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

The views expressed herein do not reflect those of the Bank of Italy. We would like to thank Marios Angeletos, Andrea Neri, Alessandro Notarpietro, Alessandro Pavan, Tiziano Ropele, Alfosno Rosolia, Fabrizio Venditti, and Roberta Zizza for their useful comments and suggestions as well as seminar participants. This project was reviewed and approved in advance by the Institutional Review Board at the Goethe University Frankfurt. It is also registered at the AEA registry under AEARCTR-0016177. After the study is accepted for publication, we will share all the code through a public repository as well as instructions on how to request the data. Yuerong Zhuang provided excellent research assistance. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

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
JEL No. C93, D22, L21, O33

ABSTRACT

We present the first large-scale field experiment test of 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 accurate information about adoption rates. Most firms substantially underestimated competitors' current adoption, and when provided with information, they updated their expectations about competitors' future adoption. The information increased firms' own intended future adoption of robotics, although we do not observe a significant effect on AI adoption. 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

A firm's decision to innovate, adopt new technologies, or set prices depends not only on internal factors but also on the perceived behavior of its competitors. 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 what is, to our knowledge, the first 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 in that for years it has been tracking the diffusion of advanced technologies, which we need 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 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 information about actual adoption rates of their competitors—firms in their same sector and size class—to half of the sample, randomly selected. 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' innovation, ruling out spurious correlations that plague observational studies. Third, the ability to re-survey the same firms a year later—and to link the INVIND data to administrative records such as VAT and customs data—may allow future iterations of the study to assess whether effects on reported intentions ultimately translate into realized behavior.

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 27% of firms used AI in some form, with even higher rates in certain sectors such as real estate services. By comparison, 23% of firms used robotic technologies in 2025. While both 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 derive our econometric design based on a simple conceptual framework of strategic interaction in competition under incomplete information (along the lines of ofAngeletos and Huo (2021) or Huo and Takayama (2024)). Firms, uncertain about the productivity state, learn about it from the adoption of their competitors. The framework leads to a belief update reduced form, which we take to the data. Furthermore, since firms compete with each other their innovation policy reacts to the one of the competitors to avoid falling behind. The equilibrium action from the model leads to a second stage reduced form, which reflects the total causal effects of beliefs on policies. The coefficient estimate of this stage capture the composite motive of learning adjustment and competitive motive. To gauge more insights on the role of competition, we rely on the matching with firms' balance sheet data and interact the causal effect of beliefs in the second stage with a dummy for high and low market share, mark-ups or concentration indeces.

From the estimates of the first stage, we document substantial differences between the beliefs about the adoption rates of competitors and the information we provided. On average, prior beliefs underestimated actual adoption by 24.6 pp. Moreover, a strong majority of firms underestimated competitors' adoption, although some firms underestimated by a much larger margin than others.[1] We observe that they updated their beliefs significantly in response to the provision of information, meaning that subjects paid attention to the information. Consistent with Bayesian learning, the direction and strength of the belief updating is a function of the direction and strength of the prior misperceptions. Given significant misperceptions, coupled with significant belief updating toward the signal, in aggregate firms learn from the signal.

Moreover, we observe significant effects of the information on the firms' own stated adoption plans about robotics. On average, the information treatment increased the intention to adopt robotics technology by 6.7 pp. This average effect is not only statistically significant (p=0.029), but also large in magnitude, corresponding to a 17.8% 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 on firms who underestimated the most the competitor's adoption. Second, we show no effects on pre-treatment outcomes, such as the current adoption level, elicited before the information-provision stage.

Using a Two-Stages-Least-Squares (2SLS) model, we estimate the firms' reaction function:

[1] We find that the prior beliefs of treated firms appear more accurate than those of control firms. We interpret this as evidence of possible contamination, as treated respondents may have revised their initial answers after being exposed to the information. To address this concern, we impute prior beliefs using an econometric model. If anything, this procedure likely adds some noise to our estimates, thereby attenuating the results.

a 1 pp increase in the share of competitors expected to adopt advanced technologies causes an increase of 0.704 pp in the firm's own adoption probability of robotics. The steepness of the reaction function highlights the strength of strategic complementarities in technology adoption, suggesting that firms' decisions are highly interdependent. Moreover, this result underscores that policies or shocks influencing a subset of firms could have large aggregate effects through strategic spillovers in adoption decisions.

In contrast to the effects on the adoption plans for robotics, we do not find any significant effects for AI adoption. This lack of effect is true both in the entire sample and when looking at the subsample of firms who underestimated beliefs the most. While we cannot rule out modest effects on AI, they appear weaker than those on robotics adoption. This difference could be partly explained by higher baseline intentions to adopt AI, which leave less scope for increases. Another likely factor is that robotics is a more mature and established technology, so learning that competitors are adopting it may trigger a stronger competitive response. In contrast, AI adoption is newer and often experimental, leading firms to react more cautiously to information about peers' adoption.

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.[2] Our evidence on the presence of significant information frictions suggests an additional policy lever that could complement financial incentives. 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.

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). Subsequent research has made substantial advances, including extensions to multi-player contexts (Angeletos and Pavan, 2004), the incorporation of dynamic settings (Weintraub et al., 2008), and the development

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

of mean-field game frameworks (Cardaliaguet et al., 2019). More recent work has introduced elements of bounded rationality into these models (Angeletos and Huo, 2021; Huo and Takayama, 2024).

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 technologies, 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.

Our study is also related to the broader literature on innovation and the adoption of new technologies. For example, Comin and Hobijn (2010) document large cross-country and cross-sector differences in the speed of technology diffusion. Bencivelli et al. (2025) analyze firm-level determinants of cloud and AI adoption in Italy. Hill and Stein (2025) study how competition to publish first shapes scientists' research decisions. Kalyani et al. (2025) discuss the spread of new technologies across firms, industries, and countries. And Atkin et al. (2017) study barriers to technology adoption among manufacturers in Pakistan. We contribute to this literature by identifying a novel mechanism: firms systematically underestimate competitors' adoption rates, and correcting these beliefs can meaningfully alter their investment plans. Thus, our experimental evidence highlights that information frictions can play a central role in shaping the adoption of new technologies.[3]

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.

[3] The role of information frictions is related to Gupta et al. (2020), who show that access to call centers for agricultural advice in rural India can improve agricultural productivity.

Section 4 discusses results on competitor's adoption, misperceptions and belief updating. Section 5 presents the effects of the information on the firm's 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.[4] On the other hand, AI adoption is growing among Italian firms, but still remains below the European Union average both in manufacturing and services.[5] At last, Italy is currently implementing a plan of subsidies for firms investing in advanced technology—the National Recovery and Resilience Plan, which in turn is funded through the EU program 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 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

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

[5] 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).

adoption of advanced technologies—allowing us to construct the information necessary for our intervention.[6]

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 or assisted by 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).[7]

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 A 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.[8] The survey data also include sampling weights to account for selection probabilities. 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.

While this version of the study includes data from the 2025 survey wave only, we may be able to collect additional data in a future version of the study. The panel component of the survey may allow us to include questions in the next wave (2026) to assess whether the effects of the information treatment—if any—persist a year later.[9] Moreover, when data for the post-

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

[7]Although the survey does not explicitly identify the respondent's role, one question indicates that 53% of firms reported that a company manager, or someone working closely with one, assisted in completing the questionnaire.

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

[9]This type of analysis may be feasible given that in the past 10 years, 83% of firms have participated in the survey for at least two consecutive years.

treatment period become available, we may be able to link the survey data to administrative sources (such as value-added tax records, customs data, and balance-sheet information) to estimate treatment effects on additional outcomes and to conduct heterogeneity analyses. For instance, we could use customs and VAT data to test whether the information treatment increases the probability of purchasing from firms that sell robots.

3 Research Design

3.1 Survey Design

The survey instrument for our tailored module in the 2025 wave is provided in Appendix D and described below.[10]

We focus on the adoption of the two most important types of automation technologies, which have experienced rapid growth in recent years: AI and robotics. The 2025 module begins by providing definitions for predictive AI (Pred.-AI), generative AI (Gen.-AI), and robotics. We ask some questions separately for Pred.-AI and Gen.-AI, rather than combining them into a single question on AI, because one of these two types—Gen.-AI—is more recent and expanding rapidly, so we wanted to be able 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.[11]

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.).

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

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

Right after eliciting the prior belief, half of the firms were randomly assigned to receive information about the share of competitors that had invested in AI. We refer to this information as the *signal*. To calculate the signal, we used responses from the previous wave of INVIND. For reference, Appendix E includes the survey module on advanced technologies from the 2024 wave. We calculated the share of firms that reported having already adopted AI or robotics, or that planned to do so by the end of 2024.[12] 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.[13] For size, we split firms into those with fewer than 50 employees and those with 50 or more.[14] 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 B.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.[15] 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

[12]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.

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

[14]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.

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

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.

3.2 Survey Implementation

A total of 3,983 firms participated in the 2025 survey wave. Since the questionnaire is relatively long, not all firms completed every module. The questions on AI and robotics were placed in the central part of the survey and were not mandatory. As a result, about 900 firms did not answer the key questions from our module (e.g., those on baseline adoption and prior beliefs). We are therefore left with a sample of approximately 3,080 firms. For consistency, unless explicitly stated otherwise, all subsequent analyses are conducted on this sample of respondents.

3.3 Descriptive Statistics and Randomization Balance

Column (1) of Table 1 reports descriptive information about the subject pool. 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 38% 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, 16.7% had already adopted Pred.-AI, 25.0% had adopted Gen.-AI, and 23.1% had adopted robotics.

3.4 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).[16] 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 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 B.

This does not mean that the prior-belief question is useless. Such responses are typically

[16] 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.

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.[17] These imputed priors are then used to perform the heterogeneity analysis. The results of this imputation method, presented in Appendix B, show that the imputed prior beliefs, as intended, are statistically indistinguishable between the treatment and control groups.[18]

4 Misperceptions and Belief Updating about Competitors' Adoption

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.[19] 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.[20] Panel (a) shows current usage in 2025,

[17]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%.

[18]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.

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

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

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 intensity is observed in Textiles, Clothing, and Footwear (13%), and the highest in Real Estate and Other Services (48%).

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.[21] These AI adoption rates measured in our survey are broadly consistent with those reported in other data sources.[22] 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 (46%), 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 is largely extensive.

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

[22]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.

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 small share (5%) 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 24.6 pp.[23]

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.[24] 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

[23]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.

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

these patterns and that they largely reflect genuine misperceptions.

Our finding that firms hold misperceptions about factors relevant to their decision-making is far from surprising. A growing body of evidence shows that firms—even large ones—can harbor substantial misperceptions. For instance, Cullen et al. (2022) demonstrate that firms misperceive the wages offered by other employers. Similarly, Kim (2025) show that firms are often uninformed about their competitors' prices. And Coibion et al. (2018) use survey data to document significant misperceptions about macroeconomic conditions such as inflation and economic growth.

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.[25] 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., 2025) and business owners (Cai and Szeidl, 2017).[26]

4.3 A Conceptual Framework with Learning and Strategic Complementarities

We now describe a simple conceptual framework that leads to the econometric specifications that characterize our two-stage strategy. There is an economy with a continuum of firms that interact strategically and are uncertain about the underlying productivity of the economy, which also affects their innovation decisions.

[25] 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).

[26] 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).

Belief Update. Firms have common prior $\xi \sim N(0,1)$ where $N \sim (0, \sigma^2)$, denotes the normal distribution with mean μ and variance σ^2. Firms' actions (technology adoption) depends on their own first order beliefs and on the competitors' reaction function. Each firm receives a signal on the past adoption decisions of other firms, which we denote $x_i = \xi + \epsilon_i$, where $\epsilon_i \sim N(0, \frac{\nu}{1-\nu})$ with some $\nu \in (0,1)$ and with ξ, ϵ_i all being independent from each other.[27]. Let $s_{i,t}$ denote firm i's belief about the share of competitors that have adopted advanced technologies up to now, and let $s_{i,t+1}$ denote the corresponding belief about the share of firms that will adopt the technology in the future. In the context of the survey, period t corresponds to 2025, and $t+1$ corresponds to 2027. We assume that firms form their future expectations by projecting their perceptions about the past (see Cavallo et al. (2017)):[28]

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

where β captures the degree of pass-through from perceived current 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 rates among competitors will exceed current levels.[29]

Let T_i be an indicator variable equal to 1 if the individual was randomly chosen to receive a signal and 0 otherwise. We allow the firm's belief about current adoption, $s_{i,t}$, to depend

[27] Note that within the models of global games with learning, such as Angeletos and Huo (2021) or Huo and Takayama (2024) equilibrium actions are a function of the higher order beliefs. Hence our signal on past competitors' adoption can also be interpreted as the equilibrium higher order beliefs from the previous period.

[28] In the literature with strategic complementarities (see Angeletos and Huo (2021) or Huo and Takayama (2024)) learning is more commonly modeled as a Kalman filter. We opt for a simpler update to the signal, but the Kalman filter would lead to a similar testable reduced form. Specifically, the Kalman filter implies that the belief about the state of the economy evolves as follows: $\mathbb{E}_{i,t}[\xi_{i,t}] = \mathbb{E}_{i,t-1}[\xi_{i,t}] + (1-\frac{\lambda}{\rho})(x_{i,t} - \mathbb{E}_{i,t-1}[\xi_{i,t}])$. This coupled with strategic complementarity leads to the following k-order beliefs: $\mathbb{E}^k_{i,t}[\xi_{i,t}] = (1-\frac{\lambda}{\rho})\frac{1}{1-\lambda L} x_{i,t}$, with $\lambda = \frac{1}{2}[(\frac{1}{\rho}+\rho+\frac{1}{\rho\sigma^2}) - \sqrt{(\frac{1}{\rho}+\rho+\frac{1}{\rho\sigma^2})^2 - 4}]$ and where $\sigma^2 = \frac{\nu}{1-\nu}$, which implies that the belief's elasticity to the signal depends on the variance of the signal, σ^2, and the speed of learning convergence, ρ.

[29] 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).

on the firm's prior belief ($s^0_{i,t}$) and on a signal ($s^T_{i,t}$), that the individual may receive in the experiment. Note that the signal carries subscript i to match our application, in which signals differ depending on the firm's sector and size. We start with the case in which the individual receives the information ($T_i = 1$). Based on the assumptions of a Bayesian learning model with Gaussian distributions,[30] after observing the information the individual is expected to update beliefs about current adoption as follows:

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

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

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

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).

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 in turn would 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

[30] 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.

information and estimate the following specification:

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

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 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. So, by taking the ratio $\frac{\gamma_1}{\gamma 3}$ we can separately identify the parameter α.

Action Update. Each firm i chooses its innovation policy to minimize a quadratic function of the innovation policy of the following form: $-(a_i - \xi - \delta\overline{a})^2$, where $\overline{a}$ is the average action in the sector, and which capture the cost of deviating from the average competitors. The first order condition leads to the following innovation policy:

$$a_{i,t} = (1-\delta)\mathbb{E}^k_i[s_{i,t}] + \delta \mathrm{E}_{i,t}[a_t], \quad \text{where} \quad a_t = \int a_{i,t}. \tag{5}$$

The parameter $\delta \in (-1, 1)$ determines the degree of strategic complementarity ($\delta > 0$) or substitability ($\delta < 0$) between agents' actions. After k-iterations, or else substitution of the competitors' reaction function, we get the equilibrium innovation policy:

$$a_{i,t} = (1-\delta)\sum_{k=0}^{\infty} \delta^k \overline{\mathrm{E}}^{k+1}_t [s_{i,t}] \tag{6}$$

The equation above links the post treatment intention to adopt with the change in beliefs estimated from the first stage. Below we discuss econometric aspects of this second stage equation and firms' heterogenous response by market power, shedding light on the motives for responding to competitor adoption.

4.4 Effect of the Information Treatment on Firms' Belief Update

Panel (b) of Figure 2 shows how the provision of information affected the formation of posterior beliefs. The y-axis represents the posterior belief ($s_{i,t+1}$) and the x-axis represents the imputed prior misperceptions ($s^T_{i,t} - s^0_{i,t}$). As predicted by the simple learning model of equation (4), the treatment and control groups exhibit significantly different slopes. Moreover, estimating equation (4) yields a parameter β that is large (1.47) and statistically significant (p<0.001). The fact that β is well 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, panel (b) of Figure 2 reports a value of α that is positive (0.26) and statistically significant ($p < 0.001$). The fact that subjects placed substantial weight on the signal suggests that they 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 (0.26) 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.

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

5.1 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 8.1 pp (p<0.001), from 19.3% to 27.4%. 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 technologies, respectively. For robotics, we find that relative to the control group, the information treatment increased the intention to adopt by 6.7 pp. This effect is both statistically significant (p = 0.029) and economically meaningful, representing a rise from 37.6% to 44.3%—a 17.8% increase ($= \frac{6.7}{37.6}$). In contrast, we find no significant effects on the intention to adopt AI technologies: the estimated impacts are close to zero (2.1 pp for Pred.-AI and -0.2 pp for Gen.-AI) and statistically insignificant (p=0.515 and p=0.955, respectively). These null results, however, should be interpreted with caution, as the estimates are imprecise and do not rule out modest positive effects.

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 (0.2 pp) and statistically insignificant (p = 0.677). 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 (0.2 pp) and statistically insignificant (p = 0.922). Similarly, for Pred.-AI and Gen.-AI, the effects are small (3.3 pp and 0 pp, respectively) and statistically insignificant (p = 0.178 and 0.998).

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 from -24% to 25.5%. The above-median group consists of firms that most strongly underestimated competitors' adoption, with prior gaps ranging from 25.6% to 56.5%.

Panel (a) of Figure 4 shows that, as expected, the effect on posterior beliefs was highly heterogeneous with respect to prior gaps: 3.2 pp (p=0.122) for the below-median group versus 13.3 pp (p<0.001) for the above-median group. Consistently, panel (d) of Figure 4 shows that the treatment effects on the intention to adopt robotics were also highly heterogeneous: 3.4 pp (p = 0.478) for the below-median group versus 11.2 pp (p = 0.001) for the above-median group.

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 (0.2 pp for Pred.-AI and 1.8 pp for Gen.-AI) and statistically insignificant (p = 0.953 and 0.619, 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. As expected, the effects on pre-treatment outcomes are generally close to zero and uniformly statistically insignificant.

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 (0.2 and 1.8 pp, respectively), the corresponding 90% confidence intervals cannot rule out effects of up to 6.1 pp and 7.6 pp, respectively. In other words, we cannot rule out AI adoption effects roughly half as large as those observed for robotics adoption (11.2 pp).

While we cannot rule out that the treatment effects on AI adoption may have been positive, the evidence suggests that these effects were clearly weaker than the corresponding effects on robotics adoption. Several plausible explanations may account for this difference. One possible factor—though unlikely to fully explain it—is the difference in baseline adoption rates. In the control group, intended AI adoption is already higher (48.1% and 52.2% of firms expect to use Pred.-AI and Gen.-AI by 2027, respectively) than intended robotics adoption (37.6%). The fact that baseline rates are substantially higher for AI may simply leave less room for further increases. A second factor may be that robotics is a more established technology, with a long history of adoption and more extensive use rather than experimental implementation. As a result, if a firm learns that competitors are using robotics more than it thought, the perceived urgency to respond may be greater. In contrast, AI technologies are more recent, and those who have adopted them often do so experimentally rather than extensively. Thus, firms may be more cautious about "following the crowd" in this domain.

Regarding the effects of information on robotics adoption, two caveats should be kept in mind. First, experimenter-demand effects may play a role. For instance, after learning that competitors' adoption rates are higher than initially believed, some respondents might feel compelled to report higher intended adoption themselves. However, this explanation seems unlikely, as we observe effects for robotics but not for AI, and it is unclear why experimenter demand would arise for one technology but not the other. A second caveat is that changes in intentions may not translate into actual behavior. Even if respondents genuinely plan to adopt robotics as a result of the treatment, they may fail to follow through—because they forget, face higher-than-expected costs, encounter financing constraints, or simply change their minds. If possible, future versions of the study may explore these mechanisms using follow-up survey data or linked administrative records.

5.2 2SLS Model

The results presented above are intention-to-treat effects—that is, the effects of providing accurate information. Equation 6 from the conceptual framework provides a reduced form equation relating changes in equilibrium posterior beliefs to changes in equilibrium adoption choices. We now describe how we estimate this equation to infer the causal impact of a shock to beliefs.

To better assess the magnitude of the causal effects of expectations, 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 ($s_{i,t+1}$), hence the posterior derived in the previous section. However, because posterior beliefs may be endogenous, such a regression would not necessarily identify the causal effect of those beliefs. To achieve causal identification, we use a 2SLS model that relies solely on the exogenous variation in posterior beliefs generated by the randomized provision of information.

We can illustrate the intuition behind the 2SLS model with a simple example. The model compares pairs of individuals with the same prior gap. For instance, consider two individuals who both underestimate competitors' 2024 adoption rate by 10 pp. We then randomly assign information about the true adoption rate. Relative to the uninformed individual, the informed individual should end up with a higher posterior belief about future adoption. Suppose the uninformed individual continues to underestimate future adoption by 10 pp, while the informed individual reacts strongly and now underestimates it 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. We can then examine how this 8 pp shift in expectations translates into behavioral differences—specifically, whether it increases the probability of adopting an advanced technology in the future.[31]

The specification of interest is:

$$a^j_{i,t+1} = \beta_0 + \beta^j_s \cdot s_{i,t+1} + \beta_2 \cdot (s^T_{i,t} - s^0_{i,t}) + \beta_3 \cdot s^0_{i,t} + X_i\beta_X + \epsilon_i \tag{7}$$

The endogenous variable is $s_{i,t+1}$ and the excluded instruments is $T_i \cdot (s^T_{i,t} - s^0_{i,t})$.[32] X_i is the vector of additional control variables, such as the firms' current adoption status and

[31]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 those whose posteriors shifted in response to the information intervention. By design, this estimate assigns greater weight to individuals with larger initial misperceptions and, given those misperceptions, to those who adjust their beliefs more strongly when exposed to the signal.

[32]Formally, the exogeneity assumption is $\mathbb{E}\left[\left(s^T_{i,t} - s^0_{i,t}\right) \cdot T_i \cdot \epsilon_i \mid X_i\right] = 0$ where X_i is the vector of control variables $\{s^T_{i,t} - s^0_{i,t}, s^0_{i,t+1}, X_i\}$. 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.

basic characteristics.[33]

Besides the baseline specification in 7, we also estimate the heterogenous change in action across firms with different mark-ups, market shares and facing different degree of market concentration. This elasticity, which we present in Appendix C, move a step forward in the direction of isolating the competition channel from the rest.

5.3 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.

Most importantly, Panel (a) of Table 2 presents the 2SLS estimates. Before discussing those, we briefly consider Panels (b) and (c). Panel (b) reports the first-stage estimates. The results indicate that, consistent with the evidence from Section 4.3, posterior beliefs change significantly in response to the provision of information. This strong first-stage relationship suggests that there is little scope for weak-instrument bias (Stock et al., 2002). For a more 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.[34] 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.

In turn, Panel (c) of Table 2 reports the reduced-form estimates—that is, the intention-

[33]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 program; 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).

[34]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.

to-treat effects of information on future adoption. Consistent with the results from Section 5 discussed above, the reduced-form regressions show significant effects of the information treatment on robotics adoption but no significant effects on either form of AI adoption.

Panel (a) of Table 2 reports the 2SLS estimates, which combine the first-stage and reduced-form effects.[35] Expected competitors' innovation has no significant effect on AI adoption intentions but exhibits a large and precisely estimated effect on robotics adoption. More specifically, a 1 pp increase in expected competitors' adoption raises the probability of adopting robotics by 0.704 pp. This coefficient can be interpreted as the slope of the reaction function, which carries two key implications. First, the steepness of the reaction function underscores the strength ofI complementarities in technology adoption, suggesting that firms' decisions are highly interdependent. In other words, when one firm expects its competitors to adopt, it becomes substantially more likely to follow suit—likely to avoid falling behind. Second, the magnitude of this response implies that local shocks—such as financial incentives for adoption—may have effects that propagate well beyond the directly treated firms. Because each firm's decision influences the expectations and incentives of others, even modest initial increases in adoption could generate sizable aggregate spillovers.

Finally, table C.1 in Appendix C presents the heterogenous response of robot adoption across firms facing different mark-ups (which we compute with a Lerner index), market shares or market concentration (computed with Herfindhal index). We find that firms with large mark-ups adjust their adoption decision the most, a sign that those firms are compelled by the investment decisions of competing firms. We also find that firms with below median market share are more responsive to investments by competing firms. Small or young firms may be especially sensitive to competition as a matter of survival.

6 Conclusions

This study set out to provide 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 examined how firms update their beliefs about competitors'

[35] Appendix C.2 provides complementary evidence based on the correlation between beliefs and adoption.

adoption of advanced technologies and whether these updated beliefs affect their own intended adoption. We documented widespread underestimation of competitor adoption, substantial belief updating in response to information, and effects on intended adoption of robotics—but not AI. These findings suggest that spillovers are stronger for more mature technologies, where competitors' choices provide clearer strategic signals.

We conclude by discussing 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 program, is devoted to supporting firms that invest in advanced technologies or innovation.

Our evidence on information frictions suggests that, as a complement to financial-incentive policies, governments could deploy information campaigns that raise firms' awareness of the productivity benefits of new technologies and the extent to which their peers are adopting them. While further research is needed to refine their design and evaluate their effectiveness, the relatively low cost of implementation indicates that such policies could be a highly cost-effective tool to stimulate technological diffusion.

References

Angeletos, G.-M. and Z. Huo (2021). Myopia and Anchoring. *The American Economic Review 111*(4), 1166–1200.

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 (2024). Annual Report for 2023. https://www.bancaditalia.it/pubblicazioni/relazione-annuale/2023/index.html.

Bank of Italy (2025). Survey of Industrial and Service Firms in 2024. 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.

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 D.P.*.

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 (2017). 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.

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

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 (2022). What's My Employee Worth? The Effects of Salary Benchmarking. *Review of Economic Studies, forthcoming*.

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*.

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

Faia, E., M. Mayer, and V. Pezone (2025). The Value of Firm Networks: A Natural Experiment on Board Connections. *Review of Corporate Finance Studies, forthcoming*.

Farboodi, M., R. Mihet, T. Philippon, and L. Veldkamp (2019). Big Data and Firm Dynamics. *American Economic Association 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), 223–59.

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 (2020). Information, technology adoption and productivity: The role of mobile phones in agriculture.

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. https://automazione-plus.it/wp-content/uploads/sites/3/2023/07/Market-Report-Italy-2023-Free-Preview-compresso.pdf.

Huo, Z. and N. Takayama (2024). Rational Expectations Models with Higher-Order Beliefs. *Review of Economic Studies*, rdae096.

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.

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

(a) Use of Artificial Intelligence, by Intensity

(b) Use of Robots, by Intensity

(c) Use of Artificial Intelligence, Over Time

(d) Use of Robots, Over Time

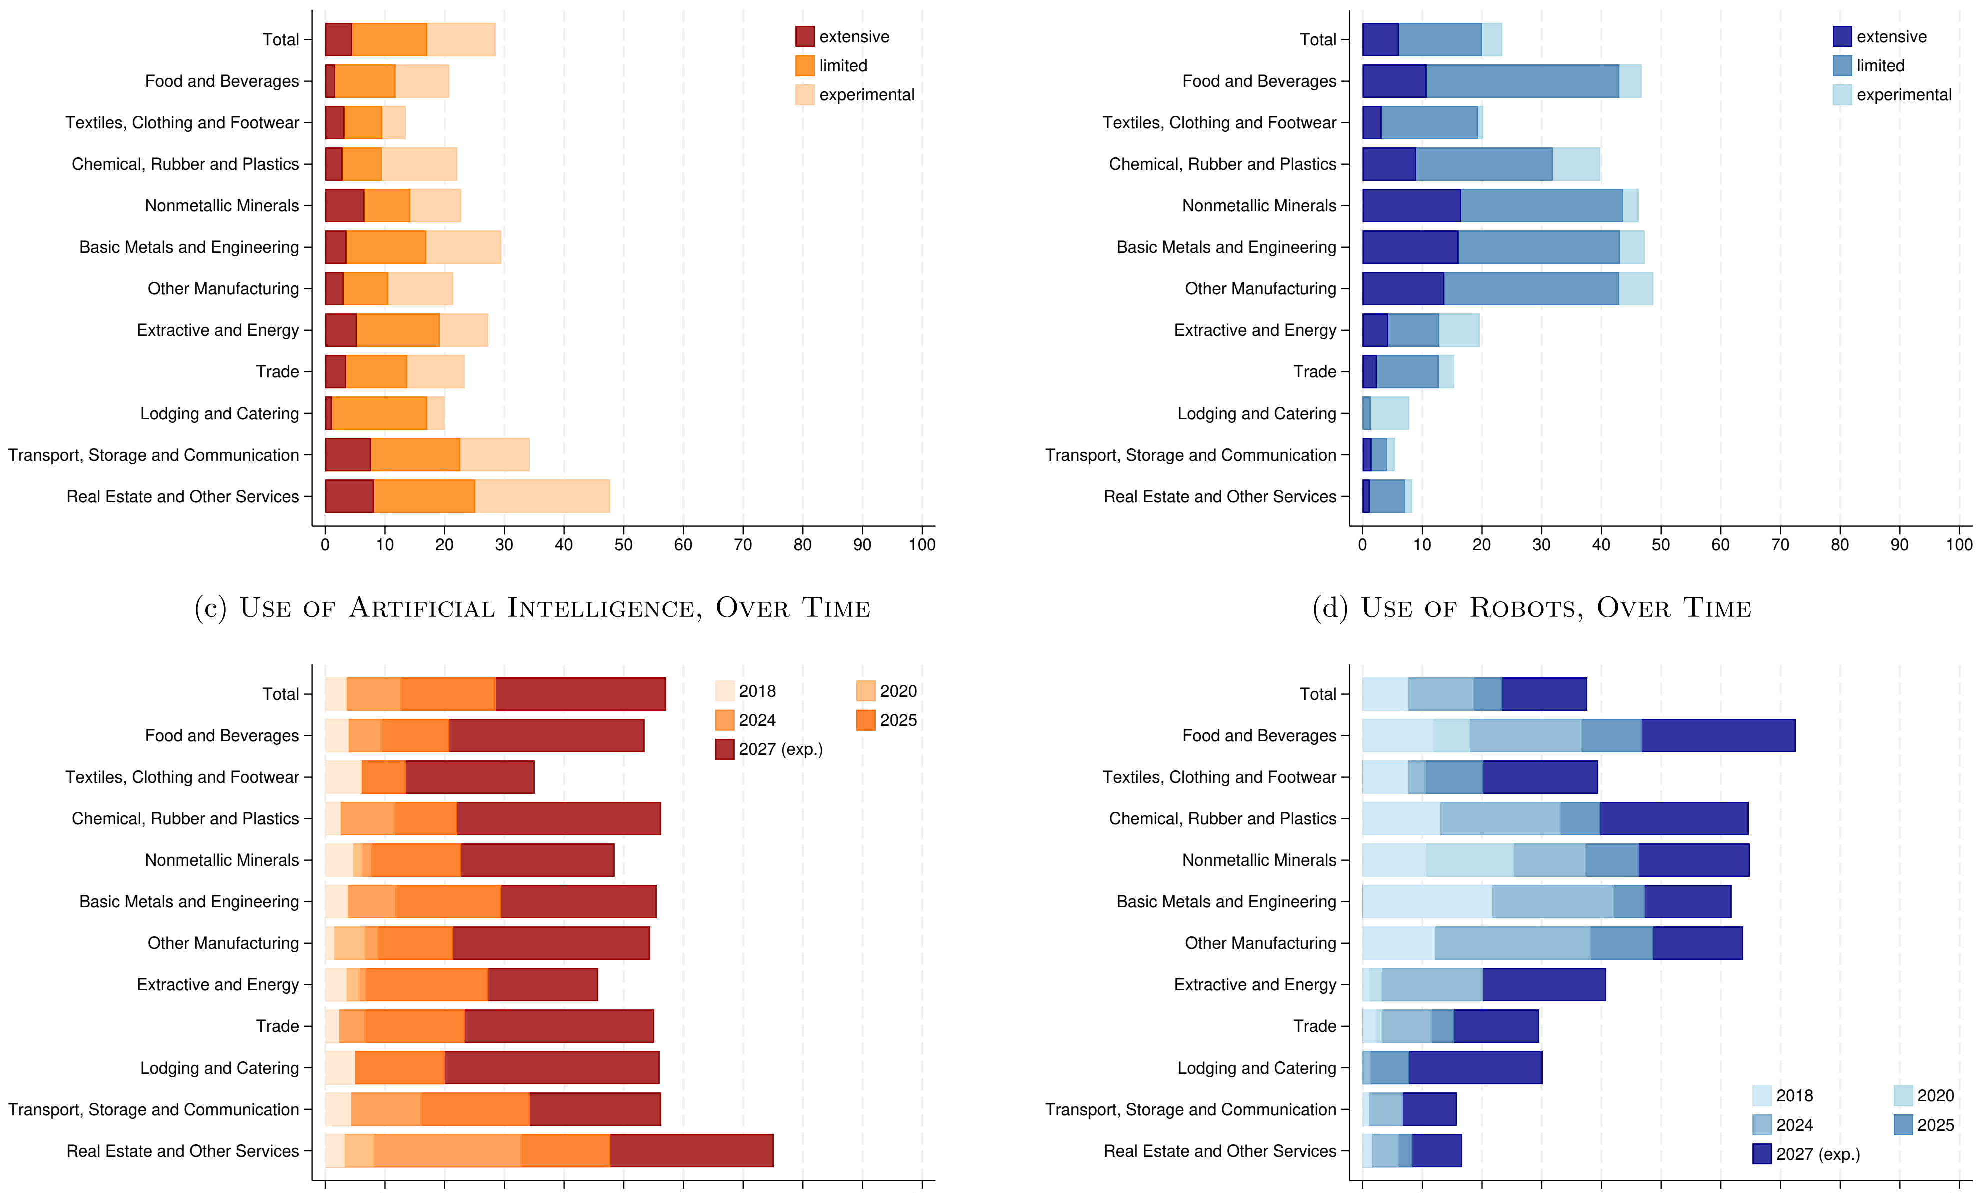

Notes: Panels (a) and (b) are based on the INVIND survey conducted in 2025 and report the intensity in the use of either predictive AI or generative AI (panel a) and robots (panel b) at the moment of the interview (the corresponding survey questions are TEC5N1, TEC5N2 and TEC11N). Panels (c) and (d) are based on the surveys conducted in 2018, 2020 and 2025 and report the share of firms that already used AI or robots at all in those years. Note, the figures for 2024 reported here may differ from the information treatments described in Table B.1, because the latter included firms that were not using these technologies but expected to introduce them by the end of 2024. The bars related to expected use of these technologies in 2027 are based on questions TEC27A, TEC27B and TEC27C in the 2025 survey wave and are computed considering only the sample of control firms.

Figure 2: Distribution of Prior Beliefs and Belief-Updating

(a) Raw perception gap (b) Expected adoption by competitors

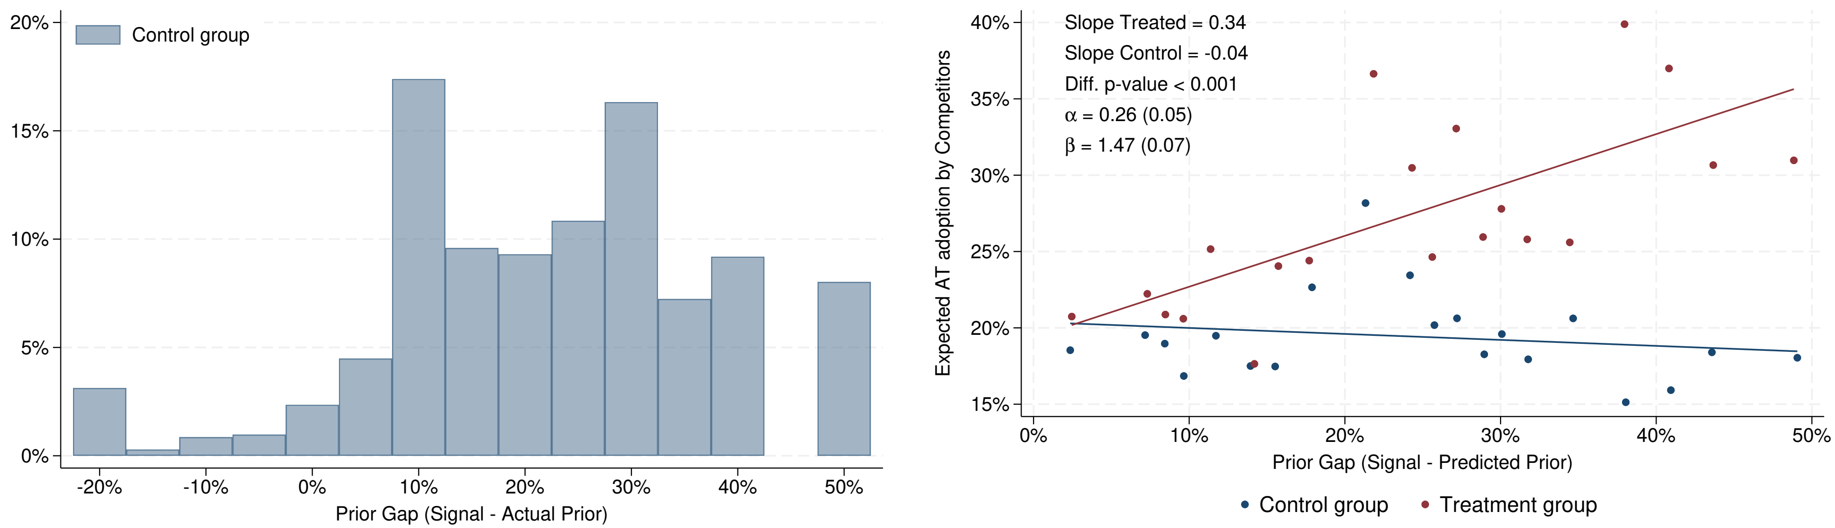

Notes: The Figure is based on the INVIND survey conducted in 2025. Panel (a) shows the empirical distribution of the gap between prior belief and the information that would have been shown if treated, for the sample of control firms. The first and the last bins group observations whose gaps are below -15 and above 45, respectively. Panel (b) displays a binscatter of the posterior beliefs on competitors' adoption of advanced technologies (questions TEC26A and TEC26B in the Survey module) and the imputed gap between prior beliefs and the information treatment shown, controlling for the level of prior beliefs. α indicates the learning weight, β the degree of extrapolation from the information received, as obtained in equation (3).

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

(a) Treatment Effects (b) Falsification Tests

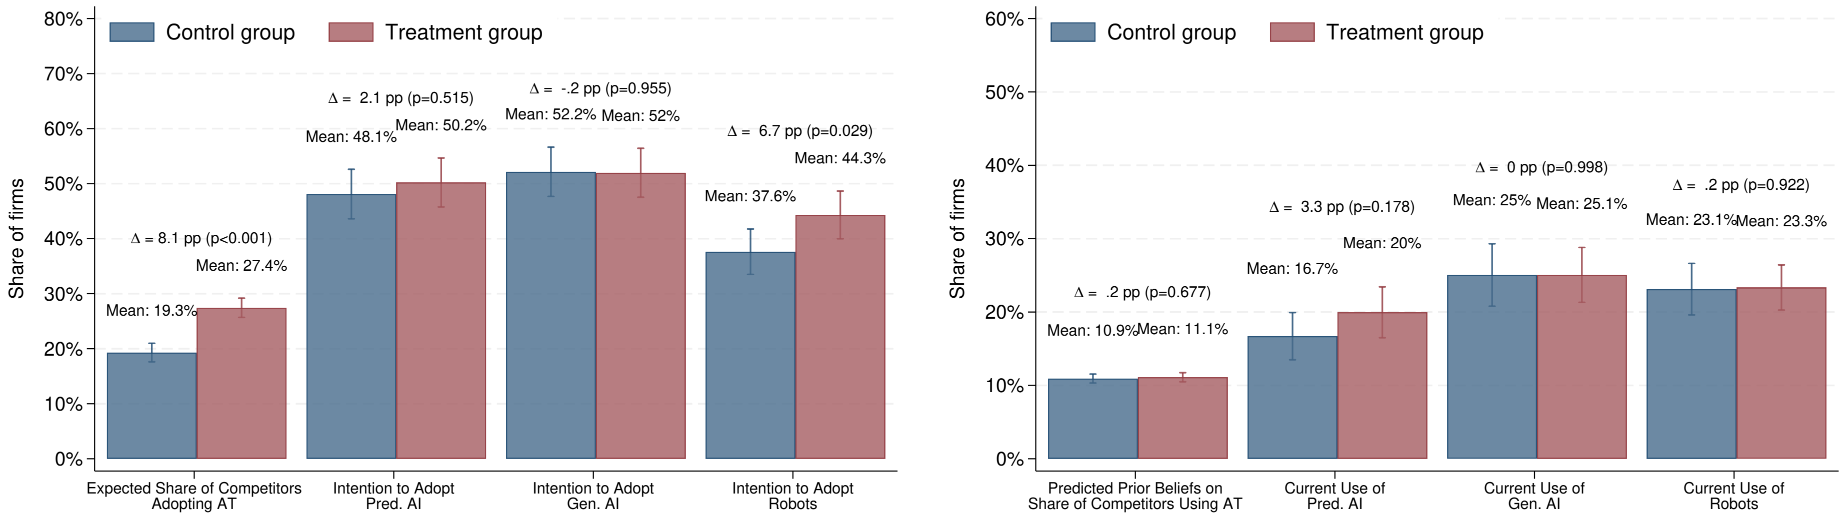

Notes: The Figure is based on the INVIND survey conducted in 2025. Panel a represents posterior beliefs on competitors' future adoption of advanced technologies (question TEC26A and TECT26B in the Survey module) and intention to adopt advanced technologies by 2027 (questions TEC27A, TEC27B and TEC27C). Panel b represents imputed prior beliefs on competitors' current use of advanced technologies (question TEC24 in the Survey module) and current use in the firms of these technologies (questions TEC5N1, TEC5N2 and TEC11N). Imputed prior beliefs are based on the estimates presented in Table B.2. Confidence intervals are at 95% level.

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

(a) Posterior Beliefs

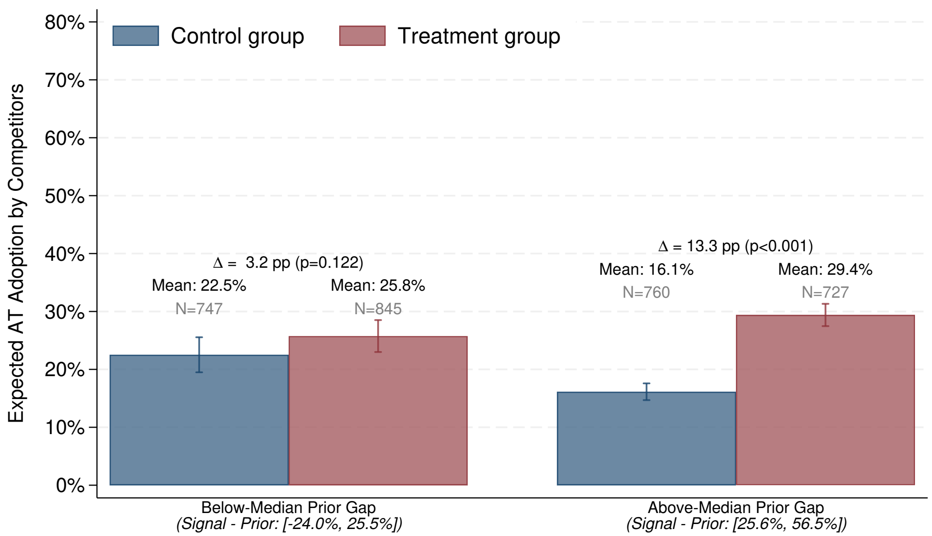

(b) Pred.-AI Future Adoption

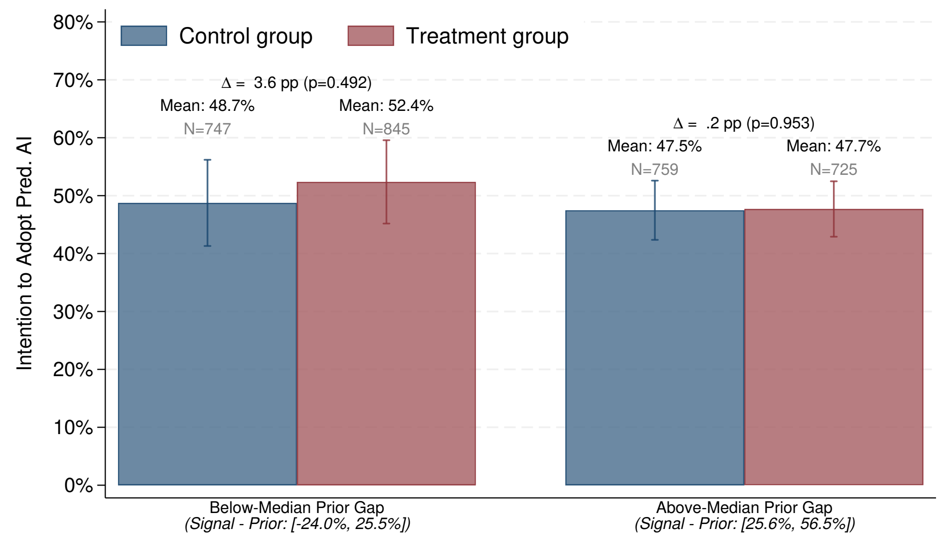

(c) Gen.-AI Future Adoption

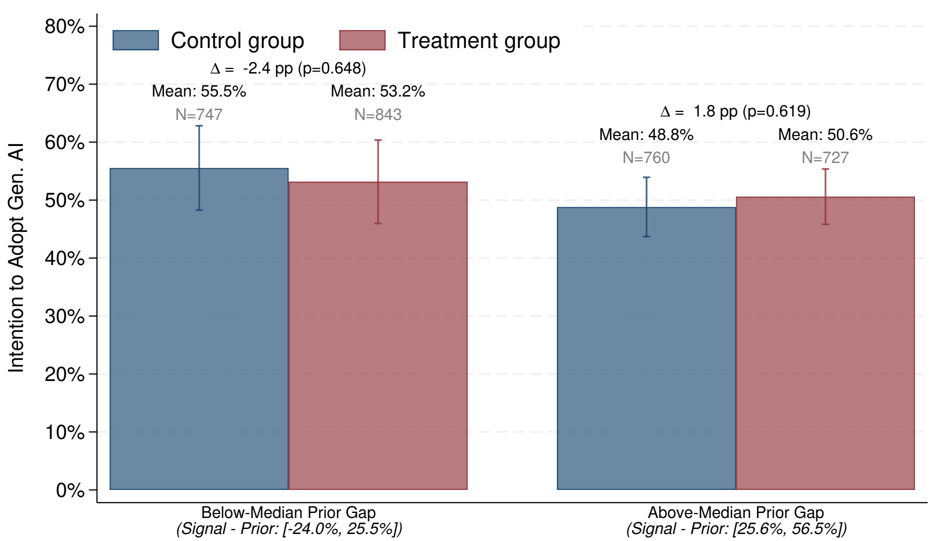

(d) Robots Future Adoption

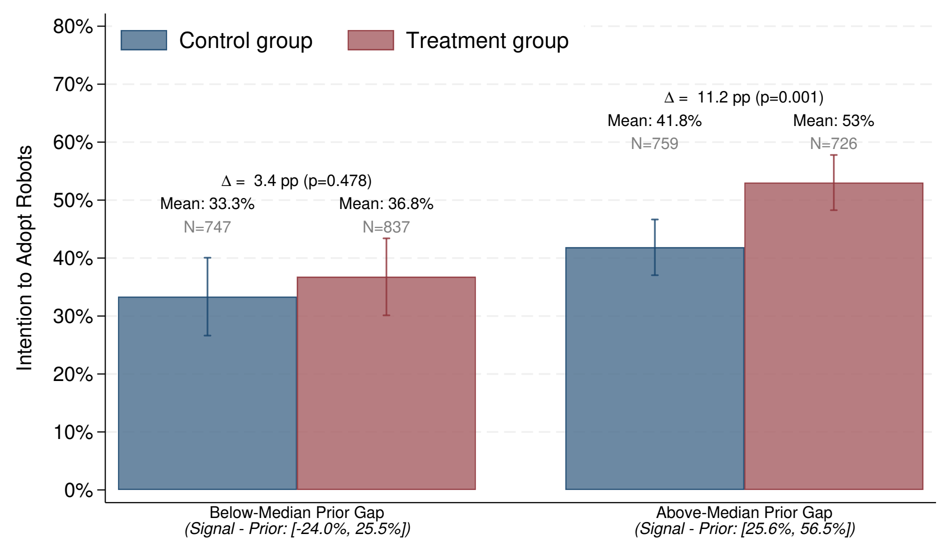

Notes: The Figure is based on the INVIND survey conducted in 2025. The graphs represent posterior beliefs on competitors' future adoption of advanced technologies (question TEC26A and TECT26B in the Survey module) and intention to adopt advanced technologies by 2027 (questions TEC27A, TEC27B and TEC27C) by imputed prior gap (below vs above median) and by treatment group. Confidence intervals are at 95% level. The imputed prior gap is based on the imputed prior, according to the estimates presented in Table B.2.

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

(a) Imputed Prior Beliefs (b) Current use of PredAI

(c) Current use of GenAI (d) Current use of Robots

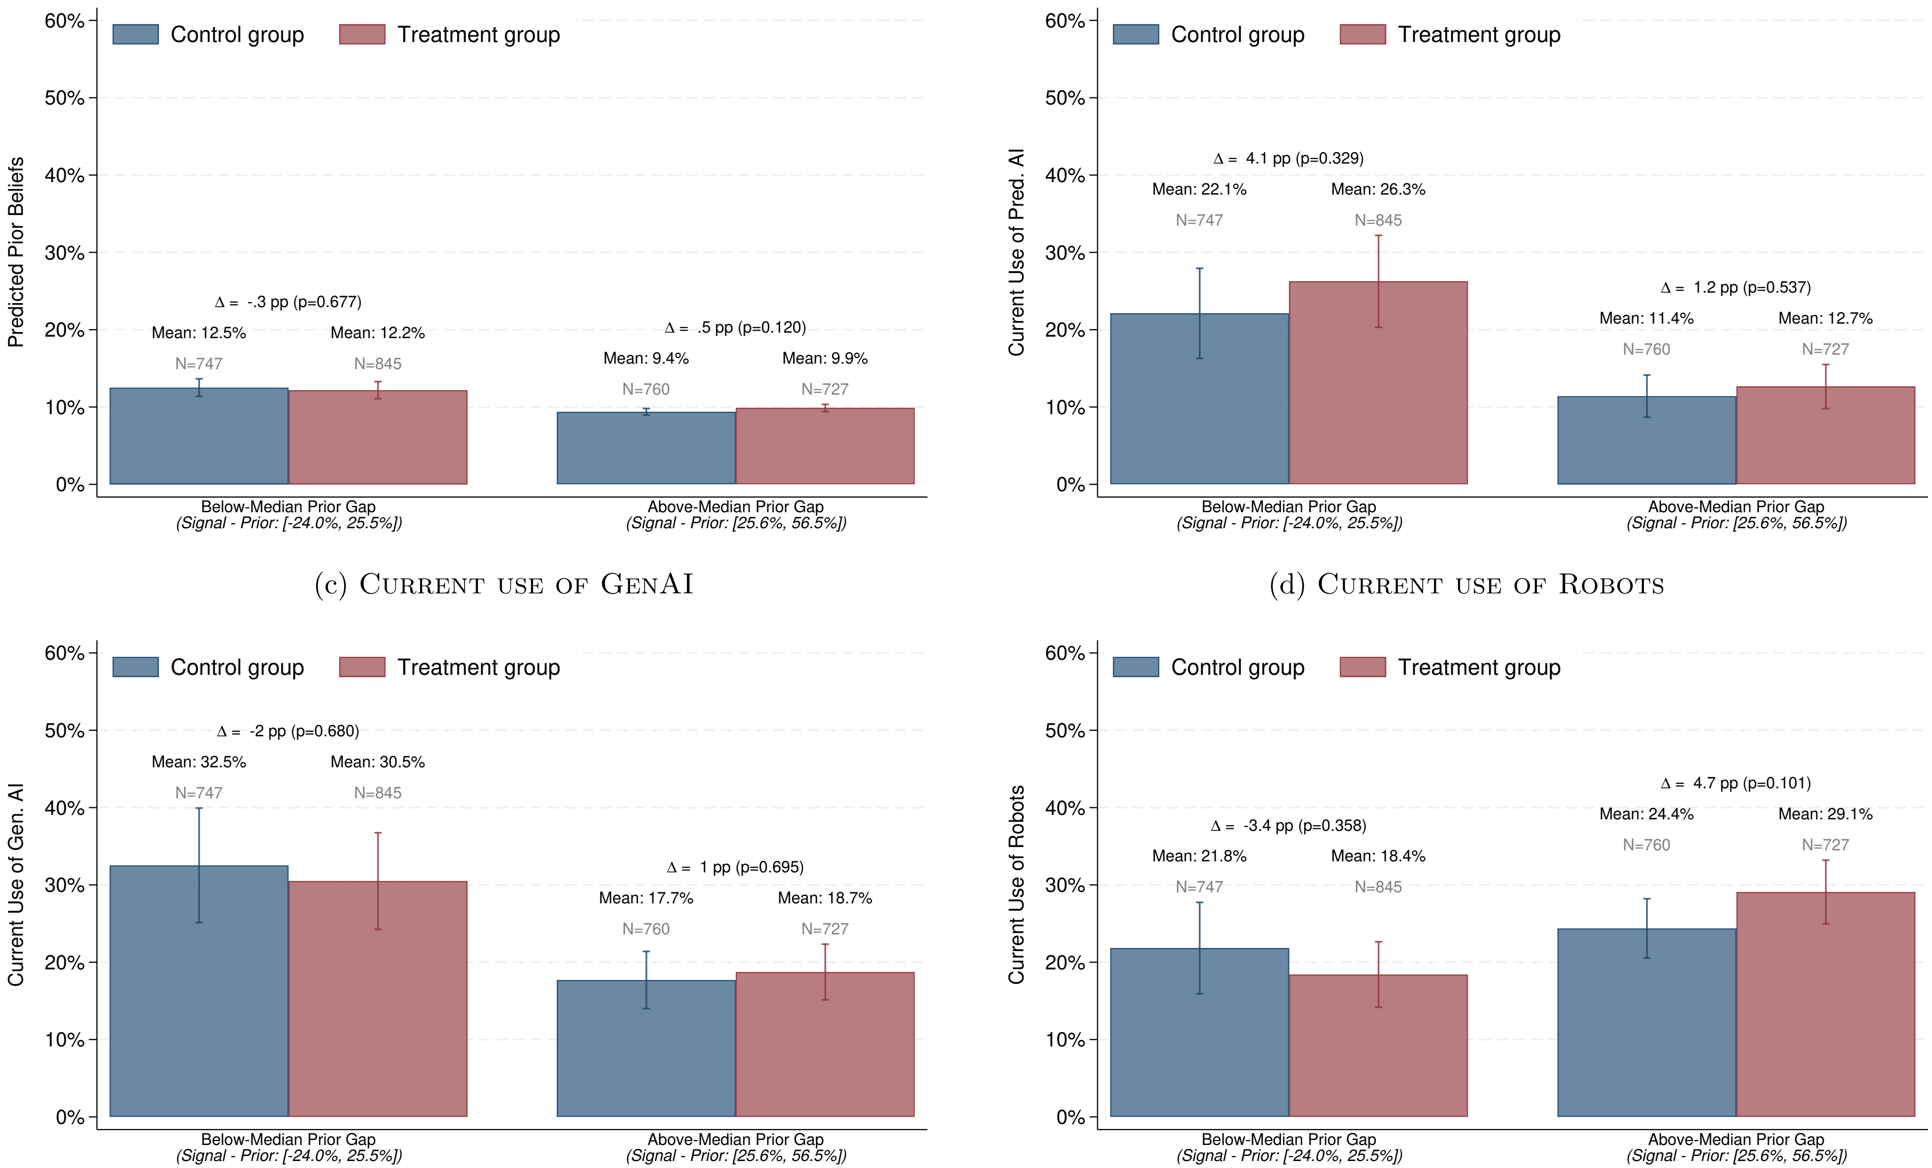

Notes: The Figure is based on the INVIND survey conducted in 2025. The graphs represent imputed prior beliefs on competitors' current use of advanced technologies (question TEC24 in the Survey module) and current use in the firms of these technologies (questions TEC5N1, TEC5N2 and TEC11N) by imputed prior gap (below vs above median) and by treatment group. Confidence intervals are at 95% level. The imputed prior gap is based on the imputed prior, according to the estimates presented in Table B.2.

Table 1: Summary Statistics and Randomization Balance

	All (1)	Control (2)	Treatment (3)	p-value (4)
Firm Age	40.450	40.368	40.530	0.898
	(0.384)	(0.537)	(0.548)	
Number of Employees in 2024	95.908	95.893	95.922	0.997
	(11.322)	(11.978)	(19.029)	
Geographical Area: North (%)	57.732	58.013	57.462	0.857
	(0.890)	(1.266)	(1.253)	
Geographical Area: Center (%)	21.280	20.860	21.683	0.751
	(0.738)	(1.042)	(1.044)	
Geographical Area: South and Islands (%)	20.988	21.127	20.856	0.893
	(0.734)	(1.047)	(1.030)	
Manufacturing Firms (%)	40.075	38.663	41.429	0.338
	(0.883)	(1.249)	(1.248)	
Turnover (Million euros)	42.786	44.286	41.348	0.537
	(5.158)	(7.021)	(7.542)	
Share of Investment on Turnover	0.067	0.064	0.071	0.522
	(0.007)	(0.005)	(0.014)	
Hourly Labor Cost	0.033	0.039	0.028	0.333
	(0.024)	(0.049)	(0.001)	
Exporters (%)	53.737	53.843	53.636	0.950
	(0.899)	(1.279)	(1.264)	
Current Use of Pred.-AI (%)	18.374	16.712	19.967	0.178
	(0.698)	(0.957)	(1.013)	
Current Use of Gen.-AI (%)	25.054	25.050	25.058	0.998
	(0.781)	(1.111)	(1.098)	
Current Use of Robotics (%)	23.233	23.115	23.347	0.922
	(0.761)	(1.081)	(1.072)	
Share of Investment in AT	6.801	6.508	7.083	0.491
	(0.262)	(0.362)	(0.378)	
Number of People Involved in the Survey	2.128	2.139	2.117	0.750
	(0.026)	(0.035)	(0.037)	
Observations	3,079	1,521	1,558	

Notes: The Table presents summary statistics form a set of firms' characteristics, location and survey responses for the INVIND survey conducted in 2025. Standard errors in parenthesis.

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

	Pred.-AI (1)	Gen.-AI (2)	Robotics (3)
Panel (a) - 2SLS Regressions			
Exp. AT Adoption by Competitors	0.020	0.066	0.704***
	(0.245)	(0.232)	(0.201)
Panel (b) - 2SLS: First Stage			
$T_i *$ Prior gap	0.339***	0.339***	0.339***
	(0.028)	(0.028)	(0.028)
Panel (c) - Reduced-Form Regressions			
$T_i *$ Prior gap	0.008	0.023	0.242***
	(0.084)	(0.076)	(0.072)
Observations	3,066	3,068	3,064
Dep. Var.: Baseline Mean	48.1	52.0	37.7
Kleibergen-Paap Wald F-statistic	134.6	132.7	132.4

Notes: Panel (a) presents the results of the 2SLS model from equation (7) and discussed in Section 5.2. The dependent variable is a dummy taking a value of 100 if firm i expects to use either predictive AI, generative AI or robots by 2027. Panel (b) presents the results of the first stage associated to the 2SLS model, i.e. the effect of the variable $T_i \cdot (s^T_{i,t} - s^0_{i,t})$ on the the posterior beliefs about competitors' future adoption of advanced technologies ($s_{i,t+1}$). Panel (c) presents the results for the corresponding reduced-form of the 2SLS model, capturing the effect of $T_i \cdot (s^T_{i,t} - s^0_{i,t})$ when it is not used to instrument posterior beliefs. All standard errors are bootstrapped.