

Department of Economics and Finance
Degree Program in Economics and Business

Course of Macroeconomics

The Impact of AI on Productivity and Growth: A Labour Market Perspective through Occupational Exposure

Prof. Francesco Lippi

SUPERVISOR

ID 279121

Giulia Bondatti

CANDIDATE

The Impact of AI on Productivity and Growth: A Labour Market Perspective through Occupational Exposure

Abstract

This thesis examines the impact of artificial intelligence (AI) on productivity and economic growth through a labour market lens, with a focus on the occupational structure of the Italian economy. Building on recent advances in the literature, it develops novel indicators of AI exposure and potential complementarity by measuring the semantic similarity between AI domain capabilities and occupational abilities, and by incorporating work context and job complexity features. These indicators are then aggregated to the sectoral level and interacted with a proxy for AI adoption—sectoral R&D intensity—to estimate their relationship with labour productivity.

Using a semi-parametric panel model over 21 Italian macro-sectors between 1996 and 2022, the analysis finds that productivity gains are not driven by exposure alone but emerge when exposure is accompanied by high complementarity and realized adoption. The results confirm that AI's effects on economic performance are conditional on the structure of work, investment in innovation, and the ability to harness AI for augmentation rather than substitution.

Table of Contents

Abstract.....	1
1. A Definition of Artificial Intelligence	4
1.1 Natural Language Processing	5
1.2 Computer Vision	5
1.3 Data Mining and Knowledge Discovery	5
1.4 Planning and Decision Systems.....	6
2. Revision of Developed Metrics for AI Exposure	6
2.1 Different Methods for AI exposure	7
2.1.1 Task-based exposure indices.....	7
2.1.2 Occupation-based Exposure Indices	10
2.1.3 Patent-based exposure indices	13
3. A Novel Exposure Measure.....	16
3.1 Introducing the need for a new indicator.....	16
3.2 Building the Exposure and Complementarity Index	18
3.2.1 Data Sources	18
3.2.2 Methodology	19
3.2.3 Exposure Score	20
3.2.4 Complementarity Effect.....	23
3.3 Results	26
3.3.1 General Review of Results.....	26
3.4 Comparison with existing indicators	29
3.4.1 Comparison with Other Indicators.....	31
3.5 Aggregation	32
3.5.1 Mapping to the Italian Classification System	33
3.5.2 Sectoral Patterns in the Italian Economy	34
3.6 Link to the Empirical Framework	34
4. The Effect on Productivity	36
4.1 Conceptual Framework and Hypotheses	37
4.1.1 From task structure to productivity.....	37
4.1.2 Hypotheses.....	38
4.2 Model Specification.....	38
4.2.1 Estimation and Diagnostics.....	39
4.3 Data and Sample Construction	40
4.3.1 Data Sources, Cleaning, Diagnostics, and Validation	40
4.4 Analytical Contribution	41
4.5 Estimation Results	41
4.6 Robustness and Limitations.....	44
4.7 Policy Implications	46
5. Conclusions	48

Appendix A – Variable Distributions and Transformations.....	49
Appendix B – Model Diagnostics	53
Appendix C – Extended Estimation Results.....	57

1. A Definition of Artificial Intelligence

Before beginning to analyse the effect of a disruptive innovation such as AI, it is of utmost importance to clearly define the various nuances of this technology in order to fully understand the source of its transformative impact. In this section, I will provide a general analysis of the different classifications and denominations of AI.

Artificial Intelligence is defined as “a technical and scientific field devoted to the engineered system that generates output such as content, forecasts, recommendations or decisions for a given level of human-defined objectives¹”. It emerged after the Second World War, with the term “Artificial Intelligence” first gaining recognition in 1956. Since then, AI has expanded significantly, permeating various aspects of human life. It can be classified into several application domains: computer vision, natural language processing (NLP), data mining and knowledge discovery, and planning and decision systems.²




Figure 1- A General Framework for AI Classification based on Functionality

The diagram in Figure 1 illustrates the hierarchical structure of AI, from foundational models (left) to practical applications (right), showing how AI models influence system types, which in turn define AI domains and real-world applications. This classification serves as a foundation for understanding how AI is integrated into various fields,

¹ ISO/IEC 22989:2022

² ISO/IEC, 2022, Clause 9

providing the necessary context for evaluating AI's impact on labour markets through exposure indicators.

1.1 Natural Language Processing

Natural Language Processing (NLP) refers to the ability of AI systems to acquire, process, and understand human language. It is widely applied across various AI fields and relies on Language Models (LMs), which predict the probability distribution of language expressions to interpret and generate text.³ NLP enables applications such as speech recognition, machine translation, sentiment analysis, and conversational AI.

1.2 Computer Vision

Computer Vision encompasses a different perception sensitivity in AI, enabling systems to recognize, interpret, and process physical objects and images. Using pattern recognition techniques and point descriptors, these systems can identify both two-dimensional and three-dimensional objects, extract meaningful features, and apply or store the gathered information. Common applications include medical imaging, facial recognition (e.g., FaceID), autonomous vehicles, and surveillance systems.

1.3 Data Mining and Knowledge Discovery

Data Mining is a subset of AI that focuses on transforming large volumes of data into meaningful knowledge by identifying patterns, correlations, and trends within extensive databases.⁴ It involves the use of statistical, machine learning, and computational techniques to uncover hidden insights that might not be immediately apparent. As described by Hand et al. (2001) "Data mining is the analysis of (often large) observational data sets to find unsuspected relationships and to summarize the data in novel ways that are both understandable and useful to the data owner."

Common applications of data mining include fraud detection, recommendation systems, market analysis, customer segmentation, and predictive analytics.

³ (Russell & Norvig, 2016)

⁴ (Han et al., 2012)

1.4 Planning and Decision Systems

As the name suggests, Planning and Decision Systems are designed to optimize processes, automate complex workflows, and enhance predictive analytics.⁵ These systems enable AI to evaluate different courses of action, anticipate outcomes, and support strategic decision-making. Just as planning in everyday life requires analysing potential consequences before taking action, AI-driven planning systems help structure approaches to complex problems, ensuring efficiency, adaptability, and informed decision-making. Common applications include logistics optimization, autonomous systems, business strategy modelling, and financial forecasting.

2. Revision of Developed Metrics for AI Exposure

The impact of AI on labour markets is a subject of increasing debate, as AI-driven technologies continue to reshape job roles, alter task structures, and influence employment dynamics. To quantify these effects, researchers have developed AI exposure metrics, which attempt to measure how different occupations and industries are affected by AI advancements.

Understanding these exposure measures is essential for evaluating how AI transforms work, whether through task automation, augmentation, or displacement. These indices offer a structured approach to assessing which jobs are most at risk and which may benefit from AI integration.

This chapter reviews the key AI exposure metrics developed in the literature, analysing their methodologies, strengths, and limitations. By critically assessing these indicators, it is possible to identify gaps and challenges that will inform the construction of a more comprehensive AI exposure index in the following chapter.

⁵ (Ghallab et al., 2004)

2.1 Different Methods for AI exposure

Different indices of AI exposure can be categorized into 3 main areas:

1. Task-based indices, which assess AI's impact by evaluating its ability to automate or complement specific tasks.
2. Occupation-based indices, which aggregate AI exposure at the job level, considering the skills and abilities required for different occupations.
3. Patent-based indices, which estimate AI's impact based on technological advancements reflected in patent filings and their relevance to occupational tasks.

Each of these approaches offers unique insights into how AI affects the labour market but also comes with methodological limitations. The following sections will examine the distinct nuances pertaining to each indicator, with particular attention to their structure, while presenting their main findings.

2.1.1 Task-based exposure indices

Task-based indices measure AI exposure by assessing how AI capabilities align with specific job tasks, typically based on occupational databases like O*NET⁶. This approach allows for a granular understanding of AI's impact but often overlooks task adaptation and workforce reallocation effects.

Frey & Osborne (2017)

Frey and Osborne assess automation risk for 702 O*NET occupations by first categorizing 70 occupations as either fully automatable (1) or not automatable (0) based on expert judgment. Consequently, they apply a Gaussian process classifier⁷ to estimate the probability of computerization for the remaining 632 occupations.

Their study is based mainly on general Machine Learning (ML) and Machine Reasoning (MR) characteristics and capabilities— as of 2017 —without a particular focus on LLMs or GenAI, which have since become major drivers of concern.

⁶ The O*NET Database comprehends information on careers and occupations, as provided by the U.S. Department of Labor Employment and Training Administration. https://www.onetcenter.org/db_releases.html

⁷ A Gaussian Process Classifier (GPC) is a probabilistic model used for classification tasks, characterized by defining a distribution over possible functions rather than assuming a fixed functional form. This allows for flexible decision boundaries and uncertainty-aware predictions.

Main Findings: AI has significant potential to substitute labour in non-routine cognitive tasks, especially in low-skill and low-wage occupations. However, jobs relying on social and creative intelligence are safer.

Brynjolfsson, Mitchell, et al. (2018)

The indicator developed by Brynjolfsson and Mitchell, called “Suitability for Machine Learning” (SML) evaluates the potential for Machine Learning to automate a set of tasks described in the O*NET database. Experts rate 23 features on a scale from 1 (strongly disagree) to 5 (strongly agree), then aggregated at the occupational level.

Unlike Frey & Osborne (2017), this study explicitly accounts for task complementarity, acknowledging that ML will transform jobs rather than completely automate them.

Main findings: Machine Learning technology has the potential to significantly transform various jobs across the economy; however, its primary impact will be in the reengineering of processes and the restructuring of tasks rather than full automation. The data supports the argument that ML is more likely to augment or reshape jobs, particularly in fields involving structured decision-making, data processing, and administrative tasks.

Tolan et al. (2021)

Tolan et al. (2021) propose a three-layer framework to assess AI exposure, consisting of tasks, cognitive abilities, and AI benchmarks. Their approach maps occupations from worker surveys and occupational databases to cognitive tasks using AI research benchmarks. Unlike previous approaches such as Brynjolfsson, Mitchell, et al. (2018), which focus on broad ML applicability, this method evaluates AI exposure based on the degree to which AI research is directed toward specific cognitive abilities required for different occupations, without assuming immediate substitution.

Building on Felten et al. (2021)’s work, their framework broadens the scope of AI benchmarks beyond perception-related tasks to include language processing, planning, information retrieval, and automated deduction/induction. They also refine the measurement of AI research intensity, addressing the issue of nonlinear performance trends that hinder comparability across benchmarks. To overcome this, they translate AI benchmarks into AI research activity, ensuring a more accurate alignment between technological advancements and occupational exposure.

Expanding on Webb (2020) methodology, they link research intensity in AI domains—such as computer vision and NLP—to the specific abilities required for performing job tasks.

Main findings: Their findings indicate that the occupations most at risk are concentrated in high-income sectors, including medical professionals, office clerks, and teachers, not due to automation but because of augmentation and transformation. Unlike previous waves of automation that primarily displaced workers, AI’s impact in these fields is more focused on reshaping and enhancing jobs rather than replacing them. The most affected abilities are those related to problem-solving and idea generation, emphasizing that AI’s role is primarily in transforming work processes rather than substituting human labour entirely.

Eloundou et al. (2023)

Eloundou et al. (2023) analyse the impact of Large Language Models (LLMs), such as Generative Pre-trained Transformers (GPTs), on occupational exposure. Their approach builds on previous work by Brynjolfsson, Mitchell, et al. (2018), Felten et al. (2021), and Webb (2020) but does not differentiate between labour-augmenting and labour-displacing effects. Instead, it measures AI exposure using the O*NET database, following Felten’s framework for automation potential and aggregating exposure estimates at both the occupational and industry levels.

The methodology is based on matching detailed work activities (DWAs) from the O*NET database with expert assessments and ChatGPT evaluations. The core metric categorizes tasks as either “Directly Exposed” or “Indirectly Exposed”, depending on whether LLMs can reduce the time required to complete a task by more than 50%. This threshold serves as a proxy for the degree of automation potential but assumes a direct correlation between task completion time and AI substitutability, which is a notable limitation.

Main Findings: Their findings suggest that higher-wage occupations are the most exposed, as tasks in these jobs tend to involve significant information processing, text generation, and decision-making, all areas where LLMs demonstrate high proficiency. Overall, the study estimates that approximately 19% of jobs in the U.S. have at least 50% of their tasks exposed to LLMs, signalling a substantial potential impact on knowledge-intensive professions. However, the approach has limitations, particularly in its reliance

on ChatGPT's own assessment of its capabilities, which introduces subjectivity and potential biases into the analysis. Moreover, the use of a 50%-time reduction threshold as the primary determinant of exposure oversimplifies the complexities of AI's impact on labour dynamics.

Briggs and Kodnani, (2023)

Briggs and Kodnani (2023) analyse AI exposure using data from the O*NET database, focusing on task content to assess the extent of labour-saving automation. Based on existing literature on AI's potential use cases, they identify 13 out of 39 work activities as exposed to automation. They then estimate occupational exposure using an importance and complexity-weighted average, which allows them to determine the share of workload AI could automate.

Main findings: Around two-thirds of current jobs in the U.S. and Europe are exposed to AI automation to varying degrees. However, up to a quarter of these jobs could face disruptive substitution effects, highlighting the uneven impact of AI across different sectors. Additionally, they estimate that Generative AI (GenAI) could increase annual U.S. labour productivity growth by just under 1.5 percentage points over a 10-year period, assuming successful adoption and realization of AI's projected capabilities. The study also suggests that AI could contribute to a 7% increase in global GDP, contingent on the pace of adoption and the accuracy of AI development projections.

2.1.2 Occupation-based exposure indices

Occupation-based indices aggregate task-level exposure into broader occupational categories, offering a more structured view of AI's labour market impact. However, these models may oversimplify job heterogeneity, failing to capture within-occupation differences.

Felten et al. (2021)

Felten, Raj, and Seamans (2021) developed the AI Occupational Exposure Index (AIOE), an indicator that maps specific AI applications to occupational abilities while remaining neutral on whether AI complements or substitutes jobs. The measure relies on expert

assessments to evaluate how AI capabilities interact with different job roles, offering a broad and flexible approach to analysing AI exposure across occupations.

Main Findings: High-skill occupations tend to have the highest AI exposure, particularly roles requiring analytical and cognitive abilities, such as financial examiners, actuaries, genetic counsellors, and mathematicians. These professions are highly susceptible to AI-driven advancements in data analysis, pattern recognition, and complex decision-making. Conversely, occupations with low AI exposure are typically non-office jobs requiring significant physical abilities, including dancers, athletes, fitness trainers, and manual labourers such as roofers, brick masons, and block masons. The study highlights that AI's impact is unevenly distributed across the labour market, with cognitive-intensive roles being more exposed to transformation while physical and manual jobs remain relatively insulated.

Bonfiglioli et al. (2024)

The indicator developed by Bonfiglioli is one of the most comprehensive in the literature, based on task exposure, job mobility, and economic transitions. It focuses on the effect of AI at a sectoral level, analysing shifts due to AI and implementation rates. The indicator combines industry-level AI adoption, tracked through job growth in AI-related occupations, with local industry employment shares to assess regional exposure across U.S. commuting zones. Unlike static task-based measures, this approach captures both displacement effects (job losses due to automation) and complementarity effects (AI augmenting labour), offering a dynamic view of AI's impact on employment and labour market adjustments.

Main findings: unlike previous technological waves, AI's effects are more prominent in services than in manufacturing. While AI adoption has a negative employment effect on low-skilled and production workers, it positively impacts high-wage earners and STEM occupations, suggesting that AI's labour market influence is highly skill-biased and sector-dependent.

Pizzinelli et al. (2023)

Pizzinelli et al. (2023) expand on the AI Occupational Exposure Index (AIOE) developed by Felten et al. (2021), introducing the Complementarity-Adjusted AI Occupational

Exposure (C-AIOE). This refined indicator incorporates AI complementarity potential at the occupational level, using O*NET data to assess how AI interacts with different job roles. Unlike traditional exposure measures that primarily focus on automation risk, the C-AIOE accounts for the extent to which AI complements, rather than replaces, human labour. A key distinction of this approach is its broader assessment of occupational exposure, considering not only task-based automation potential but also the social and physical context of work, which influences AI's role in workplace dynamics.

Main findings: The study finds that high-skilled occupations with high AI exposure also tend to have high complementarity scores, suggesting that AI will likely enhance productivity in these jobs rather than replace them. Conversely, occupations such as clerical support roles—where AI is more likely to serve as a direct substitute—are more prone to labour market disruptions. This framework underscores the heterogeneous impact of AI, where certain professions benefit from augmentation, while others face higher risks of displacement.

Arntz et al. (2017)

Arntz et al. (2017) adopt an occupational-level approach to automation risk, emphasizing within-occupation heterogeneity rather than treating entire occupations as fully automatable. Their methodology integrates automation probabilities from Frey & Osborne (2017) with job-level characteristics from the PIAAC (Programme for the International Assessment of Adult Competencies) database, allowing for a more nuanced assessment of automation exposure.

Instead of assigning automation risks to entire occupations, they apply Frey and Osborne's task-level automation probabilities to individual job tasks, then re-estimate overall automation risk while accounting for variation within occupations.

Main findings: Their adjustment reveals that only 9% of U.S. jobs are at risk of automation, a significant revision compared to the 38% projected by Frey and Osborne's task-based approach. Their findings suggest that previous models overestimated automation risk by not considering the complexity of job roles, as many occupations involve a mix of automatable and non-automatable tasks, making full displacement less likely.

2.1.3 Patent-based exposure indices

Patent-based indices estimate AI's labour market impact by analysing technological progress and its relevance to different occupations.

Webb (2020)

Webb (2020) develops a hybrid model that combines a task-based approach with AI-related patents to measure the alignment between technological advancements and job tasks. The methodology involves text-matching between job descriptions and patents, assigning task-level scores that are then aggregated into occupational scores on a percentage scale. This approach enables a quantitative assessment of AI exposure at the job level.

In addition to AI-related patents, Webb extends the analysis by examining similarities between patents linked to AI, robotics, and software and occupational task descriptions, providing insight into how these technologies affect employment growth across different occupations. A key innovation in this model is the direct patent-to-occupation mapping, which refines the task-exposure measurement by extracting verb-noun pairs from patent titles and job descriptions. This method improves the accuracy of linking AI innovations to specific job functions, offering a more detailed picture of how technological progress reshapes labour demand.

Main Findings: AI exposure is unevenly distributed across occupations, with high-skill cognitive jobs—particularly those involving data analysis, decision-making, and prediction—being the most affected. Using a patent-to-task text-matching approach, the study shows that occupations more exposed to AI-related patents tend to experience slower employment growth, indicating that AI advancements may be influencing labour demand. Additionally, Webb distinguishes between AI, robotics, and software patents, finding that AI is more closely linked to cognitive and analytical tasks, whereas robotics patents align with manual and repetitive tasks. By mapping patent text to job descriptions, the study provides a task-specific measure of AI exposure, offering a more precise understanding of how technological progress interacts with different jobs and potentially contributes to occupational shifts and employment reallocation.

Kogan et al. (2024)

Kogan et al. (2024) construct a technology exposure measure focused on capturing the displacement of existing tasks from the perspective of incumbent workers. Their approach is similar to Webb (2020) but improves on it by using natural language processing (NLP) and text embeddings instead of word hierarchies. By representing words as vectors trained on large text datasets, this method enables more precise similarity scoring between patents and occupational descriptions. Unlike prior approaches that relied on verb-noun pairs, their full-text analysis allows for more context-aware differentiation of tasks, improving the accuracy of AI exposure assessment across different occupations.

Main findings: labour-saving technologies have a negative impact on wages across all worker levels while labour-augmenting technologies produce heterogeneous effects—increasing earnings for new entrants but leading to wage declines for incumbents. Ultimately, their research indicates that technological advancements that enhance industry productivity can contribute to overall earnings growth, as aggregate labour demand rises, regardless of whether the technology is labour-saving or labor-augmenting.

Autor et al. (2022)

Autor (2022) develops an exposure measure that examines how technological progress contributes to the creation of new tasks and occupations. The methodology is patent-based, leveraging Natural Language Processing (NLP) techniques to quantify how closely patent descriptions align with occupational descriptions from the Census Alphabetical Index (CAI). Specifically, the indicator computes a similarity score based on the textual overlap between patent documentation and occupational micro-titles listed in the CAI, focusing on breakthrough innovations—patents characterized as both novel (distinct from prior innovations) and impactful (widely cited by subsequent patents).

This approach provides a comprehensive overview of the impact of “breakthrough technologies,” considering both labour-saving innovations (automated tasks) and skill obsolescence (where innovations may complement activities but require new skills that incumbent workers lack, thus making their expertise outdated). The significance of this framework lies in its influence on subsequent studies analysing the impact of AI, as it has established a theoretical foundation for measuring how new technologies transform work dynamics.

Meindl et al. (2021)

Meindl, Frank, and Mendonça (2021) develop a task-level technology exposure indicator that maps over 900 occupations to patents using Natural Language Processing (NLP), providing a real-world measure of technology diffusion rather than merely estimating theoretical automation potential. Their methodology distinguishes between traditional patent exposure and Fourth Industrial Revolution (4IR) patent exposure, reflecting how new technologies are adopted and diffused in the labour market. Instead of assessing automation potential in isolation, their approach first accounts for differences in task exposure and then aggregates these effects at the occupational level. Their method is similar to Kogan et al. (2024) in that it uses a distance matrix approach, but instead of comparing occupations to patents directly, it matches O*NET task descriptions to patent descriptions to improve the accuracy of exposure measurement.

Main Findings: Their findings indicate that manual and production occupations, such as those in construction, manufacturing, and transportation, are more exposed to traditional technologies but have low exposure to 4IR technologies. In contrast, cognitive and analytical occupations, including those in finance, marketing, and data entry, show high exposure to 4IR innovations. The study also highlights that 4IR technologies influence job growth with a lag of 10 to 20 years, reinforcing the idea that technological adoption takes time to reshape labour markets. Among the most exposed occupations are credit authorizers, statistical assistants, and computer network support specialists, while physically intensive roles like meat cutters and floor sanders remain largely unaffected. Compared to prior AI exposure measures, such as Frey & Osborne (2017) and Brynjolfsson, Mitchell, et al. (2018), which focus on theoretical automation potential, this study provides a more precise tool for labour market analysis by capturing real-world technological diffusion.

The review presented in this chapter has outlined the principal methodologies developed to measure occupational exposure to artificial intelligence. While each approach offers valuable insights, they differ in scope, assumptions, and applicability. This comparative overview has served to contextualize the diverse ways in which exposure has been

conceptualized and operationalized. Having considered these contributions, the next chapter turns to the construction of the specific indicator used in this thesis. Building on the foundations reviewed here, it introduces a bi-dimensional measure designed to capture both AI exposure and potential complementarity in a way that can be operationalized within a macroeconomic framework.

3. A Novel Exposure Measure

3.1 Introducing the need for a new indicator

As artificial intelligence (AI) technologies have advanced, so too has the effort to measure their potential effects on work. In recent years, a variety of occupational exposure indicators have been developed to quantify the alignment between AI capabilities and human tasks. These measures have helped frame the debate around job automation and the future of work. However, as discussed in Chapter 2, important methodological and conceptual limitations remain that constrain their usefulness for dynamic, productivity-oriented analysis.

Many of the most cited indicators—such as Frey & Osborne (2017) occupation-level automation probabilities—rely on static, binary classifications of task susceptibility, often extrapolated from expert judgment about the technological frontier at a fixed point in time. While this approach was foundational, subsequent critiques, such as Arntz et al. (2017), demonstrated that it significantly overstates risk by ignoring intra-occupational variation in task content. Similarly, task-based indices like the Suitability for Machine Learning (SML) measure Brynjolfsson, Mitchell, et al. (2018) and the benchmark-matching framework of Tolan et al. (2021) incorporate more granularity, but remain reliant on subjective features or limited to narrow domains like cognitive ability—excluding increasingly relevant areas such as embodied or perceptual work.

Another recurring limitation across these indicators is the merge of technological exposure with economic impact. For example, Eloundou et al. (2023) define exposure based on whether GPT-4 is estimated to reduce task completion time by 50%, but offer

no differentiation between substitution and augmentation. This distinction is critical, as emphasized in the task-based framework developed by Acemoglu & Restrepo (2019), which highlights that the effects of automation depend not only on the technical feasibility of task performance but also on the reallocation and transformation of human roles.

Some more recent approaches have begun to address this shortcoming. Pizzinelli (2023) introduce a Complementarity-Adjusted Occupational Exposure Index (C-AIOE), which combines a traditional exposure score with features intended to capture whether an occupation is more likely to be complemented or substituted by AI. While this dual-dimensional structure is conceptually valuable, their method is still limited by its use of fixed task taxonomies and lacks a mechanism for incorporating new or evolving AI capabilities over time.

These constraints are particularly problematic when the goal is not just to assess risk, but to understand how AI interacts with labour to influence productivity. A valid indicator in this context must be capable of capturing the evolving technological landscape, reflect heterogeneity within and across occupations, and distinguish between exposure and complementarity as analytically separate forces. Moreover, it must be empirically replicable, updateable, and capable of integration into sectoral frameworks.

To meet these requirements, this thesis constructs a new AI exposure and complementarity indicator. The approach departs from expert-driven and static classification methods by leveraging recent advances in semantic modelling. Following the logic of Webb (2020) and Kogan et al. (2021), the exposure component is calculated through the cosine similarity between AI domain descriptions and occupational abilities as defined in the O*NET system. These embeddings allow for unsupervised, dynamic alignment between language used in AI capabilities and the skill requirements of jobs, improving both accuracy and adaptability.

In parallel, complementarity is computed using structural features of occupations—including Job Zone classifications and selected Work Context variables—following and extending the methodology introduced by Pizzinelli (2023). Rather than collapsing

exposure and complementarity into a composite index, they are retained as separate but jointly interpretable dimensions, enabling more nuanced classification of how AI may augment, substitute, or leave unchanged different categories of work.

The indicator developed here offers three primary advantages over existing alternatives. First, it provides ability-level granularity, capturing intra-occupational variation often ignored in aggregate scores. Second, it is entirely computational and semantically based, allowing for continuous updates as AI technologies evolve. Third, its bi-dimensional structure reflects the economic insight that exposure alone does not determine outcomes—complementarity conditions are equally important, especially when productivity is the object of analysis.

Considering these methodological, theoretical, and empirical considerations, the construction of a new indicator is a necessary step toward answering the central question of this research: under what structural conditions does AI adoption translate into productivity gains?

3.2 Building the Exposure and Complementarity Index

3.2.1 Data Sources

The construction of the exposure and complementarity index relies on the following data sources:

- (i) The O*NET database, from which I extract 52 occupational abilities, their associated “importance” and “level” scores, and job-level metadata including Job Zones and Work Contexts.⁸
- (ii) A set of AI domain descriptions (e.g., natural language processing, computer vision, planning and decision-making, data mining), which are formulated in natural language based on existing literature.
- (iii) A pre-trained transformer model is employed to compute semantic similarity between AI domain descriptions and the text of O*NET ability descriptions.⁹

⁸ Available at <https://www.onetonline.org>

⁹ Specifically, the *all-mpnet-base-v2* model from HuggingFace, available at <https://huggingface.co/sentence-transformers/all-mpnet-base-v2>

3.2.2 Methodology

The construction of the final occupation impact framework follows a three-step methodology. Each step builds upon and extends the established methodologies analysed in the previous chapter.

First, I derive a general measure of exposure to AI across occupations using semantic similarity between different AI domains capabilities. Specifically, I consider four key AI Domains discussed in Chapter 1—Natural Language Processing, Computer Vision, Data Mining, Planning and Decision System—and compute a cosine similarity score between their functional description derived from existing literature and textbooks, and the 52 abilities descriptors provided in the O*NET database. These scores contribute to each occupational score ϵ by an importance and level weighted average. This step provides a general measure of the extent to which AI can perform the core cognitive, physical, psychomotor, and sensory functions required by each job.

Unlike prior studies such as Tolan et al. (2021) or Martínez-Plumed et al. (2021), which focus exclusively on cognitive abilities when measuring AI occupational exposure, my approach deliberately includes the full range of O*NET abilities—cognitive, sensory, physical, and psychomotor. This broader scope reflects the increasing capacity of AI systems, especially when integrated with robotics or advanced control technologies, to perform not only reasoning or language-based tasks but also perception-driven and embodied interactions with the physical environment. Incorporating these additional ability domains enables a more comprehensive assessment of potential AI impact across the entire occupational spectrum.

Secondly, I incorporate a measure of complementarity following the approach of Pizzinelli (2023) to differentiate between the substitution and the augmentation potential of AI. This dimension, denoted θ , is derived from two occupation-level features in the O*NET database:

- (i) Job Zones: classify occupations in 5 categories based on the level of education, training and experience needed. Categories with a higher score are

typically more complex and abstract, thus not easily substitutable by AI, but rather complemented.

- (ii) Work Contexts: describe aspects under which the work is carried out, such as communication, responsibility, physical conditions and others.

Combining these measures, the derived complementarity score distinguishes between roles likely to be enhanced by AI (high θ) and those at higher risk of automation (low θ).

Finally, I integrate the two measures in a bivariate scatter plot.

This allows for a typology of occupations along two axes: high vs. low exposure and high vs. low complementarity. Occupations are thus categorized into four quadrants representing distinct AI impact profiles to visualize the heterogeneity of AI's potential effects across the labour market.

I then compare the results of this classification with those obtained by Pizzinelli (2023), as well as with other relevant contributions in the literature.

3.2.3 Exposure Score

3.2.3.1 Preliminary Steps

As a first step in the analysis, I compute an AI exposure score for each occupation. This score reflects the extent to which AI systems can potentially perform the core abilities required by a given occupation, based on semantic similarity between AI domain capabilities and occupational abilities.

To start constructing the exposure index, I collected and organized the relevant textual descriptors for both occupational abilities and AI capabilities. I began by extracting and unifying ability descriptions from the O*NET database, focusing on the 52 abilities categorized into four main groups: cognitive, physical, psychomotor, and sensory. For each ability, I retained the official definition provided by O*NET, ensuring consistency and comparability across occupations.

In parallel, I defined textual descriptions for a selected set of AI functional domains—namely, Natural Language Processing, Computer Vision, Data Mining, and Planning and

Decision-Making Systems. These descriptions were constructed using a combination of authoritative academic sources, textbooks, and international standards. Specifically, the main references included:

- *Russell & Norvig (2016) Artificial Intelligence: A Modern Approach*
- *Hand et al, (2001), Principles of Data Mining*
- *Han et al, (2012), Data Mining: Concepts and Techniques*
- *Ghallab et al. (2004), Automated Planning: Theory and Practice*
- *Information Technology — Artificial Intelligence — Artificial Intelligence Concepts and Terminology, (2022)*
- *Mandi et al. (2024) Decision-Focused Learning: Foundations, State of the Art, Benchmark and Future Opportunities*

These sources were carefully reviewed to synthesize concise yet semantically rich descriptions for each AI domain, capturing their functional focus and core capabilities.

Once the textual descriptors were refined, I proceeded to compute semantic similarity scores using a sentence transformer model.¹⁰

The degree of similarity is measured using cosine similarity, assessing how similar the text meanings are within a high-dimensional space. This yielded a score $\epsilon_{j,n} \in [-1,1]$, representing the semantic alignment between ability j and AI domain n . The result was a matrix of exposure scores spanning all 52 abilities and four AI domains.

In the intermediate steps of the calculation, I chose to retain the disaggregation by both ability category and AI domain, to ensure interpretability and analytical flexibility before aggregating the results per occupation. This decision enables a more granular and differentiated assessment of how various AI capabilities relate to specific categories of human abilities, and how their effects may vary across different types of occupations. It

¹⁰ Specifically, I employed the pre-trained *all-mpnet-base-v2* model from HuggingFace, which belongs to a family of models known as sentence transformers. These models are designed to understand and represent the meaning of sentences or short texts in a way that allows for meaningful comparisons between them. A sentence transformer works by converting each sentence or phrase into a vector embedding—a fixed-length numerical representation that captures the semantic content of the text. <https://huggingface.co/sentence-transformers/all-mpnet-base-v2>

also facilitates domain- and ability-specific analyses, allowing for more nuanced insights into the heterogeneity of AI's potential impact across the labour market.

3.2.3.2 Computation

After obtaining the exposure scores of each ability j to each AI domain n , denoted as $\epsilon_{j,n}$, I proceeded to compute their impact at the occupational level. The first step was to determine the relative weight of each ability within each occupation k , based on O*NET's "Importance" and "Level" indicators. These two scores jointly reflect how central an ability is within an occupation and to what degree it must be mastered.

For each ability j and occupation k , I computed the raw ability weight:

$$Weight_{j,k} = I_{j,k} \cdot L_{j,k}$$

where $I_{j,k}$ is the importance score and $L_{j,k}$ is the level score of ability j in occupation k . The values were then scaled to ensure comparability across abilities within each occupation:

$$w_{j,k} = \frac{RawWeight_{j,k}}{\sum_{j=1}^{52} I_{j,k} \cdot L_{j,k}}$$

This gives a set of weights $w_{j,k}$ summing to 1 for each occupation k , representing the relative contribution of each ability to the overall skill profile of the occupation.

Next, I computed the impact of each AI domain n on each O*NET ability category c (i.e., cognitive, physical, psychomotor, sensory) by aggregating the weighted exposure scores within each category:

$$\epsilon_{k,n,c} = \sum_{j \in c} w_{j,k} \cdot \epsilon_{j,n}$$

This yields a set of disaggregated exposure scores for each occupation k , AI domain n , and ability category c . This intermediate step allows for a detailed understanding of how

each domain of AI capability is aligned with different human ability types, and how this varies across occupations.

After calculating $\varepsilon_{k,n,c}$ for all occupations, ability categories, and AI domains, I proceeded to two levels of aggregation:

1. Category-level aggregation by AI domain: For each occupation k , I computed the total exposure to each AI domain n by summing across all ability categories:

$$\varepsilon_{k,n}^{\text{total}} = \sum_c \varepsilon_{k,n,c}$$

2. Overall AI exposure: Finally, I computed the total AI exposure score for each occupation k , aggregating across all AI domains:

$$\varepsilon_k^{\text{AI}} = \sum_{n=1}^N \varepsilon_{k,n}^{\text{total}}$$

This final score $\varepsilon_k^{\text{AI}}$, subsequently normalized, reflects the overall exposure of occupation k to AI technologies, incorporating both the structure of required abilities and their semantic alignment with AI functional domains. By preserving domain-specific and ability-specific detail throughout the process, this method provides a highly granular and interpretable measure of AI exposure at the occupational level.

3.2.4 Complementarity Effect

To complement the exposure score and reflect the heterogeneity in how AI may interact with different types of work, I construct a complementarity index, building on the framework developed by Pizzinelli (2023). This index assesses whether an occupation is more likely to be augmented or substituted by AI, based on contextual and skill-based features extracted from the O*NET database.

3.2.4.1 Variables

As in Pizzinelli (2023), I rely on a set of Work Contexts and the Job Zones classification, which capture relevant non-task-specific job characteristics. Work Contexts describe the physical, interpersonal, and organizational conditions of work, while Job Zones indicate the level of education, experience, and training required for a given occupation.

Following the same framework, I group these features into six thematic components as:

1. Communication: Includes face-to-face interaction and public speaking, reflecting the importance of both spoken and written interpersonal communication.
2. Responsibility: Includes responsibility for outcomes and responsibility for others' health.
3. Physical Conditions: Includes exposure to outdoor environments and physical proximity to others.
4. Criticality: Includes consequence of errors, frequency of decision-making, and freedom to make decisions.
5. Routine: Includes degree of automation and—unlike Pizzinelli (2023)—importance of repeating the same tasks, capturing the level of task repetition and routineness (which correlates negatively with complementarity).
6. Skills: Based on Job Zones, which I re-scale by multiplying the original values by 20 to ensure consistency with the broader scale of the Work Context variables.

3.2.4.2 Computation

For each of the six thematic groups, I calculated a group-specific score for each occupation based on the selected Work Context variables (or Job Zone score, in the case of Group 6), as described above. For each group g , I computed the average of the standardized values of the relevant variables for each occupation k . This approach captures the intensity of each group's characteristics while giving equal weight to each component.

Mathematically, for a group g composed of m variables $v_1, v_2, \dots, v_m$, the group score for occupation k is calculated as:

$$\text{Group Score}_k = \frac{1}{m} \sum_{i=1}^m v_{i,k}$$

where $v_{i,k}$ is the value of variable v_i for occupation k .

As noted earlier, for Group 5 (Routine)—which conceptually works in the opposite direction (higher scores imply lower complementarity)—I apply a flipping transformation around the median to align its direction with the other groups:

$$\text{Adjusted } v_{i,5} = v_{i,5} + 2(\text{Median}_5 - v_{i,5})$$

Finally, the complementarity score for each occupation k , denoted θ_k , is computed as the sum of the six group scores:

$$\theta_k = \sum_{g=1}^6 \text{Group Score}_{k,g}$$

This unweighted sum represents the overall potential of AI to complement human labour in occupation k , based on a range of contextual and skill-based dimensions. I chose not to normalize the final scores, maintaining their raw scale to preserve interpretability and comparability with the exposure scores.

3.3 Results

3.3.1 General Review of Results




Figure 2

The bi-dimensional indicator developed in this thesis offers a structured and intuitive representation of how different occupations are positioned in relation to two core dimensions of the current wave of artificial intelligence: Exposure (ϵ), which captures the share of tasks likely to be performed by contemporary AI systems, and Potential Complementarity (θ), which measures the share of tasks rooted in human-specific skills that resist substitution. By assigning each of the 879 occupations a position in the ϵ - θ space, I aim toward a richer understanding of how the composition of work interacts with machine capabilities.

The exposure index (ϵ) ranges from highly negative values (e.g., -2.3 for Dancers) to strongly positive ones (e.g., $+1.4$ for Actuaries), reflecting a wide variation in the presence of AI-performable tasks. The complementarity index (θ) displays a narrower range of values compared to the exposure index (ϵ), with most occupations falling between 0.33

and 0.78. This more limited variation stems from how the index is built: it combines standardized measures of occupational features such as education requirements, decision-making frequency, and interpersonal interaction, all of which are bounded by design. Unlike exposure, which reflects the degree to which AI systems can technically perform a task, complementarity captures how much human input remains essential despite AI capabilities. Because nearly all jobs include some repetitive or structured elements, even highly interpersonal or judgment-based roles cannot reach the maximum possible complementarity score. As a result, the index is naturally compressed within a moderate range of values, reflecting both the structure of the underlying data and the economic reality that few occupations are entirely resistant to automation.

Graphically, the occupations do not distribute uniformly across the ε - θ plane. Rather, they cluster within a central region—approximately between $\varepsilon \in [-0.5, +0.5]$ and $\theta \in [0.45, 0.65]$ —suggesting that most occupations involve a moderate mix of both AI-exposed and human-centric tasks. However, splitting the space at the median values of each axis ($\varepsilon = 0$; $\theta = 0.55$) reveals four distinct typologies of work that reflect fundamentally different interactions with AI.

I. High Exposure – High Complementarity ($\varepsilon > 0$; $\theta > 0.55$):

This quadrant includes Lawyers, Judges, Flight Engineers, and Medical Specialists. These occupations are characterized by a substantial presence of tasks that AI can already perform—such as document review, data mining, or diagnostic support—yet they simultaneously feature a high concentration of tasks requiring interpretative judgment, ethical reasoning, interpersonal communication, and trust-building. These are core features of high-cognitive, high-stakes professions.

Such results confirm the insight of Acemoglu & Restrepo (2019), who argue that AI differs from past technological waves (e.g., robotics) by extending automation into non-routine cognitive occupations. However, rather than signalling direct displacement, high- ε , high- θ occupations represent spaces where AI is more likely to augment rather than substitute human labor. These workers might experience productivity gains, but the task structure remains resistant to full automation.

II. Low Exposure – High Complementarity ($\varepsilon < 0$; $\theta > 0.55$):

This group includes occupations such as Early Childhood Educators, Social Workers, Psychotherapists, and Creative Artists. These roles are dominated by non-routinized tasks deeply embedded in emotional intelligence, context sensitivity, creativity, and social interaction. Their low ε scores reflect the current limitations of AI in replicating such capabilities. The high θ scores reinforce the resilience of these professions in the face of automation.

These results align with the broader literature on the limits of algorithmic replication of emotional labour, confirming that complementarity is not a generic defence against AI, but is domain-specific, particularly strong in the caring, cultural, and educational professions.

III. Low Exposure – Low Complementarity ($\varepsilon < 0$; $\theta < 0.55$):

Occupations in this quadrant, such as Maids and Housekeeping Cleaners, Dishwashers, and Elementary Production Workers, are currently unexposed to AI systems—often due to the physical nature of the tasks or the low economic returns to automation—but they also lack strong human-centred features. Their low θ scores signal that the tasks are largely repetitive, low-discretion, and do not require higher-order cognitive skills.

IV. High Exposure – Low Complementarity ($\varepsilon > 0$; $\theta < 0.55$):

This is the quadrant where occupations face the greatest short-term risk of disruption. Examples include Telemarketers, Receptionists, Data Entry Clerks, and Routine Diagnostic Technicians. These jobs involve a high share of tasks that can already be handled by AI tools—such as form-filling, transcription, or scripted interactions—but offer little in terms of creativity, discretion, or person-sensitive judgment.

The high ε and low θ scores place these occupations at the core of the ongoing automation process. These are jobs where AI is not only able to perform a large number of tasks but can do so without significant loss in quality, because the complementarity with human skills is minimal. This is the segment where substitution is most advanced.

A key empirical insight from the distribution of ε and θ is that high exposure does not imply low complementarity, nor does low exposure guarantee high complementarity. The correlation between the two dimensions is modestly negative (around -0.2), indicating that many occupations contain a complex blend of replaceable and irreplaceable tasks. This suggests that future labour market adjustments will not be uniform: some jobs will be transformed, some displaced, and some reinforced, depending on their position in the ε – θ space.

In addition, cognitive intensity plays a critical role. Occupations involving high levels of information processing, problem-solving, and language comprehension tend to score higher on exposure (ε), while those that also involve high levels of discretion, decision-making under uncertainty, and social interaction score higher on complementarity (θ). This confirms that task-level analysis, rather than occupation-level typologies, is essential to capturing the true nature of AI's economic effects.

3.4 Comparison with existing indicators

The results produced by the exposure (ε) and complementarity (θ) analysis developed in this thesis exhibit broad consistency with recent contributions—most notably the Complementarity-Adjusted AI Occupational Exposure (C-AIOE) index proposed by Pizzinelli (2023)—but reveal several important empirical distinctions. In both frameworks, highly exposed occupations are concentrated in knowledge-intensive domains such as ICT, finance, law, and analytics, where core tasks include language use, information retrieval, decision-making, and planning. This shared result is expected given the alignment between these tasks and current capabilities of language-based AI systems.




Figure 3

However, when analysing my results (Figure 2), the distribution of occupations appears more dispersed and polarized than the pattern found in Pizzinelli (2023) (Figure 3). High-exposure/high-complementarity occupations such as oral surgeons, legal arbitrators, and university professors form a clear upper-right cluster, indicating roles that, while exposed, are more likely to be augmented by AI. This result contrasts with the C-AIOE, which merges exposure and complementarity into a single scalar and produces a more compressed occupational cloud, limiting the capacity to distinguish between occupations with similar exposure but differing AI interaction profiles.

The distinction is particularly salient among mid-level occupations. For example, administrative support and clerical roles display greater variation in my indicators, with some appearing moderately complementary due to contextual features like decision-making autonomy or interpersonal responsibility. In the C-AIOE index, such variation is less evident, as complementarity is fixed across five standardized dimensions and combined early in the aggregation process.

3.4.1 Comparison with Other Indicators

It is informative to contrast the results produced by the present framework with those emerging from other earlier and alternative indicators, focusing specifically on how occupations are classified in terms of exposure and potential impact.

The most widely cited benchmark, Frey & Osborne (2017), assigns high automation probabilities to a large share of middle- and even high-skill jobs, including legal clerks, diagnostic technicians, and some healthcare professionals. These classifications rely on expert assessments of task characteristics and machine learning models trained on binary “automatable” labels. However, when matched to my $\varepsilon - \theta$ index, many of these same occupations appear as high-exposure but also high-complementarity—indicating that they involve complex, judgment-intensive tasks not easily reduced to deterministic substitution. This discrepancy is especially visible for physicians and legal professionals, who appear highly exposed in both models, but are assigned much greater complementarity in my framework due to contextual factors like ethical responsibility and decision-making autonomy. Thus, while Frey & Osborne (2017) approach broadly anticipates exposure, it tends to overestimate displacement.

The Suitability for Machine Learning (SML) index by Brynjolfsson, Mitchell, et al. (2018) shifts the focus to task suitability rather than entire occupations. It classifies tasks as more or less amenable to machine learning, relying on expert ratings. In my results, occupations composed of such tasks—e.g., financial analysts or insurance underwriters—do indeed rank as highly exposed. However, my complementarity scores show greater dispersion within these roles, again highlighting heterogeneity in how AI may affect task bundles within jobs. Moreover, the SML scores are fixed in time and technological scope, while my semantic similarity model allows for dynamic reassessment as AI capabilities evolve, including emerging domains like multimodal systems and planning agents.

Tolan et al. (2021) introduce a benchmark-task alignment approach using AI system performance scores on cognitive tasks. While this method provides quantitative precision, it is limited to a narrow subset of abilities, often omitting embodied or social tasks. Accordingly, my results diverge most in low-exposure occupations: my indicator assigns

uniformly low ε scores to roles involving physical movement, dexterity, and sensory feedback (e.g., janitors, machine operators), whereas Tolan’s index lacks coverage of these dimensions. The θ index, by including interpersonal and physical work context variables, captures more clearly the non-cognitive barriers to substitution in such occupations.

Finally, Eloundou et al. (2023) present a novel approach based on the performance of GPT-4 across thousands of tasks, estimating task completion acceleration. Their measure equates time savings with exposure but does not differentiate between productivity-enhancing augmentation and displacement. My results clarify this ambiguity. Occupations flagged as exposed in Eloundou’s analysis—such as customer service agents or content creators—are confirmed in my ε scores but diverge in θ depending on whether the role involves human interaction, unpredictability, or emotional labour. This adds interpretive value and avoids conflating acceleration with substitution.

Therefore, the framework developed in this thesis produces results that are more granular, structurally flexible, and sensitive to contextual nuance. By explicitly separating exposure from complementarity, and grounding both in semantic alignment and occupational metadata, my results provide a differentiated map of risk and opportunity that better reflects the multi-dimensional nature of AI–human interaction in the workplace.

3.5 Aggregation

To analyse the potential macroeconomic impact of artificial intelligence on labour productivity, the occupation-level indicators of AI exposure and complementarity must be mapped onto a structure compatible with Italy’s economic statistics. This section details how occupational-level scores were aligned with the Italian labour market and translated into sectoral indicators, enabling empirical analysis at the macro-sector level.

Occupational characteristics determine how AI interacts with work, but productivity—the main outcome of interest—is measured at the sectoral level. As highlighted in Bonfiglioli et al. (2024), task-based indicators must be adapted to sector-level data to support empirical modelling of productivity. Aggregation is also necessary to match

other economic variables such as R&D intensity, capital investment, and value added, which are typically published by national statistical agencies using sectoral taxonomies.

Moreover, given Italy's diverse economic structure—characterized by a large services sector, strong manufacturing base, and highly fragmented employment patterns—aggregating exposure and complementarity to ATECO macro-sectors (A–R) is essential for interpreting how AI might affect the broader economy.

3.5.1 Mapping to the Italian Classification System

The first step involved mapping the original AI exposure and complementarity scores—computed for 879 SOC-2019 occupations—onto the Italian CP2021 classification system at the 4-digit level. This mapping was conducted manually and cross-referenced using ISTAT documentation to ensure the closest possible correspondence between U.S. and Italian occupations. Where multiple SOC occupations corresponded to a single CP2021 code, values were harmonized through simple averaging, due to the lack of high-resolution occupational employment data for every match.

Once occupation-level scores were assigned to CP2021 codes, a crosswalk between occupations and sectors was used to assign them to ATECO macro-sectors. This mapping was informed by publicly available employment data and documentation on occupational distribution by sector. For each sector, the corresponding scores were obtained by taking the arithmetic mean of the AI exposure and complementarity values of the associated occupations.

This method ensures that sectoral indicators remain empirically tractable and conceptually grounded. Importantly, the aggregation preserves the conceptual distinction between exposure (how AI-capable the occupational content is) and complementarity (how amenable the tasks are to human–AI collaboration)

3.5.2 Sectoral Patterns in the Italian Economy

The aggregated sector-level indicators highlight meaningful heterogeneity across the Italian economy.

High-Exposure and High-Complementarity Sectors include Information and Communication (J), Financial and Insurance Activities (K), and Professional and Scientific Activities (M). These sectors are composed of occupations that align closely with AI's functional domains, especially in areas like data analysis, planning, and language processing. Their high complementarity suggests that AI may serve an augmentative role, rather than a substitutive one, potentially reinforcing productivity dynamics when paired with adequate investment and institutional readiness.

Low-Exposure, Low-Complementarity Sectors such as Agriculture (A), Construction (F), and Accommodation and Food Services (I) remain less affected by cognitive AI technologies. These sectors typically involve manual, embodied tasks or spatial interactions that current AI systems cannot easily replicate. Consequently, both the exposure and potential for augmentation are limited, echoing international findings that embodied labour remains relatively insulated from AI-driven disruption.¹¹

Intermediate or Ambiguous Sectors, such as Wholesale and Retail Trade (G) and Transport and Storage (H), show moderate exposure but low complementarity. These sectors may face productivity challenges without organizational transformation or workforce reskilling, as AI could replace some routine tasks without substantially enhancing human roles.

3.6 Link to the Empirical Framework

These sectoral exposure and complementarity scores form the core variables in the empirical model introduced in the next chapter. By interacting them with lagged R&D intensity—a proxy for sector-level AI adoption—the model explores how AI's labour market potential translates into actual productivity performance. This structure aligns

¹¹ (Francesca Borgonovi, Flavio Calvino, Chiara Criscuolo, Julia Nania, Julia Nitschke, Layla O'Kane, Lea Samek, Helke, 2023)

with findings from Dalla Zuanna et al. (2024), who stress the importance of connecting micro-level exposure indicators with macro-level economic outcomes.

The aggregation strategy adopted in this thesis is therefore a conceptual bridge between task-based labour market insights and macroeconomic measurement—allowing AI’s sectoral impact to be studied in a way that is both empirically feasible and economically meaningful.

4. The Effect on Productivity

Chapter 3 established two task-level indicators that capture how Italian occupations relate to current artificial-intelligence capabilities. AI exposure (ε) measures the share of core abilities that can in principle be automated by existing systems, while task complementarity (θ) proxies the scope for AI to augment rather than substitute human effort.

The central objective of Chapter 4 is to determine whether this job structure helps explain the persistent heterogeneity in sectoral labour-productivity growth observed since the mid-1990s (Syverson, 2010). Despite successive waves of digitalisation—and the recent popularity gained by large-language models—Italy, like most OECD economies, has witnessed sluggish aggregate productivity (Calligaris et al., 2018). A growing literature suggests that general-purpose technologies (GPTs) such as AI raise output only when three conditions coincide:

- (i) the technology can be technically deployed in the tasks performed (high ε);
- (ii) those tasks possess characteristics—autonomy, problem-solving, interpersonal interaction—that make AI a complement rather than a substitute (high θ);
and
- (iii) firms invest in the intangible assets needed to absorb the technology (R&D, re-organisation, data infrastructure).

These arguments are formalised by (Acemoglu & Restrepo, 2019) and echoed in recent firm-level evidence (Babina et al., 2024) and sectoral analyses of absorptive capacity (Aghion et al., 2017).

To test these ideas, the chapter estimates a semi-parametric panel model in which standardised log labour-productivity is regressed on the interaction of ε and θ with lagged R&D intensity (a proxy for realised AI adoption). A penalised-spline term captures non-linear returns to capital deepening, while sector and year fixed effects control for time-invariant heterogeneity and common shocks. Full variable definitions and transformations appear in Appendix A – Variable Distributions and Transformations; the regression equations and diagnostic graphs are reported in Appendix B – Model Diagnostics.

By moving from micro-level task indicators to macro-sectoral outcomes, this chapter aims to provide new evidence on the structural pre-conditions under which AI adoption translates into productivity gains.

The empirical analysis proceeds as follows. Section 4.1 develops the conceptual framework and states the hypotheses. Section 4.2 presents the econometric model. Section 4.3 describes the data assembly. In Section 4.4, I highlight the chapter's analytical contributions. Section 4.5 reports the estimation results, and Section 4.6 examines robustness checks and discusses key limitations. Section 4.7 concludes with policy implications.

4.1 Conceptual Framework and Hypotheses

The link between artificial intelligence and aggregate performance is best understood through task-based models of technological change. In this setting, each occupation is viewed as a bundle of granular tasks, some of which are automatable by current AI, while others rely on skills that machines cannot easily replicate and may even amplify when combined with software (Acemoglu & Restrepo, 2019; Autor et al., 2018).

4.1.1 From task structure to productivity

In Acemoglu (2025), aggregate labour productivity rises when two conditions hold simultaneously: (i) the stock of effective tasks performed per worker expands, and (ii) the technology that contributes those tasks diffuses widely. Condition (i) depends on ε and θ ; condition (ii) on firms' absorptive capacity, proxied in this research by R&D intensity (Cohen & Levinthal, 1990). Without investment in complementary assets (data infrastructure, process re-design, worker training) exposure merely displaces labour without increasing output (Babina et al., 2024).

Capital deepening offers a third channel, with diminishing returns to physical capital once intangible assets are scarce (Andrews et al., 2016). Section 4.5 therefore models capital intensity with a non-parametric spline.

4.1.2 Hypotheses

The framework yields two propositions:

- (i) H1 *Conditional exposure effect*. In sectors with above-average R&D intensity, a one-standard-deviation increase in ε is associated with a positive change in subsequent labour productivity.
- (ii) H2 *Conditional complementarity effect*. Conditional on R&D, sectors scoring higher on θ enjoy higher productivity growth, even if their exposure is moderate.

Together, H1 and H2 imply that AI-induced productivity gains are neither automatic nor uniform: they materialise only where technical feasibility, human-task complementarity and absorptive investment intersect.

4.2 Model Specification

The model is estimated using both fixed-effects panel regressions and a Generalized Additive Model (GAM) to allow for nonlinear relationships. The core regression equation is:

$$\begin{aligned}\log_Productivity_{i,t} \\ &= \beta_1 Exposure_RD_{i,t-1} + \beta_2 Complementarity_RD_{i,t-1} \\ &+ s(Capital\ Intensity_{i,t}) + \alpha_i + \delta_t + \epsilon_{i,t}\end{aligned}$$

Where:

- $\log_Productivity_{i,t}$ is the log of value added per employee in sector i , year t , standardized.
- $Exposure_RD_{i,t-1}$ is the lagged interaction between AI exposure and log R&D intensity in sector i .
- $Complementarity_RD_{i,t-1}$ is the corresponding lagged interaction between AI complementarity and log R&D.
- $s(Capital\ Intensity_{i,t})$ is a smooth function (penalized spline) of log capital-to-labour ratio.
- α_i and δ_t are sector and year fixed effects, respectively.
- $\epsilon_{i,t}$ is the residual error term.

Exposure and complementarity enter the model as standardised indices to facilitate the economic interpretation of coefficients. They are interacted with the one-year-lagged logarithm of R&D intensity to reflect the diffusion lag documented by Guarascio et al. (2023).

The lag structure reflects the notion that the benefits of AI adoption, like those of any general-purpose technology, are not instantaneous but require time to diffuse through firms and institutions.

By lagging the interaction terms, this model improves temporal ordering and addresses potential concerns over simultaneity bias, although endogeneity cannot be fully ruled out due to the observational nature of the data.

The use of R&D intensity as a proxy for AI adoption follows the precedent of Calvino and Fontanelli (2023), which find that firms in R&D-intensive sectors are significantly more likely to implement AI technologies and restructure tasks accordingly. R&D thus serves as a credible and measurable proxy for latent adoption dynamics.

To control for unobserved heterogeneity, the model includes:

- Sector fixed effects (α_i), capturing time-invariant characteristics such as regulatory frameworks, institutional setups, or typical capital–labor structures;
- Year fixed effects (δ_t), which absorb macroeconomic fluctuations, technological shocks, and common policy shifts across the economy.

4.2.1 Estimation and Diagnostics

The model was estimated in multiple stages:

- (i) Fixed Effects Model (FE): the baseline included standardized interaction terms and a linear capital control. Initial results showed good fit, but residual plots and LOESS diagnostics revealed nonlinearities in capital intensity, prompting a transition to a semi-parametric model.
- (ii) Quadratic Model: Including squared terms improved fit marginally but introduced multicollinearity and reduced interpretability.
- (iii) Generalized Additive Model (GAM): The final model included a penalized spline for log capital intensity and lagged interaction terms. Residuals passed

normality tests (Shapiro–Wilk), and no major heteroskedasticity or misspecification was detected.

The next section (§ 4.3) details the construction of the dataset, variable transformations and descriptive patterns.

4.3 Data and Sample Construction

To empirically examine the relationship between AI-related occupational structures and productivity, a balanced panel dataset was constructed covering 21 Italian macro-sectors (ATECO A–R) over the period 1996–2022. The dataset integrates occupational exposure and complementarity indicators with sector-level data on productivity, R&D investment, and capital accumulation.

4.3.1 Data Sources, Cleaning, Diagnostics, and Validation

The occupational indicators for AI exposure and complementarity were computed as detailed in Aggregation (§ 3.5), ensuring that the final indicators represent the average occupational structure within each sector.

Sector-level economic variables—labour productivity, capital stock, employment, and R&D expenditure—were sourced from ISTAT¹² and the OECD STAN database¹³.

All datasets were reshaped, merged on consistent sector-year identifiers, and harmonized to ensure alignment across classification systems. Observations with missing values in any core variable were excluded. R&D intensity reporting was incomplete in early years, leading to the removal of a small number of pre-2000 observations.

¹²

https://esploradati.istat.it/databrowser/#/it/dw/categories/IT1,DATAWAREHOUSE,1.0/UP_ACC_MISPRO/IT1,98_197_DF_DCCN_PRODUTTIVITA_4,1.0

¹³ [https://data-explorer.oecd.org/vis?fs\[0\]=Topic%2C1%7CScience%252C%20technology%20and%20innovation%23INT%23%7CResearch%20and%20development%20%28R%26D%29%23INT_RD%23&fs\[1\]=Reference%20area%2C0%7CItaly%23ITA%23&pg=0&fc=Reference%20area&snb=19&vw=tb&df\[ds\]=dsDisseminateFinalDMZ&df\[id\]=DSD_ANBERD%40DF_ANBERDi4&df\[ag\]=OECD.STI.STP&df\[vs\]=1.0&dq=ITA.A...XDC.V.&pd=%2C&to\[TIME_PERIOD\]=false](https://data-explorer.oecd.org/vis?fs[0]=Topic%2C1%7CScience%252C%20technology%20and%20innovation%23INT%23%7CResearch%20and%20development%20%28R%26D%29%23INT_RD%23&fs[1]=Reference%20area%2C0%7CItaly%23ITA%23&pg=0&fc=Reference%20area&snb=19&vw=tb&df[ds]=dsDisseminateFinalDMZ&df[id]=DSD_ANBERD%40DF_ANBERDi4&df[ag]=OECD.STI.STP&df[vs]=1.0&dq=ITA.A...XDC.V.&pd=%2C&to[TIME_PERIOD]=false)

Variable transformations were motivated by visual diagnostics: **Figure A. 5** and **Figure A. 6** show a correction of skewness after log-transformation; **Figure A. 7** reveals normalization after log-scaling R&D.

Standardization was necessary for Exposure $\times$ R&D and Complementarity $\times$ R&D due to dispersion observed in **Figure A. 8** and **Figure A. 9**.

4.4 Analytical Contribution

This model specification makes four novel contributions to the AI–productivity debate:

- (i) It incorporates a dual indicator structure for AI exposure and complementarity, building on task-based and semantic AI mapping frameworks.
- (ii) It tests whether productivity effects arise conditionally on adoption, using R&D as a credible proxy.
- (iii) It models capital intensity flexibly, capturing threshold and diminishing returns that are missed in linear models.
- (iv) It combines these elements within a semi-parametric fixed-effects framework, integrating both structural theory and empirical diagnostics.

By capturing the structural, temporal, and nonlinear dynamics of AI adoption, the model enhances the understanding of when and where AI delivers productivity gains—and, crucially, where it does not. This has direct implications for policies targeting sectoral R&D incentives, workforce reskilling, and AI integration strategies.

4.5 Estimation Results

The estimation results presented in this section aim to assess whether sectors characterized by higher occupational exposure to AI—and to AI complementarity in particular—tend to exhibit systematically higher levels of productivity. The results of the Generalized Additive Model (GAM) incorporating sector and year fixed effects show that both AI exposure and complementarity, when interacted with R&D intensity as a proxy for actual AI adoption, are positively associated with labour productivity.

Variable	Coefficient	Std. Error	t-value	Significance
Lagged Exposure $\times$ R&D Intensity	3.15	0.66	4.77	***
Lagged Complementarity $\times$ R&D Intensity	0.43	0.16	3.98	***
s(log(Capital Intensity))	EDF $\approx$ 8.09 F $\approx$ 28.7			***
Sector Fixed Effects	Included			
Year Fixed Effects	Included			
Adjusted R ²	0.818			
Deviance Explained	84.6%			

Significance codes: *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

Table C.4 – GAM with Lagged Regressors and Fixed Effects

Both lagged interaction terms—lagged Exposure_RD_z and lagged Complementarity_RD_z—are positive and highly significant. The coefficient for lagged AI exposure ($\beta \approx 3.15, p < 0.001$) indicates that sectors whose workforce is more functionally aligned with AI capabilities, and where AI adoption is higher, are systematically associated with higher levels of labour productivity in the following year. Similarly, the coefficient for lagged complementarity ($\beta \approx 0.43, p < 0.001$) suggests that productivity benefits are stronger when AI exposure is coupled with high potential for augmentation, as captured by occupational context and job complexity.

The smooth term for capital intensity is also highly significant ($F \approx 28.7, p < 0.001$), confirming the presence of nonlinear effects. The estimated spline reveals a concave relationship: productivity gains increase steeply with early capital accumulation but plateau beyond a certain threshold. This pattern is consistent with prior findings on capital-skill complementarity and diminishing returns in capital-intensive industries (Bartel et al., 2005), where capital deployment must be matched by organizational and skill-based adaptation to yield sustainable gains.

The adjusted R² of the model exceeds 0.8, and the deviance explained is above 84%, indicating high explanatory power. These values suggest that occupational AI exposure and complementarity—when weighted by sectoral AI investment—explain a substantial

share of the variation in productivity across sectors, even after accounting for unobserved heterogeneity. The robustness of these results is reinforced by the diagnostic figures in **Appendix B – Model Diagnostics**, which confirm model adequacy in terms of residual behaviour, variable interactions, and multicollinearity

Notably, the results highlight the asymmetric nature of AI’s productivity effects. Sectors with low exposure and low complementarity, such as traditional manual services and low-tech manufacturing, appear less likely to benefit from AI in the near term. This echoes conclusions from the report from EY, ManPower Group, Sanoma Italia, *Il Futuro delle competenze nell’era dell’Intelligenza Artificiale* (2023), which emphasizes that many low-qualification or procedural office roles are increasingly vulnerable to stagnating demand and automation exposure.

Conversely, sectors such as information and communication technologies, finance, and professional services exhibit both high AI exposure and high complementarity, suggesting that they are ideally positioned to absorb AI productively. These findings confirm earlier macro-level insights from Acemoglu (2025), who emphasizes that productivity-enhancing gains from AI will depend heavily on sectoral context and the nature of human–machine interaction, with long-run TFP impacts only partially manifesting within a decade .

From a substantive perspective, these findings support the hypothesis that AI does not exert uniform effects across the labour market. Instead, productivity gains are most pronounced in sectors where:

- (i) AI capabilities are aligned with task profiles (high exposure),
- (ii) Work conditions enable human–AI collaboration (high complementarity),
- (iii) Sufficient capital is available to integrate and scale new technologies.

The joint inclusion of exposure and complementarity also reveals an important distinction: exposure alone, while necessary, may be insufficient to yield productivity

benefits unless occupations are structured in ways that allow for effective augmentation (Acemoglu & Restrepo, 2019).

The results highlight not only where AI is likely to have the strongest economic impact, but under what conditions—offering important implications for sector-specific investment, workforce training, and policy prioritization. Although derived from a sector-level labour productivity model, these findings have broader relevance beyond their immediate empirical scope. In particular, they raise important considerations for the macroeconomic interpretation of AI-driven growth patterns.

A possible extension of this work concerns the connection between the empirical findings and the broader macroeconomic debate on the aggregate productivity impact of artificial intelligence. While this thesis focuses on sector-level productivity and does not directly model total factor productivity (TFP) growth or long-run dynamics, the evidence presented—namely when conditions (i) to (iii) coincide—is highly relevant to recent theoretical contributions. In particular, (Filippucci et al., 2025) show that, even in the presence of strong task-level productivity gains, the aggregate effect on TFP can be severely limited by structural constraints such as low factor mobility or inelastic sectoral demand (Baumol effect)¹⁴. Although these general equilibrium channels are beyond the scope of the present model, the results obtained here provide microeconomic conditions that can inform such macro calibrations: productivity gains from AI are not automatic, but conditional on the alignment of technological capabilities with the sectoral capacity to absorb and complement them. This suggests that the labour market structure—specifically, the distribution of tasks and skills—plays a critical mediating role in translating AI potential into realized productivity growth.

4.6 Robustness and Limitations

A series of sensitivity checks confirms that the headline finding—productivity rises only where AI-exposed tasks coincide with high complementarity and substantive R&D effort—does not rest on a particular functional form.

¹⁴ (Baqee & Farhi, 2019)

I first re-estimated the benchmark with a conventional two-ways fixed-effects regression in which capital intensity enters linearly (**Table C. 1**). The coefficients on the lagged $\text{Exposure} \times \text{R\&D}$ and $\text{Complementarity} \times \text{R\&D}$ terms remain virtually unchanged in magnitude and significance, indicating that the semi-parametric treatment of capital is not driving the interaction effects. Adding quadratic terms for each regressor (**Table C. 2**) likewise leaves the core estimates intact, although the extra curvature slightly inflates standard errors and raises variance-inflation factors toward—but never beyond—two.

A preferred specification (**Table C. 3**) replaces the linear capital term with a penalized thin-plate spline, capturing the well-documented concave relationship between capital deepening and productivity once intangible complements become scarce. This adjustment lifts the adjusted R^2 to roughly 0.81 and increases the share of deviance explained to more than 84 %, yet the two interaction coefficients continue to be positive and highly significant. Alternative lag structures, reported in **Table C. 4**, show that a one-year delay between R&D effort and productivity response delivers the best fit—consistent with the diffusion lag observed in Italian manufacturing by Guarascio et al. (2023).

Diagnostic evidence collected in **Appendix B – Model Diagnostics** supports the statistical soundness of these specifications. Variance-inflation factors for all covariates—including the spline basis—stay below 2 (**Table B. 1**), ruling out multicollinearity. Residual-versus-fitted plots and Q–Q diagrams (**Figure B. 1** and **Figure B. 2**) reveal no material departures from homoskedasticity or normality, while the Shapiro–Wilk statistics confirm well-behaved errors. Finally, the smooth-term graph in **Figure B. 9** displays the expected diminishing-returns pattern in capital intensity without signs of over-fitting, validating the choice of spline complexity.

These robustness checks notwithstanding, two caveats deserve emphasis. First, the fixed-effects framework cannot fully purge contemporaneous, sector-specific shocks—such as sudden regulatory changes or regionally targeted digital-infrastructure programmes—that might correlate with both R&D spending and productivity. Future work could incorporate explicit policy dummies or difference-in-difference designs to sharpen identification. Second, using R&D spending as a proxy for AI adoption remains an indirect measure of absorptive capacity. Sector-wide surveys that record the

share of firms actively deploying machine-learning systems would offer a more granular view of diffusion intensity.

Recognising these limitations clarifies the contribution of the present analysis: it shows that the productivity dividend from AI is conditional, emerging only where technical feasibility (high ε), human-task complementarity (high θ) and sustained intangible investment intersect.

4.7 Policy Implications

The empirical exercise shows that labour-productivity gains materialize only where three ingredients overlap: a large bundle of AI-performable tasks (high ε), work organisation that lets human skills complement those tasks (high θ), and sustained intangible investment (high, lagged R&D intensity). Because those conditions are satisfied in only a handful of Italian sectors—chiefly ICT, finance and professional services—the aggregate productivity effect of AI is small and uneven. Two broad policy priorities follow.

First, government should concentrate on raising absorptive capacity in the “middle” of the economy, not simply subsidizing frontier adopters. The evidence that ε and θ yield economic payoffs only when paired with R&D implies that tax credits (or direct grants) for intangible capital—data pipelines, process redesign, training time—can unlock dormant productivity potential in otherwise exposed sectors such as machinery, transport equipment and high-end retail. The results presented justify extending incentives to organisational AI projects that embed large-language models or predictive maintenance tools in existing task flows. Crucially, evaluation criteria should privilege projects that redeploy workers into judgement or interaction-intensive tasks rather than merely automating headcount.

Second, task-level complementarity must be cultivated directly inside firms and local labour markets. For sectors sitting on high exposure but low complementarity—personal services, hospitality, traditional back-office activities—the risk is displacement without productivity gain. Active labour-market policies therefore need to move toward teaching

specific hybrid competences that raise θ : prompt engineering, human-AI supervisory skills and domain-specific data stewardship.

The results also caution against relying on AI to revive productivity unless complementary reforms accompany technological rollout. Without deeper product-market competition and faster reallocation of capital toward high- θ establishments, the sectors that already sit on the steep part of the $\varepsilon - \theta - R\&D$ surface will keep pulling ahead, widening the dual economy. Conversely, a coordinated package of intangible-investment incentives, targeted up-skilling, and competition-enhancing regulation can move more sectors onto a shared productive frontier, translating AI's technical promise into broad-based growth.

5. Conclusions

This thesis has investigated the impact of artificial intelligence on productivity and economic growth from a labour market perspective, focusing on the occupational structure of the Italian economy. By constructing novel indicators of AI exposure and complementarity and matching them to detailed Italian sectoral data, the analysis aimed to shed light on the channels through which AI interacts with the labour force and contributes to productivity dynamics.

The findings suggest that AI exposure alone is not a sufficient condition for productivity gains. Rather, sectors characterized by high complementarity—where AI systems are more likely to augment rather than replace human labour—tend to experience stronger productivity performance, particularly when matched with sustained R&D investment. These results support the growing body of literature emphasizing that the macroeconomic effects of AI depend not only on technical feasibility, but also on the structure of tasks, the adoption environment, and the potential for human–AI collaboration.

By focusing on Italy, the thesis adds empirical evidence to a context marked by digital transformation challenges, persistent productivity stagnation, and heterogeneous sectoral performance. The country’s structural features make it an important case for understanding the conditions under which AI adoption can yield macroeconomic benefits.

While the model does not intend to establish causal relationships, it provides a robust empirical framework for identifying patterns that are consistent with the theoretical literature. The exposure and complementarity indicators developed here may serve as useful tools for future research aiming to study AI’s impact on labour and growth across different national contexts.

Future work could extend this analysis by incorporating firm-level data, refining measures of AI adoption, or modelling general equilibrium effects. Nonetheless, this thesis contributes to the ongoing effort to quantify and understand the economic consequences of AI, offering a labour market lens on one of the most transformative technologies of our time.

Appendix A – Variable Distributions and Transformations

This appendix presents the distributional properties and transformation choices for the main variables included in the empirical analysis.




Figure A. 1—Histogram of Sectoral Capital Values




Figure A. 2 — Histogram of Sectoral Labor Input




Figure A. 3 — Histogram of Capital Intensity (Capital per Labor)




Figure A. 4 — Histogram of Log-Transformed Capital Intensity (z-scored)




Figure A. 5 – Histogram of Raw Productivity




Figure A. 6 – Histogram of Log-Transformed Productivity (z-scored)




Figure A. 7 – Histogram of Log R&D Intensity




Figure A. 8 – Histogram of Exposure_RD_z




Figure A. 9 – Histogram of Complementarity_RD_z

These distributions show that transformations (especially log and standardization) were necessary to address skewness and heteroscedasticity, improving comparability across sectors and time.

Appendix B – Model Diagnostics

This appendix provides graphical and statistical diagnostics to assess the robustness and validity of the estimated models.




Figure B. 1 – Residuals vs Fitted Values (GAM)




Figure B. 2 – Q-Q Plot of Residuals (GAM)




Figure B. 3– Histogram of Residuals




Figure B. 4 – Residuals vs Linear Predictor (FE Model)




Figure B. 5 – Deviance Residuals vs Quantiles (GAM)




Figure B. 6 – LOESS Fit: $\log_productivity_z \sim Exposure_RD_z$




Figure B. 7 – LOESS Fit: $\log_productivity_z \sim Complementarity_RD_z$




Figure B. 8 – LOESS Fit: $\log_productivity_z \sim \log_capital_intensity_z$




Figure B. 9 – Smooth Terms in GAM: $s(\text{Exposure_RD_z})$ and $s(\text{log_capital_intensity_z})$

Variable	VIF
Exposure $\times$ R&D Intensity	1.12
Complementarity $\times$ R&D Intensity	1.09
log(Capital Intensity)	1.18

Table B. 1 – Variance Inflation Factors (VIF) for Main Regressors

Note: Multicollinearity tests show all VIF values < 1.2 , confirming low correlation among predictors.

Appendix C – Extended Estimation Results

This appendix reports the full output of the alternative model specifications discussed in Section 4.6.

Variable	Coefficient	Std. Error	Significance
Exposure $\times$ R&D (z)	4.026	0.762	(***)
Complementarity $\times$ R&D (z)	0.574	0.119	(***)
log(Capital Intensity) (z)	0.643	0.058	(***)
Sector Fixed Effects	Included		
Year Fixed Effects	Included		
Adjusted R ²	0.249		

Note: Standard linear fixed effects model. All continuous variables z-standardized. Significance codes: *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

Table C. 1 – Fixed Effects Model Estimates

Variable	Coefficient	Std. Error	t-value	Significance
Exposure $\times$ R&D Intensity	3.592	0.792	4.54	***
(Exposure $\times$ R&D) ²	−0.707	0.400	−1.76	.
Complementarity $\times$ R&D Intensity	0.647	0.130	4.97	***
(Complementarity $\times$ R&D) ²	−0.02	0.046	−0.43	
log(Capital Intensity)	0.437	0.101	4.34	***
(log(Capital Intensity)) ²	−0.067	0.027	−2.46	*
Sector Fixed Effects	Included			
Year Fixed Effects	Included			
Adjusted R ²	0.270			

Note: Inclusion of quadratic terms leads to modest improvement in fit but introduces multicollinearity, as reflected in VIF values approaching 2. Significance codes: *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

Table C. 2 – Quadratic Fixed Effects Model Estimates

Variable	Coefficient	Std. Error	t-value	Significance
Exposure $\times$ R&D Intensity	3.176	0.68	4.44	***
Complementarity $\times$ R&D Intensity	0.435	0.17	4.02	***
s(log(Capital Intensity))	EDF $\approx$ 7.839 F $\approx$ 26.5			***
Sector Fixed Effects	Included			
Year Fixed Effects	Included			
Adjusted R ²	0.805			
Deviance Explained	83.2%			

Note: Penalized spline applied to log capital intensity. GAM estimated using REML. All predictors standardized.
Significance codes: *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

Table C. 3 – Generalized Additive Model (GAM) Estimates

Variable	Coefficient	Std. Error	t-value	Significance
Lagged Exposure $\times$ R&D Intensity	3.15	0.66	4.77	***
Lagged Complementarity $\times$ R&D Intensity	0.43	0.16	3.98	***
s(log(Capital Intensity))	EDF $\approx$ 8.09 F $\approx$ 28.7			***
Sector Fixed Effects	Included			
Year Fixed Effects	Included			
Adjusted R ²	0.818			
Deviance Explained	84.6%			

Note: This is the final model specification used in the thesis. Lag structure reflects delayed impact of AI adoption. Splines allow nonlinear capital effects. Model diagnostics reported in Appendix B.
Significance codes: *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

Table C. 4 – GAM with Lagged Regressors and Fixed Effects

All models confirm the significance of the interaction terms between AI exposure/complementarity and R\&D intensity. The spline term for capital intensity is statistically significant and displays a concave shape, consistent with non-linear capital returns.

Together, these tables and diagnostics provide evidence of model stability and robustness across specifications.

Bibliography

- Acemoglu, D. (2021). *Harms of AI*.
- Acemoglu, D. (2024). *The Simple Macroeconomics of AI*.
- Acemoglu, D. (2025). The simple macroeconomics of AI. *Economic Policy*, 40(121), 13–58. <https://doi.org/10.1093/epolic/eiae042>
- Acemoglu, D., Anderson, G., Beede, D., Buffington, C., Childress, E., Dinlersoz, E., Foster, L., Goldschlag, N., Haltiwanger, J., Kroff, Z., Restrepo, P., & Zolas, N. (2024). *Automation and the Workforce: A Firm-Level View from the 2019 Annual Business Survey*.
- Acemoglu, D., & Restrepo, P. (2016). *The Race between Machine and Man: Implications of Technology for Growth, Factor Shares and Employment*. <http://www.nber.org/papers/w22252>
- Acemoglu, D., & Restrepo, P. (2018). *Artificial Intelligence, Automation and Work*.
- Acemoglu, D., & Restrepo, P. (2019). Automation and New Tasks: How Technology Displaces and Reinstates Labor. *Journal of Economic Perspectives*, 33(2), 3–30. <https://doi.org/10.1257/jep.33.2.3>
- Acemoglu, D., & Restrepo, P. (2021). *Tasks, Automation, and the Rise in US Wage Inequality*.
- Aghion, P., Jones, B. F., & Jones, C. I. (n.d.). *Artificial Intelligence and Economic Growth*.
- Aghion, P., Jones, B. F., & Jones, C. I. (2017). *Artificial Intelligence and Economic Growth*. 23928.
- Albanesi, S., Dias Da Silva, A., Jimeno, J. F., Lamo, A., & Wabitsch, A. (2025). New technologies and jobs in Europe. *Economic Policy*, 40(121), 71–139. <https://doi.org/10.1093/epolic/eiae058>
- Anelli, M. (2024). *University to work transition*.
- Anna Milanez. (2023). *The Impact of AI on the Workplace: Evidence from OECD Case Studies of AI Implementation* (OECD Social, Employment and Migration Working Papers No. 289; OECD Social, Employment and Migration Working Papers, Vol. 289). <https://doi.org/10.1787/2247ce58-en>

- Anthony Mullen, Erick Brethenoux, Pieter den Hamer, Bettina Tratz-Ryan. (2024, February 12). *Generative AI Will Impact the Design and Control of Sociotechnical Systems*. Gartner.
- Arntz, M., Gregory, T., & Zierahn, U. (2017). *Revisiting the risk of automation*.
- Autor, D., Autor, D., Salomons, A., & Salomons, A. (2018). Is Automation Labor Share-Displacing? Productivity Growth, Employment, and the Labor Share. *Brookings Papers on Economic Activity*, 2018(1), 1–87. <https://doi.org/10.1353/eca.2018.0000>
- Autor, D., Chin, C., & Salomons, A. (2022). *New Frontiers: The Origins and Content of New Work, 1940–2018*.
- Babina, T., Fedyk, A., He, A., & Hodson, J. (2024). Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics*, 151, 103745. <https://doi.org/10.1016/j.jfineco.2023.103745>
- Baqae, D. R., & Farhi, E. (2019). *The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten's Theorem*.
- Bartel, A., Ichniowski, C., & Shaw, K. (2005). *How Does Information Technology Really Affect Productivity? Plant-Level Comparisons of Product Innovation, Process Improvement and Worker Skills* (No. w11773; p. w11773). National Bureau of Economic Research. <https://doi.org/10.3386/w11773>
- Bonfiglioli, A., Crinú, R., Gancia, G., & Papadakis, I. (2024). *Artificial Intelligence and Jobs: Evidence from US Commuting Zones*.
- Briggs, J., Kodnani, D., Pierdomenico, G., & Hatzius, J. (2023). *Global Economics Analyst The Potentially Large Effects of Artificial Intelligence on Economic Growth*.
- Brynjolfsson, E., Li, D., & Raymond, L. R. (2023). *Generative AI at Work*. 31161. <http://www.nber.org/papers/w31161>
- Brynjolfsson, E., Mitchell, T., & Rock, D. (2018). What Can Machines Learn and What Does It Mean for Occupations and the Economy? *AEA Papers and Proceedings*, 108, 43–47. <https://doi.org/10.1257/pandp.20181019>
- Brynjolfsson, E., Rock, D., & Syverson, C. (2018). *The Productivity J-Curve: How Intangibles Complement General Purpose Technologies*. https://web.stanford.edu/~mww/webb_jmp.pdf

- Brynjolfsson, E., & Unger, G. (2023). *The Macroeconomics of Artificial Intelligence. Artificial Intelligence.*
- Calligaris, S., Del Gatto, M., Hassan, F., Ottaviano, G. I. P., & Schivardi, F. (2018). The productivity puzzle and misallocation: An Italian perspective. *Economic Policy*, 33(96), 635–684. <https://doi.org/10.1093/epolic/eiy014>
- Calvano, E., & Calzolari, G. (2025). AI and policy: What makes AI different? *Economic Policy*, 40(121), 1–8. <https://doi.org/10.1093/epolic/eiae067>
- Catenacci, A. M., & Velpi, G. (2024). *IMPATTO ECONOMICO DEI NUOVI MODELLI DI INTELLIGENZA ARTIFICIALE: SFIDE E OPPORTUNITÀ PER L'UNIONE EUROPEA.*
- Cœur, B. (2025). *Benoit Cœur- discussion of: The simple macroeconomics of AI.*
- Cohen, W. M., & Levinthal, D. A. (1990). Absorptive Capacity: A New Perspective on Learning and Innovation. *Administrative Science Quarterly*, 35(1), 128. <https://doi.org/10.2307/2393553>
- Dalla Zuanna, A., Dottori, D., Gentili, E., & Lattanzio, S. (2024). *An assessment of occupational exposure to artificial intelligence in Italy.* Banca d'Italia. <https://doi.org/10.32057/0.QEF.2024.878>
- Dan Andrews, Chiara Criscuolo, & Peter N. Gal. (2016). *The Best versus the Rest: The Global Productivity Slowdown, Divergence across Firms and the Role of Public Policy* (OECD Productivity Working Papers No. 5; OECD Productivity Working Papers, Vol. 5). <https://doi.org/10.1787/63629cc9-en>
- Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2023a). *GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models* (No. arXiv:2303.10130). arXiv. <https://doi.org/10.48550/arXiv.2303.10130>
- Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2023b). *GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models* (No. arXiv:2303.10130). arXiv. <https://doi.org/10.48550/arXiv.2303.10130>
- EY, ManPower Group, Sanoma Italia. (2023). *Il Futuro delle competenze nell'era dell'Intelligenza Artificiale.*
- Ezra Greenberg, Asutosh Padhi, and Sven Smit. (2024). *2024 and Beyond: Will it be Economic Stagnation or the Advent of Productivity-driven Abundance?*

- Felten, E., Raj, M., & Seamans, R. (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, 42(12), 2195–2217. <https://doi.org/10.1002/smj.3286>
- Ferriere, A. (2025). Axelle Ferriere discussion of: AI & US jobs. *Economic Policy*, 40(121), 195–197. <https://doi.org/10.1093/epolic/eiae064>
- Filippucci, F., Gal, P., & Schief, M. (2025). *Aggregate Productivity Gains from Artificial Intelligence: A Sectoral Perspective*.
- Flavio Calvino, Luca Fontanelli. (2023). *A portrait of AI adopters across countries: Firm characteristics, assets' complementarities and productivity* (OECD Science, Technology and Industry Working Papers No. 2023/02; OECD Science, Technology and Industry Working Papers, Vol. 2023/02). <https://doi.org/10.1787/0fb79bb9-en>
- Francesca Borgonovi, Flavio Calvino, Chiara Criscuolo, Julia Nania, Julia Nitschke, Layla O’Kane, Lea Samek, Helke, S. (2023). *Emerging Trends in AI skill demand across 14 OECD Countries*. 2.
- Francesco Filippucci, Giuseppe Nicoletti, Cecilia Jona-Lasinio, Alvaro Leandro, & Peter Gal. (2024). *The impact of Artificial Intelligence on productivity, distribution and growth: Key mechanisms, initial evidence and policy challenges* (OECD Artificial Intelligence Papers No. 15; OECD Artificial Intelligence Papers, Vol. 15). <https://doi.org/10.1787/8d900037-en>
- Frey, C. B., & Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, 114, 254–280. <https://doi.org/10.1016/j.techfore.2016.08.019>
- Galasso, A. (2025). Alberto Galasso discussion of: AI harm & adoption. *Economic Policy*, 40(121), 221–224. <https://doi.org/10.1093/epolic/eiae060>
- Gans, J. S. (2025). How learning about harms impacts the optimal rate of artificial intelligence adoption. *Economic Policy*, 40(121), 199–219. <https://doi.org/10.1093/epolic/eiae053>
- Ghallab, M., Nau, D. S., & Traverso, P. (2004). *Automated planning: Theory and practice*. Elsevier/Morgan Kaufmann.

- Gmyrek, P., Berg, J., Bescond, D., & International Labour Organization. Research Department,. (2023). *Generative AI and jobs: A global analysis of potential effects on job quantity and quality*. ILO. <https://doi.org/10.54394/FHEM8239>
- Guarascio, D., Reljic, J., & Stöllinger, R. (2023). *Artificial Intelligence and Employment: A Look into the Crystal Ball*.
- Han, J., Kamber, M., & Pei, J. (2012). *Data mining: Concepts and techniques* (3rd ed). Elsevier/Morgan Kaufmann.
- Hand, D. J., Mannila, H., & Smyth, P. (2001). *Principles of data mining*. The MIT Press.
- Harrington, J. E. (2025). An economic test for an unlawful agreement to adopt a third-party's pricing algorithm. *Economic Policy*, 40(121), 261–295. <https://doi.org/10.1093/epolic/eiae054>
- Hémous, D. (2025). David Hémous discussion of: The simple macroeconomics of AI. *Economic Policy*, 40(121), 65–69. <https://doi.org/10.1093/epolic/eiae061>
- Introduction to The Simple Macroeconomics of AI. (2025). *Economic Policy*, 40(121), 9–12. <https://doi.org/10.1093/epolic/eiae066>
- ISO_IEC 22989_2022 ed.1—Id.74296 Publication PDF (en). (2022, July).
- James Manyika, Saurabh Sanghvi, Ryan Ko, Jonathan Woetzel, Jacques Bughin, Michael Chui, Susan Lund, & Parul Batra. (2017, December). *Jobs Lost, Jobs Gained: Workforce Transitions in a Time of Automation*. McKinsey&Company.
- Jeon, D.-S. (2025). Doh-Shin Jeon discussion of: AI stack competition. *Economic Policy*, 40(121), 257–260. <https://doi.org/10.1093/epolic/eiae063>
- Kogan, L., Papanikolaou, D., Schmidt, L. D. W., & Seegmiller, B. (2021). *Technology, Vintage-Specific Human Capital, and Labor Displacement: Evidence from Linking Patents with Occupations*. <http://www.nber.org/papers/w29552>
- Kogan, L., Papanikolaou, D., Schmidt, L. D. W., & Seegmiller, B. (2024). *Technology and Labor Displacement: Evidence from Linking Patents with Worker-Level Data*.
- Korinek, A., & Vipra, J. (2025). Concentrating intelligence: Scaling and market structure in artificial intelligence. *Economic Policy*, 40(121), 225–256. <https://doi.org/10.1093/epolic/eiae057>

- Mandi, J., Kotary, J., Berden, S., Mulamba, M., Bucarey, V., Guns, T., & Fioretto, F. (2024). Decision-Focused Learning: Foundations, State of the Art, Benchmark and Future Opportunities. *Journal of Artificial Intelligence Research*, 80, 1623–1701. <https://doi.org/10.1613/jair.1.15320>
- Marcin, S. (2019, July). *Economic impacts of artificial intelligence*. EPRS.
- Martínez-Plumed, F., Gómez, E., & Hernández-Orallo, J. (2021). Futures of artificial intelligence through technology readiness levels. *Telematics and Informatics*, 58, 101525. <https://doi.org/10.1016/j.tele.2020.101525>
- Meindl, B., Frank, M. R., & Mendonça, J. (2021). *Exposure of occupations to technologies of the fourth industrial revolution* (No. arXiv:2110.13317). arXiv. <https://doi.org/10.48550/arXiv.2110.13317>
- Nestor Maslej, Loredana Fattorini, Raymond Perrault, Vanessa Parli, Anka Reuel, Erik Brynjolfsson, John Etchemendy, Katrina Ligett, Terah Lyons, James Manyika, Juan Carlos Niebles, Yoav Shoham, Russell Wald, and Jack Clark. (2024). *The AI Index 2024 Annual Report*. AI Index Steering Committee, Institute for Human-Centered AI, Stanford University.
- Noy, S., & Zhang, W. (2023). *Experimental evidence on the productivity effects of generative artificial intelligence*.
- OECD. (2024). *OECD Compendium of Productivity Indicators 2024*. OECD. <https://doi.org/10.1787/b96cd88a-en>
- Peng, S., Kalliamvakou, E., Cihon, P., & Demirer, M. (2023). *The Impact of AI on Developer Productivity: Evidence from GitHub Copilot* (No. arXiv:2302.06590). arXiv. <https://doi.org/10.48550/arXiv.2302.06590>
- Pizzinelli, C. (2023). Labor Market Exposure to AI: Cross-country Differences and Distributional Implications. *IMF Working Papers*, 2023(216), 1. <https://doi.org/10.5089/9798400254802.001>
- Russell, S. J., & Norvig, P. (with Davis, E., & Edwards, D.). (2016). *Artificial intelligence: A modern approach* (Third edition, Global edition). Pearson.
- Sapi, G. (2025). Geza Sapi discussion of: Testing for unlawful pricing algorithm adoption. *Economic Policy*, 40(121), 297–300. <https://doi.org/10.1093/epolic/eiae062>

- Stuhler, J. (2025). Jan Stuhler discussion of: AI & EU Jobs. *Economic Policy*, 40(121), 141–143. <https://doi.org/10.1093/epolic/eiae068>
- Svanberg, M., Li, W., Fleming, M., Goehring, B., & Thompson, N. (2024). Beyond AI Exposure: Which Tasks are Cost-Effective to Automate with Computer Vision? *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.4700751>
- Syverson, C. (2010). *What determines productivity?*
- The European House - Ambrosetti. (2023). *AI 4 Italy: Impatti e prospettive dell'intelligenza artificiale generativa per l'Italia e il Made in Italy*.
- Tolan, S., Pesole, A., Martínez-Plumed, F., Fernández-Macías, E., Hernández-Orallo, J., & Gómez, E. (2021). Measuring the Occupational Impact of AI: Tasks, Cognitive Abilities and AI Benchmarks. *Journal of Artificial Intelligence Research*, 71, 191–236. <https://doi.org/10.1613/jair.1.12647>
- Van Noordt, C., & Misuraca, G. (2022). Artificial intelligence for the public sector: Results of landscaping the use of AI in government across the European Union. *Government Information Quarterly*, 39(3), 101714. <https://doi.org/10.1016/j.giq.2022.101714>
- Webb, M. (2020). *The Impact of Artificial Intelligence on the Labor Market*. https://web.stanford.edu/~mww/webb_jmp.pdf
- Zala, B. (2024). Should AI stay or should AI go? First strike incentives & deterrence stability. *Australian Journal of International Affairs*, 78(2), 154–163. <https://doi.org/10.1080/10357718.2024.2328805>
- Zarifhonarvar, A. (2024). Economics of ChatGPT: A labor market view on the occupational impact of artificial intelligence. *Journal of Electronic Business & Digital Economics*, 3(2), 100–116. <https://doi.org/10.1108/JEBDE-10-2023-0021>