NBER WORKING PAPER SERIES

THE INNOVATION RACE:
EXPERIMENTAL EVIDENCE ON ADVANCED TECHNOLOGIES

Zoë B. Cullen
Ester Faia
Elisa Guglielminetti
Ricardo Perez-Truglia
Concetta Rondinelli

Working Paper 34532
http://www.nber.org/papers/w34532

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2025, Revised August 2026

This draft: August 12, 2026. The views expressed herein do not necessarily reflect those of the Bank of Italy or the National Bureau of Economic Research. We would like to thank Olivier Coibion, Diego Comin, Andrea Neri, Alessandro Notarpietro, Olivier Pfauti, Tiziano Ropele, Alfonso Rosolia, Fabrizio Venditti, Roberta Zizza, and Nathan Zorzi for their useful comments and suggestions, as well as seminar participants at the NBER Productivity and Innovation Group, Dartmouth University, the University of Texas at Austin, Northwestern University, Northeastern University, the Chicago Fed, the ESCoE Conference on Economic Measurement, the University of Manchester, and the University of Sheffield. This project was reviewed and approved in advance by the Institutional Review Board at Goethe University Frankfurt. It is also registered with the AEA Registry under AEARCTR-0016177. After the study is accepted for publication, we will share all code through a public repository, along with instructions on how to request the data. Yuerong Zhuang provided excellent research assistance.

The Innovation Race: Experimental Evidence on Advanced Technologies
Zoë B. Cullen, Ester Faia, Elisa Guglielminetti, Ricardo Perez-Truglia, and Concetta Rondinelli
NBER Working Paper No. 34532
December 2025, Revised August 2026
JEL No. C93, D22, L21, O33

ABSTRACT

We present a large-scale field experiment testing for strategic complementarities in firms' technology adoption. Our experiment was embedded in a Bank of Italy survey covering around 3,000 firms. We elicited firms' beliefs about competitors' adoption of two advanced technologies: artificial intelligence (AI) and robotics. We randomly provided half of the sample with a signal about adoption among comparable firms. Relative to this signal, most firms substantially underestimated competitors' adoption, and those provided with the information updated their expectations about competitors' future adoption. The information increased firms' intended future adoption of robotics: a 1 pp increase in the share of competitors expected to adopt advanced technologies causes a 0.732 pp increase in the firm's intention to adopt robotics by 2027. Our findings provide causal evidence on coordination in innovation and illustrate how information frictions shape technology diffusion.

Zoë B. Cullen
Harvard University
Harvard Business School
and NBER
zcullen@hbs.edu

Ester Faia
Goethe University Frankfurt
and CEPR
faia@wiwi.uni-frankfurt.de

Elisa Guglielminetti
Banca d'Italia
elisa.guglielminetti@bancaditalia.it

Ricardo Perez-Truglia
University of California, Los Angeles
and NBER
ricardotruglia@gmail.com

Concetta Rondinelli
Banca d'Italia
concetta.rondinelli@bancaditalia.it

An online appendix is available at http://www.nber.org/data-appendix/w34532
A randomized controlled trials registry entry is available at AEARCTR-0016177

1 Introduction

The classic insights of Cooper and John (1988) on strategic complementarities highlight how firms' incentives often depend on the actions of others. In practice, however, clean evidence on the extent, sign, and magnitude of these spillovers has been difficult to obtain. Strategic complementarities are ubiquitous, yet causal identification is notoriously elusive, since measures of competitive interaction are typically confounded by simultaneity and unobserved shocks. Laboratory experiments have demonstrated that coordination can arise in settings where payoffs hinge on beliefs alone (e.g., Cooper et al., 1990, 1992; Nagel, 1995). But whether such mechanisms operate in high-stakes, real-world firm decisions remains largely an open question.

Against this backdrop, we provide a large-scale field experiment testing for strategic complementarities in innovation among firms. We focus on the adoption of advanced technologies. More specifically, we study two automation technologies that have seen rapid growth in recent years: AI and robotics. Investments in advanced technologies are both costly and consequential, as they require skill upgrading and organizational change, yet they have the potential to raise productivity, expand sales, and increase firm market values (Babina et al., 2024). Indeed, some policymakers regard these advanced technologies as transformative and as key drivers of growth and industrial competitiveness.

Our empirical evidence is based on Italy—a well-suited context for this research question for at least two reasons. First, it provides rich data on past adoption and the opportunity to conduct an experiment with firms. We leverage a long-standing firm survey conducted by the Bank of Italy—the Survey of Industrial and Service Firms (INVIND, from the Italian acronym). This survey is ideal because it has tracked the diffusion of advanced technologies for years, providing the data needed to compute the information treatments. Additionally, we were able to embed an information experiment into this survey. Second, Italy is a developed economy with both a strong presence and considerable growth potential in advanced technologies. Indeed, it is one of Europe's largest industrial economies: its manufacturing base is extensive, ranking second in Europe after Germany, and the country is among the leading producers of industrial robots.

Our research design is inspired by models of strategic interaction under incomplete information, including Cooper and John (1988), Angeletos et al. (2007), Angeletos and Pavan (2004), and, more recently, Huo and Takayama (2025). Through the lens of those models, firms face uncertainty when choosing actions—in our application, how much to invest in a new technology. Firms do not observe the true productivity of the technology; instead, each firm receives a private signal. These models also allow for strategic considerations: a firm's optimal investment depends not only on the technology's productivity but also on the adoption choices of its competitors. As a result, firms care about competitors' adoption decisions for at least two reasons. First, there is a direct competitive motive—if competitors adopt, a firm wants to adopt as well to avoid falling behind its rivals. Second, there is an indirect learning motive—by observing competitors' adoption choices, a firm can infer information about the private signals those competitors have received regarding the technology's productivity.[1]

Our tailored survey experiment was embedded in the 2025 wave of the INVIND survey. These data provide a unique opportunity to measure the beliefs firms hold about their competitors' adoption decisions—and to identify their causal effects on firms' own adoption behavior. The experiment first elicited firms' beliefs about the share of competitors that had already adopted advanced technologies. We then provided half of the sample, randomly selected, with a signal about adoption among comparable firms—firms in the same sector and size class. Next, we elicited expectations about competitors' future adoption—as of 2027. We can measure whether information about competitors' current adoption affects expectations about their future adoption. Most importantly, we also measured firms' own intended adoption plans for the future, enabling us to identify the causal effect of expectations about competitors' adoption on firms' own adoption plans.

Our research design offers several advantages. First, the Bank of Italy's long history with INVIND ensures high-quality data and large samples, thereby delivering the statistical power necessary to rigorously test our hypotheses. Second, because the information treatment is randomized, we can cleanly identify the causal impact of beliefs about competitors'

[1] Appendix A provides a simple model, adapted from Angeletos and Pavan (2004), that formalizes these two channels in our specific context.

innovation, ruling out spurious correlations that plague observational studies. Third, we can resurvey the same firms a year later and link the INVIND data to administrative records—such as electronic invoices and balance-sheet data—as well as to other survey data.

We begin by summarizing current adoption patterns and recent trends. Italy has experienced solid growth in the use of advanced technologies. As of 2025, around 28% of Italian firms used AI in some form. Although somewhat below the European Union average, this adoption rate is broadly comparable to other European countries (Baumann et al., 2026). Robotics adoption was also significant, with 23% of Italian firms using robotic technologies in 2025. While both robotics and AI are automation technologies and their 2025 adoption rates are relatively similar, there are important differences between them. Robotics is a more established technology, with some firms having adopted it decades ago and using it extensively. In contrast, AI is a newer technology: most firms adopting AI have done so only recently, and many are still using it experimentally. Moreover, robotics remains more deeply entrenched in specific sectors such as manufacturing.

We document substantial gaps between beliefs about competitors' adoption rates and the information we provided. Relative to the signal we provided, a strong majority of firms underestimated competitors' adoption, although some firms did so by a much larger margin than others. On average, prior beliefs fell 24.5 pp below the signal. Firms updated their beliefs significantly in response to the information, indicating that respondents paid attention to it. Consistent with Bayesian learning, the direction and strength of belief updating vary with the direction and magnitude of prior misperceptions.

Moreover, we observe significant effects of the information on firms' stated plans to adopt robotics. On average, the information treatment increased the intention to adopt robotics by 6.5 pp. This average effect is not only statistically significant (p=0.039) but also significant in magnitude, corresponding to a 17.3% increase relative to the baseline. We further provide two key robustness checks. First, we show that the effects on expectations and adoption plans are concentrated among firms that most strongly underestimated competitors' adoption.[2] Second, we find no effects on pre-treatment outcomes, such as current adoption, elicited

[2] As explained in Section 3.5, due to concerns about potential contamination, we use predicted rather than raw prior beliefs in this and the subsequent analyses.

before the information-provision stage.

Using a two-stage least-squares (2SLS) model, we estimate the firms' reaction function: a 1 pp increase in the share of competitors expected to adopt advanced technologies causes a 0.732 pp increase in the firm's intention to adopt robotics by 2027. The steepness of the reaction function highlights the strength of strategic complementarities in technology adoption, suggesting that firms' decisions are highly interdependent.

In contrast to the effects on adoption plans for robotics, we do not find significant effects on AI adoption. This lack of effect holds both in the full sample and in the subsample of firms with the largest belief underestimates. This difference should be interpreted cautiously: due to power limitations, we cannot rule out modest effects on AI adoption. With that caveat in mind, we discuss some potential explanations for why the treatment effects may have been weaker for AI adoption. First, the baseline level of AI adoption is quite high, leaving less scope for any treatment to increase it further. Second, AI is newer and often used experimentally, which may lead firms to respond more cautiously to information about peers' adoption.

The treatment effects we document could operate through at least two mechanisms: competition—firms may want to avoid falling behind their rivals—and learning—firms may learn about their own potential productivity gains by observing competitors' adoption decisions. Because these mechanisms are closely intertwined, disentangling them is challenging. However, we provide suggestive evidence on their relative strength in two ways. First, we conduct a heterogeneity analysis. Intuitively, the competition channel should be more pronounced for firms that face stronger competitive pressures —in the extreme case, a perfect monopolist would have no rivals to fall behind, shutting the channel down completely. Using administrative data, we split the sample according to firms' market shares and we observe that the effects are concentrated among firms with below-median market shares. Although these results should be interpreted cautiously because of their limited precision, they suggest that the competition channel plays a stronger role. Second, in the 2026 follow-up survey, we asked firms to rate their agreement with two statements designed to capture the learning and competition channels. The responses are consistent with the heterogeneity analysis: both channels appear to be at play, although the competition channel appears to be relatively more important.

The 2026 follow-up survey also allows us to assess whether stated adoption plans translate into actual adoption. Firms' intended adoption in 2025 predicts their actual adoption one year later, supporting the validity of adoption intentions as an outcome. At the same time, adoption is slow: only a minority of the firms that planned to adopt by 2027 had done so by early 2026. Consistent with the baseline effect on intentions, the estimated treatment effect on actual adoption one year later is positive. However, the estimate is imprecise and statistically insignificant, as the slow pace of adoption and the smaller follow-up sample leave us underpowered to detect an effect of the expected magnitude.

We discuss some policy implications of our findings. In recent years, Italy—like other EU members—has implemented financial incentives to spur innovation, particularly in digitization, robotization, and AI.[3] Our evidence on the presence of significant information frictions suggests an additional policy lever that could complement financial incentives. Consistent with this view, the estimated effects of the information treatment are larger among firms that did not benefit from the Transition 4.0 programme, although the difference between groups is statistically insignificant. Governments could deploy information campaigns to better inform firms about the productivity of new technologies and the adoption behavior of their competitors. While further research is needed to determine the design and impact of such campaigns, their relatively low cost makes them a potentially cost-effective complement to financial-incentive policies. Additionally, the steep reaction function we find underscores that policies or shocks affecting some firms can generate large aggregate effects through strategic spillovers in adoption decisions.

Our study relates to and contributes to several strands of the literature. Most importantly, it connects to the literature on strategic complementarities. On the theoretical side, this literature builds on foundational work on coordination games and global games, dating back to Cooper and John (1988) and Bulow et al. (1985). A theoretical framework closer to our application is provided by Angeletos and Pavan (2004), who apply these games to investment and technology adoption. Subsequent research has made substantial advances, including extensions to multiplayer contexts, in which each agent takes the average reaction function of others as given; see Weintraub et al. (2008) and the concept of mean-field games in

[3]For a discussion on how public funding can spur innovation, see Azoulay et al. (2019).

Cardaliaguet et al. (2019).

While theoretical advances in this area continue to emerge, progress on the empirical front has proved more challenging. Causal identification is difficult: when adoption choices are correlated among competitors, it is hard to disentangle whether this reflects strategic complementarities or shared exposure to correlated shocks. The seminal paper by Bloom et al. (2013) tackles this challenge by exploiting changes in federal and state tax incentives for research and development. To document adoption spillovers in robotics, Bilgin et al. (2025) uses an identification strategy that leverages firm-level input-output VAT network data from Turkey. Lin (2023) studies strategic complementarities in the adoption of electronic medical records by U.S. hospitals, using an instrumental-variable strategy that leverages cross-market spillovers within hospital systems. We contribute to this literature with an identification strategy based on experimental variation, which requires minimal assumptions.[4]

Our study also relates to the broader literature on innovation and the adoption of new technologies. Most closely related is the contemporaneous work of Menkhoff (2025), who shows that providing German firms with information about AI productivity gains and industry adoption rates increases their beliefs about AI's productivity potential and their adoption of these technologies. Hill and Stein (2025) study how competition to publish first shapes scientists' research decisions. Abebe et al. (2025) randomize two interventions among Ethiopian firms to assess the role of competition in firms' investment in management training, finding that firms in this setting do not feel threatened by competition.[5] We contribute to this literature by providing evidence that firms' misperceptions about competitors' adoption can be a significant source of underinvestment.

More broadly, our study contributes to the literature using surveys and randomized information provision for causal inference in macroeconomics (Coibion and Gorodnichenko, 2026). For example, using a large survey of New Zealand firms, Coibion et al. (2018) show that managers are poorly informed about aggregate conditions and that their misperceptions shape their pricing, employment, and investment decisions.

[4]While examining a setting different from innovation, our approach is methodologically related to Coibion et al. (2021), who provide experimental evidence on the role of higher-order beliefs about inflation in firms' pricing decisions.

[5]Other related studies include Comin and Hobijn (2010), Atkin et al. (2017), Gupta et al. (2024), Bencivelli et al. (2025), Kalyani et al. (2025), and Dürnecker et al. (2026).

The rest of the paper is organized as follows. Section 2 describes the institutional setting and data. Section 3 presents details about the design and implementation of the experiment. Section 4 discusses technology adoption and misperceptions about competitors' adoption. Section 5 documents the effect of the information treatment on firms' beliefs. Section 6 presents the effects of the information on firms' own adoption plans. The last section concludes.

2 Institutional Context and Data

2.1 Italian Firms

We focus on the context of Italian firms. Italy is a developed economy with both a strong presence and considerable growth potential in advanced technologies. Italy is one of Europe's largest industrial economies—its manufacturing base is extensive, ranking second in Europe after Germany. In robotics, Italy's industrial sector is the third-largest user of robots in Europe—and the first if the automotive industry is excluded (Bank of Italy, 2024). Italy is also a major producer and exporter, leading Europe in the number of robot and automation suppliers.[6] AI adoption is also growing among Italian firms but remains below the European Union average in both manufacturing and services.[7] Finally, Italy is implementing a plan of subsidies for firms investing in advanced technology—the National Recovery and Resilience Plan, which is funded through the EU programme NextGeneration EU.

2.2 The INVIND Survey

The INVIND questionnaire includes a section that remains constant each year, covering general information about the firm and its structure, investment, employment, turnover, operating results, capacity utilization, and financing. It also contains a variable section that changes annually, focusing on different thematic areas. For this study, we collaborated to

[6]Italy has 655 companies, followed by France with 628 and Germany with 540 (HowToRobot, 2023).

[7]According to the EU survey on ICT usage and e-commerce in enterprises (2025 wave), the share of firms with at least 10 employees adopting AI technologies is 13.5% for the EU and 8.2% for Italy (European Commission, 2025).

design a dedicated module featuring an information-provision experiment. A key advantage for this experiment is that past waves of the INVIND survey included questions on the adoption of advanced technologies—allowing us to construct the information necessary for our intervention.[8]

Each survey wave recruits a large sample of about 4,000 firms. Since a light-touch information treatment like the one in our study is expected to have modest effects, we require this type of large sample to have sufficient statistical power to detect plausible treatment effects. The surveys are conducted annually between February and May. Anecdotally, responses are provided by or completed with assistance from managers responsible for the firm's planning and operations—typically an owner-manager or chief executive officer in small and medium-sized firms and a chief financial officer or similar executive in larger firms (De Marco et al., 2021). Unfortunately, the survey does not explicitly record the roles of the individuals involved in completing the questionnaire.[9] It does, however, include one question that can serve as a rough proxy: "Did a company manager, or someone working closely with one, assist you in answering any of these questions?" About 54% of firms answered yes.[10]

Some of the questionnaires are completed with the assistance (in person or over the phone) of officers from the Bank of Italy's branches, which likely improves the overall quality of responses. Once collected, the survey data undergo exhaustive quality checks to ensure consistency and reliability—see Appendix B for more details. The survey has a long track record dating back to 1972 and adheres to best practices in survey design and testing. Past validation studies show, for example, that INVIND responses align closely with national accounts data (Caprara et al., 2024) and balance-sheet records (D'Aurizio and Papadia, 2016).

Since 2002, the survey has been representative of the population of non-financial firms with at least 20 employees headquartered in Italy, selected through a stratified sampling method.[11] The survey data also include sampling weights to account for selection proba-

[8]These questions on adoption of advanced technologies have appeared in several waves since 2016.

[9]According to responses to a question at the end of the survey, 38% of firms reported that the questionnaire was filled out by one person, 29% by two people, 22% by three people, and the remainder by more than three people.

[10]One caveat is that if the manager is the only respondent, that person might answer negatively because the manager completed the questionnaire directly.

[11]More detailed information about the methodology of the survey can be found in Bank of Italy (2017).

bilities. Following common practice with INVIND data (Guiso and Parigi, 1999; Cingano et al., 2016; Bottone, 2025), all results presented in this paper are reweighted unless stated otherwise.

The INVIND survey has a panel component: over the past ten years, 83% of firms have participated in at least two consecutive survey waves. We therefore included a short follow-up module in the 2026 survey wave. This follow-up survey had a few different goals, which are described in more detail in Section 3.2 below.[12]

2.3 Balance-sheet Data

We merged the survey responses with balance-sheet data at the firm level.[13] The Company Accounts Data Service (CADS) is a proprietary database maintained by Cerved Group, an information provider and one of Europe's major credit rating agencies. The database contains detailed balance sheet and income statement information for nearly the full universe of Italian limited liability companies since 1993. Firms are legally required to file these financial statements with their local Chambers of Commerce, and Cerved compiles and harmonizes these filings to construct CADS, which is updated annually.

The most recent data available correspond to fiscal year 2024. We use CADS to construct pre-treatment measures of firms' market shares and industry concentration. Specifically, we define firm-level market shares as the ratio of a firm's net revenues to the total net revenues of all firms operating in the same sector.[14] Using these market shares, we compute the Herfindahl-Hirschman Index (HHI) for each sector, defined as $HHI_j = \sum_{i=1}^{N_j} m_{i,j}^2$, where N_j is the number of firms operating in sector j and $m_{i,j}$ is the market share of firm i in sector j.

[12] Additionally, in the future, we may be able to estimate treatment effects on additional outcomes using administrative datasets. For example, we could use data on domestic invoicing and customs records to measure robotics adoption by identifying whether a firm purchased robots from domestic or foreign suppliers. Unfortunately, because these administrative datasets become available only after a lag of one or more years, access to the data may take some time.

[13] In addition to the balance-sheet data, our intention was to merge the survey responses to other administrative sources (e.g., domestic and customs transactions) to estimate treatment effects on additional outcomes. However, these datasets may become available only after a several-year lag, so it may take years before we can use them for this purpose.

[14] Sector definitions can be constructed at different levels of granularity. In most of the analysis, we adopt the same 11-sector classification used to compute the signals for the experiment.

3 Research Design

3.1 Survey Design

The survey instrument for our tailored module in the 2025 wave is provided in Appendix F and described below.[15]

We focus on the adoption of two important types of automation technologies that have experienced rapid growth in recent years: AI and robotics. The 2025 module begins by defining predictive and generative AI (hereafter, Pred.-AI and Gen.-AI, respectively), as well as robotics. We ask some questions separately for Pred.-AI and Gen.-AI rather than combining them into a single question on AI because Gen.-AI is more recent and expanding rapidly, and we wanted to measure its growth separately.

We began by asking firms whether they had adopted each of these three technologies. Mimicking the question format from previous survey waves, we used a scale that distinguishes between experimental, limited, and extensive use.[16]

We then elicited beliefs about competitors' adoption of advanced technologies. Specifically, we asked: "In your opinion, what is the share of companies similar to yours in terms of sector and size, potentially your competitors, that are currently using robotics and/or artificial intelligence (generative and/or predictive AI)?" Subjects could choose from bins of 10 pp: below 10%, between 10% and 20%, and so on, up to above 90%. We refer to the response to this question as the *prior belief*. In the following analysis, we use the midpoints of each bin (e.g., 5% for the first bin, 15% for the second, etc.).

Immediately after eliciting the prior belief, half of the firms were randomly assigned to receive information about the share of competitors that had adopted advanced technologies. We refer to this information as the *signal*. To calculate the signal, we used responses from the previous wave of INVIND. For reference, Appendix H includes the survey module on advanced technologies from the 2024 wave. We calculated the share of firms that reported

[15] For a copy of the full questionnaire—including our module as well as all other sections—see Bank of Italy (2025).

[16] As a complementary measure, we also asked firms to report the share of their total 2024 investment that was directed toward these advanced technologies.

having already adopted AI or robotics or that planned to do so by the end of 2024.[17] We divided the 2024 respondents into cells based on sector and firm size, then calculated the adoption share within each cell. For sectors, we used the INVIND taxonomy consisting of 11 sectors.[18] For size, we split firms into those with fewer than 50 employees and those with 50 or more.[19] This procedure resulted in 22 distinct sector-size cells. The average number of respondents per cell was 154. The corresponding signals for each sector-size cell used in the information treatment are reported in Table C.1.

Returning to the 2025 module, after the information-provision experiment, firms were asked to report their expectations about the share of competitors that will be using these technologies by 2027. This question, which we refer to as the *posterior belief,* was posed to all respondents, regardless of whether they received information or not. These posterior beliefs allow us to assess whether firms adjusted their expectations about future competitor adoption in response to the information provided about current competitor adoption.

After eliciting the posterior belief, we asked respondents about their own intentions to adopt each of the three advanced technologies (Pred.-AI, Gen.-AI, and robotics) by 2027.[20] The goal is to test whether the information treatment affected firms' plans to adopt advanced technologies themselves.

As with all new questions included in the INVIND survey, the items from our module were tested by the Bank's branches through small pilot studies to assess whether they were easy to understand and whether the information was generally accessible to respondents.

[17]More precisely, we calculated the share of respondents who did not answer "not currently used and not expected to be introduced by December 2024" to either question TEC5N or TEC11N from the 2024 wave.

[18]Sector classifications in INVIND follow the Italian Statistical Institute's taxonomy, Ateco 2007. The sectors included in the analysis are listed in Figure 1.

[19]We faced a data-availability limitation: although the 2024 survey asked for the exact number of employees, only the binary classification for fewer or more than 50 employees was readily available when implementing the 2025 experiment. In any case, this split at 50 employees is a reasonable choice, given the trade-off between granularity and precision: using finer groups would provide more detailed information but at the cost of smaller sample sizes and thus lower precision.

[20]We used the same scale as for baseline adoption, distinguishing between no adoption, experimental, limited, or extensive use.

3.2 Follow-up Survey

All the firms who responded to our baseline survey were invited to participate in another survey wave one year later, in 2026. We included a short module in that survey, which we refer to as the follow-up survey. This module included a few questions specifically designed for this project. A copy of the full survey instrument is attached as Appendix G and is summarized below.

We included a key pair of questions intended to provide suggestive evidence on the relative strength of the competition and learning channels. The question began with "With reference to advanced technologies (e.g. AI, robotics), to what extent do you agree with the following statements?" The question then provided two statements—the order of which was randomized to address any potential ordering effects. One statement was intended to assess the presence of the learning channel: "The use of these technologies by other non-competitors (e.g. foreign suppliers or companies active in markets other than Italy) indicates their effectiveness and increases the likelihood of our company adopting the same technologies." The other statement was intended to assess the presence of the competition channel: "The use of these technologies by our competitors increases the likelihood of our company adopting the same technologies, in order not to lose competitiveness." For each statement, the respondent could indicate a value between 1 and 5, where 1 means "fully disagree" and 5 means "fully agree." Whether firms agree more with one statement than the other can provide suggestive evidence on the relative importance of the competition and learning channels.

Additionally, the follow-up survey included some questions similar to those from the baseline survey—for example, one question on the adoption of advanced technologies. The goal of including these questions was twofold. First, from a descriptive perspective, we can measure how firms' intended adoption at the end of the baseline survey relates to their adoption levels twelve months later. Second, we can measure treatment effects on the outcomes measured one year later.

Lastly, the follow-up questionnaire contained additional questions included by the Bank of Italy for purposes outside the scope of the study but that we use for descriptive analyses. These include questions on the specific types of generative AI tools currently used by the firm, the business functions in which Pred.-AI and Gen.-AI are primarily employed, and firms' assessments of how AI adoption has affected, and is expected to affect, productivity, employment, and investment over the previous and following three years.

3.3 Survey Implementation

A total of 3,983 firms participated in the 2025 survey wave. Because the questionnaire was relatively long, not all firms completed every module. The questions on AI and robotics were placed in the middle of the survey and were not mandatory. As a result, about 900 firms did not answer the key questions in our module, including those on current adoption and prior beliefs. This leaves us with an analysis sample of approximately 3,080 firms. For consistency, unless explicitly stated otherwise, all analyses use this respondent sample.

All firms that participated in the 2025 survey were invited to take part in the 2026 follow-up survey. Among the 3,080 firms in the main analysis sample, 80% completed the AI module in the follow-up survey. Response rates were balanced across experimental groups: the response rate was 79% in the control group and 80% in the treatment group, with no statistically significant difference.

3.4 Descriptive Statistics and Randomization Balance

Column (1) of Table 1 reports descriptive information about the subject pool in our baseline survey. The average firm is 40 years old and employs 96 workers; however, more than half of the firms have fewer than 50 employees. This pattern is typical of Italy and Europe, where firms are predominantly small and firm entry rates are low.

For the randomization balance check, columns (2) and (3) of Table 1 present average pre-treatment characteristics by control and treatment groups. For each characteristic, column (4) reports the p-value from a test of the null hypothesis that the means are equal between the two groups. Table 1 shows that, consistent with successful random assignment, pre-treatment characteristics are balanced across treatment and control groups.

Our main outcome of interest is whether a firm intends to adopt advanced technologies by 2027. To gauge baseline levels, we examine average intended adoption in the control group: 48% intend to adopt Pred.-AI by 2027, 52% intend to adopt Gen.-AI, and 37.5% intend to adopt robotics. These expected adoption rates are well above current adoption levels. For reference, column (2) of Table 1 shows that among firms in the control group, 16.6% had already adopted Pred.-AI, 25.0% had adopted Gen.-AI, and 23.0% had adopted robotics. Firms with AI and robotics use are usually larger, more investment-intensive and export-oriented (see Appendix B.3 for more details).

3.5 Imputation of Prior Beliefs

One potential concern is that when individuals are provided with information, they may "go back" and change their answers to earlier questions, particularly the one on prior beliefs. That is, due to social desirability bias, individuals may not want to appear ignorant and might therefore attempt to correct previous answers if given the chance. In online surveys, this concern can be fully mitigated by simply preventing respondents from returning to previous questions. For example, if a respondent learns that competitors' adoption rates differ from what they guessed in the prior-belief question, they cannot go back to change that initial response.

Unfortunately, the survey's delivery method did not allow us to implement this standard mitigation measure for all respondents. For about 5% of subjects, the survey was completed through an in-person visit from an interviewer, which effectively prevented respondents from revising their prior beliefs after receiving the information treatment. For the remaining 95%, however, this safeguard was not in place. Respondents typically accessed the survey through an online platform, either completing the questionnaire directly on the platform (web self-administered) or by downloading and filling out a PDF version (remote self-administered).[21] As a result, after reading the information, respondents could easily edit their earlier answers—most importantly, their prior belief. Due to constraints in the survey structure, the prior-belief question was placed immediately before the information provision, making it both easier and more tempting for respondents to revise their initial guess if it appeared inaccurate.

Whether prior beliefs were contaminated is straightforward to test empirically. Under the

[21]Some firms (around 30%) began completing the questionnaire with assistance from Bank of Italy personnel via telephone but later continued it on the online platform.

absence of contamination, the distribution of prior beliefs should be indistinguishable between the treatment and control groups. By contrast, if treated subjects revised their prior beliefs after receiving the information, we would expect the treatment group's prior beliefs to be more accurate than those of the control group. The data suggest no contamination of prior beliefs among the 5% of firms that responded via in-person visits, but some contamination among the remaining firms—for details, see Appendix C.

This does not mean that the prior-belief question is useless. Such responses are typically used for two purposes. The first is to characterize the distribution of prior misperceptions. For this purpose, we can simply restrict the analysis to the control group, who did not receive any information and thus faced no risk of contamination.

The second use of prior beliefs is for heterogeneity analysis. Intuitively, when presented with information, individuals should adjust their posterior beliefs upward, downward, or not at all, depending on whether their prior beliefs were below, above, or equal to the information. Even if some contamination exists, this heterogeneity analysis can still be conducted properly by using imputed prior beliefs. Specifically, using the control group—where contamination is not possible—we estimate a model that predicts prior beliefs based on pre-treatment firm characteristics. We then apply the estimated parameters to predict prior beliefs not only for the control group but also for the treatment group.[22] These imputed priors are then used to perform the heterogeneity analysis. The results of this imputation method, presented in Appendix C, show that the imputed prior beliefs, as intended, are statistically indistinguishable between the treatment and control groups.[23]

[22]For consistency with the measurement of raw prior beliefs (defined as the midpoints of ten bins), we round the estimated priors to the midpoint of the corresponding decile. For example, an imputed prior belief of 17% is rounded to 15%.

[23]Even when using raw prior beliefs, their inherently subjective nature implies that they are measured with noise, introducing attenuation bias. The imputation process adds further noise, thereby amplifying this bias and working against our results.

4 Technological Adoption and Misperceptions

4.1 Level and Trends in the Use of Advanced Technologies

Figure 1 presents descriptive evidence on the levels and trends in the use of advanced technologies.[24] Panels (a) and (c) correspond to the results on AI adoption. Earlier modules in the INVIND survey did not ask adoption questions separately for Pred.-AI and Gen.-AI, so these panels focus on AI adoption in general.[25] Panel (a) shows current usage in 2025, based on responses from the 2025 wave. Around 28% of firms reported using AI that year. Usage intensity varies widely—most firms employ AI experimentally or in a limited capacity, while the remainder use it extensively. There is also notable variation across sectors: the lowest adoption rates are observed in Textiles, Clothing, and Footwear (14%), and the highest in Real Estate and Other Services (48%).

As additional descriptive evidence, we included at the beginning of our module a question on the share of total investment devoted to advanced technologies in 2024. Among firms that adopted Gen.-AI, Pred.-AI, or robotics in a limited or experimental way, the median share of total investment allocated to advanced technologies was 2.5%. Among firms with extensive use of AI technologies, the median share was 4.4%, while among firms with extensive use of robotics it was considerably higher, at 15.5%.

Panel (c) of Figure 1 depicts the evolution of AI usage over time. AI adoption has accelerated sharply: the average annual increase rose from 0.5 pp in 2018–2020 to 2.3 pp in 2020–2024, and then to 14 pp in 2024–2025.[26] These AI adoption rates measured in our

[24] For further descriptive analysis of advanced-technology adoption, see Bencivelli et al. (2025), which uses data from the 2024 INVIND wave.

[25] For a breakdown of results by predictive and generative AI, see Appendix D.1.

[26] The corresponding raw increases over each period are: 1 pp in 2018–2020, 9 pp in 2020–2024, and 14 pp in 2024–2025.

survey are broadly consistent with those reported in other data sources.[27] Moreover, our data on expected adoption indicate that this exponential growth will continue, with the share of companies intending to adopt AI expected to double by 2027 (reaching 56%).

Panels (b) and (d) of Figure 1 mirror panels (a) and (c) but depict robotics adoption instead of AI adoption. In 2025, around 23% of firms used robots—a rate comparable to AI adoption. However, there are notable differences between the two. First, robotics adoption varies much more across sectors: for example, Basic Metals and Engineering and Other Manufacturing exhibit the highest 2025 adoption rate (48%), while Transport, Storage, and Communication show the lowest (5%). Second, the trajectory of robotics adoption differs from that of AI. Robotics uptake has been more stable, showing only modest growth in 2024 and with slower expansion projected through 2027. Robot usage is also more entrenched: while the share of firms adopting AI is similar to that of firms adopting robots in 2025, the former is primarily experimental, whereas the latter displays a substantially larger share of extensive use.

4.2 Misperceptions about Competitors' Adoption

Panel (a) of Figure 2 presents the results on perceived competitors' adoption of advanced technologies. Specifically, it shows the distribution of perception gaps in the control group. The x-axis measures the difference between the *signal*—our best guess for the "true" adoption rate—and the firm's corresponding prior belief. This histogram is constructed using the raw prior beliefs rather than the imputed ones, as the analysis is restricted to the control group and thus free from contamination concerns. Positive x-axis values indicate underestimation, negative values indicate overestimation, and a value of zero corresponds to accurate priors. Only a small share (2%) held accurate priors (within ±2.5 pp of the truth), and a similarly

[27] According to the 2025 wave of the EU survey on ICT usage and e-commerce in enterprises, 13.5% of EU firms and 8.2% of Italian firms with at least 10 employees have adopted AI technologies. Among large enterprises (more than 249 employees), these shares rise to 41.2% and 32.5%, respectively. Similarly, the 2025Q2 Survey on the Access to Finance of Enterprises (SAFE) reports that 34% of European firms have invested in AI technologies at some point, with the figure reaching 47% among large firms. Differences across surveys likely reflect variations in question design and sampling. In particular, the EU ICT survey covers all firms with more than 10 employees, whereas INVIND focuses on relatively larger firms (over 20 employees), which tend to exhibit higher adoption rates. Moreover, the SAFE survey includes firms that have invested in AI at any point in the past, even if they are not currently using or investing in such technologies.

small share (4%) overestimated adoption by at least 2.5 pp. The vast majority (93%) underestimated competitors' adoption by at least 2.5 pp, with many underestimating by a large margin and some by as much as 50 pp or more. On average, prior beliefs underestimated actual adoption by 24.5 pp, and the Mean Absolute Error was 26.2 pp.[28] Firms with larger prior misperceptions tend to be larger, more often manufacturers, more likely to export, and less likely to invest in advanced technologies (see Appendix B.3 for more details).

When interpreting the evidence from Panel (a) of Figure 2, one caveat should be kept in mind. To measure true misperceptions, one would ideally compare prior beliefs to the actual adoption rates. In our case, however, the true adoption rates are not perfectly observable. Instead, we rely on a signal that is subject to some measurement error, since it is computed using a finite sample of respondents and therefore affected by sampling variation. Moreover, as is common in the measurement of beliefs, prior beliefs themselves may also contain measurement error—for example, due to some subjectivity in the wording of the question.[29] Consequently, some of the gaps shown in Panel (a) of Figure 2 may partly reflect measurement error in the signal and in the prior beliefs. Nonetheless, given the large magnitude of these gaps, our preferred interpretation is that measurement error is unlikely to fully explain these patterns and that they largely reflect genuine misperceptions.

Our finding that firms hold misperceptions about factors relevant to their decision-making is consistent with a growing body of evidence showing that firms—even large ones—can hold substantial misperceptions. For instance, Cullen et al. (2026) show that firms misperceive the wages offered by other employers. Similarly, Kim (2025) shows that firms are often uninformed about their competitors' prices. Finally, Coibion et al. (2018) use survey data to document significant misperceptions about macroeconomic conditions, such as inflation and economic growth.

[28] The signal about the current adoption in 2025 was computed using survey responses from the 2024 wave. In the treatment message, we explicitly mentioned that the signal was based on responses to the last survey: "The findings of the last survey...". In 2024, the question asked firms whether they had adopted or were planning to adopt by December 2024. Because the 2025 wave was conducted between February and May 2025, one could argue that the signals we provided slightly underestimate true adoption rates, to the extent that adoption likely continued to increase between January and May 2025. However, that mismatch is probably negligible in magnitude, and would only imply that the degree of underestimation is even larger—further reinforcing the main message.

[29] For example, different respondents may interpret differently what "similar to yours in terms of sector and size" means.

Given the substantial misperceptions about competitors' adoption of advanced technologies, we can discuss some of their potential sources. One likely contributor is the exponential growth of these technologies: without actively seeking information, firms' beliefs can quickly become outdated. Another source of misperceptions is friction in information diffusion. On the one hand, firms may have incentives to publicize their technological adoption—for example, to signal innovation to customers or demonstrate growth potential to investors. On the other hand, firms may prefer to conceal adoption to preserve a competitive edge.[30] In practice, these frictions mean that information about technological adoption often travels through informal networks rather than formal disclosures. For instance, related evidence suggests that information diffuses among firms through informal networks of executive board members (Faia et al., 2026) and business owners (Cai and Szeidl, 2018).[31]

5 Effect of the Information Treatment on Firms' Beliefs

To model belief updating, we apply the simple Bayesian learning framework from Cavallo et al. (2017). Let $s_{i,t}$ denote firm i's belief about the share of competitors that have adopted advanced technologies up to 2025, and let $s_{i,t+1}$ denote the expectation about the share of firms that will adopt the technology in the future. We assume that firms form their future expectations by projecting their perceptions about the past:

$$s_{i,t+1} = \eta + \beta \cdot s_{i,t}, \tag{1}$$

where β captures the degree of pass-through from perceived past adoption to expected future adoption. In this context—characterized by a period of sustained growth in adoption rates—it would be natural to expect $\beta > 1$. In other words, firms anticipate that future adoption

[30] Whether firms share information about innovation or other strategic decisions depends on the trade-off between losing customers and internalizing the benefits of mutual learning—see for example Stein (2008).

[31] A further factor may be broader uncertainty about productivity, which is often higher in the early stages of a technological breakthrough—such as AI—and tends to decline as data accumulate; see Crawford and Shum (2005) and Farboodi et al. (2019).

rates among competitors will exceed current levels.[32]

Let T_i be an indicator variable equal to 1 if the individual was randomly chosen to receive a signal and 0 otherwise. We start with the case in which the individual receives the information ($T_i = 1$).

In this case, the firm's posterior belief about current adoption, $s_{i,t}$, may depend on the firm's prior belief about current adoption ($s^0_{i,t}$) and on the signal ($s^T_{i,t}$) received via the experiment. Based on the assumptions of a Bayesian learning model with Gaussian distributions,[33] after observing the information the individual is expected to update beliefs about past adoption as follows:

$$s_{i,t} = \alpha \cdot s^T_{i,t} + (1-\alpha) \cdot s^0_{i,t} \quad \text{if } \; T_i = 1, \tag{2}$$

where $\alpha \in [0, 1]$ is the weight assigned to the new information relative to the prior belief, which depends on the accuracy of the prior belief relative to the accuracy of the signal. If we combine equations (1) and (2) we obtain the following expression:

$$s_{i,t+1} = \eta + \alpha \cdot \beta \cdot \underbrace{\left(s^T_{i,t} - s^0_{i,t}\right)}_{\text{Prior Gap}} + \beta \cdot s^0_{i,t} \quad \text{if } \; T_i = 1 \tag{3}$$

The key prediction from the model is that, for individuals who received information, the belief update should be a linear function of the prior gap. Intuitively, respondents who overestimated the true share of adopters should revise their belief downward upon receiving the signal; those who underestimated should revise it upward; and those who were already accurate should exhibit no updating. Note also that the strength of this belief updating is given by the product of the two key parameters: the learning weight (α) and the degree of extrapolation (β).

[32] Whether we expect β to be above, equal to, or below 1 is context specific. In particular, it depends on whether the variable being forecasted is expressed in levels or in growth rates. For example, in the context of home price expectations, we would expect future price levels to exceed current levels, implying $\beta > 1$. However, if the variable being forecasted is the growth rate of prices, we might expect $\beta = 1$ (if growth is expected to continue at the same pace) or $\beta < 1$ (if mean reversion is expected).

[33] These assumptions include that the priors and the signal are normally distributed, and the variance of the prior and the signal is independent of the prior's mean. See Section C of Cavallo et al. (2017) for further discussion.

One possible concern with estimating equation (3) directly is that individuals may appear to adjust their beliefs toward the signal for reasons unrelated to actually receiving the information. For instance, simply being asked the same question twice could prompt them to reflect more carefully, revise their earlier response, or fix typographical mistakes, which would mechanically bring their second answer closer to the truth. To separate genuine learning effects from these spurious sources of updating, we rely on the randomized assignment of information and estimate the following specification:

$$s_{i,t+1} = \gamma_0 + \gamma_1 \cdot T_i \cdot \left(s_{i,t}^T - s_{i,t}^0\right) + \gamma_2 \cdot \left(s_{i,t}^T - s_{i,t}^0\right) + \gamma_3 \cdot s_{i,t}^0 + \epsilon_{i,t} \tag{4}$$

Intuitively, the key parameter is γ_1, which measures whether the slope between belief updates and prior gaps is stronger for individuals who received information ($T_i = 1$) than for those who did not ($T_i = 0$). In terms of equation (3), this parameter γ_1 identifies the product $\alpha \cdot \beta$. In turn, γ_2 identifies the degree of spurious learning, while γ_3 identifies the parameter β. Thus, by taking the ratio $\frac{\gamma_1}{\gamma_3}$, we can separately identify the parameter α.

The results for the effects of the information on belief updating are presented in panel (b) of Figure 2. The y-axis represents the posterior belief ($s_{i,t+1}$) and the x-axis represents the imputed prior misperceptions ($s_{i,t}^T - s_{i,t}^0$). As predicted by the simple learning model of equation (4), the treatment and control groups exhibit significantly different slopes. Moreover, this figure reports the estimates for parameters β and α from estimating equation (4). The estimated β is large (1.19) and statistically significant ($p<0.001$). The fact that β is above 1 indicates that, on average, firms anticipate strong growth in the adoption of advanced technologies between 2025 and 2027—a reasonable expectation given the exponential growth observed in recent years.

In turn, the estimated α is positive (0.28) and statistically significant ($p < 0.001$). This estimate suggests that subjects placed significant weight on the signal, implying that subjects both understood the information and regarded it as trustworthy and relevant. Moreover, the magnitude of the weight can be compared with that reported in other experiments. For example, Cavallo et al. (2017) find that, when forming inflation expectations, the average Argentine respondent assigns a weight of 0.432 to the inflation signal. Similarly, Fuster et al. (2022) find that, when forming home-price expectations, the average subject assigns a weight of 0.380 to the signal. The weight in our study (0.28) is somewhat lower than in these other contexts (0.432 and 0.380), though not dramatically so. There are two potential explanations that could, individually or jointly, account for this difference. First, relative to the other studies, subjects in our setting may perceive the signal as less precise. The other studies rely on information derived from large samples—such as the Consumer Price Index—whereas our signal is based on a smaller sample of respondents from a previous survey wave. As a result, subjects in our context may reasonably infer that the signal is noisier and therefore place less weight on it. Second, part of the difference may reflect attenuation bias, as prior beliefs in our study are imputed—and therefore noisier—whereas in the other studies they do not need to be imputed.

6 Effect of the Information Treatment on Firms' Own Future Adoption Plans

6.1 Intention to Treat Estimates

Panel (a) of Figure 3 reports the average treatment effects across several outcomes. The first outcome is the expected share of competitors adopting advanced technologies by 2027. The treatment raises this expectation by 8.0 pp ($p<0.001$), from 19.2% to 27.3%. This result is intuitive: because individuals tend to underestimate competitors' adoption, providing them with accurate information about actual adoption leads them to revise their beliefs upward—both about current adoption and, in turn, about expected future adoption.

The other three outcomes reported in panel (a) of Figure 3 correspond to the intention to adopt Pred.-AI, Gen.-AI, and robotics, respectively.[34] For robotics, we find that, relative to the control group, the information treatment increased the intention to adopt by 6.5 pp. This effect is both statistically significant ($p = 0.039$) and economically meaningful, representing a rise from 37.5% to 44.0%—a 17.3% increase ($= \frac{6.5}{37.5}$). In contrast, we find no significant

[34] In the baseline specification, we distinguish between intended use and no intended use, regardless of intensity. For transparency, Appendix D presents the main result separately by intensity level.

effects on the intention to adopt AI technologies: the estimated impacts are close to zero (2.3 pp for Pred.-AI and -0.1 pp for Gen.-AI) and statistically insignificant (p=0.478 and p=0.985, respectively). These null results, however, should be interpreted with caution, as the estimates are imprecise and do not rule out modest positive effects—we discuss this issue in greater detail at the end of this section.[35]

One caveat to keep in mind when interpreting the effects of information on robotics adoption is the possibility of experimenter-demand effects. For example, after learning that competitors' adoption rates are higher than expected, some respondents may feel pressure to report greater adoption intentions themselves. This explanation appears unlikely, however, because we find effects for robotics but not for AI, and it is unclear why experimenter demand would influence responses for one technology but not the other.

To assess whether these findings are real or spurious, panel (b) of Figure 3 reports a series of falsification tests. It mirrors panel (a) but uses pre-treatment outcomes instead of post-treatment ones. Since the treatment had not yet been administered, we expect no effects on these pre-treatment outcomes. The first pre-treatment outcome is the (imputed) prior belief. As expected, the difference in prior beliefs is close to zero (0.1 pp) and statistically insignificant ($p = 0.777$). The other three pre-treatment outcomes correspond to the current adoption of Pred.-AI, Gen.-AI, and robotics technologies. For robotics, the effect is negligible (-0.1 pp) and statistically insignificant ($p = 0.971$). Similarly, for Pred.-AI and Gen.-AI, the effects are small (3.3 pp and 0.1 pp, respectively) and statistically insignificant ($p = 0.180$ and 0.978).

Next, we examine heterogeneity in treatment effects by prior beliefs. Intuitively, firms that more strongly underestimated their competitors' adoption should update their beliefs more sharply in response to information and, as a result, exhibit a stronger increase in their intention to adopt. Indeed, this is exactly what we find. Figure 4 presents these results, effectively splitting panel (a) of Figure 3 into two groups: firms with prior gaps below versus above the median. The below-median group includes a mix of firms that either somewhat overestimated or underestimated competitors' adoption, with prior gaps ranging

[35] For instance, although the point estimates are close to zero (2.3 pp for Pred.-AI and -0.1 pp for Gen.-AI, respectively), the corresponding 90% confidence intervals cannot rule out effects of up to 7.8 pp and 5.4 pp, respectively.

from -21.3 pp to 24.1 pp. The above-median group consists of firms that most strongly underestimated competitors' adoption, with prior gaps ranging from 24.2% to 57.7%.

Panel (a) of Figure 4 shows that, as expected, the effect on posterior beliefs was highly heterogeneous with respect to prior gaps: 3.1 pp (p=0.147) for the below-median group versus 13.1 pp for the above-median group, with the difference of coefficients being statistically significant (p<0.001). Consistently, panel (d) of Figure 4 shows that the treatment effects on the intention to adopt robotics were also highly heterogeneous: 2.7 pp (p=0.584) for the below-median group versus 11.7 pp for the above-median group, and the difference is statistically significant (p<0.001).

For individuals with above-median prior gaps, we find a large effect on robotics adoption (panel (d)). By contrast, for Pred.-AI and Gen.-AI adoption, we do not find significant effects within this same subgroup. Panels (b) and (c) of Figure 4 show that, even among individuals with above-median prior gaps, the effects of information were close to zero (-0.5 pp for Pred.-AI and 1.0 pp for Gen.-AI) and statistically insignificant (p = 0.895 and 0.780, respectively). For completeness, Figure 5 presents the heterogeneity analysis for the falsification outcomes. Specifically, it replicates panel (b) of Figure 3, splitting the sample into below-median and above-median prior gaps. The estimated differences are small and, in all but two of the eight comparisons, statistically insignificant at conventional levels.[36]

Since the treatment did not increase planned AI adoption even among those with above-median prior gaps, this might appear at first glance to be a null effect. However, this result should be interpreted with caution, as the estimates are imprecise and therefore do not allow us to rule out modest positive effects on AI adoption. For instance, although the point estimates from panels (b) and (c) of Figure 4 are close to zero (-0.5 and 1.0 pp, respectively), the corresponding 90% confidence intervals cannot rule out effects of up to 5.4 pp and 6.9 pp, respectively. In other words, we cannot rule out AI adoption effects roughly half as large as those observed for robotics adoption (11.7 pp).

While we cannot rule out positive treatment effects on AI adoption, the evidence suggests

[36]The two exceptions are imputed prior beliefs among firms with above-median prior gaps (0.6 pp, p=0.062)—a difference that is economically negligible—and current use of robotics in the same subgroup (5.3 pp, p=0.064). Given the large number of falsification tests we conduct, some are expected to be significant just by chance. Across the twelve comparisons reported in panel (b) of Figure 3 and in Figure 5, we expect 1.2 differences significant to be significant at the 10% level just by chance.

that they were probably weaker than the corresponding effects on robotics adoption. We can offer some plausible explanations for this difference. The first relies on differences in baseline adoption rates. In the control group, intended AI adoption is higher than intended robotics adoption—37.5% for robotics versus 48.0% and 52.0% for Pred.-AI and Gen.-AI, respectively. The substantially higher baseline rates for AI may simply leave less room for further increases. A second factor may be that, at the time of the experiment, robotics was a more established technology than AI, with a much longer history of adoption and more extensive rather than experimental use. As a result, if a firm learns that robotics adoption is higher among competitors, the perceived urgency to respond may be greater. Indeed, responses to a question in the follow-up survey are consistent with this interpretation. Most firms report that AI adoption has produced little or no measurable improvement in productivity over the past three years, although they are considerably more optimistic about its expected effects over the next three years. Additionally, another survey shows that firms face challenges that may be unique to AI adoption, such as a lack of skills or limited applicability.[37]

6.2 From Ex-Ante Intentions to Ex-Post Adoption

One caveat to keep in mind is that treatment effects on intentions may translate only slowly into actual adoption, or not at all. The adoption of advanced technologies is typically a gradual process, involving multi-year investments and training. Thus, even when intentions are strong, adoption may take time to materialize. Moreover, even when firms intend to adopt, implementation may fail to occur for a variety of reasons: stakeholders may slowly forget about the information treatment, lose enthusiasm, face higher-than-expected costs, encounter financing constraints, or be unable to find individuals with the required skills. Although we cannot trace the full adoption funnel, some of the data from the follow-up survey can provide suggestive evidence.

First, we find some evidence that stronger adoption intentions predict higher future adoption. Figure 6 presents this evidence by combining responses from the baseline and follow-up surveys. The sample consists of firms that responded to both waves and reported not using robotics in 2025. Among firms not using robotics in 2025, 23.0% intended to adopt it by

[37] For more details, see Appendix D.1.

2027, while the remaining 77.0% did not. One year later, adoption rates were higher among firms that had planned to adopt by 2027 than among those that had not: 19.3% versus 5.2%. Figure 6 also shows that the intensity of adoption intentions is positively correlated with ex-post adoption: firms that intended extensive use adopted at higher rates than those that intended limited use, which in turn adopted at higher rates than those that intended experimental use.

On the other hand, Figure 6 shows how slow robotics adoption can be. Although some firms that intended to adopt by 2027 had already done so by the beginning of 2026, they were a minority (19.3%). Even among those with the strongest adoption intentions (39 firms planning extensive use), only 28.1% had adopted robotics by 2026. This likely reflects a wide range of factors, such as stakeholders losing interest, unexpected costs or financing constraints, or challenges in finding suppliers.

Given the low adoption rate, we are underpowered to detect treatment effects on adoption one year later. As a back-of-the-envelope calculation, at baseline we found that intended adoption increased by 6.5 pp. If only 19.3% of that adoption was expected to have materialized at follow-up, we would expect an effect on adoption at follow-up of just 1.25 pp ($= 6.5 \cdot 0.193$). However, we are underpowered to detect effects of this magnitude. Moreover, the analysis of follow-up outcomes presents additional practical challenges. First, only a subset of firms responded to the follow-up survey, reducing the statistical power to detect significant effects. Second, we provided the information only once, and its effects on beliefs may decay over time. For example, other information experiments document that most effects decay within a few months (e.g., Cavallo et al., 2017); this decay could be even stronger in our case because we follow up a full year later. With those caveats in mind, we find a positive effect on adoption of 2 pp, in line with the expected magnitude of 1.25 pp described above, but the effect is imprecisely estimated and thus statistically insignificant. However, we find a large positive effect (13.3 pp, p=0.071) among firms with the strongest adoption plans—see Appendix E for more details. One possible interpretation is that the information treatment may have affected adoption among firms with the strongest plans to adopt, but that a one-year horizon is too short to detect effects in the full sample.

6.3 2SLS Model

The results presented above capture intention-to-treat effects of providing information, which are not the same as the effects of expectations. To measure the latter, we employ a simple 2SLS model commonly used in information-provision studies (e.g., Cavallo et al., 2017; Cullen and Perez-Truglia, 2022; Giaccobasso et al., 2025). Our main outcome of interest, $a^j_{i,t+1}$, is an indicator equal to 100 if individual i intends to adopt technology j by 2027, where j can be Gen.-AI, Pred.-AI, or Robots. We aim to estimate a regression of future adoption ($a^j_{i,t+1}$) on the expected future adoption of competitors ($s_{i,t+1}$). Conceptually, this is the reaction function that relates a firm's own adoption choice to competitors'adoption —Appendix A provides a simple framework that formalizes this. However, because beliefs may be endogenous, such a regression would not necessarily identify a causal effect. To obtain causal identification, we use the following 2SLS model that exploits only the exogenous variation in posterior beliefs generated by the randomized provision of information:

$$a^j_{i,t+1} = \psi_0 + \psi^j_s \cdot s_{i,t+1} + \psi_2 \cdot (s^T_{i,t} - s^0_{i,t}) + \psi_3 \cdot s^0_{i,t} + X_i \psi_X + \epsilon_i \tag{5}$$

The endogenous variable is $s_{i,t+1}$ and the excluded instrument is $T_i \cdot (s^T_{i,t} - s^0_{i,t})$.[38] X_i is a vector of additional control variables, such as the firms' current adoption status and basic characteristics.[39]

We can illustrate the intuition behind the 2SLS model with a simple example. Consider a pair of firms with the same prior gap. For instance, suppose the two firms underestimated competitors' 2025 adoption rate by 10 pp. Through random assignment, one firm receives information about the adoption rate, while the other does not. Suppose the uninformed

[38] Formally, the exogeneity assumption is $\mathbb{E}\left[\left(s^T_{i,t} - s^0_{i,t}\right) \cdot T_i \cdot \epsilon_i \mid \mathbb{X}_i\right] = 0$, where $\mathbb{X}_i$ is a vector that includes $\{s^T_{i,t} - s^0_{i,t}, s^0_{i,t}, X_i\}$. In plain English, we assume that heterogeneity in the effects of information is driven solely by differences in prior misperceptions, and not by heterogeneity in other unobserved factors that are correlated with those misperceptions.

[39] The full set of controls includes the following: indicator variables for the firms' current adoption; the firms' age and geographical area; an indicator variable for being an exporter; turnover and number of employees (standardized within sector-size cells); the share of total investment over turnover; the share of investment in advanced technologies (AT); an indicator variable for having benefited from or expecting to benefit from tax credits for capital goods under the Transition 4.0 programme; an indicator variable equal to 1 if the firm had or expects to receive orders linked to the National Recovery and Resilience Plan; and two variables capturing the firms' attention in answering the survey (the number of people involved in completing the questionnaire and whether any manager participated in the process).

firm continues to underestimate future adoption by 10 pp. By contrast, the informed firm finds the information persuasive and thus underestimates future competitor adoption by only 2 pp. In this case, the information provision can be interpreted as an 8 pp positive shock to the expected adoption rate of competitors. Now suppose that, relative to the firm that did not receive the information, the firm that received it has a 6 pp higher probability of adopting AI in the future. Taking the ratio of these two effects implies that, for each 1 pp increase in perceived competitor adoption, a firm's own adoption probability increases by 0.75 pp ($= \frac{6}{8}$). This analysis focuses on a pair of firms that underestimated competitors' adoption by 10 pp, but in practice few respondents fall exactly into this subgroup, so this ratio cannot be estimated precisely. However, the same logic can be applied to pairs of firms that underestimated by other margins or even to firms that overestimated, and the results can then be averaged across all groups. This is what the IV regression is intended to do.[40]

6.4 2SLS Results

Table 2 reports the results from the 2SLS model. Each column corresponds to a separate regression using the same data and specification, with the only difference being the dependent variable. In column (1), the dependent variable equals 100 if the firm expects to adopt predictive AI by 2027. In column (2), it equals 100 if the firm expects to adopt generative AI. Column (3) corresponds to robotics adoption. Each column reports the 2SLS estimate (Panel (a)), as well as the corresponding first-stage results (Panel (b)) and reduced-form results (Panel (c)).[41]

When the dependent variables are the intentions to adopt predictive AI and generative AI (columns (1) and (2)), the 2SLS coefficients are positive but small (0.015 and 0.060) and statistically insignificant. By contrast, when the dependent variable is the intention to adopt robotics (column (3)), the 2SLS coefficient is positive, large (0.732), and statistically signifi-

[40] For additional details on the 2SLS framework, see Cullen and Perez-Truglia (2022). A caveat is that the 2SLS estimate recovers a Local Average Treatment Effect (LATE)—the average effect of beliefs for firms whose posteriors shift in response to the information intervention. By design, this estimate assigns greater weight to firms with larger initial misperceptions and, conditional on those misperceptions, to those that adjust their beliefs more strongly when exposed to the signal.

[41] The slight differences in the number of observations across columns (1)–(3) reflect differences in missing observations across the three outcomes.

cant (p<0.001).[42] Taken together, these 2SLS estimates corroborate the evidence discussed above that higher expected competitor adoption increases firms' intentions to adopt robotics but has no effect on intentions to adopt predictive or generative AI.

A common concern in the estimation of 2SLS models is weak-instrument bias (Stock et al., 2002). Given the strong effect of the treatment on belief updating documented in Section 5, this is unlikely to be a concern in our application. However, for a formal assessment, the bottom of Table 2 reports the Kleibergen-Paap *rk* Wald F-statistics for each 2SLS specification—a standard measure of instrument strength.[43] Following the guideline of Staiger and Stock (1997), F-statistics of 10 or higher indicate that weak identification is not a serious concern. Our reported values are well above this threshold, confirming that weak instruments are not a concern.

The 2SLS estimate makes the magnitude of the effects easier to interpret. Equation (5) can be interpreted as a reaction function: how much more likely a firm is to intend to adopt a technology when it believes that a larger share of competitors will adopt it as well.[44] More specifically, the 0.732 estimate from column (3) of Table 2 indicates that a 1 pp increase in the expected share of competitors adopting advanced technologies causes a 0.732 pp increase in the firm's probability of intending to adopt robotics by 2027. For an even more intuitive interpretation, we can also express this effect as an elasticity. The average firm in the control group expects 19.2% of its competitors to adopt advanced technologies and has a 37.5% probability of intending to adopt robotics. Thus, a 1 pp increase in the expected share of competitors adopting advanced technologies corresponds to a 5.21% ($=\frac{1}{19.2}$) increase relative to the baseline, and the corresponding 0.732 pp increase in the probability of intended robotics adoption corresponds to a 1.95% ($=\frac{0.732}{37.5}$) increase relative to that baseline. Taken together, these imply an elasticity of 0.375 ($=\frac{1.95}{5.21}$) between expected competitor adoption and firms' own intended robotics adoption. In other words, a 10% increase in expected competitor adoption would increase the firm's probability of intending to adopt robotics by 3.75%.

[42] Consistent with these findings, the reduced-form results (panel (c) of Table 2) show significant effects of the information treatment on intended robotics adoption (column (3)) but no significant effects on intended adoption of either form of AI (columns (1) and (2)).

[43] Although the conventional rule of thumb is based on the Cragg-Donald statistic, which assumes homoskedastic errors, Baum et al. (2007) recommend using the Kleibergen-Paap statistic as its robust counterpart.

[44] In our conceptual model in Appendix A, the corresponding reaction function is given by equation (A.4).

6.5 Heterogeneity Analysis and Causal Mechanisms

Ceiling Effects. One feature of the empirical design may lead us to understate the true magnitude of the effects. The key issue is that the dependent variable is subject to a ceiling effect. For firms that have not yet adopted the technology, the intervention can raise the probability of future adoption. By contrast, for firms that have already adopted it, there may be little or no scope for the intervention to increase adoption further. Put differently, these firms are likely to continue using the technology regardless of whether they receive the intervention.

The relevant evidence is presented in columns (4) and (5) of Table 2, which estimate the specification from column (3)—which focuses on future robotics adoption—separately for two groups: column (4) corresponds to the 77% of firms that had not used robotics in the past, while column (5) corresponds to the 23% that reported some prior use of robotics, whether experimental, limited, or extensive.[45] The results point to a substantial ceiling effect. Among firms with no prior robotics use, the estimated effect on future robotics adoption is large (1.070) and statistically significant ($p<0.001$). By contrast, among firms with prior robotics use, the estimated effect is much smaller (0.170) and statistically insignificant ($p = 0.270$). The clearest indication of this ceiling effect is the baseline mean in column (5), which is 97.3%. In other words, among firms that had already adopted robotics, 97.3% are expected to continue using it in the future, leaving very little room for any treatment—whether informational or otherwise—to increase adoption further.

This ceiling effect could also be relevant for the AI outcomes. Intuitively, the small estimated effects on AI adoption might reflect a similar ceiling effect and could otherwise have been larger. For brevity, we present these additional results in Appendix D.3. However, they do not alter the interpretation of the evidence discussed above. Even when we restrict the sample to firms with no prior AI adoption, the 2SLS estimates for AI adoption remain close to zero and statistically insignificant.

Manager Involvement. Another feature of the empirical design that may lead us to understate the true magnitude of the effects is the role of the survey respondent. Whether

[45] These shares, like all others reported in the paper, are calculated using sample weights. As a result, they may not match the unweighted numbers of observations reported at the bottom of Table 2.

information affects the firm's future behavior depends critically on the decision-making authority of the person responding to the survey and thus receiving the information. Effects are more likely if this person can influence technology-adoption decisions. By contrast, if the survey is completed entirely by someone with no meaningful authority over such decisions (e.g., an assistant or accountant), the information is unlikely to affect the firm's future adoption.

As explained in Section 2.1, although we do not know the identity of the individuals involved in completing the survey, one survey question provides a proxy for whether a manager was involved. The relevant evidence is presented in columns (6) and (7) of Table 2, which estimate the specification from column (3)—which focuses on future robotics adoption—separately for two groups: column (6) corresponds to the 44.5% of firms reporting that a manager was not involved in completing the survey, while column (7) corresponds to the remaining 55.5% of firms reporting that a manager was involved. We find that this heterogeneity is strong and in the expected direction. Among firms in which a manager was involved in the survey, the estimated effect on future robotics adoption is large (1.223) and statistically significant ($p<0.001$). By contrast, among firms in which a manager was not involved, the estimated effect is much smaller (0.230) and statistically insignificant ($p=0.370$).[46]

Financial Incentives. The next heterogeneity analysis examines how information interventions may interact with existing financial incentives. The Italian Transition 4.0 programme introduced subsidies to support firms investing in advanced technologies and innovation, with the aim of increasing technology adoption. The treatment effects of information documented above suggest that information campaigns could serve as an additional policy lever to foster technology adoption. A natural question, then, is how the effects of an information campaign relate to those of more traditional financial incentives. More specifically, would an information campaign primarily affect the same firms that were reached by these subsidies, or would it instead be more effective among firms that do not benefit from financial incentives? This distinction matters for policy design. For instance, if information campaigns

[46] Whether the manager was involved may also matter for the AI outcomes. In principle, the small estimated effects on AI adoption could be stronger once attention is restricted to firms in which a manager participated in the interview. For brevity, we report these additional results in Appendix D.3. However, the conclusions remain unchanged: even when limiting the sample to firms with managerial involvement, the 2SLS estimates for AI adoption are still close to zero and statistically insignificant.

are more effective among firms not reached by subsidies, they could complement traditional financial support by broadening the set of firms that adopt.

The results of this additional heterogeneity analysis are reported in the last two columns of Table 3. Specifically, columns (6) and (7) split the sample according to whether the firm either used, or expected to use by the end of 2025, the tax credit for capital goods under the Transition 4.0 programme. We construct this measure from the following question in the 2025 wave of INVIND: "Have you used the following incentives for new investment in capital goods in 2024, or do you plan to use them in 2025?" Respondents were asked about two types of incentives; the one relevant for our analysis is the second, labeled: "Tax credit for capital goods under the Transition 4.0 programme (*new tangible and intangible capital goods for the technological and digital transformation of production processes.*)"

This Transition 4.0 incentive applied to 35.8% of firms in our sample. The estimates suggest that the effect of information is substantially stronger among firms that did not benefit from the incentive. For these firms, the coefficient is 0.918 and highly statistically significant ($p<0.001$). By contrast, among firms that did benefit from the incentive, the coefficient is 0.420—roughly half as large—and statistically insignificant ($p=0.112$). These results should, however, be interpreted with caution. Because receipt of the financial incentive was not randomized, the comparison may partly capture other differences between firms that did and did not benefit from the programme. Moreover, the difference between the two coefficients is estimated imprecisely and is itself statistically insignificant ($p=0.143$). Even so, the pattern is suggestive: it indicates that information about competitors' adoption may be especially valuable for firms not reached by standard financial incentives and may therefore serve as a complementary policy tool for broadening technology adoption across a wider set of firms.

Competition versus Learning. The effect may operate through at least two underlying causal mechanisms: competition—where firms seek to avoid lagging behind their rivals—and learning—where firms infer their own potential productivity gains from competitors' adoption decisions. Indeed, the simple model in Appendix A shows that this effect can be decomposed into two additive components, each corresponding to one of these mechanisms. Because these mechanisms are closely intertwined, disentangling them is challenging. A firm that learns that a competitor has adopted may be compelled to adopt both because it has gained information about the technology's returns and because it fears losing business. Nevertheless, we can provide suggestive evidence through a heterogeneity analysis. Intuitively, the competition channel should depend on the degree of competition in the market. In the extreme case of a perfect monopolist, a firm could still learn from the adoption decisions of other firms, but the competition channel should be completely absent.

This additional heterogeneity analysis is presented in Table 3. Column (1) reproduces the baseline result from column (3) of Table 2, which corresponds to the effect on intended robotics adoption. Columns (2) through (7) of Table 3 use the same specification as column (1) but are estimated on different subsamples.[47] In columns (2) and (3), we split the sample according to whether a firm's market share is below or above the sample median. We use balance-sheet data from CADS to compute firm-level market shares, defined as each firm's share of total revenues among all firms in the same sector. More precisely, column (2) corresponds to firms with below-median market shares, while column (3) corresponds to firms with above-median market shares. For reference, the average market share in column (3) is almost nine times as large as the corresponding average in column (2). The 2SLS coefficient is substantially larger for firms with lower market shares than for firms with higher market shares: 1.264 in column (2) versus 0.112 in column (3). This difference is not only economically large, but also statistically significant (p=0.015). The fact that the effects are stronger for firms with lower market share suggests that the competition channel may play a significant role. Moreover, this result is largely consistent under an alternative specification: if we calculate market shares using more granular sectors (4-digit instead of the broad 11-sector classification) by using electronic invoicing data instead of balance-sheet data, we find consistent evidence—results reported in Appendix D.3. On the other hand, this heterogeneity should be interpreted with caution because market share was not randomized and the analysis is therefore subject to the usual omitted-variable concerns. That is, firms with lower market shares may respond more strongly not only because of their market share,

[47]The sample splits reported in columns (2)–(5) may appear uneven even though the median is used as the splitting criterion. This is because the number of observations reported at the bottom of Table 3 is unweighted, whereas the median is calculated with sample weights. In addition, for the HHI, 22% of observations are exactly at the median; these observations are classified as above the median.

but also because of other firm characteristics correlated with market share, such as firm size.

Columns (4) and (5) of Table 3 examine the same heterogeneity through an alternative measure of market structure. Using balance-sheet data from CADS, we compute the Herfindahl–Hirschman Index (HHI) for each of the 11 sectors and then split the sample according to whether sectoral HHI is below the median (column (4)) or above the median (column (5)). For reference, the average HHI in column (5) is four times as large as the average in column (4). The results are imprecisely estimated and therefore inconclusive. On the one hand, the estimated effects are statistically significant in less concentrated markets but statistically insignificant in more concentrated markets. On the other hand, this pattern appears to be driven largely by lower precision in the estimates for more concentrated markets. These results should again be interpreted with caution because market concentration was not randomized. Moreover, the findings are quite sensitive to the measure of market concentration: they change substantially when we use a more granular definition based on electronic invoicing data—see Appendix D.3.

To provide further suggestive evidence on the relative strength of the learning and competition channels, we use responses to two questions introduced in the 2026 survey. These questions asked respondents about their agreement with one statement capturing the competition channel and another capturing the learning channel. The results are presented in Figure 7.[48] Although the figure presents the results in a fixed order, the order of the two questions was randomized to avoid question-ordering effects.[49] Panel (a) shows that the distribution of agreement is broadly similar across the two statements, though it is shifted somewhat to the right for the competition statement relative to the learning statement. Our preferred interpretation, consistent with the heterogeneity analysis, is that both the competition and learning channels are at play, but that the competition channel may be somewhat more important. To provide a more direct comparison of the relative strength of the two channels, Panel (b) shows the distribution of the gap between agreement with the competition statement and agreement with the learning statement. Most firms (65.3%) reported equal agreement with the two statements; a minority (5.9%) agreed more strongly with the

[48]Due to an implementation issue with this question, the figure is based on a subsample of respondents. The results are similar in the broader sample; see Appendix E for details.

[49]For an analysis of the effects of order randomization, see Appendix E.

learning statement; and a larger minority (28.8%) agreed more strongly with the competition statement.

7 Conclusions

This study provides causal evidence on whether firms' innovation decisions exhibit strategic complementarities. Using a large-scale survey experiment embedded in the Bank of Italy's INVIND survey, we examine how firms update their beliefs about competitors' adoption of advanced technologies and whether these revised beliefs affect their own intended adoption. We document widespread underestimation of competitor adoption, substantial belief updating in response to information, and effects on intended robotics adoption—but no significant effects for AI. The effects on robotics adoption are robust to a wide range of checks and falsification tests.

To better assess the magnitude of these effects, and guided by a model of strategic interaction, we estimate a 2SLS specification that recovers the slope of the firm's reaction function. The steep reaction function highlights the strength of strategic complementarities in technology adoption, suggesting that firms' decisions are highly interdependent: a 1 pp increase in the share of competitors expected to adopt advanced technologies raises the firm's own intended robotics adoption by 0.732 pp. We also provide suggestive evidence that the competition channel plays an important role, as the effects are concentrated among firms with lower market shares.

We end with a discussion of the policy implications. In recent years, Italy—alongside other EU countries—has introduced several initiatives aimed at fostering innovation, with a particular focus on advanced technologies and digitalization. Following the pandemic, the EU approved a € 750 billion recovery fund for eligible national projects under the NextGenerationEU (NGEU) plan. Within this framework, Italy launched the *Piano Nazionale di Ripresa e Resilienza* (PNRR), its National Recovery and Resilience Plan, allocating € 191.5 billion in planned investments plus € 30.6 billion from a complementary national fund. The PNRR directs substantial resources toward modernizing public administration—including cloud computing, digital identity, and online public services—while also enhancing digital skills and offering tax incentives to encourage firm-level innovation. A specific set of subsidies, under the Transition 4.0 programme, is devoted to supporting firms that invest in advanced technologies or innovation.

One implication of our findings is that, because firms' adoption decisions are interdependent, the effects of financial incentives may extend well beyond the firms that receive them directly. When one firm's adoption affects the expectations and incentives of others, policies that induce some firms to adopt can generate substantial spillovers. Our results also point to an important role for information policy. As a complement to financial incentives, governments could use information campaigns to increase firms' awareness of both the productivity benefits of new technologies and the extent to which peer firms are adopting them. Although more research is needed to refine the design of such interventions and to evaluate their effectiveness, their low implementation cost suggests that they could be a highly cost-effective way to promote technological diffusion. Indeed, our heterogeneity analysis shows that the effects of our information treatment were strongest among firms that did not benefit from the Transition 4.0 programme. This suggests that information campaigns could broaden the reach of government intervention by influencing firms that are not directly affected by existing financial-incentive programs.

References

Abebe, G., S. Caria, P. Dupas, M. Fafchamps, and T. Getahun (2025). Competition and Management Upgrading: Experimental Evidence from Ethiopia. *NBER Working Paper No. 33886*.

Angeletos, G.-M., C. Hellwig, and A. Pavan (2007). Dynamic Global Games of Regime Change: Learning, Multiplicity, and the Timing of Attacks. *Econometrica 75*(3), 711–756.

Angeletos, G.-M. and A. Pavan (2004). Transparency of Information and Coordination in Economies with Investment Complementarities. *The American Economic Review 94*(2), 91–98.

Atkin, D., A. Chaudhry, S. Chaudry, A. K. Khandelwal, and E. Verhoogen (2017). Organizational Barriers to Technology Adoption: Evidence from Soccer-Ball Producers in Pakistan. *The Quarterly Journal of Economics 132*(3), 1101–1164.

Azoulay, P., J. S. Graff Zivin, D. Li, and B. N. Sampat (2019). Public R&D Investments and Private-Sector Patenting: Evidence from NIH Funding Rules. *The Review of Economic Studies 86*(1), 117–152.

Babina, T., A. Fedyk, A. He, and J. Hodson (2024). Artificial Intelligence, Firm Growth, and Product Innovation. *Journal of Financial Economics 151*, 103745.

Bank of Italy (2017). Methods and Sources: Methodological Notes. Survey of Industrial and Service Firms. Technical report, Bank of Italy.

Bank of Italy (2024). Annual Report for 2023. Technical report, Bank of Italy. https://www.bancaditalia.it/pubblicazioni/relazione-annuale/2023/index.html.

Bank of Italy (2025). Survey of Industrial and Service Firms in 2024. Technical report, Bank of Italy. https://www.bancaditalia.it/pubblicazioni/indagine-imprese/2024-indagini-imprese/en_statistiche_IIS_01072025.pdf?language_id=1.

Baum, C. F., M. E. Schaffer, and S. Stillman (2007). Enhanced Routines for Instrumental Variables/GMM Estimation and Testing. *The Stata Journal 7*(4), 465–506.

Baumann, U., Z. B. Cullen, E. Faia, A. Ferrando, R. Perez-Truglia, and J. Rariga (2026). Innovation Without Borders? The Geography of Technological Diffusion. *NBER Working Paper No. 35314*.

Bencivelli, L., S. Formai, E. Mattevi, and T. Padellini (2025). Embracing the Digital Transition: The Adoption of Cloud Computing and AI by Italian Firms. *Bank of Italy's Occasional Paper No. 946*.

Bilgin, N. M., E. Faia, and G. Ottaviano (2025). Technology Spillovers, Diffusion and Rivalry in Firm Networks. *CEPR Discussion Paper* (19804).

Bloom, N., M. Schankerman, and J. Van Reenen (2013). Identifying Technology Spillovers and Product Market Rivalry. *Econometrica 81*(4), 1347–1393.

Bottone, M. (2025). Comparing Survey Measures of Firms' Expectations and Uncertainty. *Empirical Economics 69*, 393–430.

Bulow, J. I., J. D. Geanakoplos, and P. D. Klemperer (1985). Multimarket Oligopoly: Strategic Substitutes and Complements. *Journal of Political Economy 93*(3), 488–511.

Cai, J. and A. Szeidl (2018). Interfirm Relationships and Business Performance. *The Quarterly Journal of Economics 133*(3), 1229–1282.

Caprara, D., L. Infante, M. Magnani, L. Modugno, and A. Neri (2024). Linking Macro and Microdata to Produce Distributional Accounts for Non-Financial Corporations. *Bank of Italy's Occasional Paper No. 846*.

Cardaliaguet, P., F. Delarue, J.-M. Lasry, and P.-L. Lions (2019). *The Master Equation and the Convergence Problem in Mean Field Games*. Princeton, NJ: Princeton University Press.

Cattaneo, M. D., R. K. Crump, M. H. Farrell, and Y. Feng (2024). On Binscatter. *American Economic Review 114*(5), 1488–1514.

Cavallo, A., G. Cruces, and R. Perez-Truglia (2017). Inflation Expectations, Learning, and Supermarket Prices: Evidence from Survey Experiments. *American Economic Journal: Macroeconomics 9*(3), 1–35.

Cingano, F., F. Manaresi, and E. Sette (2016). Does Credit Crunch Investment Down? New Evidence on the Real Effects of the Bank-Lending Channel. *The Review of Financial Studies 29*(10), 2737–2773.

Coibion, O. and Y. Gorodnichenko (2026). *Expectations Matter: The New Causal Macroeconomics of Surveys and Experiments*. Princeton, NJ: Princeton University Press.

Coibion, O., Y. Gorodnichenko, and S. Kumar (2018). How Do Firms Form Their Expectations? New Survey Evidence. *The American Economic Review 108*(9), 2671–2713.

Coibion, O., Y. Gorodnichenko, S. Kumar, and J. Ryngaert (2021). Do You Know That I Know That You Know...? Higher-Order Beliefs in Survey Data. *The Quarterly Journal of Economics 136*(3), 1387–1446.

Comin, D. and B. Hobijn (2010). An Exploration of Technology Diffusion. *The American Economic Review 100*(5), 2031–2059.

Cooper, R., D. V. DeJong, R. Forsythe, and T. W. Ross (1992). Communication in Coordination Games. *The Quarterly Journal of Economics 107*(2), 739–771.

Cooper, R. and A. John (1988). Coordinating Coordination Failures in Keynesian Models. *The Quarterly Journal of Economics 103*(3), 441–463.

Cooper, R. W., D. V. DeJong, R. Forsythe, and T. W. Ross (1990). Selection Criteria in Coordination Games: Some Experimental Results. *The American Economic Review 80*(1), 218–233.

Crawford, G. S. and M. Shum (2005). Uncertainty and Learning in Pharmaceutical Demand. *Econometrica 73*(4), 1137–1173.

Cullen, Z., S. Li, and R. Perez-Truglia (2026). What's My Employee Worth? The Effects of Salary Benchmarking. *Review of Economic Studies 93*(4), 2531–2573.

Cullen, Z. and R. Perez-Truglia (2022). How Much Does Your Boss Make? The Effects of Salary Comparisons. *Journal of Political Economy 130*(3), 766–822.

D'Aurizio, L. and G. Papadia (2016). Using External Sources to Understand Sample Survey Bias: The Case of the Bank of Italy's Survey of Industrial and Services Firms. *Bank of Italy's Occasional Paper No. 329*.

De Marco, F., J. Sauvagnat, and E. Sette (2021). Corporate Overconfidence and Bank Lending. *CEPR Discussion Paper No. 15785*.

Dürnecker, G., M. Ignaszak, J. Jeun, and L. Kaas (2026). Firm Expectations, Innovation and Growth. mimeo.

European Commission (2025). State of the Digital Decade: 2025 Report. Technical report, European Commission. https://digital-strategy.ec.europa.eu/en/library/state-digital-decade-2025-report.

Faia, E., M. Mayer, and V. Pezone (2026). The Value of Firm Networks: A Natural Experiment on Board Connections. *The Review of Corporate Finance Studies*.

Farboodi, M., R. Mihet, T. Philippon, and L. Veldkamp (2019). Big Data and Firm Dynamics. *AEA Papers & Proceedings 109*, 38–42.

Fuster, A., R. Perez-Truglia, M. Wiederholt, and B. Zafar (2022). Expectations with Endogenous Information Acquisition: An Experimental Investigation. *Review of Economics and Statistics 104*(5), 1059–1078.

Giaccobasso, M., B. Nathan, R. Perez-Truglia, and A. Zentner (2025). Where Do My Tax Dollars Go? Tax Morale Effects of Perceived Government Spending. *American Economic Journal: Applied Economics 17*(4), 223–259.

Goerg, S. J. and J. Kaiser (2009). Nonparametric Testing of Distributions—The Epps-Singleton Two-Sample Test Using the Empirical Characteristic Function. *The Stata Journal 9*(3), 454–465.

Guiso, L. and G. Parigi (1999). Investment and Demand Uncertainty. *The Quarterly Journal of Economics 114*(1), 185–227.

Gupta, A., J. Ponticelli, and A. Tesei (2024). Language Barriers, Technology Adoption and Productivity: Evidence from Agriculture in India. *The Review of Economics and Statistics*, 1–28.

Hill, R. and C. Stein (2025). Race to the Bottom: Competition and Quality in Science. *The Quarterly Journal of Economics 140*(2), 1111–1185.

HowToRobot (2023). Market Analysis of Robot and Automation Companies in Italy. Technical report, HowToRobot. https://automazione-plus.it/wp-content/uploads/sites/3/2023/07/Market-Report-Italy-2023-Free-Preview-compresso.pdf.

Huo, Z. and N. Takayama (2025). Rational Expectations Models with Higher-Order Beliefs. *The Review of Economic Studies 92*(5), 3138–3173.

Kalyani, A., N. Bloom, M. Carvalho, T. Hassan, J. Lerner, and A. Tahoun (2025). The Diffusion of New Technologies. *The Quarterly Journal of Economics 140*(2), 1299–1365.

Kim, H. (2025). The Value of Competitor Information: Evidence from a Field Experiment. *Management Science 71*(4), 3600–3621.

Lin, J. (2023). Strategic Complements or Substitutes? The Case of Adopting Health Information Technology by U.S. Hospitals. *The Review of Economics and Statistics 105*(5), 1237–1254.

Menkhoff, M. (2025). Belief Updating and AI Adoption: Experimental Evidence from Firms. *CESifo Working Paper No. 12291*.

Nagel, R. (1995). Unraveling in Guessing Games: An Experimental Study. *The American Economic Review 85*(5), 1313–1326.

Staiger, D. and J. H. Stock (1997). Instrumental Variables Regression with Weak Instruments. *Econometrica 65*(3), 557–586.

Stein, J. C. (2008). Conversations Among Competitors. *The American Economic Review 98*(5), 2150–2162.

Stock, J. H., J. H. Wright, and M. Yogo (2002). A Survey of Weak Instruments and Weak Identification in Generalized Method of Moments. *Journal of Business & Economic Statistics 20*(4), 518–529.

Weintraub, G. Y., C. L. Benkard, and B. Van Roy (2008). Markov Perfect Industry Dynamics with Many Firms. *Econometrica 76*(6), 1375–1411.

Figure 1: Intensity and Evolution of Usage of Advanced Technologies

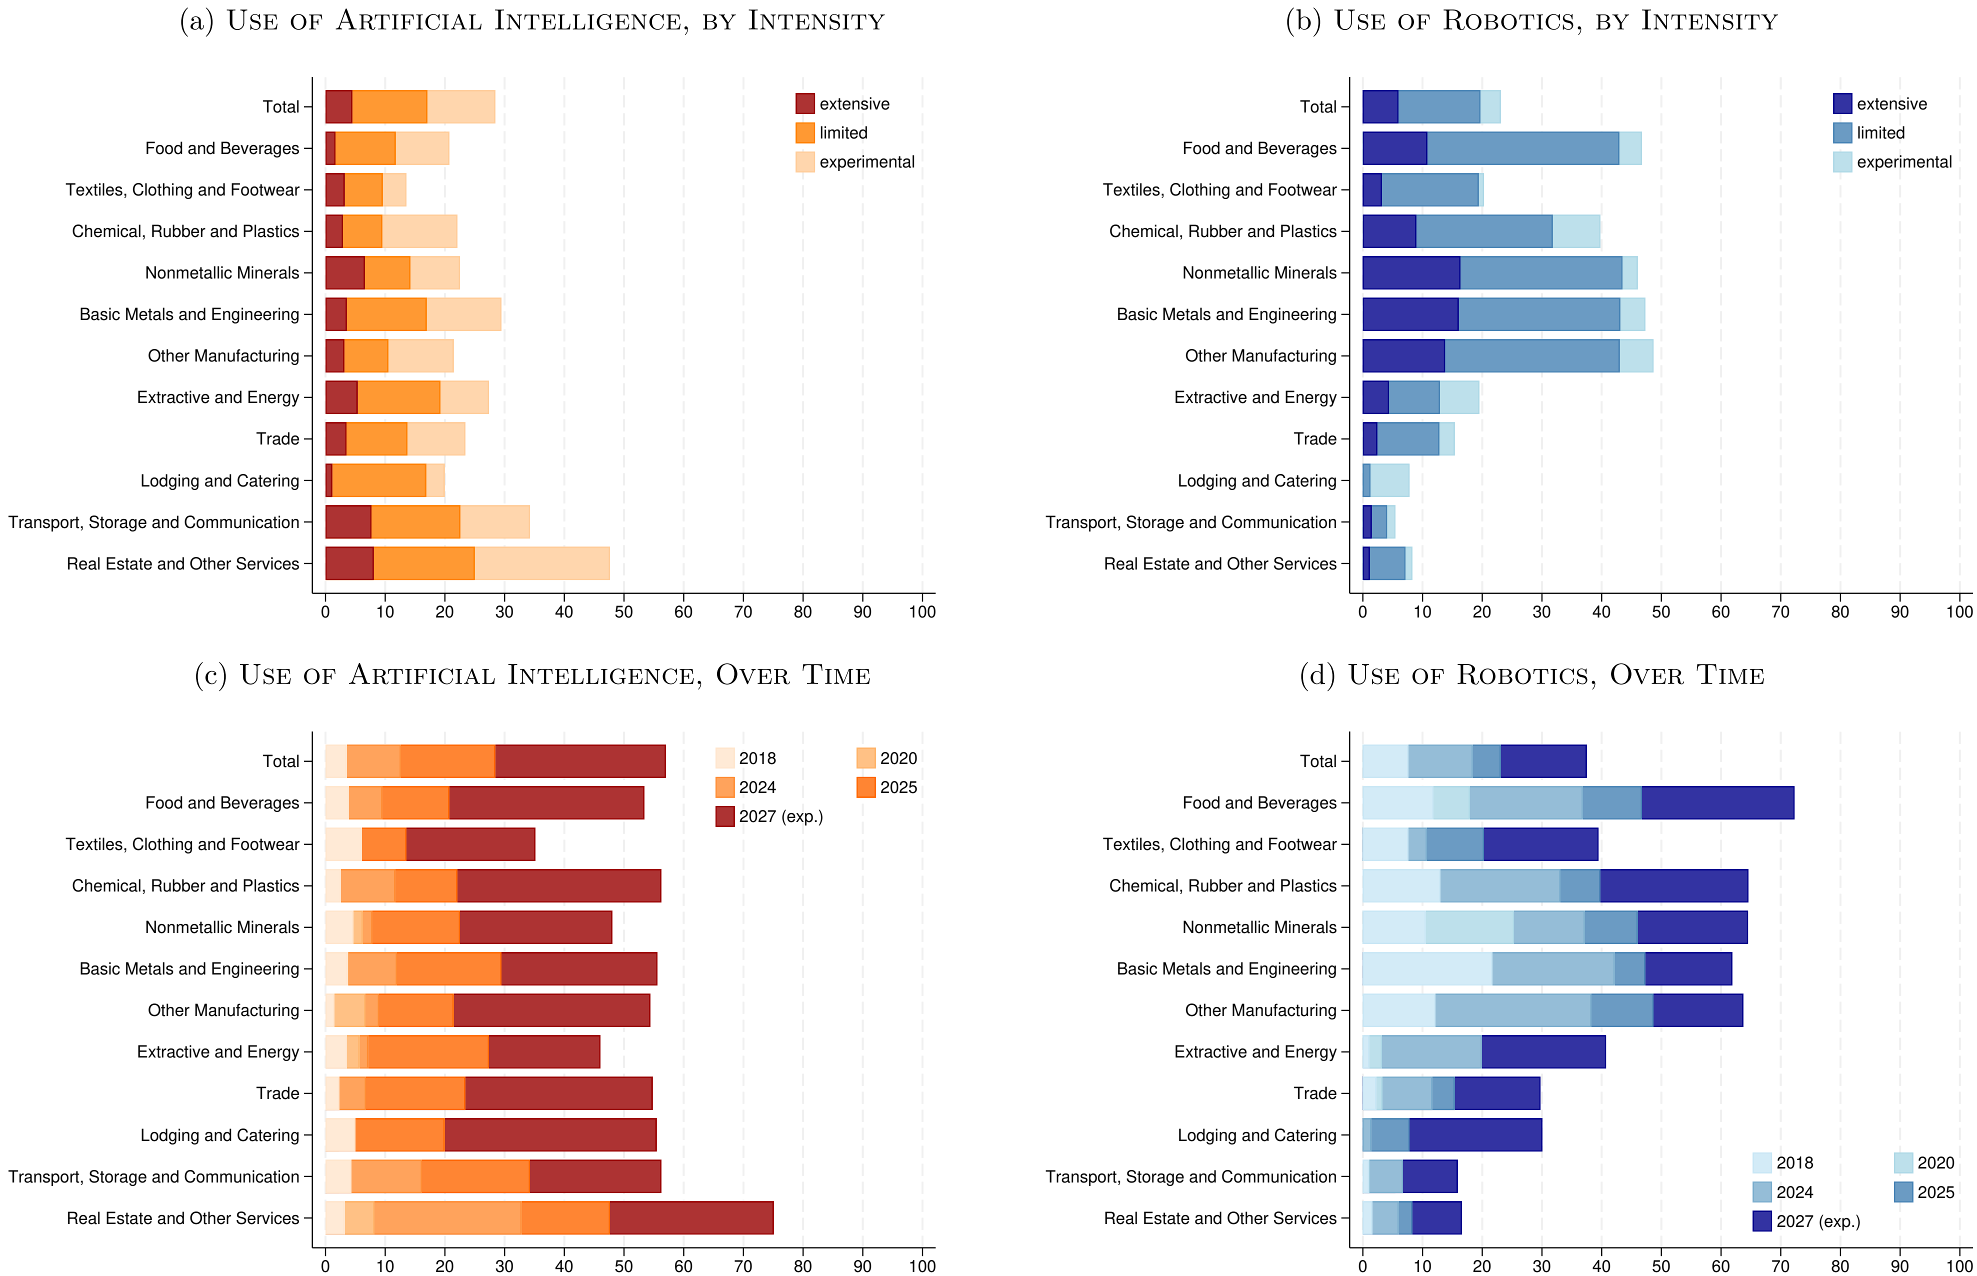

Notes: Panels (a) and (b) use the 2025 INVIND survey and report the intensity of predictive or generative AI use (panel (a)) and robotics use (panel (b)) at the time of the interview (questions TEC5N1, TEC5N2, and TEC11N). Panels (c) and (d) use the 2018, 2020, 2024, and 2025 survey waves and report the share of firms already using AI or robotics in those years. The 2024 figures shown here may differ from the information treatments described in Table C.1 because the latter also include firms that were not yet using these technologies but expected to adopt them by the end of 2024. The bars for expected use in 2027 are based on questions TEC27A, TEC27B, and TEC27C in the 2025 survey wave and are computed using only control-group firms.

Figure 2: Distribution of Prior Beliefs and Belief Updating

(a) Raw Perception Gap (b) Expected Adoption by Competitors

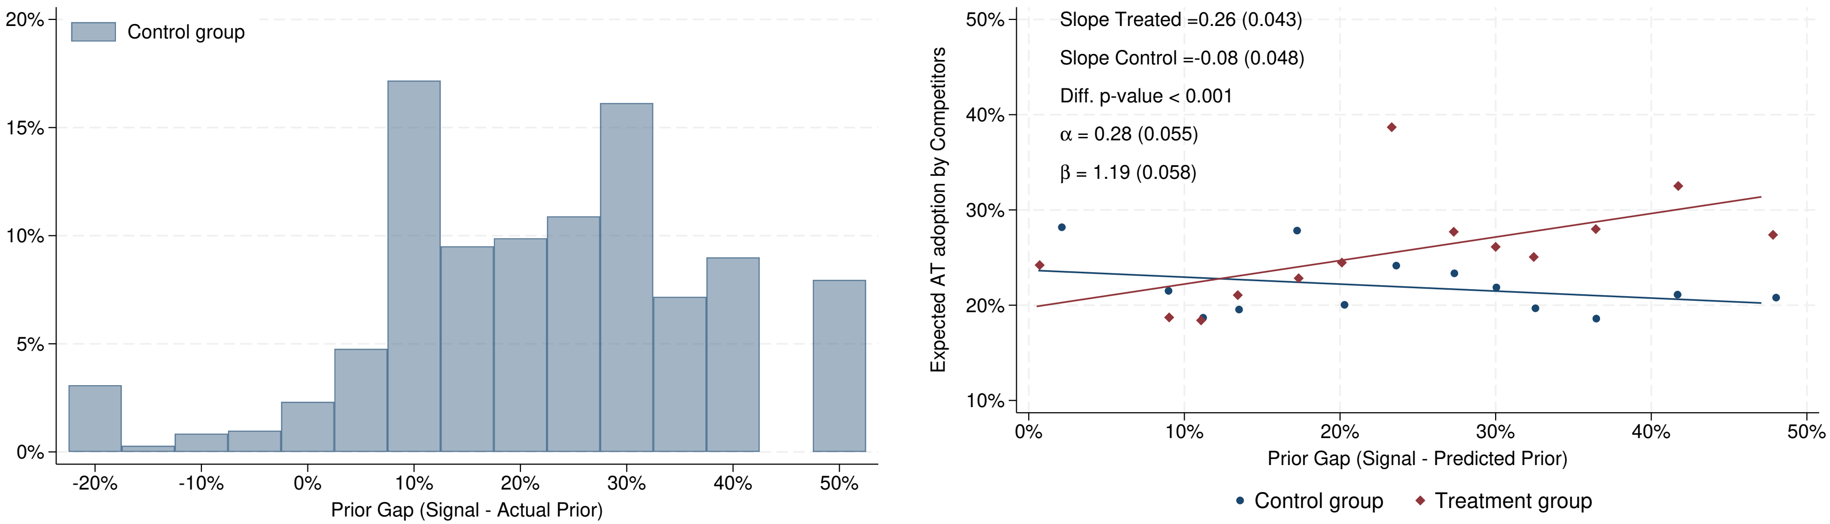

Notes: The figure is based on the INVIND survey conducted in 2025. Panel (a) shows, for the control group, the empirical distribution of the gap between prior beliefs and the information that firms would have seen if treated. The first and last bins group observations whose gaps are below -15 and above 45, respectively. Panel (b) displays a binscatter of posterior beliefs about competitors' adoption of advanced technologies (questions TEC26A and TEC26B in the survey module) against the imputed gap between prior beliefs and the information treatment, controlling for the level of prior beliefs. α indicates the learning weight and β the degree of extrapolation from the information received, as obtained in equation (3). Standard errors are reported in parentheses. The p-value for the difference between the two groups is obtained using the procedure described in Cattaneo et al. (2024).

Figure 3: Treatment Effects on Posterior Beliefs, Intention to Adopt Advanced Technologies by 2027, with Falsification Tests

(a) Treatment Effects

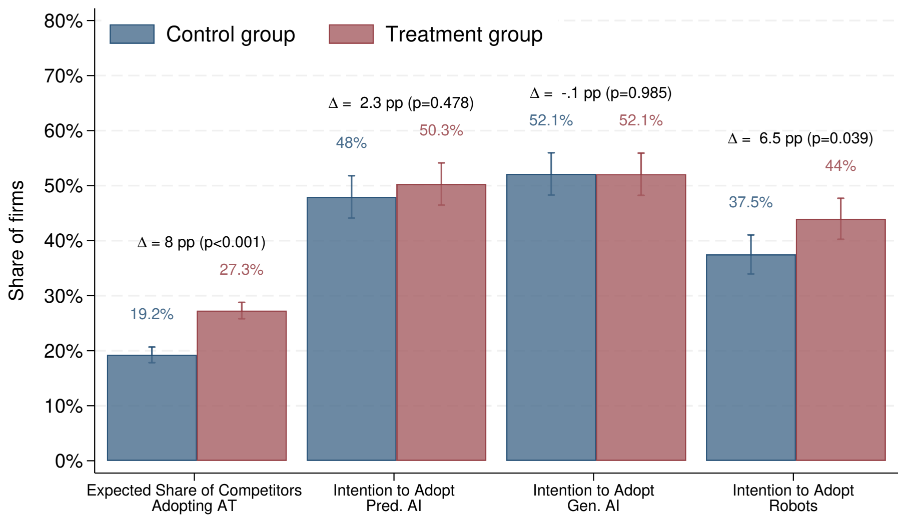

(b) Falsification Tests

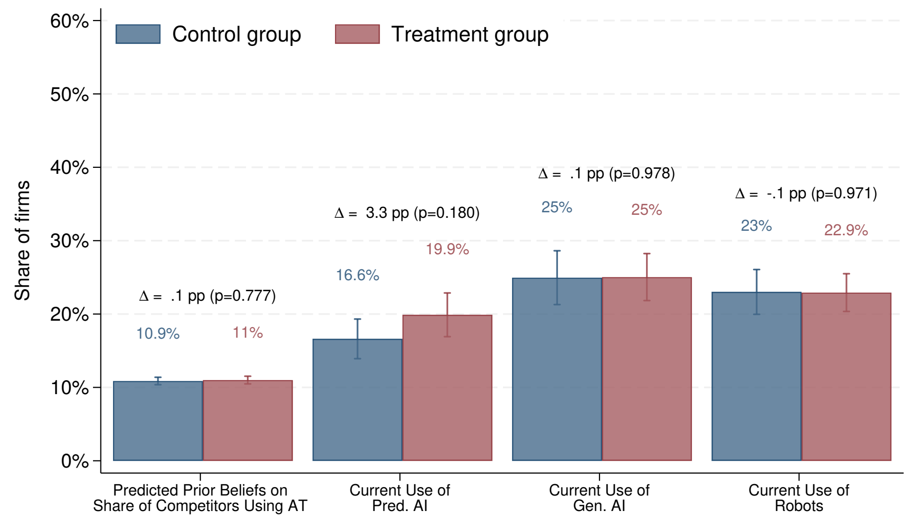

Notes: The figure is based on the INVIND survey conducted in 2025. Panel (a) presents posterior beliefs about competitors' future adoption of advanced technologies (questions TEC26A and TEC26B in the survey module) and intentions to adopt advanced technologies by 2027 (questions TEC27A, TEC27B, and TEC27C). Panel (b) presents imputed prior beliefs about competitors' current use of advanced technologies (question TEC24 in the survey module) and firms' current use of these technologies (questions TEC5N1, TEC5N2, and TEC11N). Imputed prior beliefs are based on the estimates presented in Table C.2. Labels above the bars report estimated percentages. For each outcome, the figure reports the difference between treated and control firms, expressed in percentage points, together with the p-value from a test of equality of means. The figure reports 90% confidence intervals.

Figure 4: Treatment Effects on Posterior Beliefs and Intention to Adopt Advanced Technologies by 2027, by Prior Gaps

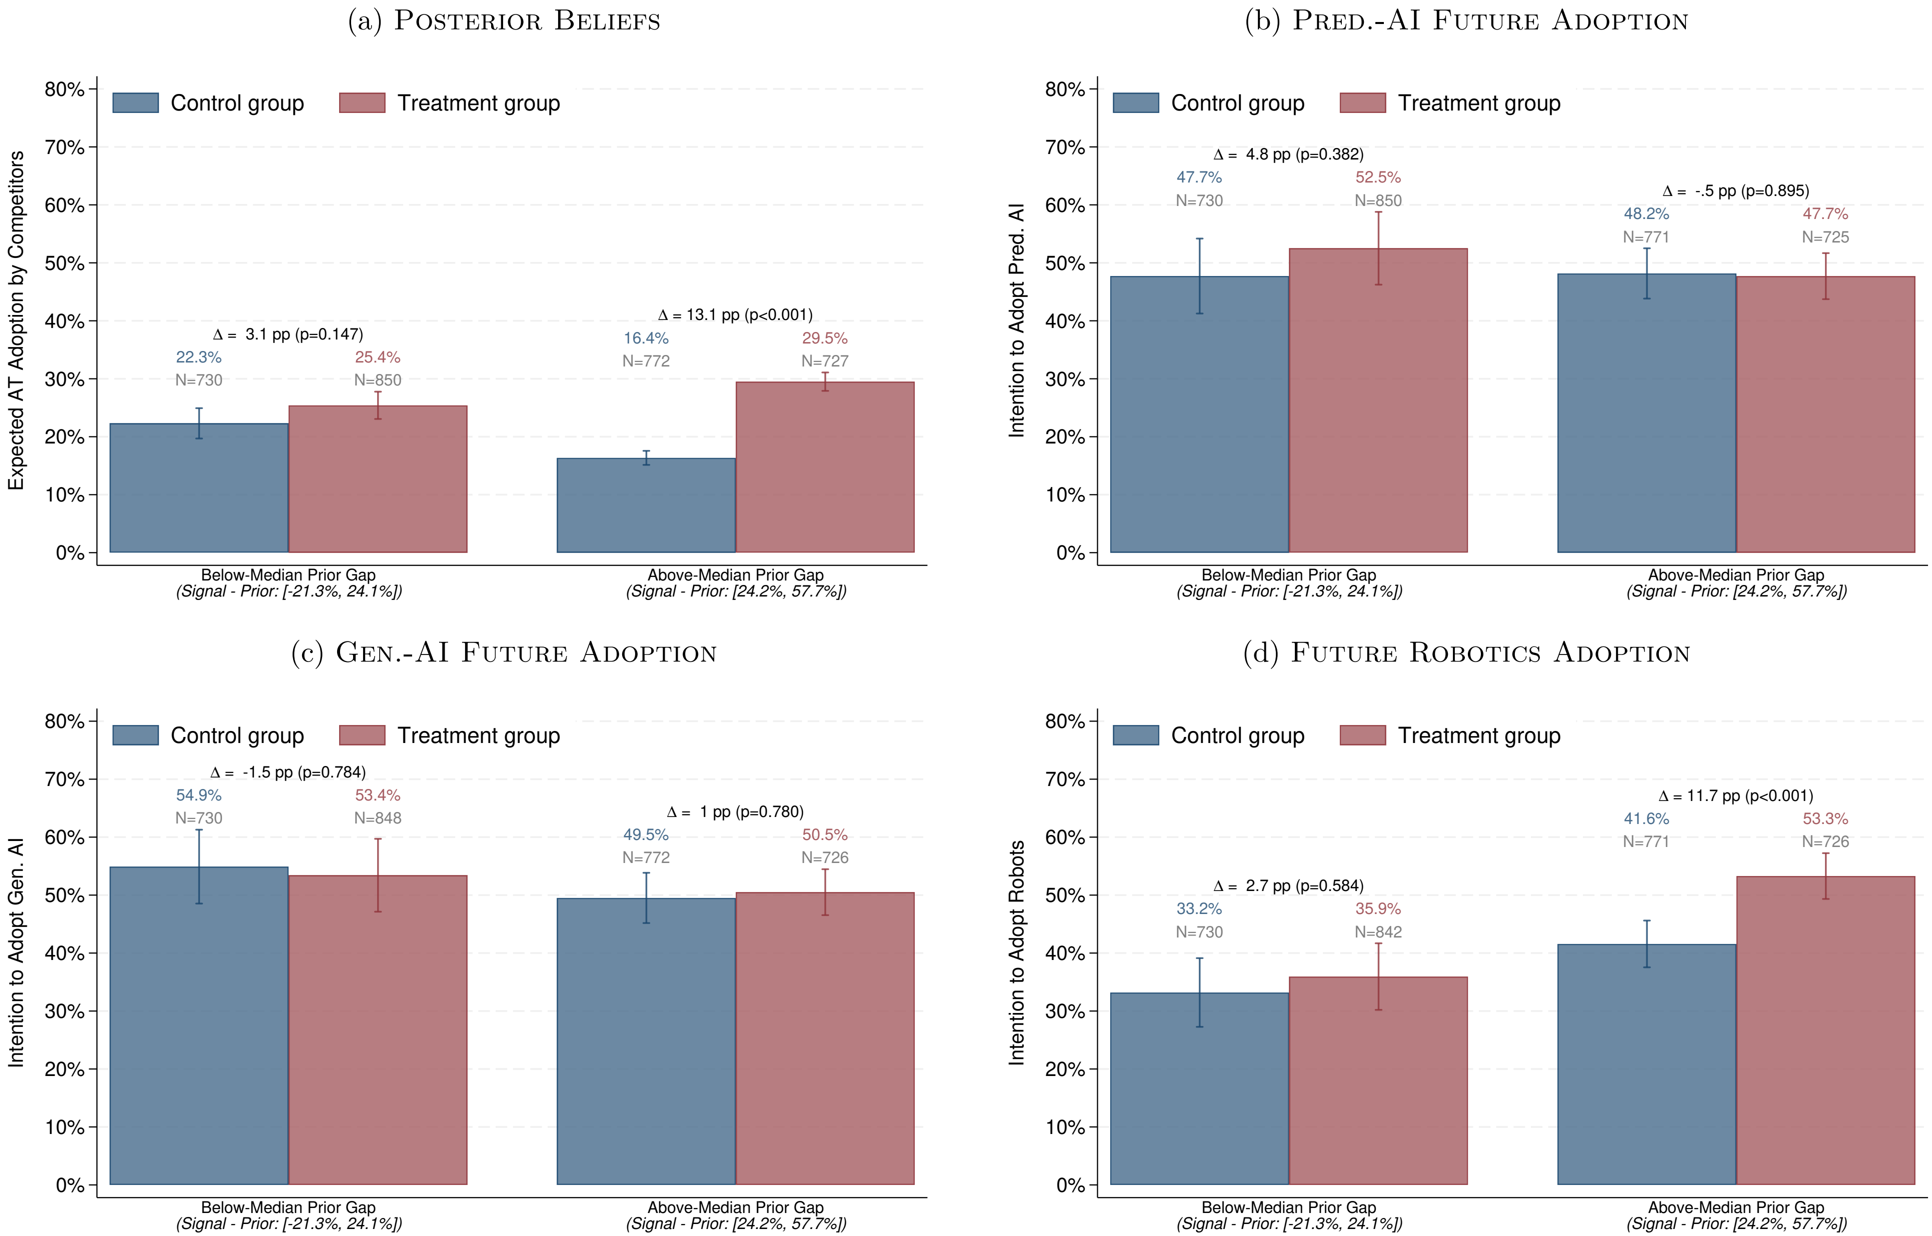

Notes: The figure is based on the INVIND survey conducted in 2025. The graphs present posterior beliefs about competitors' future adoption of advanced technologies (questions TEC26A and TEC26B in the survey module) and intentions to adopt advanced technologies by 2027 (questions TEC27A, TEC27B, and TEC27C), by imputed prior gap (below or above the median) and treatment group. The imputed prior gap is based on the imputed prior, using the estimates presented in Table C.2. Labels above the bars report estimated percentages and the corresponding number of observations. For each outcome, the figure reports the difference between treated and control firms, expressed in percentage points, together with the p-value from a test of equality of means. The figure reports 90% confidence intervals.

Figure 5: Falsification Tests: Treatment Effects on Prior Beliefs and Current Adoption, by Prior Gap

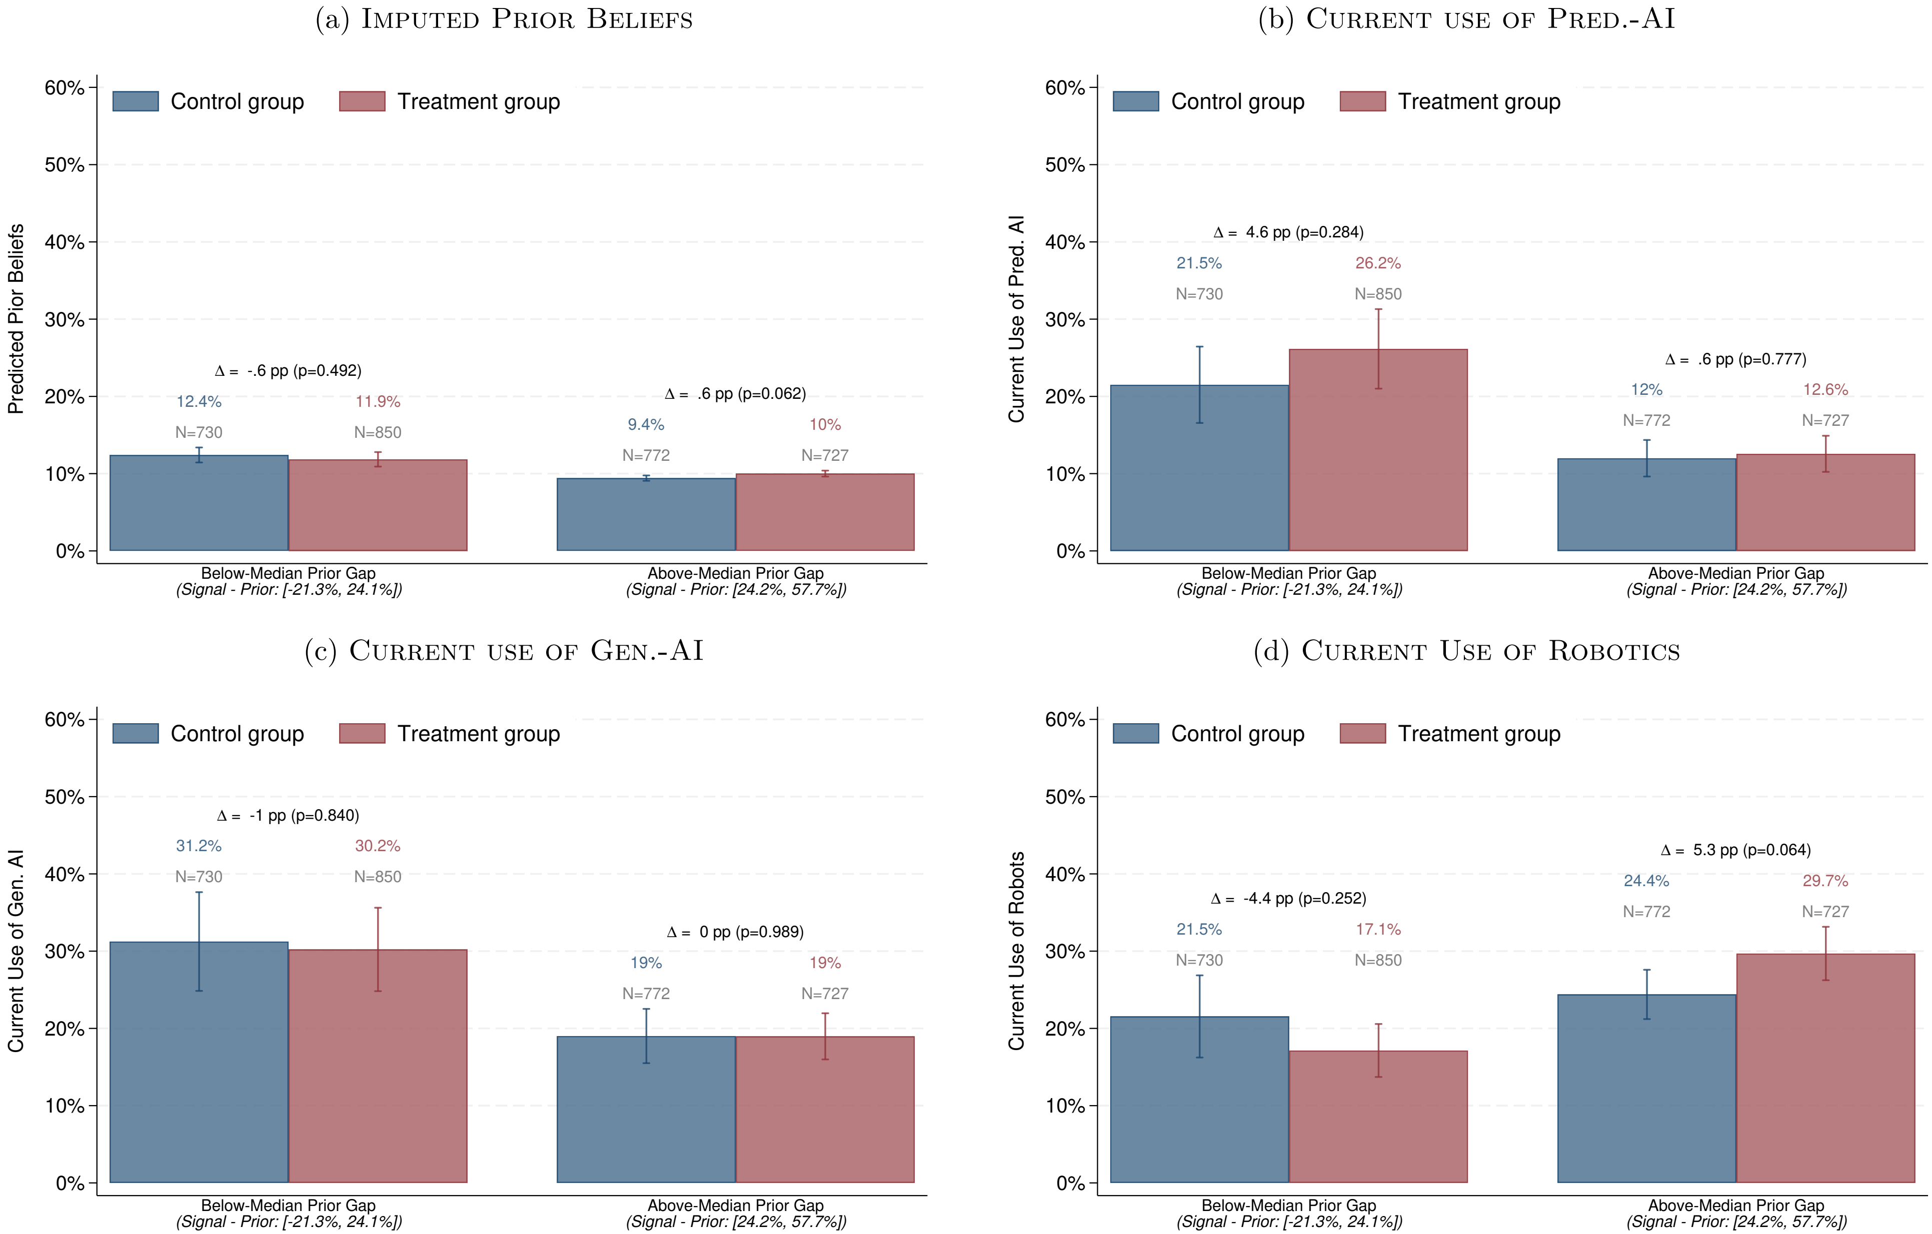

Notes: The figure is based on the INVIND survey conducted in 2025. The graphs present imputed prior beliefs about competitors' current use of advanced technologies (question TEC24 in the survey module) and firms' current use of these technologies (questions TEC5N1, TEC5N2, and TEC11N), by imputed prior gap (below or above the median) and treatment group. The imputed prior gap is based on the imputed prior, using the estimates presented in Table C.2. Labels above the bars report estimated percentages and the corresponding number of observations. For each outcome, the figure reports the difference between treated and control firms, expressed in percentage points, together with the p-value from a test of equality of means. The figure reports 90% confidence intervals.

Figure 6: Relationship between Ex Ante Intention to Adopt Robotics (2025) and Ex Post Adoption (2026)

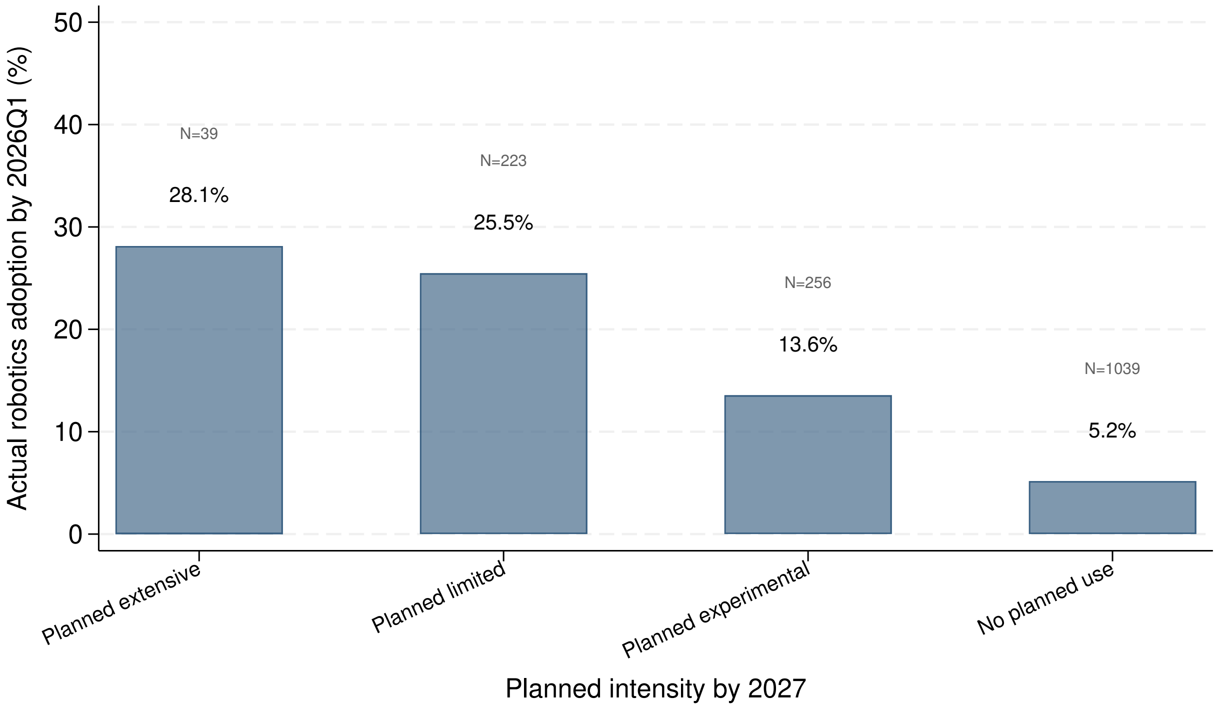

Notes: The sample includes firms surveyed in both the 2025 and 2026 INVIND waves. The figure shows the proportion of firms that had not yet adopted robotics in 2025 but reported having already adopted it by the time of the 2026 interview. Firms are grouped along the x-axis according to their planned intensity of robotics adoption in 2027, as reported in the 2025 wave.

Figure 7: Distribution of Learning, Competition, and their Gap

(a) Learning and Competition (b) Gap

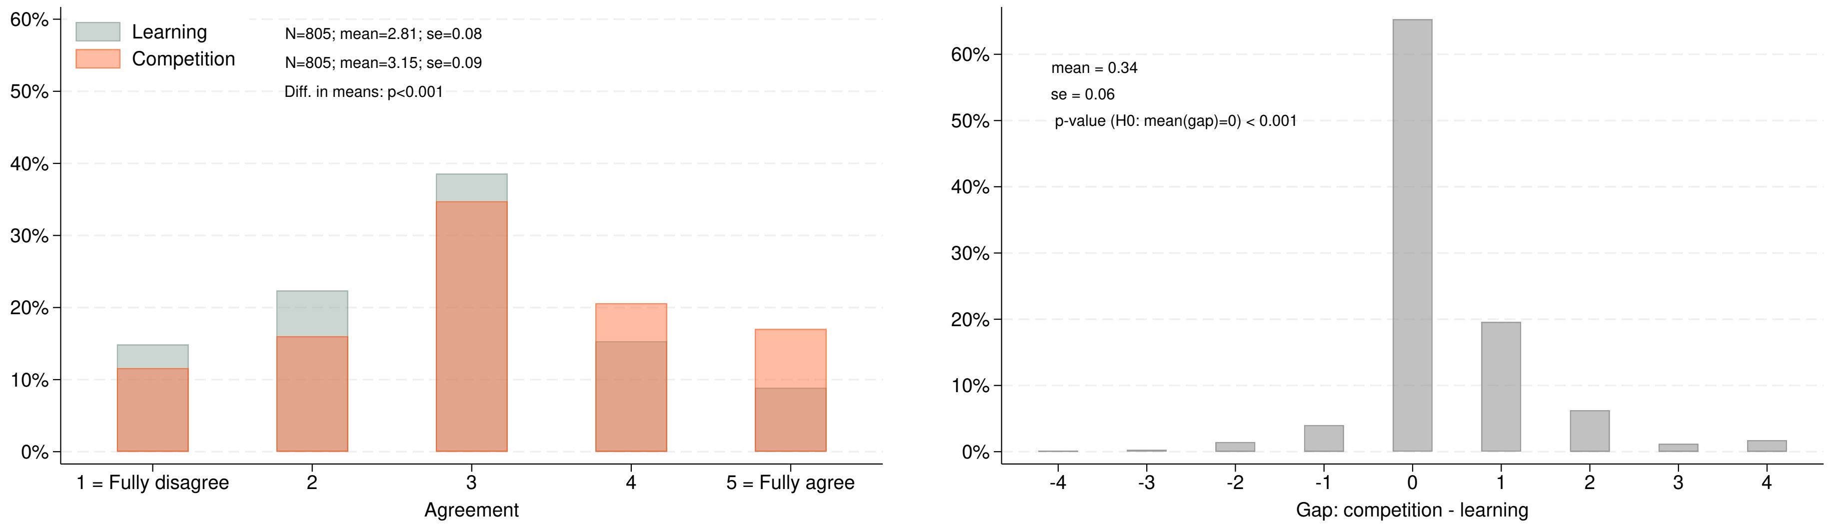

Notes: The figure is based on the 2026 INVIND survey and a subset of firms that were also interviewed in the previous year and responded to the AI module and questions TEC28AA, TEC28AB, TEC28BA, and TEC28BB. Panel (a) shows the empirical distributions of responses to the learning and competition questions. Answers range from 1 to 5, where 1 means fully disagree and 5 means fully agree. Panel (b) shows the distribution of the gap between the score assigned to the competition question and the score assigned to the learning question.

Table 1: Summary Statistics and Randomization Balance

	All (1)	Control (2)	Treatment (3)	p-value (4)
Firm Age	40.389	40.293	40.480	0.884
	(0.383)	(0.536)	(0.547)	
Number of Employees in 2024	96.005	96.071	95.941	0.986
	(11.136)	(11.849)	(18.647)	
Geographical Area: North (%)	57.289	57.382	57.201	0.954
	(0.892)	(1.270)	(1.253)	
Geographical Area: Center (%)	21.437	21.139	21.722	0.829
	(0.740)	(1.048)	(1.044)	
Geographical Area: South and Islands (%)	21.273	21.479	21.077	0.847
	(0.738)	(1.054)	(1.033)	
Manufacturing Firms (%)	39.281	38.021	40.480	0.397
	(0.880)	(1.246)	(1.243)	
Turnover (Million euros)	42.718	44.337	41.176	0.510
	(5.154)	(7.132)	(7.431)	
Share of Investment on Turnover	0.067	0.064	0.071	0.528
	(0.007)	(0.005)	(0.014)	
Exporters (%)	52.927	53.045	52.814	0.945
	(0.900)	(1.281)	(1.264)	
Current Use of Pred.-AI (%)	18.294	16.615	19.893	0.180
	(0.697)	(0.956)	(1.011)	
Current Use of Gen.-AI (%)	24.997	24.956	25.037	0.978
	(0.780)	(1.111)	(1.097)	
Current Use of Robotics (%)	22.968	23.013	22.924	0.971
	(0.758)	(1.081)	(1.064)	
Share of Investment in AT	6.740	6.476	6.992	0.531
	(0.260)	(0.361)	(0.374)	
Number of People Involved in the Survey	2.127	2.138	2.116	0.756
	(0.025)	(0.035)	(0.037)	
Manager Involved (%)	54.011	53.494	54.503	0.760
	(0.898)	(1.281)	(1.261)	
Transition 4.0 Programme (%)	35.831	35.062	36.563	0.625
	(0.864)	(1.225)	(1.219)	
National Recovery and Resilience Plan (%)	12.894	14.996	10.891	0.064
	(0.604)	(0.917)	(0.789)	
Observations	3,079	1,518	1,561	

Notes: The table presents summary statistics for firm characteristics, location, and survey responses from the INVIND survey conducted in 2025. Column (1) reports the full sample, columns (2) and (3) the control and treatment groups, and column (4) the p-value from an equality-of-means test. Standard errors are reported in parentheses.

Table 2: 2SLS Estimates: Effects of Expected Competitors' Adoption on Own Future Adoption

				Robotics			
				Previous use		Manager present	
	Pred.-AI	Gen.-AI	Robotics	No	Yes	No	Yes
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel (a): 2SLS Regressions							
Exp. AT Adoption by Competitors	0.015	0.060	0.732***	1.070***	0.170	0.230	1.223***
	(0.233)	(0.225)	(0.209)	(0.283)	(0.157)	(0.281)	(0.357)
Panel (b): 2SLS: First Stage							
$T_i *$ Prior gap	0.343***	0.342***	0.341***	0.344***	0.348***	0.397***	0.290***
	(0.028)	(0.028)	(0.028)	(0.035)	(0.049)	(0.041)	(0.041)
Panel (c): Reduced-Form Regressions							
$T_i *$ Prior gap	0.005	0.021	0.250***	0.368***	0.059	0.091	0.355***
	(0.079)	(0.076)	(0.070)	(0.094)	(0.052)	(0.107)	(0.098)
Observations	3,075	3,077	3,073	1,949	1,124	1,217	1,749
Dep. Var.: Baseline Mean	48.0	52.0	37.5	19.6	97.3	34.0	38.1
Kleibergen-Paap Wald F-statistic	134.7	132.8	132.6	87.4	46.7	96.6	48.5

Notes: Panel (a) presents the results of the 2SLS model specified in equation (5) and discussed in Section 6.3. The dependent variable is a dummy equal to 100 if firm i intends to use predictive AI, generative AI, or robotics by 2027. Panel (b) presents the first-stage results associated with the 2SLS model, i.e., the effect of the variable $T_i \cdot (s^T_{i,t} - s^0_{i,t})$ on posterior beliefs about competitors' future adoption of advanced technologies ($s_{i,t+1}$). Panel (c) presents the corresponding reduced-form results, capturing the effect of $T_i \cdot (s^T_{i,t} - s^0_{i,t})$ when it is not used to instrument for posterior beliefs. Columns (4) and (5) split the sample according to whether the firm reported using robotics at the time of the interview. Columns (6) and (7) split the sample according to whether a manager was involved in completing the questionnaire. All standard errors are bootstrapped.

Table 3: Heterogeneity by Market Concentration and Public Incentives: Effects of Expected Competitors' Adoption on Own Future Robotics Adoption (2SLS)

		Market Share		Herfindahl-Hirschman Index		Public Incentives	
	Total	Below Median	Above Median	Below Median	Above Median	No	Yes
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel (a): 2SLS Regressions							
Exp. AT Adoption by Competitors	0.732***	1.264***	0.112	0.579**	0.539	0.918***	0.420
	(0.209)	(0.358)	(0.236)	(0.233)	(0.396)	(0.284)	(0.276)
Panel (b): 2SLS: First Stage							
$T_i *$ Prior gap	0.341***	0.327***	0.327***	0.273***	0.375***	0.335***	0.334***
	(0.028)	(0.048)	(0.033)	(0.034)	(0.054)	(0.039)	(0.036)
Panel (c): Reduced-Form Regressions							
$T_i *$ Prior gap	0.250***	0.413***	0.037	0.158**	0.202	0.307***	0.140
	(0.070)	(0.116)	(0.078)	(0.080)	(0.138)	(0.101)	(0.091)
Observations	3,073	908	2,165	1,130	1,943	1,594	1,479
Dep. Var.: Baseline Mean	37.5	31.4	43.4	30.9	44.3	26.1	59.3
Het. Var. Mean		0.0018	0.0153	0.0017	0.0070		
Kleibergen-Paap Wald F-statistic	132.6	46.3	86.4	27.9	121.2	68.2	81.4

Notes: The table presents the results of the 2SLS model specified in equation (5) and discussed in Section 6.3. The dependent variable is a dummy equal to 100 if firm i intends to use robotics by 2027. Market shares and HHI are based on balance-sheet data from CADS. Column (1) replicates column (3) of Table 2. Columns (2) and (3) split the sample according to whether the firm's market share in its sector lies below or above the weighted median across sectors. Columns (4) and (5) split the sample according to whether the Herfindahl-Hirschman Index (HHI) of the firm's sector lies below or above the median. The HHI of a given sector j is defined as $HHI_j = \sum_{i=1}^{N_j} m_{ij}^2$, where N_j is the number of firms operating in sector j and m_{ij} is the market share of firm i in sector j. Columns (6) and (7) split the sample according to whether the firm used or expected to use by the end of 2025 the tax credit for capital goods under the Transition 4.0 programme. In columns (2)–(5), the sectoral classification follows the same taxonomy reported in Table C.1. All standard errors are bootstrapped.