)*. <https://eur-lex.europa.eu/eli/reg/2016/679/oj/eng>

International AI Safety Report. (2025). *International AI Safety Report 2025*. <https://internationalaisafetyreport.org/publication/international-ai-safety-report-2025>

International Monetary Fund. (2025). *Generative Artificial Intelligence for Compliance Risk Analysis*. <https://www.imf.org/-/media/files/publications/tnm/2025/english/tnmea2025013.pdf>

JPMorgan. (2025). *How Vibe Coding Can Reshape Startup Economics*. <https://www.jpmorgan.com/insights/technology/artificial-intelligence/vibe-coding-a-guide-for-startups-and-founders>

Littler. (2025). *2025 European Employer Survey Report*. <https://www.littler.com/news-analysis/littler-report/littler-european-employer-survey-report-2025>

McKinsey & Company. (2025b). *Accelerating Europe's AI Adoption: The Role of Sovereign AI*. <https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/accelerating-europes-ai-adoption-the-role-of-sovereign-ai#/>

McKinsey & Company. (2025a). *The State of AI: How Organizations Are Rewiring to Capture Value*. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-how-organizations-are-rewiring-to-capture-value>

McKinsey Global Institute. (2024). *A New Future of Work: The Race to Deploy AI and Raise Skills in Europe and Beyond*. <https://www.mckinsey.com/mgi/our-research/a-new-future-of-work-the-race-to-deploy-ai-and-raise-skills-in-europe-and-beyond>

OECD. (2015). *The OECD Job Quality Framework*. [https://www.oecd.org/en/publications/measuring-and-assessing-job-quality\\_5jrp02kpw1mr-en.html](https://www.oecd.org/en/publications/measuring-and-assessing-job-quality_5jrp02kpw1mr-en.html)

- OECD. (2023). *Artificial Intelligence and the Labour Market: Introduction*. [https://www.oecd.org/en/publications/oecd-employment-outlook-2023\\_08785bba-en/full-report/artificial-intelligence-and-the-labour-market-introduction\\_ea35d1c5.html](https://www.oecd.org/en/publications/oecd-employment-outlook-2023_08785bba-en/full-report/artificial-intelligence-and-the-labour-market-introduction_ea35d1c5.html)
- OECD. (2025a). *AI Adoption by Small and Medium-Sized Enterprises*. [https://www.oecd.org/content/dam/oecd/en/publications/reports/2025/12/ai-adoption-by-small-and-medium-sized-enterprises\\_9c48eae6/426399c1-en.pdf](https://www.oecd.org/content/dam/oecd/en/publications/reports/2025/12/ai-adoption-by-small-and-medium-sized-enterprises_9c48eae6/426399c1-en.pdf)
- OECD. (2025b). *Generative AI and the SME Workforce*. [https://www.oecd.org/content/dam/oecd/en/publications/reports/2025/11/generative-ai-and-the-sme-workforce\\_83bafdfb/2d08b99d-en.pdf](https://www.oecd.org/content/dam/oecd/en/publications/reports/2025/11/generative-ai-and-the-sme-workforce_83bafdfb/2d08b99d-en.pdf)
- Pew Research Center. (2025). *U.S. Workers Are More Worried Than Hopeful About Future AI Use in the Workplace*. <https://www.pewresearch.org/social-trends/2025/02/25/u-s-workers-are-more-worried-than-hopeful-about-future-ai-use-in-the-workplace/>
- Regulation (EC) No 883/2004 on the Coordination of Social Security Systems*. (2004). <https://eur-lex.europa.eu/eli/reg/2004/883/oj/eng>
- Regulation (EC) No 987/2009 Laying Down the Procedure for Implementing Regulation (EC) No 883/2004 on the Coordination of Social Security Systems*. (2009). <https://eur-lex.europa.eu/eli/reg/2009/987/oj/eng>
- Regulation (EU) 2019/1149 Establishing a European Labour Authority*. (2019). <https://eur-lex.europa.eu/eli/reg/2019/1149/oj/eng>
- Regulation (EU) 2021/1057 Establishing the European Social Fund Plus (ESF+)*. (2021). <https://eur-lex.europa.eu/eli/reg/2021/1057/oj/eng>
- Regulation (EU) 2021/691 on the European Globalisation Adjustment Fund for Displaced Workers (EGF)*. (2021). <https://eur-lex.europa.eu/eli/reg/2021/691/oj/eng>
- Regulation (EU) 2024/1689 (Artificial Intelligence Act)*. (2024). <https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng>

United Nations Economic Commission for Europe. (2015). *Quality of Employment Framework*.  
[https://unece.org/statistics/quality-employment-framework#accordion\\_1](https://unece.org/statistics/quality-employment-framework#accordion_1)

World Economic Forum. (2025). *Future of Jobs Report 2025*. <https://www.weforum.org/publications/the-future-of-jobs-report-2025/in-full/>