The Tyranny of Corporate Control: Subsidies and Ownership Structure

Silvan Hofer*[1], Philipp-Leo Mengel[1], and Giorgio Presidente[1]

[1] *Institute for European Policymaking, Bocconi University*

April 28, 2026

Abstract

This paper provides causal estimates of the long-run impact of European R&D grants to small and medium-sized enterprises and shows that the effects are mediated by recipients' ownership structure. We exploit a quality label awarded by the European Commission to high-scoring R&D grant proposals that could not be funded due to budget constraints and find that subsidies have a positive and significant long-run impact on productivity when the recipient has no corporate shareholder or is owned by individuals. In contrast, the effect is at best muted for controlled subsidiaries owned by other corporations. The mediating role of the ownership structure is not driven by firm size, and the results are robust across a battery of alternative performance measures.

*All authors contributed equally to this work.

1 Introduction

Over the past decade, the fiscal cost of research and development (R&D) subsidies has approached €1 trillion in the European Union. During the same period, the United States—where publicly funded R&D represents a similar share of GDP—spent at least as much.

The existing causal evidence suggests that R&D subsidies can be effective in boosting recipients' performance (Bloom et al., 2019; Howell, 2017; Howell et al., 2021; Santoleri et al., 2022), especially among young and small companies. In light of its immediate policy application, teasing out size-related heterogeneity is a major contribution of this literature.

The aim of this paper is extending the empirical evidence to another policy-relevant and targetable dimension of heterogeneity: ownership structure. Specifically, we ask whether the impact of R&D grants vary depending on whether the recipient is a controlled subsidiary or an independent company.

Ownership structure may affect the additionality of a subsidy through several channels. Unlike independent firms, subsidiaries of corporate groups can rely on internal capital markets and may therefore face weaker financing constraints (Gertner et al., 1994; Shin and Stulz, 1998). Firms embedded in corporate networks might be in a better position to "game" the system, securing support without meaningfully changing their operations (Criscuolo et al., 2019). More broadly, controlled subsidiaries and independent companies differ markedly in shareholder composition and governance (Demsetz and Lehn, 1985; Bertrand and Mullainathan, 2003; Shleifer and Vishny, 1986), which can have an impact on the allocation of subsidy-funded resources.

To study the issue empirically, we build a novel dataset linking the universe of R&D grants disbursed by the European Commission since 2014, to rich firm-level data on recipient companies, including detailed information on ownership structure and shareholder characteristics.

We focus on a specific program that supports small and medium-sized enterprises (SMEs), for which we have a list of companies awarded the *Seal of Excellence*—a quality label granted to high-ranking project proposals that ultimately did not receive funding due to budget constraints.[1] This institutional feature provides a credible comparison group and enables us to estimate the causal effects of receiving a grant.

Crucially, for each beneficiary firm we observe ownership and shareholders' characteristics. This allows us to assess how the causal impact of the grant on firm performance varies with ownership structure.

Our baseline metric of performance is Total Factor Productivity (TFP), which we measure using revenue-based data. In our context, the advantage of revenue-based TFP is that it compactly captures efficiency, product quality, variety and price/markup effects more broadly. R&D grants can plausibly affect all these dimensions through innovation-related activities, which are otherwise difficult to measure independently.

Based on this setup, we study the causal impact of R&D grants on productivity.

We make three main contributions. The main contribution is providing causal evidence that impact of R&D grants on performance is mediated by ownership structure: grants have a positive and significant long-run impact on productivity for independent companies, but not for controlled subsidiaries. The result is robust across a battery of alternative TFP and performance measures, such as patents and revenue growth.

The second contribution is extending the causal evidence on R&D grants to the European context, as most existing papers focus on the United States (Bloom et al., 2019; Howell, 2017; Howell et al., 2021). One exception is Santoleri et al. (2022), which studies the same SME support program examined in this paper. However, we have access to longer time series. Consequently, the third contribution of this paper is to provide long-term causal estimates of the grant.

The rest of the paper is organized as follows: Section 2 presents the institutional setting;

[1] https://research-and-innovation.ec.europa.eu/funding/funding-opportunities/seal-excellence_en

Section 3 and 4 describe the data and ownership structure respectively; Section 5 presents the empirical approach; Section 6 presents the results, and Section 7 concludes.

2 Institutional Framework

The Framework Programmes are the European Union's multi-year funding cycles for research and innovation. Beginning with Horizon 2020 (2014–2020), the EU introduced dedicated instruments to support SMEs more directly.[2]

The main programs supporting SMEs are the *SME Instrument* (2014–2020) and the *EIC Accelerator* (2021–present), both providing EU support for innovation-driven SMEs.[3]

There are close similarities between the two programs. In both cases, the objective is developing a product, process, or service that can be readily commercialized.

The programs are endowed with a similar budget of approximately €3-4 billion over seven years. SME Instrument was rolled out by the Executive Agency for SMEs (EASME), which evolved into the European Innovation Council and SMEs Executive Agency (EISMEA) now administrating the EIC Accelerator.

Under both SME Instrument and EIC Accelerator, firms can access grants up to €2.5 million. Eligible R&D activities covered a broad set of late-stage innovation tasks, including prototyping, testing, design, performance evaluation, demonstration, piloting, validation for market replication, scale-up, and application development.

One notable difference is that previously under the SME Instrument, applicants could receive limited grant support for preparatory activities of €50,000 (Phase 1), while full proposals could be awarded grants of up to €2.5 million to fund activities such as demonstration, piloting, validation, and scale-up (Phase 2). Phase 1 grants have been discontin-

[2] Before Horizon 2020, SME support was fragmented and largely embedded within other programs rather than a consolidated accelerator-like scheme.

[3] Another program, *Fast Track to Innovation* (2015–2020) provided grants to consortia of at least three entities—often including SMEs—supporting close-to-market technologies through collaborative development rather than single-firm acceleration. Because funding is awarded to multi-beneficiary consortia, it is difficult to isolate the causal impact of the grant on any single firm, so we exclude this program from our analysis.

ued under the EIC Accelerator, which only provides grants after a successful full proposal and interview stage.

The program duration is typically between six and a maximum of twenty-four months, although there is no minimum program duration for the EIC Accelerator.

Under the SME Instrument, firms applied to specific calls by submitting a thirty-page proposal containing a business plan and a description of the planned activities; proposals were then evaluated remotely by four independent experts appointed by EASME. The EIC Accelerator follows a broadly similar logic—supporting high-impact innovations close to market—but uses a multi-stage evaluation process. Applicants first submit a short proposal (including a pitch deck and video), which is screened remotely by expert evaluators. Successful applicants are invited to submit a full proposal with a detailed business plan and work plan, evaluated again by independent experts. Projects that pass this stage proceed to an interview with an EIC Jury before funding decisions are made. Unlike the SME Instrument's fixed calls, the EIC Accelerator operates a continuous submission system with periodic cut-off dates.

Submissions are deemed eligible when three criteria are met: (i) the applicant must be a profit-oriented SME, (ii) it must be legally based in an EU Member State or a Horizon-associated country, and (iii) it cannot be simultaneously submitting or carrying out another project under the SME Instrument.

For our framework, it is important to focus on the definition of SME as part of criterion (i). The European Commission defines SMEs as firms with fewer than 250 persons employed and either an annual turnover of up to €50 million or a balance-sheet total of no more than €43 million. Under the EU's SME definition, ownership structure is a key determinant of SME status, distinguishing between autonomous, partner, and linked enterprises. An enterprise is considered autonomous if no other enterprise holds 25% or more of its capital or voting rights; holdings at or above this threshold generally create a partner relationship, requiring proportional aggregation of financials and employment. Full linkage arises when there is control—such as majority voting rights, the power to ap-

point management, or dominant influence through contractual arrangements—in which case 100% of the linked entities' data must be consolidated. Importantly, certain investors (e.g. venture capital firms, business angels below a specified investment ceiling, universities, and regional development funds) may exceed the 25% threshold without compromising autonomy, provided they do not exercise control. By contrast, enterprises that are 25% or more owned by public bodies are excluded from SME status altogether. This ownership-based framework is designed to ensure that only genuinely independent and economically constrained firms benefit from SME-targeted policies.

3 Data

The analysis is based on a novel dataset that links information on European Commission funding—extracted from CORDIS, the official repository covering all grants awarded between 2014 and 2024—to balance sheet data on recipient firms, sourced from Bureau van Dijk's ORBIS IP database (ORBIS henceforth). Online Appendix B presents the details of the matching procedure between CORDIS and ORBIS. We clean balance sheet data from ORBIS by following the standard procedure of Kalemli-Özcan et al. (2024).

We focus our analysis on the SME instruments, i.e. EU programs funding SMEs. As mentioned in Section 2, there are two main instruments devoted to SMEs funded by the European Commission. We combine these into a single consolidated program spanning 2014 through the last year of our sample.[4] To ensure that our sample includes only beneficiaries that received a single grant, we drop the 587 firms that obtained both Phase 1 and Phase 2 funding under the SME Instrument (10% of the sample).

3.1 Ownership and Shareholder Data

We rely on two sources of ownership information provided by ORBIS. First, each company is assigned a single "entity type", a categorical variable distinguishing between

[4] In this paper, *SME Instrument* (singular) refers exclusively to the specific programme implemented in 2014–2020. By contrast, SME instruments (plural) denotes the broader set of EU support instruments targeting SMEs, including the EIC Accelerator.

"independent company" and "controlled subsidiary". The labels of controlled subsidiary is assigned by Bureau van Dijk depending on whether there is an identified shareholder with over 50% direct ownership, or the company is identified as a branch or foreign. Independent companies are those with no identified shareholder with over 25% ownership share.

We complement the data with ORBIS' ownership module, which provides information on beneficiaries' complete shareholder structure and their corresponding ownership shares, distinguishing by shareholder type. Details on the shareholder data cleaning process can be found in Appendix B.

Online Appendix Table A3 presents ownership shares by type of shareholder. The most frequent shareholder types are "Corporate" and "One or more individuals or families". We focus on contrasting these two types of ownership because individual shareholders such as founders, entrepreneurs, or family owners typically reflect a form of entrepreneurship that differs fundamentally from the structured innovation processes typical of large corporate R&D labs. Individuals usually operate with fewer resources, face higher personal and financial risks, and rely more on tacit knowledge and experimentation, whereas corporate shareholders channel innovation through formal R&D departments embedded in broader strategic and organizational frameworks.

3.2 Descriptive Statistics

Our sample consists of an unbalanced panel of 1,940 companies observed for 9 years on average, operating in 17 NACE 1-digit industries and 35 EU27 and non-EU27 countries.[5]

Online Appendix Figure A1 presents the sample distribution by industry. It shows that the vast majority of companies in the sample belongs to Professional, technical & support services (34%), Manufacturing (31%) and Information and Communication (22%).

[5]In addition to EU Member States, Horizon programs include non-EU Associated Countries that have concluded an association agreement with the EU. Legal entities established in these countries may generally participate under conditions equivalent to those for Member-State entities, including eligibility for EU funding.

Online Appendix Figure A2 presents statistics by country. The sample includes EU27 and non-EU27 companies. The largest number of companies in the sample is found in Italy (18%), Spain (13%), United Kingdom (10%), and Germany (6%).

Table 1 presents summary statistics of the beneficiaries for the overall sample and distinguishing by type of entity. Panel A presents statistics for the full sample. The average revenue is €4 million, and the median is €560,000—indicating that the size distribution is skewed by a few larger companies. The maximum revenue is above €65 million, which is above the threshold allowed to be qualified as SME by the Commission. This is due to the fact that companies can grow over the years of our sample. In terms of employees, the companies tend to be small. The average number of employees is 21 and the median is 8.

The companies in our sample received a total of approximately €1 billion in R&D grants, evenly distributed between controlled subsidiaries and independent companies. Table 1 shows that on average firms receive €500,000 in grants, while the median grant is €50,000. In line with the program design described in Section 2, the maximum grant received in the sample is €2.5 million, while the minimum is €50,000.

The average program duration is 10 months, while the median duration is 5 months. We notice that some companies participate in the program for as long as 44 months.

Panel A1 and A2 present statistics for controlled subsidiaries and independent companies separately. The average revenue is more than €5 million for controlled subsidiaries and slightly less than €2 million for independent companies. Also in this case, the size distribution is skewed, with median revenue being lower than the mean in both cases. This can be also observed in Online Appendix Figure A1, which plots the distribution of log-revenue for the two entity types.

Despite the significant differences in size, controlled subsidiaries and independent companies receive similar grants and stay in the program for the same amount of time.

Table 1: Summary Statistics.

Variable†	Obs	Firms	Mean	Median	Std. Dev.	Min	Max
Panel A: All firms							
Revenue ('000)	11,205	1,459	4,003	559	9,401	0	65,969
Employees	12,487	1,799	21	8	36	0	220
Grant ('000)	1,634	1,634	472	50	804	50	2,500
Duration (months)	1,635	1,635	10	5	11	2	44
Breakdown of Panel A by entity type:							
A.1 Controlled subsidiaries							
Revenue ('000)	6,793	855	5,476	935	11,231	0	65,969
Employees	7,551	1,043	26	10	41	0	220
Grant ('000)	984	984	409	50	750	50	2,500
Duration (months)	984	984	9	5	10	2	44
A.2 Independent companies							
Revenue ('000)	4,412	604	1,736	350	4,665	0	65,969
Employees	4,936	756	14	6	25	0	220
Grant ('000)	650	650	566	50	873	50	2,500
Duration (months)	651	651	11	5	11	2	44

†All variables are winsorized at the 1st and 99th percentiles. Monetary values are in euros unless stated otherwise.
Source: IEP-COMPET (Gros et al., 2025)

4 Ownership Structure

Before examining the causal impact of R&D grants on the beneficiaries, we examine the differences between controlled subsidiaries and independent companies using shareholder information. Since the shareholder data are time-varying, we use the ownership information corresponding to the year in which the companies in our sample receive the grant.

Figure 1 presents the distribution of ownership shares for each entity. The key messages of Panel (a) is that controlled subsidiaries are more likely to be run by other companies. Almost all controlled subsidiaries have more than 80% corporate shares.

The ownership share distributed for independent companies is skewed on the left. The bulk of independent companies has less than 50% shareholding by corporation, in line with the definition provided by ORBIS. There is still a positive mass on the right, which could be because ORBIS assigns a time-invariant entity type, while ownership status can change over time.

To characterize further differences in ownership structure between controlled subsidiaries and independent companies, we turn to logit regressions of a dummy equal to 1 if a company is a controlled subsidiary on the ownership share of corporations, individuals, and other two aggregate shareholder categories: financial institutions, and public and research entities. Online Appendix B presents details about the shareholder types aggregation used to construct the broad groups.

We include in the logit regression log-revenue as a proxy of firm size, and an index of ownership concentration to account for the different shapes of the distributions in Panel (a) of Figure 1.[6] We also add industry and country fixed effects in order to absorb characteristics that can systematically shape firms' ownership structures.[7]

The results are presented in Panel (b) of Figure 1, which presents the average marginal effects of each variable, normalized to have zero mean and unitary standard deviation in the sample. Relative to independent companies (the reference group), controlled subsidiaries are less likely to be owned by individuals but more likely to be owned by other firms, or public or research entities. They are also larger and have more concentrated ownership, with an average concentration index double the magnitude of independent companies. We do not find any significant difference in terms of ownership by financial institutions.

5 Empirical Approach

We base our causal analysis on the following event study regression:

$$Y_{ist} = \alpha_i + \gamma_{st} + \sum_{k=-T}^{T} \beta_k D_{ik} + \epsilon_{ist}, \tag{1}$$

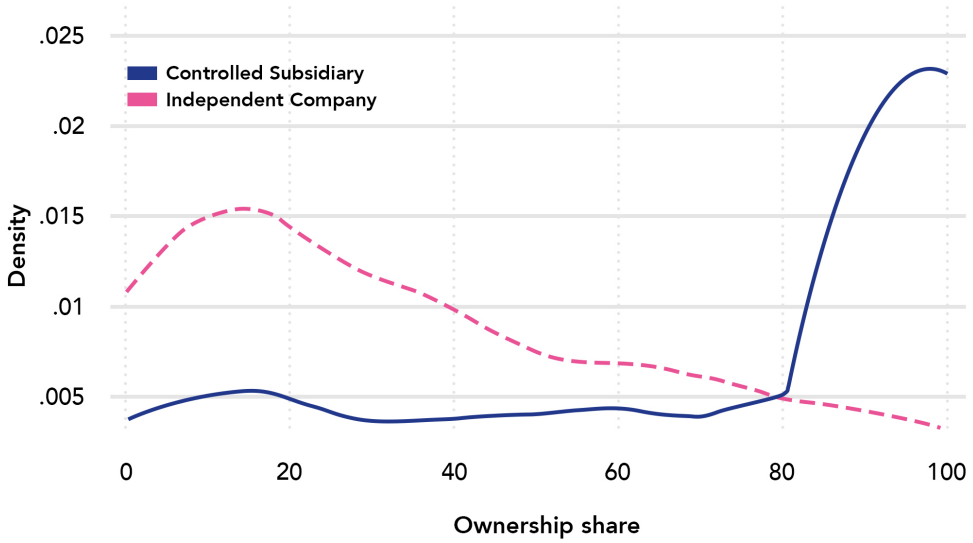
where Y_{ist} is the outcome variable of firm i in industry s at time t. The indicator function

[6] Specifically, we construct a Herfindahl–Hirschman Index based on the ownership shares of all shareholders.

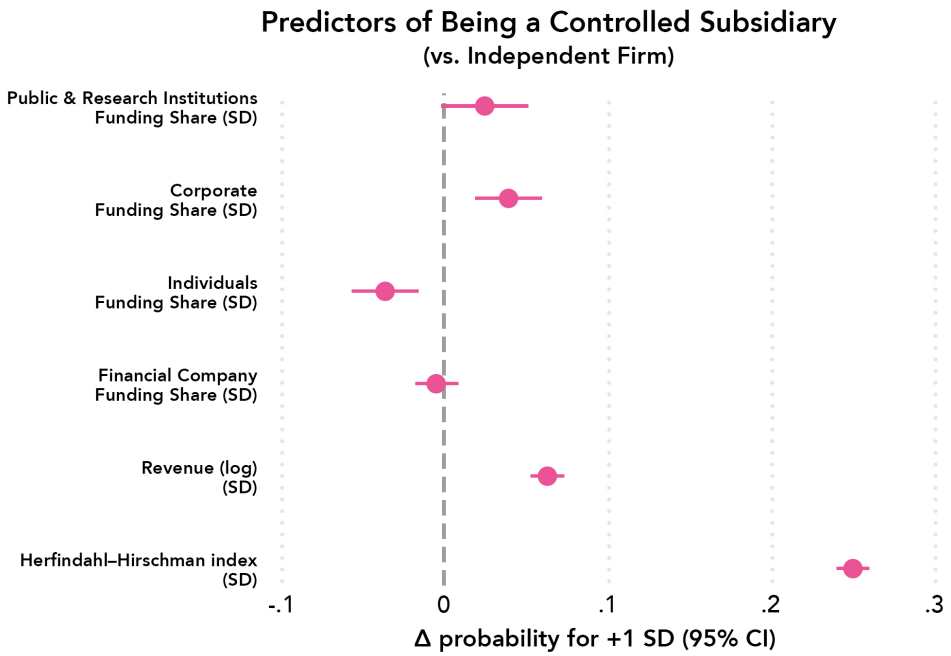
[7] For instance, firms in industries with high capital requirements, or countries with stronger investor protection or more developed capital markets might be more more likely to host controlled subsidiaries.

Figure 1: Ownership shares by shareholder and entity type.

(a)



(b)



Notes: this figure present the sample distribution of ownership shares by shareholder and entity type (a), and a characterization of controlled subsidiaries by their shareholder characteristics (b). In Panel (b), the coefficients are from a logit regressions of a dummy equal to 1 if a company is a controlled subsidiary on the ownership share of corporations, individuals, and other two aggregate shareholder categories: financial institutions, and public and research entities. The regressions include country, industry and year fixed effects.

D_{ik} captures the time to/from the grant and β_k is the associated coefficient. ϵ_{ist} denotes the error term.

In (1), α_i is firm-fixed effect and γ_{st} denotes industry-year fixed effects. The inclusion of γ_{st} allows to purge the estimates from the impact of unobserved sector-specific shocks.

The parameter T denotes the length of the pre/post treatment period and is chosen to be the largest possible depending on data availability. In order to avoid picking T arbitrarily, we set the earliest pre-treatment such that at least 500 observations are available. This results in the event windows [-5,6].

5.1 Endogeneity Concerns

The primary identification issue facing Equation (1) is that treatment—i.e., receiving a grant—is not randomly assigned across firms. For example, the European Commission might intentionally target companies based on their performance, potentially biasing the estimates. Firms treated last could be systematically different from those treated earlier, as well as from those never treated, compromising comparability. All such issues might result in a violation of the parallel trend assumption.

Some of these issues can be partially mitigated by including firm fixed effects, which control for firm-specific, time-invariant characteristics. However, fixed effects alone cannot address biases stemming from time-varying firm characteristics. Therefore, we adopt a strategy leveraging a sample of firms that applied to specific funding programs, were considered eligible by the evaluators, but ended up not receiving the funding due to budgetary reasons. Such firms are awarded by the Commission a quality label named *Seal of Excellence* (SoE).

Unfortunately, the European Commission does not disclose the complete list of SoEs. Therefore we compile a non-exhaustive list of 530 SoEs from Bpifrance (2025) and Demolin et al. (2025), which we manually match through their company name to ORBIS' balance sheet and ownership data.[8] Of these SoE firms, 368 are either controlled sub-

[8] A few SoE firms in our list are also in CORDIS. This suggests that they first applied unsuccessfully,

sidiaries or independent companies, and we use them as the control group.

The validity of our identification strategy hinges on grant recipients and SoEs to be comparable, ideally differing only in whether they received a grant. Therefore, Table 2 presents a comparison of key observable characteristics of SoEs and treated firm.

The table shows that controlled subsidiaries receiving the grant have an average revenue of €5.4 million, average employment of 27 and average age of 18 years. The SoEs counterparts have an average revenue of €6.5 million, 26 employment and age 14. Despite being somewhat younger on average, SoEs that are controlled subsidiaries are thus comparable to the treated firms.

For independent companies, grant recipients have average revenue of €1.8 million, employment 14 and age 16; the SoEs counterparts have average revenue of €1.1 million, 12 employment and age 14. Therefore, on average the treated have larger revenue, but they are comparable in the other observable dimensions.

Table 2 also shows that the size of the control group tends to be smaller than the treatment group. For controlled subsidiaries, the control group is 11% of the total; for independent companies, the control group accounts for 22% of the sample.

Table 2: Comparison of grant recipients and SoEs.

Entity type	**Obs**	**Firms**	**Revenue ('000)**	**Employees**	**Age**
Panel A: Sample excl. SoE (treatment group)					
ContrSubs	8,829	984	5,386	26.5	17.5
IndepCo	5,372	650	1,833	14.1	16.2
Panel B: Seal of Excellence (SoE) (control group)					
ContrSubs	866	126	6,501	25.8	13.9
IndepCo	1,068	180	1,115	11.9	14.1

Note: All variables are winsorized at the 1st and 99th percentiles. Monetary values are in euros unless stated otherwise.
Source: IEP-COMPET (Gros et al., 2025)

A final concern with our strategy is that since the list of SoEs is incomplete and gathered from public sources, there might be some sample selection. In particular, there might be unobserved factors correlated to being publicly disclosed as a SoE and company character-

but then obtained a grant in later funding rounds. We drop such firms from the dataset.

istics. For this reason, we implement an Inverse Probability Weighting (IPW) procedure assigning each firm a weight equal to the inverse of its probability of receiving the grant, conditional on observed characteristics: revenue, employment, total assets, NACE 1-digit sector classification, and year of grant receipt. To estimate such probabilities, we use a logit regression.

5.2 Productivity Measurement

Our main metric of performance is Revenue-Total Factor Productivity (TFP). To obtain TFP measures while addressing transmission bias (Griliches and Mairesse, 1995), we apply the two-stage proxy variable estimator of Ackerberg et al. (2015), and the dynamic panel method of Blundell and Bond (2000). In both cases, we estimate translog production functions separately for six 1-digit industries based on industry-deflated revenue/value added, capital, employment, and materials.[9] Online Appendix C presents the details.

An advantage of using revenue-based measures of TFP is that they capture the combined impact of several channels related to the performance of the grant recipient. This includes process and product innovation, but also more broadly the firm's ability to appropriate returns through higher markups or expanded market reach. An R&D grant may influence all these dimensions simultaneously, and while they are individually difficult to disentangle empirically, their joint contribution is precisely what matters for assessing the overall effectiveness of the grant. Revenue-based TFP therefore provides a compact, informative indicator of performance that aggregates these channels in a meaningful way.

6 The Causal Impact of R&D Grants

Given that firms receive grants at different times, we estimate Equation (1) using the estimator of Sun and Abraham (2021), which is designed for situations in which the

[9] We do not observe firm-level prices, a common issue in the literature (De Loecker and Warzynski, 2012; Klette and Griliches, 1996; Foster et al., 2008; De Loecker et al., 2016). For this reason, in our analysis we employ a battery of alternative performance indicators.

treatment time is staggered across units. The treatment group is composed by grant recipients, while the control group is composed by SoE firms, which never receive the grant.

Panel (a) of Figure 2 depicts the estimated impact of the grant on recipients' TFP. The figure presents point estimates and 95% confidence intervals from Equation (1), distinguishing between controlled subsidiaries and independent companies. There is no evidence of significant pre-trends, which reassures us on the validity of the parallel trend assumption required for a causal interpretation of the estimates. This suggests that *Seal of Excellence* awardees form a reasonable control group.

We find a positive, statistically significant and persistent impact of the grant on independent companies, but not on controlled subsidiaries. The grant increases TFP of independent companies by 25% during the first two years of the program, and then by 50% three-to-six years later. Controlled subsidiaries exhibit a milder increase in productivity, roughly 25%, but the effect dies out three years after the grant reception.

Panel (b) presents the estimated impact on log-revenue, which closely resemble the TFP estimates. The estimates show a positive and significant +50% impact of the grants on recipients' revenue, but only for independent companies. For controlled subsidiaries, the estimates are again hump-shaped, with a positive and significant +25% impact for the first two years after the reception of the grant. However, the estimates become small and not statistically significant thereafter.

We compare our estimates with the existing evidence from Santoleri et al. (2022). They find that winning an SME grant leads to a large and statistically significant increase in firm revenues *one year* after the award. Using a regression discontinuity design, the estimated impact on log revenues ranges from about 37% to 66%, depending on specification. The magnitude of our estimates is thus in line with the existing evidence, but they shed light on two important differences. The first is that the average effect masks heterogeneity in effects depending on ownership structure. The second is that the effect of the grant can change over longer time horizons, and therefore measuring the short-run

impact might be misleading.

6.1 Robustness

To check the robustness of our main findings, we examine the causal impact of the grant using alternative outcome variables.

The evidence strongly suggests that the pattern of Figure 2 is genuine, as it carries over to several other variables arguably related to research and innovation. Online Appendix Figure A2 shows the estimated impact of a grant on: (a) patents, (b) value-added per employee, and (c) a TFP measure based on different estimation assumptions. In all cases, the point estimates for independent companies are larger and more precisely estimated.

The impact of R&D on patents is presented in Panel (a). The post-treatment estimates for independent companies are large and statistically significant, suggesting that the grant increases the number of patents filed per year by about two. The base value of the patent count variable is roughly three (for both entity types). Although the estimates must be interpreted with caution due to the marginally significant pre-trend coefficients, the estimated impact of the grant on patents is quite large. Grants have a positive and significant impact on controlled subsidiaries, but the size of the coefficients is roughly one-fourth to those for independent companies, and it becomes precisely zero after six years.

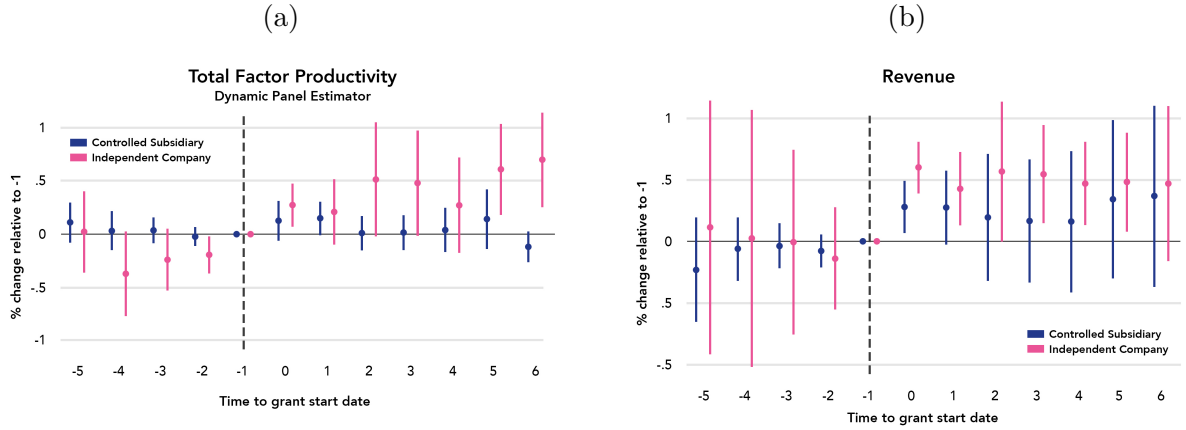
Panel (b) presents the estimates for value-added per worker, which is an important variable especially for manufacturing firms, constituting 30% of the sample (Online Appendix Table A1). Also in this case, the grant has a positive and significant impact on controlled subsidiaries, but the impact on independent companies is much larger.

Panel (c) presents the results for an alternative TFP measures based on the Proxy Variable Estimator. The estimates are fully consistent with the main productivity estimates.

Finally we check that the differential impact for the two entity types is not simply driven by systematic differences in size between controlled subsidiaries and independent compa-

nies. Online Appendix Figure A3 presents results based on a split sample of companies above and below the median value of log-revenue. If size was driving the result, than we should see reflected the pattern of Figure 2. However, this is not the case, as we do not find significant effects for both categories.

Figure 2: The causal impact of R&D grants on TFP and revenue.



Notes: this figure shows point estimates and 95% confidence intervals of the impact of R&D grants on TFP (a), and log-deflated revenue (b). Blue estimates correspond to controlled subsidiaries and red estimates concern independent companies.

The treated group is composed by firms receiving a grant in some year of the sample, while the controlled group is composed by firms that received the Seal of Excellence certification but never got a grant.

Estimates are based on the estimator by Sun and Abraham (2021). All regressions include firm fixed effect and 2-digit industry-year fixed effects. Standard errors are clustered at the firm level.

7 Conclusions

This paper provides causal estimates of the impact of R&D grants to SMEs and documents that it is mediated by recipients' ownership structure. Grants increase long-run productivity and performance more broadly, but only for independent companies. The effect is at best muted for controlled subsidiaries.

We characterize ownership in terms of the size of shareholders' ownership shares and their characteristics. Controlled subsidiaries tend to have highly concentrated ownership by corporations, while independent companies are more likely to have no shareholders or be owned by individuals. These differences can shape the impact of a subsidy over and above the impact of firm size, which is extensively documented in the literature.

Pinpointing the specific mechanism behind our causal effects goes beyond the scope of

this paper. Nonetheless, our evidence suggests that giving more weight to ownership characteristics in program requirements might substantially improve the efficiency of support schemes for research and innovation.

References

Ackerberg, D. A., K. Caves, and G. Frazer (2015). Identification properties of recent production function estimators. *Econometrica 83*(6), 2411–2451.

Bertrand, M. and S. Mullainathan (2003). Enjoying the quiet life? corporate governance and managerial preferences. *Journal of political Economy 111*(5), 1043–1075.

Bloom, N., J. Van Reenen, and H. Williams (2019). A toolkit of policies to promote innovation. *Journal of Economic Perspectives 33*(3), 163–184.

Blundell, R. and S. Bond (2000). Gmm estimation with persistent panel data: an application to production functions. *Econometric reviews 19*(3), 321–340.

Bpifrance (2025). Seal of excellence: High-quality project proposals under eic accelerator. Accessed via EuroQuity, in partnership with Wallonie Entreprendre.

Criscuolo, C., R. Martin, H. G. Overman, and J. Van Reenen (2019). Some causal effects of an industrial policy. *American Economic Review 109*(1), 48–85.

De Loecker, J., P. K. Goldberg, A. K. Khandelwal, and N. Pavcnik (2016). Prices, markups, and trade reform. *Econometrica 84*(2), 445–510.

De Loecker, J. and F. Warzynski (2012). Markups and firm-level export status. *American economic review 102*(6), 2437–71.

Demolin, M., T. Povel, and D. Germano (2025). Seals of excellence. In *Dealroom.co*.

Demsetz, H. and K. Lehn (1985). The structure of corporate ownership: Causes and consequences. *Journal of political economy 93*(6), 1155–1177.

Foster, L., J. Haltiwanger, and C. Syverson (2008). Reallocation, firm turnover, and efficiency: Selection on productivity or profitability? *American Economic Review 98*(1), 394–425.

Gertner, R. H., D. S. Scharfstein, and J. C. Stein (1994). Internal versus external capital markets. *The quarterly journal of economics 109*(4), 1211–1230.

Griliches, Z. and J. Mairesse (1995). Production functions: The search for identification. *NBER Working Paper No. 5067*.

Gros, D., S. M. Hofer, P.-L. Mengel, M. Molteni, G. Presidente, and F. Schimmel (2025). Compet dataset. *IEP Working Paper*.

Howell, S. T. (2017). Financing innovation: Evidence from r&d grants. *American Economic Review 107*(4).

Howell, S. T., J. Rathje, J. Van Reenen, and J. Wong (2021). Opening up military innovation: Causal effects of reforms to u.s. defense research. Technical Report 28700, NBER. Working Paper.

Kalemli-Özcan, Ş., B. E. Sørensen, C. Villegas-Sanchez, V. Volosovych, and S. Yeşiltaş (2024). How to construct nationally representative firm-level data from the orbis global database: New facts on smes and aggregate implications for industry concentration. *American Economic Journal: Macroeconomics 16*(2), 353–374.

Klette, T. J. and Z. Griliches (1996). The inconsistency of common scale estimators when output prices are unobserved and endogenous. *Journal of applied econometrics 11*(4), 343–361.

Rosati, N., P. Bomprezzi, M. Ferrraresi, A. Frigo, and M. Nardo (2020). Common shareholding in europe. *JRC Technical Report*.

Santoleri, P., A. Mina, A. Di Minin, and I. Martelli (2022). The causal effects of r&d grants: Evidence from a regression discontinuity. *Review of Economics and Statistics*, 1–42.

Shin, H.-H. and R. M. Stulz (1998). Are internal capital markets efficient? *The Quarterly Journal of Economics 113*(2), 531–552.

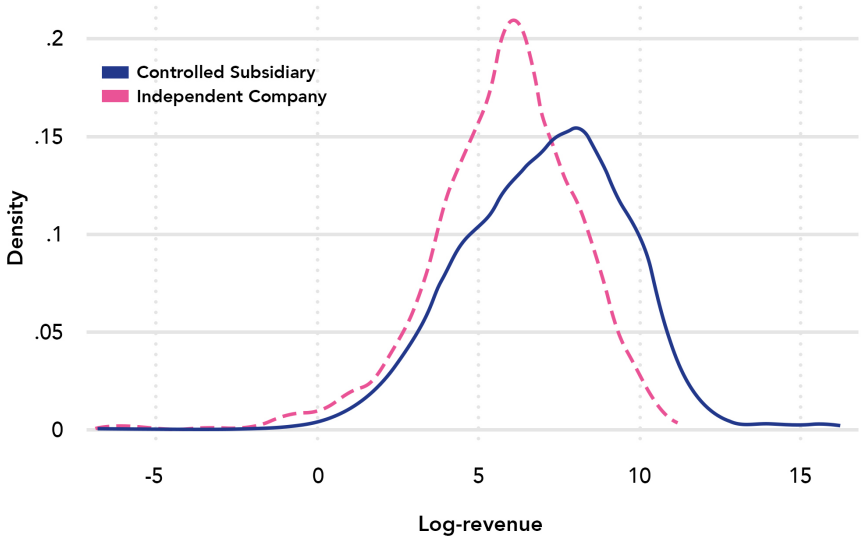
Shleifer, A. and R. W. Vishny (1986). Large shareholders and corporate control. *Journal of political economy 94*(3, Part 1), 461–488.

Sun, L. and S. Abraham (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of econometrics 225*(2), 175–199.

Yeh, C., C. Macaluso, and B. Hershbein (2022). Monopsony in the us labor market. *American Economic Review 112*(7), 2099–2138.

A Figures and Tables

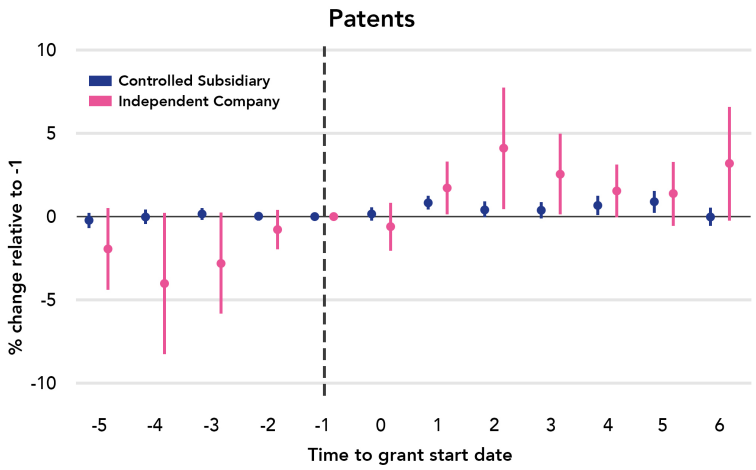
Figure A1: Size Distribution by Entity Type.



Notes: this figure presents the log-revenue distribution of firms in the sample, distinguishing between controlled subsidiaries (solid line) and independent companies (dash line).

Figure A2: The causal impact of R&D grants on alternative outcomes.

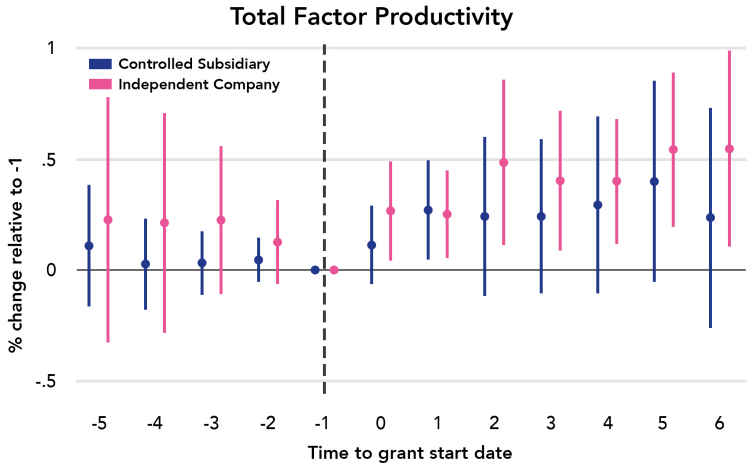
(a)



(b)



(c)



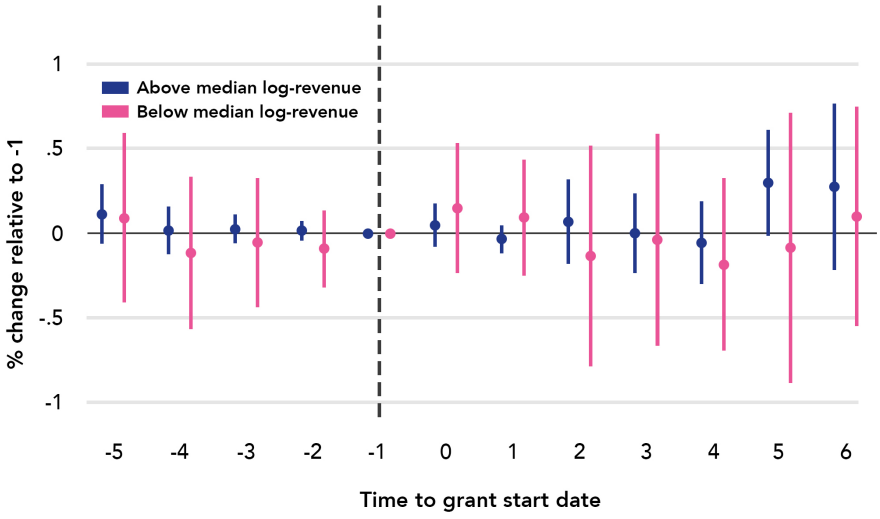
Notes: this figure shows point estimates and 95% confidence intervals of the impact of R&D grants on alternative outcomes. Blue estimates correspond to controlled subsidiaries and red estimates concern independent companies. The treated group is composed by firms receiving a grant in some year of the sample, while the controlled group is composed by firms that received the Seal of Excellence certification but never got a grant.
Estimates are based on the estimator by Sun and Abraham (2021). All regressions include firm fixed effect and 2-digit industry-year fixed effects. Standard errors are clustered at the firm level.

Table A1: Summary Statistics by Industry.

NACE sector	Obs	Firms	% sample	Avg. Revenue ('000)
A — Agriculture, forestry and fishing	123	15	0.8	1,913
C — Manufacturing	5,170	594	30.6	7,500
D-H — Utilities, construction, transport	570	68	3.5	5,051
G — Wholesale & retail trade	1,010	116	6.0	3,813
J — Information & communication	3,454	434	22.4	2,834
K — Financial & insurance activities	121	16	0.8	4,203
L — Real estate activities	66	7	0.4	2,246
M-N — Professional, tech. & support services	5,288	649	33.5	1,209
P — Education	47	5	0.3	948
Q — Human health & social work services	175	21	1.1	1,607
R — Arts, entertainment & recreation	20	2	0.1	.
S — Other service activities	82	11	0.6	1,782
NA — Not available	9	2	0.1	2
Total	16,135	1,940	100.0	

Note: All variables are winsorized at the 1st and 99th percentiles. Monetary values are in euros unless stated otherwise.
Source: IEP-COMPET (Gros et al., 2025)

Figure A3: Causal Impact of R&D Grants on TFP by Size Class.



Notes: this figure shows point estimates and 95% confidence intervals of the impact of R&D grants on TFP. Blue estimates correspond to companies with log-revenue above the sample median, and red estimates concern those below the median.
The treated group is composed by firms receiving a grant in some year of the sample, while the controlled group is composed by firms that received the Seal of Excellence certification but never got a grant.
Estimates are based on the estimator by Sun and Abraham (2021). All regressions include firm fixed effect and 2-digit industry-year fixed effects. Standard errors are clustered at the firm level.

Table A2: Summary Statistics by Country.

Country	**Obs**	**Firms**	**% sample**	**Avg. Revenue ('000)**
Austria	399	46	2.4	15,237
Belgium	224	28	1.4	7,265
Bulgaria	102	12	0.6	841
Croatia	29	3	0.2	434
Cyprus	10	1	0.1	19,550
Czechia	163	20	1.0	559
Denmark	378	59	3.0	356
Estonia	270	31	1.6	1,844
Finland	672	77	4.0	3,013
France	672	105	5.4	5,616
Germany	863	124	6.4	7,003
Greece	20	2	0.1	2,390
Hungary	547	58	3.0	1,414
Iceland	199	24	1.2	749
Ireland	487	63	3.2	6,657
Israel	84	28	1.4	20,984
Italy	3,208	344	17.7	5,718
Latvia	97	10	0.5	503
Lithuania	66	7	0.4	965
Luxembourg	42	5	0.3	7,926
Malta	6	1	0.1	10
Netherlands	523	68	3.5	6,711
Norway	471	58	3.0	1,438
Poland	388	50	2.6	2,591
Portugal	477	56	2.9	2,621
Romania	118	13	0.7	1,003
Serbia	37	4	0.2	5,460
Slovakia	152	16	0.8	1,143
Slovenia	374	39	2.0	3,667
Spain	2,124	248	12.8	2,412
Sweden	772	86	4.4	2,984
Switzerland	400	47	2.4	2,720
Turkey	3	1	0.1	34
Ukraine	10	1	0.1	15
United Kingdom	1,699	195	10.1	13,038
Total	16,135	1,940	100.0	

Note: All variables are winsorized at the 1st and 99th percentiles. Monetary values are in euros unless stated otherwise.
Source: IEP-COMPET (Gros et al., 2025)

Table A3: Share of ownership by type of shareholder type.

Variable†	Obs	Firms	Mean	Median	Std. Dev.	Min	Max
Panel A: All firms							
Corporate share (%)	7,357	891	51.30	50.00	39.20	0.00	100.00
Finance share (%)	3,710	452	25.50	16.10	28.00	0.00	100.00
Individual share (%)	10,884	1,258	67.10	92.00	39.50	0.00	100.00
Public research share (%)	413	49	19.00	10.00	26.50	0.00	100.00
Panel B: Controlled subsidiaries							
Corporate share (%)	3,537	403	68.00	89.20	37.30	0.00	100.00
Finance share (%)	1,265	144	30.70	19.40	33.10	0.00	100.00
Individual share (%)	5,098	560	75.70	100.00	35.90	0.00	100.00
Public research share (%)	116	14	30.60	12.30	36.90	0.00	100.00
Panel C: Independent companies							
Corporate share (%)	2,018	260	30.20	22.00	30.00	0.00	100.00
Finance share (%)	1,384	182	19.60	13.00	22.10	0.00	95.90
Individual share (%)	3,455	414	58.00	69.70	41.20	0.00	100.00
Public research share (%)	201	24	9.30	6.80	11.70	0.00	49.90

Note: All variables are winsorized at the 1st and 99th percentiles. Monetary values are in euros unless stated otherwise.
Source: IEP-COMPET (Gros et al., 2025)

B Data Appendix

B.1 Matching CORDIS to ORBIS

In order to combine the company-grant information from CORDIS with firms' financial data from OBRIS we rely on name matching. For details, consult Gros et al. (2025).

In short, we identify unique matches in three steps. First, we use the VAT codes which Corids provides to find entities in ORBIS. This identifies candidate matches in ORBIS, often multiple ORBIS entities per Cordis entry. To narrow down potential matches, we use additional information like country, postcode or street adress. After this procedure, we find 28,606 ORBIS entities for 22,930 entries. We add a variable with the number of matches for each Cordis entry. For the remaining Cordis entries, we use ORBIS 'batch-search' tool to identify unique entities. We accept high-confidence matches (A-score), and review the remaining matches manually. We end up with 8,828 one-to-one matches. For 1,071 Cordis entries we could not find any match in ORBIS. Of the total 37,434 companies we then select the relevant subset for our case of study. We merge the funding information with the firms' financial data requires using the project start date provided by CORDIS.

B.2 Shareholders Data

To get shareholder data for our ORBIS matches, we download shareholdership from 2002 to 2024 data from ORBIS and clean it in line with Rosati et al. (2020). This includes the shareholder id, the shareholder share, and the update date of this information. We then impute the link existence like Rosati et al. (2020), where we assume that a link persists if it is only missing for one year. We also transform non-numerical shares, e.g. 'MO', or 'Majority owned' into numbers like Rosati et al. (2020). Details can be found in the code.

One issue of note is that shares can sum up to more than 100% in ORBIS. This is because of how ORBIS updates their data. In essence, they add new data when they

find it, together with an update-date. So, if a new share is added (e.g. because of a press announcement) which brings the sum above 1, no action is taken and ORBIS simily shows all shares, even if they sum to more than one. To remedy this, we assume that, within a year, the newest shares are the most reliable ones. To then compute the yearly shareholdership, we check if the share sums to more than one. If it does, we remove the oldest share in that year until we get below a sum of 1. The remaining shares are then taken as shareholdership for that year (Note, we download the data always for the end of the year, so the shareholdership for a given year should be ready as end-of-year shareholdership). In our final dataset there are then no observations where the sum of all shareholder-shares for a given year is above 1.

In the end, we have a dataset with, for every company, the year, shareholder name, shareholder bvdid, and the share. We do have entries with 'empty' shares, i.e. no information. However, we keep these to be able to count the number of shareholders. Note, we have this dataset for the whole COMPET dataset (Gros et al., 2025), which we later subset to only include company that fit our purpose of study. From this dataset we can then construct indicators for ownership. For example, we group shareholder types and form three categories: i) Corporations; ii) Financial institutions (Bank; Financial company; Hedge fund; Insurance company; Mutual and pension fund, nominee, trust, trustee; Private equity firm; Venture capital); iii) Individuals (One or more named individuals or families; Self ownership), and iv) Public and research entities.[10]

C Production Function Estimation

C.1 Proxy Variable Estimator

We assume that in each year t, firm f produces value added v_{ft} with the following production function (in logs):

[10] We drop the categories left as they constitute less than 2% of the sample.

$$v_{ft} = f(k_{ft}, l_{ft}; \beta) + \omega_{ft} + \epsilon_{ft} \tag{C1}$$

The term ω_{ft} represents the log of Hicks-neutral productivity, which is known by plants' managers but not by us. The variable ϵ_{ft} is an i.i.d. error term that captures factors such as measurement errors.

We are interested in estimating the vector of the production function parameters β. To recover unbiased and consistent estimates of firms' production function (C9), we need to address the well-known simultaneity problem deriving from the fact that ω_{ft} is correlated to labour and materials but not to capital, which is chosen one period ahead. We build on the methodology of Ackerberg et al. (2015) and treat capital as a dynamic input, and labour and materials as static inputs. In particular, we make the following timing assumptions concerning inputs' decisions: i) capital k_{ft} is chosen at $t-1$ and ii) l_{ft} is chosen at $t-b$ after observing ω_{ft}. We further assume that iii) materials m_{ft} are chosen at $t-a$, with $1 < b < a$.

We specify the materials demand function, $m_{ft} = \tilde{h}(\omega_{ft}, k_{ft}, l_{ft})$. Assuming that the materials' demand function of the plant, $\tilde{h}$ is monotonically increasing and invertible in ω, we obtain a control function that proxies for unobserved productivity:

$$\omega_{ft} = h(k_{ft}, l_{ft}, m_{ft}) \tag{C2}$$

where $h \equiv \tilde{h}^{-1}$. Adding $h(\cdot)$ to Equation (C9), we get

$$v_{ft} = f(k_{ft}, l_{ft}; \beta) + h(k_{ft}, l_{ft}, m_{ft}) + \epsilon_{ft} \tag{C3}$$

We follow Ackerberg et al. (2015) by approximating the right-hand-side of (C3) with a third-order polynomial in all its elements.

From the first stage, we obtain expected output $\hat{q}_{ft}$ and the residuals $\hat{\epsilon}_{ft}$.[11] The next

[11] It should be noticed that in the first stage, none of the production function parameters are identified,

step is specifying a law of motion for productivity ω_{ft}. We assume that ω_{ft} follows a Markov process that can be shifted by plant managers' action:

$$\omega_{ft} = g(\omega_{f,t-1}) + \xi_{ft} \tag{C4}$$

The variable ξ_{ft} denotes the innovation to productivity.[12] Current expected productivity is then expressed as a function of the data and parameters:

$$\omega(\beta)_{ft} = \hat{v}_{ft} - f(k_{ft}, l_{ft}; \beta) \tag{C5}$$

To estimate β, we form moments based on the innovation ξ_{ft} in the law of motion (C4),

$$\xi(\beta)_{ft} = \omega(\beta)_{ft} - E[\omega(\beta)_{ft} | \omega(\beta)_{f,t-1}] \tag{C6}$$

The moments that identify the parameters are:

$$E[\xi(\beta)_{ft} M_{ft}] = 0 \tag{C7}$$

where the vector M_{ft} includes current capital and lagged labor, and materials consumption.

For our empirical application, we use a flexible Translog specification to approximate $f(\cdot)$.

The parameters of the production functions $\hat{\beta}$ are estimated separately for six 1-digit industries[13] with the Generalised Method of Moments (GMM), which we implement bootstrapping errors over hundred repetitions.[14]

because they enter both $f(\cdot)$ and $h(\cdot)$.

[12] In our application we impose a simple AR(1) form for (C4).

[13] TFP is estimated only for firms in the retained 1-digit NACE sectors (C; D–F and H; G; I; J; and M–N). Sectors are included only if they contain at least two distinct years of observations. We exclude sectors A, B, K, L, P, Q, R, S, and NA.

[14] Estimates are robust to using different numbers of bootstrap replications.

C.2 Dynamic Panel Estimator

We assume that in each year t, producer f produces gross output Q_{ft} with the following production function:

$$Q_{ft} = F(K_{ft}, L_{ft}, M_{ft}) \cdot \Omega_{jt} \tag{C8}$$

where M_{ft} are intermediate inputs, K_{ft} the capital stock and L_{ft} labor. The term Ω_{jt} represents Hicks-neutral productivity. Taking logs, we have:

$$q_{ft} = f(k_{ft}, l_{ft}, m_{ft}; \beta\mathbf{0}) + \omega_{ft} + \epsilon_{ft} \tag{C9}$$

In (C9), ϵ_{ft} is an error term, described further below.

We follow Blundell and Bond (2000) and assume that the productivity term ω_{ft} follows an AR(1) process:

$$\omega_{ft} = \bar{\omega}_f + \rho^{\omega}\omega_{ft-1} + \eta_{ft}$$

where η_{ft} is a MA(0) productivity shock.

Our objective is estimating $\beta\mathbf{0}$.[15]

The dynamic representation of the model is

$$\begin{aligned} q_{ft} &= f(k_{ft}, l_{ft}, m_{ft}; \beta\mathbf{0}) \\ &+\rho^{\omega}\Big(q_{ft-1} - f(k_{ft-1}, l_{ft-1}, m_{ft-1}; \beta\mathbf{0})\Big) \\ &+\epsilon_{ft} - \rho^{\omega}\epsilon_{ft-1} + \bar{\omega}_f + \eta_{ft} \end{aligned} \tag{C10}$$

We assume that the error term ϵ_{ft} includes a component reflecting shocks common to

[15] We also estimate ρ^{ω}, but it is not of primary importance in our framework.

all producers (u_t), a time-invariant firm-specific shifter (u_f), and an i.i.d. error term capturing measurement errors (ε_{ft}):

$$\epsilon_{ft} = u_t + u_f + \varepsilon_{ft}$$

We define the following terms:

$$u_t^* \equiv u_t - \rho^\omega u_t$$

$$u_f^* \equiv u_f - \rho^\omega u_f + \bar{\omega}_f$$

$$\varepsilon_{ft}^* \equiv \varepsilon_{ft} - \rho^\omega \varepsilon_{ft-1}$$

Then, we can rewrite (C10) as

$$\begin{aligned} q_{ft} &= f(k_{ft}, l_{ft}, m_{ft}; \beta\mathbf{0}) \\ &+\rho^\omega \Big(q_{ft-1} - f(k_{ft-1}, l_{ft-1}, m_{ft-1}; \beta\mathbf{0}) \Big) \\ &+u_t^* + u_f^* + \varepsilon_{ft}^* + \eta_{ft} \end{aligned} \tag{C11}$$

where η_{ft} is MA(0) and ε_{ft}^* is MA(1).

As in Blundell and Bond (2000), we further assume that $E[\Delta\tilde{x}_{ft} u_f^*] = 0$, where $\tilde{x}_{ft} = \{k_{ft}, l_{ft}, m_{ft}\}$, and $E[\Delta q_{f2} u_f^*] = 0$

Letting $w_{ft} = \varepsilon_{ft}^* + \eta_{ft}$, which is thus MA(1), and defining $x_{ft} = (q_{ft}, \tilde{x}_{ft})$, we can write the following moment conditions:[16]

[16] These are the same moment conditions used by Yeh et al. (2022).

$$E[x_{ft-s}\Delta w_{ft}] = 0,\ s \geq 3$$

$$E[\Delta x_{ft-s}(u_f^* + w_{ft})] = 0,\ s = 2$$

We choose a flexible Translog specification to approximate $f(\cdot)$. The parameters of the production functions $\hat{\beta}\mathbf{0}$ are estimated separately for 16 2-digit industries with the Generalised Method of Moments (GMM). We calculate standard errors with the sandwich estimator, which are thus robust to heteroskedasticity.[17]

[17] The cost of robust standard errors is that by relaxing the assumption of i.i.d. residuals, one cannot compute the Sargan test for the validity of over-identifying restrictions.