Drivers of AI Adoption and Investment Intentions: Insights from Irish SMEs

April 2026

ESRI Working Paper No. 830

Marta Alvaro-Taus[a,b], Keith Fitzgerald[c], Janez Kren[a,b], Cynthia O'Regan[d] and Conor O'Toole[a,b]

a: Economic and Social Research Institute, Dublin, Ireland
b: Trinity College Dublin, Dublin, Ireland
c: Department of Enterprise, Tourism and Employment
d: Department of Finance

* Corresponding author:
Dr Marta Alvaro-Taus
Economic and Social Research Institute,
Whitaker Square, Sir John Rogerson's Quay,
Dublin, Ireland
Email: Marta.Alvaro-Taus@esri.ie

Acknowledgements:
We gratefully acknowledge funding from the Department of Enterprise, Tourism and Employment through the DETE–ESRI Research Programme.

Abstract

In this paper, we use novel survey data to consider factors impacting the adoption and intentions to invest in artificial intelligence (AI) by Irish micro-, small- and medium-sized enterprises (SMEs). More specifically, we explore the effects of firm characteristics, barriers to adoption and investment (such as access to finance, security and cost concerns) as well as self-reported knowledge, awareness and understanding of the benefits of AI. We find knowledge-related barriers are informative in explaining differences in current adoption and investment intention rates, whereas operational barriers such as access to finance, security concerns or skills, have a strong impact on shaping future intentions to invest. From a policy-perspective, our results suggest that raising firms' awareness of the benefits of AI to their business and sector, as well as providing a stable and certain regulatory and security environment, can enhance AI take up among SMEs.

Keywords: AI adoption ; AI investment intentions; Small and Medium-Sized Enterprises (SMEs) ; barriers to adoption ; knowledge gap
JEL codes:D22, D25, M2, O3, L2, J2

1 Introduction

The usage and integration of artificial intelligence (AI) into daily household and business life has come into increased focus since the release of ChatGPT 3.5 into the public domain in late 2023. Since then, the spread of generative AI has increased rapidly. For proponents of AI, it is seen as a major general-purpose technology which can enhance efficiency and boost flagging productivity levels, especially in advanced economies (Baily et al., 2023; Trammell & Korinek, 2023; Korinek & Suh, 2024; Benassi et al., 2022). However, others are more gloomy about the effects noting that AI has the potential to substitute capital for labour and have implications for inequality and the wage distribution (Williamson et al., 2025; Eloundou et al., 2023; Albanesi et al., 2025). Understanding which of these effects are likely to predominate is an empirical question and is expected to vary notably across economies.

For those proponents of the productivity benefits of AI, the impact depends on both the adoption and investment in the technology and its complementarity with the existing factors of production at the firm level. Indeed, there is still considerable debate around the extent to which AI tools can be integrated into the business activities of smaller firms, in particular those in domestically non-traded services sectors, and whether the enabling conditions are in place for these firms to gain from the digital transition. There is also a gap in the evidence base as to the extent to which smaller firms use these technologies and view them as important for their growth and development.

Given the relatively short timescale since the widespread availability of AI technologies, data on the usage of these technologies is limited. For Ireland, and the broader European context, there are data available from the CSO/Eurostat survey on enterprise ICT. However, this survey contains a limited number of observations when specific cohorts of the SME population are of interest, such as smaller and micro enterprises or specific age groups, while access to the underlying firm-level data is restricted. Therefore, new information is required to explore the usage of these technologies across enterprises to inform policy.

This research addresses this gap by leveraging new survey evidence on AI usage among Irish SMEs from the Department of Finance Credit Demand Survey (CDS). In the most recent survey (which went into the field in January 2025), a number of additional questions were included to gauge the penetration of AI across enterprises and to understand barriers to further investment in AI. To our knowledge, this is one of the first surveys that captures firms intentions to invest in AI along with the barriers that they are facing.

More specifically, we exploit this unique firm-level data from the CDS to examine the determinants of AI adoption and firms' intentions to invest in AI. We use new indicators capturing knowledge-based barriers, such as whether firms understand AI, whether they are aware of other firms in their sector using it, and other measures of AI awareness. These are coupled with data on operational barriers, including skills shortages, access to finance, security concerns, regulatory issues, and cost-related constraints.

To understand heterogenous effects across firms, we look at differences across firms by size, sector, region, and age structure, as well as characteristics associated with higher productivity, financial dynamism, and employment growth. It is important to note that this study does not analyse employment effects directly, as such outcomes lie beyond the scope of the available data. Nevertheless, we are able to document correlations suggesting that, at this early stage, expanding SMEs are more likely to express an interest in investing in AI.

The paper begins by analysing the raw data to extract key stylised facts, providing an overview of the dominant patterns and disparities in AI adoption across different types of firms. It then proceeds with the empirical analysis, which aims to statistically identify the salient factors associated with either AI adoption or future diffusion prospects, after jointly controlling for all potentially relevant determinants. To do so, we model firms' AI adoption decisions and their investment intentions as two distinct yet interrelated processes governed by unobserved latent variables. Specifically, we estimate these relationships using discrete-choice binary probit models, alongside a system of bivariate seemingly unrelated probit regressions that accounts for the direct association between existing AI usage and anticipated investment plans.

A number of findings emerge from this research. First, we find a general effect across both AI adoption and investment intentions, whereby the predominant factors reducing the likelihood of adopting AI or planning future investments are a limited understanding of the benefits of AI and insufficient knowledge of its effectiveness. This lack of understanding about the feasibility of AI represents a key barrier to firms adopting digital enhancements, which can be a catalyst for improving the productivity of their enterprise. Policy supports to boost awareness of the functionality of AI should help to improve understanding in this context, and ultimately, foster more positive attitudes towards these technologies. Other factors also having a negative influence on investment intentions are access to finance and security constraints. Thus, designing targeted policy measures that address both financial constraints and security related-risks are important to foster greater AI take-up and diffusion rates.

There are also notable sectoral differences with firms in professional, technical and scientific as well as financial, insurance and high-tech manufacturing all exhibiting higher levels of AI adoption. The broad based productivity benefits from AI on a macroeconomic level are likely to be best felt the wider the scope of use of the technology. Therefore, widening the adoption of AI outside these first-mover sectors is likely to help deepen the overall impact on the Irish economy. In this regard, the sectoral effects on firms' intentions to invest do not differ significantly across each other, indicating a favourable environment for broad-based future AI adoption, if the right policy levers are put in place to enable SMEs to realise intended investments.

Focusing further on the determinants of intention to invest in AI, we find that this is highly correlated with having existing AI technologies in use by the enterprise. We also find that firms reporting higher barriers relating to staff skills or costs have a higher rate of intention to invest. While this relationship is likely to be endogenous, the reported positive coefficient may reflect that firms actively seeking to invest in AI are worried about the skills absorption capacity of their business, as they consider increasing digitalisation efforts. As such, those firms already engaging with the technology tend to be more aware of the practical challenges and related operational barriers associated with these type of innovation.

Regarding other firm characteristics, we find positive associations between investment intentions and being a high-productivity firm, a firm experiencing positive employment growth, or an exporter. The evidence related to employment growth is particularly noteworthy. It indicates that, at this early stage of AI adoption, there is no sign of detrimental effects on employment; rather, more dynamic firms appear to display a greater appetite for innovation.

Finally, to assess the robustness of our findings, we cross-validate the results using machine-learning methods. To ensure that the patterns we identify are not driven by model-specific assumptions, we employ a Random Forest classifier combined with SHAP analysis. The results from this exercise reinforce the central insights from our main econometric analysis that knowledge-related barriers exert a broad and consistent influence across both outcomes, and that early exposure to AI emerges as a strong catalyst for subsequent investment.

From a policy perspective, two priority areas arise. First, a targeted information campaign is suggested which can help improve firms' understanding of the use and benefits of AI to their enterprise and their sector. Second, policy should focus on ensuring the enabling conditions are in place including access to finance, robust regulatory and security frameworks, and targeted workforce upskilling. This should ensure that the uncertainty and risk is lowered for enterprises seeking to procure AI inputs.

The rest of this paper is structured as follows. First, we discuss the policy background and literature. Second, we present the data and summary statistics. Third we present the empirical model followed by the empirical findings and the robustness check. We then conclude and draw out the implications for policy.

2 Policy background and related literature

2.1 Policy context

Ireland's approach to AI and enterprise policy has developed rapidly over the past decade, shaped by both the emergence of the EU as a standard-setter for AI governance, and the acceleration of the capabilities of AI technology, particularly with respect to recent innovation in generative AI.

Ireland's policy framework for AI adoption by enterprise is anchored in the Government's National Digital and AI Strategy: Digital Ireland – Connecting our People, Securing our Future (published in 2026). This updated strategy aims to consolidate Ireland's position as a digital leader and AI hub within Europe. The strategy sets out 20 high level objectives supported by 90 deliverables, structured across five strategic ambitions (apply, grow, invest, lead, and empower).

Digital Ireland outlines an ambition to deepen Ireland's position as a digital leader while ensuring that SMEs can effectively participate in, and benefit from, the accelerated adoption of AI and digital technologies. Acknowledging the barriers facing small firms in integrating advanced digital tools, the strategy outlines a series of supports aimed at improving adoption across SMEs. The strategy includes a suite of initiatives aimed at building digital capability among SMEs, including an AI and digital literacy campaign, sector specific AI adoption roadmaps, and the appointment of AI Sector Champions to help smaller firms to understand transformative technologies.

Another key aspect of Digital Ireland is the creation of the AI Office of Ireland, which will act as the national coordinating authority for the EU AI Act. This office will also host an AI Regulatory Sandbox, a mechanism facilitating SMEs to trial AI systems in a controlled environment, aimed at reducing the cost and risk of compliance. The strategy also emphasises workforce capability as a foundation for SME competitiveness. A national Digital and AI skilling campaign, complemented by an AI skills platform, seeks to address the skills gap that can constrain smaller firms.

These are positive developments, which should go some way towards encouraging greater take-up of AI technology by smaller firms. These measures aim to support EU aligned objectives (EU Digital Decade 2030) such as achieving 90 per cent basic digital intensity among SMEs by 2030 and boosting adoption of cloud, AI, and data driven technologies across the enterprise base.

The European Commission's "Annual Report on the State of the Digital Decade" (June 2025), measures digitalisation in Ireland relative to other EU member states, tracking progress towards these targets. The 2025 report shows that Irish enterprises compare well with their EU counterparts:

- On measures of the digital intensity of SMEs, Ireland ranks 11th, at 73.4 per cent compared to the EU average of 72.71 per cent;
- On AI use, 13.8 per cent of Irish SMEs report using AI, compared to an EU average of 12.6 per cent (this has increased significantly from 2023, when only 8 per cent of SMEs reported using AI).

AI adoption is generally rising across the EU, but progress is uneven, with large firms remaining two-to-three times more likely than SMEs to use AI. This dual-track pattern, with rapid uptake among large enterprises and slower diffusion among smaller firms, has implications for competitiveness and broad-based productivity growth. Eurostat data shows that patterns of adoption in Ireland broadly reflect this European divide. Large-firm adoption is relatively high at 51 per cent, well above the EU average, while SME adoption remains modest at 14 per cent.

The digital maturity of enterprise appears to be a key determinant of AI use. Among Irish SMEs, AI adoption ranges from 3.5 per cent for digitally intensive firms to just 0.9 per cent for low-intensity firms, the latter below the EU average. A similar pattern holds for large firms: digitally advanced enterprises report adoption rates of 13.8 per cent, while low-intensity firms remain at 0.5 per cent. The divergence between low- and high-digital-intensity firms points to the risk of a two-tier digital economy, where already digitally sophisticated firms are best positioned to benefit from AI-driven productivity gains.

Other recent, significant policy developments in Ireland include the establishment of a Minister of State for AI and Digital Transformation in 2024, signalling the political prioritisation of AI as an economic and regulatory issue. Furthermore, Ireland established the AI Advisory Council, tasked with providing evidence-based advice on emerging AI risks, enterprise adoption, and international competitiveness. The Council is chaired by Ireland's first AI Ambassador. Furthermore, the current Programme for Government commits to updating Ireland's National Digital Strategy, to realise the full benefits of digitalisation, which includes AI, and to increase the productivity of Irish enterprises. This updated strategy will maintain alignment with the EU Digital Decade targets, while introducing more ambitious domestic goals, with a strong emphasis on AI adoption. The updated strategy is being developed on a whole-of-Government basis.

Among the initiatives included in the updated strategy, is a proposal to establish an AI Office of Ireland, which will act as the coordinating authority for the EU AI Act, and will be a focal point for responsible AI innovation. This new office will also host an AI Regulatory Sandbox, which aims to provide a safe, supervised environment for innovators to test new digital and AI solutions in partnership with compliance experts. The updated strategy also outlines plans to raise awareness of the full potential of digital and AI tools among SMEs, through the launch of a Nationwide AI and Digital literacy campaign. These are positive developments, which should go some way towards encouraging greater take-up of AI technology by smaller firms.

For small open economies such as Ireland – where international competitiveness is closely tied to innovation capacity and the performance of high-value export sectors – a slower pace of AI adoption represents a key risk. The institutional developments outlined above, are indicative of the emergence of a policy consensus regarding the importance of AI technology, as a determinant of firm-level productivity and innovation. This is emphasised in recent national strategies, including the Action Plan on Competitiveness and Productivity (2025).

Within this context, the uneven diffusion of digital and AI capabilities across the enterprise base – particularly between multinationals and SMEs – raises questions about the capacity of different cohorts of firms to benefit from emerging technologies and the opportunities they present. Ensuring that firms of varying sizes and digital maturity can effectively adopt and exploit AI is central to sustaining productivity growth and maintaining Ireland's competitive position in an increasingly technology-driven global economy.

2.2 Related literature

Technological progress has historically delivered substantial gains in productivity and living standards. However, previous waves of innovation have also presented significant transition costs. AI is distinct, in its capacity to affect cognitive and non-routine tasks across a wide range of occupations, extending the scope of technological change beyond routine or manual work. At the firm level, this creates significant potential for productivity gains through improved decision-making, faster information processing, and more efficient task allocation.

A central principle underpinning much of the recent literature, is the task-based view of employment, which frames occupations as bundles of tasks with AI affecting specific skills rather than entire jobs. This perspective underpins much of the empirical work on AI exposure, and provides an analytical framework for distinguishing between labour substitution and labour complementarity (Felten et al., 2021, 2023; Pizzinelli et al., 2023; Cazzaniga et al., 2024). Under this framework, AI may automate certain tasks while simultaneously increasing the productivity of others, implying that the economic impact of AI depends on how tasks are reorganised within firms rather than on the simple replacement of workers. Applying this framework to Ireland, Williamson et al. (2025) find that approximately 63 per cent of employment is relatively highly exposed to AI. Within this group, around 30 per cent of jobs fall into an "at-risk" category, defined as occupations that are exposed but exhibit low complementarity, including telephone salespersons, chartered and certified accountants, and laboratory technicians. Overall, Ireland's exposure profile appears broadly in line with that of other advanced economies (Pizzinelli et al., 2023; Cazzaniga et al., 2024). Consistent with these findings, AI adoption may exacerbate income and wealth inequality, as higher-skilled and higher-paid workers tend to occupy roles with greater AI complementarity (Williamson et al., 2025; Eloundou et al., 2023; Albanesi et al., 2025).

Micro-level evidence on AI adoption by firms remains limited, reflecting the relatively recent and uneven diffusion of the technology. Nevertheless, the available studies generally point to positive, albeit heterogeneous, firm-level productivity effects. Studies using patent data and related proxies find that firms with greater AI-related activity tend to exhibit higher total factor productivity and, in some cases, higher wages (Bessen & Righi, 2019; Bassetti et al., 2020). Related work shows that the accumulation of AI-related and broader "Fourth Industrial Revolution" knowledge is positively correlated with labour and total factor productivity (Benassi et al., 2022; Alderucci et al., 2019; Damioli et al., 2021). Czarnitzki et al. (2023) uses survey data on AI adoption and similarly finds a positive and significant association between AI use and firm productivity.

The broader macroeconomic effects will depend on the extent to which AI adoption diffuses across firms. These effects may be observed through several possible channels. The process of (intangible) capital deepening can drive aggregate productivity growth (Acemoglu & Restrepo, 2019). However, model-based studies show that, in the absence of sufficient labour-augmenting investment, AI adoption can reduce the labour share even as output grows (Qian et al., 2023; Trammell & Korinek, 2023). At the same time, AI adoption may reinforce market concentration due to high fixed costs, first-mover advantages, and control over data, giving rise to "winner-takes-most" dynamics (Brynjolfsson & Unger, 2023).

Reflecting the limited availability of empirical data on AI adoption, the literature on the macroeconomic effects is dominated by model-based scenario analyses. Some authors argue that AI may compress the traditional productivity "J-curve" by leveraging existing digital infrastructure, allowing productivity gains to materialise more quickly than in previous technological waves (Baily et al., 2023). More speculative contributions explore scenarios ranging from modest productivity gains to explosive growth outcomes under the emergence of Artificial General Intelligence (AGI), with profound implications for wages, the labour share, and income distribution (Trammell & Korinek, 2023; Korinek & Suh, 2024). A recurring conclusion in the literature, is that policy will be decisive in influencing whether AI adoption favours labour-augmenting or labour-replacing dynamics, shaping market concentration, and determining access to AI technologies across firms of different sizes (Brynjolfsson & Unger, 2023). As a result, the macroeconomic impact of AI will depend critically on how firms, workers, and policymakers shape diffusion of the technology.

Recent work by the ESRI adds Ireland-specific evidence on enterprise digitalisation more broadly, providing relevant context for how AI adoption may ultimately diffuse across firms. Using firm-level microdata, Kren & O'Toole (2024) show that investment in digital assets by Irish micro, small and medium-sized enterprises is highly heterogeneous and shaped by firm characteristics, with profitability, growth prospects and firm size playing a central role in determining adoption. The paper finds that smaller and older firms are significantly less likely to invest in digitalisation. Complementing this analysis, Hogan et al. (2024) places Irish enterprise digitalisation in a cross-country perspective, drawing on Eurostat and European Investment Bank data to compare Ireland's uptake of digital technologies with European peers. They find that Ireland performs comparatively well in terms of e-commerce and advanced technologies, with low levels of digitalisation in the construction sector and among smaller firms.

From the perspective of AI adoption, these findings are particularly relevant, as the literature consistently identifies baseline digital capabilities (such as data infrastructure, digital processes and ICT investment) as important prerequisites for the effective deployment of AI technologies. It stands to reason then, that promoting greater levels of enterprise digitalisation, and tackling the barriers that smaller firms face in digitalising, will be key to driving greater overall take up of digital technologies.

Building on these insights, the analysis that follows contributes to the literature by providing AI-specific micro-level evidence on the factors constraining the diffusion of AI across firms.

3 Descriptive statistics

3.1 Data

Throughout this paper, we draw on data from the Department of Finance's 2024 Credit Demand Survey (CDS) to provide novel insights into AI adoption among Irish SMEs. Data for this wave of the survey were collected in January 2025, with questions referring to activities undertaken in 2024. It includes a section dedicated to gauge the penetration of AI across enterprises and to understand their most recent attitudes towards these technologies. In total, 1,503 enterprises were interviewed. The study sample is representative of the overall SME population based on size and sector (as defined by NACE codes), as it was selected at random from an industry-leading database of all SMEs across Ireland. The classification of companies as micro, small, or medium enterprises in the dataset follows the standard EU definition: micro enterprises employ between 1 and 9 people, small enterprises between 10 and 49, and medium-sized enterprises between 50 and 250.

For the purpose of the survey, the definition of AI used during the interview was "AI technologies can be: i) Embedded in devices, like autonomous robots or drones, ii) or, software-based, such as chatbot virtual assistants based on natural language processing, object and face recognition, speech recognition, machine translation, or data analysis using machine learning."

Based on this definition, companies were asked about their recent experience using AI technologies in the production of goods or the delivery of services, including obstacles to AI adoption and their understanding of the potential benefits and knowledge of applications for their business. Among these questions, key for our analysis are understanding the extent of AI adoption and plans for future investment, as well as the reported barriers to adoption, which provide essential insights into the factors shaping AI intake patterns — central to address the policy questions this study seeks to explore.

To complement the analysis of this novel data, we use data from Eurostat's 2024 ICT Usage Survey. The ICT Usage Survey is collected on a yearly basis by National Statistical

Institutes based on the annual Eurostat model questionnaires on ICT usage and e-commerce in enterprises. The aim of the survey is to collect and disseminate harmonised and comparable information on the use of digital innovations at the European level, including artificial intelligence. The definition of AI used in the Eurostat questionnaire is slightly more specific than the CDS's, it is essentially the same: "Artificial intelligence systems can be purely software based, e.g.: [...]; or embedded in devices, e.g. [...]". A key advantage of the Eurostat data is its coverage of large enterprises, which allows us to clearly portray the presence of a dual-track adoption pattern, i.e. the divide between large firms and small and medium-sized enterprises. A limitation, however, is that micro-enterprises are not covered. As such, we use this data to compare our primary indicators of AI adoption by firm size (derived from the CDS) with adoption rates from this more established data source, including those for large enterprises. The analysis of this data also allows us to benchmark Ireland against the broader European landscape and, in doing so, provide a comprehensive overview of the AI adoption across Europe.

3.2 AI adoption patterns among Irish SMEs: stylised facts

This section presents the key stylised facts derived from the data, providing an overview of the dominant patterns and disparities in AI adoption that underlie the subsequent empirical analysis.

3.2.1 An international comparison

First, we present comparative analysis of AI adoption rates across European countries based on Eurostat's ICT Usage Survey, which places Ireland's adoption rate in the wider European context.

Figure 1 ranks countries by their AI adoption rates—defined as the share of firms reporting the use of at least one type of AI tool—distinguishing between large firms (Panel a) and SMEs (Panel b). These figures highlight a pronounced and systematic gap in AI uptake between large enterprises and their smaller counterparts across countries, i.e. the so-called "digital divide" across firm size. In the case of Ireland, more than 50% of large companies report using at least one type of AI technology — significantly above the EU-27 average of 41% in 2024 – placing Ireland 7th among EU countries. In contrast, the AI adoption rate among Irish SMEs stands at 13.8%, only one percentage point above the EU average, ranking 11th across member states.

3.2.2 AI adoption rates from the Credit Demand Survey

Using responses from the CDS on the extent to which companies employed AI technologies between 2022 and 2024, we calculate our own overall AI adoption rate for SMEs, as well as separate rates for each enterprise size category – micro, small, and medium firms. We then compare these figures with the corresponding adoption rates reported in the ICT Usage Survey to cross-validate the findings from both sources.

The left-hand panel of Figure 2 illustrates the overall AI adoption rate for SMEs, and the rates for individual firm size categories. The overall adoption rate among Irish SMEs

Figure 1: European Cross-Country Comparison of AI Adoption rates

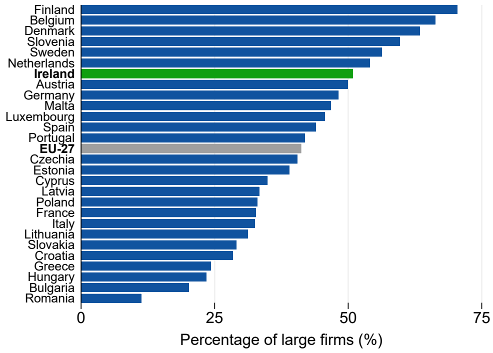

(a) Large Companies

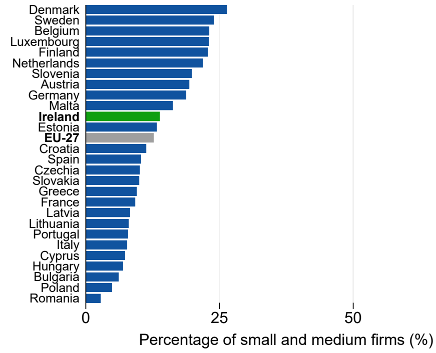

(b) SMEs

Source: Eurostat ICT Usage Survey (2024)

is 16.5%, slightly above the 13.8% rate encountered by Eurostat's ICT usage survey. This figure clearly indicates that smaller firms exhibit lower adoption rates, with 14.3% of micro enterprises, 15.2% of small companies, and 22.3% of medium enterprises reporting using AI. The right-hand panel contrasts these figures with those from the ICT Usage Survey. Overall, the figures are broadly consistent. While Eurostat does not report data for micro enterprises, the adoption rate for small firms is somewhat lower (11.8% versus 15.2%), while the rate for medium-sized firms is slightly higher (24.9% compared to 22.3% in the CDS). Most importantly, we observe that adoption rates increase with firm size in both cases.

Figure 2: AI Adoption rates by firm size, 2024

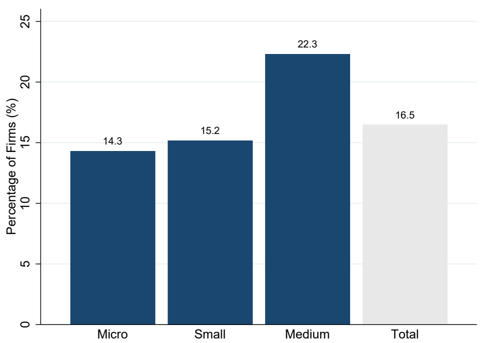

(a) Credit Survey Demand (CDS)

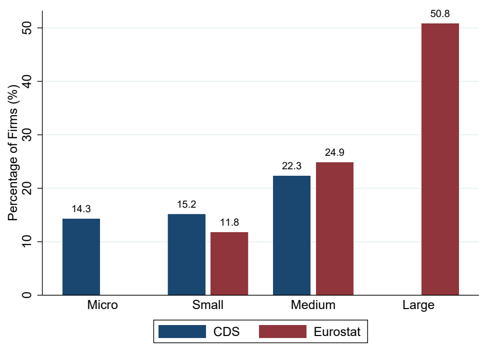

(b) CDS vs Eurostat

Source: Credit Demand Survey and Eurostat ICT Usage Survey (2024)

Next, we examine data on whether adopting firms integrate AI into their operations on a formal basis (Table 1), and their overall satisfaction with the performance of these technologies (Table 2).

Table 1 breaks down adopters by usage level and tabulates this against whether AI was adopted on a formal or ad-hoc basis.[1] Overall, the figures reveal that most companies use AI on an ad-hoc basis (55.5%). Nevertheless, a large share of firms with high AI intensity usage report having put in place formal strategies to integrate AI tools into their workflows.

Table 1: Formal vs Ad-hoc use of AI technologies

Usage level	**Formal**	**Ad-hoc**	**Don't know**	**Total**
Tested	9.3	67.2	23.5	100.0
Low	11.0	53.3	35.6	100.0
Medium	38.3	50.8	10.9	100.0
High	54.0	39.0	7.1	100.0
Total	16.0	55.5	28.5	100.0

Source: Credit Demand Survey (2025) and authors' calculations

Table 2 cross-examines AI usage levels against satisfaction relative to expectations. A key takeaway is that the proportion of firms reporting better-than-anticipated outcomes increases with usage intensity. Notably, firms with high AI usage levels overwhelmingly report that AI technologies exceeded their expectations (93%); while firms whose experience with AI is limited to the testing stage predominantly report mixed results (55.8%). Nevertheless, below-expectation ratings are rare, i.e. 7.6% in total. Overall, this paints a positive outlook for AI integration, suggesting that when effectively implemented, AI tools deliver clear benefits for work performance.

Table 2: Satisfaction relative to expectations

Usage level	**Below**	**As expected**	**Better**	**Mixed**	**Total**
Tested	8	14	22	56	100
Low	6	34	23	37	100
Medium	13	38	33	16	100
High	2	0	93	5	100
Total	7.6	29	26	37	100

Source: Credit Demand Survey (2025) and authors' calculations

[1] The question asked in this regard was "Are you aware of individual staff using AI in their work, and if so, do they use it on an ad hoc or formal basis?". Given this question, employers responding "ad-hoc" may be including the use of "shadow" AI use — that is, the use of personal AI tools, like personal ChatGPT, Claude subscriptions, etc., to make part of their job more productive.

3.2.3 Intentions-to-invest in AI

Next, we report the intentions-to-invest rate based on responses to whether firms plan to invest in AI between 2025 and 2027. This forward-looking indicator complements the measure of AI adoption by providing insights into firms' considerations beyond current usage levels. This is particularly relevant given that many AI innovations are recent and not yet widely diffused. Moreover, in light of the widespread discussion about AI and its potential impact on the workplace, understanding firms' investment intentions offers a valuable perspective on how quickly adoption may progress based on current expectations.

Figure 3 depicts the overall intentions-to-invest rate and the disaggregation by firm-size (Panel a) and adoption status (Panel b). Overall, 23.4% of firms plan to invest in AI within the next three years. Medium enterprises have much larger intentions to invest (39%) than micro (16%) and small (21%) firms. Of all those planning to invest in AI, 8.7 percentage points consist of firms that have already implemented some form of AI and are therefore expected to deepen or expand their usage. The remaining 14.6 percentage points correspond to firms that have not yet adopted AI. This group reflect the group of prospective new users as things stand. Adding this figure to the current adoption rate suggests a tentative adoption rate of 31% in three years' time. Any outcome exceeding this estimate could be interpreted as an unexpectedly rapid acceleration in AI uptake.

Figure 3: AI Investment Intention rates by Firm Size and Adoption Status, 2024

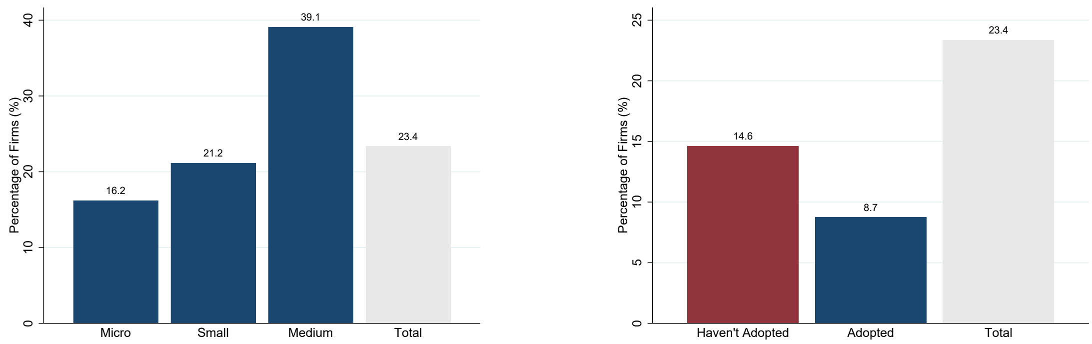

(a) Firm Size (b) Adoption Status

Source: Credit Demand Survey (2025) and authors' calculations

3.2.4 Firm attributes, AI usage and intentions-to-invest

In this section of the descriptive analysis, we examine AI adoption rates and investment prospects by firm characteristics, highlighting several emerging patterns from the data (Figure 4).

We find that exporters are more prone to be using some type of AI, and also express stronger intentions to invest in AI in the near future: 22.4% of exporters have adopted AI compared to 15.1% of non-exporters, while 37.8% of exporters plan to invest in AI compared to 19.9% of non-exporters (Figure 4a).

By age, there is less of an outstanding pattern (Figure 4b). While younger firms show slightly higher adoption rates than older companies (i.e. 19% for firms younger than 10 years, 17% for firms that existed between 10 and 19 years, and 16% for firms older than 20 year), all age brackets show larger intention-to-invest ratios, with a less distinguishable trend.

A sectoral breakdown reveals that business services —particularly Professional, Scientific, and Technical Activities, as well as Administrative and Support Services — and high-tech manufacturing lead in intentions-to-invest (Figure 4c). Financial Services stands out with the highest outstanding adoption rate at 40%. Interestingly, however, its intentions-to-invest rate is significantly lower than these other three industries (26% versus 44%, 36%, and 35%).

By region, the South-East leads in terms of current AI uptake, with an adoption rate of 28% — nearly 10 percentage points higher than the next-highest region (Figure 4d). This pattern likely reflects the region's strong concentration of pharmaceutical and high-tech manufacturing firms. In contrast, Dublin and the Mid-East display a more favourable outlook for future adoption: over 33% of SMEs in Dublin and 25% in the Mid-East report intentions to invest in AI within the next three years, compared with 19% in the South-East.

Figure 4: AI Adoption rates by Key Firm Attributes

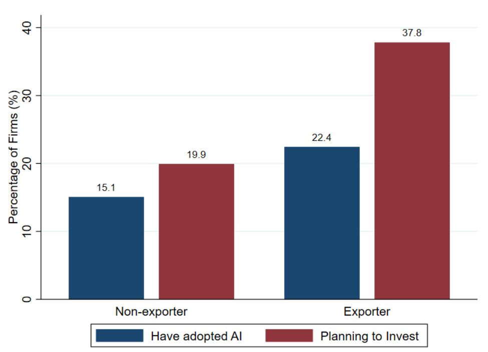

(a) Exporters

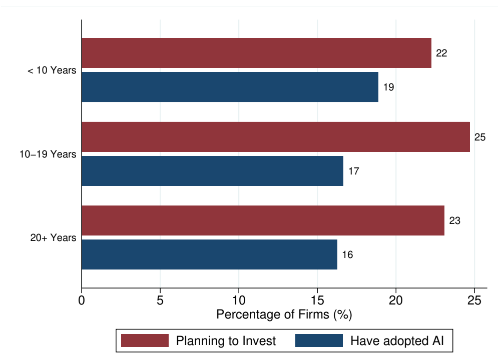

(b) Age

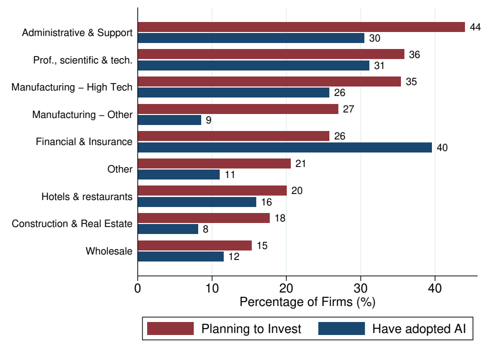

(c) Sectors

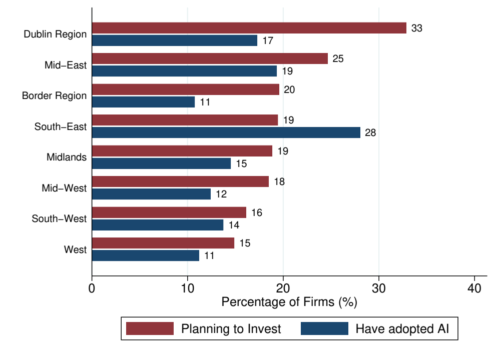

(d) Regions (NUTS-3)

Source: Credit Demand Survey (2025) and authors' calculations

3.2.5 Operational barriers and knowledge-related indicators

To conclude this section, we examine firms' attitudes towards AI, focusing on their responses to questions concerning perceived obstacles, including operational barriers and a lack of knowledge about how to derive value from AI applications. Specifically, our analysis is based on ten statements from the CDS questionnaire, for which firms reported whether they agreed or disagreed. Six of these statements are most closely related to 'Operational barriers', while the remaining four can be best associated with 'Knowledge-related barriers'. We adopt this distinction between types of obstacles and use it throughout the report.

Figures 5 and 6 display both sets of indicators, broken down by intentions-to-invest and adoption status. The left-hand panels illustrate the percentage of non-adopters reporting negatively on each statement, disaggregated by investment intentions. The right-hand panels present the corresponding statistics for AI adopters.

Starting with the operational barriers (Figure 5), security concerns and regulatory uncertainty emerge among the most frequently cited obstacles among firms with no intention to invest in AI. In particular, these barriers are reported by approximately 36% and 35%, respectively, of non-adopters with no plans to invest in AI (Figure 5a). Among firms that are already using AI but do not intend to invest further, security concerns again emerge as the most prominent barrier, followed by staff training costs and regulatory issues (47%, 35%, and 33%, respectively - Figure 5b). Notably, the share of AI adopters with no plans for further investment reporting security concerns (47%) is substantially higher than that observed among non-adopters without investment intentions (36%).

In contrast, firms planning to invest in AI place greater emphasis on operational barriers related to staff skills and procurement costs. Among non-adopters intending to invest, 37% and 33% report each of these barriers, respectively (Figure 5a). Among current AI users planning further investment, the corresponding shares are 30% and 34% (Figure 5b).

Turning to knowledge-related barriers (Figure 6), the most frequently reported indicators are the perception that other firms in the same industry are not using AI and the difficulty of finding relevant information on what AI adoption would imply for the firm. These concerns are voiced by both current AI users and non-users, regardless of whether they intend to invest in AI. However, the prevalence of these barriers is highest among non-adopters with no intention to invest: 61% report that other firms in their sector do not use AI, and 48% cite difficulties in finding adequate information (Figure 6a).

Figure 5: Operational Barriers by Adoption Status & Intentions-to-Invest

(a) Non-Adopters (b) Adopters

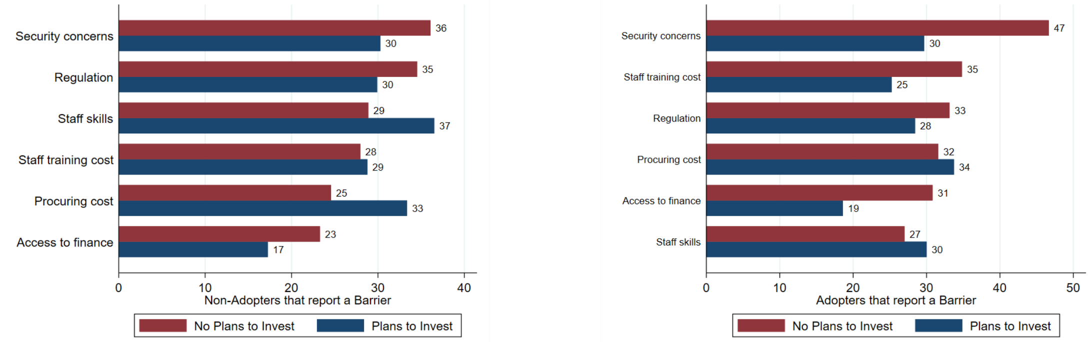

Figure 6: Knowledge-related Barriers by Adoption Status & Intentions to Invest

(a) Non-Adopters (b) Adopters

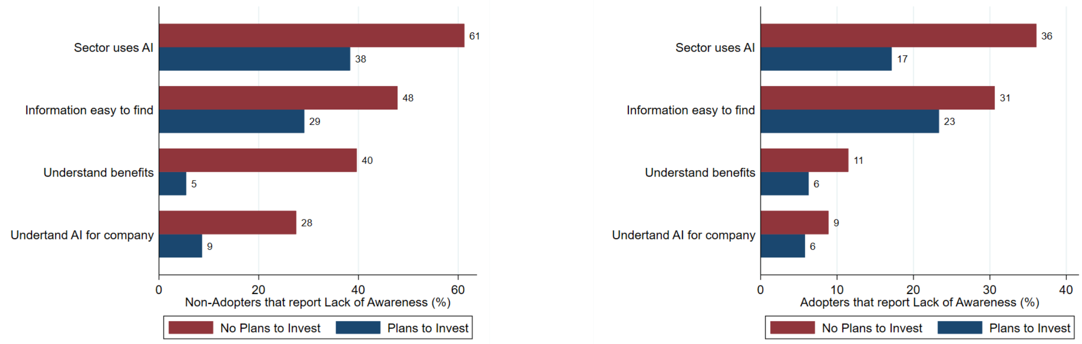

Source: Credit Demand Survey (2025) and authors' calculations

4 Methodological approach

Having established the main stylised facts from the data, we now turn to an empirical analysis aimed at identifying important factors for AI adoption and investment intentions. While the preceding section documented correlations, the empirical analysis seeks to evaluate their statistical significance, controlling jointly for all relevant determinants.

4.1 Empirical framework

In this analysis, we model firms' AI adoption decisions and investment intentions as two different but interrelated processes. From an economic standpoint, the basic intuition driving firms' behaviour is that a profit-maximising enterprise adopts AI and entertains investments in AI technologies only if the respective net benefits are positive. Since net benefits are not observed, we model the decision-making process as a bivariate discrete-choice model; that is, a system of two latent-variable equations in which the net benefits *adopt** and *invest** represent the unobservable latent variables. In formal terms,

$$
\begin{aligned}
&ai\_adopt_i = 1 \quad \text{if } ai\_adopt_i^* = x_i'\alpha + \eta_i > 0, && (1)\\
&ai\_adopt_i = 0 \quad \text{otherwise;}\\
&plans\_to\_invest_i = 1 \quad \text{if } invest_i^* = y_i'\beta + \gamma ai\_adopt_i + \varepsilon_i > 0, && (2)\\
&plans\_to\_invest_i = 0 \quad \text{otherwise.}
\end{aligned}
$$

In this model, only the binary choices for actual adoption of AI ($ai\_adopt_i = 1$) and intention to invest ($plans\_to\_invest_i = 1$) are observed provided that the latent variables exceed certain thresholds, set at zero.

The unobservable net benefits $ai\_adopt_i^*$ and $invest_i^*$ depend on observable factors captured by the vectors x_i and y_i, as well as unobservable factors summarised by the disturbance η and ε. In the estimation of the model, both sets of regressors x_i and y_i include variables belonging to four categories, namely perceived operational barriers to adoption, firms' knowledge and familiarity with AI tools, attitudes towards external financing and other firm characteristics. Among the regressors in equation (2), we also include ai_adopt_i to capture the potential direct association between existing AI usage and anticipated investment plans. This is as it is sensible to expect that plans to invest in AI are influenced by companies contemporary experience with AI technologies.

The disturbances η and ε are assumed to be normally distributed vectors with zero mean. An important assumption is that these disturbances may be correlated such that $Cov(\eta, \varepsilon) = \rho I \neq 0$, given that firms' decisions on AI adoption and investment intentions might be jointly determined by unobservable factors. As such, ρ reflects the non-idiosyncratic correlation of both decisions and I denotes the unity matrix. In our application, the correlation between η and ε could be capturing preferences and attitudes towards "innovation" of a company's owners or staff members. Such unobservable factors can affect both early adoption decisions and investment intentions simultaneously, and thus it may induce endogeneity in our model. That is, if such correlation is present and not taken into account, the parameter estimate of the effects of AI adoption on future investment plans could be capturing the influence of these unobserved drivers rather than the truthful impact of current AI usage on the decision-making process of investments.

It is important to note that the estimated coefficients cannot be interpreted causally. In particular, they may be affected by measurement error arising from perception-based responses, as well as by potential simultaneity bias. Nevertheless, as we show below, the results produce important insights into the AI adoption process and investment behaviour that go beyond the evidence from the descriptive analysis.

4.2 On the interpretation of coefficients

The functional form for a standard probit model estimating the probability that a binary dependent variable is equal to one is

$$\text{Prob}(Y = 1/x) = F(x, \beta), \quad (3)$$

where $F(.)$ corresponds to the standard normal distribution function. The set of parameters β reflects the impact of changes in x on the probability. However, given that the probit is a nonlinear regression model, the coefficients are not necessarily the marginal effects as they vary with the values of x,

$$\frac{\partial F(x'\hat{\beta})}{\partial x} = f(x'\hat{\beta}) \times \hat{\beta} \quad (4)$$

where $f(t)$ is the standard normal density.

Thus, when computing marginal effects, we follow the common practice of evaluating the expressions at every observation and use the sample average of the individual marginal effects (i.e. average partial effects) as our measure of interest when analysing the results (see, e.g. Wooldridge, 2002).

$$A\hat{P}E = \frac{1}{n} \sum_{i=1}^{n} f(x_i'\hat{\beta})\hat{\beta}. \quad (5)$$

5 Empirical results

In this section, we present the main findings for the empirical analysis. This includes the presentation of marginal effects, predicted probabilities and a discussion of the main findings emerging from the econometric model.

5.1 Main estimation results

Table 3 presents the results of our estimations. We estimate the models in equations (1) and (2) as separate probit models (Columns 1 and 2) and as a seemingly unrelated bivariate probit system of equations (Columns 3 and 4).

The unrelated bivariate probit specification estimation is meant to tackle the potential issues with simultaneity bias; that is, it accounts for the fact that AI adoption choices might endogenously affect investment intentions (a recursive structure). As we need to simultaneously estimate the probit models, the sets of explanatory variables x_i and y_i must differ by at least one variable in order to achieve identification, which in turn is required to test whether the disturbance terms ϵ and η in equation (1) and (2) are correlated (Maddala, 1983). To satisfy this requirement, we rely on different attitudes to external financing variables in each equation. Specifically, in the AI adoption equation (Column 1 and 3), we include indicators for whether a firm had requested bank financing and non-bank financing in previous years. In contrast, for the AI investment intentions equation (Column 2 and 4), we include a forward-looking measure capturing firms' intentions to seek future bank financing, under the assumption that intended AI investments may be influenced by expected access to external finance.

The estimated ρ coefficient, reported at the bottom of the bivariate probit model in Columns (3) and (4) in Table 3, is not statistically significant; therefore, the null hypothesis $H_0 : \rho = 0$ cannot be rejected. This suggests that unobserved factors affecting the two equations are uncorrelated, and it implies that considering the separate probit models is appropriate. Nevertheless, we report both sets of estimates to demonstrate that the results are broadly consistent across specifications.

Looking at the impact of AI adoption on the probability of planning to invest in AI, we find that the AI adoption coefficient is positive and statistically significant (Column 2). This indicates that firms that have already adopted AI are more likely to express an intention to expend further in AI technologies, as expected.

The first set of explanatory variables presented in Table 3 captures the effects of operational barriers on each of the outcomes. We find that procuring costs is the only statistically significant barrier for explaining the probability of AI adoption, with a positive association: firms reporting concerns about procurement costs are more likely to adopt AI. Procurement costs are also positively associated with intentions to invest in AI, suggesting that implementation costs are an important concern when considering AI innovations. For intentions to invest, the coefficient on staff skills is also statistically significant and positively correlated. These patterns are consistent with the correlations observed in the stylised facts and suggests that, overall, the investment process goes hand in hand with practical considerations related to the implementability and scalability of AI technologies. In turn, we find that the security and access to finance concerns are detrimental for putting in place an investment plan. The coefficients to both variables are negative and statistically significant in Columns (2) and (4). While the importance of security concerns was already evident from the descriptive analysis, the relevance of access to finance as a main obstacle to future investment in AI is new to the empirical analysis.

Next, we turn to the knowledge-related barriers. The coefficients on 'Do not understand the benefits of using AI' and 'Other firms in my industry do not use AI' are both negative and statistically significant for both models. In particular, perceptions that others in the sector do not use AI play a prominent role in explaining the likelihood of adopting AI. This factor is less important for explaining investment intentions, for which a lack of understanding of AI's benefits appears to be more salient. This latter pattern is intuitive from an economic standpoint, and in-line with our system of latent-variables model, which states that innovation decisions are primarily driven by expected profitability. It is also worth noting that the coefficient on the variable 'Information not easy to find' — which is frequently cited by firms, as discussed in the descriptive analysis — is not statistically significant.

Among the firm characteristics included in the model, firm performance indicators exhibit the strongest explanatory power for AI investment decisions. In particular, firms with higher productivity within their sector (i.e. those in the top quartile) and firms that have experienced employment growth over the past five years show a significantly higher likelihood of investing in AI. This is in line with related literature establishing a positive correlation between firm's productivity and AI activity (Bessen & Righi, 2019; Bassetti et al., 2020; Benassi et al., 2022; Alderucci et al., 2019; Damioli et al., 2021; Czarnitzki et al., 2023). Although these results do not allow us to draw conclusions about employment effects, the evidence does not support a detrimental relationship with employment; rather, it suggests that more dynamic firms exhibit a greater appetite for innovation.

Likewise, medium-sized firms are more likely to invest in AI than smaller firms, while exporting firms display a higher probability of both adopting AI and having an AI investment strategy in place compared with non-exporters. By contrast, although the coefficients on firm age are negative—suggesting that older firms may be less inclined to adopt AI—they are not statistically significant, indicating that firm age is not a key determinant of AI adoption or investment decisions. These findings are also consistent with existing evidence on firm-level digitalisation in Ireland (Kren & O'Toole, 2024; Hogan et al., 2024).

Looking at the sector dummy coefficients, we find some heterogeneity in the impacts over the probability of AI adoption across sectors, while differences in investment intentions are less pronounced. On the whole, this may indicate less sectoral divergence in firms' forward-looking attitudes toward AI relative to outstanding adoption status. With respect to AI adoption, our results largely confirm the findings from the descriptive analysis. That is, firms operating in Administrative and Support Services, Professional, Scientific and Technical Activities, High-Technology Manufacturing, and Financial and Insurance Activities exhibit a higher propensity to adopt AI. Regarding the effect on investment intentions, there is no discernable pattern other than for Administrative and Support Services, and a hint that the prospects for investment are gloomy in Wholesale, and in the Financial and Insurance sectors, relative to the baseline.

We also examine geographical differences. Using NUTS-2 regional classifications, we find that SMEs located in the South of Ireland present higher levels of AI deployment, while firms in the Eastern and Midlands region display a more forward-looking investment profile, characterised by stronger intentions to invest in AI, relative to the Eastern and Midland region — also in alignment with the descriptive analysis.

Finally, we look at whether firms seeking external financing, which can be indicative of enterprises with greater financial dynamism, is associated with more positive attitudes toward technological innovation. We find that only the variable representing those seeking bank-financing, the more traditional of the external financing options, is statistically significant at 10 percent. The coefficients of the variables capturing those looking for non-bank financing opportunities or expected to look for future bank financing are insignificant. Thus, this provides modest support to the argument that firms that may be more financially sophisticated are also more inclined to invest in AI technologies.

Table 3: Probit Models

			Recursive Probit	
	(1)	(2)	(3)	(4)
	Adopted AI	Intend to invest in AI	Adopted AI	Intend to invest in AI
Operational Barriers				
Staff skills	0.08	0.36***	0.08	0.36***
	(0.81)	(3.80)	(0.74)	(3.77)
Procuring cost	0.22+	0.44***	0.24*	0.48***
	(1.90)	(3.87)	(2.03)	(4.17)
Staff training cost	-0.12	-0.16	-0.11	-0.17
	(-0.93)	(-1.27)	(-0.82)	(-1.39)
Security concerns	0.02	-0.32**	-0.00	-0.31**
	(0.16)	(-3.06)	(-0.03)	(-2.93)
Access to finance	-0.03	-0.28*	-0.05	-0.28*
	(-0.21)	(-2.22)	(-0.38)	(-2.21)
Regulation	-0.14	0.10	-0.11	0.08
	(-1.26)	(0.98)	(-1.01)	(0.79)
Knowledge-related Barriers				
Do not Understand AI for company	-0.29+	-0.19	-0.24	-0.23
	(-1.93)	(-1.32)	(-1.60)	(-1.55)
Do not Understand benefits	-0.46**	-0.80***	-0.46**	-0.84***
	(-3.27)	(-5.51)	(-3.18)	(-5.81)
Other in sector Do Not use AI	-0.71***	-0.27*	-0.72***	-0.36*
	(-5.92)	(-2.52)	(-5.73)	(-2.57)
Information Not easy to find	0.01	-0.05	-0.02	-0.05
	(0.12)	(-0.49)	(-0.22)	(-0.51)
Firm Characteristics:				
High-productivity dummy	0.00	0.16+	-0.00	0.15+
	(0.02)	(1.79)	(-0.04)	(1.77)
Employment growth > 0 dummy	0.01	0.25**	0.04	0.24**
	(0.13)	(2.72)	(0.45)	(2.68)
Exporter	0.27*	0.27*	0.29**	0.32**
	(2.52)	(2.56)	(2.58)	(2.85)
Firm Size (baseline Micro)				
Small	0.05	0.02	0.04	0.02
	(0.50)	(0.14)	(0.39)	(0.23)
Medium	0.15	0.45***	0.13	0.46***
	(1.21)	(3.67)	(1.00)	(3.88)
Firm age: (baseline <10 years)				
10-19 years	-0.20	-0.00	-0.21	-0.04
	(-1.01)	(-0.01)	(-1.02)	(-0.23)
20+ years	-0.28	-0.17	-0.26	-0.22
	(-1.57)	(-0.93)	(-1.41)	(-1.19)
Sector: (baseline Manufacturing - other)				

Construction & real estate	0.13	-0.13	0.23	-0.10
	(0.61)	(-0.68)	(0.97)	(-0.50)
Wholesale	0.25	-0.34*	0.28	-0.29+
	(1.42)	(-2.14)	(1.42)	(-1.73)
Hotels & restaurants	0.44*	-0.22	0.45*	-0.15
	(2.08)	(-1.13)	(1.97)	(-0.71)
Prof., scientific & tech.	0.85***	0.23	0.92***	0.39+
	(4.64)	(1.38)	(4.52)	(1.73)
Administrative & support	0.82***	0.35+	0.91***	0.50*
	(3.75)	(1.74)	(3.86)	(2.07)
Manufacturing - high tech	0.86**	0.24	0.90**	0.37
	(2.87)	(0.81)	(2.92)	(1.18)
Financial & insurance	0.97***	-0.66*	1.05***	-0.43
	(3.75)	(-2.45)	(3.86)	(-1.20)
Other	0.15	-0.15	0.23	-0.12
	(0.70)	(-0.82)	(1.01)	(-0.62)
NUTS-2 Regions: (baseline: Northern & Western)				
Southern	0.26+	0.03	0.30*	0.08
	(1.93)	(0.22)	(2.13)	(0.57)
Eastern & Midland	0.16	0.29*	0.17	0.31*
	(1.23)	(2.32)	(1.23)	(2.49)
External Financing Requested				
Bank financing	0.20+		0.18	
	(1.93)		(1.45)	
Non-bank financing	0.07		0.04	
	(0.52)		(0.29)	
Future bank financing		0.04		0.04
		(0.33)		(0.40)
Adopted AI		0.77***		0.11
		(7.49)		(0.17)
Constant	-1.25***	-1.04***	-1.34***	-0.94***
	(-4.76)	(-4.19)	(-4.80)	(-3.46)
ρ			0.38	
			(0.94)	
N	1,503	1,428	1,428	
ll	-574.71	-599.73	-1142.86	
AIC	1,209.43	1,259.47	2,407.72	
BIC	1,368.89	1,417.39	2,728.83	

The Wald test of $\rho = 0$ gives $\chi^2(1) = .25$ with Prob > $\chi^2 = 0.61$. t statistics in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

5.2 Disentangling barriers across key firm characteristics

In this section, we first present predicted probabilities of AI adoption and investment intentions across key firm characteristics, based on the probit estimates reported in

Columns (1) and (2) of Table 3. We then examine the marginal effects of operational and informational barriers, first at the overall sample level and subsequently by disaggregating the results by sector and firm size.

5.2.1 Predicted probabilities

We begin by examining the predicted probabilities of AI adoption and intentions to invest across four key firm characteristics: size, sector, exporting status, and age (Figure 7). Overall, predicted probabilities are consistently higher for investment intentions than for current AI adoption. This gap suggests that, across all firm characteristics, a substantial share of firms are considering the implementation of AI but have not yet adopted it. Such a pattern is consistent with AI being a technology that remains at a relatively early stage of diffusion.

Medium-sized enterprises exhibit a substantially higher probability of intending to invest in AI, around 30 percent, compared with micro and small enterprises. For this group, this figure is almost twice as large as the probability of having already adopted AI, which stands at approximately 18 percent (Figure 7a).

Across sectors, the predicted probability of future AI investment exceeds 30% in Administrative and Support Activities, Professional, Scientific and Technical Activities, and High-Tech Manufacturing. While the probability of current adoption is somewhat lower than this threshold, the overlap in confidence intervals between investment intentions and adoption indicates that the difference between the two outcomes is not statistically significant (Figure 7b). Although a high probability of AI adoption does not translate one-for-one into a high probability of intending to invest, this pattern suggests that investment prospects in these sectors are more likely to reflect a deepening or expansion of AI use among existing adopters rather than a substantial inflow of first-time users.

In contrast, firms in Other manufacturing sectors, Construction and Real Estate Activities, and Other Industries display a noticeably higher probability of future investment relative to their current levels of AI adoption. This gap indicates that a sizeable share of firms in these sectors are considering becoming first-time AI users.

A noteworthy exception emerges in Financial and Insurance activities, where SMEs exhibit the highest rate of AI adoption but comparatively weaker intentions to invest further in the future. This distinctive pattern was already highlighted in the descriptive statistics. That it also emerges in the predicted outcome should make it relevant for policy considerations. A deeper understanding of the constraints and needs of firms in this sector may be particularly relevant for the design of targeted policy interventions.

Turning to exporter status, the gap between the probabilities of investment and current adoption is largest among non-exporting firms (Figure 7c). Similarly, by firm age, the difference between investment intentions and adoption is most pronounced for older firms, whereas the gap is considerably smaller for younger firms (Figure 7d).

Figure 7: Predicted Probabilities by Key Firm Characteristics

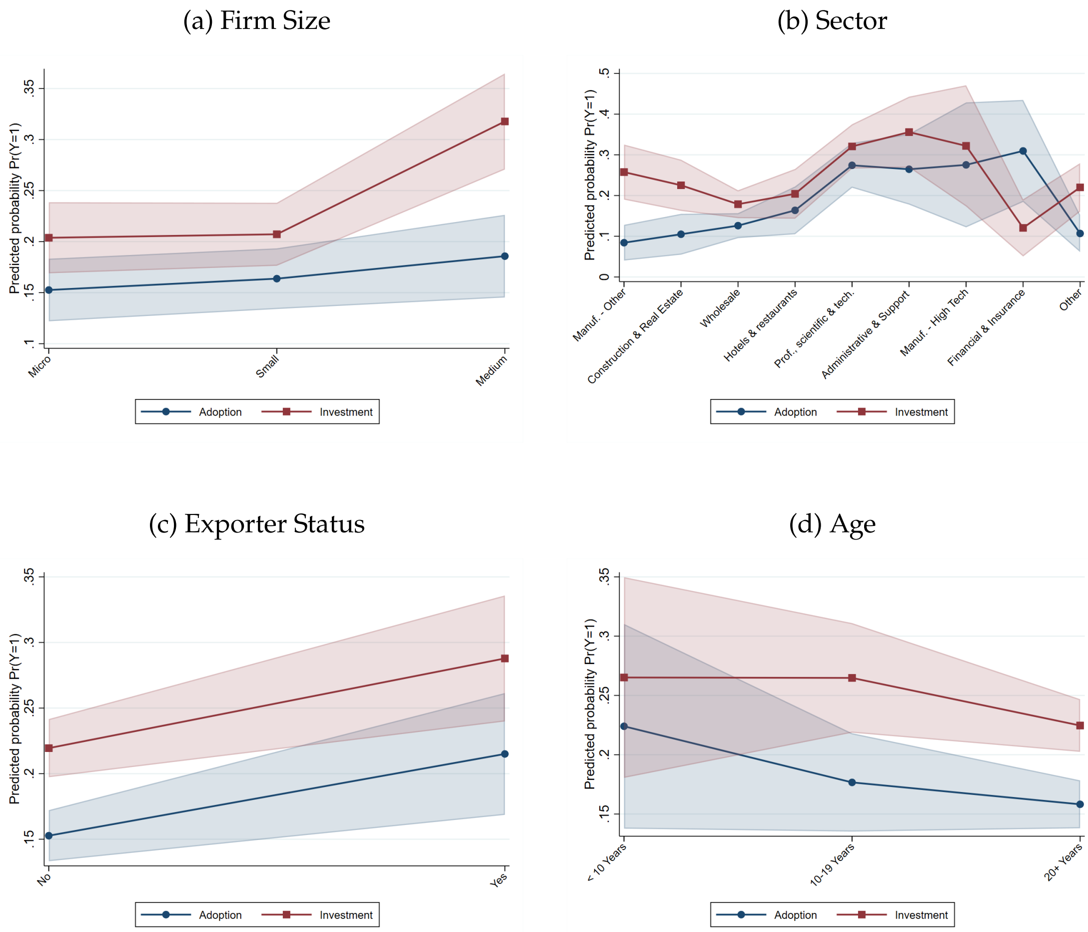

Note: Shaded areas represent 95% confidence intervals around the probabilities.

Taken together, these results highlight a clear distinction between firms' prevailing AI adoption behaviours and their forward-looking attitude towards AI diffusion, as captured by firms' broader investment intentions. While a clear pattern of sectoral laggards emerges, their investment intentions are more closely aligned with those of the first-mover sectors.

5.2.2 Average Marginal Effects

Having examined the predicted probabilities, we now turn to the magnitude of the marginal effects associated with barriers to AI assimilation. We first focus on the average marginal effects of knowledge-related and operational barriers. We then explore heterogeneity in these effects by sector and firm size in order to disentangle differential impacts across firm types.

At a high level, the results indicate that knowledge-related barriers are important for explaining both the probability of current AI adoption and firms' intentions to invest, while operational barriers are more strongly associated with investment intentions. Firms reporting that they do not understand the overall benefits of AI have, on average, a 10 percentage-point lower probability of having already adopted it and a 20 percentage-point lower probability in intentions-to-invest (Figure 8). Similarly, firms that report not believing that AI is used by others in their sector exhibit a 15 percentage-point lower probability of AI adoption and around a 10 percentage-point lower probability of intending to invest.

Turning to operational barriers, and consistent with the main results discussed above, the marginal effects of most operational constraints — aside from procurement costs — are not significantly associated with the probability of current AI adoption, on average. As shown in Figure 9a, the estimates are small in magnitude, ranging from approximately –5 to +5 percentage points, and present relatively wide standard errors.

In contrast, operational barriers play a more substantial role in shaping intentions to invest in AI (Figure 9b). For future investment, estimated marginal effects range from around –10 percentage points for security concerns and access to finance to approximately +10 percentage points for procurement costs. Notably, firms that report procurement costs or staff skills as concerns are, on average, 10 percentage points more likely to plan AI investments. Rather than discouraging investment, these constraints may signal that firms are already engaged in the planning phase of AI adoption and are identifying concrete implementation challenges. By contrast, security concerns and access to finance emerge as the two most important operational barriers hindering AI investment intentions. Holding other factors constant, addressing security concerns or alleviating financial constraints could increase a firm's probability of investing in AI by around 10 percentage points.

5.2.3 Heterogeneous effects of AI barriers

Next, we break down the marginal effects by sector and firm size, effectively examining interaction effects between these dimensions. The results are exhibited in Figure 10 and 11. Overall, we find that there is limited heterogeneity in the effects of barriers across sectors and firm size. On average, the estimated marginal effects move in the same direction, with statistical significance and magnitudes varying only modestly.

A modest degree of heterogeneity is observed for knowledge-related barriers, particularly across sectors (Figure 10a). For instance, the marginal effect of AI adoption associated with the indicator capturing lack of AI use among peers ("Other in the sector") ranges from around –20 percentage points in Financial and Insurance Activities to approximately –10 percentage points in Other Manufacturing sectors (lhs panel). Similarly, the marginal effect of lacking understanding of the overall benefits of AI on investment intentions (rhs panel) is largest in Administrative and Support Activities (around –25 percentage points) and smallest in Financial and Insurance Activities (approximately –12 percentage points). By contrast, these differences are considerably less pronounced when the results are disaggregated by firm size (Figure 10b). In particular, we find

Figure 8: Marginal Effect of Knowledge Barriers to AI

(a) Probability of AI Adoption (b) Probability of Future AI Investments

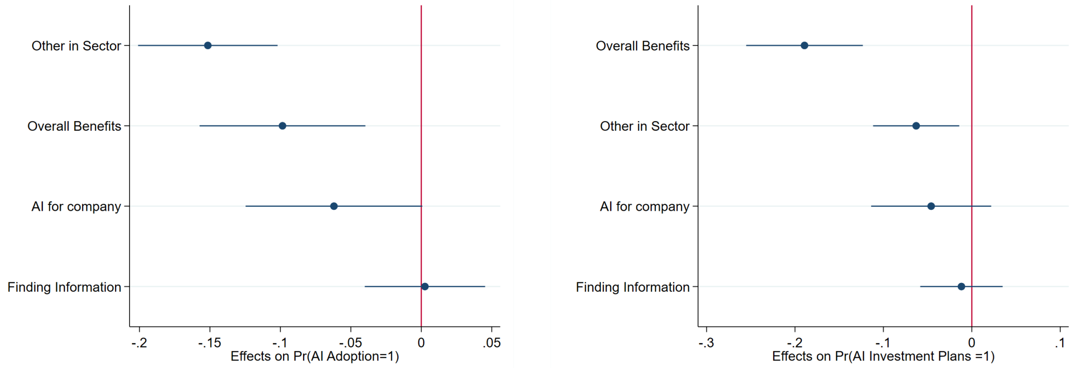

Figure 9: Marginal Effects of Operational Barriers to AI

(a) Probability of AI Adoption (b) Probability of Future AI Investments

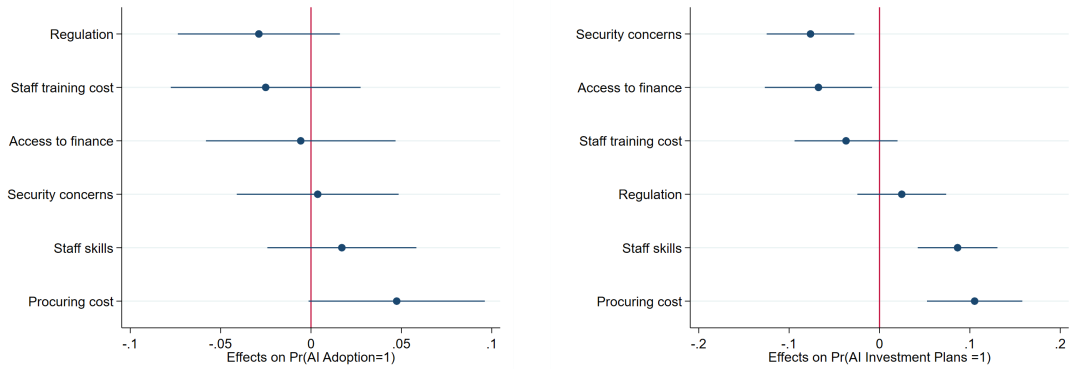

Notes: Points represent estimated marginal effects; horizontal lines denote 95% confidence intervals. The vertical line at zero indicates no effect.

that the effects on micro and small enterprises are virtually indistinguishable, while the marginal effects on medium-size firms are slightly larger in absolute terms.

Turning to operational barriers, heterogeneity across sectors and firm size is more limited, particularly across the marginal effects on AI adoption reflecting the lesser relevance in determining this outcome, as already discussed (Figure 11). The largest variation arises in the marginal effects on investment intentions across sectors. In particular, the sensitivity of investment intentions to staff skills and procurement costs is strongest in Administrative and Support Activities and High-Tech Manufacturing, and weakest in Financial and Insurance Activities. A similar pattern to that observed for knowledge-related barriers emerges across firm size, whereby medium-sized enterprises exhibit the largest marginal responses to operational barriers, especially with respect to staff skills and procurement costs, while micro and small enterprises display nearly identical effects.

Figure 10: Heterogeneity in Marginal Effects of Knowledge-related Barriers by Size and Sector

(a) Sector

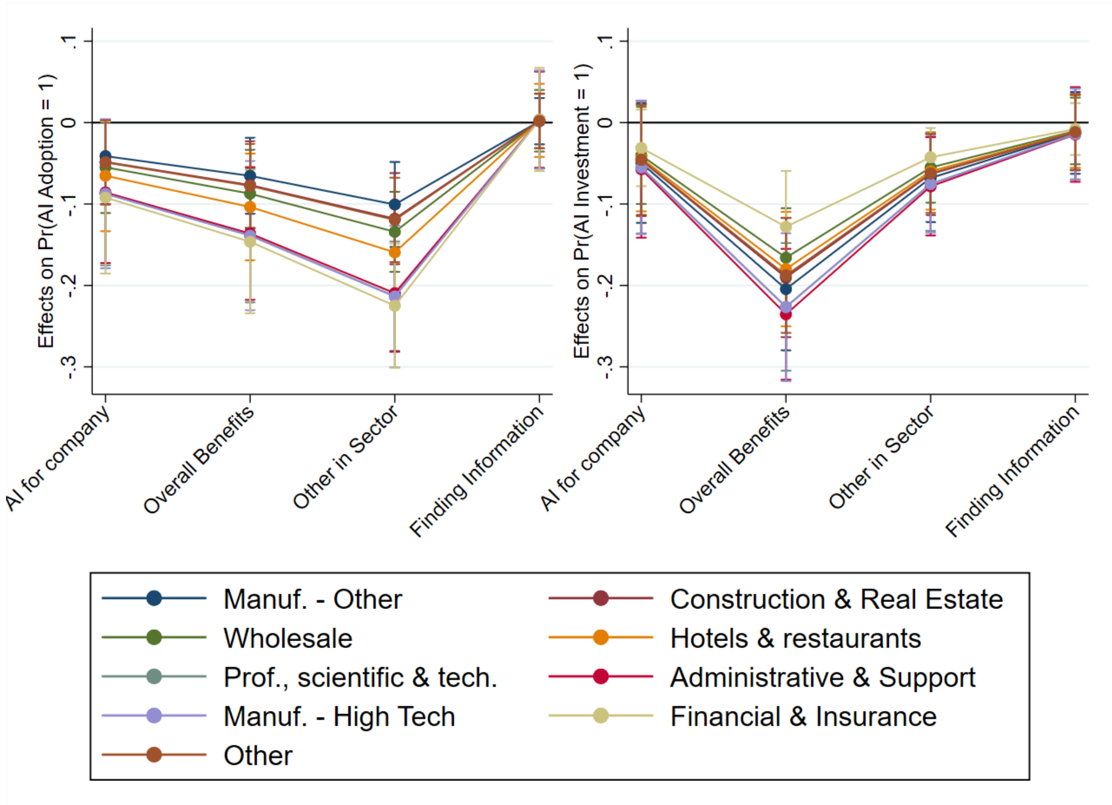

(b) Firm Size

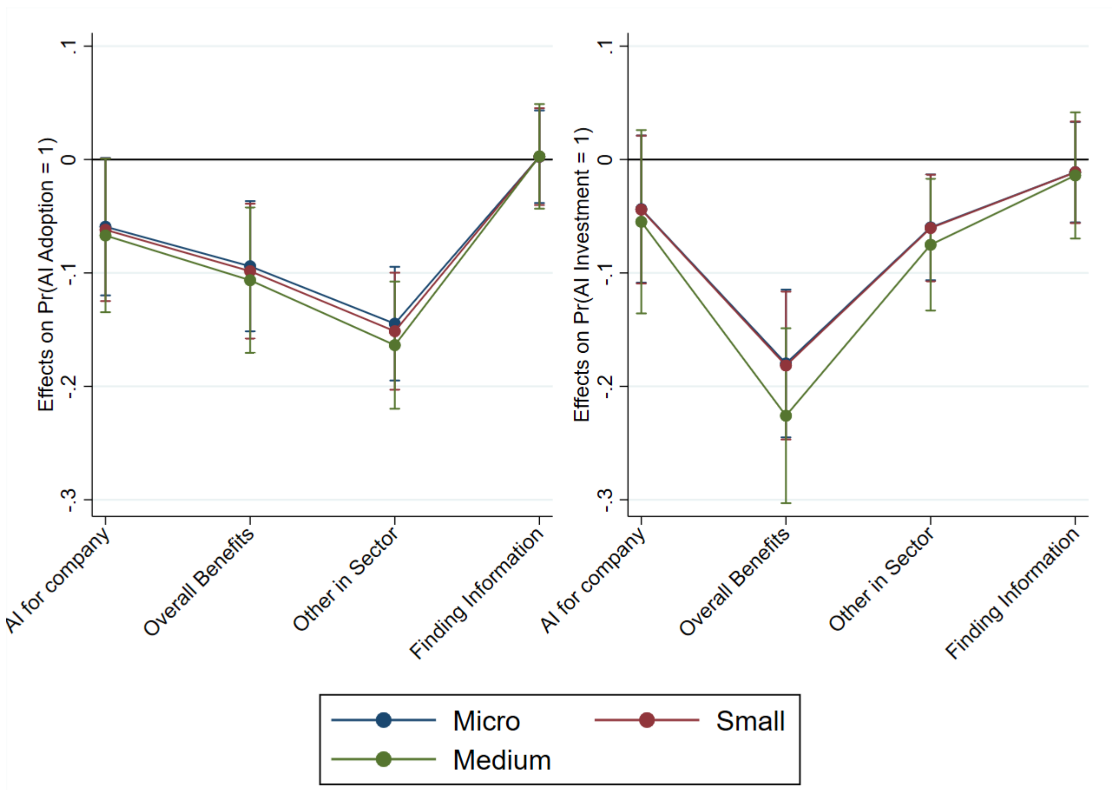

Notes: Points represent estimated marginal effects; vertical lines denote 95% confidence intervals. The vertical line at zero indicates no effect.

Figure 11: Heterogeneity in Marginal Effects of Operational Barriers by Size and Sector

(a) Sectors

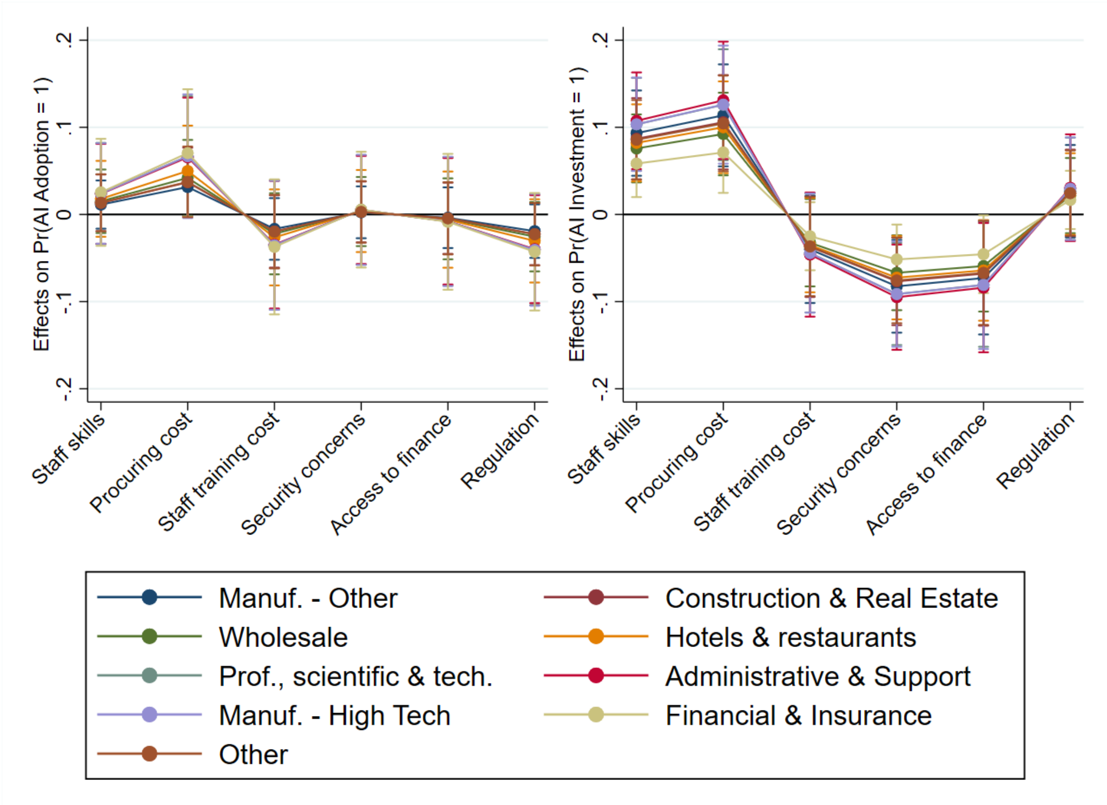

(b) Firm size

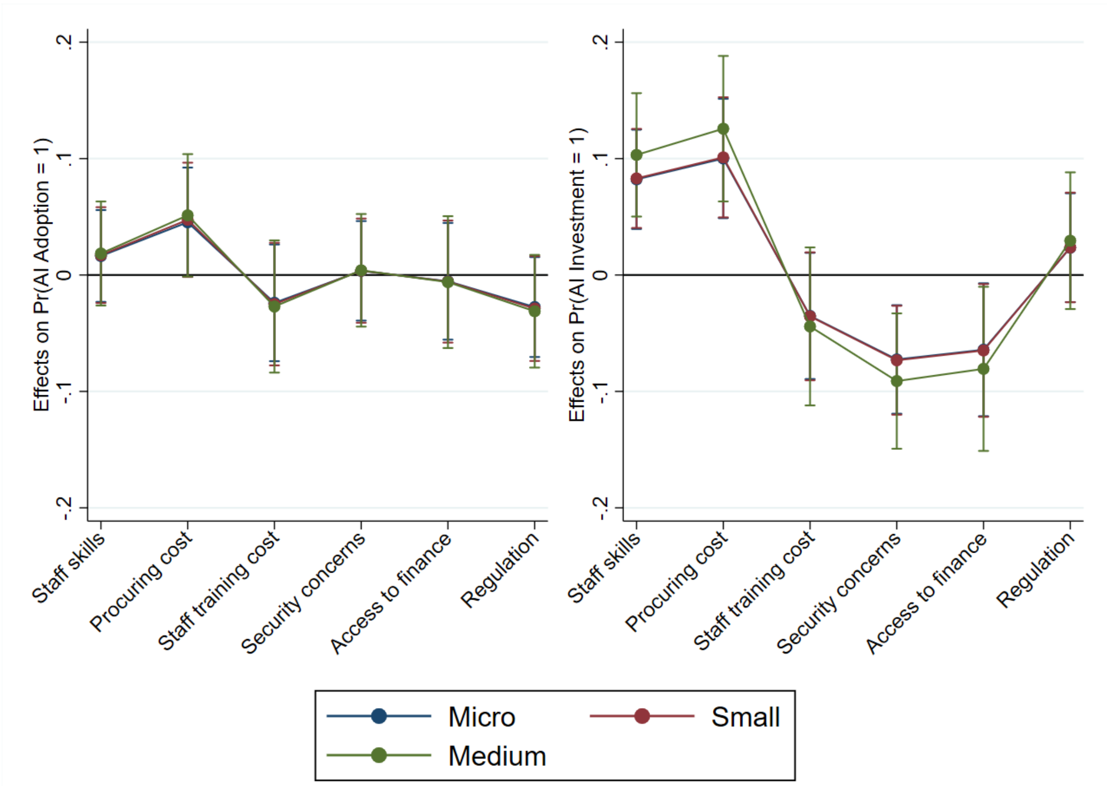

Notes: Points represent estimated marginal effects; vertical lines denote 95% confidence intervals.

5.3 Robustness checks using Machine Learning

The purpose of this section is twofold. First, we use a set of machine learning algorithms to cross-validate the findings from our probit estimations, ensuring that our main results are not model-dependent. Second, this exercise illustrates the value of incorporating machine-learning tools into economic analysis, showing how these methods can complement traditional econometric approaches by uncovering nonlinearities, interaction patterns, and robustness in predictive performance.

5.3.1 Random Forest with SHAP analysis

To complement the probit estimations and assess the robustness of our findings, we employ a supervised machine-learning framework based on Random Forest classification models. Random Forests are ensemble methods that construct multiple decision trees on bootstrapped samples of the data and aggregate their predictions. By aggregating the decision trees, Random Forests reduce overfitting and capturr complex nonlinearities and interaction patterns that standard parametric models may miss. This makes them particularly suited for high-dimensional firm-level data where relationships between covariates and outcomes may not be strictly linear.

Given the sensitivity of tree-based models to their hyperparameters — such as the number of trees, maximum tree depth, or the subsample fraction used for splitting — we optimise model performance using random search. Random search explores the hyperparameter space by sampling parameter combinations at random, which has been shown to be substantially more efficient than exhaustive grid search when only a subset of hyperparameters has a strong influence on model performance. This approach enables robust tuning while ensuring computational tractability.

To interpret model predictions and connect them back to the economic mechanisms driving AI adoption and investment intentions, we rely on SHAP (SHapley Additive exPlanations) values. SHAP provides a theoretically grounded decomposition of each prediction into additive feature (i.e. variable in the economic jargon) contributions, helping to assess both the global importance of covariates and their marginal effect patterns.[2] Unlike raw feature importances, SHAP values allow us to visualise non-linear effects, heterogeneous responses across firms and feature interactions in a transparent and model-agnostic way (Lundberg & Lee, 2017; Strumbelj & Kononenko, 2010; Shapley, 1953).

Overall, the combination of Random Forests, random search hyperparameter tuning, and SHAP-based interpretation provides a powerful robustness check for our econometric results. It allows us to validate whether the key drivers identified in the probit

[2] SHAP analysis is rooted in Shapley values, a concept from collaborative game theory. Shapley values provide a fair distribution of a payout among players in a collaborative game where players work together for a common goal, even if the players may have contributed unequally. These values can be interpreted as the weighted average of a player's marginal contributions across all possible coalitions (Ponce-Bobadilla et al., 2024).

analysis remain influential in a flexible non-parametric setting, while also uncovering additional structure in the data that traditional models may not detect.

5.3.2 Results

We show the so-called "SHAP beeswarm" plot in Figure 12 . This plot offers a global summary of how each variable (i.e. "feature") contributes to the model's predictions across all observations. Features are ranked vertically by their mean absolute SHAP value, placing the most influential variables at the top. Each point represents one observation's SHAP value for a given feature, with its horizontal position indicating the magnitude and direction of the feature's contribution to the predicted probability. Point colour encodes the underlying feature value, enabling the visualization of how different levels of a feature relate to its SHAP contribution. Given that most variables in our analysis are binary, red points correspond to observations taking the value 1, while blue points correspond to observations taking the value 0. This makes it straightforward to detect whether the presence (1) or absence (0) of a characteristic systematically increases or decreases the predicted probability of AI adoption. The joint use of colour and position enables the identification of global patterns (such as monotonic or non-linear relationships), differences in the spread of contributions across features, clusters of similar SHAP distributions, and potential outliers.

Figures 12a and 12b display the SHAP beeswarm plots for the probability of AI adoption and investment intentions, respectively. Across both outcomes, the SHAP patterns point to a generalised and dominant role of knowledge-related barriers, as in our main results. These barriers consistently emerge as the most influential predictors, particularly in the case of AI adoption. The visual clustering of red points on the negative side of the SHAP axis shows a clear divide: firms reporting difficulties in understanding what AI can do for them — or how it could be integrated into their processes —contribute strongly toward lowering the predicted likelihood of adoption; conversely, firms not reporting these barriers (blue points) tend to have positive SHAP contributions, signalling a higher probability of adopting AI.

The beeswarm plot for investment intentions further reveals that prior AI adoption constitutes the second-most influential predictor, corroborating the econometric results that early engagement with AI is a critical catalyst for additional investment activity.

The SHAP analysis further shows that firms experiencing positive employment growth are more likely to intend to invest than those with stagnant employment. Firm characteristics also matter: as in the main results, firms in Professional, Scientific and Technical activities sector and exporters are strong predictors of both adoption and investment intentions, while those in the Wholesale sector — and in Construction or Real Estate for adoption — display consistently negative SHAP values, indicating that they are more likely to lag behind. The SHAP values further reveal notable regional disparities, with SMEs in the North and West exhibiting lower predicted probabilities of adoption, whereas those in the East and Midlands show higher predicted probabilities

Figure 12: Variable impacts on output (SHAP values). Top ten features

(a) AI Adoption

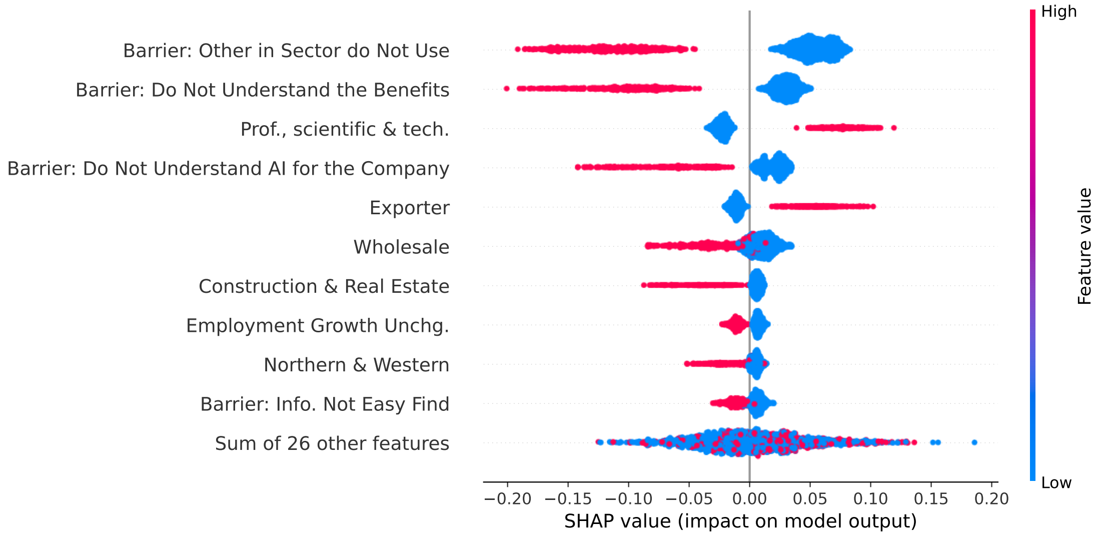

(b) Investment Intentions

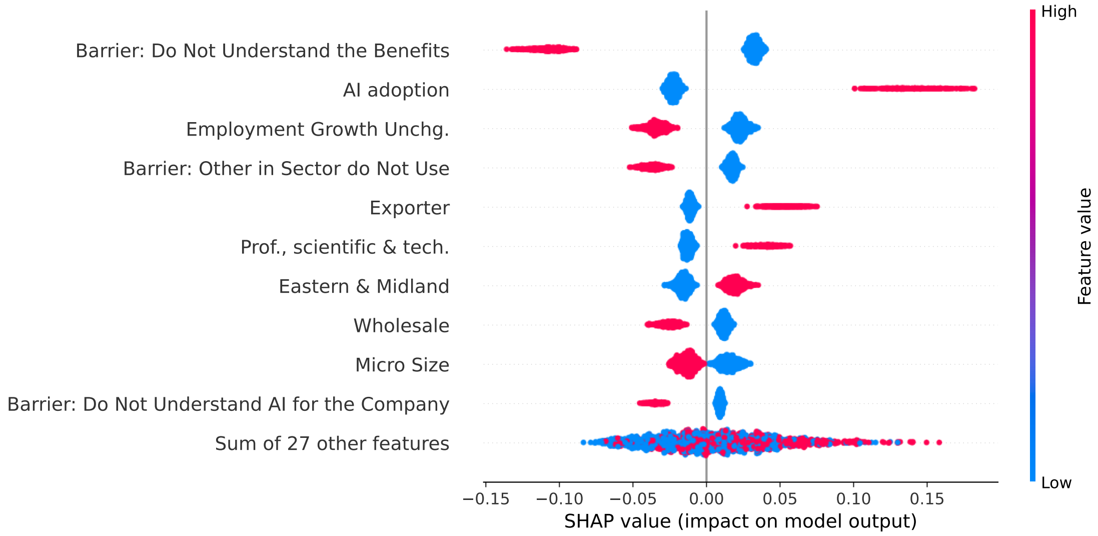

of intending to invest, mirroring the geographic imbalances identified in the descriptive analysis.

6 Conclusions and implications for policy

6.1 Discussion of main findings

Throughout the empirical analysis, we have sought to provide a deeper understanding of the factors associated with firms' AI adoption process. Leveraging unique survey data, we examine key contributors from two complementary perspectives: (i) prevailing AI adoption decisions and (ii) firms' forward-looking attitudes toward AI deployment. Although the adoption decision-making process is closely intertwined with future investment intentions, analysing how various factors relate to each of these dimensions offers valuable insights into the broader dynamics of AI diffusion, which in turn provide a useful baseline for informing policy interventions.

Among the factors examined, we focus on operational and knowledge-related barriers as our primary variables of interest, alongside a broad set of firm-level attributes that enable us to assess heterogeneity in adoption patterns across different types of firms.

At a high level, the results reveal a clear distinction: ultimate AI adoption decisions appear to be primarily influenced by informational barriers, whereas investment intentions are explained by both knowledge-related concerns and operational barriers. When these barriers are broken down by firm characteristics, their effects generally point in the same direction, although some heterogeneity emerges across sectors. Moreover, we find that there are more firms willing to invest in AI that have already adopted AI. This has important policy implications, as it indicates a strong appetite for AI adoption and suggests that a common framework for intervention can be effective, provided it is tailored to address sector-specific needs and applications.

Given the early stage of AI technologies at the time these enterprises were surveyed, the discrepancies between firms' adoption choice and attitudes towards future implementation are understandable and economically intuitive. For example, we find that perceptions that peers within the same sector are not using AI strongly influences decisions not to adopt AI. Similarly, a lack of understanding about the benefits of AI or its relevance to a firm's business emerges as other key factors behind non-adoption. These findings underscore the importance of targeted information campaigns that clarify sector-specific applications of AI and the ways in which these technologies can benefit a firm. Such initiatives could be highly influential, particularly at encouraging a significant number of SMEs to begin using widely available AI tools.

Regarding investment intentions, we find that the factors influencing spending plans closely align with considerations over standard inputs in the capital formation processes. First and foremost, not understanding the benefits of AI plays a significant role in hindering the probability of investments, which can be associated with the uncertainty related to the profitability of the investment. From an operational standpoint, key factors include staff skill requirements, procurement costs, security concerns, and access to finance. Such determinants reflect a more structured approach to AI adoption, consistent with firms planning formal investment strategies. For these strategies to materialise, firms require a deep and technical understanding of available AI tools and how they can be integrated into organizational workflows and production systems.

This latter consideration can play a critical role in determining the long-term success of AI investments. Existing research on the particularities of AI suggests that companies that do not benefit from adopted AI technologies do so because they become locked into static tools that can't adapt to their evolving workflows. Meanwhile, business needs and the trajectory of the next-generation AI technologies (i.e. agentic AI) are shifting toward systems capable of continuous learning. Thus, advanced AI national strategies should place emphasis on investments in technologies that can integrate into workflows and deliver sustained value (Challapally et al., 2025). To support this transition, policy initiatives offering guidance and education from an early-stage may be instrumental in enabling the deployment of fully integrated adaptative systems as they become available.

6.2 Implications for policy

Given the evidence in this research on the different factors that have affected current levels of AI adoption and future intentions to invest, there are a number of implications for the calibration and operationalisation of any policy framework for AI. Our results suggest the following two priority areas, which could be incorporated into any policy platform, in both specific terms but also in terms of the overall operating conditions for investment and absorption of AI-specific technologies.

First, our research clearly points towards a need to bridge knowledge gaps, enhance learning about the technology and build AI experiences across firms. In line with current policy in the area, focusing on the speed of diffusion across firms within sectors to help reduce knowledge deficiencies and improve information flows is a key element. In this regard, the priority should be to close informational gaps and raise firms' awareness of AI opportunities, applications, and benefits. Designing sector-specific information campaigns that showcase positive experiences and concrete use cases can serve as an effective and straightforward pathway to initial adoption through readily available AI tools. Allowing employees to become familiar with standard AI tools and understand their capabilities and limitations is crucial for later transitioning to workflow-integrated systems. Developing AI proficiency within the company can create a competitive advantage in workflow optimization. As workforce transformation is likely to occur gradually rather than through abrupt displacement events, the opportunity to boost AI literacy across all staff should also alleviate distributional constraints.

Secondly, enhancements and developments of the broader operating environment can create the conditions for firms to safely invest in AI technologies when the opportunity arises. Key pillars of any regulatory framework should cover issues around security and regulatory compliance, procurement costs and access to finance, and staffing needs, encompassing workforce skills and training costs. Targeted instruments and legislation which address these aforementioned enabling conditions should provide a stable environment for enterprises to increase investment in digitalisation and in doing so raise productivity in the broader SME sector.

As a final consideration, we would like to stress how crucial it is to not just encourage business to adopt AI, but to prepare them for the evolving nature of AI technologies. Policy programmes should place great emphasis on educating firms about the importance of flexibility and agility when investing in AI, as sustained value will increasingly be generated by successive waves of AI developments, including agentic AI systems, which are expected to be the truly transformative ones. Strategic policy frameworks designed today must not only address current capabilities but also anticipate future technological shifts, making sure that firms develop the capacity to adapt over time.

References

Acemoglu, D. & Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2), 3–30.

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.

Alderucci, D., Branstetter, L. G., Hovy, E. H., Runge, A., Ryskina, M., & Zolas, N. (2019). Quantifying the impact of AI on productivity and labor demand: Evidence from U.S. Census Microdata 1.

Baily, M., Brynjolfsson, E., & Korinek, A. (2023). *Machines of mind: The case for an AI-powered productivity boom*. Brookings.

Bassetti, T., Galvez, Y. B., Pavesi, F., & Del Sorbo, M. (2020). *Artificial intelligence: impact on total factor productivity, E-commerce & fintech*. Publications office of the European Union Luxembourg.

Benassi, M., Grinza, E., Rentocchini, F., & Rondi, L. (2022). Patenting in 4ir technologies and firm performance. *Industrial and Corporate Change*, 31(1), 112–136.

Bessen, J. & Righi, C. (2019). *Shocking technology: what happens when firms make large IT investments?* Boston University School of Law.

Brynjolfsson, E. & Unger, G. (2023). *The macroeconomics of artificial intelligence*. International Monetary Fund.

Cazzaniga, M., Jaumotte, M. F., Li, L., Melina, M. G., Panton, A. J., Pizzinelli, C., Rockall, E. J., & Tavares, M. M. M. (2024). *Gen-AI: Artificial intelligence and the future of work*. International Monetary Fund.

Challapally, A., Pease, C., Raskar, R., & Chari, P. (2025). The GenAI Divide: State of AI in business 2025. *MIT NANDA*.

Czarnitzki, D., Fernández, G. P., & Rammer, C. (2023). Artificial intelligence and firm-level productivity. *Journal of Economic Behavior & Organization*, 211, 188–205.

Damioli, G., Van Roy, V., & Vertesy, D. (2021). The impact of artificial intelligence on labor productivity. *Eurasian Business Review*, 11(1), 1–25.

Department of Enterprise, Tourism and Employment (2026). Digital Ireland – connecting our people, securing our future.

Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2023). Gpts are gpts: An early look at the labor market impact potential of large language models. *arXiv preprint arXiv:2303.10130*, 10.

European Commission (2025). State of the digital decade 2025.

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.

Felten, E. W., Raj, M., & Seamans, R. (2023). Occupational heterogeneity in exposure to generative AI. *SSRN Working Paper 4414065*.

Government of Ireland (2025). Action plan on competitiveness and productivity.

Hogan, K., Kren, J., & O'Toole, C. (2024). Irish enterprise digitalisation: A cross-country exploration. *ESRI, QEC Research Note RN20240201*.

Korinek, A. & Suh, D. (2024). *Scenarios for the Transition to AGI*. Technical report, National Bureau of Economic Research.

Kren, J. & O'Toole, C. (2024). What drives sme investment in digitalisation? Micro-data evidence for Ireland. *ESRI Working Paper 777*.

Lundberg, S. M. & Lee, S.-I. (2017). A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems*, volume 30: Curran Associates, Inc.

Maddala, G. (1983). *Limited-Dependent and Qualitative Variables in Econometrics*. Cambridge University Press.

Pizzinelli, C., Panton, A. J., Tavares, M. M. M., Cazzaniga, M., & Li, L. (2023). *Labor market exposure to AI: Cross-country differences and distributional implications*. International Monetary Fund.

Ponce-Bobadilla, A. V., Schmitt, V., Maier, C. S., Mensing, S., & Stodtmann, S. (2024). Practical guide to SHAP analysis: Explaining supervised machine learning model predictions in drug development. *Clinical and Translational Science*, 17(11), e70056.

Qian, C., Zhu, C., Huang, D.-H., & Zhang, S. (2023). Examining the influence mechanism of artificial intelligence development on labor income share through numerical simulations. *Technological Forecasting and Social Change*, 188, 122315.

Shapley, L. S. (1953). A value for n-person games. In H. W. Kuhn & A. W. Tucker (Eds.), *Contributions to the Theory of Games II* (pp. 307–317). Princeton: Princeton University Press.

Strumbelj, E. & Kononenko, I. (2010). An efficient explanation of individual classifications using game theory. *The Journal of Machine Learning Research*, 11, 1–18.

Trammell, P. & Korinek, A. (2023). *Economic growth under transformative AI*. Technical report, National Bureau of Economic Research.

Williamson, H., Coates, D., Daly, K., FitzGerald, K., & Gannon, N. (2025). Occupational exposures, complementarity and the potential consequences of ai for the labour market: Some evidence from Ireland. *Journal for Labour Market Research*, 59(1), 30.

Wooldridge, J. M. (2002). *Econometric Analysis of Cross Section & Panel Data*. MIT Press.