Innovation without Borders?

The Geography of Innovation Diffusion

Ursel Baumann
European Central Bank

Zoe Cullen
Harvard University

Ester Faia
Goethe University Frankfurt

Annalisa Ferrando
European Central Bank

Ricardo Perez-Truglia
University of California, Los Angeles

Judit Rariga
European Central Bank

Pending approval for release—do not circulate or cite

Abstract

How well does the diffusion of innovation spread 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 elicited firms' perceptions of the share of similar firms in their own country that had invested in AI, as well as the share among firms in the three largest European countries. We randomly provided half of the sample with accurate signals about both domestic and foreign AI adoption. We show that firms substantially underestimated competitors' current AI investment, both domestically and abroad, and that they updated their expectations about competitors' future adoption in response to the information treatment. The treatment also caused a statistically significant ($p < 0.001$) increase in firms' own expected AI investment rate. We find strong strategic complementarities within borders: a 1 pp increase in the expected share of domestic competitors investing in AI raises a firm's own expected AI investment rate by 0.567 pp. These complementarities are absent across borders: the effect of an increase in the expected share of *foreign* competitors investing in AI on a firm's own adoption is close to zero and statistically insignificant. In sum, our evidence shows that both the diffusion of innovation and strategic complementarities in AI investment are much stronger at the domestic than at the foreign level.

JEL Classification: O33, D22, C93, L21.
Keywords: Innovation, diffusion, artificial intelligence, field experiment.

---

*This draft: March 26, 2026. Ursel Baumann: ursel.baumann@ecb.europa.eu, Zoe Cullen: zcullen@hbs.edu, Ester Faia: faia@wiwi.uni-frankfurt.de, Annalisa Ferrando: annalisa.ferrando@ecb.europa.eu, Ricardo Perez-Truglia: ricardotruglia@gmail.com, Judit Rariga: Erzsebet-Judit.Rariga@ecb.europa.eu. 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.

1 Introduction

The Schumpeterian view of innovation emphasizes that firms' incentives to adopt new technologies depend not only on their own returns, but also on the competitive environment in which they operate (see, e.g., Aghion and Howitt, 2006). This insight has motivated 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; Aghion et al., 2024). More broadly, the literature suggests that knowledge spillovers from abroad may be especially valuable for growth because they are more likely to embody genuinely novel ideas. Yet those same 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 trigger very different responses depending on whether those competitors are domestic or foreign.

Against this backdrop, we provide a large-scale field experiment on strategic complementarities in firms' AI investment. Our goal is to measure whether firms respond differently to domestic and foreign competitor investment. The key mechanisms are the same in both cases. Firms with imperfect information observe a signal about what other firms are doing. From that signal, they may learn that the technology is valuable, and they may also feel pressure to invest in order to avoid falling behind competitors. However, these forces need not operate symmetrically across domestic and foreign markets. This distinction is central for understanding the geography of innovation diffusion.

Our empirical setting is particularly well suited to this question. We study firms in twelve European countries using the Survey on the Access to Finance of Enterprises (SAFE), a long-running survey conducted by the European Central Bank and the European Commission. This setting is appealing for at least two reasons. First, the survey provides us with the necessary data to develop information treatments, as well as the opportunity to embed an experiment directly within a uniquely harmonized cross-country survey. Second, Europe is 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, the diffusion of AI remains highly uneven across countries and firms, suggesting that barriers to diffusion remain important even within an integrated economic area.

Our research design is motivated by models of strategic interaction under incomplete information, such as Cooper and John (1988), Angeletos and Pavan (2004), Angeletos et al. (2007), and more recently Huo and Takayama (2024). Through the lens of these models, firms

choose how much to invest in a technology without fully observing its underlying productivity. They instead form beliefs from noisy information and from what they think competitors are doing. A firm's optimal investment therefore depends not only on its own assessment of the returns to AI, but also on its expectations about competitors' AI investment. These expectations matter for at least two reasons. First, there is a competition channel: if competitors invest more, a firm may wish to invest more as well to avoid falling behind. Second, there is a learning channel: competitors' behavior may reveal information about the underlying profitability of the technology. In our setting, these channels may differ sharply depending on whether the relevant competitors are domestic or foreign. We developed a simple model that formalizes these ideas and yields testable implications.

The experiment was implemented in the fourth quarter of 2025. Using data from the Q2 2025 round of SAFE, we construct country-sector-size-specific signals about the share of comparable firms that had invested in AI up to June 2025. In the Q4 2025 wave, we first elicit firms' domestic prior beliefs and foreign prior beliefs: their beliefs about the current share of comparable firms investing in AI in their own country and in the largest three EU countries (Germany, France, and Italy). We then randomly disclose factual signals about both groups. After the information treatment, we elicit firms' domestic posterior beliefs and foreign posterior beliefs: their expectations about the future share of comparable firms investing in AI domestically and in the largest three EU countries. Finally, we measure the firm's own expected AI investment rate, defined as the share of its 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 research design has several advantages. First, the scale of the survey provides substantial statistical power and unusually broad cross-country variation. Second, because the information treatment is randomized, we can cleanly identify the causal effect of beliefs about competitor AI investment, ruling out many of the confounding forces that plague observational studies of strategic complementarities. Third, the simultaneous provision of domestic and foreign signals allows us to compare the strength of local and international spillovers within a common empirical framework.

We begin by documenting substantial heterogeneity in AI investment across European countries. AI investment is much more prevalent in some countries than in others, and AI use remains far from uniform even among adopters. We then show that firms hold large and systematically biased beliefs about competitor AI investment. Firms tend to underestimate both domestic and foreign competitor investment, and foreign prior beliefs are especially misaligned with the signals we provide. Moreover, firms often appear to have a distorted sense of their country's relative position within Europe, systematically perceiving AI investment

as more concentrated in a small set of prestigious economies than it actually is.

Next, we show that firms update their beliefs in response to the information treatment. Domestic posterior beliefs and foreign posterior beliefs both rise when firms receive factual information, and the updating pattern is broadly consistent with a Bayesian learning model. Firms that initially underestimate competitor AI investment revise their beliefs upward, firms that overestimate revise downward, and firms with accurate priors update little. This evidence indicates that respondents paid attention to the information and viewed it as informative.

Most importantly, the information treatment affects firms' own planned behavior. On average, treated firms report a higher own expected AI investment rate than control firms. We then use a 2SLS framework to estimate the causal effect of beliefs on intended investment. The results indicate that the firm's own expected AI investment rate responds strongly to domestic posterior beliefs but not to foreign posterior beliefs. In other words, firms react to learning that domestic competitors are investing in AI, but they react much less to learning that foreign competitors are doing so. This finding is consistent with the view that international knowledge spillovers are subject to a form of gravity: ideas may travel across borders, but their effect on firm behavior weakens with economic, cultural, and informational distance.

These findings have several implications. First, they suggest that informational frictions can materially slow the diffusion of new technologies. Second, they imply that foreign competition may be less effective than domestic competition in pushing firms to invest in AI, even in an integrated economic region such as Europe. Third, they help 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. 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 paper relates to several strands of the literature. Most directly, it contributes to the literature on strategic complementarities and coordination. On the theoretical side, this literature builds on foundational work such as Cooper and John (1988) and the global-games framework in Angeletos and Pavan (2004). Our empirical strategy is also closely related to recent work on information frictions and dynamic learning under incomplete information (Huo and Takayama, 2024). On the empirical side, identifying strategic complementarities is difficult because correlated behavior across firms may reflect shared shocks rather than causal spillovers. We contribute to this literature by bringing experimental variation to the study of firms' technology investment decisions.

The paper is also related to recent work on innovation, technology adoption, and international knowledge spillovers. Most closely related is Cullen et al. (2025b), who use an infor-

mation experiment with Italian firms to study strategic complementarities in the adoption of advanced technologies within countries and sectors. We extend that logic to a multi-country setting and use it to compare domestic and foreign spillovers directly. More broadly, our findings connect to the literature on international knowledge diffusion and its geographic frictions (Coe and Helpman, 1995; Eaton and Kortum, 2002; Keller, 2004; Keller and Yeaple, 2013), as well as to work on technology adoption and 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). Finally, given our focus on AI, our paper speaks to the growing literature on the determinants, consequences, and geography of advanced-technology adoption (Zolas et al., 2021; Agrawal et al., 2019; Kalyani et al., 2025).

The rest of the paper is organized as follows. Section 3 describes the institutional setting and data. Section 4 presents details about the design and implementation of the experiment. Section 5 discusses AI investment and firms' misperceptions about competitor AI investment. Section 6 documents the effect of the information treatment on firms' beliefs. 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). 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, $c \in \{1, 2\}$. 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 adopting a technology i, $k_{i,c}$, in country c features 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, θ_i, but also on some knowledge spillovers that occur across countries. Those are captured by the parameters A_i, which we characterize next.

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

$$A_1 = (1 - \lambda_d - \lambda_f)\theta_1 + \lambda_d K_1 + \lambda_f K_2, \tag{3}$$

$$A_2 = (1 - \lambda_d - \lambda_f)\theta_2 + \lambda_d K_2 + \lambda_f K_1. \tag{4}$$

The definition above make clear that the total return to investing in the technology in country i depend on the country level productivity, θ_i, but also on the aggregate investment made in the other country. To capture this, we model the functional form of the complementarities as follows:

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

where $\lambda \in [0, 1)$ governs the overall strength of complementarities, and $\mu \in \left[0, \frac{1}{2}\right]$, which we refer to as *connectedness*, governs the share of cross-country complementarities. At one extreme, if $\mu = \frac{1}{2}$, 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 $\mu = 0$, the countries are fully disconnected: complementarities exist between firms within the same country but are absent across countries.

It is important to put the assumption $\mu \leq \frac{1}{2}$ 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 $Var(\theta_1) = Var(\theta_2) = 1$. Fundamentals are jointly normal:

$$\begin{pmatrix}\theta_1\\ \theta_2\end{pmatrix} \sim \mathcal{N}\left(\begin{pmatrix}0\\ 0\end{pmatrix}, \begin{pmatrix}1 & \rho_\theta\\ \rho_\theta & 1\end{pmatrix}\right), \qquad 0 \le \rho_\theta < 1. \tag{6}$$

The case $\rho_\theta = 0$ provides a benchmark with no learning between fundamentals across countries. Positive correlation induces knowledge spillover across countries. Practically this maybe due to proximity: countries can share technological advanced 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:

$$z_1 = \theta_1 + \tau_z^{-1/2}\varepsilon_1, \qquad z_2 = \theta_2 + \tau_z^{-1/2}\varepsilon_2, \tag{7}$$

with $\varepsilon_1, \varepsilon_2 \sim \mathcal{N}(0,1)$ i.i.d. and independent of (θ_1, θ_2). Each firm receives a private signal about its own country's fundamental:

$$x_{i,1} = \theta_1 + \tau_x^{-1/2}\xi_{i,1}, \qquad x_{i,2} = \theta_2 + \tau_x^{-1/2}\xi_{i,2}, \tag{8}$$

where $\xi_{i,1}, \xi_{i,2} \sim \mathcal{N}(0,1)$ 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:

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

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:

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

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

$$\mathbb{E}_{i,1}[\theta_1] = a_x x_{i,1} + a_1 z_1 + a_2 z_2, \tag{11}$$

$$\mathbb{E}_{i,1}[\theta_2] = b_x x_{i,1} + b_1 z_1 + b_2 z_2, \tag{12}$$

with:

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

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

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 cross of the reaction function. 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:

$$k_{i,1} = \mathbb{E}_{i,1}[A_1] = (1-\lambda)\mathbb{E}_{i,1}[\theta_1] + \lambda(1-\mu)\mathbb{E}_{i,1}[K_1] + \lambda\mu\,\mathbb{E}_{i,1}[K_2], \tag{15}$$

$$k_{i,2} = \mathbb{E}_{i,2}[A_2] = (1-\lambda)\mathbb{E}_{i,2}[\theta_2] + \lambda(1-\mu)\mathbb{E}_{i,2}[K_2] + \lambda\mu\,\mathbb{E}_{i,2}[K_1]. \tag{16}$$

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 (2024).

functions as follows:

$$k_{i,1} = \beta x_{i,1} + \gamma_d z_1 + \gamma_f z_2, \tag{17}$$

$$k_{i,2} = \beta x_{i,2} + \gamma_d z_2 + \gamma_f z_1. \tag{18}$$

The individual guesses for the policy functions in 18 allow us to derive the aggregate investment, which we use to obtain the equilibrium firms' adoptions and which read as follows:

$$K_1 = \beta\theta_1 + \gamma_d z_1 + \gamma_f z_2, \qquad K_2 = \beta\theta_2 + \gamma_f z_1 + \gamma_d z_2. \tag{19}$$

The equilibrium is fully summarized by $(\beta, \gamma_d, \gamma_f)$ in (A.2)–(A.5). These coefficients depend on primitives only through $(\lambda, \mu, \rho_\theta, \tau_x, \tau_z)$ 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 solution for the $(\beta, \gamma_d, \gamma_f)$ coefficients in Lemma 1. We also derive the closed form solution for the term $\gamma_d - \gamma_f$. This statistics 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 $\mu \in [0, 1/2]$, the matrix in Step 2 below is automatically invertible.[2] We therefore impose the single remaining condition needed for a unique symmetric linear equilibrium:

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

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

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

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

This restriction ensures that the informational advantage of the domestic public signal for θ_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 $(\beta, \gamma_H, \gamma_F)$, we can decompose the effect of the domestic and foreign signals in four different components: the learning effects (and subsequently between domestic learning foreign learning) and the competition effects (domestic and foreign).

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

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:

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

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

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

Proposition 2.1 (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 matter 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 2.2 (The response gap shrinks with cross-country complementarities)**.** *Maintain Assumptions 1–5, so that $\gamma_H - \gamma_F > 0$. The gap*

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

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

(See Appendix A for the proof.)

Intuitively, when the economy is less open (lower μ), 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 2.3 (The response gap shrinks with correlated fundamentals)**.** *When $\mu = 0$, the gap $\gamma_d - \gamma_f$ is weakly decreasing in $\rho_\theta \in [0, 1)$ and strictly decreasing on $\rho_\theta \in (0, 1)$. Moreover, by continuity there exists $\varepsilon > 0$ such that for all $\mu \in [0, \varepsilon)$ 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 θ_1, while z_2 matters for θ_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

3.1 European Firms

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 particular 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 Mario Draghi brought renewed

[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).

attention to these concerns, emphasizing Europe's slow productivity growth, its innovation gap, and the difficulty of translating ideas into scalable commercial success.[4]

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,[5] 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. A recent report by Enrico Letta highlights that the EU internal market continues to face significant segmentation across several dimensions.[6] Particularly relevant for our setting, that report emphasizes that Europe still lacks a truly integrated market for knowledge, innovation, and research.[7] 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.[8] 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

[4]From the Draghi report: "The problem is not that Europe lacks ideas or ambition. We have many talented researchers and entrepreneurs filing patents. But innovation is blocked at the next stage: we are failing to translate innovation into commercialization, and innovative companies that want scale up in Europe are hindered at every stage by inconsistent and restrictive regulations...". Between 2008 and 2021, close to 30% of the "unicorns" founded in Europe—startups that went on to be valued over USD 1 billion—relocated their headquarters abroad, with the vast majority moving to the US. With the world on the cusp of an AI revolution, Europe cannot afford to remain stuck in the "middle technologies and industries" of the previous century. This will be key not only to lead in new technologies, but also to integrate AI into our existing industries so that they can stay at the front. The report also points to investment needs on the order of 4–5% of EU GDP per year if Europe is to remain competitive relative to the US and China.

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

[6]European markets are particularly fragmented in areas such services, financial markets, energy markets, telecommunications and digital networks. These challenges are further intensified by national administrative and regulatory burden as well as from limited labour mobility.

[7]Letta proposes a "fifth freedom" to complement the single market in goods, services, people and capital.

[8]Source: https://reforms-investments.ec.europa.eu/recovery-and-resilience-facility-1_en.

set of locations? Despite harmonized regulation through the EU AI Act and centralized funding initiatives, large disparities remain. Adoption 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.

3.2 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 select 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 and Q4 2025 covered questions on AI. Firms were first asked about the extent of their AI adoption, and then asked to report the share of their total investment that they expected to devote to AI over the subsequent twelve months.

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.

Our sample includes firms from the twelve major euro area 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.

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 D and outlined below.[9]

The module begins by eliciting whether the firm had adopted AI technologies, using a scale that distinguishes between no current use, very infrequent or experimental use, moderate use, and significant use. To better understand adoption patterns, firms reporting some use of AI were then asked to identify the two main motivations behind their adoption, 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 adoption, 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.

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

In Germany, France and Italy, [Y%] of firms similar to yours had made investments in AI technologies up to June 2025."

And 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 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 E 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 was 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 35 respondents per cluster, a minimum of 1 respondents, and a maximum of 174 respondents. Table C.1 reports 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*.

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 Adoption and Misperceptions

5.1 Baseline AI Adoption and Investment Rates

Panel A of Figure 1 reports the average signal on AI adoption 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 C.1. 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 a wide spread variation across the distribution.[10] The share of firms investing in AI ranges from 20% in Greece to 48% in Austria. The Big-3 average is 32.6%, almost exactly in the middle of the cross-country distribution, and is itself composed of 40% in Germany, 27% in France, and 23% 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. 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, even among adopters, 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 for AI adoption. Firms that reported using AI were

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

asked to select the two main reasons for adoption 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, the two most common reasons for *not* adopting AI are lack of usefulness and lack of skills, with the corresponding results reported in Appendix C.

5.2 Misperceptions about Competitors' Adoption

Panel A plots the distribution of firms' misperceptions about competitors' AI investment, measured as the difference between the signal and the 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.

Panel B 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

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.[11] More broadly, the figure points to a center-periphery misconception: firms may perceive AI investment as concentrated in a few prestigious economies even though, in reality, it is more widely dispersed.

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 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 ($p<0.001$).

Panel B of Figure 4 provides a falsification test. 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%

[11] 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).

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:

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

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.[12]

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

[12]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.

of a Bayesian model with Gaussian priors and signals,[13] 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:

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

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 adoption 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 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

[13] 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.

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:[14]

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

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

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

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

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:

$$\begin{aligned} s^d_{i,t+1} = {} & \gamma^d_0 + \gamma^d_{1d}\, T_i \left(s^{d,T}_{i,t} - s^{d,0}_{i,t}\right) + \gamma^d_{1f}\, T_i \left(s^{f,T}_{i,t} - s^{f,0}_{i,t}\right) \\ & + \gamma^d_{2d} \left(s^{d,T}_{i,t} - s^{d,0}_{i,t}\right) + \gamma^d_{2f} \left(s^{f,T}_{i,t} - s^{f,0}_{i,t}\right) + \gamma^d_3\, s^{d,0}_{i,t} + \gamma^d_4\, s^{f,0}_{i,t} + \epsilon^d_{i,t}. \end{aligned} \tag{35}$$

[14] 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.

$$s^f_{i,t+1} = \gamma^f_0 + \gamma^f_{1f}\, T_i \left(s^{f,T}_{i,t} - s^{f,0}_{i,t}\right) + \gamma^f_{1d}\, T_i \left(s^{d,T}_{i,t} - s^{d,0}_{i,t}\right) + \gamma^f_{2f} \left(s^{f,T}_{i,t} - s^{f,0}_{i,t}\right) + \gamma^f_{2d} \left(s^{d,T}_{i,t} - s^{d,0}_{i,t}\right) + \gamma^f_3\, s^{f,0}_{i,t} + \gamma^f_4\, s^{d,0}_{i,t} + \epsilon^f_{i,t}. \quad (36)$$

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.[15] These binned scatterplots also support the linear functional-form assumptions implicit in the Bayesian learning model.

Table C.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 C.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 C.2 yields $\beta^{f,f} = 1.07$, which is statistically significant, and $\beta^{f,d} = 0.09$, which is statistically insignificant. Thus, when forming foreign posterior beliefs, firms

[15] 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 C.2.

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 of <0.001 and 0.007).[16] 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, 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. (2025b) 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 0.42 in column (2)—we cannot reject the null that they are equal (p=0.885). 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=0.719). Taken together, these comparisons support the internal consistency of the learning model.

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

7 Effect of the Information Treatment on Firms' Own Future Adoption 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.

Panel B of Figure 4 provides the corresponding falsification test. Its third pair of columns uses a pre-treatment outcome: whether the firm had already adopted AI. 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 $T_i \cdot (s^{f,T}_{i,t} - s^{f,0}_{i,t})$.[17] The vector X_i includes additional covariates, such as employment and sales.[18]

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.

[17] 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 $\{s^{d,T}_{i,t} - s^{d,0}_{i,t}, s^{d,0}_{i,t+1}, s^{f,T}_{i,t} - s^{f,0}_{i,t}, s^{f,0}_{i,t+1}, X_i\}$. In plain English, we assume that heterogeneity in the effects of information is driven solely by differences in prior misperceptions, and not by heterogeneity in other unobserved factors that are correlated with those misperceptions.

[18] 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.

The 2SLS regression implements exactly this comparison.[19]

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 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 2 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 2, where the dependent variable is the firm's own expected AI investment rate. Consistent with the binned-scatterplot evidence

[19]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.

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 0.567 and statistically significant (p-value=0.001), while the coefficient on the foreign posterior belief is -0.221 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.

In particular, the estimate of 0.567 in column (1) of Table 2 implies that a 1 pp increase in the domestic posterior belief raises the firm's own expected AI investment rate by 0.567 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.

The 2SLS coefficient on the domestic posterior belief can be made even more interpretable by expressing it as an elasticity. The average firm in the control group has a domestic posterior belief of 23.59% and an own expected AI investment rate of 8.33%. Thus, a 1 pp increase in the domestic posterior belief represents a 4.24% increase relative to the baseline ($= \frac{1}{23.59}$), while the associated 0.567 pp increase in the firm's own expected AI investment rate corresponds to a 6.81% increase relative to its baseline level ($= \frac{0.567}{8.33}$). Together, these figures imply an elasticity of 1.606 ($= \frac{6.81}{4.24}$) between the domestic posterior belief and the firm's own expected AI investment rate. Put differently, a 10% increase in the expected share of domestic competitors investing in AI would raise the firm's own expected AI investment rate by 16.06%.

Before interpreting these results further, we 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 2 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 2 reports the Kleibergen-Paap *rk* Wald F-statistic for each specification, a standard measure of instrument relevance.[20] 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.

[20]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.

Columns (2) through (6) of Table 2 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 adoption. 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.[21]

Column (5) of Table 2 presents results for the extensive margin. It is identical to column (1), 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 (5) is positive but imprecisely estimated and therefore statistically insignificant.

Finally, column (6) of Table 2 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 had already adopted AI 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 2. The 2SLS coefficients may reflect at least two mechanisms: competition, as firms seek to avoid being left behind by rivals, and learning, as firms draw inferences about their own returns to AI from competitors' investment choices. Because these channels are closely intertwined, it is difficult to disentangle them completely; for related discussion and evidence, see Cullen et al. (2025a). Nonetheless, we can provide suggestive evidence.

Our approach is to examine treatment-effect heterogeneity. For domestic posterior beliefs, the competition channel should be stronger in domestic markets that are less concentrated—in the limiting case where a firm faces no domestic competitors, this channel would be absent. For foreign posterior beliefs, the competition channel should be stronger when firms are more exposed to international competition—in the limiting case where a firm faces no foreign competitors, this channel would also be absent.

Table 3 presents the heterogeneity results. Column (1) reproduces the baseline estimates

[21] 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.

from column (1) of Table 2. The remaining columns apply the same specification to different subsamples. For example, columns (2)–(3) split the sample into small and medium-sized firms in column (2) and large firms in column (3). The effects are stronger for smaller firms than for large firms. To the extent that smaller firms are more concerned about domestic competition, this pattern provides suggestive evidence that the competition channel may be important. The final heterogeneity exercise is reported in the last two columns, which split the sample by the importance of exports in firm revenues. Column (8) corresponds to firms with below-median export shares, while column (9) corresponds to firms with above-median export shares. Firms with greater export exposure may be more exposed to foreign competition. Even in that subsample, however, we do not find a significant effect of foreign posterior beliefs.

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 imply 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. By contrast, we do not find comparable effects for foreign posterior beliefs. This asymmetry helps explain why AI diffusion may remain geographically uneven even in an integrated economic area such as Europe: firms may learn from and respond to foreign competitors, but the behavioral response is much weaker than for domestic competitors.

We conclude by discussing the policy implications. Europe has devoted substantial resources to fostering innovation and AI adoption, 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 adoption and innovation, lowering these invisible barriers to cross-border spillovers may be as important as expanding financial support.

References

Aghion, P., A. Bergeaud, M. Lequien, and M. J. Melitz (2024). The heterogeneous impact of market size on innovation: Evidence from french firm-level exports. *Review of Economics and Statistics 106*(3), 608–626.

Aghion, P. and P. Howitt (2006). Joseph schumpeter lecture appropriate growth policy: A unifying framework. *Journal of the european Economic Association 4*(2-3), 269–314.

Agrawal, A., J. Gans, and A. Goldfarb (2019). *The economics of artificial intelligence: An agenda*. University of Chicago Press.

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., C. Hellwig, and A. Pavan (2007). Dynamic global games of regime change: Learning, multiplicity, and the timing of attacks. *Econometrica 75*(3), 711–756.

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

Atkin, D., A. K. Khandelwal, and A. Osman (2017). Exporting and firm performance: Evidence from a randomized experiment. *The quarterly journal of economics 132*(2), 551–615.

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.

Bilgin, N. M., E. Faia, and G. Ottaviano (2024). Technology spillovers, diffusion and rivalry in firm networks. Technical report, CEPR Discussion Paper.

Cavallo, A., G. Cruces, and R. Perez-Truglia (2017). Inflation Expectations, Learning, and Su-

permarket 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.

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 (2025a). The Innovation Race: Experimental Evidence on Advanced Technologies. *NBER Working Paper No. 34532*.

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.

Cullen, Z. B., E. Faia, E. Guglielminetti, R. Perez-Truglia, and C. Rondinelli (2025b). The innovation race: Experimental evidence on advanced technologies. Technical report, National Bureau of Economic Research.

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.

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

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.

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 (2024). Rational Expectations Models with Higher-Order Beliefs. *Review of Economic Studies*, rdae096.

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

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.

Lind, N. and N. Ramondo (2023). Global innovation and knowledge diffusion. *American Economic Review: Insights 5*(4), 494–510.

Pavcnik, N. (2002). Trade liberalization, exit, and productivity improvements: Evidence from chilean plants. *Review of Economic Studies 69*(1), 245–276.

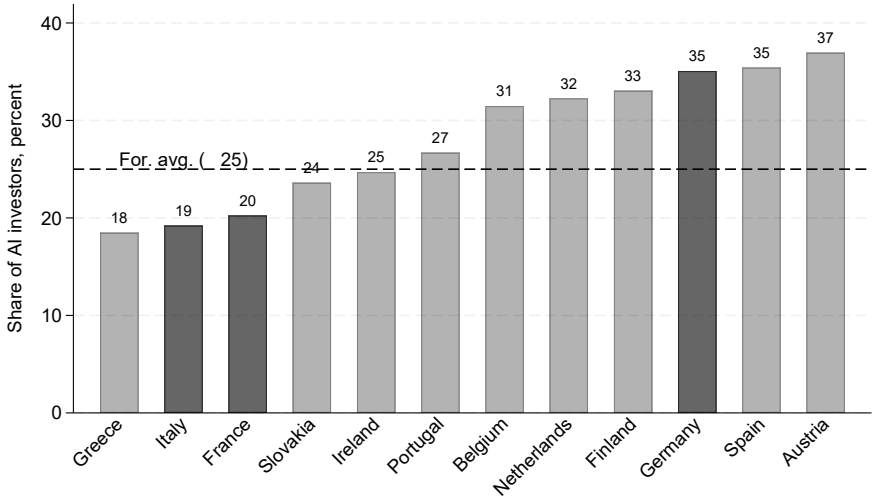
Staiger, D. and J. H. Stock (1997). Instrumental Variables Regression with Weak Instruments. *Econometrica 65*(3), 557–586.

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.

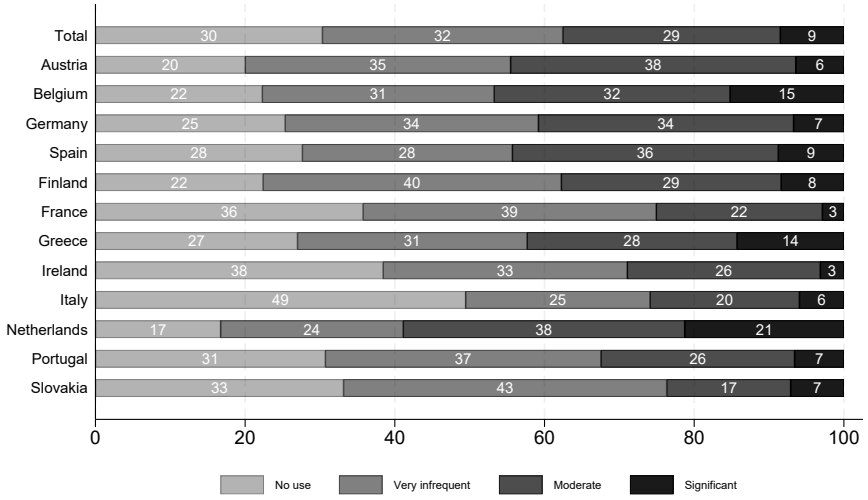
Zolas, N., Z. Kroff, E. Brynjolfsson, K. McElheran, D. N. Beede, C. Buffington, N. Goldschlag, L. Foster, and E. Dinlersoz (2021). Advanced technologies adoption and use by us firms: Evidence from the annual business survey. Technical report, National Bureau of Economic Research.

Figure 1: Descriptive Statistics about AI Investment and Adoption rate

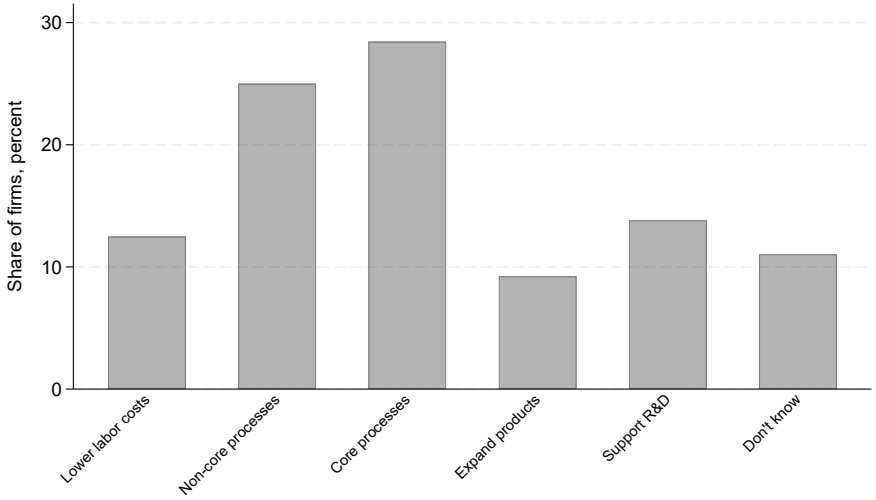
PANEL A: Average Signals by Country



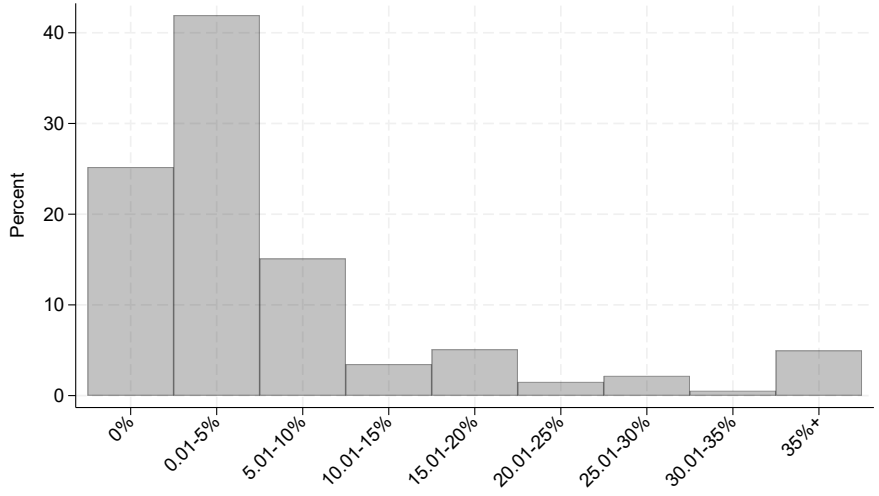
PANEL B: Adoption Rates by Country



PANEL C: Reasons for AI Adoption



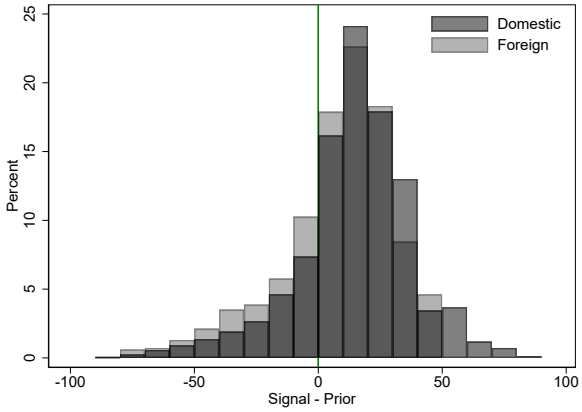
PANEL D: Expected AI Investment Rate (Control Group)



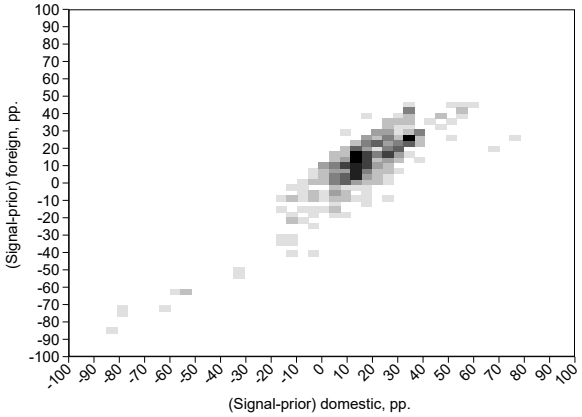
Notes: Panel A reports average AI investment signals by country. Panel B shows AI adoption rates by country. Panel C presents firms' reported reasons for AI adoption. Panel D displays the distribution of expected AI investment rates in the control group.

Figure 2: Distribution of Misperceptions about Adoption of Domestic and International Competitors

PANEL A: Distribution of Misperceptions

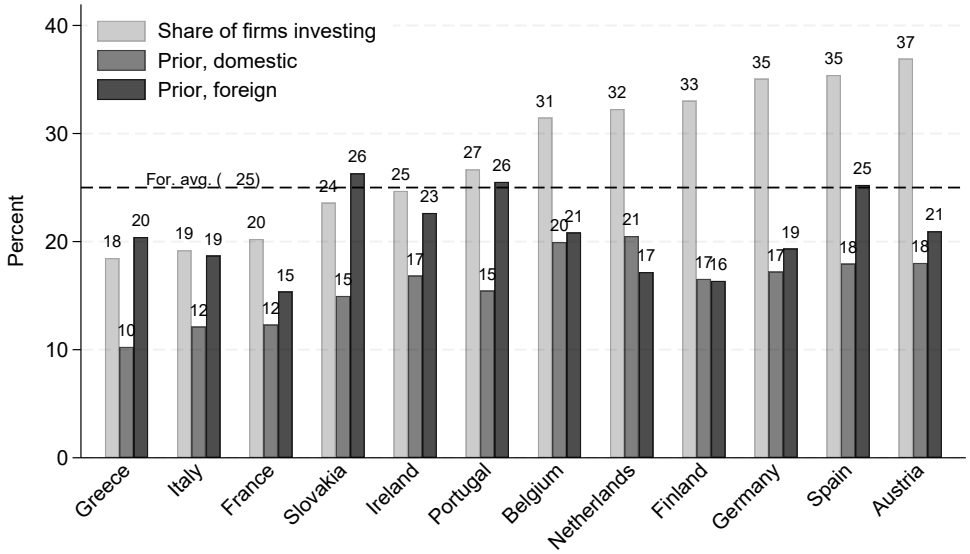


PANEL B: Correlation of Misperceptions



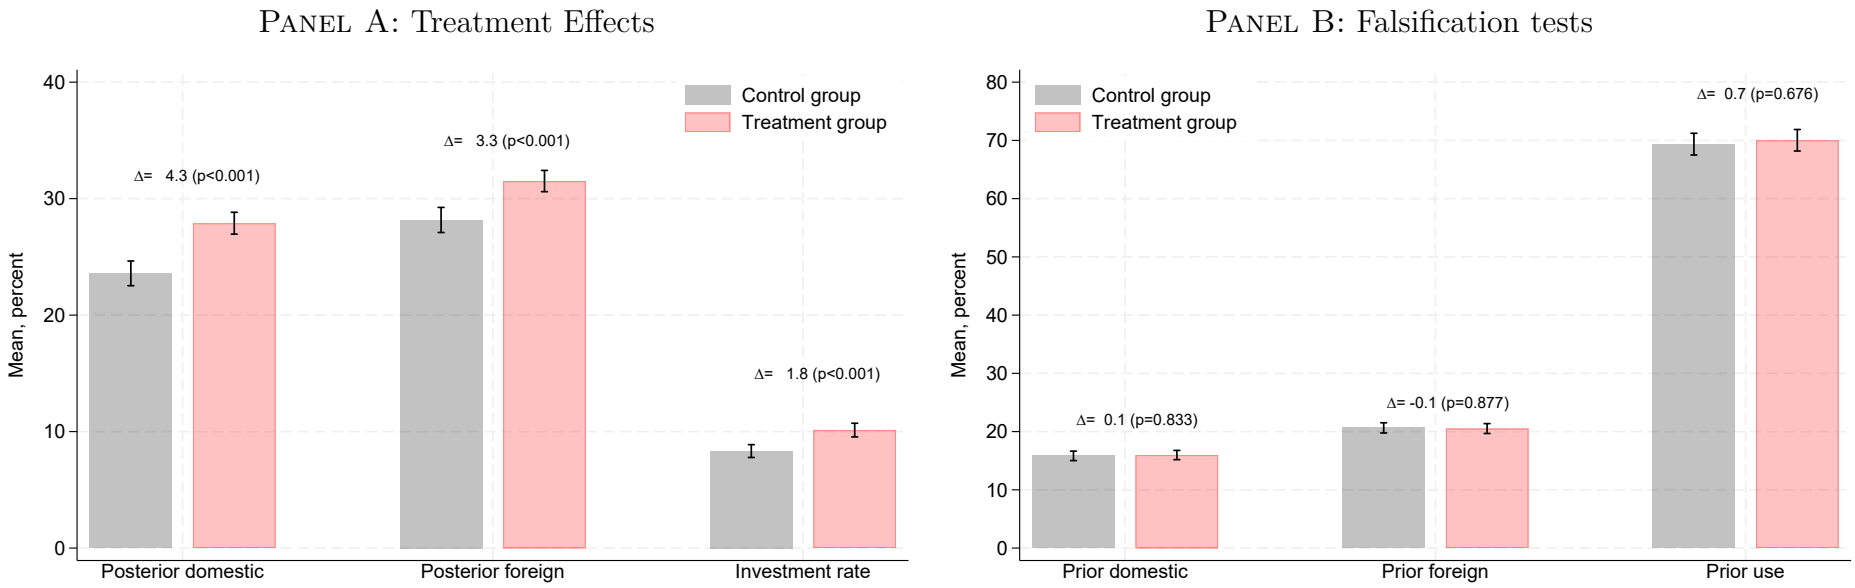
Notes: Panel A shows the distribution of misperceptions about adoption by domestic and foreign competitors, measured as signal minus prior. Panel B plots the relationship between domestic and foreign misperceptions.

Figure 3: Actual and Perceived Domestic and Foreign Adoption by Country



Notes: This figure compares the actual share of firms investing in AI with perceived domestic and foreign competitor adoption rates across countries.

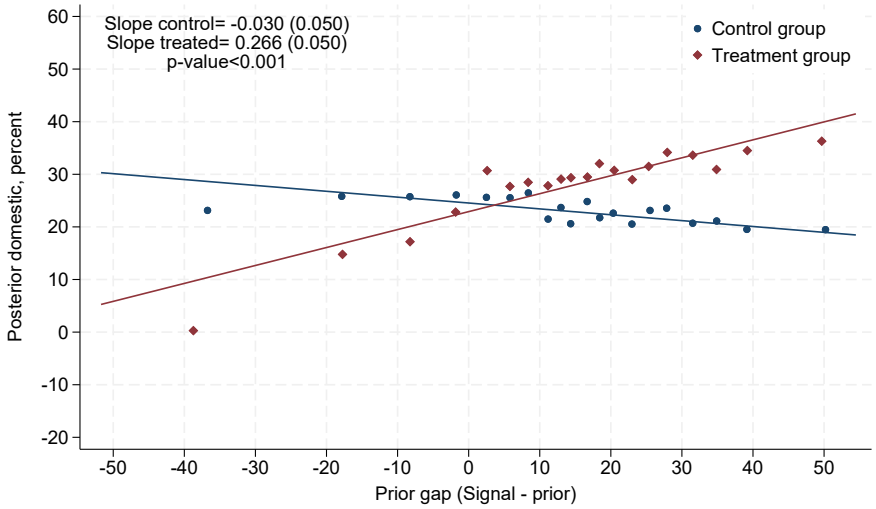
Figure 4: Treatment Effects on Posterior Beliefs and AI Investment Rate, with Falsification Tests



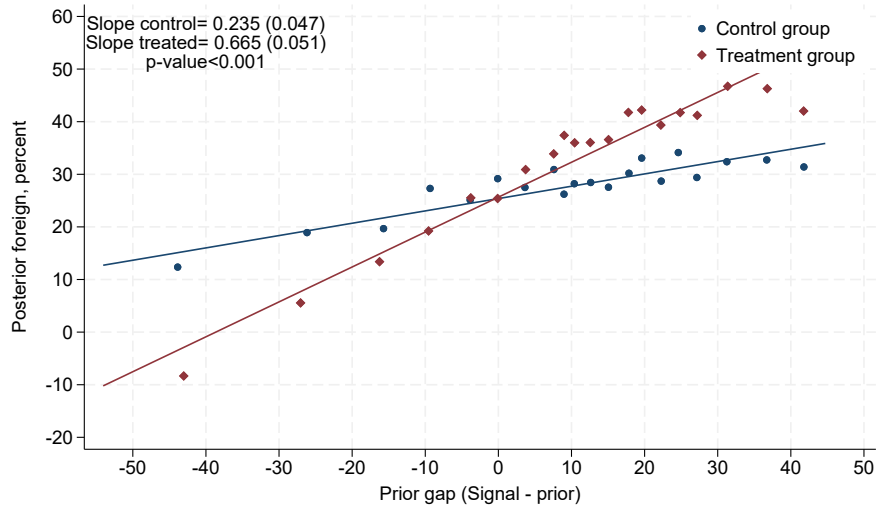
Notes: Panel A presents treatment effects on posterior beliefs about domestic and foreign adoption and on the AI investment rate, comparing control and treatment groups. Panel B reports falsification tests comparing control and treatment groups for prior domestic beliefs, prior foreign beliefs, and prior AI use.

Figure 5: Belief Updating: Perceptions about Adoption of Domestic and International Competitors

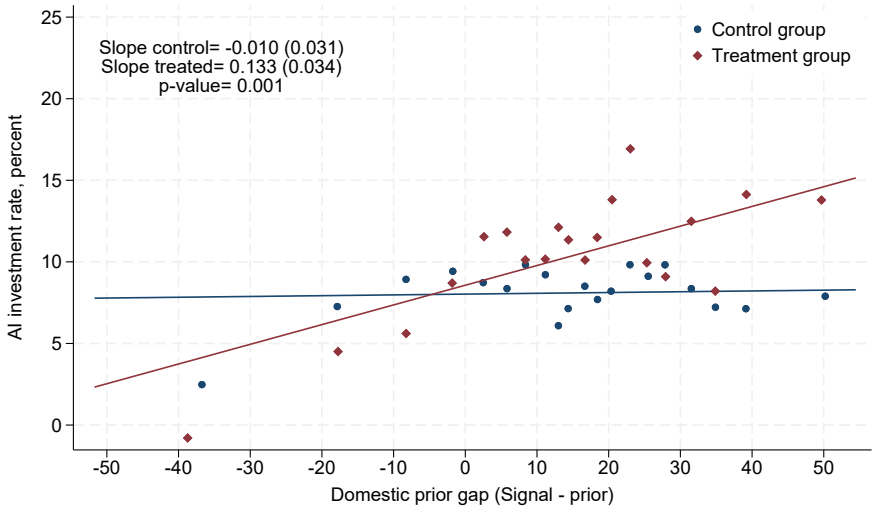
PANEL A: Effect on Domestic Posterior by Domestic Gap



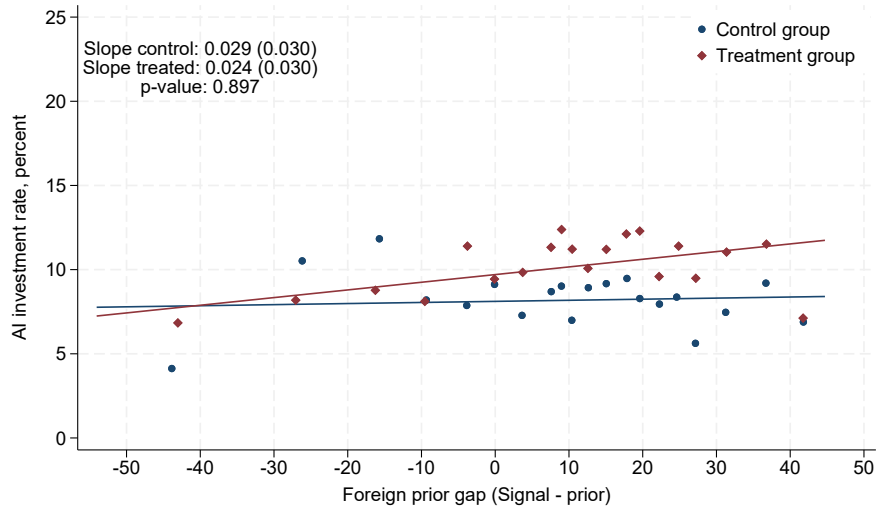
PANEL B: Effect on Foreign Posterior by Foreign Gap



PANEL C: Effect on Investment Rate by Domestic Gap



PANEL D: Effect on Investment Rate by Foreign Gap



Notes: Panel A plots 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 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	Control	Treatment	p-value
	(1)	(2)	(3)	(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: The Table presents summary statistics for a set of firm characteristics based on survey replies in Q4 2025 (share of firms by industry, financial autonomy, age groups, ownership and annual sales). Standard errors in parenthesis.

Table 2: 2SLS estimates: Effects of Competitors Adoption on Own Future Investment

	Inv. rate (1)	Inv. rate (2)	Inv. rate (3)	Inv. rate (4)	Inv. dummy (5)	Prior use (6)
			Panel A: 2SLS			
$s^d_{i,t+1}$	0.567***	0.560***	0.539***	0.542***	0.640	0.262
	(0.166)	(0.165)	(0.162)	(0.162)	(0.356)	(0.410)
$s^f_{i,t+1}$	-0.221	-0.210	-0.192	-0.200	-0.132	-0.485
	(0.157)	(0.155)	(0.151)	(0.152)	(0.326)	(0.379)
			Panel B: First-Stage (Dep. Var: Posterior Domestic)			
$T_i \cdot \left(s^{d,T}_{i,t} - s^{d,0}_{i,t}\right)$	0.260***	0.257***	0.279***	0.281***	0.260***	0.260***
	(0.049)	(0.048)	(0.048)	(0.048)	(0.049)	(0.049)
$T_i \cdot \left(s^{f,T}_{i,t} - s^{f,0}_{i,t}\right)$	0.153**	0.160**	0.139*	0.139*	0.153**	0.153**
	(0.057)	(0.056)	(0.055)	(0.055)	(0.057)	(0.057)
			Panel C: First-Stage (Dep. Var: Posterior Foreign)			
$T_i \cdot \left(s^{d,T}_{i,t} - s^{d,0}_{i,t}\right)$	0.038	0.035	0.062	0.065	0.038	0.038
	(0.047)	(0.046)	(0.047)	(0.047)	(0.047)	(0.047)
$T_i \cdot \left(s^{f,T}_{i,t} - s^{f,0}_{i,t}\right)$	0.432***	0.438***	0.415***	0.414***	0.432***	0.432***
	(0.056)	(0.055)	(0.056)	(0.057)	(0.056)	(0.056)
			Panel D: Reduced-Form			
$T_i \cdot \left(s^{d,T}_{i,t} - s^{d,0}_{i,t}\right)$	0.139***	0.137***	0.138***	0.139***	0.161*	0.050
	(0.031)	(0.030)	(0.031)	(0.031)	(0.081)	(0.089)
$T_i \cdot \left(s^{f,T}_{i,t} - s^{f,0}_{i,t}\right)$	-0.009	-0.002	-0.005	-0.007	0.041	-0.169
	(0.037)	(0.036)	(0.037)	(0.037)	(0.091)	(0.101)
Control for Prior AI Use Dummies	No	Yes	Yes	Yes	No	No
Cou.-Ind.-Size FE	No	No	Yes	Yes	No	No
Control for Other Firm Characteristics	No	No	No	Yes	No	No
Observations	3,316	3,316	3,316	3,316	3,316	3,316
Dep. Var.: Baseline Mean (%)	8.33	8.33	8.33	8.33	74.82	69.36
Kleibergen-Paap Wald F-statistic	12.13	11.89	12.35	12.22	12.13	12.13

Notes: Sample of firms reporting all AI variables of interest. In all specifications, control for gaps and priors. AI use int: categorical var (no, limited, moderate, significant) in the FE: size: as for the treatment, below, above 50 empl, Firm controls: employment, export share - both continuous variables, categorical vars for: financial autonomy, sales_annual, age and ownership. Column 6 and 7: outcome variable is a dummy, taking value 100, if firm expects to invest in AI or uses AI.

Table 3: 2SLS estimates: Heterogeneous Effects of Competitors Adoption on Own Future Investment

	All firms (1)	SME (2)	Large (3)	No financing problems (4)	Financing problems (5)	Low union share (6)	High union share (7)	Low exp. share (8)	High exp. share (9)
				Panel A: 2SLS					
$s^d_{i,t+1}$	0.567***	0.664***	0.148	0.549**	0.656	0.829*	0.319	0.709*	0.484*
	(0.166)	(0.197)	(0.339)	(0.192)	(0.346)	(0.326)	(0.172)	(0.284)	(0.202)
$s^f_{i,t+1}$	-0.221	-0.291	0.053	-0.237	-0.179	-0.466	0.001	-0.313	-0.180
	(0.157)	(0.182)	(0.381)	(0.185)	(0.297)	(0.316)	(0.160)	(0.264)	(0.193)
				Panel B: First-Stage (Dep. Var: Posterior Domestic)					
$T_i \cdot \left(s^{d,T}_{i,t} - s^{d,0}_{i,t}\right)$	0.260***	0.225***	0.541***	0.286***	0.203*	0.210**	0.326***	0.187**	0.343***
	(0.049)	(0.052)	(0.135)	(0.060)	(0.083)	(0.067)	(0.068)	(0.066)	(0.072)
$T_i \cdot \left(s^{f,T}_{i,t} - s^{f,0}_{i,t}\right)$	0.153**	0.193**	-0.207	0.153*	0.129	0.162*	0.139	0.263***	0.027
	(0.057)	(0.059)	(0.165)	(0.069)	(0.101)	(0.078)	(0.080)	(0.073)	(0.086)
				Panel C: First-Stage (Dep. Var: Posterior Foreign)					
$T_i \cdot \left(s^{d,T}_{i,t} - s^{d,0}_{i,t}\right)$	0.038	0.013	0.255*	0.052	0.006	0.004	0.082	-0.027	0.104
	(0.047)	(0.050)	(0.113)	(0.059)	(0.075)	(0.072)	(0.057)	(0.067)	(0.065)
$T_i \cdot \left(s^{f,T}_{i,t} - s^{f,0}_{i,t}\right)$	0.432***	0.465***	0.132	0.432***	0.421***	0.407***	0.461***	0.516***	0.342***
	(0.056)	(0.060)	(0.144)	(0.070)	(0.093)	(0.086)	(0.068)	(0.079)	(0.079)
				Panel D: Reduced-Form					
$T_i \cdot \left(s^{d,T}_{i,t} - s^{d,0}_{i,t}\right)$	0.139***	0.146***	0.093	0.145***	0.132*	0.172***	0.104**	0.141**	0.147***
	(0.031)	(0.033)	(0.095)	(0.037)	(0.056)	(0.047)	(0.039)	(0.044)	(0.044)
$T_i \cdot \left(s^{f,T}_{i,t} - s^{f,0}_{i,t}\right)$	-0.009	-0.007	-0.024	-0.018	0.009	-0.055	0.045	0.025	-0.048
	(0.037)	(0.039)	(0.120)	(0.044)	(0.068)	(0.055)	(0.048)	(0.052)	(0.053)
Observations	3,316	3,031	285	2,460	816	1,835	1,481	1,644	1,672
Dep. Var.: Baseline Mean (%)	8.33	8.19	9.86	7.98	9.30	7.91	8.87	7.95	8.72
Kleibergen-Paap Wald F-statistic	12.13	9.86	6.14	7.04	5.82	3.88	12.47	4.15	12.11

Notes: Specification like column (1) of Table 2. SMEs are firms with less than 250 employees. Financing problems defined based on firms reply to the survey question on whether access to finance is a problem. Median country level union share: 14%. Median firm size-industry-country level export share: 12%. Firms replying at least 7 on a scale of 1 to 10 are considered to have problems with access to finance. Trade Union: Proportion of employees who are member of a trade union as a percentage of the total number of employees.

Online Appendix (For Online Publication Only)

Innovation without Borders?

U. Baumann, Z. Cullen, E. Faia, A. Ferrando,
R. Perez-Truglia and J. Rariga (March 26, 2026)

A Proofs of Propositions

Proof Lemma 1. We can derive the equilibrium policies in two steps:

Step 1: solve for β

A country-1 firm forms

$$\mathbb{E}_{i,1}[K_1] = \beta\mathbb{E}_{i,1}[\theta_1] + \gamma_d z_1 + \gamma_f z_2, \qquad \mathbb{E}_{i,1}[K_2] = \beta\mathbb{E}_{i,1}[\theta_2] + \gamma_f z_1 + \gamma_d z_2. \tag{A.1}$$

Substitute these and (11)–(12) into (15) and match coefficients on $x_{i,1}$:

$$\beta = \Big((1-\lambda) + \lambda(1-\mu)\beta\Big)a_x + \lambda\mu\,\beta\,b_x \quad \Rightarrow \quad \beta = \frac{(1-\lambda)a_x}{1-\lambda(1-\mu)a_x - \lambda\mu\,b_x}. \tag{A.2}$$

Step 2: solve for (γ_d, γ_f)

Matching coefficients on z_1 and z_2 in (15) yields

$$\begin{aligned}\Big(1-\lambda(1-\mu)\Big)\gamma_d - \lambda\mu\,\gamma_f &= \Big((1-\lambda) + \lambda(1-\mu)\beta\Big)a_1 + \lambda\mu\,\beta\,b_1,\\ -\lambda\mu\,\gamma_d + \Big(1-\lambda(1-\mu)\Big)\gamma_f &= \Big((1-\lambda) + \lambda(1-\mu)\beta\Big)a_2 + \lambda\mu\,\beta\,b_2.\end{aligned}$$

Let:

$$\Delta \equiv \Big(1-\lambda(1-\mu)\Big)^2 - (\lambda\mu)^2 > 0. \tag{A.3}$$

Then:

$$\gamma_H = \frac{\Big(1-\lambda(1-\mu)\Big)\Big[\Big((1-\lambda) + \lambda(1-\mu)\beta\Big)a_1 + \lambda\mu\,\beta\,b_1\Big] + \lambda\mu\Big[\Big((1-\lambda) + \lambda(1-\mu)\beta\Big)a_2 + \lambda\mu\,\beta\,b_2\Big]}{\Delta}, \tag{A.4}$$

$$\gamma_F = \frac{\lambda\mu\Big[\Big((1-\lambda) + \lambda(1-\mu)\beta\Big)a_1 + \lambda\mu\,\beta\,b_1\Big] + \Big(1-\lambda(1-\mu)\Big)\Big[\Big((1-\lambda) + \lambda(1-\mu)\beta\Big)a_2 + \lambda\mu\,\beta\,b_2\Big]}{\Delta}. \tag{A.5}$$

Lemma A.1 (Closed form for $\gamma_H - \gamma_F$). Define

$$P \equiv a_1 - a_2, \qquad Q \equiv b_2 - a_2, \qquad c \equiv b_x P + a_x Q, \qquad d_0 \equiv 1 - \lambda(1-\mu)a_x. \tag{A.6}$$

Under Assumptions 1–4,

$$\gamma_H - \gamma_F = \frac{(1-\lambda)\Big(P - \lambda\mu\, c\Big)}{\Big(1-\lambda+2\lambda\mu\Big)\Big(d_0 - \lambda\mu\, b_x\Big)}. \tag{A.7}$$

Proof of Lemma A.1. Subtract the z_2-coefficient equation from the z_1-coefficient equation in the 2×2 system defining (γ_d, γ_f) (Step 2). This yields:

$$\Big(1 - \lambda(1-\mu) + \lambda\mu\Big)(\gamma_d - \gamma_f) = R_1 - R_2, \tag{A.8}$$

where:

$$R_1 = \Big((1-\lambda) + \lambda(1-\mu)\beta\Big)a_1 + \lambda\mu\,\beta\, b_1, \qquad R_2 = \Big((1-\lambda) + \lambda(1-\mu)\beta\Big)a_2 + \lambda\mu\,\beta\, b_2. \tag{A.9}$$

Using $b_1 = a_2$:

$$R_1 - R_2 = \Big((1-\lambda) + \lambda(1-\mu)\beta\Big)(a_1 - a_2) - \lambda\mu\,\beta\,(b_2 - a_2). \tag{A.10}$$

Substitute β from (A.2). Writing $P = a_1 - a_2$ and $Q = b_2 - a_2$ and using b_x and a_x from (13)–(14), algebra gives:

$$R_1 - R_2 = \frac{(1-\lambda)\Big(P - \lambda\mu(b_x P + a_x Q)\Big)}{1 - \lambda(1-\mu)a_x - \lambda\mu b_x}. \tag{A.11}$$

Finally, note that $d_0 - \lambda\mu b_x = 1 - \lambda(1-\mu)a_x - \lambda\mu b_x$ and $1 - \lambda(1-\mu) + \lambda\mu = 1 - \lambda + 2\lambda\mu$, yielding (A.7). ■

Proof of Proposition 2.1. By Lemma A.1, the denominator of (A.7) is positive under Assumption 2. Under Assumption 2, we have $P - \lambda\mu c > 0$, so (A.7) implies $\gamma_d - \gamma_f > 0$. Since $\partial k_{i,1}/\partial z_1 = \gamma_d$ and $\partial k_{i,1}/\partial z_2 = \gamma_f$, the result follows. ■

Proof of Proposition 2.2. From Lemma A.1,

$$\gamma_d - \gamma_f = \underbrace{\frac{1-\lambda}{1-\lambda+2\lambda\mu}}_{\text{strictly decreasing in } \mu} \cdot \underbrace{\frac{P - \lambda\mu\, c}{d_0 - \lambda\mu\, b_x}}_{\equiv f(\mu)}. \tag{A.12}$$

Since $P, c, d_0, b_x > 0$ are constant in μ, we have

$$f'(\mu) = \frac{\lambda\Big(b_x P - c d_0\Big)}{(d_0 - \lambda\mu\, b_x)^2}. \tag{A.13}$$

Using $c = b_x P + a_x Q$ and $Q > 0$,

$$b_x P - c d_0 = b_x P(1 - d_0) - d_0 a_x Q. \tag{A.14}$$

If $d_0 \geq \frac{1}{2}$, then $1 - d_0 \leq d_0$. Also, under Assumptions 2–3 and $0 \leq \rho_\theta < 1$ one has $b_x \leq a_x$ and $P \leq Q$ (since $b_2 > a_1$ and $b_1 = a_2$). Hence

$$b_x P(1 - d_0) \leq a_x P d_0 \leq a_x Q d_0, \tag{A.15}$$

so $b_x P - c d_0 < 0$ and therefore $f'(\mu) < 0$ on the admissible region where $d_0 - \lambda\mu\, b_x > 0$. Since both factors are positive and the first is strictly decreasing, the product is strictly decreasing in μ on $[0, \bar{\mu})$. ■

Derivative checks for posterior weights (used in Proposition 2.3). For $\rho_\theta \in (0, 1)$, the posterior coefficients in (13)–(14) satisfy

$$\frac{\partial a_1}{\partial \rho_\theta} < 0, \qquad \frac{\partial a_x}{\partial \rho_\theta} < 0, \qquad \frac{\partial a_2}{\partial \rho_\theta} > 0, \qquad \frac{\partial b_x}{\partial \rho_\theta} > 0, \qquad \frac{\partial b_2}{\partial \rho_\theta} < 0,$$

and at $\rho_\theta = 0$ these derivatives are weakly zero except $\partial a_2 / \partial \rho_\theta > 0$ and $\partial b_x / \partial \rho_\theta > 0$. (These signs follow by differentiating the closed forms (13)–(14) and using that $\kappa(\rho_\theta) = 1/(1 - \rho_\theta^2)$ increases in ρ_θ.)

Proof of Proposition 2.3. If $\mu = 0$, then $\lambda_f = 0$ and the equilibrium coefficients satisfy

$$\gamma_d = \frac{a_1}{1 - \lambda a_x}, \qquad \gamma_f = \frac{a_2}{1 - \lambda a_x}, \qquad \Rightarrow \qquad \gamma_d - \gamma_f = \frac{a_1 - a_2}{1 - \lambda a_x}. \tag{A.16}$$

For $\rho_\theta \in (0, 1)$, the derivative checks above imply a_1 decreases, a_2 increases, and a_x decreases in ρ_θ, hence $(a_1 - a_2)$ strictly decreases and $(1 - \lambda a_x)$ strictly increases in ρ_θ. Therefore the ratio is strictly decreasing on $(0, 1)$ and (by continuity) weakly decreasing on $[0, 1)$. The extension to sufficiently small $\mu > 0$ follows by continuity of $(\gamma_d - \gamma_f)$ in (μ, ρ_θ) on the interior of the stability region. ■