LUISS

Research Center
for European Analysis
and Policy

The Three Ws of Horizon Participation:
Whether, When and for Whom Does It Raise Productivity among Italian SMEs? Empirical Evidence from Matrix Completion

Marco Pini

Working Paper 11/2026

November 9, 2026

The Three Ws of Horizon Participation: Whether, When and for Whom Does It Raise Productivity among Italian SMEs? Empirical Evidence from Matrix Completion

Marco Pini*

Abstract

With research and innovation central to the EU's competitiveness agenda, assessing the economic returns to Horizon participation is a policy priority. This paper examines whether, when and for whom Horizon participation improves labour productivity among Italian SMEs. We link European Commission participation records with firm-level panel data (2015–2024), following 348 SMEs entering Horizon in 2020 and three alternative untreated donor pools of 6,960 firms each. Treatment effects are estimated using Matrix Completion, with robustness checks based on propensity-score matching and Difference-in-Differences. Average productivity gains range from 5.7% to 7.5% over 2021–2024, reaching 9.2–11.7% four years after programme entry (t+4, 2024). Effects are concentrated among firms with fewer than 50 employees (7.9–9.2%). Pre-treatment diagnostics identify a statistically significant widening of the treated–untreated productivity gap in 2019, consistent with anticipatory capability-building among future Horizon beneficiaries during the proposal-preparation phase. Combining effect magnitude, timing and firm-size heterogeneity across distinct donor pools and estimators, the study adds evidence on Horizon's economic additionality. Three policy priorities emerge: assess impacts over innovation cycles rather than grant durations; reduce SME participation barriers through information, application and consortium-building support; and foster entrepreneurial ambition, innovation culture and project readiness before the application stage.

I am grateful to Emanuele Pugliese for his valuable guidance, support and insightful suggestions. Special thanks are due to Paolo Pasimeni for making this research possible. I would also like to thank Anton Giulio Bottoni for his valuable support in database preparation and the constructive discussions on statistical analysis.

Keywords: Horizon programmes; Research and innovation policy; Labour productivity; Small and medium-sized enterprises (SMEs); Matrix Completion; Causal impact evaluation; Dynamic and heterogeneous treatment effects.

* Centro Studi delle Camere di Commercio Guglielmo Tagliacarne.

1. Introduction

Research and innovation (R&I) policy has returned to the centre of the European economic agenda as productivity growth, technological capabilities and innovation increasingly shape competitiveness and resilience (Draghi, 2024). This renewed emphasis also forms part of a broader revival of industrial policy, driven by technological change and strategic objectives (Aiginger & Rodrik, 2020; Juhász et al., 2024; Millot & Rawdanowicz, 2024). Within this framework, Horizon is the European Union's main instrument for supporting collaborative research and innovation, combining financial resources with access to knowledge, technological capabilities and international research networks. Its potential contribution is particularly relevant for SMEs, which often face greater financial, technological and organisational constraints in undertaking R&D and exploiting external knowledge. Assessing whether such support translates into stronger firm-level productivity is therefore central to evaluating its economic additionality and informing the design of European R&I policy. Previous research provides evidence that participation in European R&I programmes can affect firm performance and productivity (Barajas et al., 2012; Aguiar & Gagnepain, 2017; Mitra & Niakaros, 2023; Kátay et al., 2025). However, the available firm-level evidence remains relatively limited on the dynamic evolution of productivity effects around programme participation and their heterogeneity across firm size, while evidence focused specifically on Italian firms remains scarce. These questions are especially relevant in Italy, where strong participation in Horizon coexists with persistent productivity weakness and a productive structure dominated by SMEs.

This paper addresses this gap by organising the analysis around the "3 Ws" of policy evaluation – Whether, When and for Whom. Specifically, it asks: RQ1: To what extent does Horizon participation improve the labour productivity of Italian SMEs? RQ2: How does the productivity effect evolve over time following programme entry? RQ3: How do these effects differ between small and medium-sized firms?

The empirical analysis combines European Commission administrative data on Horizon participation with firm-level panel data for 2015–2024. Treatment effects are estimated using the Matrix Completion approach of Athey et al. (2021) across three largely distinct untreated donor pools, with additional robustness checks based on propensity-score matching and Difference-in-Differences estimation on the same matched sample.

The results show a positive and statistically significant average productivity effect of around 6–8%, remarkably stable across donor-pool constructions. Effects are strongest at the longest observed horizon, reaching around 10–12% in 2024, and are substantially larger among small firms than among medium-sized firms. Pre-treatment diagnostics also reveal an anticipation-type effect in 2019, suggesting that economically relevant capability-building processes may begin before formal programme entry, although this pre-award effect cannot be interpreted as a causal treatment effect.

These findings suggest that R&I policy evaluation should adopt sufficiently long horizons to capture delayed productivity gains, preserve effective access for SMEs, and strengthen firms' capabilities to engage in ambitious European R&I projects already before formal programme entry.

The remainder of the paper is organised as follows. Section 2 presents the institutional background, Section 3 reviews the literature and develops the research questions, Sections 4 and 5 describe the data and empirical strategy, Section 6 presents the results and robustness analyses, and Section 7 concludes the paper with policy implications and directions for further research.

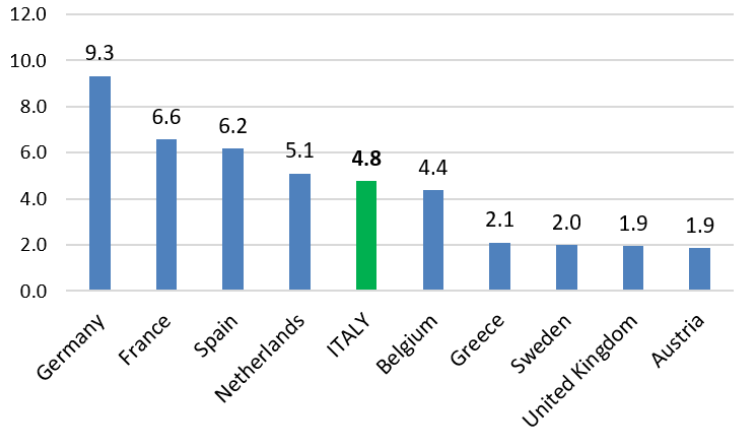
2. Institutional background

Horizon is the EU's main framework for supporting research and innovation, aimed at strengthening Europe's scientific and technological base, industrial competitiveness and the economic and societal impact of R&I (European Parliament & Council of the European Union, 2013, 2021). Beyond project finance, it operates as a supranational policy instrument that combines resources, capabilities and innovation actors across countries, generating European added value.

This role has become increasingly central to the EU competitiveness agenda. Horizon Europe explicitly links R&I investment to economic resilience, open strategic autonomy and technological leadership, including the capacity to develop and deploy critical technologies (European Commission, 2024). This orientation is reinforced by the current policy agenda, which identifies closing the innovation gap as essential to restoring productivity and growth (European Commission, 2025a).

Within this broader policy framework, Italy represents a particularly relevant case. According to the European Commission's Horizon Europe dashboard, as of July 2026, Italian participants had been allocated approximately €4.8 billion in net EU contributions (Figure 0), accounting for about 8% of the total and ranking fifth after Germany, France, Spain and the Netherlands (European Commission, 2026a).

Figure 0. Top ten participating countries by net EU contribution under Horizon Europe



Note: Cumulative contribution recorded in the Horizon Europe Dashboard as of July 2026. Net EU contribution refers to the EU contribution allocated to signed grants and does not necessarily correspond to payments already disbursed.
Source: European Commission (2026a), Horizon Europe Dashboard.

Horizon supports activities across the research and innovation process, from frontier research to technological development and market-oriented innovation. Under Horizon Europe, these activities are organised around three main pillars—Excellent Science; Global Challenges and European Industrial Competitiveness; and Innovative Europe—alongside partnerships, missions, research infrastructures and instruments managed by bodies such as the European Research Council and the European Innovation Council (European Commission, 2021).

Projects are selected through competitive calls and often involve cross-country collaboration among firms, universities, research organisations and public institutions. For firms, participation can therefore provide more than funding: it can reduce the cost and risk of R&D, facilitate access to expertise and infrastructures, and connect firms to international knowledge networks. Horizon participation should thus be viewed as exposure to a complex research and innovation policy intervention rather than as a standardised business subsidy.

3. Literature review

3.1 Institutions, innovation and the rationale for public intervention

Innovation depends not only on firms' internal choices but also on the institutional environment in which they operate. New Institutional Economics defines institutions as the formal and informal rules that structure economic and social interactions (North, 1990). By reducing uncertainty and transaction costs, including those related to information acquisition, institutions make innovative investment more predictable and improve the allocation of knowledge (Dobler, 2009; Greenwald & Stiglitz, 2013). This role becomes especially important during periods of rapid change or external shocks, when economic performance depends on the capacity of institutions to adapt existing rules and create new ones in response to emerging challenges (North, 2005).

Institutions can therefore strengthen firms' willingness to take risks and their ability to mobilise the knowledge, skills and complementary resources needed for innovation.

This institutional perspective is reinforced by the market-failure literature, which explains why private incentives alone may not generate an adequate level of innovative activity. Since knowledge is partly non-rival and difficult to appropriate, firms may be unable to capture the full returns generated by innovation and may therefore invest less than is socially desirable (Nelson, 1959; Arrow, 1962). Innovation can also generate learning and informational externalities. Firms experimenting with new technologies may reveal information about technical feasibility, production costs and market opportunities that benefits other actors but cannot be fully appropriated by the innovating firm. As a result, private investment may again fall below the socially desirable level, providing an additional rationale for public support (Juhász et al., 2024; OECD, 2024). Innovation projects are additionally characterised by high uncertainty, limited collateral and strong information asymmetries between firms and external investors, which can raise the cost of capital and restrict access to finance (Hall & Lerner, 2010).

The rationale for intervention also extends beyond knowledge and financing externalities. Innovation increasingly depends on complementary investments, specialised skills, research infrastructures and cooperation among multiple organisations. Coordination problems may therefore arise when firms cannot identify suitable partners or combine the resources required to develop and apply complex technologies. Aiginger and Rodrik (2020) describe these obstacles as network failures, which may prevent potentially productive collaborations from emerging even when they would generate collective benefits. Public action may also support activity-specific inputs that are essential to particular technologies or innovative activities but are unlikely to be supplied adequately through general horizontal policies alone (Juhász et al., 2024).

Public intervention in innovation should therefore not be understood only as a response to insufficient private investment. The innovation-systems literature views innovation as an interactive process involving firms, universities, research organisations, financial actors and public authorities (Lundvall, 1992; Nelson, 1993). From this perspective, institutions support the circulation of knowledge, the coordination of different actors and the combination of complementary capabilities. This broader interpretation is especially relevant to European research and innovation programmes, whose potential effects depend not only on funding but also on their capacity to connect firms with partners, expertise and technological resources across national borders. At the same time, these rationales do not imply that public intervention is always effective: its justification depends on the importance of the underlying failures, the suitability of the chosen instrument and the institutional capacity to limit misallocation and other unintended effects (Garcia-Macia & Sollaci, 2024; OECD, 2024).

3.2 How public support can strengthen firms' innovation capabilities

Public support can influence firms' innovation through several connected mechanisms. First, financial assistance increases liquidity and provides organisational slack, allowing firms to undertake innovation projects that would otherwise be considered too uncertain or financially demanding (Zona, 2012). Second, cost sharing reduces firms' exposure to technological and commercial risk and may enable projects of greater scale, complexity or ambition (Keizer & Halman, 2007; Roper et al., 2008; European Commission, 2025b). Third, public intervention can shape emerging markets and technological trajectories by supporting activities whose economic potential is not yet fully recognised by private actors (Mazzucato, 2016). Fourth, public programmes may provide access to external knowledge, research infrastructures and specialised expertise that firms cannot develop internally (OECD, 2010). Fifth, competitive funding may signal project or firm quality, helping recipients attract complementary private finance and mitigate information asymmetries (Feldman & Kelley, 2006; Meuleman & De Maeseneire, 2012).

Innovation policy therefore does more than transfer public resources: it can expand firms' opportunities and strengthen their ability to combine internal capabilities with external knowledge. However, these effects depend on programme design, the stage of the innovation process and the characteristics of participating firms (Becker, 2019).

3.3 Previous evidence on European research and innovation programmes

Building on the mechanisms discussed above, previous studies have examined whether European Framework Programmes translate institutional support into measurable firm-level outcomes. Horizon combines competitive funding with cross-border access to complementary knowledge, expertise and technological capabilities, thereby addressing financial and coordination barriers (Nepelski et al., 2019; Barajas et al., 2012). Participation should therefore be understood not simply as the receipt of financial resources, but as exposure to a broader policy intervention that can reduce project risk and strengthen firms' innovation capabilities.

Existing evidence reflects this complexity. Barajas et al. (2012) find that participation in research joint ventures supported by EU Framework Programmes increased Spanish firms' intangible assets and indirectly improved productivity. Aguiar and Gagnepain (2017), using variation in funding under the Fifth Framework Programme, identify substantial long-run labour-productivity effects for firms participating in industry-oriented collaborative projects, with stronger effects for projects closer to market application.

The institutional design of European programmes also matters. Hünermund and Czarnitzki (2019) find no average effect of grants from a pan-European programme on employment or sales under the allocation rule actually applied, but identify positive effects for high-quality projects and show that alternative budget-allocation rules could have generated stronger outcomes. This suggests that programme effectiveness depends not only on the availability of funding, but also on selection procedures, project quality and resource-allocation rules.

More recent evidence is broadly consistent with positive firm-level effects. The European Commission's counterfactual evaluation of Horizon 2020 finds increases of around 20% in employment and around 30% in assets and revenues among funded firms relative to comparable applicants, with effects persisting beyond the average duration of funded projects (Mitra & Niakaros, 2023). Similarly, Kátay et al. (2025) report substantial positive effects of FP7 grants on firm sales and labour productivity, with particularly large gains among smaller, initially less productive and financially constrained firms.

This evidence also points to an important temporal dimension. The European Commission's macroeconomic evaluation of Horizon Europe finds that the gains associated with R&I funding extend into the medium term and that business investment continues to generate economic benefits beyond the intervention period (European Commission, 2025c).

3.4 From innovation support to firm productivity

Horizon participation affects labour productivity through a transmission process rather than directly through funding. EU support can strengthen technological and organisational capabilities, innovation and intangible investment, thereby improving firm performance over time (Roper et al., 2008; European Commission, 2025b). Evidence from EU Framework Programmes also links participation to higher intangible assets and, indirectly, labour productivity (Barajas et al., 2012).

These effects are not automatic. Firms differ in their capacity to absorb and exploit external knowledge according to prior capabilities and human capital (Cohen & Levinthal, 1990). Policy effects may therefore vary across firms and contexts (Juhász et al., 2024), with stronger firms translating project results into productive applications more rapidly (Guisado-González et al., 2018; Vanino et al., 2019). Recent evidence confirms heterogeneity by initial productivity, financial constraints, technological intensity and project characteristics (Kátay et al., 2025).

Timing is also part of the transmission mechanism. Public support may initially raise R&D, investment, specialist employment and intangible assets before affecting productivity, although some effects may emerge earlier. Previous evidence shows that impacts differ across firms and over time and that increases in output, investment or employment do not always translate immediately into higher productivity (Millot & Rawdanowicz, 2024). Productivity gains may become stronger once technologies are integrated into production or commercialised, consistent with evidence that the effects of European research support can persist or become more visible beyond the funding period (Aguiar & Gagnepain, 2017; European Commission, 2025c).

The relevant questions are therefore not only whether Horizon participation improves labour productivity, but also when these effects emerge and how they vary across firms with different structural characteristics.

3.5 Research gap and research questions

The literature suggests that Horizon participation can strengthen firms' technological and organisational capabilities, but productivity effects may vary over time and across firms. Evidence remains limited on how productivity evolves after project entry and how effects differ across firm characteristics.

This is particularly relevant for Italy, where weak productivity growth is associated with limited innovation, small firm size and low R&D investment (European Commission, 2026b). Focusing on SMEs is therefore important given their central role in the Italian economy and the constraints they face in undertaking R&D and exploiting external knowledge. The paper therefore addresses the following research questions:

RQ1. To what extent does participation in Horizon-funded research and innovation projects improve the labour productivity of Italian SMEs?

RQ2. How does the productivity effect of Horizon participation evolve over time following programme entry?

RQ3. How do the productivity effects of Horizon participation differ between small and medium-sized firms?

Together, these questions move beyond a static average effect by examining both the timing and the heterogeneity of programme outcomes. This provides evidence relevant to the evaluation of European research and innovation policy and to the identification of the firms and conditions under which public support is more likely to translate into productivity gains.

4. Data and research design

4.1 Data sources and research design

The empirical analysis combines administrative information on participation in Horizon programmes with longitudinal firm-level data covering the period 2015–2024. Firm-level information on productivity and structural characteristics is drawn from the Italian Statistical Business Register and related business data provided by the Italian National Statistical Institute (ISTAT).

Treatment definition and timing

Horizon participation is identified using administrative data provided by the European Commission (DG Research and Innovation). The available data cover Italian organisations participating in Horizon 2020 and Horizon Europe projects whose grant agreements were signed between 2020 and 2024, which is the full period available in the Horizon dataset used for this study. The records identify the participating organisation, the funded project and the year in which the grant agreement was signed.

Because firms may participate in more than one Horizon project, project-level participation records are consolidated at the firm level. For each firm, programme entry is identified by the earliest observed Project Signature Year. Subsequent project participations do not generate a new treatment event.

The signature year provides a clear administrative marker of formal entry into Horizon, but it does not identify the exact timing or intensity of economic exposure within the calendar year. Under the EU grant framework, project implementation can begin after the grant agreement is signed, and funding is generally provided over the duration of the project through an initial pre-financing payment, followed, where applicable, by interim payments and a final balance. At the same time, the grant agreement defines activities that are implemented over a project-specific period and are subject to subsequent reporting and deliverables. The year of signature therefore identifies the beginning of formal participation, but not a common point within the year at which all firms become fully exposed to the programme.

This distinction is particularly important because the available data report only the year, rather than the exact date, of signature. Firms recorded as entering Horizon in 2020 may therefore have signed their grant agreements at very different points during that year, implying substantially different periods of potential exposure within the 2020 accounting year. Treating the entire 2020 accounting year as either fully untreated or fully treated would impose an exposure assumption that cannot be supported by the administrative information available. The empirical design consequently treats 2020 as a transition year: firms whose first observed Horizon grant agreement was signed in 2020 define the treated cohort, but 2020 itself is excluded from both the pre-treatment and post-treatment windows. Excluding 2020 is also appropriate because the COVID-19 shock and extraordinary employment-support measures may have generated atypical movements in labour productivity, potentially distorting the relationship between value added and employment.

The pre-treatment period is therefore 2015–2019, while treatment effects are evaluated over 2021–2024. Focusing on firms entering Horizon in 2020 also provides the longest available post-treatment window, allowing effects to be followed from 2021 to 2024.

Outcome and firm characteristics

The Horizon records are linked to longitudinal firm-level information on labour productivity and structural characteristics. Labour productivity is measured as value added per employee and is expressed in natural logarithms in the main analysis. The logarithmic transformation reduces the influence of the strongly skewed firm-level productivity distribution and facilitates the interpretation of treatment effects in percentage terms.

The dataset also contains information on firm size, sector of activity and geographical location, which are used to account for observable heterogeneity across firms. Firm size is measured before treatment, using employment in 2019, and the corresponding size classification is held fixed throughout the observation period. Consequently, changes in employment following Horizon participation do not affect the classification of firms into size groups or the estimates of treatment-effect heterogeneity by firm size. The resulting firm-year panel allows the productivity trajectories of Horizon beneficiaries and untreated firms to be observed before and after programme entry.

Focus on SMEs. The main empirical analysis focuses on small and medium-sized enterprises. This choice reflects both the structure of the Italian productive system and the central role assigned to SMEs in European research and innovation policy. Under Horizon Europe, the European Innovation Council explicitly supports start-ups and SMEs in developing and scaling high-risk and potentially market-creating innovations, recognising the particular financing and capability constraints that these firms may face (European Commission, 2026c). Focusing on SMEs therefore provides a policy-relevant setting for assessing whether participation in Horizon can translate research and innovation support into stronger firm-level productivity.

4.2 Analytical sample construction

The empirical analysis is based on firm-level panel data covering 2015–2024. The treatment cohort consists of SMEs (hereafter referred to simply as firms) whose first observed Horizon grant agreement was signed in 2020, which defines their entry into the programme. As discussed in Section 4.1, 2020 is treated as a transition year and is excluded from the estimation window. Treatment effects are therefore evaluated over 2021–2024, using 2015–2019 as the pre-treatment period.

Construction of the untreated donor pools

The initial pool of untreated firms is drawn from a population of approximately 1.5 million SMEs that did not receive any of the public incentives considered in the analysis. Restricting the comparison pool to firms not receiving other forms of public support helps isolate the effect associated with Horizon participation and reduces the risk that estimated productivity differences reflect the influence of other incentive schemes.

This population was identified from the universe of Italian firms by excluding firms recorded as beneficiaries of public support in the Italian National Register of State Aid (*Registro Nazionale degli Aiuti di Stato, RNA*).

To assess whether the estimated effects are sensitive to the construction of the comparison group, the analysis considers three alternative untreated donor pools based on different sampling strategies, each constructed using a 1:20 treated-to-untreated sampling ratio. The objective is to verify whether the estimated productivity

effects remain stable when the untreated comparison group is constructed according to different principles (Table 1).

- **Treated-proportional – DB2.** Untreated firms are sampled using a stratified design aimed at reproducing the distribution of Horizon-treated firms across eight strata defined by the interaction of sector (industry/services), firm size (small/medium-sized) and geographical area (Centre-North/South). However, the limited availability of eligible untreated firms in some strata prevents the final donor pool from exactly reproducing the treated firms' distribution. The resulting sample composition is reported in Appendix Table A1.
- **Random – DB4.** Untreated firms are randomly sampled from the eligible untreated-firm population.
- **Population-proportional – DB5.** Untreated firms are sampled using the same eight strata adopted for DB2—sector × firm size × geographical area—but reproducing the corresponding distribution observed in the overall SME population rather than among Horizon-treated firms.

For brevity, the three specifications are hereafter referred to as DB2, DB4 and DB5, respectively.

The definitions and coding of the sector, firm-size and geographical classifications used for donor-pool stratification are reported in Appendix Table A0.

The resulting donor pools are substantially different at the firm level. Pairwise overlap is only around 11–12%, approximately 1.5% of untreated firms are common to all three samples, and roughly 78% of each donor pool is unique to that specification. This limited overlap strengthens the usefulness of the three-database design as a robustness exercise, since similar treatment effects across specifications would be obtained from largely distinct sets of untreated firms.

Table 1. Main analytical specifications and sample size

Specification	Donor-pool construction	Treated SMEs	Untreated SMEs	Total SMEs
Treated-proportional - DB2	Proportional to treated firms	348	6,960	7,308
Random - DB4	Random extraction	348	6,960	7,308
Population-proportional - DB5	Proportional to the SME population	348	6,960	7,308

Notes: The SME universe refers to the combined population of SMEs that never received any of the public incentives considered in the analysis and SMEs participating in Horizon. All three specifications focus on SMEs. For each treated firm, untreated firms were sampled using a 1:20 treated-to-untreated ratio, yielding 6,960 untreated firms for 348 treated firms in each specification.

The composition of treated and untreated firms across the three baseline donor pools (DB2, DB4 and DB5) is reported in Appendix Table A1. The table presents distributions by firm size, geographical area (Centre-North and South) and OECD sector, together with average pre-treatment labour productivity.

Across all donor-pool constructions, Horizon beneficiaries differ systematically from untreated firms. Most notably, their average labour productivity over 2015–2019 is approximately €77,600 per employee, compared with about €57,700 in DB2, €49,600 in DB4 and €49,200 in DB5. This corresponds to a pre-treatment productivity advantage of around 34% in DB2 and 56–58% in DB4–DB5. Differences are also visible in firm size and sectoral composition, with treated firms more frequently represented among medium-sized enterprises, high- and medium-high-technology activities and knowledge-intensive services. These patterns confirm that Horizon participants are not directly comparable with untreated SMEs on observable pre-treatment characteristics and reinforce the need for an estimation strategy capable of accounting for both observed heterogeneity and differences in pre-treatment productivity trajectories.

5. Empirical strategy: Matrix Completion estimator

5.1 Matrix Completion estimator

Estimating the causal effect of Horizon participation requires addressing the non-random nature of programme entry, since beneficiaries may differ systematically from untreated firms in characteristics that also affect subsequent productivity (Hünermund & Czarnitzki, 2019; Vanino et al., 2019). To address this challenge, the analysis applies the Matrix Completion approach developed by Athey et al. (2021), using it to reconstruct the untreated post-treatment outcomes of Horizon beneficiaries from their pre-treatment trajectories, outcomes observed for untreated firms and structural covariates.

Unlike conventional Difference-in-Differences, Matrix Completion allows untreated outcomes to depend on a latent low-rank structure, providing a more flexible representation of unobserved heterogeneity and estimating the latent outcome matrix from its observed entries, thereby allowing also the use of unbalanced panel data.

Let Y_{it} denote the logarithm of labour productivity for firm *i* in year *t*. For untreated firm-year observations, Y_{it} corresponds to the observed outcome under no Horizon participation. For treated firms after programme entry, the untreated potential outcome $Y_{it}(0)$ is not observed and must therefore be estimated. Following Athey et al. (2021), the analysis treats these post-treatment counterfactual outcomes as missing elements of the firm-year outcome matrix and reconstructs them from the observed data.

The untreated outcome is represented as:

$$Y_{it}(0) = L_{it} + \alpha_i + \delta_t + X'_{it}\beta + \varepsilon_{it}$$

where α_i and δ_t capture firm and year fixed effects, respectively, X_{it} contains observed firm characteristics, β denotes the associated coefficient vector, L_{it} represents the latent low-rank component, and ε_{it} is an idiosyncratic error term. Firm fixed effects absorb persistent differences across enterprises, while year effects capture shocks common to all firms.

In the baseline Matrix Completion specifications (DB2, DB4 and DB5), the observed structural covariates include firm size (small: 0–49 employees, including micro firms; medium: 50–249 employees), five OECD sector groups (high- and medium-high-technology manufacturing; low- and medium-low-technology manufacturing; knowledge-intensive services; less knowledge-intensive services; and other activities), and four geographical macro-areas (North-West, North-East, Centre, and South), as detailed in Appendix Table A0. These structural covariates are included as categorical indicators interacted with year dummies for 2016–2019 and 2021–2024, with 2015 as the reference year and 2020 excluded. These interactions allow productivity dynamics to differ across firm-size classes, sectors and geographical macro-areas over time. In particular, macro-area-by-year interactions account for heterogeneous productivity trends across North-West, North-East, Centre and South. We adopt this level of geographical aggregation rather than more disaggregated region- or province-specific year interactions to balance model flexibility and parsimony. Although the baseline donor pools are relatively large, the matched sample comprises only 290 untreated firms. Finer geographical disaggregation would increase model dimensionality and may result in sparsely populated geographical cells, potentially reducing estimation precision.

The key element of the model is the latent component L_{it}. Matrix Completion relies on the idea that an incomplete matrix can be reconstructed by exploiting an underlying low-rank structure (Candès & Recht, 2009). Rather than imposing a predetermined number of factors, the estimator uses nuclear-norm regularisation to recover a parsimonious representation of the outcome matrix from its observed entries (Mazumder et al., 2010; Athey et al., 2021). This regularisation controls the complexity of the latent component without requiring the number of latent factors to be fixed in advance.

The estimator is obtained by minimising the squared prediction error over the observed untreated entries of the outcome matrix while penalising the nuclear norm of the latent component:

$$\left(\hat{L}, \hat{\alpha}, \hat{\delta}, \hat{\beta}\right) = \arg\min_{L,\alpha,\delta,\beta} \left\{ \frac{1}{|\mathcal{O}|} \sum_{(i,t)\in\mathcal{O}} (Y_{it} - L_{it} - \alpha_i - \delta_t - X_{it}'\beta)^2 + \lambda \parallel L \parallel_* \right\}$$

Where $\mathcal{O}$ denotes the set of observed untreated firm-year outcomes, $\parallel L \parallel_*$ is the nuclear norm of the latent matrix, and λ is the regularisation parameter controlling the complexity of the estimated latent structure. The notation $|\mathcal{O}|$ indicates the number of observations in $\mathcal{O}$, while hats denote estimated quantities.

The nuclear norm, defined as the sum of the singular values of the matrix, provides a convex way of favouring low-rank solutions in Matrix Completion problems (Candès & Recht, 2009; Mazumder et al., 2010). Higher values of λ impose stronger regularisation, whereas lower values allow the model to reproduce more complex patterns in the observed data.

The regularisation parameter λ is selected separately for each Matrix Completion specification by cross-validation, using the root mean squared error (RMSE) as the measure of predictive performance. The corresponding validation-error curves are reported in Appendix Figure A1.

Once the model is estimated, the counterfactual productivity of each treated firm in each post-treatment year is predicted as:

$$\hat{Y}_{it}(0) = \hat{L}_{it} + \hat{\alpha}_i + \hat{\delta}_t + X_{it}'\hat{\beta}$$

The firm-year treatment effect, denoted by $\hat{\tau}_{it}$, is then obtained as the difference between observed productivity and its estimated untreated counterfactual:

$$\hat{\tau}_{it} = Y_{it} - \hat{Y}_{it}(0)$$

This procedure allows the counterfactual productivity path of Horizon beneficiaries to reflect persistent firm heterogeneity, common time shocks, observable structural characteristics and latent productivity patterns recovered from the panel. The resulting treatment effects therefore measure the difference between observed post-entry productivity and the productivity that the same firms are predicted to have experienced in the absence of Horizon participation.

Identification assumptions. The causal interpretation of Matrix Completion estimates requires that untreated productivity be adequately represented by firm and year fixed effects, observed structural covariates and a

low-rank latent component, with sufficient observed untreated outcomes to recover this structure (Athey et al., 2021). Treatment selection may depend on the systematic component of productivity but must be independent of idiosyncratic productivity shocks conditional on that component. Causal identification also requires the absence of spillovers to untreated firms and that pre-treatment outcomes provide a valid basis for reconstructing untreated post-treatment productivity, with potential anticipatory adjustments requiring careful interpretation. In this application, the five pre-treatment years and the large donor pools provide information for reconstructing heterogeneous productivity trajectories. The stability of estimates across alternative donor pools further supports the robustness of the estimated productivity effects on differences in control-group composition. The 2019 pre-treatment shift, discussed in Section 6.5, may reflect pre-award capability-building processes, although its underlying causes cannot be established with the available data.

5.2 Treatment effects and inference

Treatment effects are summarised through the Average Treatment Effect on the Treated (ATT), computed both over the full post-treatment period and separately for each year from 2021 to 2024. Since labour productivity is expressed in logarithms, log-point estimates are reported in percentage terms using the standard exponential transformation.

Year-specific ATTs are calculated as the average difference between observed and estimated counterfactual log productivity among treated firms in each year. These are period-specific effects on productivity levels, rather than cumulative or year-on-year changes. The overall ATT for 2021–2024 is calculated as the average of these differences across all treated firm-year observations in the post-treatment period.

To assess whether the estimated ATT is unusually large relative to apparent effects that may arise in the absence of treatment, we adopt a placebo approach. Placebo-based inference has already been employed as a complementary inferential procedure in empirical applications of Matrix Completion (Alrababa'h et al., 2021), drawing on the well-established tradition of in-space placebo tests in the synthetic control literature (Abadie et al., 2010). The resulting placebo distribution provides an empirical benchmark for assessing how unusual the observed ATT is compared with the effects estimated by applying Matrix Completion to hypothetical treatment assignments.

Matrix Completion accounts for heterogeneity across firms by using pre-treatment productivity trajectories, observed covariates and a latent component to estimate counterfactual outcomes.

Nevertheless, because Horizon participation is not randomly assigned and beneficiaries differ systematically from untreated firms in size composition and pre-treatment productivity (Appendix Table A1), the placebo distribution may not adequately approximate the distribution of estimated effects that would arise for Horizon beneficiaries under the null of no treatment effect.

To address these differences in observable characteristics, we also estimate Matrix Completion on a propensity-score-matched sample (Section 6.4 and Table 5), which substantially improves covariate balance (Appendix Table A4). However, improved balance does not, by itself, guarantee the validity of placebo-based inference.

Accordingly, placebo p-values are interpreted as empirical measures of how unusual the observed ATT is relative to the placebo estimates, rather than as exact randomization p-values. Their inferential relevance depends on whether the estimated effects for actual beneficiaries would be comparable to those obtained for placebo-treated firms under the null hypothesis of no treatment effect (Abadie et al., 2010).

For the baseline Matrix Completion estimates reported in Table 2 (DB2, DB4 and DB5), placebo treatment is randomly assigned among originally untreated firms, preserving the observed number of treated firms and excluding actual Horizon participants. The Matrix Completion model is re-estimated for 1,000 placebo assignments using the same covariates, panel structure, treatment timing and regularisation parameter. Two-sided empirical p-values are then computed as the share of placebo ATTs at least as large in absolute value as the observed ATT.

The main analysis reports two-sided empirical p-values, calculated as:

$$p = \frac{1 + \Sigma_{b=1}^{B} I(\left|\widehat{ATT}_{placebo}^{b}\right| \geq \left|\widehat{ATT}_{obs}\right|)}{B + 1}$$

Here, B = 1,000 denotes the number of placebo assignments, b indexes individual assignments, $I(\cdot)$ is the indicator function, and the placebo and observed ATT terms denote the corresponding estimated treatment effects.

The same two-sided empirical placebo p-value calculation, based on 1,000 assignments and re-estimation of Matrix Completion for each assignment, is applied to the overall and year-specific estimates in Tables 2 and 3, as well as to the matched Matrix Completion estimate in Table 5.

Unlike the baseline procedure (Table 2), placebo treatment labels for the firm-size-specific estimates (Table 3) are randomly reassigned within each size group, preserving the observed numbers of treated firms (263 small and 85 medium-sized firms). This procedure also includes actual Horizon beneficiaries, since DB4 and DB5 contain fewer untreated medium-sized firms than treated ones. Similarly, for the matched Matrix Completion specification (see Section 6.4 and Table 5), 290 firms are randomly selected from the full matched sample of 580 firms (290 beneficiaries and 290 untreated firms), since restricting the assignments to untreated firms would allow only one possible placebo group. In both cases, including actual beneficiaries may introduce genuine treatment effects into the placebo distributions, warranting caution in interpreting the resulting p-values.

For the Matrix Completion estimates reported in Tables 2 and 3, and for the matched Matrix Completion estimate in Table 5, standard errors are calculated as the standard deviation of the estimated firm-level treatment effects divided by the square root of the number of treated firms, conditional on the fitted counterfactual. For the overall 2021–2024 ATT, treatment effects are first averaged within each firm across the observed post-treatment years. Approximate 95% confidence intervals are constructed on the log scale using the Student's t distribution with N–1 degrees of freedom, where N denotes the number of treated firms. ATT estimates and confidence interval bounds are then transformed into percentage terms, while percentage standard errors are obtained using the delta method.

These standard errors and confidence intervals are not derived from the 1,000 placebo assignments, which are used separately to compute empirical p-values. Consequently, the confidence intervals and placebo p-values are based on different inferential procedures and need not yield identical significance assessments. Moreover, the conditional confidence intervals do not account for the full uncertainty associated with estimating the Matrix Completion counterfactual.

6. Results

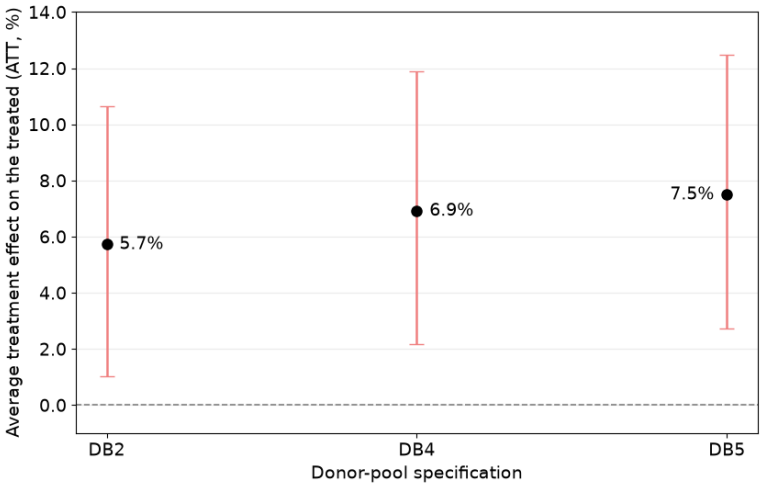
The empirical results are organised around the three research questions introduced in Section 3.5. Section 6.1 addresses RQ1 (*Whether*) by examining whether Horizon participation increases labour productivity and whether the average effect remains stable across the three Matrix Completion specifications described in Section 4.2. Section 6.2 addresses RQ2 (*When*) by analysing the evolution of treatment effects from 2021 to 2024, while Section 6.3 addresses RQ3 (*For Whom*) by examining heterogeneity between small and medium-sized firms. Section 6.4 assesses the robustness of the main results by comparing an alternative matched Difference-in-Differences (DiD) specification with a Matrix Completion estimate obtained on the same matched sample of treated and untreated firms identified through propensity score matching. Finally, Section 6.5 examines pre-treatment dynamics and identification through descriptive and TWFE-based diagnostics.

6.1 *Whether*: average productivity effects across alternative specifications

The first research question asks whether participation in Horizon-funded projects improves labour productivity. To assess the robustness of the estimated effect on the composition of the comparison group, the Matrix Completion analysis is implemented using three alternative untreated donor pools specifications: i) **DB2** selects untreated firms proportionally to the composition of treated firms; ii) **DB4** uses a randomly selected untreated donor pool; iii) while **DB5** selects untreated firms proportionally to the composition of the overall SME universe.

The Matrix Completion estimates provide a remarkably consistent answer across the three main specifications: the estimated average effect over 2021–2024 ranges from 5.7% in DB2 to 7.5% in DB5, with DB4 yielding an intermediate estimate of 6.9%. (Figure 1). Placebo-based inference provides support for the overall estimated effect, although statistical significance varies across donor pools. The overall ATT is significant at the 5% level in DB4 and DB5, while the DB2 estimate is significant at the 10% level. (Table 2; Appendix Figure A2).

Figure 1. Average treatment effect of Horizon participation on labour productivity across alternative Matrix Completion specifications



Notes: Donor-pool specifications: Treated-proportional – DB2; Random – DB4; Population-proportional – DB5.

The narrow range of estimates is notable given the substantial differences in donor-pool construction (see Table A1 in Appendix), panel structure and the final number of untreated firms available across specifications, especially considering that only 1.5% of untreated firms are common to all three samples. Despite these

differences, the estimated average treatment effects remain close in magnitude. This suggests that the estimated productivity gain is not driven by a specific donor-pool design.

The positive productivity effect is consistent with the literature on public support for research and innovation. Horizon participation can ease financial constraints, reduce R&D risk and improve access to external knowledge and specialised expertise. These channels may strengthen firms' technological and organisational capabilities and support higher value added per worker (Barajas et al., 2012; Vanino et al., 2019; Kátay et al., 2025).

Overall, these estimates provide a consistent answer to RQ1. Horizon participation is associated with higher labour productivity among beneficiary SMEs over 2021–2024, with both the magnitude and statistical significance of the estimated effect remaining relatively stable across alternative donor-pool constructions.

6.2 *When*: the evolution of treatment effects after programme entry

The second research question concerns the timing of the productivity effect. The most robust result concerns the longest observed post-treatment horizon. In all three donor-pool specifications, 2024 is the year with the largest estimated ATT. The effect ranges from 9.2% in DB2 to 11.7% in DB5, with an estimate of 11.0% in DB4 (Table 2, Figure 2). In each specification, the 2024 effect is substantially larger than the corresponding estimates for the preceding post-treatment years.

Table 2. Overall and annual treatment effects of Horizon participation on labour productivity

Period	ATT (%)	SE (%)	p-value	95% CI
Treated-proportional - DB2				
2021-2024	**5.7***	**2.440**	0.070	**[1.02, 10.63]**
2021	4.9	2.440	0.151	[0.21, 9.80]
2022	3.8	2.790	0.283	[-1.50, 9.47]
2023	5.0	2.770	0.182	[-0.34, 10.58]
2024	9.2**	3.360	0.021	[2.83, 16.04]
Random - DB4				
2021-2024	**6.9****	**2.470**	0.047	**[2.16, 11.89]**
2021	4.4	2.430	0.259	[-0.26, 9.30]
2022	5.0	2.810	0.189	[-0.40, 10.66]
2023	7.4*	2.840	0.070	[1.92, 13.12]
2024	11.0**	3.420	0.015	[4.48, 17.94]
Population-proportional - DB5				
2021-2024	**7.5****	**2.480**	0.029	**[2.71, 12.49]**
2021	7.5*	2.490	0.063	[2.76, 12.55]
2022	5.2	2.830	0.191	[-0.21, 10.93]
2023	5.6	2.810	0.141	[0.22, 11.27]
2024	11.7**	3.420	0.011	[5.20, 18.65]

Obs.: The post-treatment panel is balanced, with the same 348 treated and 6,960 untreated firms observed in each year from 2021 to 2024. Each annual ATT is based on 348 treated firm-year observations, for a total of 1,392 over the four post-treatment years. The outcome variable is the natural logarithm of labour productivity. Treatment effects and confidence interval bounds are reported as percentage changes using the exponential transformation. For details on Matrix Completion standard errors, confidence intervals and placebo p-values, see Section 5.2. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

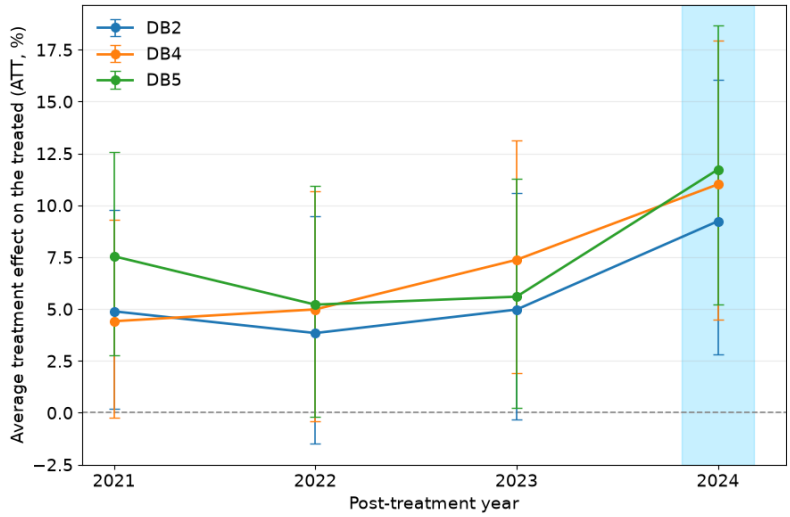
The year-by-year estimates nevertheless show some variation in the timing of this increase. In DB2, the estimated effect is 4.9% in 2021, declines to 3.8% in 2022, rises to 5.0% in 2023 and reaches 9.2% in 2024. DB4

displays a more progressive pattern, with estimated effects increasing from 4.4% in 2021 to 5.0% in 2022, 7.4% in 2023 and 11.0% in 2024. DB5 starts from a comparatively higher estimate of 7.5% in 2021, falls to 5.2% in 2022, remains broadly similar at 5.6% in 2023, and then rises markedly to 11.7% in 2024.

The results do not indicate a common monotonic year-by-year treatment path. This non-monotonic pattern is consistent with the positive overall ATT, which summarizes the average productivity effect over 2021–2024 rather than implying a continuous increase in the effect from one year to the next.

However, the 2024 estimates are the most robust across specifications: their 95% confidence intervals lie entirely above zero, and the placebo-based inference confirms statistical significance at the 5% level in all three cases, unlike the estimates for most earlier post-treatment years (Figure 2, Table 2).

Figure 2. Evolution of annual treatment effects of Horizon participation on labour productivity across alternative Matrix Completion specifications, 2021–2024



Notes: Donor-pool specifications: Treated-proportional – DB2; Random – DB4; Population-proportional – DB5.

Overall, the results provide a nuanced answer to RQ2. Across all Matrix Completion specifications, the largest estimated productivity gain is observed in 2024, four years after programme entry, consistent with a delayed strengthening of the productivity effect. At the same time, there is no clear evidence of a monotonic year-by-year increase in the post-treatment effect, and whether the effect beyond 2024 continues to increase, stabilises, or eventually attenuates remains an open question that cannot be assessed with the currently available data.

A delayed productivity response is also economically plausible, as Horizon-funded projects may initially affect R&D, knowledge acquisition and organisational change, with gains materialising only after these inputs are integrated into production and commercial activity. This interpretation is consistent with previous evidence that the effects of collaborative R&D support may take time to emerge (Aguiar & Gagnepain, 2017; European Commission, 2025c).

6.3 For *Whom*: heterogeneous productivity effects

The third research question asks whether the productivity effects of Horizon participation differ across firms of different sizes within the SME population. Across the three Matrix Completion specifications, the estimated effect is consistently and substantially larger for small firms than for medium-sized firms. For small firms, the average ATT over 2021–2024 ranges from 7.9% to 9.2% and is statistically significant at the 5% level in all three

specifications. By contrast, the corresponding estimates for medium-sized firms range from –0.8% to 2.3% and are not statistically significant (Table 3, Figure 3).

The year-specific estimates reinforce this pattern. For small firms, the productivity effect becomes particularly pronounced in 2024: the ATT rises to approximately 14–15% and is statistically significant at the 1% level across all three specifications, confirming at the firm-size level the longer-horizon pattern observed for the overall SME sample. By contrast, the annual effects for medium-sized firms remain substantially weaker and are not statistically significant in any post-treatment year. However, the relatively small number of treated medium-sized firms (85 firms) implies lower statistical precision and warrants caution in interpreting these estimates.

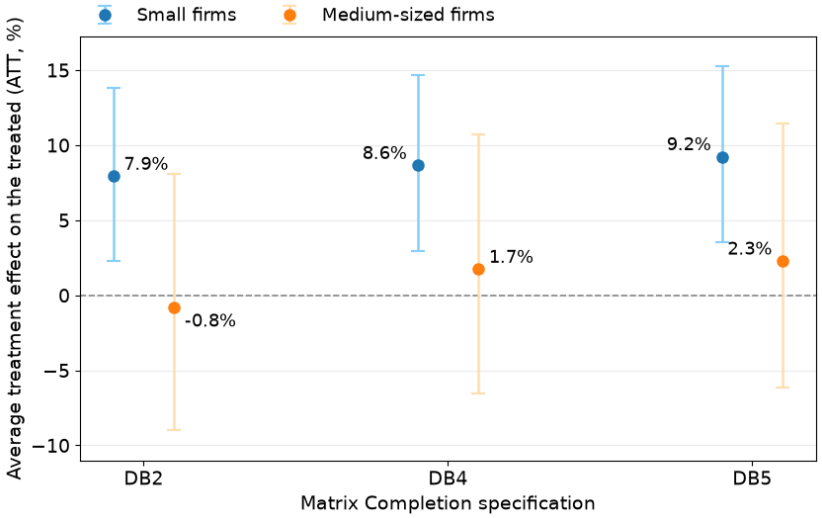
Overall, the results provide a consistent answer to RQ3. The productivity gains associated with Horizon participation are concentrated primarily among small firms rather than being uniformly distributed across the SME population.

Table 3. Overall and annual treatment effects of Horizon participation on labour productivity by firm size

	Small (<50 employees) firms				Medium (50-249 employees) firms			
Period	**ATT (%)**	**SE (%)**	**p-value**	**95% CI**	**ATT (%)**	**SE (%)**	**p-value**	**95% CI**
Treated-proportional - DB2								
2021-2024	**7.9****	**2.925**	0.026	**[2.32, 13.84]**	**-0.8**	**4.294**	0.876	**[-9.01, 8.08]**
2021	5.4	2.773	0.196	[0.07, 10.99]	3.4	5.112	0.623	[-6.32, 14.04]
2022	5.0	3.252	0.224	[-1.26, 11.55]	0.5	5.423	0.928	[-9.74, 11.87]
2023	7.6*	3.361	0.079	[1.17, 14.42]	-2.7	4.674	0.660	[-11.59, 7.03]
2024	14.0***	4.058	0.003	[6.28, 22.28]	-4.3	5.642	0.544	[-14.87, 7.62]
Random - DB4								
2021-2024	**8.6****	**2.963**	0.022	**[2.96, 14.64]**	**1.7**	**4.337**	0.822	**[-6.54, 10.73]**
2021	6.2	2.802	0.184	[0.80, 11.84]	-0.9	4.846	0.907	[-10.04, 9.26]
2022	5.9	3.271	0.188	[-0.32, 12.57]	2.1	5.508	0.821	[-8.27, 13.68]
2023	8.6*	3.416	0.052	[2.06, 15.52]	3.7	4.974	0.698	[-5.72, 14.09]
2024	14.1***	4.109	0.003	[6.28, 22.47]	2.0	5.959	0.846	[-9.18, 14.57]
Population-proportional - DB5								
2021-2024	**9.2****	**2.971**	**0.016**	**[3.53, 15.24]**	**2.3**	**4.408**	**0.760**	**[-6.13, 11.43]**
2021	7.4*	2.810	0.087	[1.99, 13.06]	8.0	5.345	0.330	[-2.09, 19.20]
2022	6.6	3.303	0.132	[0.27, 13.29]	1.1	5.482	0.897	[-9.24, 12.61]
2023	8.4*	3.414	0.059	[1.87, 15.33]	-2.6	4.651	0.733	[-11.43, 7.10]
2024	14.8***	4.101	0.003	[6.95, 23.12]	2.9	5.985	0.767	[-8.38, 15.47]

Obs.: 263 treated small firms and 85 treated medium-sized firms in each database. Untreated firms: DB2, 6,540 small and 420 medium-sized; DB4, 6,917 small and 43 medium-sized; DB5, 6,911 small and 49 medium-sized firms. The outcome variable is the natural logarithm of labour productivity. Treatment effects and confidence interval bounds are reported as percentage changes using the exponential transformation. For details on Matrix Completion standard errors, confidence intervals and placebo p-values, see Section 5.2.
*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

Figure 3. Average treatment effects of Horizon participation on labour productivity by firm size across alternative Matrix Completion specifications



Notes: Donor-pool specifications: Treated-proportional – DB2; Random – DB4; Population-proportional – DB5.

One possible explanation is that smaller firms face stronger financial, technological and organisational constraints in undertaking R&D and exploiting external knowledge (Meuleman & De Maeseneire, 2012; Vanino et al., 2019). Public funding and participation in collaborative research networks may therefore generate greater additionality when these constraints are more binding. This interpretation is consistent with recent evidence showing that the effects of EU research and innovation grants tend to be larger among smaller and financially constrained firms (Kátay et al., 2025).

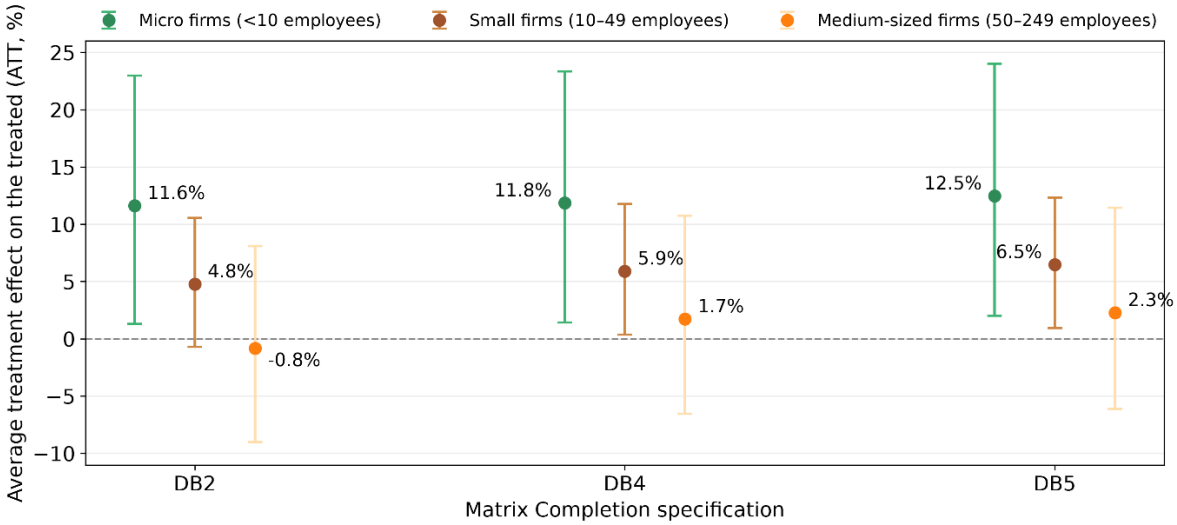
A more detailed analysis reveals additional heterogeneity within the group of firms with fewer than 50 employees, which in the baseline analysis was treated as a single "small-firm" category. When this broad group is decomposed into micro firms (fewer than 10 employees) and small firms in the narrower sense (10–49 employees), their relative composition differs markedly between treated and untreated firms. Among Horizon beneficiaries with fewer than 50 employees, micro firms account for about 47%, while small firms with 10–49 employees account for about 53%. By contrast, among untreated firms, micro firms represent more than 94% of the corresponding group across all three donor pools (Table 4). This compositional difference suggests that the aggregate effect estimated for firms below 50 employees may conceal substantial heterogeneity within the small-firm category.

Table 4. Distribution of micro and small firms among firms with fewer than 50 employees

Group	N Micro (<10 employees)	N Small (exc. micro) 10-49 employees	N firms <50 employees	Micro (<10 employees) (%)	Small (exc. micro) 10-49 employees (%)	Total (%)
Treated-proportional - DB2						
Treated	123	140	263	46.8	53.2	100.0
Untreated	6,163	377	6,540	94.2	5.8	100.0
Random - DB4						
Treated	123	140	263	46.8	53.2	100.0
Untreated	6,531	386	6,917	94.4	5.6	100.0
Population-proportional - DB5						
Treated	123	140	263	46.8	53.2	100.0
Untreated	6,544	367	6,911	94.7	5.3	100.0

We therefore further distinguish between micro firms (<10 employees) and small firms excluding micro firms (10–49 employees). The results reveal a clear size gradient. Over 2021–2024, the estimated ATT for micro firms ranges from 11.6% to 12.5% across the three Matrix Completion specifications, compared with 4.8% to 6.5% for firms with 10–49 employees (Figure 4). The effect for micro firms is statistically significant at 5% across all three specifications. The estimated effect for firms with 10–49 employees is smaller but remains statistically significant in all three donor pools, at 5% levels in DB4 and DB5 and at the 10% level in DB2.

Figure 4. Average treatment effects of Horizon participation on labour productivity for micro and small and medium firms across alternative Matrix Completion specifications



Notes: Donor-pool specifications: Treated-proportional – DB2; Random – DB4; Population-proportional – DB5.

Their strong estimated effect may therefore reflect both greater policy additionality and the greater underlying heterogeneity of this segment. Micro firms may differ more markedly from untreated firms along both observed and unobserved dimensions, making comparability more challenging than for larger SMEs. Further analysis using richer pre-treatment firm characteristics would help distinguish more clearly between these mechanisms. The following robustness analysis addresses this concern by constructing a more balanced comparison group through propensity score matching on pre-treatment firm characteristics and then re-

estimating the treatment effect using both Matrix Completion and Difference-in-Differences on the same matched sample.

6.4 Robustness across donor-pool composition and estimation methods: Matrix Completion and Difference-in-Differences after propensity score matching

The baseline Matrix Completion results are already robust across three largely distinct untreated donor pools (Figure 1, Table 2), with positive and statistically significant effects in all specifications. However, Section 6.3 highlights substantial differences between treated and untreated firms, particularly in firm-size composition besides pre-treatment productivity (Table A1). We therefore conduct an additional robustness exercise using a more closely balanced comparison sample.

For this purpose, we focus on DB4, whose untreated donor pool is originally constructed through random sampling from the eligible population of untreated SMEs. Starting from this random donor pool, treated and untreated firms are matched using one-to-one nearest-neighbour propensity score matching without replacement. The propensity score is estimated through a probit model including four sets of pre-treatment firm characteristics: firm size, distinguishing micro (0–9 employees), small (10–49 employees) and medium-sized firms (50–249 employees); OECD sector, aggregating high- and medium-high-technology manufacturing, low- and medium-low-technology manufacturing, knowledge-intensive services (KIS), less knowledge-intensive services (LKIS), and other activities into five groups; geographical macro-area, distinguishing North-West, North-East, Centre and South; and mean log labour productivity over 2015–2019. Detailed definitions and measurement periods of these covariates are reported in Appendix Table A0. The explicit separation of micro firms from other small firms reflects the substantial heterogeneity documented in Section 6.3 (Table 4). Matching is performed on the log-odds of the estimated propensity scores using a caliper equal to 0.2 times their standard deviation.

The matching procedure substantially improves the covariate balance. The initial DB4 sample contains 348 treated and 6,960 untreated firms, while the matched sample, hereafter referred to as "DB Matched", retains 290 treated firms and identifies 290 untreated firms with the closest propensity scores. As reported in Appendix Table A4, the standardized mean differences (SMD) decline sharply after matching. The maximum absolute SMD falls from 1.544 before matching to 0.102 after matching, while the mean absolute SMD decreases from 0.501 to 0.060. With the exception of one sectoral category that remains marginally above the conventional 0.10 threshold,[1] post-matching imbalances are small across the observed covariates. Particularly large improvements are observed for firm size and pre-treatment productivity, with the SMD for micro-sized firms decreasing in absolute value from 1.544 to 0.021 and that for mean log labour productivity from 1.011 to 0.057.

To enhance comparability across estimation methods, Matrix Completion is first re-estimated on the same propensity-score-matched sample of 290 treated and 290 untreated firms (hereafter referred to as "Matched Matrix Completion"). This specification uses the same structural covariates as the propensity-score model—firm size (micro, small and medium-sized firms), OECD sector and geographical macro-area, as detailed in Appendix Table A0. As in the baseline Matrix Completion specifications, these covariates are interacted with year indicators. Unlike the propensity-score model, however, it does not include mean pre-treatment log labour productivity as a separate covariate, since pre-treatment productivity trajectories enter directly through the observed outcome matrix.

[1] The Standardized Mean Difference is defined as the difference in means between treated and untreated groups divided by the pooled standard deviation. In absolute terms, values below 0.10 are commonly interpreted as indicating negligible residual imbalance.

Matched Matrix Completion yields an average ATT of 7.6% over 2021–2024, compared with 6.9% in the baseline DB4 specification based on 348 treated and 6,960 untreated firms.

A two-way fixed-effects Difference-in-Differences (DiD) model is then estimated on the same propensity-score-matched sample of 290 treated and 290 untreated firms, using the following specification:

$$Y_{it} = \alpha_i + \delta_t + \beta(\text{Treated}_i \times \text{Post}_t) + \varepsilon_{it}$$

where Y_{it} denotes log labour productivity, α_i and δ_t are firm and year fixed effects, respectively, Treated_i identifies Horizon beneficiaries entering the programme in 2020, and Post_t equals one for 2021–2024 and zero for 2015–2019, with 2020 excluded. The coefficient β represents the estimated Difference-in-Differences treatment effect in log points.

No additional covariates are included in the DiD regression, as pre-treatment firm characteristics are used in the preceding matching procedure. Standard errors are clustered at the firm level. The estimated coefficient and confidence interval bounds are transformed into percentage terms using $100[exp(\beta) - 1]$, while percentage standard errors are obtained using the delta method.

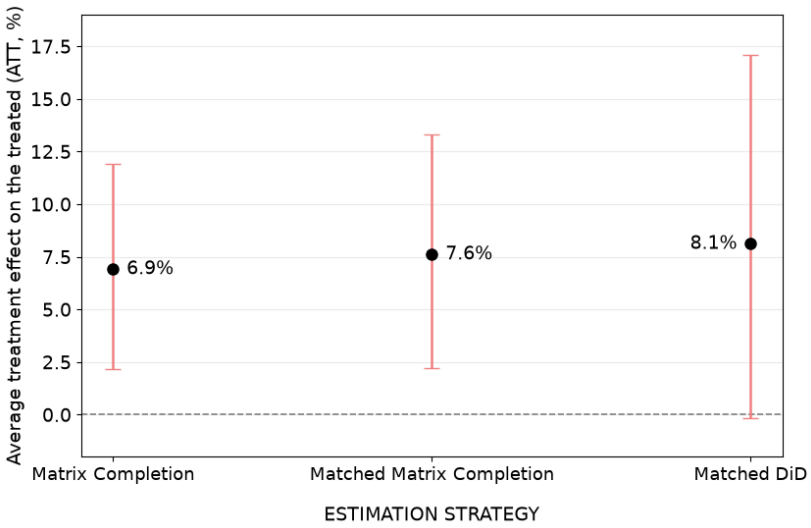
The causal interpretation of the DiD estimate relies on the parallel-trends assumption, requiring that treated and matched untreated firms would have followed similar productivity trends in the absence of Horizon participation, as well as the absence of treatment-related anticipation and spillover effects. The pre-treatment diagnostics in Section 6.5 show no statistically significant differential trends over 2015–2018 (joint test p = 0.505), followed by a distinct shift in 2019 (0.115 log points, p = 0.007). This pattern may reflect pre-award adjustments, although it limits the evidence supporting parallel trends over the full pre-treatment period.

Table 5. Overall treatment effects on labour productivity, 2021–2024: Matrix Completion and matched Difference-in-Differences after propensity score matching

Estimation strategy	N. treated	N. untreated	ATT (%)	SE (%)	p-value	95% CI
			Random - DB4			
Matrix Completion	348	6,960	6.9**	2.470	0.047	[2.16, 11.89]
Matched Matrix Completion	290	290	7.6*	2.816	0.060	[2.21, 13.30]
Matched DiD	290	290	8.1*	4.393	0.055	[-0.17, 17.08]

Notes: a) Matrix completion estimation strategy refers to DB4 donor pool specification; b) The outcome variable is the natural logarithm of labour productivity; c) Treatment effects and confidence interval bounds are reported as percentage changes using the exponential transformation. For the matched DiD specification, standard errors are clustered at the firm level. For details on Matrix Completion standard errors, confidence intervals and placebo p-values, see Section 5.2.

Figure 5. Robustness of the estimated productivity effect across donor-pool composition and estimation methods



Notes: a) Points report average treatment effects over 2021–2024; vertical bars indicate 95% confidence intervals. Matched Matrix Completion and Matched DiD are estimated on the same 1-NN propensity-score-matched DB4 sample (290 treated and 290 untreated firms); b) Matrix completion estimation strategy refers to DB4 donor pool specification.

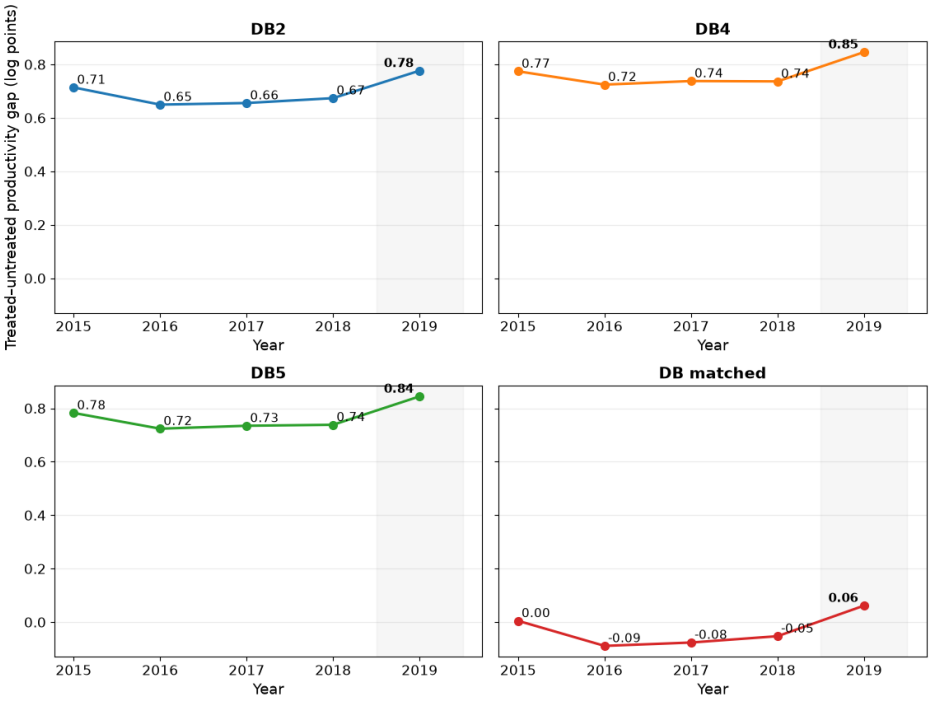
In summary, two main findings emerge. First, the positive productivity effect remains when the comparison group is substantially rebalanced through propensity score matching. Second, the increase from the baseline to the matched estimates appears to be driven mainly by the use of a more comparable control group, while the choice between Matrix Completion and DiD has only a limited effect on the point estimate.

6.5 Pre-treatment dynamics and identification diagnostics

Although Matrix Completion does not rely on the conventional parallel-trends assumption underlying Difference-in-Differences estimators, pre-treatment dynamics provide an important diagnostic of the credibility of the counterfactual design. They are particularly informative here because the robustness analysis also applies a TWFE Difference-in-Differences estimator to the propensity-score-matched sample. We therefore examine both the evolution of the treated–untreated productivity gap and formal TWFE pre-treatment tests.

Figure 6 shows a consistent pattern across specifications. In the three baseline donor pools (DB2, DB4, DB5) the productivity gap remains broadly stable through 2018 and widens in 2019. The same pattern persists after propensity score matching (DB matched): although matching largely removes the level difference between treated and untreated firms, a clear 2019 increase remains. This suggests that the 2019 shift is not simply driven by differences in donor-pool composition.

Figure 6. Pre-treatment productivity gap between Horizon-treated and untreated firms across alternative donor-pool specifications



Notes: a) The figure reports the annual difference in mean log labour productivity between future Horizon participants and untreated firms; b) Donor-pool specifications: DB2 = treated-proportional donor pool; DB4 = random donor pool; DB5 = population-proportional donor pool; DB Matched = propensity-score-matched sample constructed from DB4.

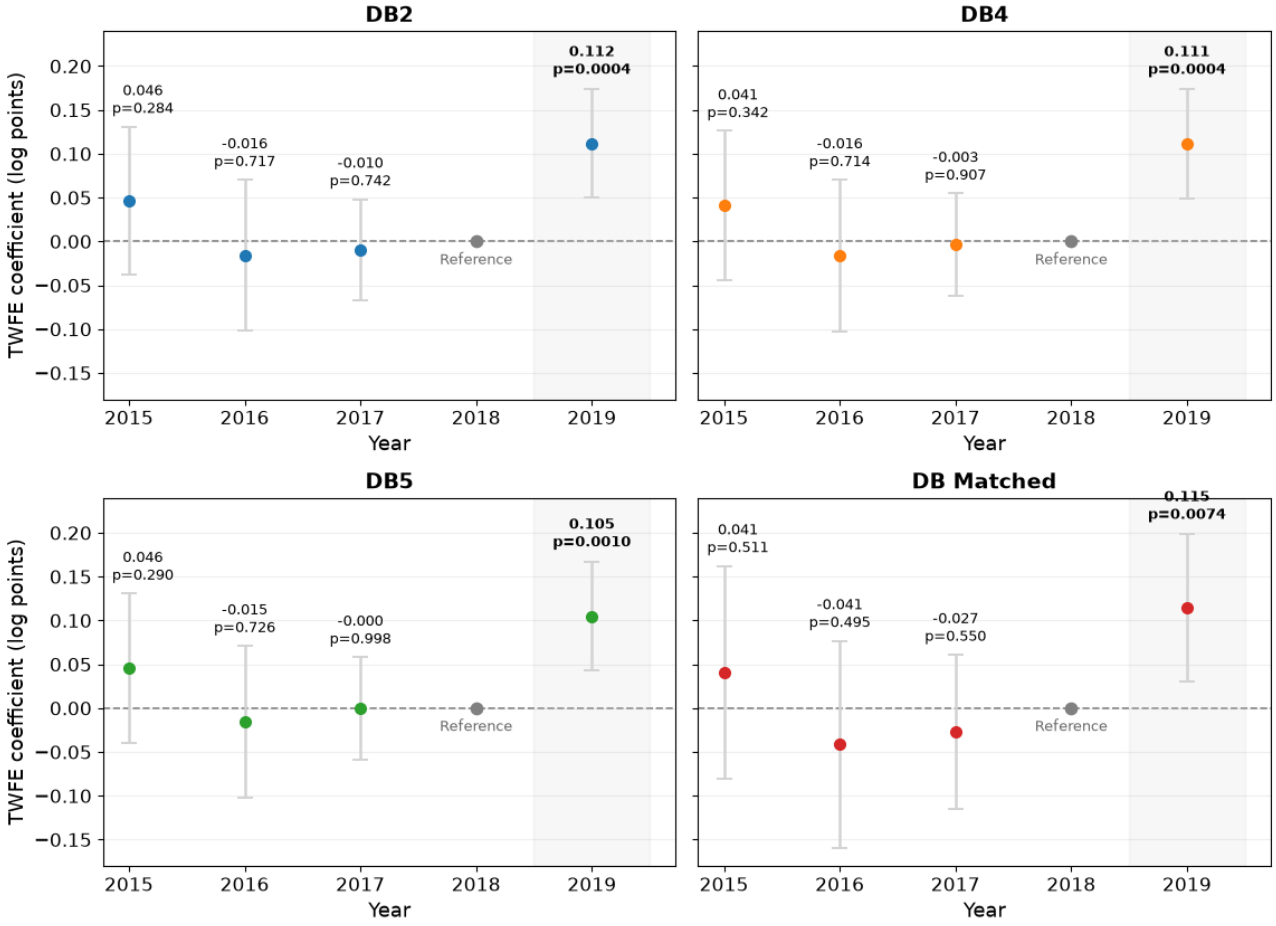
The formal tests confirm this pattern. With 2018 as the reference year, the joint null of no differential pre-treatment dynamics is not rejected in any specification, whereas using 2019 as the reference year it is rejected across all four samples (Table 6). Consistently, the 2019 coefficient relative to 2018 is positive and significant in every specification, including the matched sample (Figure 7). Full estimates are reported in Tables A2–A3.

Table 6. Comparison of TWFE joint pre-treatment tests using alternative reference years

Database	2015-2018 vs 2019 p-value	2015-2018 vs 2019 null	2015-2017 vs 2018 p-value	2015-2017 vs 2018 null
Treated-proportional - DB2	0.004	Rejected	0.501	Not rejected
Random - DB4	0.004	Rejected	0.618	Not rejected
Population-proportional - DB5	0.010	Rejected	0.563	Not rejected
DB Matched	0.050	Rejected	0.505	Not rejected

Notes: DB Matched = propensity-score-matched sample constructed from DB4.

Figure 7. TWFE pre-treatment coefficients relative to 2018 and the 2019 productivity-gap shift



Notes: Donor-pool specifications: DB2 = treated-proportional donor pool; DB4 = random donor pool; DB5 = population-proportional donor pool; DB Matched = propensity-score-matched sample constructed from DB4.

Overall, the evidence indicates that the rejection obtained when benchmarking against 2019 does not appear to reflect a persistent differential pre-trend, but rather a distinct widening of the productivity gap between future Horizon beneficiaries and untreated firms immediately before programme entry (i.e., from 2018 to 2019), equivalent to higher relative productivity growth among future beneficiaries over that interval. Since the underlying cause of this pre-treatment shift cannot be established, it warrants caution in the causal interpretation of post-treatment estimates.

Several mechanisms may plausibly account for the 2019 pre-treatment shift, although they cannot be separately identified with the available data. First, firms preparing competitive Horizon proposals may strengthen specialised human capital before formal programme entry by mobilising researchers, engineers, R&D managers or external technical expertise. Such resources may contribute not only to proposal preparation but also to technological problem solving and the firm's ability to absorb external knowledge, consistent with the absorptive-capacity literature and the OECD evidence on the productivity role of skills and organisational capital (Cohen & Levinthal, 1990; Cammeraat et al., 2021). Second, firms may strengthen R&D and innovation capabilities during project preparation. While the productivity returns to R&D generally require time to materialise (Ugur et al., 2016), activities such as technology scouting, testing and prototyping, specialised software adoption or improved R&D management may generate earlier efficiency gains. Third, the preparation of collaborative projects may strengthen research and innovation networks, providing access to universities, research organisations, technological partners and complementary knowledge even before formal

implementation. Fourth, proposal preparation may improve organisational capabilities, including project management, resource allocation, internal coordination and innovation-management practices.

7. Conclusion

Research and innovation policy is central to the European competitiveness agenda, with productivity growth increasingly linked to innovation (Draghi, 2024; European Commission, 2025a). This paper evaluates the causal effect of Horizon participation on labour productivity among Italian SMEs, focusing on three questions: *whether* participation raises productivity, *when* effects emerge, and *for whom* they are stronger.

Using Horizon administrative data linked to firm-level panel data for 2015–2024, the analysis applies Matrix Completion across three donor pools, complemented by matched Difference-in-Differences and matched Matrix Completion robustness checks.

Four main findings emerge. First, Horizon participation raises labour productivity by about 6–8% on average over 2021–2024. Second, effects are strongest in 2024, reaching about 10–12%, although no common monotonic pattern appears in earlier post-treatment years. Third, an anticipation-type effect emerges in 2019, suggesting that relevant adjustments may begin before formal programme entry, although this cannot be interpreted as a causal treatment effect. Fourth, effects are stronger for small firms, with average gains of about 8–9%, while estimates for medium-sized firms are smaller and not statistically significant.

Overall, the results are consistent with a two-stage process: a pre-award phase of capability building, followed by post-award R&D, innovation and intangible investment whose productivity returns may take time to materialise.

A – Policy recommendations

Firstly, a longer evaluation horizon is needed for R&I policies. Since the strongest productivity effects emerge only several years after programme entry, evaluations focused on the project period or the first post-award years may underestimate Horizon's economic impact. Monitoring frameworks should therefore combine implementation indicators with firm-performance outcomes measured over a longer horizon, following the innovation cycle rather than the administrative duration of the grant.

Secondly, effective SME access should remain a central priority of European R&I policy. The stronger effects observed among small firms suggest that SMEs may derive particularly high returns from Horizon participation. Policy should therefore focus not only on the availability of funding, but also on reducing barriers to participation through information, application support and consortium-building assistance.

Thirdly, policy should foster firms' innovation ambition and readiness to engage in European R&I programmes. The anticipation effect observed before programme entry suggests that firms may already undertake economically relevant capability-building processes during the preparation phase by mobilising specialised skills, strengthening R&D activities and developing external collaborations. Although these pre-award dynamics cannot be interpreted as causal effects of Horizon participation, they highlight the potential value of policy interventions that act before the application stage. Besides raising awareness of funding opportunities, support should therefore aim to strengthen entrepreneurial ambition, innovation culture and firms' capacity to conceive and pursue ambitious European R&I projects.

B – Directions for future research and limitations

First, the relatively short post-treatment window does not allow us to establish whether the productivity effect observed at t+4 persists, increases or fades over time. Longer firm-level panels are therefore needed to assess the durability of Horizon's productivity effects beyond the current evaluation horizon.

Second, the analysis treats Horizon participation as a binary treatment and does not identify how treatment intensity, project characteristics or intermediate outcomes shape the productivity effect. Future research should therefore examine whether impacts vary with funding size, project duration, consortium structure and firms' role in the project, while also tracing the channels through which support affects productivity, including R&D investment, intangible assets, specialised employment and innovation outputs.

Third, future research should further strengthen identification. Specifically, richer pre-treatment covariates should be incorporated to better account for observable differences, such as in innovation capacity, R&D intensity, human capital, financial conditions, management practices, ownership and governance structure, and prior participation in public support programmes.

Fourth, future research should exploit unsuccessful Horizon applicants as an alternative control group. Applicants represent a particularly informative counterfactual because they are eligible firms that have revealed a comparable willingness and capacity to engage in European R&I programmes, thereby sharing many observable and unobservable characteristics with funded firms. This comparison would provide a demanding robustness check of the estimated treatment effects and would be especially valuable for assessing the anticipation-type effect, since both funded and unsuccessful applicants are exposed to the application and proposal-preparation process prior to the funding decision.

Fifth, future research should decompose labour-productivity gains into output and employment effects. Distinguishing whether higher productivity reflects stronger output growth or lower employment is important for assessing whether Horizon supports productive expansion or labour-saving adjustment.

Sixth, an important extension concerns the additionality of European funding. Future research should assess whether Horizon crowds in additional private R&D and complement national innovation incentives or instead substitutes investment that firms or governments would have undertaken anyway. Distinguishing complementarity from crowding-out is essential for evaluating the true added value of EU intervention and the interaction between European, national and private innovation finance.

Finally, extending the same causal framework to other EU Member States would provide an important external-validity test. Starting with countries receiving substantial Horizon funding would make it possible to assess whether the productivity effects observed for Italian SMEs are common across European firms or depend on national institutional and productive structures.

References

Abadie, A., Diamond, A., & Hainmueller, J. (2010). Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program. Journal of the American Statistical Association, 105(490), 493–505. https://doi.org/10.1198/jasa.2009.ap08746

Aguiar, L., & Gagnepain, P. (2017). European cooperative R&D and firm performance: Evidence based on funding differences in key actions. *International Journal of Industrial Organization, 53*, 1–31. https://doi.org/10.1016/j.ijindorg.2016.12.007

Aiginger, K., & Rodrik, D. (2020). Rebirth of industrial policy and an agenda for the twenty-first century. *Journal of Industry, Competition and Trade, 20*(2), 189–207. https://doi.org/10.1007/s10842-019-00322-3

Alrababa'h, A., Marble, W., Mousa, S., & Siegel, A. A. (2021). Can exposure to celebrities reduce prejudice? The effect of Mohamed Salah on Islamophobic behaviors and attitudes. American Political Science Review, 115(4), 1111–1128. https://doi.org/10.1017/S0003055421000423

Arrow, K. J. (1962). Economic welfare and the allocation of resources for invention. In National Bureau of Economic Research (Ed.), *The rate and direction of inventive activity: Economic and social factors* (pp. 609–626). Princeton University Press.

Athey, S., Bayati, M., Doudchenko, N., Imbens, G. W., & Khosravi, K. (2021). Matrix completion methods for causal panel data models. *Journal of the American Statistical Association, 116*(536), 1716–1730.

Barajas, A., Huergo, E., & Moreno, L. (2012). Measuring the economic impact of research joint ventures supported by the EU Framework Programme. *The Journal of Technology Transfer, 37*(6), 917–942. https://doi.org/10.1007/s10961-011-9222-y

Becker, B. (2019). The impact of innovation policy on firm innovation and performance: A review of recent research developments. *ifo DICE Report, 17*(4), 10–15.

Cammeraat, E., Samek, L., & Squicciarini, M. (2021). Management, skills and productivity. *OECD Science, Technology and Industry Policy Papers, No. 101*. OECD Publishing. https://doi.org/10.1787/007f399e-en

Candès, E. J., & Recht, B. (2009). Exact matrix completion via convex optimization. *Foundations of Computational Mathematics, 9*, 717–772. https://doi.org/10.1007/s10208-009-9045-5

Cohen, W. M., & Levinthal, D. A. (1990). Absorptive capacity: A new perspective on learning and innovation. *Administrative Science Quarterly, 35*(1), 128–152. https://doi.org/10.2307/2393553

Dobler, C. (2009). *The impact of formal and informal institutions on per capita income* (Hohenheimer Diskussionsbeiträge No. 319). University of Hohenheim. https://doi.org/10.2139/ssrn.1636715

Draghi, M. (2024). *The future of European competitiveness: A competitiveness strategy for Europe*. European Commission.

European Commission. (2021). *Horizon Europe: The EU research and innovation programme 2021–2027—General overview*. Directorate-General for Research and Innovation.

European Commission. (2024). *Horizon Europe strategic plan 2025–2027*. Directorate-General for Research and Innovation.

European Commission. (2025a). *A competitiveness compass for the EU* (COM(2025) 30 final).

European Commission. (2025b). *Interim evaluation of the Horizon Europe Framework Programme for Research and Innovation (2021–2024): Horizon Europe—Research and innovation at the heart of competitiveness* (SWD(2025) 110 final).

European Commission. (2025c). *The RHOMOLO and FIDELIO interim evaluation of the impact of Horizon Europe* (Research and Innovation Paper Series, Working Paper 2025/06). Publications Office of the European Union.

European Commission. (2026a). *Horizon Europe dashboard: Net EU contribution by country or territory*. Data extracted on 28 July 2026.

European Commission. (2026b). *2026 country report – Italy* (Commission Staff Working Document SWD(2026) 212 final). European Commission.

European Commission. (2026c). *European Innovation Council work programme 2026*. European Innovation Council and SMEs Executive Agency.

European Parliament, & Council of the European Union. (2013). Regulation (EU) No 1291/2013 of 11 December 2013 establishing Horizon 2020—The Framework Programme for Research and Innovation (2014–2020). *Official Journal of the European Union, L 347*, 104–173.

European Parliament, & Council of the European Union. (2021). Regulation (EU) 2021/695 of 28 April 2021 establishing Horizon Europe—The Framework Programme for Research and Innovation. *Official Journal of the European Union, L 170*, 1–68.

Feldman, M. P., & Kelley, M. R. (2006). The ex ante assessment of knowledge spillovers: Government R&D policy, economic incentives and private firm behavior. *Research Policy, 35*(10), 1509–1521. https://doi.org/10.1016/j.respol.2006.09.019

Garcia-Macia, D., & Sollaci, A. (2024). *Industrial policies for innovation: A cost-benefit framework* (IMF Working Paper No. 2024/176). International Monetary Fund. https://doi.org/10.5089/9798400288401.001

Greenwald, B. C., & Stiglitz, J. E. (2013). Industrial policies, the creation of a learning society, and economic development. In J. E. Stiglitz & J. Y. Lin (Eds.), *The industrial policy revolution I: The role of government beyond ideology* (pp. 43–71). Palgrave Macmillan. https://doi.org/10.1057/9781137335173_4

Guisado-González, M., González-Blanco, J., Coca-Pérez, J. L., & Guisado-Tato, M. (2018). Assessing the relationship between R&D subsidy, R&D cooperation and absorptive capacity: An investigation on the manufacturing Spanish case. *The Journal of Technology Transfer, 43*(6), 1647–1666. https://doi.org/10.1007/s10961-017-9579-7

Hall, B. H., & Lerner, J. (2010). The financing of R&D and innovation. In B. H. Hall & N. Rosenberg (Eds.), *Handbook of the economics of innovation* (Vol. 1, pp. 609–639). Elsevier. https://doi.org/10.1016/S0169-7218(10)01014-2

Hünermund, P., & Czarnitzki, D. (2019). Estimating the causal effect of R&D subsidies in a pan-European program. *Research Policy, 48*(1), 115–124. https://doi.org/10.1016/j.respol.2018.08.001

Juhász, R., Lane, N. J., & Rodrik, D. (2024). The new economics of industrial policy. *Annual Review of Economics, 16*, 213–242. https://doi.org/10.1146/annurev-economics-081023-024638

Kátay, G., Mosberger, P., & Tucci, F. (2025). *The impact of EU grants for research and innovation on firms' performance* (European Economy Discussion Paper No. 238). European Commission, Directorate-General for Economic and Financial Affairs. https://doi.org/10.2765/0105659

Keizer, J. A., & Halman, J. I. M. (2007). Diagnosing risk in radical innovation projects. *Research-Technology Management, 50*(5), 30–37. https://doi.org/10.1080/08956308.2007.11657459

Lundvall, B.-Å. (Ed.). (1992). *National systems of innovation: Towards a theory of innovation and interactive learning*. Pinter Publishers.

Mazumder, R., Hastie, T., & Tibshirani, R. (2010). Spectral regularization algorithms for learning large incomplete matrices. *Journal of Machine Learning Research, 11*, 2287–2322.

Mazzucato, M. (2016). From market fixing to market-creating: A new framework for innovation policy. *Industry and Innovation, 23*(2), 140–156. https://doi.org/10.1080/13662716.2016.1146124

Meuleman, M., & De Maeseneire, W. (2012). Do R&D subsidies affect SMEs' access to external financing? *Research Policy, 41*(3), 580–591. https://doi.org/10.1016/j.respol.2012.01.001

Millot, V., & Rawdanowicz, Ł. (2024). *The return of industrial policies: Policy considerations in the current context* (OECD Economic Policy Papers No. 34). OECD Publishing. https://doi.org/10.1787/051ce36d-en

Mitra, A., & Niakaros, K. (2023). *The Horizon effect: A counterfactual analysis of EU research & innovation grants* (Research and Innovation Paper Series No. 2023/08). Publications Office of the European Union.

Nelson, R. R. (1959). The simple economics of basic scientific research. *Journal of Political Economy, 67*(3), 297–306. https://doi.org/10.1086/258177

Nelson, R. R. (Ed.). (1993). *National innovation systems: A comparative analysis*. Oxford University Press.

Nepelski, D., Van Roy, V., & Pesole, A. (2019). The organisational and geographic diversity and innovation potential of EU-funded research networks. *The Journal of Technology Transfer, 44*(2), 359–380. https://doi.org/10.1007/s10961-018-9692-2

North, D. C. (1990). *Institutions, institutional change and economic performance*. Cambridge University Press. https://doi.org/10.1017/CBO9780511808678

North, D. C. (2005). *Understanding the process of economic change*. Princeton University Press.

OECD. (2010). *The OECD innovation strategy: Getting a head start on tomorrow*. OECD Publishing. https://doi.org/10.1787/9789264083479-en

OECD. (2021). *Understanding firm growth: Helping SMEs scale up*. OECD Studies on SMEs and Entrepreneurship. OECD Publishing.

OECD. (2024). *How do governments direct support for innovation? Lessons from recent OECD measurement and impact analysis (MABIS) work* (OECD Policy Briefs No. 4). OECD Publishing. https://doi.org/10.1787/c1d93d1c-en

Roper, S., Du, J., & Love, J. H. (2008). Modelling the innovation value chain. *Research Policy, 37*(6–7), 961–977. https://doi.org/10.1016/j.respol.2008.04.005

Ugur, M., Trushin, E., Solomon, E., & Guidi, F. (2016). R&D and productivity in OECD firms and industries: A hierarchical meta-regression analysis. *Research Policy, 45*(10), 2069–2086. https://doi.org/10.1016/j.respol.2016.08.001

Vanino, E., Roper, S., & Becker, B. (2019). Knowledge to money: Assessing the business performance effects of publicly funded R&D grants. *Research Policy, 48*(7), 1714–1737. https://doi.org/10.1016/j.respol.2019.04.001

Zona, F. (2012). Corporate investing as a response to economic downturn: Prospect theory, the behavioural agency model and the role of financial slack. *British Journal of Management, 23*(S1), S42–S57.

Appendix

Table A0. Variable definitions, coding and measurement in the empirical analysis

Variable	Definition and coding	Reference period	Empirical application
Labour productivity	Value added per employee, expressed in natural logarithms	Annual, 2015–2024; 2020 excluded from estimation	Outcome variable
Horizon participation	Binary indicator identifying firms whose first observed Horizon grant agreement was signed in 2020	2020 (program entry); 2021–2024 (post-treatment period)	Treatment definition
Firm size – baseline	Small firms: 0–49 employees (including micro); medium-sized firms: 50–249 employees	2019	Baseline Matrix Completion; size-heterogeneity analysis
Firm size – matched	Micro: 0–9 employees; small: 10–49; medium-sized: 50–249	2019	PSM and Matched Matrix Completion
Sector – donor-pool stratification	Binary classification distinguishing Industry (NACE Rev. 2 divisions 10–43, covering manufacturing, utilities and construction) and Services (NACE Rev. 2 divisions 45 onwards, starting with wholesale and retail trade).	2019	Stratification of untreated firms in DB2 and DB5
OECD sector	Five groups: high- and medium-high-technology manufacturing; low- and medium-low-technology manufacturing; knowledge-intensive services; less knowledge-intensive services; other activities	2019	Matrix Completion (Baseline and Matched) and PSM
Geographical area – stratification	Centre-North; South (included Islands)	2019	Construction of DB2 and DB5 donor pools
Geographical macro-area	North-West; North-East; Centre; South (included Islands)	2019	Matrix Completion (Baseline and Matched) and PSM
Pre-treatment labour productivity	Mean natural logarithm of value added per employee	2015–2019	PSM

Notes: a) Centre-North comprises North-West (Piedmont, Aosta Valley, Lombardy, Liguria), North-East (Trentino-Alto Adige, Veneto, Friuli-Venezia Giulia, Emilia-Romagna) and Centre (Tuscany, Umbria, Marche and Lazio). South comprises Abruzzo, Molise, Campania, Apulia, Basilicata, Calabria, Sicily and Sardinia. The geographical classifications cover all 20 Italian regions. b) the OECD-based sectoral classification (OECD, 2021, Annex C, Methodological notes) is constructed using two-digit NACE Rev. 2 codes and aggregated into five groups: (i) High- and medium-high-technology manufacturing: 20–21, 26–30; (ii) Low- and medium-low-technology manufacturing: 10–19, 22–25, 31–33; (iii) Knowledge-intensive services (KIS): 50–51, 58–66, 69–75, 78, 80, 84–93; (iv) Less knowledge-intensive services (LKIS): 45–47, 49, 52, 55–56, 68, 77, 79, 81–82, 95; and (v) Other activities: residual category comprising activities not included in the preceding groups.

Table A1. Composition of treated and untreated firms used in the Matrix Completion specifications

	Treated-proportional - DB2		Random - DB4		Population-proportional - DB5	
	Treated	**Untreated**	**Treated**	**Untreated**	**Treated**	**Untreated**
N. firms	348	6,960	348	6,960	348	6,960
Firm size (%)						
Small	75.6	94.0	75.6	99.4	75.6	99.3
Medium	24.4	6.0	24.4	0.6	24.4	0.7
Macro-area (%)						
Centre-North	85.3	87.5	85.3	81.9	85.3	81.5
South	14.7	12.5	14.7	18.1	14.7	18.5
OECD sector (%)						
High & medium-high tech manufacturing	14.7	2.3	14.7	1.3	14.7	1.4
Knowledge-intensive services (KIS)	63.2	37.9	63.2	37.9	63.2	38.6
Less knowledge-intensive services (LKIS)	9.5	38.5	9.5	41.8	9.5	40.3
Low & medium-low tech manufacturing	8.9	11.0	8.9	9.0	8.9	9.2
Others	3.7	10.3	3.7	10.0	3.7	10.6
Labour productivity						
Mean 2015-2019	77,571	57,735	77,571	49,573	77,571	49,241
Mean log, 2015-2019	10.982	10.285	10.982	10.214	10.982	10.213

Notes: Entries report percentage distributions within treated and untreated firms in each analytical specification. Firm size, OECD sector and macro-area correspond to the structural covariates included in the Matrix Completion model. Mean log labour productivity refers to the mean natural logarithm of value added per employee over 2015–2019.

Table A2. TWFE pre-treatment coefficients relative to 2019

Year	Coefficient	SE	p-value	95% CI
Treated-proportional - DB2				
2015	-0.061	0.042	0.149	[-0.1437, 0.0219]
2016	-0.123***	0.047	0.008	[-0.2147, -0.0321]
2017	-0.123***	0.038	0.001	[-0.1963, -0.0492]
2018	-0.109***	0.029	0.000	[-0.1669, -0.0516]
Random - DB4				
2015	-0.066	0.042	0.122	[-0.1486, 0.0175]
2016	-0.122***	0.047	0.009	[-0.2129, -0.0301]
2017	-0.112***	0.038	0.003	[-0.1861, -0.0383]
2018	-0.111***	0.030	0.000	[-0.1687, -0.0530]
Population-proportional - DB5				
2015	-0.052	0.042	0.219	[-0.1352, 0.0310]
2016	-0.111**	0.047	0.017	[-0.2025, -0.0198]
2017	-0.106***	0.038	0.005	[-0.1798, -0.0321]
2018	-0.103***	0.030	0.001	[-0.1612, -0.0454]
DB Matched				
2015	-0.055	0.061	0.366	[-0.1744, 0.0643]
2016	-0.139**	0.064	0.030	[-0.2650, -0.0138]
2017	-0.134**	0.054	0.012	[-0.2393, -0.0292]
2018	-0.107***	0.040	0.007	[-0.1849, -0.0288]

*** p < 0.01; ** p < 0.05; * p < 0.1. Coefficients and standard errors are expressed in natural-log points.

Table A3. TWFE pre-treatment coefficients relative to 2018

Year	Coefficient	SE	p-value	95% CI
Treated-proportional - DB2				
2015	0.047	0.043	0.273	[-0.0370, 0.1320]
2016	-0.014	0.046	0.765	[-0.1040, 0.0760]
2017	-0.014	0.036	0.694	[-0.0850, 0.0570]
2019	0.109***	0.029	0.000	[0.0520, 0.1670]
Random - DB4				
2015	0.043	0.043	0.322	[-0.0420, 0.1280]
2016	-0.013	0.046	0.774	[-0.1030, 0.0770]
2017	-0.003	0.037	0.941	[-0.0750, 0.0690]
2019	0.111***	0.03	0.000	[0.0530, 0.1690]
Population-proportional - DB5				
2015	0.051	0.043	0.236	[-0.0330, 0.1350]
2016	-0.008	0.046	0.866	[-0.0980, 0.0820]
2017	-0.003	0.037	0.944	[-0.0750, 0.0690]
2019	0.103***	0.03	0.001	[0.0450, 0.1610]
DB Matched				
2015	0.041	0.062	0.511	[-0.0808, 0.1623]
2016	-0.041	0.060	0.495	[-0.1590, 0.0769]
2017	-0.027	0.045	0.550	[-0.1145, 0.0610]
2019	0.115***	0.043	0.007	[0.0308, 0.1982]

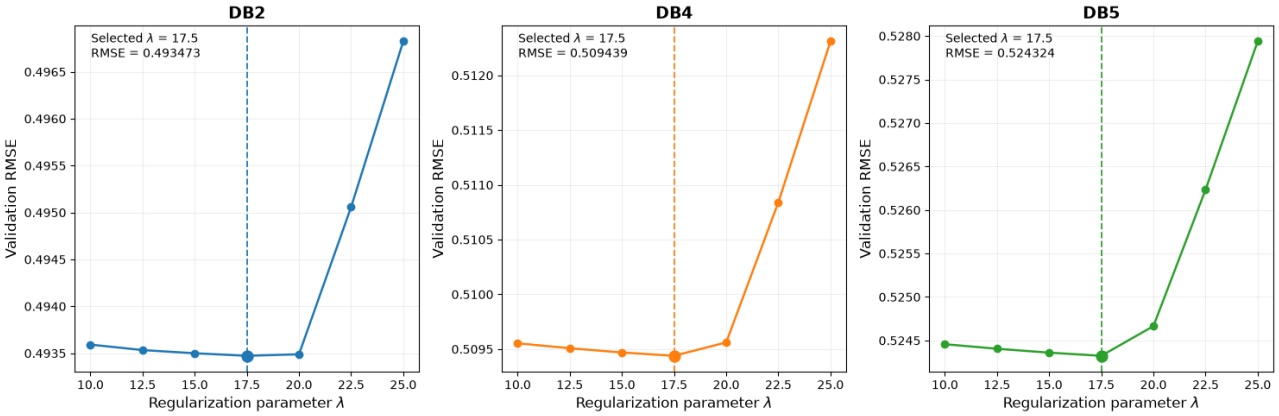
Notes: a) DB Matched = propensity-score-matched sample constructed from DB4; b) Coefficients and standard errors are expressed in natural-log points. *** p < 0.01; ** p < 0.05; * p < 0.1.

Table A4. Covariate balance before and after propensity score matching: Standardized Mean Differences

	Before	After
N treated firms	348	290
N untreated firms	6,960	290
Firm size		
Micro	-1.544	0.021
Small	0.905	-0.077
Medium	0.770	0.077
Macro-area		
North-West	0.036	-0.007
North-East	0.060	0.078
Center	-0.017	-0.098
South	-0.093	0.048
OECD sector		
High & medium-high tech manufacturing	0.509	0.079
Low & medium-low tech manufacturing	-0.002	0.087
Knowledge-intensive services (KIS)	0.524	-0.007
Less knowledge-intensive services (LKIS)	-0.797	-0.102
Others	-0.251	-0.049
Labour productivity		
Mean log, 2015-2019	1.011	-0.057

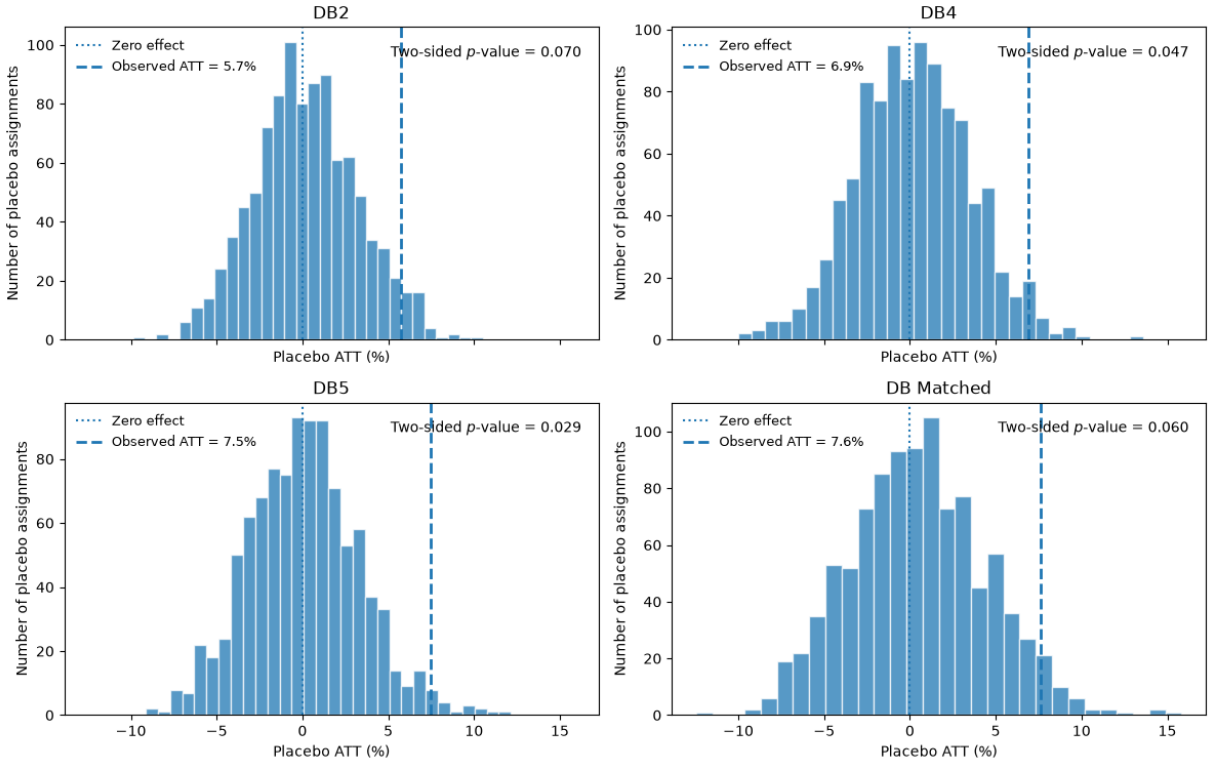
Note: The Standardized Mean Difference is defined as the difference in means between treated and untreated groups divided by the pooled standard deviation. In absolute terms, values below 0.10 are commonly interpreted as indicating negligible residual imbalance.

Figure A1. Cross-validation curves for the regularization parameter λ across Matrix Completion specifications



Notes: a) The figure reports cross-validation root mean squared errors (RMSE) across candidate values of the regularization parameter λ. The dashed vertical line indicates the λ selected for each Matrix Completion specification; b) Donor-pool specifications: Treated-proportional – DB2; Random – DB4; Population-proportional – DB5.

Figure A2. Placebo distributions of the overall ATT across Matrix Completion specifications



Notes: a) The histogram reports the distribution of overall ATT estimates obtained from placebo treatment assignments. The solid vertical line indicates the observed ATT, while the dashed line denotes zero treatment effect. The reported p-value is the two-sided empirical placebo p-value; b) Donor-pool specifications: DB2 = treated-proportional donor pool; DB4 = random donor pool; DB5 = population-proportional donor pool; DB Matched = propensity-score-matched sample constructed from DB4.