NBER WORKING PAPER SERIES

INNOVATION WITHOUT BORDERS? THE GEOGRAPHY OF TECHNOLOGICAL DIFFUSION

Ursel Baumann
Zoë B. Cullen
Ester Faia
Annalisa Ferrando
Ricardo Perez-Truglia
Judit Rariga

Working Paper 35314
http://www.nber.org/papers/w35314

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2026

We are thankful for comments by colleagues and seminar participants. This project was reviewed and approved in advance by the Institutional Review Board at Goethe University Frankfurt and was pre-registered with the AEA RCT Registry under AEARCTR-0017121. After the study is accepted for publication, we will make all code publicly available in a repository and provide instructions for accessing the anonymized microdata. This paper should not be reported as representing the views of the European Central Bank (ECB). The views expressed are those of the authors and do not necessarily reflect those of the ECB. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Innovation without Borders? The Geography of Technological Diffusion
Ursel Baumann, Zoë B. Cullen, Ester Faia, Annalisa Ferrando, Ricardo Perez-Truglia, and Judit Rariga
NBER Working Paper No. 35314
June 2026
JEL No. C93, D22, L21, O33

ABSTRACT

How well does innovation diffuse across geographic boundaries? To shed light on this question, we present a large-scale field experiment involving 3,300 firms across twelve European Union countries. We elicit firms' perceptions of the share of similar firms in their own country that had invested in artificial intelligence (AI), as well as the corresponding share among similar firms in Germany, France, and Italy. We randomly provide half of the sample with accurate information about both domestic and foreign AI investment. We show that firms substantially underestimate competitors' current AI investment, both domestically and abroad, and that they update their expectations about competitors' future AI investment in response to the information treatment. The treatment also causes a statistically significant increase in firms' own expected AI investment rate (p-value < 0.001). We find strong strategic complementarities within borders: a 1 pp increase in the expected share of domestic peers investing in AI raises a firm's own expected AI investment rate by 0.570 pp. These complementarities are absent across borders: the effect of an increase in the expected share of foreign peers investing in AI on a firm's own expected AI investment rate is statistically insignificant. Overall, our evidence shows that innovation diffusion and strategic complementarities in AI investment are much stronger domestically than internationally.

Ursel Baumann
European Central Bank
ursel.baumann@ecb.europa.eu

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

Annalisa Ferrando
European Central Bank
annalisa.ferrando@ecb.europa.eu

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

Judit Rariga
European Central Bank
erzsebet-judit.rariga@ecb.europa.eu

An online appendix is available at http://www.nber.org/data-appendix/w35314
A randomized controlled trials registry entry is available at EARCTR-0017121

1 Introduction

The rise of Artificial Intelligence and the expansion of the knowledge economy can have profound consequences for growth. This makes it vital to understand the underlying determinants of innovation and its diffusion in space. Strategic complementarities - the idea that the marginal benefit of one innovation activity increases when another is present - are a key force since AI is a general-purpose technology that rarely generates value in isolation. Its effective adoption depends on learning complementarities with other adopters; furthermore pressure from competitors acts as a strategic catalyst forcing firms to innovate. Those insights have shaped influential work on technology diffusion, from the adoption of hybrid corn in response to local competitive pressures (Griliches, 1960, 1987) to the role of trade and cross-border knowledge flows in shaping innovation and firm dynamics (Keller and Yeaple, 2013). More broadly, this literature suggests that knowledge spillovers from abroad may be especially valuable for growth because they are more likely to embody genuinely novel ideas. Yet the same foreign spillovers may exert weaker pressure on firms' behavior if foreign competitors face cultural, linguistic, or practical barriers to entering local markets. In practice, then, the same signal—for example, that competitors are investing in a new technology—may lead firms to respond differently depending on whether those competitors are domestic or foreign.

We begin with a simple model of strategic interaction under incomplete information to guide our research design. Building on Angeletos and Pavan (2004), we adapt their framework to the context of technology investment and extend it to a multi-country setting. Firms choose how much to invest in a new technology without fully observing its underlying productivity. Instead, they form beliefs from noisy information and from their expectations about competitors' choices. A firm's optimal investment therefore depends not only on its own assessment of the returns to AI, but also on what it expects competitors to invest. These expectations matter through two main channels. First, there is a competition channel: if competitors invest more, a firm may also want to invest more to avoid falling behind. Second, there is a learning channel: the behavior of similar firms may reveal information about the underlying profitability of the technology.

The model predicts that firms should respond to information about other firms' investment decisions, and that the size of this response may depend critically on whether those firms are domestic or foreign. If domestic firms are closer competitors than foreign firms, the competition channel implies that a firm should respond more strongly to investment by domestic firms than foreign firms. If productivity is more correlated among domestic firms than among foreign firms, the learning channel yields the same prediction: domestic investment should be more informative than foreign investment. Thus, the relative effects of information about domestic and foreign investment provide a test of the technological and competitive frictions that separate markets across borders.

We study AI investment and use, an important setting in which firms' investment decisions are both costly and consequential. Using AI effectively often requires skill upgrading and organizational change, but it can also raise productivity, expand sales, and increase firm market value (Babina et al., 2024). Policymakers therefore increasingly view AI as a key driver of growth and industrial competitiveness. Our empirical setting is a large-scale field experiment with firms in twelve European Union (EU) countries. This setting provides a natural laboratory for studying domestic versus foreign competitive pressure. The European Single Market was explicitly designed to expose firms to cross-border competition and thereby foster innovation and growth. At the same time, AI diffusion remains highly uneven across countries and firms, suggesting that important barriers to diffusion persist even within an integrated economic area.

We embedded our information-provision experiment in the Survey on the Access to Finance of Enterprises (SAFE), a long-running survey conducted by the European Central Bank and the European Commission. SAFE provides a uniquely harmonized cross-country survey infrastructure for implementing the experiment. Because our design randomizes tailored information about domestic and foreign firms' investment decisions, we included a series of questions in an earlier survey wave to construct the information that would subsequently be randomized. More precisely, we use data from an earlier SAFE wave to compute, for each country-sector-size cell, the share of firms that had invested in AI by June 2025. We then used these data in the information-provision experiment, which was implemented in the experimental wave.

The experiment first elicits firms' prior beliefs about domestic investment: that is, the share of similar firms—those in the same sector and size class and from the same country—that had invested in AI. It then elicits firms' prior beliefs about foreign investment: that is, the share of similar firms in the same sector and size class from the three largest EU economies—Germany, France, and Italy, hereafter referred to as the big-3—that had invested in AI. In the information-provision stage, half of the firms are randomly assigned to a control group and receive no additional information. The other half are assigned to a treatment group and receive information about the share of domestic and foreign firms that had invested in AI, as measured before the experiment. We then elicit firms' posterior beliefs about both domestic and foreign investment rates. Finally, we measure each firm's own expected AI investment rate, defined as the share of total investment that it expects to allocate to AI over the next twelve months. This design allows us to trace the full chain from information to beliefs to intended behavior.

Our main analysis sample consists of 3,316 firms that completed our module and are located in the following twelve EU countries: Austria, Belgium, Finland, France, Germany, Greece, Ireland, Italy, the Netherlands, Portugal, Slovakia, and Spain. The scale of the survey provides substantial statistical power to document firms' perceptions of their domestic and foreign peers, and to distinguish the effects of domestic and foreign beliefs on firms' own investment choices.

We begin by documenting substantial unevenness in AI investment and use across these twelve EU countries. As of June 2025, AI investment is much more prevalent in some countries than in others. The share of firms investing in AI averages about 28%, but it ranges from 20% in Greece to 48% in Austria. The share of firms investing in AI is slightly lower among the big-3 (25%) than among the other nine countries (29%). Current AI use is similarly uneven across countries. Firms differ sharply in their intensity of use: 30% report no AI use, 32% report very infrequent use, 29% report moderate use, and only 9% report significant AI use. Indeed, AI use and investment in the twelve EU countries we study are broadly comparable to those in the United States, both in overall levels and in the unevenness of diffusion across firms (Baslandze et al., 2026; Bonney et al., 2026; Yotzov et al., 2026).

We then show that firms hold large and systematically biased beliefs about competitors' AI investment. Most firms underestimate both domestic and foreign competitors' AI investment, often by substantial margins. On average, firms underestimate domestic competitors' AI investment by 14 pp and foreign competitors' AI investment by 7 pp. While domestic and foreign misperceptions are positively correlated, they are far from perfectly correlated. This indicates that each contains substantial independent variation to allow us to separately identify their effects.

A country-level breakdown reveals two additional patterns. First, underestimation is especially severe in countries where actual AI investment is high, suggesting that firms' beliefs are compressed relative to the true cross-country variation. Second, firms systematically overstate the big-3's relative position: they believe that big-3 firms invest more in AI than firms in their own country, even though the average investment share is similar between the big-3 and the other nine countries. These patterns suggest a broader center-periphery misconception, in which firms perceive AI investment as concentrated in a small set of technologically advanced economies even though it is more widely dispersed in practice.

We show that firms update their beliefs in response to the information treatment. Firms that initially underestimate competitor AI investment revise their beliefs upward, firms that overestimate revise downward, and firms with accurate priors update little. Consistent with baseline underestimation, factual information raises both domestic and foreign posterior beliefs, by 4.3 pp and 3.3 pp respectively. We use the experimental data to estimate a simple Bayesian learning model of expectation formation. The implied learning weights are 0.35 for domestic expectations and 0.36 for foreign expectations—meaning that firms place roughly one-third weight on the signal and two-thirds weight on their prior beliefs. This evidence indicates that respondents paid attention to the information and viewed it as credible and relevant. We also document asymmetric cross-market extrapolation: firms use foreign investment to forecast domestic investment, but not vice versa. This suggests that firms expect domestic markets to partly catch up to the big-3, but do not expect investment in their own country to affect investment in the big-3.

The information intervention changed not only firms' beliefs about competitors, but also their own intended AI investment over the next twelve months. Treated firms expect to allocate 10.13% of their investment to AI, compared with 8.33% among control firms. This 1.8 pp treatment effect is statistically significant (p-value < 0.001). We show that the treatment effect is robust to alternative specifications and falsification tests. For example, the treatment has no effect on pre-treatment AI use or prior beliefs. Moreover, as predicted by Bayesian learning, the treatment effects on future AI investment are concentrated among firms that initially underestimated competitor AI investment the most.

The above treatment effects correspond to the joint provision of information about the AI investment of domestic and foreign firms. To disentangle between the domestic and foreign channels, we estimate a Two-Stage Least Squares (2SLS) model. Following Cullen and Perez-Truglia (2022), this model uses the exogenous variation in domestic and foreign posterior beliefs induced by the randomization of information. The estimates show a sharp asymmetry, with the behavioral response operating primarily through beliefs about domestic competitors. On the one hand, firms' own expected AI investment rate responds strongly to domestic posterior beliefs: a 1 pp increase in the perceived share of domestic competitors investing in AI raises the firm's own expected AI investment rate by 0.570 pp. This effect is statistically significant (p-value < 0.001) and also large in magnitude: it implies an elasticity of 1.6 between own investment and the belief about domestic investment. By contrast, the effect of a 1 pp increase in the belief in foreign investment is closer to zero and statistically insignificant. Moreover, we can reject the null hypothesis that the effects of domestic and foreign beliefs are equal to each other (p-value < 0.001).

Our model points to two channels that could explain the sharp difference between the effects of domestic and foreign beliefs. The first is a competition channel. According to this channel, firms respond to domestic peers' AI investment because they view those peers as close competitors and do not want to fall behind. By contrast, they respond less to foreign peers' investment if they are not perceived as direct competitors. The second is a learning channel. According to this channel, firms respond to domestic peers' AI investment because domestic peers operate in similar cultural, institutional, and organizational environments. Thus, learning that these peers are investing in AI signals that the technology may be more valuable than the firm previously believed. By contrast, if foreign peers operate in sufficiently different environments, then their AI investment would be less informative about the value of AI for firms outside those foreign markets.

Because the competition and learning channels are closely intertwined, disentangling them is challenging. Nevertheless, we provide suggestive evidence through heterogeneity analysis. Intuitively, the competition channel should vary with the degree of competition in the market. In the extreme case of a perfect monopolist, a firm could still learn from other firms' AI investment decisions, but the competition channel would be completely absent. Consistent with the competition channel, we find that the effect of beliefs about domestic peers is stronger when the domestic market is less concentrated and when the firm is less exposed to foreign competitors.

Taken together, the results point to an important limit on cross-border diffusion. An increase in perceived AI investment by foreign firms does not cause firms to increase their own AI investment, suggesting that foreign competition may be less effective than domestic competition in encouraging technology adoption, even in an integrated economic region such as Europe. This finding helps quantify one channel through which innovation may remain geographically concentrated: if firms respond less to foreign competitors, regional gaps in adoption may be harder to close. In this sense, international knowledge spillovers appear to be subject to a form of gravity: ideas can travel across borders, but their behavioral impact weakens with economic, cultural, and informational distance. More broadly, our results speak to the extent to which the European project has succeeded in exposing firms to cross-border competitive pressure and thereby fostering innovation.

Our study relates to, and contributes to, several strands of literature. Most importantly, our findings speak to work on international knowledge diffusion and its geographic frictions (Coe and Helpman, 1995; Eaton and Kortum, 2002; Keller, 2004; Comin and Hobijn, 2010; Keller and Yeaple, 2013; Kalyani et al., 2025), as well as to research on technology adoption and firms' exposure to foreign firms and markets (Pavcnik, 2002; Amiti and Konings, 2007; Atkin et al., 2017; Alfaro-Urena et al., 2022; Bilgin et al., 2024). We contribute to this literature by providing, to our knowledge, the first experimental measurement of innovation diffusion within and across national borders.

More broadly, our paper relates to the study of strategic complementarities. While this literature is mainly theoretical, a growing set of empirical studies has begun to test for strategic complementarities in firm behavior (e.g., Bloom et al., 2013; Lin, 2023; Bilgin et al., 2024).

In particular, our study builds on experimental work by Cullen et al. (2025) and Menkhoff (2025), who study the effects of providing firms with information about peer adoption of advanced technologies. Beyond the innovation context, another notable example is Coibion et al. (2021), who study higher-order beliefs among New Zealand firms, such as firms' beliefs about other firms' inflation expectations. Our paper also relates to the growing literature on the determinants and consequences of AI adoption (e.g., Zolas et al., 2021; Bencivelli et al., 2025).

The rest of the paper is organized as follows. Section 2 presents the multi-country model of strategic complementarities in technological adoption. Section 3 describes the institutional setting and data. Section 4 discusses the research design. Section 5 discusses AI investment and firms' misperceptions about competitor AI investment. Section 6 documents the effect of the information treatment on firms' expectations. Section 7 presents the effects of the information on the firm's own expected AI investment rate. The last section concludes.

2 Conceptual model

To motivate the experimental design and the key hypotheses, we provide a simple two-country model of AI Adoption with learning and strategic complementarities. Firms have to decide whether to adopt a technology that features strategic complementarities across firms, which may capture aspects such as innovation races and network or scale externalities (where certain investments are more productive when more firms adopt them; see Cooper and John (1988)). Firms have incomplete information about the fundamentals, namely the return or productivity of the technology, but can learn from their own signals, from public signals, and also from others. Social learning is a second form of strategic complementarities across firms. This model is an extension of the global game class of models and more specifically of the model in Angeletos and Pavan (2004). The main modification is the extension to a multi-country context, so that strategic and learning complementarities (which we will sometimes refer to as knowledge spillovers) occur not only across firms but also across countries.

There are two countries, $c \in \{1, 2\}$. The countries do not have to be symmetric. For example, both countries could be similar in size, country 1 could be a small economy and country 2 a large economy, or vice versa. Indeed, in the empirical section we study a context where country 2 corresponds to a large economy (Germany, France, and Italy). In each country, there is a continuum of firms of measure one, indexed by $i \in [0, 1]$. Each country is characterized by a fundamental process, namely the return to investment, which we characterize below: on average, the technology in one country may be more productive than in another due to the institutional or production environment. Each firm chooses a single adoption level: $k_{i,c} \in \mathbb{R}$. Let:

$$K_c \equiv \int_0^1 k_{i,c}\, di \tag{1}$$

denote aggregate adoption in country c. Firms are risk-neutral and have the following profit function:

$$\pi_{i,c} = A_c\, k_{i,c} - \tfrac{1}{2} k_{i,c}^2. \tag{2}$$

The returns from adopting a technology, $k_{i,c}$, in country c feature some quadratic adjustment costs, $\frac{1}{2}k_{i,c}^2$. This is compatible with the lumpy nature of investing. Second, the returns depend on some country-specific productivity level, ω_c, but also on some knowledge spillovers that occur across countries. Those are captured by the parameters A_c, which we characterize next.

Assumption 1 (symmetric complementarities and reparameterization). Returns feature within- and cross-country complementarities:

$$A_1 = (1 - \lambda_d - \lambda_f)\omega_1 + \lambda_d K_1 + \lambda_f K_2, \tag{3}$$

$$A_2 = (1 - \lambda_d - \lambda_f)\omega_2 + \lambda_d K_2 + \lambda_f K_1. \tag{4}$$

$$\lambda_d \equiv \lambda(1 - \chi), \qquad \lambda_f \equiv \lambda\chi, \tag{5}$$

where $\lambda \in [0, 1)$ governs the overall strength of complementarities, and $\chi \in \left[0, \frac{1}{2}\right]$, which we refer to as *connectedness*, governs the share of cross-country complementarities. At one extreme, if $\chi = \frac{1}{2}$, the two economies are perfectly connected, so complementarities are equally strong between firms within the same country and between firms in different countries. At the other extreme, if $\chi = 0$, the countries are fully disconnected: complementarities exist between firms within the same country but are absent across countries. Note, once again, that complementarities are proportional to country size: if the two countries differed in size, aggregate capital would reflect this difference, which in turn would feed into different reaction functions.

It is important to put the assumption $\chi \leq \frac{1}{2}$ in the context of our empirical application. As explained in Section 3 below, firms in Europe operate under the European Single Market. Firms and goods can freely enter foreign countries, effectively neutralizing regulatory barriers. The single currency also eliminates exchange-rate-related costs. Despite the lack of regulatory barriers, there are still frictions in international trade. First, the degree of tradability varies across goods and services. Transport costs, delivery times, and logistical constraints limit firms' effective reach, particularly for bulky, perishable, or time-sensitive goods. Second, informational and cultural frictions can segment markets. Language differences, business practices, and heterogeneous consumer preferences give domestic firms an advantage in understanding and serving local demand, reducing substitutability between domestic and foreign varieties. Third, network effects and local embeddedness may favor domestic incumbents. Established user bases, supplier and distribution networks, reputation, and after-sales services can disadvantage foreign entrants. Search and matching frictions, limited awareness of foreign suppliers, and greater trust in domestic transactions reinforce this asymmetry.

Assumption 2 (symmetric correlated fundamentals; normalization). Without loss of generality, we normalize $Var(\omega_1) = Var(\omega_2) = 1$. Fundamentals are jointly normal:

$$\begin{pmatrix} \omega_1 \\ \omega_2 \end{pmatrix} \sim \mathcal{N}\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & \rho_\omega \\ \rho_\omega & 1 \end{pmatrix} \right), \qquad 0 \leq \rho_\omega < 1. \tag{6}$$

The case $\rho_\omega = 0$ provides a benchmark with no learning between fundamentals across countries. Positive correlation induces knowledge spillovers across countries. In practice, this may be due to proximity: countries can share technological advances through FDI, human capital migration, etc. (see Lind and Ramondo (2023) or Akcigit et al. (2024), among others).

Assumption 3 (symmetric information; precisions). Firms have incomplete information about the returns of the technology, but can learn about it by observing public and private signals. Both public signals are observed by all firms in both countries:

$$z_1 = \omega_1 + \tau_z^{-1/2}\varepsilon_1, \qquad z_2 = \omega_2 + \tau_z^{-1/2}\varepsilon_2, \tag{7}$$

with $\varepsilon_1, \varepsilon_2 \sim \mathcal{N}(0,1)$ i.i.d. and independent of (ω_1, ω_2). Each firm receives a private signal about its own country's fundamental:

$$x_{i,1} = \omega_1 + \tau_x^{-1/2} v_{i,1}, \qquad x_{i,2} = \omega_2 + \tau_x^{-1/2} v_{i,2}, \tag{8}$$

where $v_{i,1}, v_{i,2} \sim \mathcal{N}(0,1)$ is i.i.d. across firms and independent of all other shocks. Let $\tau_x > 0$ and $\tau_z > 0$ denote the private- and public-signal precisions, respectively, and define:

$$\kappa(\rho_\omega) \equiv \frac{1}{1 - \rho_\omega^2}. \tag{9}$$

In our empirical setting, we examine how a firm's adoption decision responds to changes in its beliefs about average adoption by domestic and foreign firms. To keep the model tractable, we study this question indirectly. Firms observe public signals about the exogenous payoff to adoption, and we analyze how their decisions vary with these signals. In the model, a higher public signal affects behavior through two channels. First, it leads firms to update their beliefs about the underlying return to adoption, increasing expected profitability (the learning channel). Second, it raises competitors' adoption levels, thereby intensifying complementarities (the competition channel). These correspond to the same mechanisms we expect empirically. Observing higher adoption by competitors both conveys information about the technology's fundamental value and increases competitive pressure to keep pace.

Firms combine their private signals with domestic and foreign signals to form posterior beliefs. Consistent with the empirical framework described in Section 6, firms update these beliefs according to Bayes' rule.[1] Next, we calculate the posterior beliefs from the country 1 perspective. A country-1 firm observes $(x_{i,1}, z_1, z_2)$. Gaussian priors and signals imply linear posteriors. Define:

$$\mathcal{A} \equiv \kappa + \tau_x + \tau_z, \qquad \mathcal{C} \equiv \kappa + \tau_z, \qquad \mathcal{D} \equiv \mathcal{A}\mathcal{C} - \rho_\omega^2 \kappa^2. \tag{10}$$

Then, given all the assumptions made above, firms' posterior beliefs in each country are given by:

$$\mathbb{E}_{i,1}[\omega_1] = a_x x_{i,1} + a_1 z_1 + a_2 z_2, \tag{11}$$

$$\mathbb{E}_{i,1}[\omega_2] = b_x x_{i,1} + b_1 z_1 + b_2 z_2, \tag{12}$$

with:

$$a_x = \frac{\mathcal{C}\tau_x}{\mathcal{D}}, \qquad a_1 = \frac{\mathcal{C}\tau_z}{\mathcal{D}}, \qquad a_2 = \frac{\rho_\omega \kappa \tau_z}{\mathcal{D}}, \tag{13}$$

$$b_x = \frac{\rho_\omega \kappa \tau_x}{\mathcal{D}}, \qquad b_1 = \frac{\rho_\omega \kappa \tau_z}{\mathcal{D}} = a_2, \qquad b_2 = \frac{\mathcal{A}\tau_z}{\mathcal{D}}. \tag{14}$$

By symmetry, a country-2 firm's posterior means have the same form with indices $1 \leftrightarrow 2$.

Next, we need to derive the equilibrium innovation policies. Given the complementarities outlined in Assumption 1, each firm will choose its own adoption policy to maximize its profits (see 2). This will result in a reaction function in which the innovation policy of firm i in country c depends on the investments of all other firms in its own country, as well as on investments by firms in other countries. The Nash equilibrium will then result from the intersection of the reaction functions. To guarantee existence and uniqueness, we introduce some additional assumptions.

To obtain the Nash equilibrium adoption policies, we first derive the first-order conditions of firms choosing adoption to maximize 2 in each country. This leads to the individual reaction functions:

$$k_{i,1} = \mathbb{E}_{i,1}[A_1] = (1-\lambda)\mathbb{E}_{i,1}[\omega_1] + \lambda(1-\chi)\mathbb{E}_{i,1}[K_1] + \lambda\chi\, \mathbb{E}_{i,1}[K_2], \tag{15}$$

$$k_{i,2} = \mathbb{E}_{i,2}[A_2] = (1-\lambda)\mathbb{E}_{i,2}[\omega_2] + \lambda(1-\chi)\mathbb{E}_{i,2}[K_2] + \lambda\chi\, \mathbb{E}_{i,2}[K_1]. \tag{16}$$

We can rely on the assumption of a symmetric equilibrium and guess the adoption policy

[1] In dynamic extensions of this class of models, Bayesian learning can be carried out using a Kalman filter, see Huo and Takayama (2025).

functions as follows:

$$k_{i,1} = \varphi_x x_{i,1} + \varphi_d z_1 + \varphi_f z_2, \tag{17}$$

$$k_{i,2} = \varphi_x x_{i,2} + \varphi_d z_2 + \varphi_f z_1. \tag{18}$$

The individual guesses for the policy functions in 18 allow us to derive aggregate investment, which we then use to obtain the firms' equilibrium adoption levels:

$$K_1 = \varphi_x \omega_1 + \varphi_d z_1 + \varphi_f z_2, \qquad K_2 = \varphi_x \omega_2 + \varphi_f z_1 + \varphi_d z_2. \tag{19}$$

The equilibrium is fully summarized by $(\varphi_x, \varphi_d, \varphi_f)$ in (A.2)–(A.5). These coefficients depend on primitives only through $(\lambda, \chi, \rho_\omega, \tau_x, \tau_z)$ via the posterior weights $(a_x, a_1, a_2, b_x, b_2)$ in (13)–(14). In Appendix A we derive the closed-form solutions for the $(\varphi_x, \varphi_d, \varphi_f)$ coefficients and for the response gap $\varphi_d - \varphi_f$. This statistic is a crucial part of the proposition derived below as it dictates the response gap to local versus global competition. We discuss this more in detail below.

Existence and Uniqueness. Next, we report the parameter regions under which the equilibrium exists and is unique, as well impose some restrictions to align the equilibrium properties with realistic observations.

Assumption 4. Stability/ uniqueness. Since $\lambda < 1$ and $\chi \in [0, 1/2]$, the matrix in Step 2 below is automatically invertible.[2] We therefore impose the single remaining condition needed for a unique symmetric linear equilibrium:

$$1 - \lambda(1-\chi)a_x - \lambda\chi\, b_x > 0, \tag{20}$$

where a_x and b_x are defined in (13)–(14).

Assumption 5 (cross-country complementarities are small enough). Assume χ is low enough that:

$$\chi < \bar{\chi} \equiv \frac{a_1 - a_2}{\lambda\Big(b_x(a_1 - a_2) + a_x(b_2 - a_2)\Big)}. \tag{21}$$

This restriction ensures that the informational advantage of the domestic public signal for ω_1 (measured by $a_1 - a_2 > 0$) is not overturned by amplification of the foreign signal through cross-country complementarities.

Effects Decomposition. Given the closed form solution for the policy functions, obtained by solving for the parameters $(\varphi_x, \varphi_d, \varphi_f)$, we can decompose the effect of the domestic and foreign signals in four different components: the learning effects (and subsequently between domestic learning and foreign learning) and the competition effects (domestic and

[2] Indeed, with $\lambda_d = \lambda(1-\chi)$ and $\lambda_f = \lambda\chi$, one has $(1-\lambda_d)^2 - \lambda_f^2 = (1-\lambda)^2 + 2\lambda\chi(1-\lambda) > 0$.

foreign). From (15), substituting (11)–(12) and the equilibrium expressions for $\mathbb{E}_{i,1}[K_1]$ and $\mathbb{E}_{i,1}[K_2]$, we can write:

$$
\begin{aligned}
k_{i,1} = x_{i,1}\Big[a_x\Big((1-\lambda)+\lambda\varphi_x(1-\chi)\Big)+\lambda\varphi_x\chi\, b_x\Big] \\
+ \quad \underbrace{(1-\lambda)a_1 z_1}_{\text{domestic learning from } z_1} \quad + \quad \underbrace{(1-\lambda)a_2 z_2}_{\text{foreign learning from } z_2} \\
+ \underbrace{\lambda\Big[(1-\chi)(\varphi_x a_1+\varphi_d)+\chi(\varphi_x b_1+\varphi_f)\Big]z_1}_{\text{domestic competition effect of } z_1} \\
+ \underbrace{\lambda\Big[(1-\chi)(\varphi_x a_2+\varphi_f)+\chi(\varphi_x b_2+\varphi_d)\Big]z_2}_{\text{foreign competition effect of } z_2}
\end{aligned}
\tag{22}
$$

In particular, $\partial k_{i,1}/\partial z_1 = \varphi_d$ and $\partial k_{i,1}/\partial z_2 = \varphi_f$.

The next proposition shows that domestic public news has a larger effect on domestic adoption than foreign public news:

Proposition 1 (Domestic public signal has a larger marginal effect)**.** *Under Assumptions 1–5,*

$$\frac{\partial k_{i,1}}{\partial z_1} > \frac{\partial k_{i,1}}{\partial z_2}. \tag{23}$$

(See Appendix A for the proof.)

The intuition is simple. Domestic news is, by construction, the most relevant piece of information for domestic conditions, while foreign news matters for domestic conditions only indirectly (to the extent that the two countries move together). In addition, our restriction on connectedness rules out cases in which spillovers from foreign adoption are so strong that foreign news would be amplified enough to dominate domestic news. As a result, domestic news both teaches firms more about the environment they face and has more scope to shape expectations about what their main competitors are doing, making its overall effect larger.

The next proposition examines how connectedness affects the difference between the impact of domestic and foreign public news on domestic adoption:

Proposition 2 (The response gap shrinks with cross-country complementarities)**.** *Maintain Assumptions 1–5, so that* $\varphi_d - \varphi_f > 0$*. The gap*

$$\varphi_d - \varphi_f = \frac{\partial k_{i,1}}{\partial z_1} - \frac{\partial k_{i,1}}{\partial z_2} \tag{24}$$

is decreasing in χ *on* $[0, \bar{\chi})$ *provided* $d_0 \geq \frac{1}{2}$*, where* $d_0 \equiv 1 - \lambda(1-\chi)a_x$*.*

(See Appendix A for the proof.)

Intuitively, when the economy is less open (lower χ), domestic firms are less influenced by adoption abroad, so foreign news has less scope to spill over into domestic incentives through the cross-country complementarity (or "foreign competition") channel. By contrast, domestic news continues to matter both because it is more directly relevant for domestic fundamentals and because it coordinates expectations about what domestic competitors are doing. As a result, reducing connectedness widens the gap between the domestic- and foreign-signal effects.

The next proposition focuses on how the domestic–foreign gap varies with the correlation of country fundamentals:

Proposition 3 (The response gap shrinks with correlated fundamentals)**.** *When $\chi = 0$, the gap $\varphi_d - \varphi_f$ is weakly decreasing in $\rho_\omega \in [0,1)$ and strictly decreasing on $\rho_\omega \in (0,1)$. Moreover, by continuity there exists $\varepsilon > 0$ such that for all $\chi \in [0, \varepsilon)$ the gap remains weakly decreasing on $[0,1)$ and strictly decreasing on $(0,1)$.*

(See Appendix A for the proof.)

Intuitively, the domestic signal has a stronger effect because it is constructed to be more informative about domestic fundamentals: z_1 is a direct signal of ω_1, while z_2 matters for ω_1 only insofar as the two fundamentals move together. This intuition is most transparent when connectedness is zero (or close to zero), because foreign news then affects domestic adoption mainly through what it teaches firms about domestic fundamentals.[3]

3 Institutional Context and Data

We focus on firms across twelve euro area countries, drawing on data from the Survey on Access to Finance of Enterprises (SAFE) conducted across Europe. This is a particularly relevant context, as innovation has emerged as a key policy priority throughout the region. After an initial period of growth and convergence in the euro area, productivity growth and innovation have been comparatively weak in recent years, especially in the aftermath of the financial and sovereign-debt crises. A recent report by Draghi (2024) brought renewed attention to these concerns, emphasizing Europe's slow productivity growth, its innovation gap, and the difficulty of translating ideas into scalable commercial success. Indeed, the report explicitly mentioned the need to integrate AI into European industries so that they can stay at the frontier.

[3]When connectedness is high, however, foreign news also moves domestic incentives through strategic spillovers: z_2 shifts beliefs about foreign adoption K_2, which feeds back into domestic returns through cross-country complementarities, and z_1 can simultaneously become more informative about foreign adoption when fundamentals are correlated. These additional feedback and coordination effects can complicate the comparison, which is why we emphasize the low-connectedness case (or small deviations from it).

This policy debate naturally raises the question of whether competitive pressure is sufficient to drive innovation. Since the creation of the European Single Market in 1993,[4] firms have, in principle, been exposed to competition from both domestic and foreign markets within Europe. Yet this integration has not fully eliminated market fragmentation. Particularly relevant for our setting, Letta (2024) emphasizes that Europe still lacks a truly integrated market for knowledge, innovation, and research. Our survey experiment is designed in part to shed light on exactly this issue: whether perceived competition helps explain the spread of innovation across European firms.

Lack of funding is less likely to be among the main constraints. In 2021, the European Union approved the Next Generation EU plan to support post-pandemic recovery and longer-run growth. Its core component, the Recovery and Resilience Facility (RRF), provides more than 577 billion euros in loans and grants to finance national reforms and investments.[5] The targeted areas include green transition, digital transformation, health, and education. More specifically for AI, in January 2024 the European Commission launched the AI innovation package to support AI startups and SMEs and to promote the development of trustworthy AI consistent with EU rules and values.

These policy efforts coincided with a sharp increase in AI investment in Europe in 2025, including a reported record 17.5 billion in AI venture funding. Investment has been especially concentrated in a few hubs, such as France and Germany. This raises a second question that is central to our study: even if Europe succeeds in mobilizing investment, how can innovation diffuse more evenly across countries and firms rather than remaining concentrated in a small set of locations? Despite harmonized regulation through the EU AI Act and centralized funding initiatives, large disparities remain. AI use is higher in countries such as Sweden and the Netherlands than in much of Eastern Europe. There is also substantial heterogeneity by firm size: while AI use is widespread among large firms, it remains much less common among smaller ones.

It is also worth noting that, by international standards, the European competition-policy environment is relatively strict. The European Commission is well known for antitrust enforcement aimed at limiting monopoly power, excessive concentration, and distortive support to national firms. In principle, this should create an environment conducive to entry, experimentation, and diffusion of innovation among smaller firms as well.

[4] The European Single Market was created to ensure the “four freedoms”—free movement of goods, services, people and capital.

[5] For more information on the Recovery and Resilience Facility, see European Commission (2021).

3.1 The SAFE Survey

The experiment was embedded in the Survey on Access to Finance of Enterprises (SAFE), a firm-level survey conducted in Europe by the European Central Bank and the European Commission. Originally launched in 2009, the survey was conducted twice a year, targeting large samples of firms across European member states. Starting in 2024, its frequency was increased to quarterly. The composition of countries included in the survey varies by round: three rounds (those conducted in Q1, Q2, and Q4) focus on the twelve largest euro area countries, while the Q3 round encompasses all EU member states along with selected neighboring countries.

Since its inception the SAFE was designed to measure firms' financing conditions and credit frictions as well as a wider range of other topics, including employment, turnover, and other business conditions. Two ad hoc modules conducted in Q2–2025 and Q4–2025 covered questions on AI. The Q2–2025 module measured firms' past and current AI investment: firms reported whether they had ever invested or were currently investing in AI, and, if so, the average share of overall investment allocated to AI over the previous 12 months. The Q4–2025 module measured current AI use and forward-looking investment plans: firms reported their current intensity of AI use and the share of total investment they expected to allocate to AI over the next 12 months.

Our sample includes firms from the twelve major EU countries: Austria, Belgium, Germany, Spain, Finland, France, Greece, Ireland, Italy, the Netherlands, Portugal, and Slovakia. This selection represents a mix of both large and small economies within the euro area. The survey is broadly representative of the European firm landscape. It covers micro (1 to 9 employees), small (10 to 49 employees), medium-sized (50 to 249 employees), and large firms (250 or more employees). In addition to firm size, it includes information on sector, country, firm age, financial autonomy, and ownership. Sectoral stratification is based on the one-digit NACE classification of non-financial sectors, focusing on four major sectors: construction, industry, trade and services.

The fieldwork typically lasts between three and five weeks. Interviews are conducted primarily by telephone using computer-assisted telephone interviewing (CATI). Since the 2014H1 round, respondents have also had the option of completing the questionnaire online through computer-aided web interviewing (CAWI). The respondent within each firm is a senior executive, such as a general manager, financial director, or chief accountant.

4 Research Design

4.1 Survey Design

The experiment was embedded in the SAFE survey round covering the fourth quarter of 2025. The survey instrument for our tailored module is presented in Appendix C and outlined below.[6]

The module begins by eliciting firms' current AI use, using a scale that distinguishes between no current use, very infrequent or experimental use, moderate use, and significant use. To better understand AI use patterns, firms reporting some use of AI were then asked to identify the two main motivations for using AI, including cost reduction, improvements in core or non-core processes, product and service expansion, and support for research and innovation. Conversely, firms with no or limited use of AI were asked to report the two main barriers to using AI, such as lack of skills, high costs, data or ethical concerns, distrust of AI outputs, or technological incompatibility.

Next, we elicited firms' beliefs about competitors' AI investment in the same country as well as in the big-3 European countries (Germany, France, and Italy). Specifically, we asked: "Thinking about other firms like yours, in the same sector and of a similar size, what percentage of these do you think had invested in AI technologies up to June 2025?" Respondents could respond any number between 0% and 100%, up to one decimal place. We refer to these two measures as the *domestic prior beliefs* and the *foreign prior beliefs*. Immediately after eliciting these prior beliefs, half of the firms were randomly assigned to receive information about the share of competitors that had invested in AI. More precisely, the surveyors read the following message:

> "You have been randomly selected to receive the following information about the behaviour of potential competitors in your sector and size class.
> In your country, [X%] of firms similar to yours had made investments in AI technologies up to June 2025.
> In Germany, France and Italy, [Y%] of firms similar to yours had made investments in AI technologies up to June 2025."

If needed, the surveyors were instructed to provide the following information: "These figures come from a previous survey round, when firms were asked about investment in AI technologies up to June 2025." We refer to the two pieces of information provided in this

[6]For the complete questionnaire—including both our module and the remaining sections—see European Central Bank (2025).

experiment—one about domestic AI investment and one about AI investment in the big-3—as the *signals*. These signals were constructed using responses from the previous survey wave (Q2–2025). For completeness, Appendix D reproduces the Q2–2025 survey module. The key input is the question asking firms about their investments in AI. We grouped respondents from the Q2–2025 survey into country-sector-size cells and computed investment rates for each cell. The data were structured around the four sectors (construction, industry, trade, and services), with firms further classified by size into two categories: those with fewer than 50 employees and those with 50 or more employees. This procedure resulted in 96 unique country-sector-size clusters, with an average of 56 respondents per cluster, a minimum of 3 respondents, and a maximum of 252 respondents. Tables B.2 and B.3 report the resulting signals used in the information treatment.

Returning to the Q4–2025 wave, after the information treatment, firms were asked to state their expectations regarding the share of competitors that would invest in AI technologies in the next 12 months, both domestically and in the big-3. We refer to these two measures as the *domestic posterior beliefs* and the *foreign posterior beliefs*. The questions were presented to all firms, regardless of treatment status. These posterior beliefs allow us to examine whether firms updated their expectations about future competitor investment in response to information about current competitor investment. Finally, after eliciting posterior beliefs, we asked all respondents about their future plans to invest in AI technologies: "Over the next 12 months, what percentage of your firm's total investment do you expect to allocate to AI technologies?" This question allows us to measure the firm's *own expected AI investment rate*.

One design choice merits further discussion. While our survey elicits beliefs and provides information about firms in the big-3 foreign economies, an alternative design would have elicited and randomized information about AI investment in one of the other eleven countries, selected at random. We focused on the big-3 because we wanted to give information about foreign firms the largest possible chance to affect firms' beliefs and investment decisions. Our assumption was that if any foreign peers were likely to be salient, relevant, or influential for firms in our sample, they would be firms from the largest economies. Under that assumption, the fact that we find no effect of beliefs about AI investment in the big-3 suggests that information about a randomly selected foreign country would be even less likely to move firms' own investment plans.

4.2 Survey Implementation

A total of 5,067 firms participated in the Q4–2025 SAFE wave, but roughly 1,751 firms did not answer at least one of the key questions in our module: domestic prior beliefs, foreign prior beliefs, domestic posterior beliefs, foreign posterior beliefs, or the firm's own expected AI investment rate. To keep the sample consistent throughout the analysis, we exclude firms with missing data on any of these outcomes. This leaves an analysis sample of about 3,316 firms. Unless noted otherwise, all results reported below are based on this sample.

4.3 Descriptive Statistics and Randomization Balance

Column (1) of Table 1 summarizes the characteristics of the firms in the sample. Most firms are at least 10 years old (91% of firms) and have on average 113 employees, although more than half employ fewer than 15 workers. This distribution is in line with expectations for Europe, where small and old firms are common. To assess randomization balance, columns (2) and (3) of Table 1 report mean pre-treatment characteristics for the control and treatment groups, respectively. Column (4) presents the p-value for the test of equal means across the two groups for each characteristic. The results indicate that pre-treatment characteristics are well balanced between treatment and control groups, consistent with successful random assignment.

5 AI Investment, Use, and Misperceptions

5.1 Baseline AI Investment and Use

Panel A of Figure 1 reports the average signal on AI investment as of June 2025 for each of the twelve countries. Within each country, the exact signal shown to a given firm depends on its country-sector-size cell, so the country-level values in the figure are summary averages—for the full list of signals, see Table B.2. For reference, the dashed horizontal line from Panel A corresponds to the big-3 average. AI investment appears to be substantially more prevalent in some European countries than in others, with wide variation across the distribution.[7] The share of firms investing in AI ranges from 20% in Greece to 48% in Austria. The big-3 average is 25%, almost exactly in the middle of the cross-country distribution, and is itself composed of 35% in Germany, 20% in France, and 19% in Italy.

As a complementary measure, we can also examine AI use directly. The Q4–2025 survey began with a question on current AI use, and Panel B of Figure 1 reports the distribution of responses. The figure shows the share of firms that, at the end of 2025, reported no use, very infrequent use, moderate use, or significant use. Overall, 70% of firms report some use of AI (hence 30% report no current AI use), and 38% report moderate or significant use (Ferrando

[7] The signal is computed as the share of firms investing in AI. If we instead focus on the average investment rate—reported in Panel B of Figure B.1—we find similarly wide cross-country variation.

et al., 2026). Furthermore, AI use is concentrated among large firms (up to 45%), but also young firms (up to 56%). Here too there is substantial heterogeneity across countries. For example, in Italy 49% of firms report no AI use, whereas in the Netherlands that share is only 17%. The figure also shows that, among firms reporting some AI use, intensity varies considerably: very infrequent and moderate use are both more common than significant use. In the pooled sample, 32% report very infrequent use, 29% report moderate use, and only 9% report significant use.

Next, we summarize the reasons firms report for using AI. Firms that reported using AI were asked to select the two main reasons for use from a predetermined list. The results are presented in Panel C of Figure 1. The most frequently selected reason is improving core processes (28%), followed by improving non-core processes (25%), supporting R&D and innovation (14%), reducing personnel costs (12%), and expanding the range of products (9%). For comparison, among firms that do not use AI or use it less intensively, the two most common barriers are lack of perceived usefulness and lack of AI skills; these results are reported in Panel D of Figure B.1.

The patterns of AI use and investment in the twelve EU countries we study are broadly comparable to those in the United States, both in overall levels and in the unevenness of diffusion across firms. Several recent surveys, mostly fielded during 2025–2026, report evidence on firms' AI adoption in the United States. For instance, Baslandze et al. (2026) use a survey of 750 CFOs and report that roughly 40% of firms did not invest in AI, that adoption is concentrated among large firms in high-skill sectors, and that the main reason for non-adoption is a lack of skilled labor, cited by 36% of firms. These figures compare with a 30% no-current-use rate and a 20% share citing skill shortages among firms with no or limited AI use in Europe, as reported by Ferrando et al. (2026). Baslandze et al. (2026) also document substantial heterogeneity, similar to the patterns documented above for Europe. Lower reported AI-use rates for the United States are reported by Bonney et al. (2026), who use the recent Census survey of AI adoption and find that AI use rates in the United States are around 18%, or 23% when employment-weighted. They also find that AI use is concentrated among large and highly productive firms. Finally, Yotzov et al. (2026) field four harmonized questions across the United States, the United Kingdom, Germany, and Australia. They find that 69% of firms currently use some form of AI, almost identical to the 70% reported for Europe, and that AI use is highest among larger and more productive firms.

5.2 Misperceptions about Competitors' AI Investment

In Panel A of Figure 2, we measure prior misperceptions as the difference between the signal about competitors' AI investment and the corresponding prior belief. The dark-gray bars correspond to domestic prior beliefs. A value of zero would indicate accurate expectations, but relatively few firms are near that point. Instead, misperceptions are both large and systematically directional: while a small minority of firms overestimate domestic AI investment by competitors (negative values), most firms underestimate it (positive values), in some cases by as much as 50 pp. This pattern indicates that firms typically perceive domestic competitor AI investment as substantially lower than it actually is. The same pattern emerges for foreign prior beliefs—shown by the light-gray bars—which are likewise often far from the signal and skewed toward underestimation. Thus, both domestic and foreign prior beliefs exhibit large and systematically biased misperceptions of similar direction and magnitude.

When interpreting this evidence, it is important to note one caveat. Ideally, true misperceptions would be measured by comparing prior beliefs with actual AI investment shares. In our setting, however, actual AI investment shares are not observed without error. We instead use a signal that may contain measurement error, as it is constructed from a finite sample of respondents and is therefore subject to sampling variation. As a result, the gaps reported in Panel (a) of Figure 2 may partly capture measurement error in the signal.[8] However, given the size of the gaps and their systematic downward bias, our preferred interpretation is that measurement error likely plays a minor role.

Panel B of Figure 2 examines the relationship between misperceptions in domestic prior beliefs and foreign prior beliefs. The figure shows a strong positive correlation: firms that underestimate domestic competitor AI investment also tend to underestimate foreign competitor AI investment, and the same is true for overestimation. This suggests that misperceptions are not isolated to a single reference group, but instead reflect a broader tendency to misjudge the diffusion of AI among competitors. At the same time, the relationship is far from perfect. The dispersion around the upward-sloping pattern reveals substantial independent variation in the two misperception measures. This matters for identification because it implies that the data contain enough separate variation in domestic and foreign beliefs to help disentangle their causal effects.

Figure 3 provides a more detailed breakdown of these misperceptions by country. The x-axis lists the 12 countries in our sample. For each country, the light-gray bar shows the actual share of firms that invested in AI, the medium-gray bar shows the average domestic prior belief, and the dark-gray bar shows the average foreign prior belief. The y-axis reports percentages. For reference, the horizontal dashed line marks the actual average investment rate across the big-3 countries.

Figure 3 yields several additional insights. First, the basic pattern holds within each of the

[8] Additionally, prior beliefs themselves may be measured with some error too. This could arise, for instance, from subjective interpretations of the survey question.

12 countries: firms systematically underestimate both domestic and foreign competitor AI investment. Second, the underestimation is more severe in countries where actual investment is higher. There is almost no relationship between the average domestic prior belief in a country and the actual investment rate. True investment rates range from about 20% to 48%, whereas domestic prior beliefs are compressed into a much narrower interval, roughly 12% to 22%.

Figure 3 also suggests that the big-3 benefit from a kind of reputation mirage. In every country, firms believe that the big-3 invest more than their own country, even though that is not always true: in practice, six countries invest more than the big-3 average, one invests about the same, and five invest less. Firms therefore seem to have a poor sense of their own country's relative standing in AI investment.[9] More broadly, the figure points to a center-periphery misconception: firms may perceive AI investment as concentrated in a few technologically advanced economies even though, in reality, it is more widely dispersed.

Our evidence on firms' misperceptions about competitor AI investment aligns with a broader literature documenting that firms—including large ones—often have imperfect information about decision-relevant variables. Most closely related, Cullen et al. (2025) show that Italian firms have large misperceptions about the adoption of advanced technologies (AI and robotics) among their competitors. Cullen et al. (2022) provide evidence that firms have inaccurate beliefs about wages at competing employers. Kim (2025) document limited knowledge of competitors' prices. Finally, Coibion et al. (2018) find substantial misperceptions about aggregate economic conditions, including inflation and economic growth.

6 Effect of the Information Treatment on Firms' Beliefs

6.1 Average Treatment Effects

We next examine whether the information treatment shifted firms' posterior beliefs about competitors' AI investment. Since Section 5 showed that firms, on average, underestimated both domestic and foreign competitor AI investment at baseline, the natural prediction is that providing accurate information should move posterior beliefs upward on average. Panel A of Figure 4 provides a first visual test of this mechanism. The first pair of bars corresponds to average domestic posterior beliefs. The gray bars represent firms in the control group, while the red bars represent firms in the treatment group. The average domestic posterior

[9]Misperceptions of relative standing are common in other contexts as well, such as relative income (see e.g., Cruces et al., 2013; Fehr et al., 2022).

belief rises from 23.59% in the control group to 27.89% in the treatment group, a difference of 4.3 pp that is statistically significant ($p<0.001$). A similar pattern emerges for foreign posterior beliefs: the average rises from 28.17% to 31.51%, corresponding to a statistically significant treatment effect of 3.3 pp ($p<0.001$). While these estimates summarize the average treatment effects, Panels C and D of Figure B.2 provide a more detailed view by comparing the full distribution of posterior beliefs in the treatment and control groups.

Panel B of Figure 4 provides a falsification test.[10] It has the same structure as Panel A, except that it uses firms' prior beliefs instead of their posterior beliefs. Since the prior beliefs were elicited before the provision of information, the treatment should have no effect on them. That is exactly what Panel B shows. For domestic prior beliefs, the average is 15.83% in the control group and 15.98% in the treatment group, a small difference of 0.1 pp that is statistically insignificant ($p=0.833$). For foreign prior beliefs, the corresponding means are 20.64% and 20.52%, implying a difference of -0.1 pp that is statistically insignificant ($p=0.877$). In other words, the upward shifts in posterior beliefs documented in Panel A do not appear to be driven by pre-existing differences across treatment status.

6.2 Bayesian Updating

So far, we have shown that the treatment had significant effects on posterior beliefs. In this section, we study these treatment effects in more detail through the lens of a simple Bayesian learning model. The model helps interpret the magnitude of the updating and also provides the basis for the first stage of the 2SLS specification presented in Section 7.2.

We adopt the simple Bayesian learning framework of Cavallo et al. (2017). Let subscript i denote firms and t time horizons. Let superscript $k \in \{d, f\}$ denote whether the belief corresponds to domestic or foreign competitors, respectively. For the sake of brevity, we start with a simpler model with only domestic beliefs and then introduce foreign beliefs. Let $s^d_{i,t}$ denote firm i's domestic prior belief, that is, its belief about the share of domestic competitors that had invested in AI up to June 2025, and let $s^d_{i,t+1}$ denote firm i's domestic posterior belief, that is, its expectation of the share of domestic competitors that will invest in the next 12 months. We assume that firms form expectations about future competitor investment by extrapolating from their perceptions of current competitor investment:

$$s^d_{i,t+1} = \mu + \beta^d \cdot s^d_{i,t}, \tag{25}$$

[10] As additional falsification tests, Panels A and B of Figure B.2 provide a more detailed view by comparing the full distribution of prior beliefs in the treatment and control groups. In addition, Figure B.3 reproduces the same analysis using prior misperceptions rather than raw prior beliefs.

where β^d measures the extent to which the domestic prior belief passes through to the domestic posterior belief. In our setting—marked by sustained growth in investment rates—it is natural to expect $\beta^d > 1$. That is, firms anticipate that future competitor investment will exceed current competitor investment.[11]

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

In this case, the firm's updated domestic prior belief, $s^d_{i,t}$, may depend on both its untreated domestic prior belief ($s^{d,0}_{i,t}$) and the experimental signal ($s^{d,T}_{i,t}$). Under the assumptions of a Bayesian model with Gaussian priors and signals,[12] the updated domestic prior belief after observing the signal is given by:

$$s^d_{i,t} = \alpha^d \cdot s^{d,T}_{i,t} + (1 - \alpha^d) \cdot s^{d,0}_{i,t} \quad \text{if } T_i = 1, \tag{26}$$

where $\alpha^d \in [0, 1]$ represents the weight placed on the new information relative to the prior, and depends on the relative precision of the prior and the signal. Combining equations (25) and (26) yields:

$$s^d_{i,t+1} = \mu + \alpha^d \cdot \beta^d \cdot \underbrace{\left(s^{d,T}_{i,t} - s^{d,0}_{i,t}\right)}_{\text{Prior Gap}} + \beta^d \cdot s^{d,0}_{i,t} \quad \text{if } T_i = 1 \tag{27}$$

The central prediction of the model is that, among treated individuals, belief revisions are a linear function of the prior gap. Intuitively, respondents who initially overestimated the true AI investment share should revise downward after receiving the signal; those who underestimated should revise upward; and those whose priors were accurate should display little or no updating. The strength of updating is governed by the product of the two key parameters: the learning weight (α^d) and the extrapolation parameter (β^d).

A potential concern with directly estimating equation (27) is that respondents may appear to move their beliefs toward the signal even if they were not provided with the signal. For example, being asked the same question twice may prompt closer reflection, correction of typographical errors, or mechanical adjustments toward more plausible values. Such forces would generate spurious updating. To isolate genuine learning from these artifacts, we exploit

[11] Whether β^d is greater than, equal to, or less than 1 depends on the context. In particular, it hinges on whether the forecasted variable is expressed in levels or growth rates. For example, when forecasting home prices in levels, one would typically expect future prices to exceed current prices, implying a $\beta > 1$. By contrast, if forecasting price growth, one might expect a $\beta = 1$ if growth is projected to continue at the same rate, or a $\beta < 1$ if mean reversion is anticipated.

[12] Specifically, priors and signals are assumed to be normally distributed, and their variances are independent of the prior mean. See Section C of Cavallo et al. (2017) for details.

the randomized assignment of information and estimate:

$$s^d_{i,t+1} = \gamma^d_0 + \gamma^d_1 \cdot T_i \cdot \left(s^{d,T}_{i,t} - s^{d,0}_{i,t}\right) + \gamma^d_2 \cdot \left(s^{d,T}_{i,t} - s^{d,0}_{i,t}\right) + \gamma^d_3 \cdot s^{d,0}_{i,t} + \epsilon^d_{i,t} \tag{28}$$

The coefficient of interest is γ^d_1, which captures whether the relationship between belief revisions and prior gaps is steeper for treated individuals ($T_i = 1$) than for controls ($T_i = 0$). In terms of equation (27), γ^d_1 identifies the product $\alpha^d \cdot \beta^d$. The parameter γ^d_2 captures spurious updating, while γ^d_3 identifies β^d. Taking the ratio $\frac{\gamma^d_1}{\gamma^d_3}$ allows us to recover α^d separately.

We now incorporate foreign beliefs. Let $s^f_{i,t}$ denote the foreign prior belief, $s^f_{i,t+1}$ the foreign posterior belief, $s^{f,0}_{i,t}$ the untreated foreign prior belief, and $s^{f,T}_{i,t}$ the signal about foreign competitors. Note that T_i does not have a domestic or foreign superscript because we randomized subjects to either receive the signal about both domestic and foreign competitors or neither.

For the sake of simplicity, we use a model in which the cross-learning arises at the *extrapolation stage* rather than at the *learning stage*. Specifically, the domestic and foreign forecasts depend on both current domestic and foreign beliefs:

$$s^d_{i,t+1} = \mu^d + \beta^{d,d}\, s^d_{i,t} + \beta^{d,f}\, s^f_{i,t}, \tag{29}$$

$$s^f_{i,t+1} = \mu^f + \beta^{f,f}\, s^f_{i,t} + \beta^{f,d}\, s^d_{i,t} \tag{30}$$

Here, $\beta^{d,d}$ and $\beta^{f,f}$ capture within-market extrapolation, while $\beta^{d,f}$ and $\beta^{f,d}$ capture cross-market spillovers in extrapolation. In turn, the Bayesian updating equations for foreign and domestic beliefs are:[13]

$$s^d_{i,t} = s^{d,0}_{i,t} + \alpha^d \left(s^{d,T}_{i,t} - s^{d,0}_{i,t}\right) \quad \text{if } T_i = 1 \tag{31}$$

$$s^f_{i,t} = s^{f,0}_{i,t} + \alpha^f \left(s^{f,T}_{i,t} - s^{f,0}_{i,t}\right) \quad \text{if } T_i = 1 \tag{32}$$

We now combine updating and extrapolation to obtain the forecasting equations. Substituting (31) and (32) into (29), we obtain:

[13] A more general model would allow $\alpha^{d,d}$ to capture how strongly the domestic posterior responds to the domestic signal, and $\alpha^{d,f}$ how strongly it responds to the foreign signal (with $\alpha^{f,f}$ and $\alpha^{f,d}$ defined analogously for foreign beliefs). However, this added flexibility would make the model substantially more complicated for little gain.

$$s^{d}_{i,t+1} = \mu^{d} + \beta^{d,d} s^{d,0}_{i,t} + \beta^{d,f} s^{f,0}_{i,t} + \beta^{d,d}\alpha^{d}\left(s^{d,T}_{i,t} - s^{d,0}_{i,t}\right) + \beta^{d,f}\alpha^{f}\left(s^{f,T}_{i,t} - s^{f,0}_{i,t}\right) \quad \text{if } T_i = 1 \tag{33}$$

Analogously, substituting (32) and (31) into (30):

$$s^{f}_{i,t+1} = \mu^{f} + \beta^{f,f} s^{f,0}_{i,t} + \beta^{f,d} s^{d,0}_{i,t} + \beta^{f,f}\alpha^{f}\left(s^{f,T}_{i,t} - s^{f,0}_{i,t}\right) + \beta^{f,d}\alpha^{d}\left(s^{d,T}_{i,t} - s^{d,0}_{i,t}\right) \quad \text{if } T_i = 1 \tag{34}$$

After allowing for spurious updating, we end up with the following system of two equations:

$$\begin{aligned} s^{d}_{i,t+1} = \gamma^{d}_{0} &+ \gamma^{d}_{1d}\, T_i \left(s^{d,T}_{i,t} - s^{d,0}_{i,t}\right) + \gamma^{d}_{1f}\, T_i \left(s^{f,T}_{i,t} - s^{f,0}_{i,t}\right) \\ &+ \gamma^{d}_{2d}\left(s^{d,T}_{i,t} - s^{d,0}_{i,t}\right) + \gamma^{d}_{2f}\left(s^{f,T}_{i,t} - s^{f,0}_{i,t}\right) + \gamma^{d}_{3}\, s^{d,0}_{i,t} + \gamma^{d}_{4}\, s^{f,0}_{i,t} + \epsilon^{d}_{i,t}. \end{aligned} \tag{35}$$

$$\begin{aligned} s^{f}_{i,t+1} = \gamma^{f}_{0} &+ \gamma^{f}_{1f}\, T_i \left(s^{f,T}_{i,t} - s^{f,0}_{i,t}\right) + \gamma^{f}_{1d}\, T_i \left(s^{d,T}_{i,t} - s^{d,0}_{i,t}\right) \\ &+ \gamma^{f}_{2f}\left(s^{f,T}_{i,t} - s^{f,0}_{i,t}\right) + \gamma^{f}_{2d}\left(s^{d,T}_{i,t} - s^{d,0}_{i,t}\right) + \gamma^{f}_{3}\, s^{f,0}_{i,t} + \gamma^{f}_{4}\, s^{d,0}_{i,t} + \epsilon^{f}_{i,t}. \end{aligned} \tag{36}$$

We can recover the structural parameters as follows. From equation (35) we can recover: $\beta^{d,d} = \gamma^{d}_{3}$, $\beta^{d,f} = \gamma^{d}_{4}$, $\alpha^{d} = \frac{\gamma^{d}_{1d}}{\gamma^{d}_{3}}$ and $\alpha^{f} = \frac{\gamma^{d}_{1f}}{\gamma^{d}_{4}}$. And from equation 36 we can recover: $\beta^{f,f} = \gamma^{f}_{3}$, $\beta^{f,d} = \gamma^{f}_{4}$, $\alpha^{d} = \frac{\gamma^{f}_{1d}}{\gamma^{f}_{4}}$ and $\alpha^{f} = \frac{\gamma^{f}_{1f}}{\gamma^{f}_{3}}$. Moreover, the model delivers two independent estimates of α^{d} and two independent estimates of α^{f}. This allows us to test the validity of the model through the null hypothesis that the corresponding coefficients are equal across regressions.

Before presenting the estimates from this learning model, we provide semi-parametric evidence using binned scatterplots. Panels A and B of Figure 5 correspond to domestic and foreign beliefs, respectively. For domestic beliefs, equation (35) predicts that the relationship between prior misperceptions and posterior beliefs should be steeper in the treatment group than in the control group. In Panel A, the y-axis reports domestic posterior beliefs ($s^{d}_{i,t+1}$), while the x-axis reports domestic prior misperceptions ($s^{d,T}_{i,t} - s^{d,0}_{i,t}$). For consistency with equation (35), the binned scatterplot residualizes the other covariates in that specification. The Bayesian model predicts that domestic posterior beliefs should move with the domestic prior gap: firms that overestimated the signal should revise downward, firms that underestimated it should revise upward, and firms with accurate priors should update little if at all. That is exactly the pattern shown in Panel A of Figure 5. The difference in slopes between treatment (0.266) and control (-0.030) is statistically significant ($p<0.001$).

The corresponding results for foreign beliefs in Panel B show the same qualitative pattern.[14] These binned scatterplots also support the linear functional-form assumptions implicit in the Bayesian learning model.

Table 2 reports the regressions from equations (35) and (36), along with the implied estimates of the β's and α's. Column (1) of Table 2 yields $\beta^{d,d} = 0.75$ and $\beta^{d,f} = 0.42$, both statistically significant. This implies that, when forming domestic posterior beliefs, firms extrapolate from both domestic prior beliefs and foreign prior beliefs. However, the domestic extrapolation effect is about twice as large as the foreign one. Combined with the misperception results above, one interpretation is that firms perceive domestic markets as lagging foreign markets but expect some partial catch-up over the next twelve months.

Column (2) of Table 2 yields $\beta^{f,f} = 1.07$, which is statistically significant, and $\beta^{f,d} = 0.09$, which is statistically significant but very close to zero. Thus, when forming foreign posterior beliefs, firms extrapolate from foreign prior beliefs but not from domestic prior beliefs. In other words, they do not expect AI investment in the leading countries to be affected by AI investment in their own country. That is intuitively sensible, since most countries in the sample are small and would not be expected to influence aggregate AI investment in the European leaders.

Column (1) also yields $\alpha^d = 0.35$ and $\alpha^f = 0.36$. Both are statistically significant (p-values $<$0.001 and 0.007).[15] For example, $\alpha^d = 0.35$ means that when forming domestic posterior beliefs, the average firm places a weight of 0.35 on the information about current domestic competitor investment and the remaining 0.65 on its domestic prior belief. The magnitudes of α^d and α^f suggest that respondents understood the signals and regarded them as credible and relevant.

We can compare the magnitudes of these α's with estimates from other information-provision experiments. However, there is no reason to expect α to be constant across studies. In a Bayesian learning framework, its value depends on both the strength of respondents' prior beliefs and the precision of the information they receive. For example, in studies of inflation expectations, the information typically comes from the Consumer Price Index, which is constructed from tens of thousands of price observations. In our setting, by contrast, the signal is derived from a much smaller sample of respondents in an earlier wave of our survey. Respondents may therefore perceive our signal as less precise than those used in inflation-expectations studies and, accordingly, place less weight on it. With that caveat in mind,

[14]For brevity, Panels A and B do not show the full set of relationships. We can also examine how individuals update domestic posterior beliefs in response to information about foreign competitor investment, and vice versa. Those additional results are reported in Table 2.

[15]Moreover, the two estimates ($\alpha^d = 0.35$ and $\alpha^f = 0.36$) are statistically indistinguishable from each other (p-value = 0.927).

our estimates can still be usefully benchmarked against the existing literature. Cavallo et al. (2017) report that Argentine households place a weight of 0.432 on inflation signals when forming expectations. Similarly, Fuster et al. (2022) find that subjects assign a weight of 0.380 to home-price signals. In the context of AI adoption, Cullen et al. (2025) estimate that the average firm places a weight of 0.26 on the information it receives. Our estimates of α^d (0.35) and α^f (0.36) thus fall squarely within the range of comparable estimates in the literature (0.26, 0.38, and 0.432).

Recall that the Bayesian learning model yields two independent estimates of α^d, which allows us to compare them across the two regressions. Although the point estimates differ somewhat—0.35 in column (1) versus 0.40 in column (2)—we cannot reject the null that they are equal (p=0.717). The same is true for α^f: the two independent estimates, 0.36 in column (1) and 0.40 in column (2), are not statistically distinguishable (p=0.912). Taken together, these comparisons support the internal consistency of the learning model.

7 Effect of the Information Treatment on Firms' Own Future AI Investment Plans

7.1 Intention to Treat Estimates

For this analysis, the main outcome of interest is the firm's own expected AI investment rate, defined as the share of total investment that the firm expects to devote to AI during the next twelve months. Panel D of Figure 1 shows the distribution of this outcome in the control group. Most of the variation is on the intensive margin rather than the extensive margin: a large majority of firms (78%) expect to devote a positive share of their investment to AI, but conditional on positive investment there is considerable heterogeneity in the amount. Most firms expect to allocate between 0% and 10% of investment to AI, but the distribution has a long right tail, with 5.3% of firms expecting to allocate 35% or more.

Section 6 showed that the information treatment led firms to revise upward both their domestic posterior beliefs and their foreign posterior beliefs. The conceptual framework in Section 2 predicts that a higher expected level of competitor investment should, in turn, raise the firm's own expected AI investment rate. This hypothesis is tested in Panel A of Figure 4. The third pair of columns reports the average own expected AI investment rate over the next twelve months: the gray bar shows the average in the control group, while the red bar shows the corresponding mean in the treatment group. The treatment increased the average own expected AI investment rate from 8.33% in the control group to 10.13% in the treatment group, a difference of 1.8 pp that is statistically significant ($p<0.001$). Thus, the treatment not only shifted beliefs but also increased firms' own expected AI investment rate. While this estimate summarizes the average treatment effect, Figure B.4 provides a more detailed view by comparing the full distribution of expected AI investment rates in the treatment and control groups.

Panel B of Figure 4 provides the corresponding falsification test. Its third pair of columns uses a pre-treatment outcome: whether the firm reported any current AI use before the information-provision stage. Because this outcome was elicited before the information-provision stage, the treatment should have no effect on it. That is exactly what we observe. The mean is 69.36% in the control group and 70.02% in the treatment group, implying a small and statistically insignificant difference of 0.7 pp (p=0.676). This supports the interpretation that the treatment effects in Panel A reflect responses to the information rather than pre-existing differences between groups.

7.2 2SLS Model

The estimates reported above are intention-to-treat effects of being offered information. These need not coincide with the causal effect of firms' expectations. To recover the latter, we adopt a standard 2SLS approach used in other information-provision experiments (e.g., Cavallo et al., 2017; Cullen and Perez-Truglia, 2022; Giaccobasso et al., 2025). Our primary outcome, $r_{i,t+1}$, is the firm's own expected AI investment rate. For example, it takes the value 0 if the firm expects to allocate 0% of its investment to AI and 10 if it expects to allocate 10%. We seek to estimate how domestic posterior beliefs ($s^d_{i,t+1}$) and foreign posterior beliefs ($s^f_{i,t+1}$) affect the firm's own expected AI investment rate. A simple regression of $r_{i,t+1}$ on $s^d_{i,t+1}$ and $s^f_{i,t+1}$ would likely suffer from omitted-variable bias. We therefore identify the causal effect by instrumenting posterior beliefs with the randomized information shock, that is, the component of belief updating induced solely by treatment assignment:

$$\begin{aligned} r_{i,t+1} = \theta_0 + \theta^d_1 \cdot s^d_{i,t+1} + \theta^f_1 \cdot s^f_{i,t+1} + \\ + \theta^d_2 \cdot (s^{d,T}_{i,t} - s^{d,0}_{i,t}) + \theta^d_3 \cdot s^{d,0}_{i,t} + \theta^f_2 \cdot (s^{f,T}_{i,t} - s^{f,0}_{i,t}) + \theta^f_3 \cdot s^{f,0}_{i,t} + X_i \theta_X + \epsilon_i, \end{aligned} \tag{37}$$

where $s^d_{i,t+1}$ and $s^f_{i,t+1}$ are the endogenous variables and the excluded instruments are $T_i \cdot (s^{d,T}_{i,t} - s^{d,0}_{i,t})$ and $T_i \cdot (s^{f,T}_{i,t} - s^{f,0}_{i,t})$.[16] The vector X_i includes additional covariates, such as

[16] Formally, the exogeneity assumptions are $\mathbb{E}\left[\left(s^{d,T}_{i,t} - s^{d,0}_{i,t}\right) \cdot T_i \cdot \epsilon_i \mid \mathbb{X}_i\right] = 0$ and $\mathbb{E}\left[\left(s^{f,T}_{i,t} - s^{f,0}_{i,t}\right) \cdot T_i \cdot \epsilon_i \mid \mathbb{X}_i\right] = 0$, where $\mathbb{X}_i$ is the vector that includes $\{s^{d,T}_{i,t} - s^{d,0}_{i,t}, s^{d,0}_{i,t}, s^{f,T}_{i,t} - s^{f,0}_{i,t}, s^{f,0}_{i,t}, X_i\}$. 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.

employment and sales.[17]

The logic of the 2SLS strategy can be conveyed with a simple thought experiment. Consider two firms that have the same prior error in their domestic prior belief. Suppose both initially underestimate domestic competitor AI investment by 10 pp, but randomization assigns one firm to treatment and the other to control. Assume the control firm does not revise its belief and therefore continues to underpredict competitors' future investment by 10 pp, while the treated firm finds the signal credible and revises so that it underpredicts by only 2 pp. The treatment then generates an 8 pp upward shift in the treated firm's domestic posterior belief. If, relative to the control firm, the treated firm's own expected AI investment rate rises by 6 pp, the implied effect is 0.75 pp higher own expected AI investment for each 1 pp increase in perceived competitor investment ($= \frac{6}{8}$). While this back-of-the-envelope ratio cannot be estimated for a single exact prior gap, the same logic can be applied across all levels of initial under- and overestimation and then aggregated. The 2SLS regression implements exactly this comparison.[18]

7.3 2SLS Results

Before presenting the estimates from this 2SLS model, we again provide semi-parametric evidence using binned scatterplots. These results are shown in Panels C and D of Figure 5. Intuitively, if the effect of information on posterior beliefs depends on whether the corresponding prior belief underestimates, overestimates, or matches the signal, then the effect of information on the firm's own expected AI investment rate should display a similar pattern of heterogeneity.

For example, Panel A of Figure 5, discussed in Section 6, shows that the effect of information on domestic posterior beliefs is negative for firms that overestimated domestic competitor AI investment, positive for firms that underestimated it, and close to zero for firms with accurate domestic prior beliefs. Panel C of Figure 5 is identical except that the y-axis reports the firm's own expected AI investment rate rather than its domestic posterior belief. Panel C displays the same pattern of heterogeneity: the effect of information on the firm's own expected AI investment rate is negative for firms that overestimated domestic competitor AI investment, positive for firms that underestimated it, and approximately zero

[17] The full set of controls includes firm-level variables from the survey: employment, annual sales, ownership type, age, financial autonomy, and export share. Most of these are categorical variables; for more details, see Table 1.

[18] For further discussion of the 2SLS setup, see Cullen and Perez-Truglia (2022). Note that the 2SLS coefficient identifies a Local Average Treatment Effect (LATE): the effect of beliefs for firms whose posteriors move in response to the information. Mechanically, this places more weight on firms with larger initial misperceptions and, conditional on those misperceptions, on those who update more strongly when treated.

for firms with accurate domestic prior beliefs.

Panel D presents the analogous exercise for foreign beliefs. It is identical to Panel B except that the y-axis reports the firm's own expected AI investment rate rather than its foreign posterior belief. The pattern in Panel D suggests that the treatment effect on the firm's own expected AI investment rate is much less related to whether the firm under- or overestimated foreign competitor AI investment.

Taken together, Panels C and D suggest that the information affected firms' own expected AI investment rate primarily through domestic posterior beliefs rather than through foreign posterior beliefs. We next use the 2SLS estimates to disentangle these two channels more directly. Table 3 reports the results. In addition to the 2SLS estimates in Panel A, the table presents the first-stage results in Panels B and C and the reduced-form results in Panel D. Each column corresponds to a different 2SLS specification, varying either the controls or the dependent variable.

The baseline results are presented in column (1) of Table 3. In column (1), the dependent variable is the firm's own expected AI investment rate (from 0–100%). Consistent with the binned-scatterplot evidence in Figure 5, the 2SLS estimates indicate that the firm's own expected AI investment rate rises significantly with the domestic posterior belief but does not respond to the foreign posterior belief. More precisely, the coefficient on the domestic posterior belief is 0.570 and statistically significant (p-value = 0.001), while the coefficient on the foreign posterior belief is -0.224 and statistically insignificant (p-value = 0.157).

The 2SLS estimates also admit a more intuitive interpretation in terms of magnitudes. Equation (37) can be viewed as the empirical counterpart to the reaction function in the conceptual model (equation (22)). Under that interpretation, the domestic 2SLS coefficient measures how the firm's own expected AI investment rate responds to an increase in its domestic posterior belief. The estimate of 0.570 in column (1) of Table 3 implies that a 1 pp increase in the domestic posterior belief raises the firm's own expected AI investment rate by 0.570 pp. By contrast, the coefficient on the foreign posterior belief is close to zero and statistically insignificant, suggesting that firms do not, on average, respond to increases in the expected share of foreign competitors investing in AI.

Before interpreting the magnitude of the results further, we first present a set of robustness checks. A standard concern in 2SLS estimation is weak instruments (Stock et al., 2002). Panels B and C of Table 3 show that, consistent with Section 6, the treatment had a strong effect on both domestic posterior beliefs and foreign posterior beliefs. To assess instrument strength more formally, the bottom section of Table 3 reports the Kleibergen-Paap *rk* Wald

F-statistic for each specification, a standard measure of instrument relevance.[19] According to the rule of thumb in Staiger and Stock (1997), F-statistics above 10 indicate that weak identification is unlikely to be a serious concern. The reported values exceed that threshold, suggesting that weak instruments are not a problem in our setting.

Columns (2) through (7) of Table 3 present robustness checks and falsification tests. Column (2) is identical to the baseline specification in column (1) except that it adds indicators for prior AI use. Column (3) builds on column (2) by further including fixed effects for country-industry-size cells. Column (4) augments column (3) with additional firm-level controls, including indicators for financial autonomy, sales, age, and ownership. Because identification relies on random assignment, one would not expect the point estimates to change materially when controls are added. That is exactly what we find: the 2SLS estimates in columns (2) through (4) are nearly identical to those in the baseline specification.[20]

Column (5) of Table 3 repeats the baseline specification after dropping observations in the bottom and top 0.5% of the domestic and foreign prior-gap distributions. This trimmed specification addresses the possibility that the baseline estimate is driven by firms with extreme initial misperceptions. The estimates remain very similar to the baseline results.

Column (6) of Table 3 presents results for the extensive margin. It repeats the baseline specification, except that the dependent variable is an indicator equal to 0 if the firm's own expected AI investment rate is zero and 100 if it is positive. As discussed in Section 7.1, only a minority of the variation in the firm's own expected AI investment rate comes from the extensive margin, so we expect lower power for this outcome. That is what the estimates suggest: the 2SLS coefficient on the domestic posterior belief in column (6) is positive but imprecisely estimated and therefore statistically insignificant.

Finally, column (7) of Table 3 presents a falsification test. As in the intention-to-treat analysis in Section 7.1, we use as a placebo outcome an indicator equal to 100 if the firm reported any current AI use before the information-provision stage and 0 otherwise. Because this outcome was elicited before the information-provision stage, the treatment should have no effect on it. That is exactly what we find: the 2SLS estimates for both the domestic posterior belief and the foreign posterior belief are statistically insignificant.

We now return to the interpretation of the baseline results in column (1) of Table 3. According to the model in Section 2, each 2SLS coefficient can be interpreted as reflecting a combination of two channels: competition, as firms seek to avoid being left behind by

[19]While the conventional benchmark relies on the Cragg-Donald statistic under homoskedasticity, Baum et al. (2007) recommend the Kleibergen-Paap statistic as a robust alternative.

[20]Adding controls can still improve precision if they reduce the variance of the error term. Consistent with this, the standard errors are somewhat smaller in columns (2)–(4) than in column (1), although the gains in precision are modest.

rivals, and learning, as firms draw inferences about their own returns to AI. According to the competition channel, the sharp difference between the effects of domestic and foreign beliefs suggests that firms view domestic peers as much closer competitors than foreign peers. According to the learning channel, this difference suggests that firms learn about their own potential returns from domestic peers, but find information about foreign peers less useful.

Because the competition and learning channels are closely intertwined, it is difficult to disentangle them precisely. Nonetheless, following Cullen et al. (2025), we provide suggestive evidence based on heterogeneity analysis. Intuitively, the strength of the competition channel should vary with the degree of competition in the market. In the extreme case of a perfect monopolist, a firm could still learn from other firms' AI investment decisions, but the competition channel should be shut down.

These results are presented in Table 4.[21] Column (1) reproduces the baseline estimates from column (1) of Table 3. The remaining columns apply the same specification to different subsamples. The first split uses firm size as a proxy for competitive exposure, comparing firms with below-median turnover (column (2)) to those with above-median turnover (column (3)). The logic is that smaller firms are likely to face stronger competitive pressure than larger firms. However, the Kleibergen-Paap F-statistic is low in column (2), so the low-turnover estimate may be affected by weak instruments and should be interpreted with caution. With that caveat, the pattern is consistent with the competition channel: the effect of beliefs about domestic peers is stronger for smaller firms than for larger firms, with coefficients of 0.835 and 0.430, respectively.

Columns (4)–(5) of Table 4 then split firms by market concentration, using the Herfindahl-Hirschman Index. More precisely, we use the latest available data (2023) from Orbis on firm revenues to compute the HHI in each country-industry-size class, and then split firms into those in below-median-concentration markets (column (4)) and above-median-concentration markets (column (5)). Column (4) also has a low Kleibergen-Paap F-statistic, so this comparison should likewise be taken with a grain of salt. The point estimates nevertheless provide suggestive evidence consistent with the competition channel: the effect of beliefs about domestic peers is larger for firms in less concentrated markets than for firms in more concentrated markets, with coefficients of 0.803 and 0.418, respectively.

Columns (6)–(7) of Table 4 further split the sample by sector tradability, classifying two-digit NACE industries as non-tradable (column (6)) or tradable (column (7)). The effect of beliefs about domestic peers is stronger in non-tradable sectors than in tradable sectors, with coefficients of 0.680 and 0.465, respectively. This pattern is also consistent with the competition channel: firms in non-tradable sectors are likely to face more localized

[21] For additional heterogeneity analyses, see Table B.1.

competition and therefore respond more strongly to information about domestic competitors. However, because the F-statistic is low in column (7), the estimate for tradable sectors should also be interpreted cautiously.

Finally, columns (8)–(9) of Table 4 split firms by export exposure, comparing firms with below-median export shares (column (8)) to those with above-median export shares (column (9)). Once again, the domestic-belief coefficient is larger among firms that are less exposed to foreign markets: 0.710 for low-export firms, compared with 0.490 for high-export firms. Because column (8) has a low F-statistic, this comparison is also subject to the weak-instruments caveat. Taken as suggestive evidence, the pattern reinforces the same interpretation: firms that rely less on exports appear more responsive to what their domestic competitors are doing, consistent with stronger strategic complementarities within domestic markets.

8 Conclusions

This study set out to provide causal evidence on how the geography of competition shapes firms' AI investment decisions. Using a large-scale survey experiment embedded in the ECB's SAFE survey, we examined how firms update their beliefs about domestic and foreign competitor AI investment and whether these updated beliefs affect their own expected AI investment rate. We documented large underestimation of competitor AI investment, substantial belief updating in response to information, and a clear asymmetry in how firms react to domestic versus foreign competition. Firms update both domestic and foreign beliefs when informed, but their own expected AI investment rate responds primarily to domestic posterior beliefs. These findings suggest that strategic complementarities in innovation weaken with distance, broadly understood to include not only geography but also informational, cultural, and market frictions.

Our estimates suggest that firms' investment decisions are highly interdependent within domestic markets. The slope of the estimated reaction function indicates that when firms expect a larger share of domestic competitors to invest in AI, they substantially increase their own expected AI investment as well. In contrast, we find no comparable effects for foreign posterior beliefs. This asymmetry helps explain why AI diffusion may remain geographically uneven, even within an integrated economic area like Europe. While firms may observe and learn from foreign competitors, their behavioral response to such foreign signals is much weaker compared to domestic competitors.

We conclude by discussing the policy implications. Europe has devoted substantial resources to fostering innovation and AI diffusion, through both centralized EU initiatives and national programs. Yet our evidence suggests that financial incentives may not be sufficient on their own. If firms systematically underestimate the extent of AI investment among their competitors, especially abroad, then informational frictions may be an important source of under-investment. As a complement to financial-incentive policies, governments and public institutions could deploy information campaigns that improve firms' awareness of peer investment and the diffusion of new technologies. While more work is needed to determine the most effective design of such interventions, their relatively low cost suggests that they could be a cost-effective tool for accelerating diffusion.

More broadly, our findings speak to a central ambition of the European project. One of the key goals of the Single Market was to expose firms to stronger cross-border competition and thereby stimulate innovation, productivity, and convergence. Our results suggest that this force is real but incomplete: firms respond strongly to domestic competitive pressure, yet much less to foreign competitive pressure, even within Europe. To the extent that policymakers seek to reduce regional disparities in AI investment, use, and innovation, lowering these invisible barriers to cross-border spillovers may be as important as expanding financial support.

References

Akcigit, U., S. T. Ates, J. Lerner, R. R. Townsend, and Y. Zhestkova (2024). Fencing Off Silicon Valley: Cross-Border Venture Capital and Technology Spillovers. *Journal of Monetary Economics 141*, 14–39.

Alfaro-Urena, A., I. Manelici, and J. P. Vasquez (2022). The Effects of Joining Multinational Supply Chains: New Evidence from Firm-to-Firm Linkages. *The Quarterly Journal of Economics 137*(3), 1495–1552.

Amiti, M. and J. Konings (2007). Trade Liberalization, Intermediate Inputs and Productivity: Evidence from Indonesia. *American Economic Review 97*(5), 1611–1638.

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. K. Khandelwal, and A. Osman (2017). Exporting and Firm Performance: Evidence from a Randomized Experiment. *The Quarterly Journal of Economics 132*(2), 551–615.

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

Baslandze, S., Z. Edwards, J. Graham, T. McClure, B. H. Meyer, M. Sparks, S. R. Waddell, and D. Weitz (2026). Artificial Intelligence, Productivity, and the Workforce: Evidence from Corporate Executives. NBER Working Paper No. 34984.

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 (2024). Technology Spillovers, Diffusion and Rivalry in Firm Networks. CEPR Discussion Paper No. 19804.

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

Bonney, K., C. L. Breaux, E. Dinlersoz, L. S. Foster, J. C. Haltiwanger, and A. A. Pande (2026). The Microstructure of AI Diffusion: Evidence from Firms, Business Functions, and Worker Tasks. NBER Working Paper No. 35141.

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.

Coe, D. T. and E. Helpman (1995). International R&D Spillovers. *European Economic Review 39*(5), 859–887.

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

Cruces, G., R. Perez-Truglia, and M. Tetaz (2013). Biased Perceptions of Income Distribution and Preferences for Redistribution: Evidence from a Survey Experiment. *Journal of Public Economics 98*, 100–112.

Cullen, Z., E. Faia, E. Guglielminetti, R. Perez-Truglia, and C. Rondinelli (2025). The Innovation Race: Experimental Evidence on Advanced Technologies. *NBER Working Paper No. 34532*.

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.

Draghi, M. (2024). The Future of European Competitiveness: In-Depth Analysis and Recommendations. Report, European Commission.

Eaton, J. and S. Kortum (2002). Technology, Geography, and Trade. *Econometrica 70*(5), 1741–1779.

European Central Bank (2025). Survey on the Access to Finance of Enterprises, 2025 Q4. https://www.ecb.europa.eu/stats/accesstofinancesofenterprises/pdf/questionnaire/ecb.safeq202602.en.pdf.

European Commission (2021). The Recovery and Resilience Facility. https://reforms-investments.ec.europa.eu/recovery-and-resilience-facility-1_en.

Fehr, D., J. Mollerstrom, and R. Perez-Truglia (2022). Your Place in the World: Relative Income and Global Inequality. *American Economic Journal: Economic Policy 14*(4), 232–68.

Ferrando, A., S. Lamboglia, J. Rariga, and M. Schmidt (2026). Adopting and Investing in AI: Evidence from Euro Area Firms in the SAFE. *Economic Bulletin Boxes 2*.

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). escftest: Stata Module to Perform Epps–Singleton Two-Sample Test. *The Stata Journal 9*(3), 454–465.

Griliches, Z. (1960). Hybrid Corn and the Economics of Innovation: Geographic Differences in the Use of Hybrid Corn Are Explained by Differences in the Profitability of That Use. *Science 132*(3422), 275–280.

Griliches, Z. (1987). R&D and Productivity: Measurement Issues and Econometric Results. *Science 237*(4810), 31–35.

Huo, Z. and N. Takayama (2025). Rational Expectations Models with Higher-Order Beliefs. *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.

Keller, W. (2004). International Technology Diffusion. *Journal of Economic Literature 42*(3), 752–782.

Keller, W. and S. R. Yeaple (2013). The Gravity of Knowledge. *American Economic Review 103*(4), 1414–1444.

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

Letta, E. (2024). Much More Than a Market. Technical report, Council of the European Union and European Commission, https://www.consilium.europa.eu/media/ny3j24sm/much-more-than-a-market-report-by-enrico-letta.pdf.

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.

Lind, N. and N. Ramondo (2023). Global Innovation and Knowledge Diffusion. *American Economic Review: Insights 5*(4), 494–510.

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

Pavcnik, N. (2002). Trade Liberalization, Exit, and Productivity Improvements: Evidence from Chilean Plants. *Review of Economic Studies 69*(1), 245–276.

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

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.

Yotzov, I., J. M. Barrero, N. Bloom, P. Bunn, S. J. Davis, K. M. Foster, A. Jalca, B. H. Meyer, P. Mizen, M. A. Navarrete, et al. (2026). Firm Data on AI. Technical report, National Bureau of Economic Research.

Zolas, N., Z. Kroff, E. Brynjolfsson, K. McElheran, D. N. Beede, C. Buffington, N. Goldschlag, L. Foster, and E. Dinlersoz (2021). Advanced Technologies Adoption and Use by U.S. Firms: Evidence from the Annual Business Survey. NBER Working Paper No. 28290.

Figure 1: Descriptive Statistics about AI Investment and Use

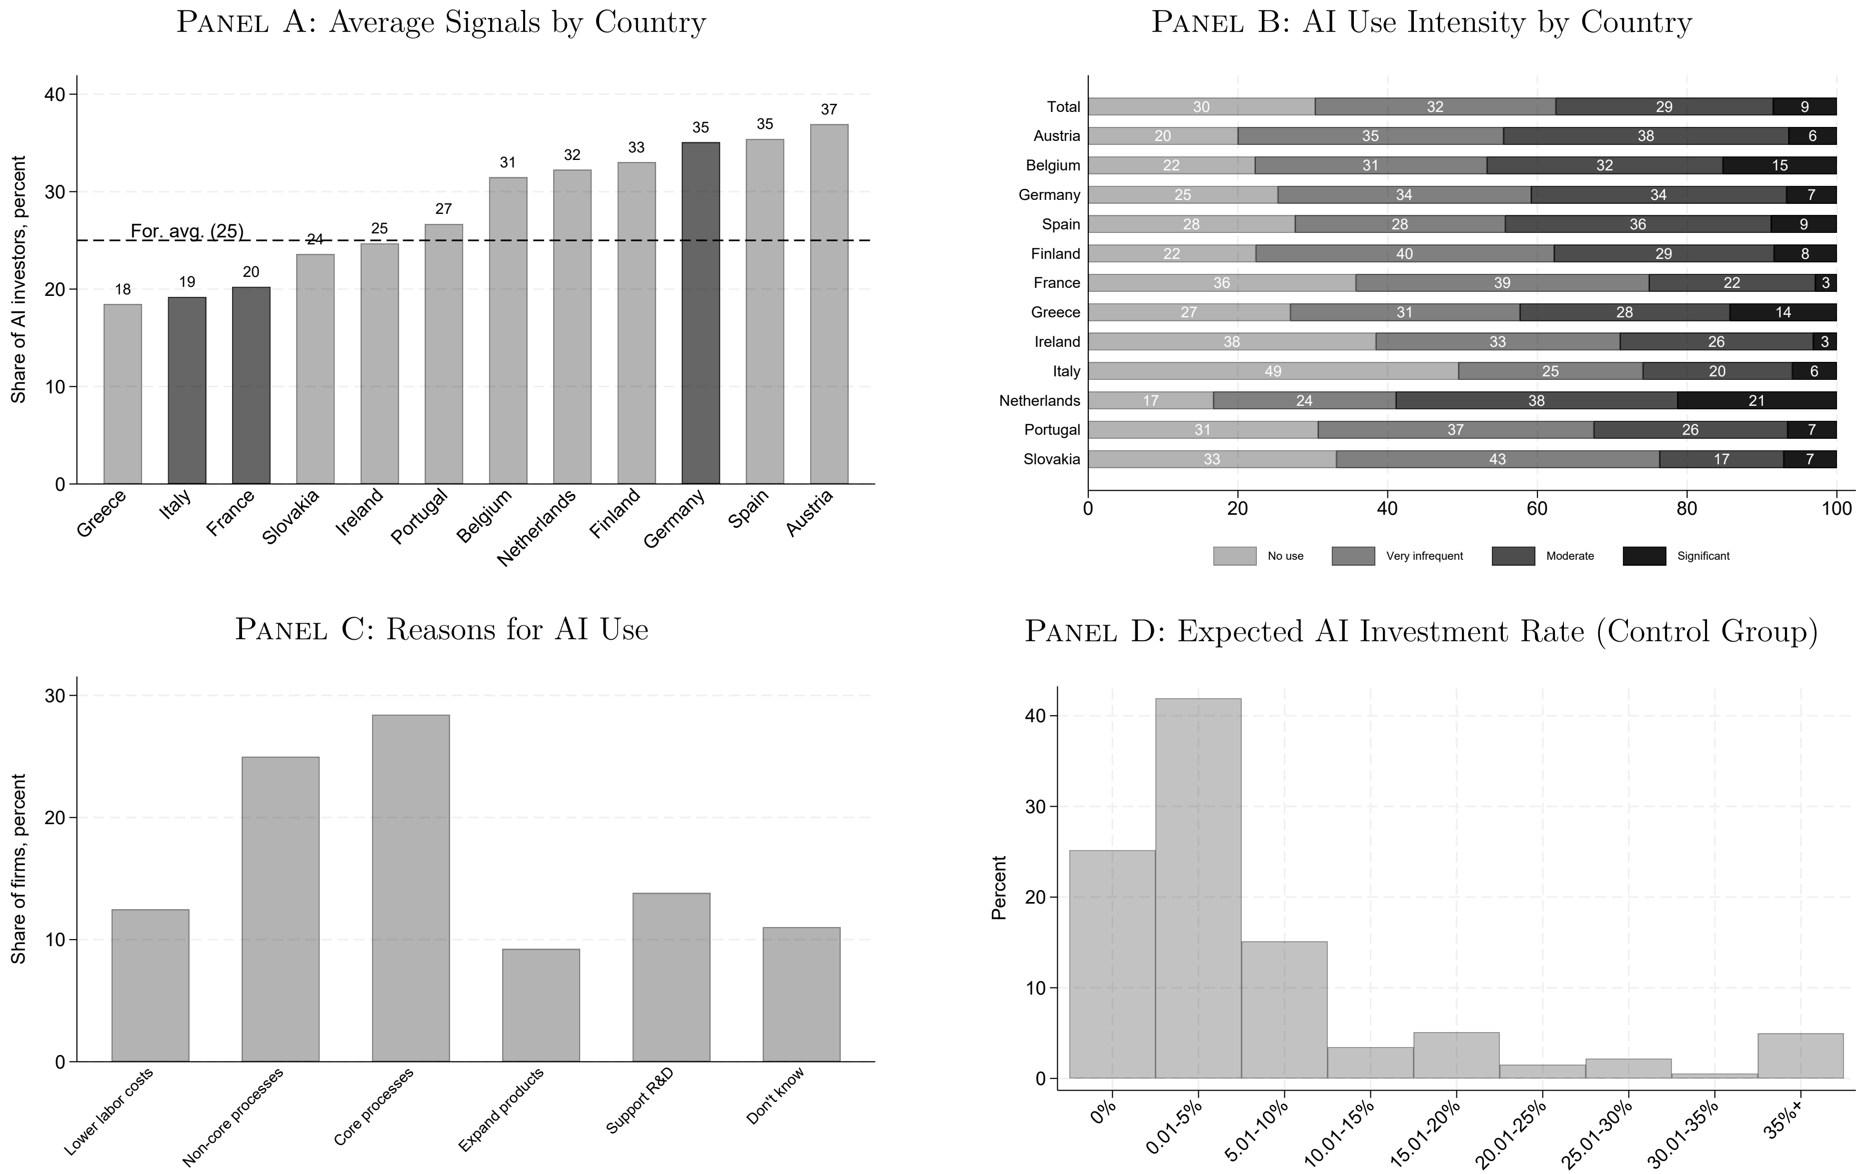

Notes: Panel A uses the Q2–2025 survey and reports the average share of firms investing in AI up to Q2–2025 by country. Foreign average denotes the average share of firms investing across Germany, France, and Italy. Panels B–D use the Q4–2025 survey and report the share of firms by current AI use intensity at the country level, firms' reported reasons for AI use, and the expected share of total investment allocated to AI over the next twelve months. Panel D shows the binned distribution of the AI investment rate for the control group.

Figure 2: Distribution of Misperceptions about Domestic and Foreign Competitor AI Investment

Panel A: Distribution of Misperceptions

Panel B: Correlation of Misperceptions

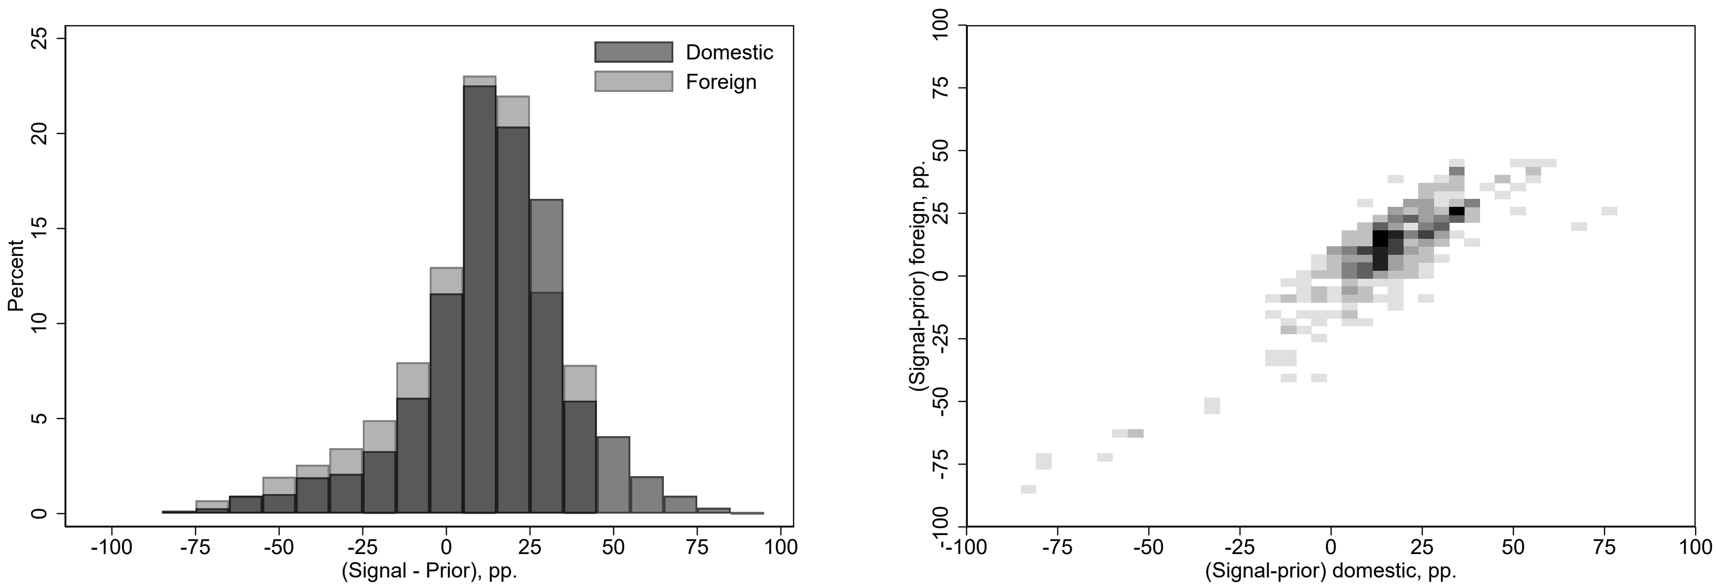

Notes: Panel A shows the distribution of the gap between the actual share of firms investing in AI and firms' beliefs about AI investment shares by domestic and foreign competitors. Panel B plots the relationship between domestic and foreign misperceptions. The figure is based on the Q2–2025 and Q4–2025 waves of the survey.

Figure 3: Actual and Perceived Domestic and Foreign AI Investment by Country

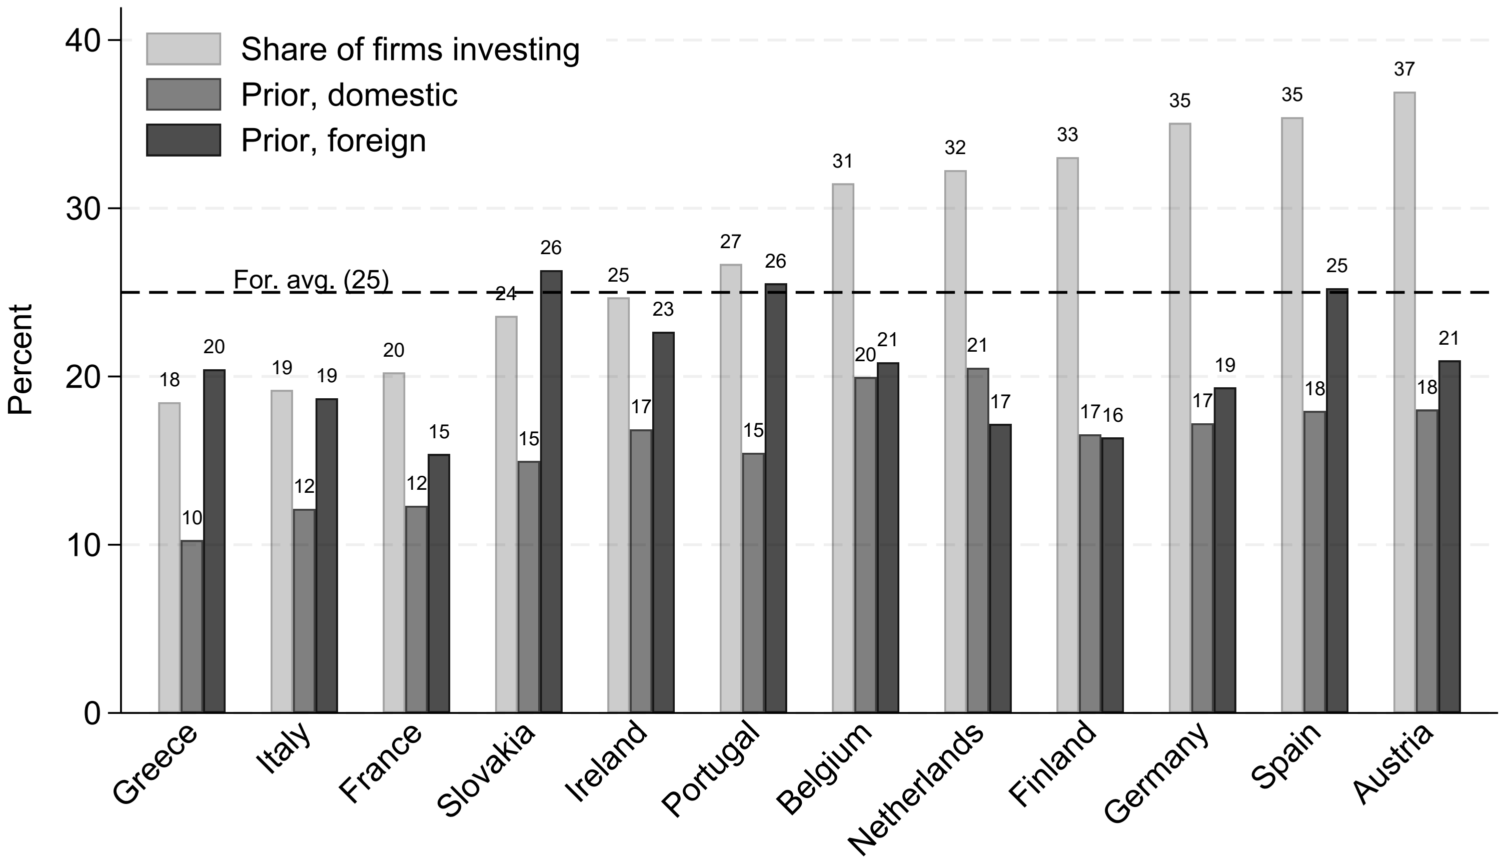

Notes: This figure shows the actual share of firms investing in AI and firms' beliefs about AI investment shares by domestic and foreign competitors at the country level. Foreign average refers to the actual share of firms investing in Germany, France, and Italy. The figure is based on the Q2–2025 and Q4–2025 waves of the survey.

Figure 4: Treatment Effects on Posterior Beliefs and AI Investment Rate, with Falsification Tests

PANEL A: Treatment Effects

PANEL B: Falsification tests

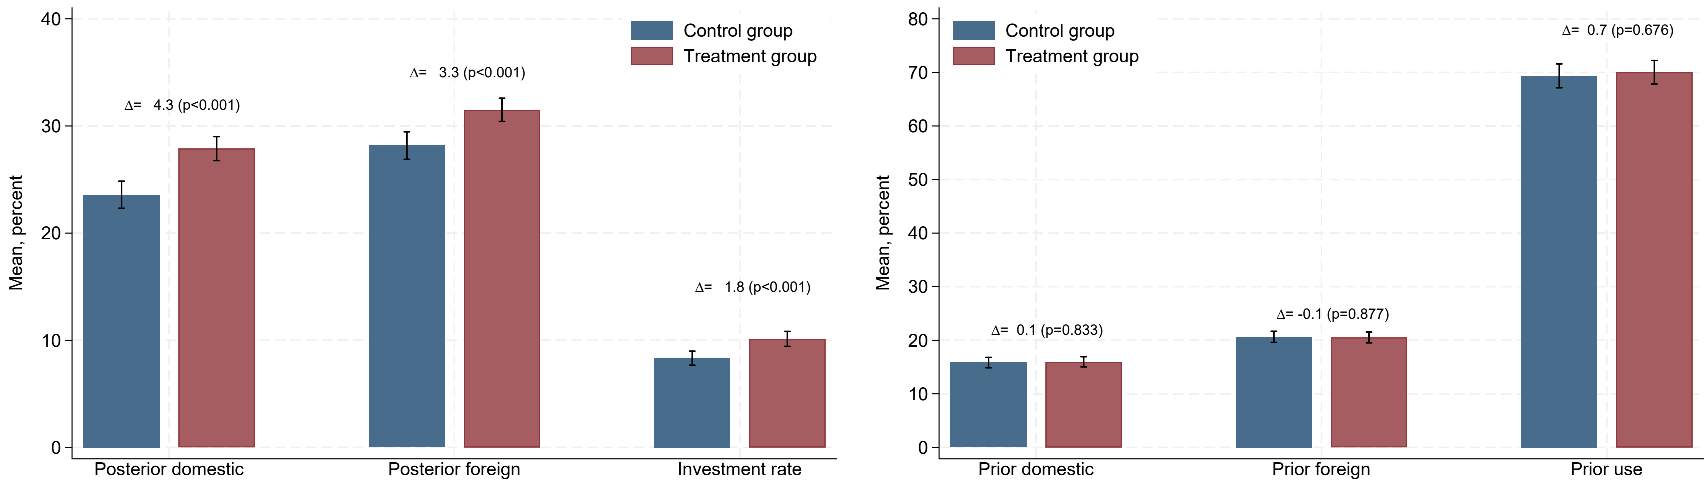

Notes: Panel A presents treatment effects on posterior beliefs about domestic and foreign AI investment and on the AI investment rate, comparing control and treatment groups. Panel B reports falsification tests comparing control and treatment groups for prior domestic beliefs, prior foreign beliefs, and prior AI use. Confidence intervals are at the 95% level. The figure is based on the Q4–2025 wave of the survey.

Figure 5: Belief Updating: Perceptions about Domestic and Foreign Competitor AI Investment

Panel A: Effect on Domestic Posterior by Domestic Gap

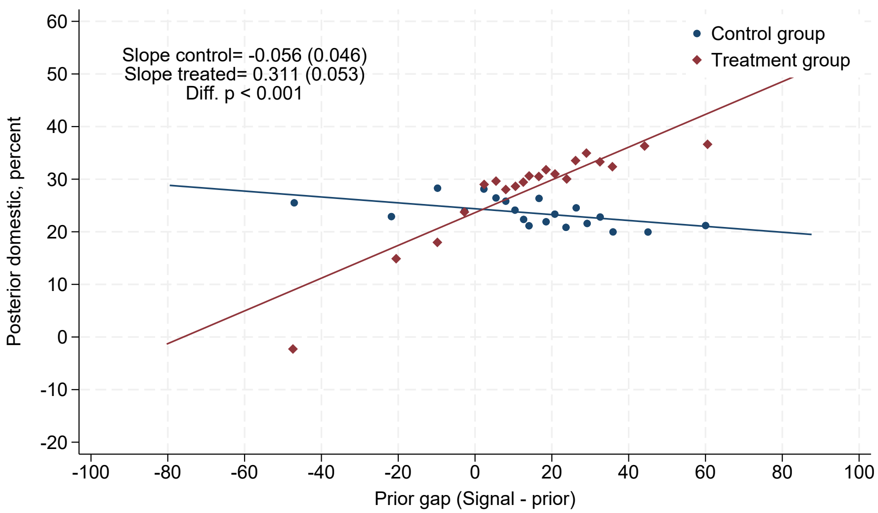

Panel B: Effect on Foreign Posterior by Foreign Gap

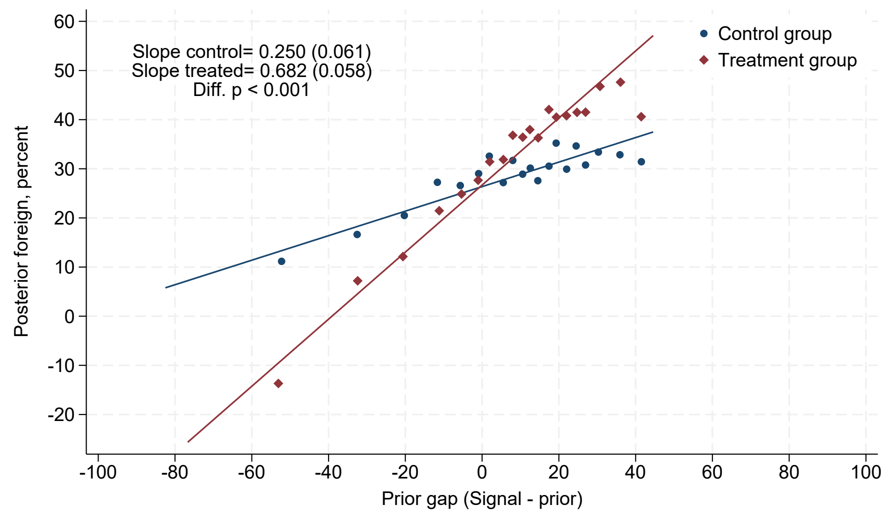

Panel C: Effect on Investment Rate by Domestic Gap

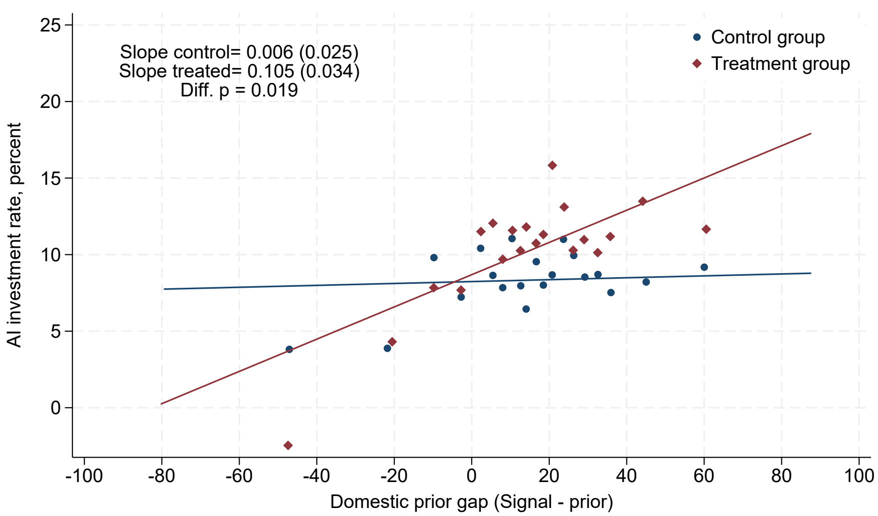

Panel D: Effect on Investment Rate by Foreign Gap

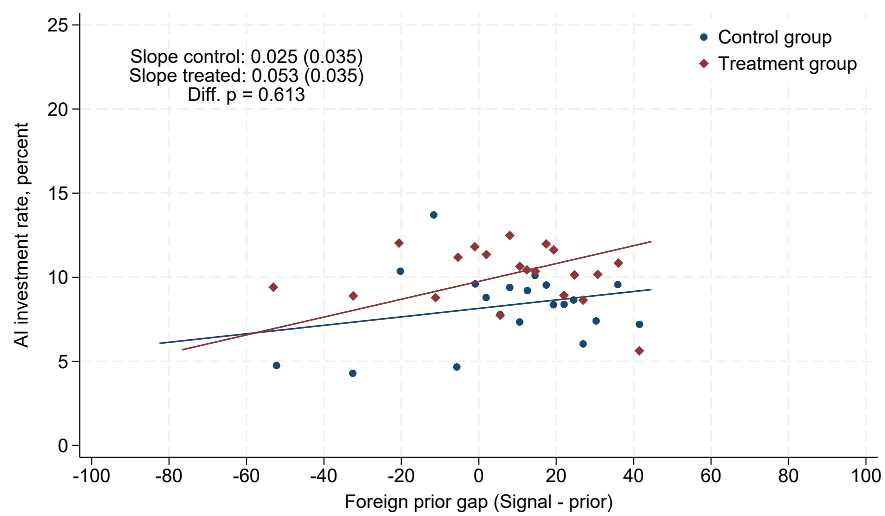

Notes: Panel A displays a binscatter of the domestic posterior beliefs against the domestic prior gap, and Panel B plots foreign posterior beliefs against the foreign prior gap, separately for control and treatment groups. Panel C shows a binscatter of AI investment rates by the domestic prior gap, and Panel D shows AI investment rates by the foreign prior gap, with fitted lines for each group.

Table 1: Summary Statistics and Randomization Balance

	All (1)	Control (2)	Treatment (3)	p-value (4)
Employment	112.66	126.44	99.05	0.555
	(1329.00)	(1837.66)	(418.91)	
Industry	26.15	25.00	27.28	0.136
	(43.95)	(43.31)	(44.55)	
Construction	13.12	13.29	12.95	0.772
	(33.77)	(33.96)	(33.58)	
Trade	21.29	20.69	21.88	0.402
	(40.94)	(40.52)	(41.36)	
Services	39.45	41.02	37.89	0.065
	(48.88)	(49.20)	(48.53)	
Autonomous	80.91	80.95	80.88	0.958
	(39.31)	(39.28)	(39.34)	
Subsidiary	18.15	18.20	18.11	0.941
	(38.55)	(38.60)	(38.52)	
Branch	0.54	0.61	0.48	0.619
	(7.35)	(7.77)	(6.91)	
Age <2 years	0.66	0.67	0.66	0.977
	(8.12)	(8.15)	(8.10)	
Age >=2, <5 years	1.60	1.82	1.38	0.311
	(12.54)	(13.37)	(11.66)	
Age >=5, <10 years	6.42	5.70	7.13	0.093
	(24.52)	(23.20)	(25.75)	
Age >=10 years	91.25	91.75	90.77	0.318
	(28.25)	(27.52)	(28.96)	
Public shareholders	1.57	1.46	1.68	0.606
	(12.43)	(11.98)	(12.85)	
Family/entrepreneurs	46.08	46.48	45.68	0.645
	(49.85)	(49.89)	(49.83)	
Other firm	14.78	14.93	14.63	0.808
	(35.49)	(35.65)	(35.35)	
Venture capital	0.72	0.79	0.66	0.661
	(8.48)	(8.85)	(8.10)	
One owner, natural p.	33.90	33.50	34.29	0.628
	(47.34)	(47.21)	(47.48)	
Sales up to 0.5 EUR ml	12.52	12.20	12.83	0.582
	(33.09)	(32.73)	(33.45)	
Sales [0.5,1] EUR ml	13.36	13.23	13.49	0.825
	(34.03)	(33.89)	(34.17)	
Sales (1,2] EUR ml	0.54	0.61	0.48	0.619
	(7.35)	(7.77)	(6.91)	
Sales (2,10] EUR ml	19.12	18.26	19.96	0.213
	(39.33)	(38.65)	(39.98)	
Sales (10,50] EUR ml	10.19	10.07	10.31	0.820
	(30.26)	(30.11)	(30.42)	
Sales above 50 EUR ml	0.19	0.20	0.18	0.220
	(0.39)	(0.40)	(0.39)	
Exp. to sales, %	16.69	16.36	17.02	0.510
	(28.71)	(28.47)	(28.95)	
Observations	3,316	1,648	1,668	3,316

Notes: Summary statistics for selected firm characteristics. These include the average number of employees, shares by industry, financial autonomy, age group, ownership type, 2024 annual sales, and the ratio of export to total sales. Data are provided for the full sample, as well as for the control and treatment groups. Standard errors are shown in parentheses. The sample is restricted to firms with non-missing observations across all AI variables.

Table 2: Bayesian Learning Estimates

	(1) $s^{d}_{i,t+1}$	(2) $s^{f}_{i,t+1}$
Panel A: Raw Regression Results		
$T_i \cdot \left(s^{d,T}_{i,t} - s^{d,0}_{i,t}\right)$	0.259***	0.036
	(0.038)	(0.038)
$T_i \cdot \left(s^{f,T}_{i,t} - s^{f,0}_{i,t}\right)$	0.154***	0.434***
	(0.042)	(0.043)
$\left(s^{d,T}_{i,t} - s^{d,0}_{i,t}\right)$	-0.002	-0.056*
	(0.033)	(0.033)
$\left(s^{f,T}_{i,t} - s^{f,0}_{i,t}\right)$	0.166***	0.250***
	(0.045)	(0.045)
$s^{d,0}_{i,t}$	0.749***	0.090**
	(0.038)	(0.039)
$s^{f,0}_{i,t}$	0.423***	1.077***
	(0.047)	(0.047)
Constant	1.636*	3.444***
	(0.850)	(0.854)
Panel B: Implied Model Parameters		
	β^{df}: 0.423***	β^{ff}: 1.077***
	(0.047)	(0.047)
	β^{dd}: 0.749***	β^{fd}: 0.090***
	(0.038)	(0.039)
	α^{d}: 0.346***	α^{d}: 0.402
	(0.053)	(0.447)
	α^{f}: 0.364***	α^{f}: 0.403***
	(0.108)	(0.043)
P-value $\alpha^{f}_{(1)} = \alpha^{f}_{(2)}$	0.912	
P-value $\alpha^{d}_{(1)} = \alpha^{d}_{(2)}$	0.717	
Observations	3,316	3,316

Notes: Column (1) shows results from estimating Equation 35, while Column (2) shows results for Equation 36. The p-value rows report tests of equality for each learning weight across the two columns. The sample is restricted to firms with non-missing observations across all AI variables.

Table 3: 2SLS estimates: Effects of Competitors' Expected AI Investment on Own Future Investment

	Inv. rate (1)	Inv. rate (2)	Inv. rate (3)	Inv. rate (4)	Inv. rate (5)	Inv. dummy (6)	Prior use (7)
			Panel A: 2SLS				
$s^d_{i,t+1}$	0.570***	0.563***	0.542***	0.544***	0.522**	0.544	0.254
	(0.167)	(0.166)	(0.163)	(0.163)	(0.170)	(0.392)	(0.410)
$s^f_{i,t+1}$	-0.224	-0.213	-0.195	-0.202	-0.206	-0.083	-0.477
	(0.158)	(0.156)	(0.152)	(0.152)	(0.159)	(0.355)	(0.379)
	Panel B: First-Stage (Dep. Var: Posterior Domestic)						
$T_i \cdot \left(s^{d,T}_{i,t} - s^{d,0}_{i,t}\right)$	0.259***	0.257***	0.278***	0.280***	0.268***	0.268***	0.259***
	(0.049)	(0.048)	(0.048)	(0.048)	(0.049)	(0.049)	(0.049)
$T_i \cdot \left(s^{f,T}_{i,t} - s^{f,0}_{i,t}\right)$	0.154**	0.161**	0.140*	0.140*	0.143*	0.143*	0.154**
	(0.057)	(0.056)	(0.056)	(0.056)	(0.056)	(0.056)	(0.057)
	Panel C: First-Stage (Dep. Var: Posterior Foreign)						
$T_i \cdot \left(s^{d,T}_{i,t} - s^{d,0}_{i,t}\right)$	0.036	0.034	0.061	0.063	0.059	0.059	0.036
	(0.047)	(0.046)	(0.047)	(0.047)	(0.048)	(0.048)	(0.047)
$T_i \cdot \left(s^{f,T}_{i,t} - s^{f,0}_{i,t}\right)$	0.434***	0.440***	0.416***	0.416***	0.417***	0.417***	0.434***
	(0.056)	(0.055)	(0.057)	(0.057)	(0.057)	(0.057)	(0.056)
			Panel D: Reduced-Form				
$T_i \cdot \left(s^{d,T}_{i,t} - s^{d,0}_{i,t}\right)$	0.140***	0.137***	0.139***	0.140***	0.128***	0.141	0.049
	(0.031)	(0.030)	(0.031)	(0.031)	(0.032)	(0.086)	(0.089)
$T_i \cdot \left(s^{f,T}_{i,t} - s^{f,0}_{i,t}\right)$	-0.009	-0.003	-0.005	-0.008	-0.011	0.043	-0.168
	(0.037)	(0.036)	(0.037)	(0.037)	(0.039)	(0.097)	(0.101)
Control for Prior AI Use Dummies	No	Yes	Yes	Yes	No	No	No
Cou.-Ind.-Size FE	No	No	Yes	Yes	No	No	No
Control for Other Firm Characteristics	No	No	No	Yes	No	No	No
Observations	3,316	3,316	3,316	3,316	3,267	3,267	3,316
Dep. Var.: Baseline Mean (%)	8.33	8.33	8.33	8.33	8.23	74.82	69.36
Kleibergen-Paap Wald F-statistic	12.03	11.80	12.27	12.14	10.57	10.57	12.03

Notes: Panel A reports 2SLS estimates of equation 37. The dependent variable is firm-level expected investment in AI, taking values between 0 and 100. All specifications control for domestic and foreign priors and domestic and foreign gaps. Column (1) presents the baseline specification. Columns (2)–(4) add further controls as indicated at the bottom of the table. Firm characteristics are the following: employment, export share, and categorical variables for financial autonomy, annual sales, age, and ownership. Column (5) repeats the baseline specification after dropping the bottom and top 0.5% of the domestic and foreign gaps. Columns (6) and (7) repeat the baseline specification for alternative outcome variables: a dummy taking value 100 if the firm plans to invest in AI in the future and a dummy taking value 100 if the firm is currently using AI. Robust standard errors are in parentheses.

Table 4: Heterogeneity Analysis: 2SLS Estimates

	All firms (1)	Low sales (2)	High sales (3)	Low conc. (4)	High conc. (5)	Non-tradable (6)	Tradable (7)	Low exp. sh. (8)	High exp. sh. (9)
					Panel A: 2SLS				
$s^d_{i,t+1}$	0.570***	0.835*	0.430*	0.803*	0.418*	0.680**	0.465	0.710*	0.490*
	(0.167)	(0.403)	(0.174)	(0.397)	(0.177)	(0.256)	(0.268)	(0.286)	(0.204)
$s^f_{i,t+1}$	-0.224	-0.465	-0.092	-0.458	-0.042	-0.331	-0.097	-0.314	-0.184
	(0.158)	(0.377)	(0.165)	(0.356)	(0.178)	(0.233)	(0.264)	(0.265)	(0.194)
					Panel B: First-Stage (Dep. Var: Posterior Domestic)				
$T_i \cdot \left(s^{d,T}_{i,t} - s^{d,0}_{i,t}\right)$	0.259***	0.229**	0.303***	0.171*	0.331***	0.255***	0.238**	0.186**	0.342***
	(0.049)	(0.076)	(0.063)	(0.071)	(0.068)	(0.068)	(0.084)	(0.066)	(0.072)
$T_i \cdot \left(s^{f,T}_{i,t} - s^{f,0}_{i,t}\right)$	0.154**	0.249**	0.067	0.246**	0.067	0.116	0.181	0.265***	0.029
	(0.057)	(0.084)	(0.075)	(0.077)	(0.084)	(0.081)	(0.096)	(0.074)	(0.086)
					Panel C: First-Stage (Dep. Var: Posterior Foreign)				
$T_i \cdot \left(s^{d,T}_{i,t} - s^{d,0}_{i,t}\right)$	0.036	0.029	0.058	-0.014	0.081	0.054	0.021	-0.030	0.103
	(0.047)	(0.088)	(0.054)	(0.079)	(0.057)	(0.055)	(0.087)	(0.068)	(0.065)
$T_i \cdot \left(s^{f,T}_{i,t} - s^{f,0}_{i,t}\right)$	0.434***	0.491***	0.379***	0.492***	0.374***	0.386***	0.437***	0.518***	0.343***
	(0.056)	(0.097)	(0.067)	(0.089)	(0.073)	(0.069)	(0.102)	(0.079)	(0.080)
					Panel D: Reduced-Form				
$T_i \cdot \left(s^{d,T}_{i,t} - s^{d,0}_{i,t}\right)$	0.140***	0.178**	0.125**	0.144**	0.135**	0.156***	0.109*	0.142**	0.149***
	(0.031)	(0.056)	(0.038)	(0.048)	(0.041)	(0.041)	(0.051)	(0.045)	(0.044)
$T_i \cdot \left(s^{f,T}_{i,t} - s^{f,0}_{i,t}\right)$	-0.009	-0.020	-0.006	-0.028	0.012	-0.049	0.042	0.025	-0.049
	(0.037)	(0.061)	(0.046)	(0.055)	(0.052)	(0.049)	(0.064)	(0.052)	(0.053)
Observations	3,316	1,496	1,802	1,491	1,825	1,575	1,220	1,644	1,672
Dep. Var.: Baseline Mean (%)	8.33	8.50	8.22	8.36	8.31	6.85	9.98	7.95	8.72
Kleibergen-Paap Wald F-statistic	12.03	2.19	15.82	2.57	14.53	8.69	2.54	4.11	12.04

Notes: Panel A reports 2SLS estimates of Equation 37. The dependent variable is firm-level expected investment in AI, taking values between 0 and 100. All specifications control for domestic and foreign priors and domestic and foreign gaps. Column (1) repeats the baseline specification from Table 3. Columns (2) and (3) split the sample by sales categories, with a cutoff of EUR 2 million in 2024. Columns (4) and (5) split the sample depending on whether market concentration in the firm's country-industry-size class is below or above the median. Market concentration is defined using the Herfindahl-Hirschman Index (HHI) based on the latest available Orbis sales data, from 2023, for firms in the SAFE. Columns (6) and (7) split the sample by classifying the 2-digit NACE code into tradables/non-tradables. Columns (8) and (9) split the sample by export sales to total sales for 2024. Panels B and C report the corresponding first-stage estimates, and Panel D reports the reduced-form results. The sample is restricted to firms with non-missing observations across all AI variables. Robust standard errors are in parentheses.