Production Network Spillovers from EU Funded R&D Grants*

Melko Dragojević[1], Ebehi Iyoha[2], and Stjepan Srhoj[3,4]

[1]Faculty of Economics and Business, University of Dubrovnik, Dubrovnik, Croatia

[2]Harvard Business School, Entrepreneurial Management Unit, Boston, MA

[3]Department of Economics, Faculty of Economics, Business and Tourism, University of Split, Split, Croatia

[4]Department for European Policies, Institute for Development and International Relations, Zagreb, Croatia

First version: March 9 2026; This version: June 1 2026

Abstract

Firms are embedded in production networks, so the returns to innovation policy may extend far beyond direct recipients. We quantify upstream propagation of R&D grants using administrative data from Croatia's €200 million program and comprehensive firm-to-firm transaction records. Grants generate large direct effects: recipient revenue rises by about 17%, material costs by about 22%, employment by about 8%, and tangible assets by about 27%. These gains transmit strongly to suppliers. Exposed suppliers expand persistently, with revenue increasing by roughly 28–46% and material costs by roughly 30–51%; employment rises more modestly but durably (about 9–12% from year 1 onward), while tangible assets accumulate gradually (roughly 17–40% over years 1–4). Spillovers are strongest for geographically proximate ties but remain meaningful well beyond local relationships. They accrue disproportionately to smaller firms and are more pronounced among financially constrained suppliers. Accounting for supplier responses raises the implied total impact to about four times the direct recipient effect. Production networks are therefore central to assessing the economic returns to innovation policy.

JEL Classification: O38, L14, H25, D22, R11

Keywords: R&D grants, production network, suppliers

*This study was supported by the Croatian Science Foundation under the project IP-2022-10-8560. We are particularly grateful to Scott Stern. We are also grateful for comments received from Pietro Santoleri, Sabrina Howell, Matt Elliott, Jesse Shapiro, Gaurav Khanna, Paul Niehaus, Karthik Muralidharan, Kyle Myers, Josh Krieger, Alex Coad, Verena Christina Wiedemann, Davide Samuele Luzzati, Alessandro Giannozzi, Ondřej Dvouletý, Jakša Puljiz, Vladimir Arčabić, Zoran Aralica, Slavo Radošević, Jadranka Švarc, Bruno Ćorić, Vinko Zaninović, Saša Čegar, Saša Drezgić, Boris Cota, and Hrvoje Šimović. We also appreciate feedback from participants at the UC San Diego Development Seminar, 2026 World Bank Group – IE University Conference, Harvard Business School Entrepreneurial Management Unit Faculty Lunch, and Duke Junior Strategy Conference. All errors remain our own.

1 Introduction

Do public R&D subsidies generate spillovers through supplier networks? Governments devote vast resources to business R&D support, yet we know surprisingly little about how these interventions propagate through production networks. While a rich literature documents knowledge spillovers from R&D investments (Bloom et al., 2019; Jaffe et al., 1993; Lucking et al., 2019), the role of input–output linkages in transmitting the effects of innovation policy remains underexplored. This paper provides evidence on network spillovers from R&D grants by tracing their effects through the production network using comprehensive firm-to-firm transaction data.

The economic rationale for R&D subsidies rests on well-established market failures. Arrow's (1962) seminal insight that knowledge is non-rivalrous and partially non-excludable implies that private returns to R&D fall short of social returns. This gap motivates public intervention, yet most evaluations focus narrowly on direct recipients (Bronzini and Iachini, 2014; Dechezleprêtre et al., 2023; Howell, 2017; Santoleri et al., 2024), potentially missing substantial indirect effects (Cai and Szeidl, 2024). Recent advances in production-network theory suggest that idiosyncratic shocks to individual firms can have aggregate consequences through input–output linkages (Acemoglu et al., 2012; Carvalho, 2014). If an R&D grant relaxes constraints and enables recipients to expand activity, the resulting increase in demand for intermediates can generate upstream spillovers to suppliers. This channel is distinct from the knowledge-spillover mechanisms emphasized in the innovation literature.

We study this mechanism using Croatia's €200 million R&D grant program, financed under the EU's Cohesion Policy framework and administered by Croatian national authorities. We combine grant records with administrative firm outcomes and transaction-level firm-to-firm sales data, which together allow us to track recipients, their pre-existing suppliers, and the evolution of buyer–supplier relationships before and after the grant decision. The 174 funded firms represent 0.1% of Croatia's firm population of about 140,000 firms, making aggregate crowding-out or general equilibrium displacement of competitors implausible at this scale. Our primary identification strategy is a matched-sample difference-in-differences design that compares funded applicants to observationally similar non-funded firms drawn from the universe of Croatian firms, conditioning on roughly sixty pre-treatment characteristics including multiple lagged outcomes and detailed transaction-network measures. For upstream spillovers, the primary specification is a matched-sample event-study at the supplier level that traces how firms linked to funded applicants through pre-existing buyer–supplier relationships adjust their outcomes following the grant decision. We additionally exploit the judges' evaluation scores assigned to each application, which provide quasi-experimental variation

in grant receipt at the program's eligibility threshold. The scoring data support two complementary cutoff-based robustness checks for direct effects on recipients. The first is a panel-based difference-in-discontinuities specification that compares marginally eligible to marginally ineligible applicants, using firm fixed effects to absorb time-invariant heterogeneity. We complement this with a local randomization design (Cattaneo et al., 2015) that validates identification through direct covariate balance tests within a narrow window around the funding threshold. Both designs are inherently local, and the discrete support and asymmetric density of the running variable leave only a small effective sample of applicants near the eligibility frontier. We use both rather than either alone: the panel structure of the difference-in-discontinuities specification recovers power by pooling repeated observations on each firm, while local randomization inference avoids reliance on asymptotic approximations that may be unreliable in small samples. For upstream spillovers, the same scoring data support a network-exposure design that uses cross-sectional variation in suppliers' pre-award sales shares to applicants near the cutoff, exploiting the same quasi-experimental variation in grant assignment. All cutoff-based estimates yield results consistent in sign and broadly similar in magnitude to the matched-sample specifications, and we report them as appendix sensitivity analyses.

Our results show that grant receipt generates economically meaningful direct effects on recipients. In our preferred matched-sample estimates, grant receipt increases recipients' revenue by about 17%, material costs by about 22%, employment by about 8%, and capital stock by about 27%.[1] These direct effects propagate strongly upstream. Exposed suppliers experience sustained post-award increases in revenue and material costs (roughly 28–46% and 30–51%, respectively), alongside more modest but persistent employment gains (about 9–12% from year 1 onward) and gradual capital accumulation (roughly 17–40% over years 1–4). Importantly for interpretation, spillovers accrue to firms that are substantially smaller than applicants: median employment is 4 for suppliers versus 40 for recipients, a tenfold difference. This is consistent with grants increasing demand that reaches further "downstream" in the firm-size distribution rather than merely amplifying outcomes among similarly large firms, and it implies that the large percentage effects documented for suppliers partly reflect a smaller baseline scale rather than a larger absolute response.

Spillovers are strongest for geographically proximate buyer–supplier ties, but they are not confined to immediate local areas. When we split supplier links by applicant–supplier distance terciles, we find statistically meaningful effects not only among nearby suppliers ($<$23.55 km) but also among suppliers located much farther away, including links in the top tercile ($\geq$171.39 km). This pattern suggests that knowing "*who trades with whom*" is critical: production-network connections transmit policy-induced demand across space in ways that purely place-based exposure measures can miss.

[1] Capital stock is proxied by tangible assets from firm balance sheets.

A final set of results uncovers a structural feature of upstream propagation that is both theoretically revealing and practically consequential. Examining heterogeneity by network centrality, we find a striking asymmetry between the role of the grant recipient and the role of the supplier. Supplier centrality strongly predicts spillover magnitude: suppliers in the top tercile of eigenvector centrality — those embedded in the core of the production network through connections to other important firms — exhibit the largest and most persistent upstream responses across all outcomes. Applicant centrality, by contrast, predicts little: if anything, suppliers connected to *less* central recipients respond more strongly, suggesting that a grant delivers a larger marginal demand stimulus precisely where the recipient is a smaller player in the network. The asymmetry implies that what governs the network multiplier of innovation policy is not the prominence of the firm that receives the grant, but the structural position of the firms through which the resulting demand shock travels. This finding speaks directly to the targeting of innovation subsidies: policymakers seeking to maximize upstream propagation should attend not only to recipient characteristics but to the network embeddedness of their suppliers.

Our findings contribute to three literatures. For the innovation policy literature, we document a novel channel through which R&D subsidies generate social returns beyond the funded firm. While existing work emphasizes knowledge spillovers through technological proximity (Bloom et al., 2019), geographic clustering (Jaffe et al., 1993), or labor mobility (Myers and Lanahan, 2022), we show that production-network spillovers operate through a distinct demand channel — grant-induced scale expansion at recipients reaches suppliers as increased procurement demand — and are quantitatively large: incorporating supplier responses raises the implied total effect to roughly four times the direct recipient effect. Standard program evaluations that focus on recipients alone therefore substantially understate the social returns to R&D subsidies.

For the production-networks literature, we provide the first firm-level causal evidence on how a positive policy shock propagates upstream through buyer–supplier links. Existing firm-level evidence comes almost exclusively from negative shocks, from natural disasters (Carvalho et al., 2021; Boehm et al., 2019) or financial disruptions (Barrot and Sauvagnat, 2016; Demir et al., 2024), and the closest evidence on positive shocks studies the specific context of domestic firms joining multinational supply chains (Alfaro-Urena et al., 2022), rather than how public funds allocated for innovation purposes propagate through pre-existing relationships between domestic firms. Sectoral-level evidence shows that positive public spending shocks transmit upstream through input-output linkages (Acemoglu et al., 2016; Barattieri et al., 2023), but with many firms operating within each sector, sectoral analysis cannot track which firms absorb and amplify these shocks, nor identify the mechanisms governing transmission. We fill this gap using administrative firm-to-firm transaction data that allow us to observe actual buyer–supplier relationships rather than imputed

sectoral exposure. Spillovers are largest among financially constrained suppliers, consistent with grant-induced demand relaxing binding liquidity constraints, and among smaller suppliers for whom the demand expansion represents a larger proportional shock. A grant generates larger network spillovers when it reaches suppliers that are well-embedded in the production network, regardless of how prominent the recipient is. Well-connected suppliers have more trading relationships on both the buying and selling sides, giving them greater flexibility to reallocate inputs, scale production, and absorb the demand expansion, generating a larger response from the same grant-induced demand shock than a more peripheral supplier would.

Finally, the paper speaks to the spatial incidence of innovation policy. Place-based evaluations typically find indirect effects confined to within a few kilometers of the treated location (Lerche, 2025; Neumark and Simpson, 2015; Bartik, 2020). We show that R&D grants generate economically meaningful spillovers to suppliers located well beyond the recipient's immediate vicinity, because what governs incidence is not where firms are located but which production relationships connect them. This implies that the benefits of innovation policy may accrue to regions and firms far from the funded recipient, a consideration absent from standard place-based cost–benefit frameworks and one that broadens the case for R&D subsidies as instruments of spatially inclusive growth.

Several features of our empirical setting support the external relevance of the results. Croatia's economy resembles many middle-income and transition contexts, combining manufacturing and services, a large SME sector, and meaningful regional disparities. The program's award sizes (€270,000–€1.4 million) are comparable to grant intensities commonly observed in innovation-support schemes, and the requirement of recipient co-financing helps ensure that funded projects reflect genuine commercial incentives rather than purely administrative participation. While the magnitude of spillovers may vary with the structure of supply chains and credit markets, the qualitative patterns we document—upstream propagation through observed buyer–supplier links, stronger effects for financially constrained firms, and economically meaningful spillovers that are not confined to recipients' immediate vicinity—are likely to be relevant in other settings where firms operate in dense production networks and face similar financing frictions.

The remainder of the paper is organized as follows. Section 2 describes the institutional setting and the R&D grant program, introduces the data, and explains how we construct applicant outcomes and production-network linkages from transaction records. Section 3 develops a conceptual framework that characterizes how firm-level grant responses depend on the joint distribution of innovative capacity and financial constraints in the host economy, and how grant-induced expansion at successful-project recipients propagates upstream through pre-existing buyer–supplier links. Section 4 presents the empirical framework, with the matched-

sample design as the main identification strategy for both direct effects and supplier spillovers, and two cutoff-based approaches—a difference-in-discontinuities specification and local randomization inference—reported in Appendices B and C as sensitivity analyses for direct effects. Section 5 reports the direct recipient effects and dynamic spillovers to suppliers, including heterogeneity by distance, supplier financial constraints, supplier size, and buyer project scores. Sections 6 and 7 aggregate the estimates into cost–benefit terms when supplier responses are incorporated and discuss implications for innovation policy evaluation in networked economies. Section 8 concludes.

2 Institutional Background and Data

Croatia's transition from a socialist economy to an EU member state provides a useful setting for studying how R&D support propagates through production networks. As a small open economy that joined the European Union in 2013, Croatia gained access to substantial structural funds aimed at narrowing the innovation gap with Western Europe. Prior to the expansion of EU-funded innovation instruments, R&D activity was low: in 2018, gross domestic expenditure on R&D remained below 1% of GDP, reflecting limited private-sector investment and weak innovation capacity. Beginning in 2019 and 2020, however, Croatia experienced a marked increase in R&D spending, coinciding with the rollout of large-scale EU-funded grant programs targeted at private firms. These programs are financed under the EU's Cohesion Policy framework, whereby EU funds are allocated to Member States and implemented by national authorities. In this context, the R&D grants we study are designed and administered by Croatian institutions, which define the evaluation criteria, scoring rules, and funding thresholds. This increase was driven primarily by the business sector, whose R&D expenditures rose sharply relative to 2018, signaling a shift toward greater private engagement in innovation activity. By 2023, total R&D spending reached approximately €1.1 billion, corresponding to 1.39% of GDP, with the business sector accounting for more than half of total R&D expenditures (DZS, 2025).

Our analysis centers on Croatia's flagship EU-funded R&D grant call, which was launched in late 2019 and accepted applications in early 2020.[2] The program's total budget was subsequently expanded to approximately HRK 1.56 billion (about €200 million). A defining feature of the scheme is substantial private co-financing: on average, the grant covered about 64% of total eligible project costs (grant intensity), well below the maximum co-financing rate of 80%, implying meaningful recipient skin-in-the-game and a large

[2]Original call title: *"Povećanje razvoja novih proizvoda i usluga koji proizlaze iz aktivnosti istraživanja i razvoja – faza II"* (call code `KK.01.2.1.02`).

combined investment shock rather than a marginal subsidy. The program received 295 applications, of which 174 firms were awarded funding and 121 were rejected, corresponding to an approval rate of 59%. Among funded firms, the mean grant amount was €1.37 million (median €1.00 million), underscoring the scale of support in the R&D arm.

The program allocates funding through a structured evaluation process based on multiple criteria rather than simple rank ordering. Each application receives points on several evaluator dimensions (e.g., *Project Design and Maturity*, *Implementation Capacity*, and *Value for Money*, alongside additional criteria such as *Novelty and Significance*, *Sustainability*, and *Potential for Commercialization*), as well as an overall score. Eligibility requires meeting minimum thresholds on each criterion and on the total score. Applicants that clear all thresholds are funded, while those failing any threshold are not funded, regardless of proximity to passing on other dimensions. This threshold-based assignment mechanism generates meaningful variation in funding status among otherwise comparable firms and provides a transparent institutional structure for constructing credible counterfactuals and conducting sensitivity analyses, even though our main empirical strategy relies on matching rather than exclusively on regression-discontinuity design. Detailed criterion-level scores are available for a subset of applicants (including both awarded and rejected firms), which we use when implementing cutoff-based robustness exercises.

Funded projects typically span 12 to 42 months and cover a broad set of eligible R&D-related expenditures, including wages for research and technical staff, investments in equipment and infrastructure, materials and intermediate inputs used in R&D activities, and contracted research services (including collaboration with universities and research institutes). Disbursements are staged and conditional on verified progress toward project milestones, limiting scope for mechanical bunching of expenditures around the award date. The funded portfolio includes technologically ambitious projects across multiple sectors. For example, *Rimac Automobili d.o.o.* received support for a project aimed at developing a software platform for autonomous driving under extreme and crisis conditions by improving perception, prediction, localization, path planning, and driver-behavior monitoring. This project is salient because it is systems-intensive frontier R&D that draws on specialized engineering and intermediate inputs, making upstream propagation through domestic supplier links especially plausible. Other illustrative projects include *PASTOR – Tvornica vatrogasnih aparata d.d.*'s development of an autonomous robotic firefighting device for prevention, early detection, and extinguishing of fires, and *SAMPLE CONTROL*'s *FoodForensics* project combining advanced analytical methods with machine learning for forensic testing of food, supplements, and medicinal herbs. We use these examples to convey the program's scope and the plausibility of network transmission through procurement

of specialized inputs and services.[3]

The flagship R&D call operated alongside other EU-funded business support schemes during the same period, making careful treatment of overlapping subsidies essential. We focus on two additional groups of programs: innovation commercialization grants and digital transformation grants.[4] These schemes are substantially smaller than the R&D program in both total volume and per-firm support. Innovation commercialization awarded 238 firms with an average grant of €271 thousand (median €182 thousand) and an average grant intensity of about 65%; 44 firms were denied funding, implying an approval rate of 84%. Digital transformation awarded 989 firms with an average grant of €48 thousand (median €37 thousand) and an average grant intensity of about 68%; 111 firms were denied funding, implying an approval rate of 90%. Because firms could participate in more than one scheme, we incorporate these programs directly into the main analysis: firms receiving other grants are excluded from the control group of non-funded R&D applicants, and supplier exposure measures are constructed so that suppliers linked to firms receiving other grants are not mechanically classified as untreated. This ensures that both direct comparisons among R&D applicants and estimated upstream spillovers reflect exposure to the flagship R&D instrument rather than concurrent policy interventions.

We combine several confidential administrative data sources to construct a panel of applicants, non-applicants, and their production-network connections. First, we use administrative grant records from the Ministry of Regional Development and EU Funds and the Croatian Agency for SMEs and Investments (HAMAG-BICRO), which include applicants (grant recipients and non-recipients), application timing, evaluation scores by criterion, threshold values, funding decisions, project duration, and grant amounts. These records are linked to annual firm financial statements (FINA), which provide standardized measures of firm outcomes and inputs, including revenue, material costs, employment, and balance-sheet items such as tangible assets. We additionally use the multidimensional firm-level default prediction metric (Omega Score) that captures financial difficulties faced by each firm in the economy (Altman et al., 2023). Finally, we use confidential transaction-level administrative data from the Tax Administration of the Ministry of Finance of the Republic of Croatia that capture firm-to-firm sales within the VAT system, allowing us to reconstruct buyer–supplier relationships and measure suppliers' exposure to treated applicants based on pre-award production links.[5] All data are accessed in secure environments and analyzed in anonymized

[3]For a full description of R&D projects see: https://efondovi.mrrfeu.hr/MISPublicApi/poziv/download/c8e242a1-e797-420e-8ace-a77187a779e4

[4]Innovation commercialization calls: *"Inovacije novoosnovanih MSP-ova – II faza"* (call code `KK.03.2.2.04`) and *"Inovacije u S3 područjima"* (call code `KK.03.2.2.06`). Digital transformation call: *"Poboljšanje konkurentnosti i učinkovitosti MSP-a kroz informacijske i komunikacijske tehnologije (IKT) – 2"* (call code `KK.03.2.1.19`).

[5]The microdata are accessed only in secure environments. Firm-to-firm transaction data are first merged using probabilistic record linkage and subsequently pseudoanonymized. The analysis reports only aggregate outputs; no individual firm identities

form.

Recipients are mature, mid-sized firms by Croatian standards, with a mean age of 19 years, median revenue of roughly €4 million, material inputs of roughly €2 million, tangible assets of roughly €460 thousand, and around 47 employees at the time of application (Table A1), and are heavily internationally oriented, with 86% being exporters and 73% being importers. Nearly 80% carry intellectual assets on their balance sheet, though only 5% hold a patent, suggesting that recipients are predominantly engaged in applied rather than frontier innovation. They are also well-embedded in production networks, maintaining on average around 65 supplier relationships and 25 customer relationships, substantially exceeding the population averages of roughly 11 suppliers and 4 customers. Notably, supplier purchases are geographically dispersed: recipients source more from suppliers located in other regions than from those in the same municipality, indicating that grant-induced upstream demand propagates through supply chains that already extend well beyond local boundaries. The sectoral composition reinforces this picture: nearly two-thirds of recipients operate in high- or medium-high digital-intensity industries and roughly 60% belong to higher-technology sectors (Appendix Table A2), consistent with the program's focus on commercially ambitious R&D projects with significant intermediate input requirements.

3 Conceptual Framework

This section develops a conceptual framework for understanding how R&D grants generate effects within a production network. The framework serves three purposes. First, it characterizes the structural conditions at the level of the firm and of the program that determine how much real activity a grant generates and through which channels it propagates. Second, it formalizes how imperfect information at the allocation stage maps into a distribution of recipient types, producing an ATT decomposition whose components are individually identified from the empirical pattern of treatment effects across outcomes. Third, it derives how grant-induced expansion at recipients propagates upstream to suppliers and characterizes the firm and network dimensions along which that propagation varies. We focus on upstream propagation throughout: the theoretical and data-boundary reasons for this choice are developed in Section 3.1.

or transaction links are disclosed.

3.1 Economy, Firms, and the Upstream Channel

Consider an economy of N firms indexed by $i = 1, \ldots, N$, each embedded in a fixed production network of pre-existing buyer–supplier relationships. Firm i produces output using capital K_i, labor L_i, and intermediate inputs M_i sourced from an upstream supplier set $\mathcal{S}_i$. Its output is sold to downstream firms as an intermediate input and to final demanders — households, government, and foreign buyers. Any shock to firm i can therefore propagate in four directions: upstream through procurement links to suppliers; downstream through input–output links to customers; horizontally to competitors in overlapping input and output markets (Bigio and La'o, 2020); and outward through knowledge diffusion across technological and geographic space to firms with no direct trading relationship (Myers and Lanahan, 2022). This paper is about the first channel.

For an R&D grant, the four channels operate on different timelines and with different degrees of certainty. The downstream and knowledge channels are slow and conditional. They activate only after the funded project succeeds commercially, and the recipient's revenue increase reflects the seller-side return to that success. Whether and how much downstream customers benefit is a separate question, governed by the nature of the innovation, the extent to which improvements are passed through in prices or embedded in product quality, and the importance of the input in the customer's own production function (Hulten, 1978; Acemoglu et al., 2012; Bigio and La'o, 2020; Iyoha, 2023). Knowledge spillovers to technologically proximate firms are similarly contingent on the type and scope of the innovation, the absorptive capacity of potential beneficiaries, and the strength of intellectual property protection (Jaffe et al., 1993; Bloom et al., 2019; Myers and Lanahan, 2022). The upstream channel is different in kind. By relaxing the recipient's financing constraint, the grant funds an expansion of input purchases — wages, intermediate inputs, capital — that flows to suppliers through pre-existing procurement links regardless of whether the funded project ultimately succeeds. This upstream demand channel is sign-unambiguous, first-order in magnitude, and activates the moment grant funds are disbursed. It is also the channel the existing evaluation literature has not measured. That is what this paper does.

The upstream channel is also causally cleaner to identify. The key exposure variable, whether supplier j transacted with buyer i before the grant decision, is predetermined and orthogonal to the grant outcome by construction. The other channels could in principle be studied but not credibly with the data available: the downstream channel requires price and quality variation absent from our transaction records; the horizontal channel requires a competitive set defined in product or technology space that our data do not support; and knowledge spillovers require patent data and a methodology for linking funded R&D to inventions by unrelated firms. Our data boundary points in the same direction. Given that 86% of recipients are exporters,

a substantial share of their downstream activity falls outside our firm-to-firm records entirely, while upstream supplier relationships are fully and symmetrically observed. The upstream channel is thus the direction our question, our data, and our design are jointly equipped to answer.

3.2 Firm Heterogeneity, Institutional Targeting, and the Recipient Population

The size and composition of the effects documented in this paper depend on two sources of heterogeneity that operate at different levels: firm-level characteristics that govern how a given recipient responds to funding, and a program-level characteristic that governs which firms are selected into the recipient population in the first place. Together, the joint position of a country on these dimensions, the prevalence of financial constraints, the distribution of innovative capacity, and the targeting quality of the allocation mechanism, constitutes a portable diagnostic for assessing where grants are likely to generate large direct effects and large upstream spillovers, and therefore how far the magnitudes documented in this paper are likely to travel across settings.

Financial constraints (λ_i). Firms differ in their access to external finance. Let λ_i denote firm i's financial constraint intensity, with lower values indicating tighter constraints. This parameter is the primary determinant of the upstream channel's size. When λ_i binds, the grant generates genuinely additional input expenditure, wages, materials, capital, that would not otherwise occur, producing a real upstream demand stimulus. When it does not bind, the grant substitutes for private funding and the upstream channel is muted: a financially unconstrained firm may well have the capacity and intention to undertake a high-potential R&D project regardless of public support, whereas that same project would not proceed at a firm without adequate funds. Financial constraints arise from two distinct sources, both of which must be kept separate from the knowledge-appropriability problem that motivates public R&D support. arrow1962 appropriability insight, that non-rivalrous, partially non-excludable knowledge generates a wedge between private and social returns, justifies intervention even for financially unconstrained firms; it is a market failure in the *price* of knowledge, not in the *supply* of capital. The financing frictions we focus on are distinct. The first is static: innovating firms tend to be asset-light and cash-flow-constrained, making debt costly and equity subject to adverse selection arising from the intangibility and uncertainty of R&D outputs (Hall and Lerner, 2010). The second is dynamic: at early innovation stages, technological uncertainty is maximal and no outside investor can price the project regardless of the firm's balance-sheet position, because no party outside the firm can assess the probability of technical success (Howell, 2017). This constraint loosens as

the project progresses and uncertainty resolves, creating a role for early-stage grants that private markets cannot replicate: Howell (2017) shows they unlock subsequent venture capital investment, and Santoleri et al. (2024) confirm across more than forty countries that the financing itself, not the quality signal conveyed by selection, drives real effects, establishing constraint relaxation as the operative mechanism.

Innovative capacity (A_i). Firms also differ in their ability to translate R&D inputs into commercially viable outputs, summarized by A_i (Furman et al., 2002; Aghion et al., 2005). The concept draws on the national and firm-level innovation systems literature, which identifies the stock of scientific and technical knowledge, the quality of human capital, and the strength of linkages between research and commercial activity as the core determinants of a firm's capacity to innovate productively (Furman et al., 2002). This parameter determines whether the recipient's funded project reaches the market. High A_i means the project succeeds commercially and the recipient generates revenue from the innovation. Whether and how much downstream customers subsequently benefit is a further question, governed by the nature of the innovation, the extent to which improvements are passed through in prices or embedded in product quality, and the importance of the input in the customer's own production function (Hulten, 1978; Acemoglu et al., 2012). Low A_i means input spending rises through the upstream channel but no commercially viable output materializes, so the recipient's revenue does not grow and the downstream and knowledge channels never fire. The two parameters are therefore complementary: λ_i determines whether the grant generates genuinely additional activity at all, while A_i determines what form that activity ultimately takes. Their interaction characterizes two empirically distinguishable configurations. When λ_i binds but A_i is low, input costs and assets expand but revenue does not, the upstream channel activates in isolation. When both λ_i binds and A_i is sufficient, input expansion is immediate while revenue growth and follow-on private investment follow as the funded innovation reaches commercial maturity, consistent with Howell (2017), with the asset-heavy, revenue-noisier effects documented by Santoleri et al. (2024), and with the finding by Myers and Lanahan (2022) that for every patent produced by a grant recipient three more are produced by others through knowledge diffusion. The relative prevalence of these configurations in any given program is an empirical question, and is precisely what the type decomposition in the next subsection identifies.

Institutional targeting quality (ϕ). The distribution of (λ_i, A_i) among firms in the host economy sets the *potential* for grants to generate real effects. How much of that potential is actually realized depends on a third parameter, $\phi \in [0, 1]$, which summarizes the *institutional targeting quality* of the program's allocation mechanism. High ϕ means the scoring and selection procedure accurately identifies firms for

which λ_i binds and A_i is sufficient; low ϕ means selection is corrupted by noise, political considerations, or criteria correlated with the wrong firm attributes. This parameter operates at the program level rather than the firm level: the same underlying distribution of (λ_i, A_i) can generate very different recipient pools, and therefore very different ATTs and upstream spillovers, depending on whether the allocation process is rule-based and transparent or dominated by discretion. The empirical importance of ϕ is established by **?**, who show that Italian investment subsidies allocated through objective, rule-based scoring generate substantially larger and more cost-effective employment gains than subsidies awarded through political discretion, even after conditioning on firm characteristics. The gap arises not from differences in firm quality, as both groups exhibit above-average performance and the two scoring components are conditionally independent of outcomes, but because political discretion systematically selects smaller firms requesting larger subsidies per worker, inflating the share of recipients who neither need the funding nor can translate it into market success.

The recipient population. The parameters (λ_i, A_i, ϕ) jointly determine who ends up in the recipient pool and what type of response they generate. A grant program awards funds to firms whose applications clear an evaluation threshold, but the granting authority cannot observe either a firm's counterfactual financing decision or the probability that its project will succeed commercially. Application materials reveal project content, implementation capacity, and commercialization potential, but even careful multi-criteria evaluation cannot fully separate firms that genuinely need public support from those that would have proceeded regardless, nor screen out projects that will fail in the market (Takalo et al., 2013; David et al., 2000; Dvouletỳ et al., 2021; Dimos and Pugh, 2016). This informational limitation means that grant recipients are a mixture of three latent types whose population shares (π_C, π_F, π_S) depend on (λ_i, A_i) in the host economy and on the targeting quality ϕ of the allocation mechanism.

Type C (*crowded-out*) firms have sufficient financing capacity to undertake the project without public support. Grants substitute for private funding and generate negligible additional real activity; their prevalence rises when λ_i is high or when ϕ is low and the program fails to screen them out.

Type F (*project-failure*) firms are financially constrained but lack the innovative capacity to translate inputs into commercially viable outputs. Grants induce additional spending that does not translate into revenue growth; their prevalence rises when A_i is low or when ϕ is low and the program funds commercially marginal projects.

Type S (*successful-project*) firms are both financially constrained and capable. Grants finance genuinely

additional R&D leading to marketable innovation, increasing input use, revenue, and potentially crowding in private financing (Howell, 2017); their prevalence rises when λ_i is low, A_i is high, and ϕ is high.

Type membership is unobserved to both the granting authority and the analyst. The population shares, however, are identified from the empirical pattern of treatment effects across outcomes, as the next subsection shows.

3.3 Type-Specific Responses

Let G_i denote the grant amount received by firm i, and let $\theta_i^M, \theta_i^K, \theta_i^L$ denote the shares of grant funds nominally directed to material inputs, capital, and labor, with $\theta_i^M + \theta_i^K + \theta_i^L \leq 1$. We characterize type-specific treatment effects on four observable outcomes: material costs (M), capital (K), labor (L), and revenue (R).

For *crowded-out* firms, the grant has no marginal effect on any outcome:

$$\Delta_C^M \approx 0, \quad \Delta_C^K \approx 0, \quad \Delta_C^L \approx 0, \quad \Delta_C^R \approx 0. \tag{1}$$

The disbursement appears in the firm's accounts but does not change input use, output, or revenue relative to the no-grant counterfactual, because the firm would have undertaken the project regardless.

For *project-failure* firms, the grant is spent on inputs in proportion to the planned expenditure shares, but revenue does not rise:

$$\Delta_F^M \approx \theta_i^M G_i, \quad \Delta_F^K \approx \theta_i^K G_i, \quad \Delta_F^L \approx \theta_i^L G_i, \quad \Delta_F^R \approx 0. \tag{2}$$

The defining empirical signature of project failure is that input use rises mechanically with the grant while revenue does not.

For *successful-project* firms, both the direct grant-funded expenditure and the equilibrium expansion of the firm's scale contribute to the observed response. Let $\xi_i^M, \xi_i^K, \xi_i^L \geq 0$ denote the additional input expansion beyond the directly grant-funded amount, and let $\xi_i^R > 0$ denote the revenue gain from successful commercialization:

$$\Delta_S^M = \theta_i^M G_i + \xi_i^M, \quad \Delta_S^K = \theta_i^K G_i + \xi_i^K, \quad \Delta_S^L = \theta_i^L G_i + \xi_i^L, \quad \Delta_S^R = \xi_i^R. \tag{3}$$

This is the case through which the grant generates genuine scale expansion at the recipient and induces

the upstream demand that propagates to suppliers above and beyond the mechanical pass-through of grant funds. Crucially, for firms with high innovative potential, the grant can trigger a project that pays off not only for the recipient but for the entire program. A single type-S firm with high A_i that is financially constrained, low λ_i, and correctly identified by the allocation mechanism, high ϕ, can generate revenue gains ξ_i^R and induced input expansion ξ_i^M large enough to offset, in aggregate welfare terms, the deadweight loss from crowded-out and project-failure recipients in the same cohort. Whether this occurs in practice depends on the same three parameters that govern the composition of the recipient population: the distribution of innovative capacity A_i in the host economy, which determines how large ξ_i^R can be at the upper tail; the prevalence of financial constraints λ_i, which determines how many high-A_i firms are genuinely blocked from proceeding without support; and the targeting quality ϕ of the allocation mechanism, which determines whether those firms are actually selected. A program operating in an economy with a thick upper tail of A_i, widespread financial constraints, and high ϕ is therefore one where the expected return to the marginal grant is not the average recipient effect but the potential payoff from funding the next high-potential project that would otherwise not have happened.

The observed average treatment effect on outcome Y is a population-weighted average of the three type-specific effects:

$$\text{ATT}^Y = \pi_C \cdot \mathbb{E}[\Delta_C^Y] + \pi_F \cdot \mathbb{E}[\Delta_F^Y] + \pi_S \cdot \mathbb{E}[\Delta_S^Y]. \tag{4}$$

Substituting the type-specific equations, and letting $\bar{\theta}^Y \bar{G} = \mathbb{E}[\theta_i^Y G_i]$ denote the population-mean grant allocation to outcome Y and $\bar{\xi}^Y = \mathbb{E}[\xi_i^Y]$ the population-mean induced expansion at successful-project firms:

$$\text{ATT}^M = (\pi_F + \pi_S) \cdot \bar{\theta}^M \bar{G} + \pi_S \cdot \bar{\xi}^M, \tag{5}$$

$$\text{ATT}^K = (\pi_F + \pi_S) \cdot \bar{\theta}^K \bar{G} + \pi_S \cdot \bar{\xi}^K, \tag{6}$$

$$\text{ATT}^L = (\pi_F + \pi_S) \cdot \bar{\theta}^L \bar{G} + \pi_S \cdot \bar{\xi}^L, \tag{7}$$

$$\text{ATT}^R = \pi_S \cdot \bar{\xi}^R. \tag{8}$$

Two features of this decomposition are central. First, crowded-out firms drop out of every input-side ATT entirely, so input-side effects are determined by the combined share $\pi_F + \pi_S$ of non-crowded-out recipients multiplied by their average grant-funded expenditure. Second, the ATT on revenue is determined solely by π_S and the average revenue gain conditional on success $\bar{\xi}^R$. The revenue ATT is therefore the cleanest population-level diagnostic for the prevalence of successful projects in the recipient population: a large and significant revenue effect implies that π_S is non-trivial and that the recipient population is not dominated

by crowding-out or project failure.

3.4 Microfoundation of Successful-Project Expansion

The induced expansion terms $\xi_i^M, \xi_i^K, \xi_i^L, \xi_i^R$ reflect the firm's optimal scale response to a positive demand shift generated by the new product or process. Consider a successful-project firm facing downward-sloping demand

$$Y_i = D_i \cdot p_i^{-\epsilon}, \quad \epsilon > 1, \tag{9}$$

and producing output according to

$$Y_i = F(K_i, L_i, M_i), \tag{10}$$

where $M_i = \left(\sum_{j \in \mathcal{S}_i} \omega_{ij} X_{ij}^{\rho}\right)^{1/\rho}$ is a constant-elasticity-of-substitution composite of inputs purchased from the firm's pre-existing supplier set $\mathcal{S}_i$, with substitution parameter $\rho \in (0, 1)$ and supplier-specific weights ω_{ij}. A successful R&D project shifts effective demand:

$$D_i' = (1 + \tau_i) D_i, \quad \tau_i > 0, \tag{11}$$

where τ_i summarizes the proportional demand shift induced by the innovation. Under standard cost minimization, optimal capital, labor, and intermediate input demand all rise in response:

$$\frac{\partial M_i^*}{\partial D_i'} > 0, \quad \frac{\partial K_i^*}{\partial D_i'} > 0, \quad \frac{\partial L_i^*}{\partial D_i'} > 0. \tag{12}$$

The induced expansion terms in equations (3) are these equilibrium adjustments: ξ_i^Y is the difference between the firm's optimal input use after the demand shift and its optimal use absent the shift.

Two features of the microfoundation matter for what follows. First, additional input demand at successful-project firms flows predominantly through the firm's existing supplier relationships, with shares determined by the weights ω_{ij}. We abstract from endogenous network responses such as supplier switching or entry, an approximation that implies the pre-existing buyer–supplier network provides a sufficient representation for tracing upstream propagation. Second, capital and labor adjust more slowly than intermediate inputs due to investment and hiring frictions. Following a sustained demand shift, firms first scale variable inputs and only gradually expand capacity through investment and hiring. The expected dynamic profile of a successful-project firm is therefore an immediate increase in revenue and material costs, followed by a more gradual rise in tangible assets and employment.

3.5 Propagation to Suppliers

Successful-project firms generate upstream spillovers through their pre-existing buyer–supplier relationships. We distinguish two channels through which a grant propagates to supplier j connected to treated buyer i.

The first is the *direct procurement channel*. To the extent that the grant directly funds input purchases from supplier j, those purchases mechanically appear in j's revenue. Letting s_{ji} denote the share of buyer i's pre-award material expenditure flowing to supplier j, the direct procurement contribution to j's revenue is approximately $s_{ji} \cdot \theta_i^M G_i$ summed across treated buyers. This channel activates for both project-failure and successful-project firms whenever grant funds are spent on inputs from j.

The second is the *induced demand channel*. When a successful-project firm scales production, its additional input purchases, captured by ξ_i^M in the microfoundation, flow through the same pre-existing supplier shares s_{ji}. This channel exists only at successful-project firms; project-failure firms generate no induced expansion because the underlying project does not reach the market. Total spillovers to supplier j are the sum of the two channels:

$$\Delta R_j \approx \sum_{i \in \mathcal{B}_j^T} s_{ji} \left[\theta_i^M G_i \cdot \mathbb{1}\{c_i \in \{F, S\}\} + \xi_i^M \cdot \mathbb{1}\{c_i = S\} \right] + \varepsilon_j, \tag{13}$$

where $\mathcal{B}_j^T$ is the set of treated buyers of supplier j and ε_j captures idiosyncratic supplier-side variation.

The two channels have different dynamic signatures. The direct procurement channel activates immediately when grant funds are disbursed and tracks the project timeline, typically 12 to 42 months in our setting. The induced demand channel activates only as the new output begins to generate market sales and persists as long as the innovation continues to generate demand, potentially well beyond the formal end of the grant project. A supplier response that grows over time and persists after the project's nominal end is the signature of the induced demand channel; a response concentrated in the project period and dissipating afterward is more consistent with direct procurement alone.

The theory predicts four dimensions of heterogeneity in supplier responses. First, financially constrained suppliers, those with low λ_j, exhibit larger proportional responses because the additional cash flow from buyer-side expansion relaxes binding liquidity constraints and enables investment and hiring that less constrained suppliers can already finance through other channels (Fazzari et al., 1988; Hubbard, 1998). Second, smaller suppliers exhibit larger proportional responses because a given absolute increase in buyer demand represents a larger share of baseline scale, $\frac{\Delta Y_j}{Y_j} \propto \frac{\sum_i s_{ji} \xi_i^M}{Y_j}$, consistent with granular perspectives on firm heterogeneity (Gabaix, 2011). Third, geographic distance attenuates propagation through transportation costs,

information frictions, and relationship-specific adjustment costs (Lerche, 2025; Allen and Arkolakis, 2018; Ramondo et al., 2016), but only conditional on a buyer–supplier link already existing; network connectedness is the core economic margin, distance a secondary moderating force (Acemoglu et al., 2012; Carvalho, 2014; Baqaee and Farhi, 2019). Fourth, network centrality generates an asymmetric prediction across the two sides of the buyer–supplier pair. For supplier j, high κ_j amplifies spillovers: a well-connected supplier can absorb an incoming demand expansion by reallocating inputs across its many relationships and re-transmit the shock further upstream, generating multiplier effects consistent with the role of the influence vector in governing aggregate shock propagation (Acemoglu et al., 2012; Golub, 2025). For grant recipient i, high κ_i attenuates the upstream stimulus: a firm already large and prominent as a buyer generates a smaller proportional demand increment per euro of grant, so suppliers see a smaller shock than they would from an equivalently funded but less central recipient (Criscuolo et al., 2019). Supplier centrality amplifies upstream spillovers while recipient centrality attenuates them, a prediction we term the *centrality asymmetry*.

These four heterogeneity predictions have direct empirical counterparts that the analysis in Section 5.4 is designed to test. Each reflects a distinct theoretical mechanism: constraint relaxation on the supplier side, scale effects in the demand shock, the primacy of network connectedness over proximity, and the asymmetric role of structural position in governing propagation intensity. Together they provide a set of sharp, falsifiable restrictions that go beyond the average spillover effect and speak to the economic channels through which innovation policy generates upstream activity.

3.6 Empirical Implications

The framework yields five sharp empirical implications.

First, a positive and economically large ATT^R on recipient revenue implies that π_S is non-trivial: the recipient population cannot be dominated by crowding-out or project failure. Via equation (8), a significant revenue effect requires both that successful projects exist in the population and that they generate meaningful market revenue.

Second, the relative magnitudes of input-side and revenue ATTs identify the composition of the recipient population. From equations (5)–(8), input-side effects scale with $\pi_F + \pi_S$ while revenue effects scale only with π_S. Input-side effects that are large relative to revenue effects signal a meaningful share of project failure; roughly proportional input and revenue effects signal that successful projects dominate.

Third, the dynamic profile of supplier spillovers, immediate revenue and material-cost responses followed by gradual capital and employment adjustment, is the empirical signature of the induced demand channel operating through successful-project expansion. Suppliers first scale variable inputs in response to the demand increase, then expand capacity through investment and hiring as the shock persists. This dynamic pattern is what the matched-sample event study in Section 5.3 is designed to detect.

Fourth, if upstream spillovers are several times larger than what direct grant-funded procurement could plausibly explain, given G_i, θ_i^M, θ_i^K, and s_{ji}, they must be operating predominantly through the induced demand channel. Both material inputs and capital goods are purchased from suppliers and therefore generate upstream demand through the production network: grant-funded capital expenditure flows to machinery and equipment suppliers just as intermediate input expenditure flows to material suppliers. The relevant benchmark is therefore total observed spillovers relative to the full direct procurement footprint across both input categories. Spillovers that substantially exceed this benchmark constitute additional evidence that π_S is meaningful in the recipient population, since the induced demand channel that drives the excess exists only at successful-project firms.

Fifth, the centrality asymmetry generates a distinctive cross-sectional prediction: spillover intensity should be increasing in supplier centrality κ_j and decreasing in recipient centrality κ_i, with the two gradients pointing in opposite directions. This is testable by comparing event-study estimates across terciles of supplier and recipient centrality separately, as in Section 5.4. A finding of opposite-signed gradients constitutes direct evidence that what governs the network multiplier of innovation policy is not the prominence of the node where the shock originates but the structural position of the nodes through which it travels.

4 Empirical Strategy

The empirical strategy is designed to map directly into the three mechanisms highlighted in the conceptual framework: (i) heterogeneous firm-level treatment effects at recipients, where the magnitude of the response may differ because of unobserved innovative capacity and financing frictions; (ii) selective grant allocation among observably strong but counterfactually heterogeneous applicants; and (iii) upstream propagation through pre-existing buyer–supplier links. Importantly, the conceptual framework does not require these sources of heterogeneity to be directly measured at the firm level. Rather, it provides an economic rationale for why treatment effects may vary across recipients and why the estimated average effect should be interpreted with external-validity considerations in mind. In this sense, whether these latent margins translate

into economically meaningful post-award differences is ultimately an empirical question. The European Innovation Scoreboard (European Commission, 2025) classifies Croatia as an "Emerging Innovator," below the EU average on most indicators of innovative capacity, while financial frictions remain salient in private credit markets. This places Croatia in an intermediate position: conditions for financial additionality are favorable, but heterogeneous innovative capacity limits the share of firms that can translate funding into commercially successful innovation. This heterogeneity shapes both the demand for R&D grants in our setting and the variation in real effects across recipients, motivating the targeting problem we analyze next.

This creates two empirical challenges. First, as described in Section 2 (and Appendix B), funding eligibility is determined by a multi-criteria evaluation rule, and applicants are strongly positively selected relative to the broader population of firms along virtually every observable dimension. Second, the conceptual framework predicts dynamic responses and spillovers that unfold over time through production relationships, which requires a design capable of tracing treatment effects relative to the award year and along observed network links. Our identification strategy is built to address both challenges simultaneously.

We work within the Rubin potential-outcomes framework (Rubin, 1974; Holland, 1986). Let $Y_{it}(1)$ and $Y_{it}(0)$ denote firm i's potential outcome at time t under grant receipt and under no grant receipt, respectively, and let $D_i \in \{0, 1\}$ indicate observed treatment. The object of interest is the average treatment effect on the treated (ATT),

$$\mathbb{E}[Y_{it}(1) - Y_{it}(0) \mid D_i = 1,\ t > T_i],$$

where T_i is firm i's award year. In the language of the conceptual framework, this ATT captures the realized firm-level expansion generated by relaxing financing frictions for treated firms.

Identification under selection on observables rests on two standard assumptions (Rosenbaum and Rubin, 1983; Imbens and Rubin, 2015). The first is the *conditional independence assumption* (CIA): conditional on a rich vector of pre-treatment covariates X_i, potential outcomes are independent of treatment assignment, $\{Y_{it}(1), Y_{it}(0)\} \perp D_i \mid X_i$. This assumption is especially natural in light of the conceptual framework, which implies that treatment assignment and post-award outcomes are both related to rich observable differences in pre-treatment firm characteristics and prior trajectories, even if the deeper sources of heterogeneity remain latent. We therefore construct X_i to approximate precisely the state variables emphasized by the framework, using approximately sixty pre-treatment characteristics drawn from financial statements, productivity measures, trade records, and the full firm-to-firm transaction graph. These include multiple lagged realizations of the four headline outcomes and detailed measures of buyer and supplier network position.

The credibility of CIA is therefore tied to whether treated and matched control firms are comparable not only in levels but also in the pre-award trajectories that capture the observable manifestations of these latent differences in growth potential and demand conditions. We assess this directly. Table A1 shows that post-matching balance is achieved on both the full covariate vector and the complete four-year pre-treatment paths of revenue, material costs, tangible assets, and employment. This is the empirical counterpart to the framework's requirement that, absent the grant, comparable firms would have followed similar dynamic paths.

The second identifying assumption is the *stable unit treatment value assumption* (SUTVA). In our setting, SUTVA must be interpreted through the network lens of the conceptual framework. For the direct-effects analysis, the relevant condition is that one applicant's treatment does not materially affect another applicant's outcomes through aggregate equilibrium channels. Given the small scale of the program (174 funded firms in an economy of more than 100,000 firms), this assumption is plausible.

For the spillover analysis, however, the framework explicitly predicts interference through buyer–supplier relationships. Rather than treating this as a violation, we redefine the unit of analysis so that the estimand aligns with the mechanism. The unit becomes the supplier, and treatment is replaced by *exposure* to a treated buyer through a pre-award transaction link. In other words, the empirical design follows the same propagation channel formalized in the conceptual framework: buyer expansion raises upstream demand, which then transmits to suppliers in proportion to relationship intensity and supplier dependence. To preserve this interpretation, we hold the buyer–supplier graph fixed at its pre-award realization, ensuring that estimated spillovers reflect transmission along existing links rather than endogenous post-treatment network rewiring.

Our main identification strategy is therefore a matched-sample panel design implemented symmetrically across the two margins implied by the framework. This symmetry between recipient-level and supplier-level designs is important for internal validity: it allows the empirical analysis to mirror the conceptual distinction between internal firm expansion and external network propagation while keeping identification assumptions transparent and comparable across estimands.

4.1 Estimating Direct Effects

For direct recipient effects, we compare funded applicants to matched non-funded firms with similar pre-award characteristics and estimate the average post-award treatment effect using a two-way fixed-effects

difference-in-differences specification:[6]

$$Y_{it} = \alpha_i + \gamma_t + \beta D_{it} + \varepsilon_{it},$$

where $D_{it} = 1\{t \geq T_i\}$ indicates the post-treatment period for funded firms, α_i are firm fixed effects, and γ_t are calendar-year fixed effects. The coefficient β identifies the average treatment effect on the treated over the post-award period within the matched sample. Internal validity rests on the standard parallel-trends requirement: absent grant receipt, treated and matched control firms would have experienced similar average outcome changes over time. We support this assumption through the rich pre-treatment matching procedure, balance on multiple lagged outcomes, and the close alignment of pre-award trajectories documented in Table A1. This static post-treatment estimand aligns with the conceptual framework's central prediction that grant funding relaxes financing frictions and shifts recipient firms onto a persistently higher outcome path, without requiring identification of fine-grained event-time dynamics.

4.2 Spillovers to Suppliers

We next turn to upstream propagation through the production network. A first identification challenge is that some suppliers serve a very large share of firms in the economy, including both treated recipients and their matched controls. This is especially relevant for utility-like providers, such as the main electricity generation company or dominant telecommunications firms, whose revenues may rise regardless of which downstream firms receive grants. Because these ubiquitous suppliers are mechanically exposed in both treatment states, they are not informative about relationship-specific propagation. We therefore begin by excluding shared high-degree suppliers and restrict attention to *non-shared supplier situations*, where a control supplier cannot simultaneously be a treated supplier through another buyer in the matched comparison set. This restriction ensures that the supplier-side treatment contrast is well defined and, as a byproduct, shifts the estimation sample toward smaller and more specialized suppliers — since high-degree suppliers tend to be large firms — which is consistent with the size differences between recipients and suppliers documented in Section 5.4.

A second challenge is comparability. Even if recipient firms are balanced on more than sixty pre-treatment variables, their suppliers need not be. Suppliers connected to treated buyers may differ systematically from

[6] Our setting does not raise the staggered treatment timing concerns documented in Baker et al. (2026) and related work (Callaway and Sant'Anna, 2021; Sun and Abraham, 2021): all treated firms are drawn from a single grant call with a single decision round, and matched controls are non-recipients rather than later-treated firms, ruling out the heterogeneous-timing bias and "forbidden comparisons" that motivate that critique. Firm and time fixed effects are retained alongside matching on over sixty pre-treatment characteristics because, even with rich observable balance, fixed effects absorb any residual unobserved time-invariant heterogeneity between treated and control firms.

suppliers connected only to control buyers in size, productivity, network centrality, and pre-trends. We therefore implement a dedicated supplier-level matching procedure before estimation, constructing exposed and non-exposed supplier samples that are balanced on rich pre-treatment observables and prior outcome paths. This additional matching layer is central for internal validity because the spillover estimand is defined at the supplier level, not inherited mechanically from recipient-side balance.

Having established a clean and balanced supplier comparison sample, we define supplier exposure as

$$E_{jt} = \sum_i s_{ji} D_i 1\{t \geq T_i\},$$

where s_{ji} measures the pre-award intensity of the supplier–buyer relationship. The key economic margin is therefore the strength of an *existing trade link*, rather than geographic proximity or sectoral coincidence.

We then embed the matched supplier sample in a panel event-study design:

$$Y_{jt} = \alpha_j + \gamma_t + \sum_{k \neq -1} \theta_k \, 1\{t - T_i = k\} \times E_j + u_{jt},$$

where E_j is fixed using the pre-treatment buyer–supplier graph. The coefficient sequence θ_k traces how supplier outcomes evolve after connected buyers receive the grant. Equivalently, the dynamic spillover estimand is

$$\theta_k = \mathbb{E}[Y_{j,T_i+k}(E = 1) - Y_{j,T_i+k}(E = 0) \mid E_j > 0],$$

which captures the average causal effect of network exposure among linked suppliers.

The main internal-validity requirement is that, conditional on supplier matching, supplier fixed effects, and common time effects, exposed and non-exposed suppliers would have followed parallel trends absent buyer-side treatment. Together with the prior exclusion of shared suppliers and the freezing of the pre-treatment network, this ensures that the event-study estimates recover dynamic propagation through existing trade relationships rather than contamination from common shocks, universal input provision, or endogenous rewiring. Appendix D reports a complementary supplier-side sensitivity analysis that uses cross-sectional variation in suppliers' pre-award sales shares to applicants near the eligibility cutoff as an alternative source of identifying variation.

5 Results

This section presents the empirical findings. We begin with the direct effects of grant receipt on funded applicants, then turn to upstream spillovers transmitted through pre-existing buyer–supplier relationships, and conclude with heterogeneity along three dimensions of the supplier population. Throughout, we focus on four headline outcomes: revenue, material costs, tangible assets, and employment. Sensitivity analyses—including the cutoff-based identification strategies described in Section 4—are reported in Section 5.5 and Appendices B, C, and D.

5.1 Matching R&D Grant Recipients

The R&D grant recipients are strongly positively selected relative to the universe of non-funded firms in the Croatian economy. Table A1 compares the 164 funded firms in our sample with the 106,027 non-funded firms in the population. Funded firms are older and substantially larger along essentially every dimension. The funded group has a much lower share of young firms (age < 4: 0.09 vs. 0.30) and a higher share of mature firms (age > 20: 0.49 vs. 0.22), with mean age of 19.03 versus 10.83 years. Scale differences are large and persistent across all four pre-award years. At $t = -1$, funded firms exhibit higher employment (3.85 vs. 1.14), tangible assets (13.04 vs. 7.33), revenue (15.22 vs. 10.37), and material costs (14.51 vs. 9.83), with similarly large gaps at $t = -2$, -3, and -4. Differences extend beyond scale: funded firms display higher labor productivity and higher ACF TFP, are substantially more engaged in international trade (exporter dummy at $t = -1$: 0.86 vs. 0.16; importer dummy: 0.73 vs. 0.16), and occupy markedly different positions in the production network, with larger supplier purchases (10.22 vs. 5.95), more suppliers (4.17 vs. 2.37), larger customer sales (9.85 vs. 4.69), and more customers (3.21 vs. 1.36). These pre-existing differences imply that any naive comparison between funded firms and the broader firm population would confound grant effects with large baseline disparities in firm fundamentals and network position, making selection on observables an essential precondition for credible inference.

A distinctive feature of our setting is that the administrative data support an exceptionally high-dimensional matching design. We construct a matched control sample drawn from the population of non-funded firms, conditioning on approximately sixty pre-treatment characteristics measured prior to the award decision. These span firm scale and input use, performance and productivity, innovation-related balance-sheet items, financial structure, sector and region, and detailed transaction-network structure on both the input and output sides. Crucially, key outcomes and trading variables are observed repeatedly prior to the award

decision, and the procedure conditions on *multiple* pre-treatment realizations of core outcomes and network measures rather than only the most recent year, which helps absorb persistent differences in firm trajectories and latent heterogeneity that would otherwise threaten identification.

Table A1 demonstrates that this matching procedure substantially alters the treated–control comparison. The matched analysis sample contains 159 funded firms paired with 159 observationally similar non-funded firms, and balance is achieved on essentially every dimension used in the matching, including age, size class indicators across multiple definitions, productivity, profitability, trade engagement, and granular measures of supplier and customer composition. Critically for identification, balance also extends to the entire pre-treatment trajectory of the four headline outcomes. Revenue differences between funded and matched control firms are 0.07, 0.04, −0.37, and −0.19 log points at $t = -1$, -2, -3, and -4 respectively; material cost differences are 0.08, 0.04, −0.33, and −0.18; tangible asset differences are −0.16, −0.16, −0.29, and −0.11; and employment differences are 0.07, 0.05, −0.09, and −0.05. None of these year-by-year differences is statistically significant, and none of the four outcomes displays a systematic monotonic drift toward (or away from) zero across the pre-period that would suggest divergent trajectories between funded and matched control firms in the absence of the grant. The same balance pattern holds for value added (IHS), profits (IHS), labor productivity, ACF TFP, exports, imports, and the full set of intellectual-asset and patent indicators across all four pre-award years, and it extends to transaction-network variables: supplier purchases (difference −0.01), top-1 supplier share (0.01), top-4 buyer share (0.01), purchases HHI, and the full distribution of sales and purchases across customer and supplier size and quartile groups.

Two implications follow. First, because year-by-year balance in pre-period outcomes is documented directly in Table A1 for each of the four pre-treatment years, this single table substitutes for an event-study pre-period plot in diagnosing parallel trends—and does so more comprehensively, by also balancing value added, profits, productivity, trade, and detailed transaction-network measures alongside the four headline outcomes. Second, the absence of systematic differences in both pre-treatment levels and pre-treatment trends, together with balance on a high-dimensional set of observables, supports the identifying assumptions of the matched-sample difference-in-differences framework and sharply reduces concerns that estimated post-award effects reflect pre-existing differences in levels or trajectories.

5.2 Direct Effects of R&D Grants on Winners Outcomes

Table 1 reports the matched-sample estimates of the direct effect of the R&D grant on applicant outcomes. Grant receipt increases recipients' revenue by 0.157 log points and material costs by 0.195 log points, corresponding to approximately 17.0% and 21.5% respectively. Tangible assets rise by 0.242 log points (about 27.4%), and employment increases by 0.076 log points (about 7.9%). These estimates are obtained in a panel specification with firm fixed effects and NUTS2–NACE 2-digit by year fixed effects, estimated on roughly 2,700 firm-year observations depending on the outcome. They summarize average post-award adjustments across the full distribution of funded firms in the matched sample.

Table 1: Direct effects of R&D grants on applicant firm outcomes

	Revenue	Material Costs	Tangible Assets	Employees
	(1)	(2)	(3)	(4)
R&D Grant	0.157***	0.195***	0.242**	0.076**
	(0.058)	(0.066)	(0.097)	(0.038)
Observations	2,697	2,697	2,677	2,697
R^2	0.962	0.960	0.944	0.972
R^2 adj.	0.898	0.893	0.851	0.925
RMSE	0.38	0.44	0.64	0.26
Mean (t-1)	44.490	33.538	19.117	156.333
Median (t-1)	2.374	1.310	0.441	40.000

Notes: Outcomes are in logs. Specifications include firm fixed effects and NUTS2 × NACE 2-digit × year fixed effects. Mean and median of the dependent variable refer to the $t = -1$ pre-period and are reported in millions of euros (employees in raw counts). $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Two patterns stand out. First, the joint increase in revenue and material costs indicates that grants translate into higher production and greater intermediate input demand, rather than purely financial or accounting changes at the firm level. The relative magnitudes—material costs rising slightly more than revenue—are consistent with grant-funded projects expanding scale through increased input intensity rather than through pure efficiency gains alone. Second, the substantial response on the tangible assets margin (about 27%) suggests that recipients direct a meaningful share of grant resources toward capacity expansion via investment in equipment and infrastructure, consistent with the program's eligibility rules that explicitly permit such expenditures. The employment response is more modest in proportional terms (about 8%) but is precisely estimated, indicating that hiring adjusts in response to grant receipt even as the bulk of the scale expansion operates through capital and intermediate inputs.

We treat these matched-sample estimates as our main results because they leverage the full distribution of

funded firms in the matched sample and deliver the statistical power needed to estimate dynamic responses and trace upstream spillovers through the production network. We assess the robustness of these estimates to alternative identification strategies in Section 5.5 and Appendices B and C, where we show that two complementary cutoff-based approaches exploiting the program's evaluation scores yield results consistent in sign and broadly similar in magnitude.

5.3 Spillovers to Suppliers

Suppliers connected to funded applicants are not randomly drawn from the universe of firms. They are jointly selected by (i) non-random formation of buyer–supplier links and (ii) positive selection of buyers into applying for, and winning, R&D grants. Unconditional comparisons between suppliers linked to grant recipients and suppliers in the broader firm population would therefore confound spillovers with persistent differences in supplier fundamentals and network position. We address this by extending the matched-sample logic of Section 5.2 from applicants to their suppliers.

Table A3 documents the magnitude of pre-existing selection. Before matching, treated suppliers—those linked in the pre-award period to firms that subsequently received R&D grants—are systematically larger and older than potential controls along essentially every dimension, including own scale, productivity, trade engagement, and position in the production network. We construct the supplier counterfactual using a matching procedure that conditions on a high-dimensional vector of pre-treatment supplier characteristics, including demographics, multiple lagged realizations of the four headline outcomes, performance and productivity measures, trading behavior, and granular transaction-network measures capturing the scale and composition of supplier sales. Conditioning on multiple pre-treatment realizations of the headline outcomes is essential for absorbing persistent differences in supplier trajectories that would otherwise threaten the parallel-trends assumption.

Post-matching, treated and matched control suppliers are balanced on all covariates and across all four pre-award years for each of the four headline outcomes (Table A3). We then estimate dynamic supplier responses using the event-study specification of Section 4.

Figure 1 reports the resulting event-time coefficients, with year -1 omitted as the reference period. Pre-award coefficients at $k = -4$, -3, and -2 are small and statistically indistinguishable from zero across all four outcomes, validating the parallel-trends assumption.

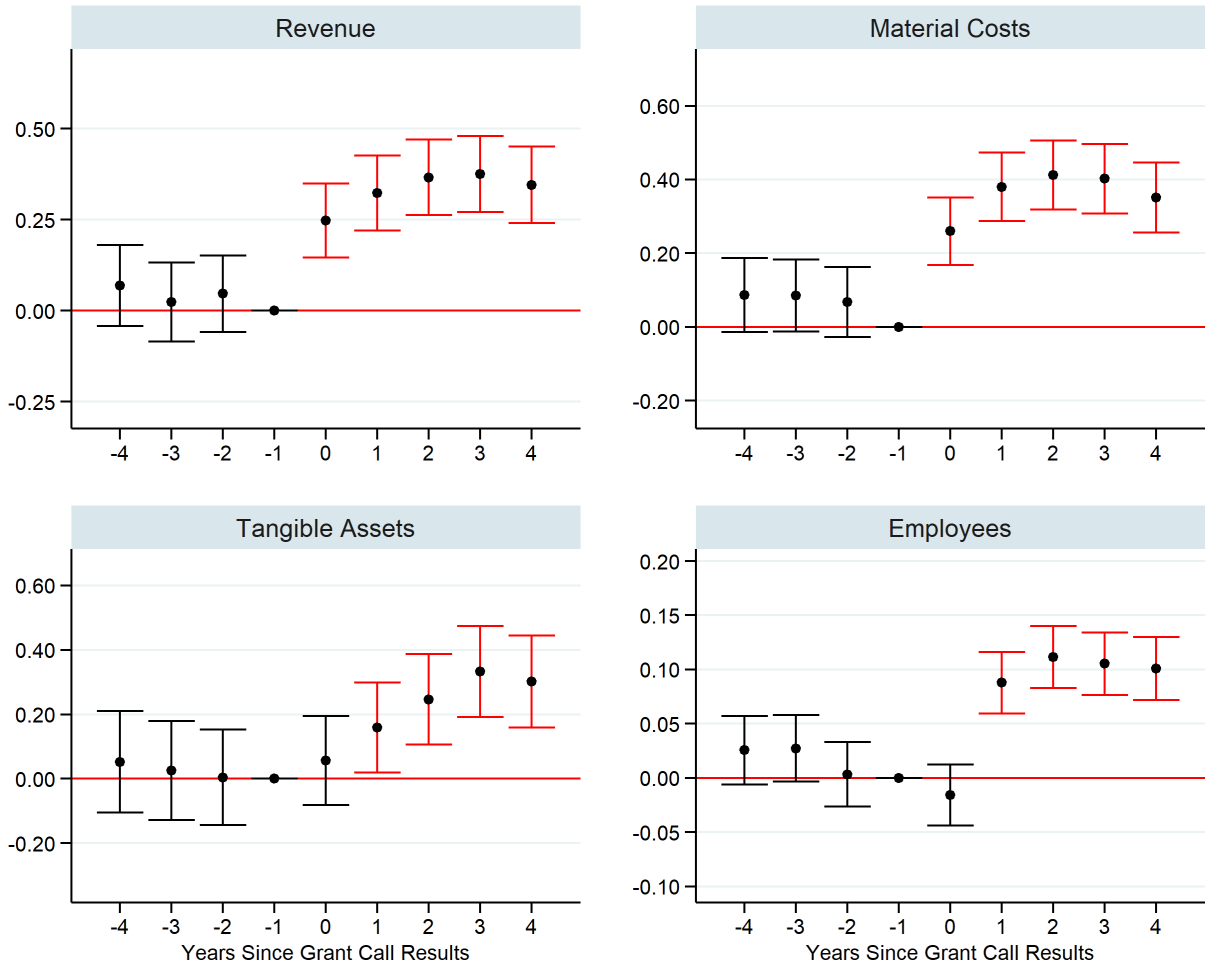


Figure 1: Dynamic upstream spillovers to suppliers after applicant grant awards

Note: Event-study estimates of supplier outcomes by event time (years since grant-call results). The omitted reference period is year -1. Whiskers denote 95% confidence intervals. Outcomes are supplier revenue, material costs, tangible assets, and employment. Pre-award supplier links are defined using relationships observed in the baseline (pre-treatment) period.

Following the grant-call results, suppliers in the matched sample exhibit sharp and persistent upstream adjustments. Supplier revenue and material costs rise immediately and continue to grow over the post-period: revenue increases by 0.248 log points (about 28%) at $k = 0$ and reaches 0.376 log points (about 46%) by $k = 3$, while material costs follow a closely aligned trajectory with slightly larger magnitudes (0.260 log points at $k = 0$ rising to 0.412 by $k = 2$). The joint evolution of supplier revenue and intermediate input use is consistent with real scale expansion among suppliers in response to a sustained increase in demand from their grant-funded buyers. Capital and labor adjustments lag the immediate scale response: tangible assets show no significant movement at $k = 0$ (0.057 log points) but accumulate from $k = 1$ onward, peaking at

0.333 log points (about 40%) at $k = 3$, while employment is essentially flat at $k = 0$ and rises to between 0.088 and 0.112 log points (about 9–12%) over $k = 1$ through $k = 4$. This timing is consistent with adjustment frictions: suppliers first scale variable inputs, then expand capacity through investment and hiring as the demand shock persists. Table A4 reports the underlying coefficients in detail.

Two features of the supplier response argue against a mechanical interpretation in which spillovers simply reflect grant resources flowing through to upstream firms via direct procurement. First, the magnitude of the supplier expansion is too large to be explained by direct pass-through alone. The cumulative supplier-side increase in revenue and material costs is several times larger than the share of grant funds that could plausibly be spent on intermediate inputs from any single supplier, particularly for the substantial fraction of suppliers whose share of sales to applicants is small. Second, the persistence and dynamic structure of the response point to an indirect channel operating through grant-induced expansion of recipient activity. Section 5.2 documents that grant receipt produces large and lasting increases in recipient revenue (about 17%), material costs (about 22%), employment (about 8%), and tangible assets (about 27%); these responses imply that grant winners not only spend the grant itself but also expand their broader operations and procurement durably over time. The timing pattern in the supplier event study—immediate revenue and input-use response, followed by gradual capital and labor adjustment that builds and persists for years after the formal grant period—matches what we would expect if upstream propagation operates predominantly through this indirect channel of expanded business at recipient firms, rather than through one-off grant disbursements. We interpret the results accordingly: the spillovers we document are the network footprint of grant-induced scale expansion at recipients, not merely the direct procurement footprint of the grant itself.

5.4 Heterogeneous Effects on Suppliers

Having established that suppliers expand in response to buyer-side grant receipt, we now examine how these upstream adjustments vary along three dimensions: geographic distance to the buyer, supplier financial constraints, and supplier size. Each dimension is informative about a distinct feature of how policy-induced demand propagates through the production network. Distance speaks to whether spillovers are spatially contained, with implications for the place-based incidence of innovation policy. Financial constraints speak to whether grant-induced demand operates as a substitute for missing external finance among suppliers, generating disproportionately large effects where credit frictions are most binding. Supplier size speaks to whether spillovers concentrate among small firms (where a given demand shock represents a larger share of baseline scale) or extend across the supplier size distribution. Each split uses the same matched-sample

event-study specification as the baseline analysis, estimated separately within tercile groups defined ex ante.

Geographic Distance to Buyer. A central question for the policy interpretation of upstream spillovers is whether they are spatially contained or whether grant-induced demand propagates across distance. If spatial frictions in production networks—transportation costs, information barriers, and relationship-specific adjustment costs—are quantitatively important, upstream responses should be concentrated among suppliers located close to funded buyers, and the geographic incidence of innovation policy would mirror that of place-based interventions. Recent evidence on related place-based programs documents strong spatial attenuation, with effects often confined to within a few kilometers of the treated location (Lerche, 2025). Whether the same pattern holds for spillovers transmitted through observed buyer–supplier links rather than through generic agglomeration channels is an open question, and one that matters directly for whether the benefits of innovation policy accrue to the immediate region of recipients or extend across broader spatial scales.

We compute straight-line distance between the municipalities of each funded applicant and its pre-award suppliers, and split links into terciles: bottom (< 23.55 km), middle (23.55–171.39 km), and top (≥ 171.39 km), with roughly 1,850–1,930 unique suppliers in each group. Figure 2 reports the event-study estimates by tercile.

Two patterns stand out. First, spillovers appear in all three distance groups, but the response is not monotone in distance: revenue and tangible asset effects are concentrated in the bottom and middle terciles (with revenue gains of 0.35 and 0.23 log points at $k = 0$ rising to roughly 0.40–0.54 over $k = 1$ to $k = 4$), while the top tercile revenue response is smaller and less precisely estimated. Second, several margins propagate even at long distances. Material costs and employment rise significantly in all three terciles, including the top distance group, and remain elevated through $k = 4$; for example, employment at $k = 2$ increases by 0.135, 0.112, and 0.099 log points across the bottom, middle, and top terciles respectively. Capital adjustment, by contrast, is more spatially concentrated, with sustained tangible asset accumulation only in the bottom and middle terciles. Tables A5 and A6 report the underlying coefficients.

The pattern indicates that upstream spillovers are not purely local. Geographic frictions attenuate some margins, particularly capital accumulation, but do not prevent grant-induced demand shocks from propagating through buyer–supplier links across substantial distances. This stands in contrast to evidence on place-based investment incentives, which often show effects concentrated within a few kilometers of the treated location (Lerche, 2025), and underscores that what matters for the spatial reach of innovation policy is not geographic proximity per se but the structure of the production network connecting firms. The non-local reach of these spillovers is consistent with the pre-existing geographic structure of Croatian supply chains, where firms

source inputs extensively from suppliers across regions (Srhoj and Mikulić, 2025). Importantly, both Dhyne et al. (2021) for Belgium and Srhoj and Mikulić (2025) for Croatia show that domestic supplier transactions embed substantial indirect import content, with most firms relying heavily on foreign inputs through domestic intermediaries rather than direct importing. This implies that the upstream spillovers we document reflect domestic value added net of such leakages, and that the true domestic multiplier of grant-induced demand may be somewhat smaller than the gross supplier revenue response suggests.

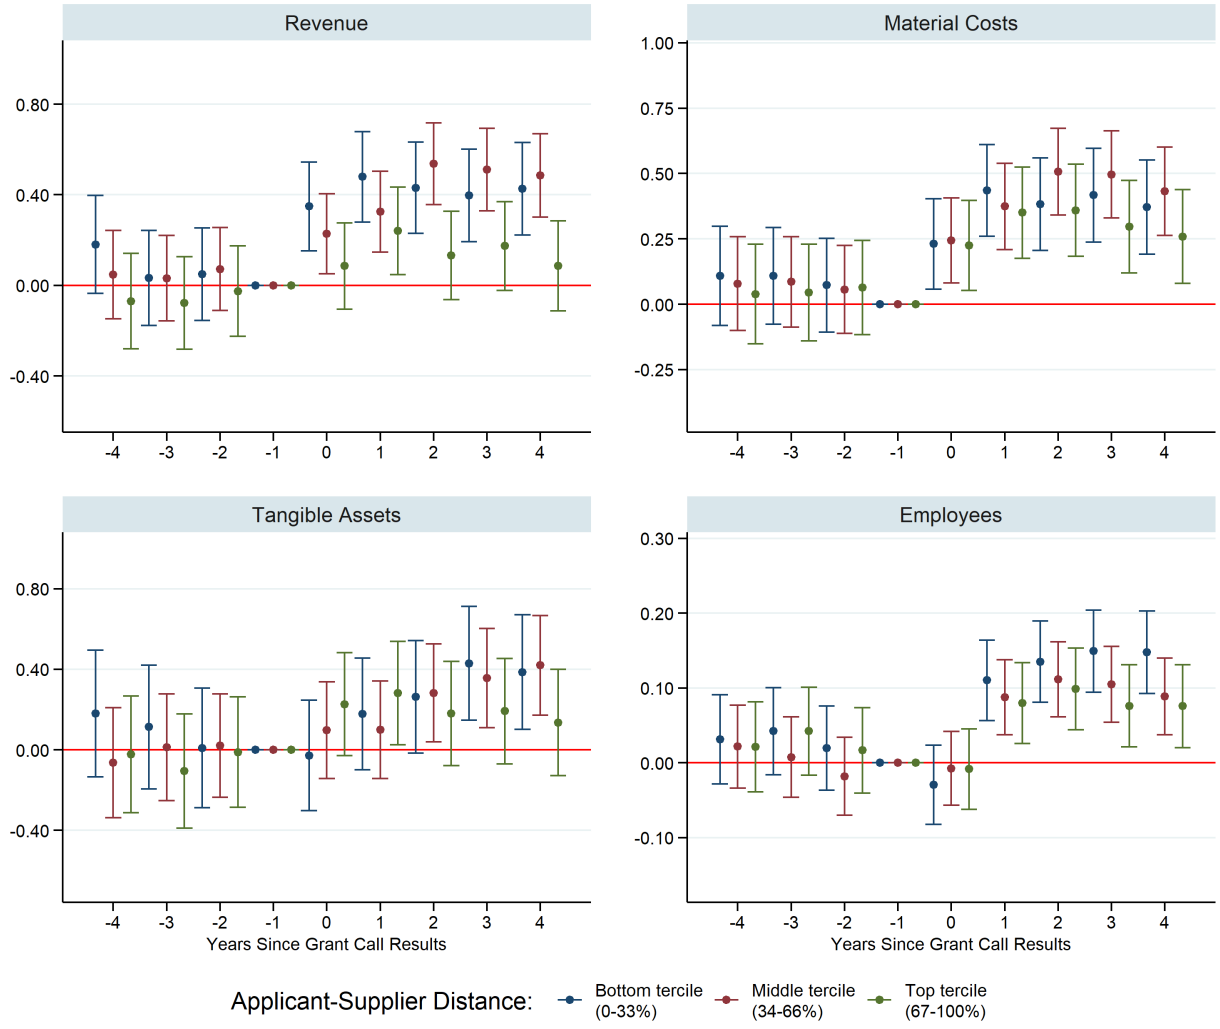


Figure 2: Heterogeneous supplier spillovers by applicant–supplier geographic distance

Note: Event-study spillover estimates (as in Figure 1) separately by terciles of straight-line distance between the municipalities of the applicant's and supplier's headquarters. Tercile definitions: bottom $<$ 23.55 km; middle 23.55–171.39 km; top $\geq$ 171.39 km. Whiskers denote 95% confidence intervals.

Supplier Financial Constraints. Whether upstream spillovers are amplified among financially constrained

suppliers is informative about the channel through which grant-induced demand propagates. Financial constraints arise when firms face limited or expensive external financing, so that internal cash flow and existing borrowing capacity bind their ability to expand production, hire workers, or invest in capacity. Under standard models of financing frictions, a positive demand shock from a creditworthy buyer is particularly valuable for a constrained supplier because it generates incremental cash flow that can be used to fund further expansion without recourse to costly external finance. If grant-induced demand operates partly through this mechanism, we should observe disproportionately large supplier responses among the most constrained firms, and the dynamic profile of these responses should reflect the gradual relaxation of binding liquidity constraints as cash flow accumulates over time.

We measure supplier financial constraints using the Omega Score of Altman et al. (2023), a granular default-risk metric that combines financial-statement information with payment-behavior and operational indicators and is calibrated and validated for Croatia. The Omega Score takes lower (more negative) values for firms with higher predicted likelihood of default, and we treat firms with higher default risk as more financially constrained on the basis that the same observable conditions that elevate default probability—weaker liquidity, higher leverage, payment delays, and tighter operating margins—also limit a firm's ability to access external finance on favorable terms. We split treated suppliers into terciles of the Omega Score; in Figure 3, the bottom tercile contains the least constrained suppliers and the top tercile contains the most constrained suppliers.

Pre-award coefficients are small and statistically indistinguishable from zero across all terciles, supporting the parallel-trends assumption. Post-award responses are sharply heterogeneous and monotone in financial constraints. The most constrained tercile exhibits the largest immediate response: revenue rises by 0.489 log points at $k = 0$ and remains in the 0.68–0.76 range through $k = 4$, with material costs following an essentially identical trajectory. The least constrained tercile shows substantially smaller and more transient effects, with revenue gains peaking around 0.13–0.14 log points and attenuating by $k = 4$. The middle tercile lies between these extremes. Capital and labor adjustments are also more pronounced for the most constrained tercile and materialize with the expected lag relative to the immediate scale response. Tables A7 and A8 report the full set of coefficients.

The immediate and disproportionately large joint increases in revenue and material costs among the most constrained suppliers, followed by gradual expansion of capital and employment, are consistent with the financing-frictions interpretation outlined above: a sustained demand expansion from grant-funded buyers improves cash flow and capacity utilization at constrained suppliers, relaxing binding liquidity constraints

and enabling subsequent investment and hiring that less constrained suppliers can already finance through other channels. The monotone pattern across the three terciles, with effect sizes rising as we move from least to most constrained, provides additional support for this interpretation: the more binding the ex ante credit friction, the larger the relative gain from buyer-driven cash flow.

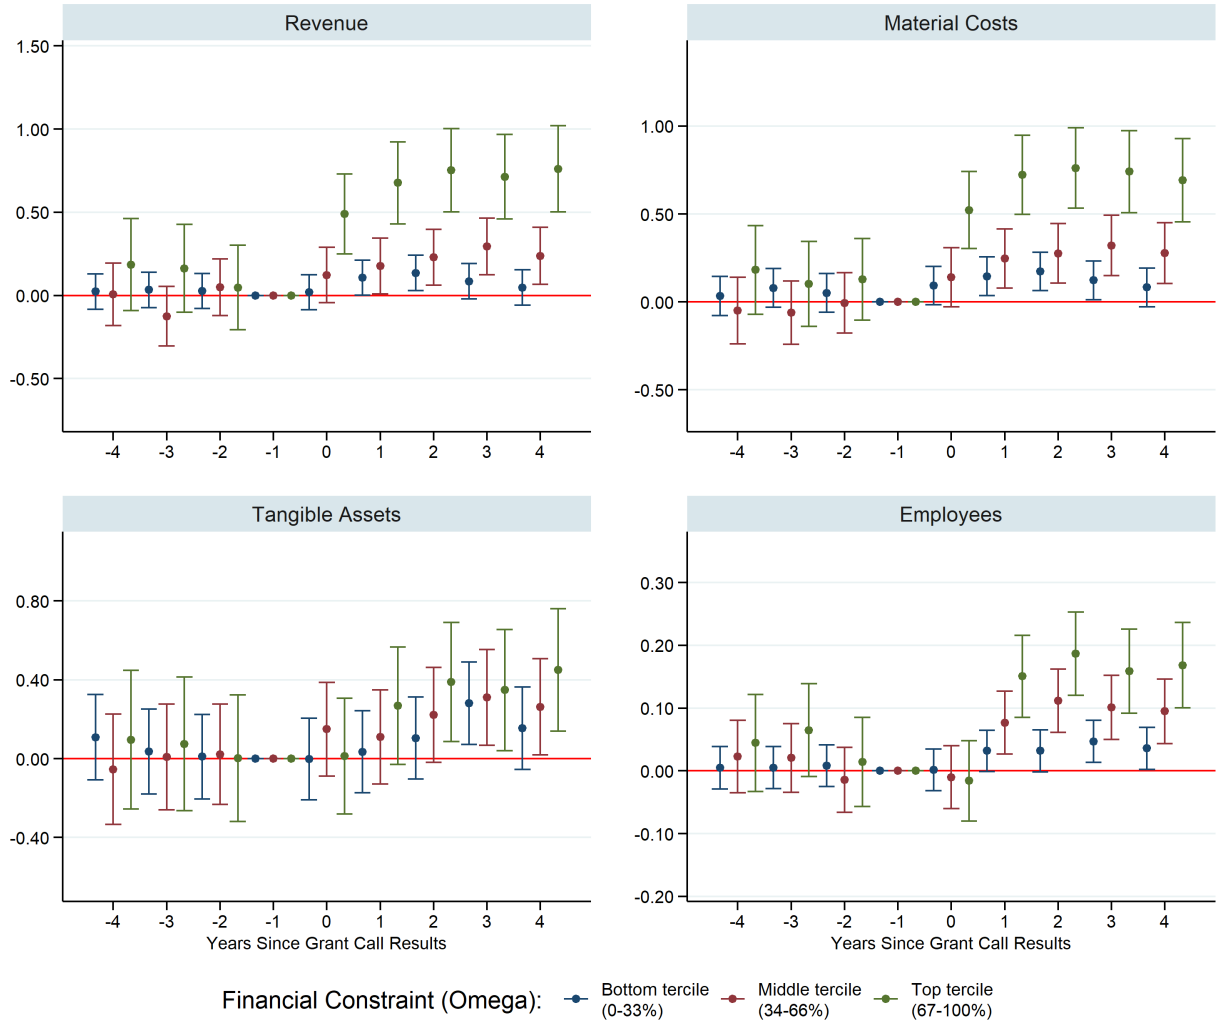


Figure 3: Heterogeneous supplier spillovers by supplier financial constraints

Note: Event-study spillover estimates (as in Figure 1) separately by terciles of the supplier Omega Score (Altman et al., 2023), used as a proxy for financial constraints. The bottom tercile (Omega $\geq$ 24.37) contains the least constrained suppliers; the top tercile (Omega $<$ $-$76.01) contains the most constrained suppliers. The middle tercile contains values between $-$24.37 and $-$76.01. Whiskers denote 95% confidence intervals.

Supplier Size. Whether spillovers concentrate among small suppliers or extend across the size distribution is informative about both the mechanism and the policy incidence. A pure mechanical scaling argument predicts that a given absolute increase in buyer demand represents a larger proportional shock for smaller suppliers,

who should therefore exhibit larger percentage adjustments. But if spillovers also reach larger suppliers, the policy footprint extends beyond the smallest firms in the network and the distributional implications of innovation policy correspondingly broaden. Figure 4 examines this margin by splitting treated suppliers into terciles of pre-award revenue. The revenue thresholds (€89,440 and €381,046) imply that all three terciles consist of small firms in absolute terms; for reference, the median recipient in our sample has pre-award revenue of approximately €2.3 million (Table 2), meaning that even suppliers in the top tercile operate at roughly one-sixth of the typical recipient's baseline scale. The exercise therefore reflects heterogeneity within a small-supplier environment rather than differences between small and large enterprises, and the large percentage effects documented below are in part a mechanical consequence of this size asymmetry.

Pre-award coefficients are balanced across terciles, supporting parallel trends. Post-award responses are larger in proportional terms for smaller suppliers, consistent with the mechanical scaling channel. Spillovers are not, however, confined to the smallest suppliers: economically meaningful responses appear in all three terciles, indicating that grant-induced upstream demand reaches across the supplier size distribution rather than concentrating exclusively at the bottom. Tangible asset and employment dynamics are more gradual than revenue and material costs across all size groups, consistent with the adjustment-friction interpretation already documented in the baseline analysis.

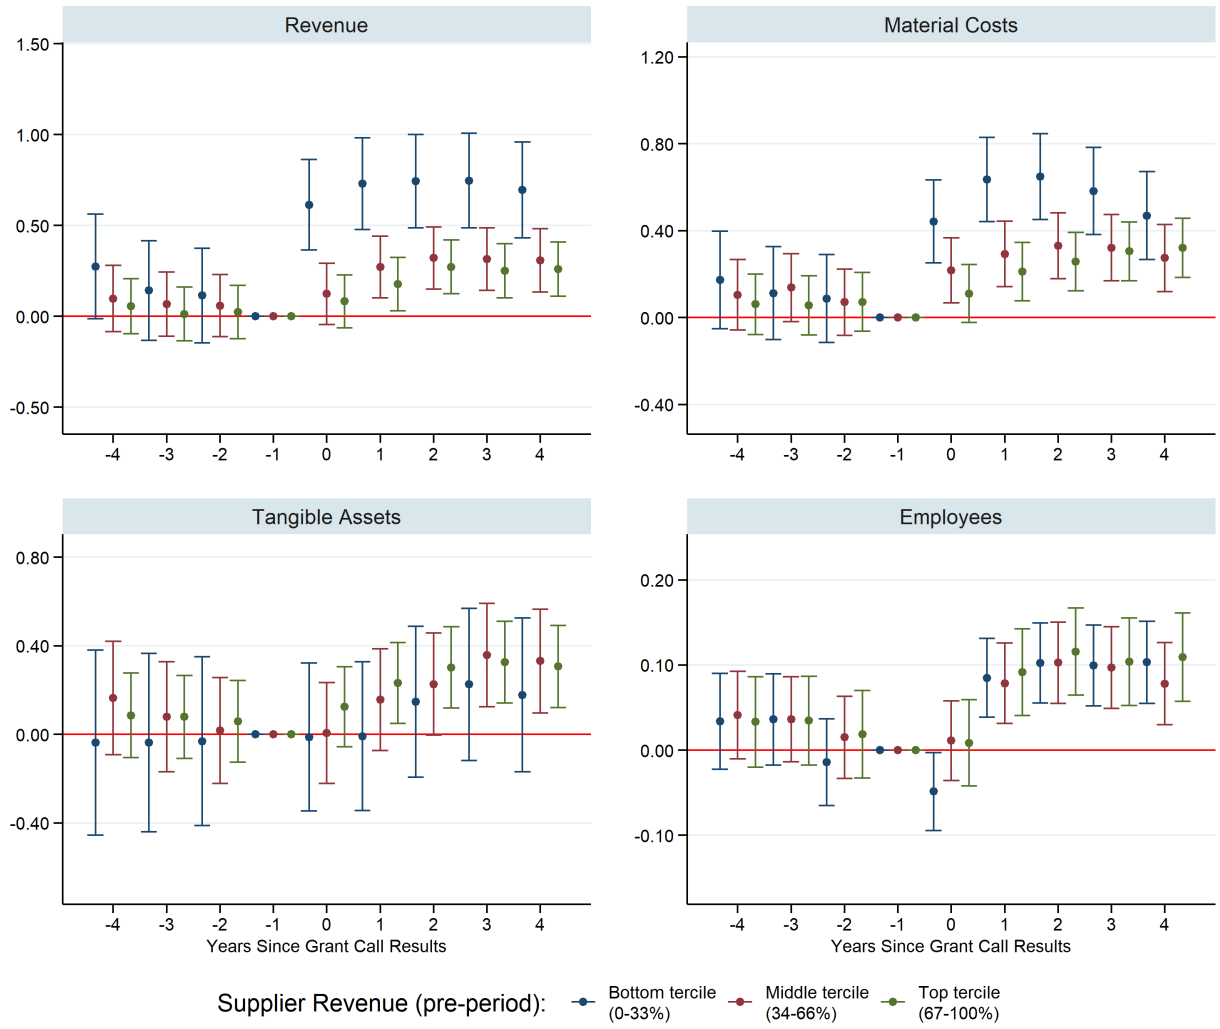


Figure 4: Heterogeneous upstream spillovers by supplier size

Note: Event-study spillover estimates (as in Figure 1) separately by terciles of baseline supplier size, measured by pre-award revenue. Tercile definitions: bottom $<$ €89,440; middle €89,440–€381,046; top $\geq$ €381,046. Whiskers denote 95% confidence intervals.

Network Centrality. A central prediction of the production networks literature is that the propagation and aggregate impact of firm-level shocks depends on the structural position of affected firms within the network (Acemoglu et al., 2012; Carvalho, 2014). We test this prediction directly by examining whether upstream spillovers vary with the network centrality of either the grant recipient or the supplier, using four measures that each capture a distinct dimension of network position. *Degree centrality* counts a firm's direct trading partners and reflects the breadth of its market relationships (Figures A3 and A7). *Betweenness centrality* measures how often a firm lies on the shortest path between other firms, capturing its role as a broker in the flow of goods and demand across the economy (Figures A1 and A5). *Closeness centrality* is the inverse of a

firm's average distance to all other firms, reflecting how quickly shocks originating elsewhere in the network can reach it (Figures A2 and A6). *Eigenvector centrality* weights a firm's connections by the centrality of its neighbors recursively ((Figures 5); Golub (2025) establishes that it is the limiting case of Katz-Bonacich centrality when indirect network effects receive maximum weight, making it the firm-level analogue of the influence vector that Acemoglu et al. (2012) show governs aggregate shock transmission (Figures A4 and 5).

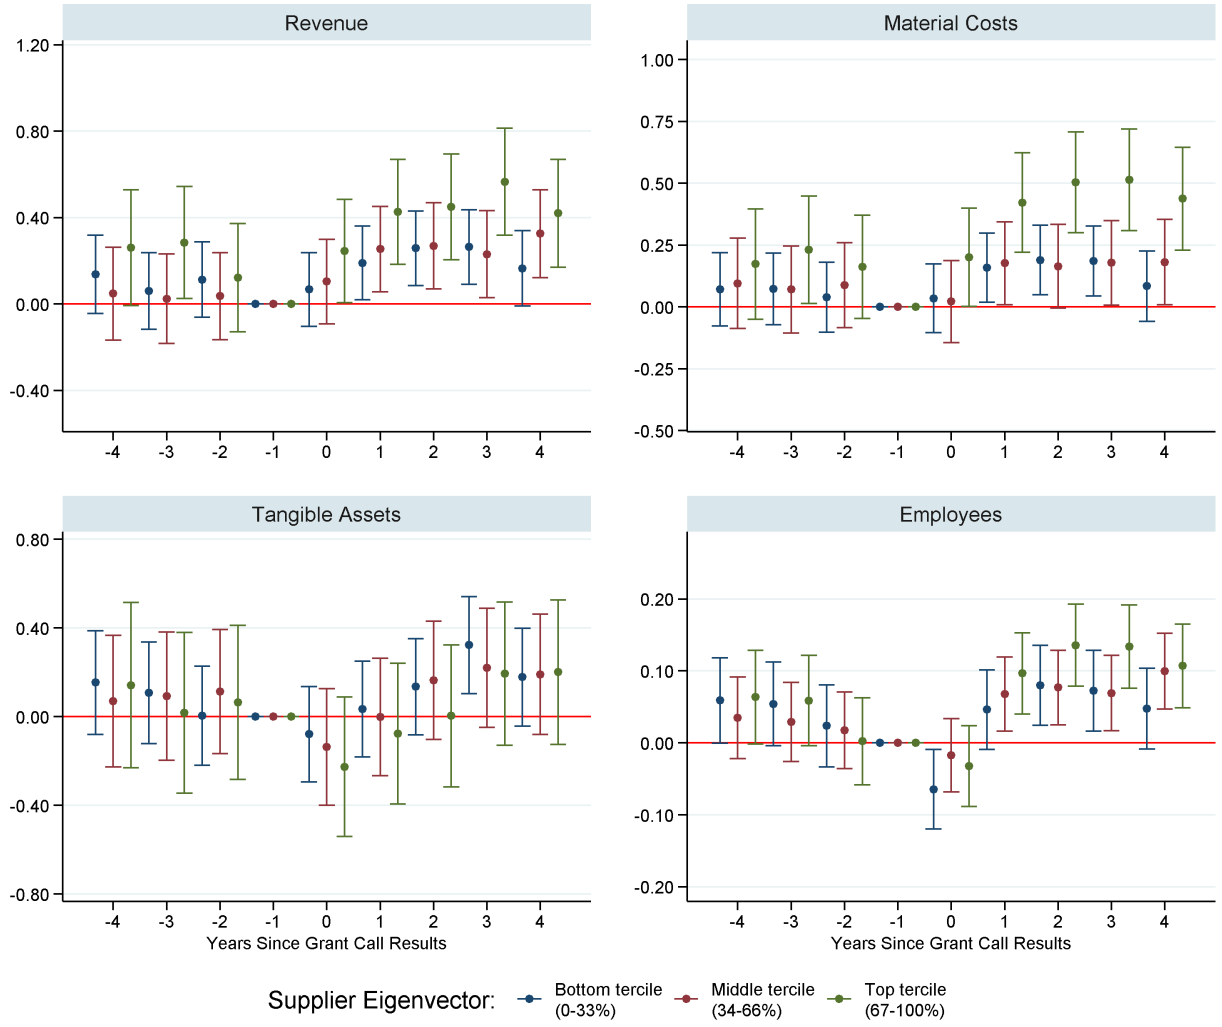


Figure 5: Upstream spillovers on suppliers, split by supplier eigenvector centrality

Note: Event-study estimates of supplier outcomes (revenue, material costs, tangible assets, and employment) as in Figure 1, estimated separately by terciles of the supplier's eigenvector centrality in the pre-award firm-to-firm transaction network. Eigenvector centrality weights a firm's connections by the centrality of its neighbors recursively, and corresponds to the limiting case of Katz-Bonacich centrality as indirect network effects receive maximum weight (Golub, 2025); it is the firm-level analogue of the influence vector that Acemoglu et al. (2012) show governs the aggregate transmission of shocks through input-output linkages. Suppliers in the top eigenvector tercile trade with other structurally important firms and are therefore best positioned to re-transmit grant-induced demand further through the network. The omitted reference period is year −1. Tercile definitions: bottom (0–33%), middle (34–66%), top (67–100%) of the pre-award supplier eigenvector centrality distribution. Whiskers denote 95% confidence intervals.

A striking and non-obvious pattern emerges. For applicant centrality, suppliers connected to the *least* central grant recipients exhibit the largest and most persistent post-award responses across revenue and material costs, with this ordering consistent across all four measures (Figures A3, A1, A2, and A4) and flat pre-award trends across terciles confirming it is not a pre-existing difference. This runs counter to the intuition that more prominent, better-connected recipients should generate larger spillovers, and is consistent with the finding in the supplier size analysis above that spillovers are proportionally larger for smaller firms: if more central applicants are also larger firms, the marginal demand shock from a given grant is smaller relative to their baseline scale, generating a weaker upstream stimulus per euro of funding (Criscuolo et al., 2019). The pattern for supplier centrality is sharply different and theoretically revealing. Spillovers concentrate among the most central suppliers across all four measures (Figures A7, A5, A6, and 5), with the top eigenvector and degree tercile exhibiting the largest and most persistent revenue, material cost, and employment responses, growing monotonically over the post-award period.

This is the firm-level empirical counterpart of the Acemoglu et al. (2012) prediction: suppliers that trade with other important firms are structurally positioned to re-transmit grant-induced demand further through the network, amplifying the initial shock beyond the immediate buyer–supplier pair. The degree result reinforces this through a complementary mechanism — suppliers with more direct trading partners can absorb demand expansions by leveraging broader capacity across their customer base and are more likely to themselves source from diverse upstream suppliers, facilitating further propagation. For betweenness and closeness, the directional pattern is consistent but somewhat less monotone, particularly for tangible assets. Together, the centrality results reveal an asymmetric role of network position in shock propagation: what matters is not the prominence of the node where the shock originates, but the prominence of the nodes through which it travels. Grant-induced demand propagates furthest and most persistently when it flows through structurally central suppliers, regardless of whether the initial treated buyer is itself centrally positioned. This finding complements the financial constraints and size results above — central suppliers tend to be larger and less financially constrained, yet the centrality gradient survives controlling for these characteristics implicitly through the tercile splits, suggesting that network position captures a distinct margin of propagation capacity beyond scale and liquidity alone. For policy, this implies that the network multiplier of a grant depends less on selecting prominent recipients and more on whether those recipients are connected to suppliers that are themselves well-embedded in the production network.

5.5 Sensitivity Analysis

This section evaluates the robustness of our results along four dimensions. We first vary the matched-sample specification used to estimate direct effects on applicants, and then do the same for the supplier event study. We then complement the matched-sample evidence on direct effects with two cutoff-based identification strategies that exploit the program's evaluation scores: a panel-based difference-in-discontinuities specification and cross-sectional local randomization inference. Finally, we extend the cutoff-based logic to the upstream side by reporting a network-exposure design that uses pre-award buyer–supplier links combined with the same scoring variation. The four exercises taken together test whether our main estimates are sensitive to (i) the richness of the matched-sample specification, (ii) the source of identifying variation, and (iii) whether evidence for upstream propagation survives at the cutoff.

Matched-sample alternatives: direct effects. We first examine whether the matched-sample direct-effect estimates of Table 1 are robust to alternative specification choices on the same matched applicant sample. Figure A8 compares the baseline specification with and without firm fixed effects. Removing firm fixed effects yields systematically larger post-award effects, consistent with residual time-invariant differences between funded and non-funded applicants that inflate estimates when they are not absorbed. Figure A9 then varies the richness of the matching procedure itself, comparing the baseline (which conditions on roughly sixty pre-treatment characteristics) with more parsimonious alternatives. Coarser matching produces larger estimated impacts, consistent with reduced control for positive selection. Both exercises indicate that our baseline direct-effect estimates are conservative relative to plausible alternatives, and that further conditioning—through firm fixed effects, richer matching, or both—reduces rather than amplifies the estimated treatment effect.

Matched-sample alternatives: supplier spillovers. We next apply the same robustness exercises to the supplier event study reported in Figure 1 and Table A4. Figure A10 compares the baseline supplier specification with and without supplier fixed effects. Point estimates and dynamic profiles are very similar across specifications, while supplier fixed effects tighten confidence intervals. This stability reduces concerns that the estimated spillovers reflect time-invariant supplier characteristics correlated with exposure to funded buyers. Figure A11 re-estimates the supplier event study using a substantially more parsimonious matching design at the buyer level. Estimated spillovers remain close to the baseline, which is consistent with the interpretation that conditional on pre-award buyer–supplier links and observed controls, grant receipt is determined by the buyer's evaluation outcome and is therefore plausibly orthogonal to most supplier characteristics. As a third check, Figure A12 addresses whether spillovers might be driven by a small set

of exceptional projects rather than by grant receipt as such, by re-estimating the supplier event study restricting attention to suppliers linked to the top 25% and top 50% best-graded funded projects. Effects are essentially identical across these score restrictions, providing no evidence that the upstream results are driven primarily by project quality at the top of the score distribution.

Cutoff-based sensitivity: direct effects. We complement the matched-sample direct-effect estimates with two cutoff-based identification strategies that exploit the program's evaluation scores as a source of quasi-experimental variation in grant receipt. The first is a panel-based difference-in-discontinuities specification that compares marginally eligible to marginally ineligible applicants over multiple post-award years, using firm fixed effects to absorb time-invariant heterogeneity (Appendix B). The second is a cross-sectional local randomization inference design (Cattaneo et al., 2015) that delivers exact finite-sample p-values within a window where covariate balance can be validated directly (Appendix C). Both designs are inherently local and operate on a substantially smaller effective sample of marginal applicants, but they test the same underlying causal claim from a different identification angle. The difference-in-discontinuities estimates are consistent in sign and broadly similar in magnitude to the matched-sample estimates of Table 1. The local randomization analysis detects a positive employment effect that strengthens monotonically across post-treatment horizons, from approximately 14% at $t = +1$ to 34% at $t = +4$, and remains statistically significant at the 10% level under the difference-in-means test at every horizon. Effects on revenue and tangible assets are positive across horizons and reach magnitudes by $t = +4$ that are comparable to the matched-sample estimates, but they remain statistically imprecise in the cross-section, reflecting the limited power of finite-sample exact inference at $N = 92$. The convergence of two methodologically distinct cutoff-based designs with the matched-sample estimates supports the interpretation that grant receipt causally raises recipient activity.

Cutoff-based sensitivity: upstream spillovers. Finally, we extend the cutoff-based logic from direct effects to upstream spillovers using a network-exposure design that combines pre-award buyer–supplier links with the same scoring variation. For each supplier, we measure pre-treatment exposure as the share of pre-award sales accounted for by applicants in the cutoff-based sample, and we identify spillovers from the interaction of exposure with the post-award period. The design exploits cross-sectional variation in suppliers' pre-existing trading relationships rather than the matched-sample event-study structure, and therefore relies on a different identification logic from the baseline supplier analysis. The first-order coefficient on exposure interacted with the post-award period is positive and statistically significant for revenue, material costs, tangible assets, and employment, with magnitudes broadly consistent with the matched-sample event-study estimates of Section 5.3. We report this network-exposure design and its results in Appendix D; in addition, we find little evidence of propagation beyond directly connected suppliers, with estimated second-order effects

that are statistically indistinguishable from zero. The cutoff-based supplier evidence therefore reinforces the matched-sample finding that spillovers operate through directly connected suppliers, while remaining silent on more distant network propagation given the limited statistical power available at the cutoff.

Taken together, the four sensitivity exercises support the interpretation that the direct and upstream effects we document arise from grant-induced scale expansion at recipients propagating through pre-existing buyer–supplier relationships, rather than from selection on observables, residual unobserved heterogeneity, or a small set of exceptional projects.

6 Cost-Benefit Analysis and Policy Implications

This section provides a transparent back-of-the-envelope cost–benefit exercise for Croatia's flagship EU-funded R&D grant instrument. The purpose is not to monetize innovation outputs in full, but to translate our estimated effects into two policy-relevant metrics that can be computed with minimal assumptions: the cost per job created and the VAT generated. Table 2 implements this accounting using the program's total public spending envelope of €199.82 million and combines it with baseline scale measures for recipients and suppliers and our estimated proportional effects on employment and revenue.

Table 2: R&D grant scheme: cost per job and VAT generated

	Direct	Upstream	Total	Change (%)
Aggregate R&D grant value (mil. €)			199.82	
Recipients	153	3370		
Panel A: Cost per job				
Median employees	40	4		
Employee estimate	0.076	0.101		
Jobs created	463	1359	1823	393.3%
Cost per job	€431,000	€147,000	€110,000	-74.5%
Panel B: Cost per one euro of VAT				
Median Revenue	2.340	0.137		
Revenue estimate	0.157	0.346		
VAT generated in mil. €	14.052	39.893	53.945	384%
Cost per one euro of VAT	14.220	5.009	3.704	-74%

Note: This table presents a back-of-the-envelope cost–benefit accounting based on the point estimates of the estimated proportional effects of the R&D grant on recipients (direct; Table 1) and their suppliers (upstream; Table A4) employment. Baseline firm characteristics are measured in the $t-1$ (pre-award) period. To remain conservative, we scale effects using the median of the baseline distributions of employment and revenue for recipients and suppliers. Jobs created are computed by applying the estimated employment effects to baseline employment levels; VAT generated is computed analogously from estimated revenue effects and baseline revenue. Calculations are illustrative and intended to quantify orders of magnitude rather than provide a full welfare evaluation.

The first panel of Table 5 reports implied job creation among (i) grant recipients (direct), (ii) upstream suppliers (indirect), and (iii) the combined total. Using recipients' baseline employment (median 40 employees) and the estimated proportional employment response, the table implies 463 direct jobs created among recipients. Applying the same logic to suppliers, using suppliers' baseline employment (median 4 employees) and the estimated proportional supplier employment response, implies 1,359 upstream jobs. Aggregating these effects yields 1,823 total jobs, so incorporating supplier responses increases the implied employment impact to about four times the recipient-only effect (1,823 versus 463).

This scaling has immediate implications for cost-effectiveness. Under a recipient-only accounting, the implied cost per job is about €431,000. When supplier responses are included, the implied cost per job falls to about €110,000, a reduction of roughly 74.5%. The key point is mechanical but important. Because supplier effects are large relative to recipient effects, ignoring production networks materially overstates the cost per job created by the program. (For completeness, the upstream-only cost per job is about €147,000.)

The second panel of Table 5 provides an analogous fiscal accounting based on revenue. Using recipients' baseline revenue (median €2.34 million) and the estimated proportional revenue response, the table implies €14.052 million in VAT generated from recipient-side activity. Incorporating upstream revenue responses raises implied VAT generation to €53.945 million in total, with €39.893 million attributed to suppliers. Expressed as "*cost per euro of VAT*," the implied cost falls from 14.220 under the recipient-only calculation to 3.704 when suppliers are included.

Two implications follow. First, evaluations that focus exclusively on direct recipients can substantially understate the footprint of innovation spending when production networks transmit grant-induced expansions beyond treated firms. In this setting, accounting for upstream spillovers changes both the implied magnitude of program effects and the implied cost-effectiveness by a factor close to four. Second, because the instrument requires substantial private co-financing, the policy should be interpreted as leveraging private resources rather than operating as a pure transfer. This makes it particularly important to measure incidence beyond the recipient firm, since a meaningful share of the induced activity is realized among upstream suppliers that are not directly targeted by the policy but benefit through increased demand.

We emphasize that this accounting is deliberately conservative. It focuses on observable expansions in activity and associated fiscal flows, and it does not attempt to value longer-run innovation outcomes, knowledge spillovers, or general equilibrium adjustments. Even under this conservative approach, the main lesson is clear. Evaluations that ignore production networks can materially mismeasure the benefits of innovation policy and can lead to misleading conclusions about program returns.

7 Discussion

This paper shows that R&D grants can generate economically meaningful spillovers through production networks, and that these spillovers materially change how innovation policy should be evaluated. The classic rationale for R&D subsidies is that private firms underinvest in innovation because knowledge is non-rival and imperfectly excludable, so private returns fall short of social returns (Arrow, 1962). Most empirical evaluations, however, focus on outcomes at the recipient firm, including evidence on grant programs and related innovation support (Bronzini and Iachini, 2014; Howell, 2017; Dechezleprêtre et al., 2023). A large complementary literature emphasizes knowledge spillovers operating through technological proximity, geography, and worker mobility (Jaffe et al., 1993; Bloom et al., 2019; Lucking et al., 2019; Myers and Lanahan, 2022). Our contribution is to document a distinct channel that is quantitatively first-order. Innovation spending can propagate through buyer–supplier linkages, generating upstream effects at firms that never receive funding and that are typically smaller than recipients. This transmission mechanism is not primarily about the diffusion of ideas. It is about the network consequences of grant-induced scale expansion and intermediate input demand.

The magnitude of network amplification matters directly for program evaluation. Table 2 translates our preferred estimates into a conservative accounting based on jobs and VAT. Counting only direct effects at recipients implies 463 jobs, while incorporating supplier responses adds 1,359 jobs, for a total of 1,823 jobs. This accounting reduces the implied cost per job from about €431,000 under a recipient-only evaluation to about €110,000 when supplier responses are included. The VAT-based exercise displays a similar scaling. Recipient-side activity implies €14.052 million in VAT, while incorporating upstream responses raises implied VAT to €53.945 million, and the implied cost per euro of VAT falls from 14.220 to 3.704. The broader point is that recipient-only evaluations can substantially understate program returns when production networks transmit grant-induced expansions beyond treated firms. This implication complements quasi-experimental evidence that R&D support often produces sizable expansions in firm activity at the margin, including revenue and employment effects, in settings such as the U.S. SBIR program (Howell, 2017) and European R&D grant environments (Santoleri et al., 2024). The cross-country consistency suggests that financing frictions limiting innovation and scale are widespread and that public support can partially relieve them.

The dynamic patterns in supplier outcomes also align closely with the production networks literature. Theory emphasizes that idiosyncratic shocks can have aggregate consequences when production is networked (Acemoglu et al., 2012; Carvalho, 2014). Empirical work has traced propagation of negative shocks, such as disruptions and financial distress, through supplier relationships (Barrot and Sauvagnat, 2016; Boehm et al.,

2019; Carvalho et al., 2021). Our setting isolates a policy-induced positive shock and shows that propagation works through the same input-output structure. In this respect, our findings are especially consistent with Acemoglu et al. (2016), who emphasize that demand-side shocks propagate predominantly upstream. The R&D grants we study are best interpreted as demand-like at the horizons we observe. Grants inject resources and require private co-financing, which mechanically supports higher spending on labor, intermediate inputs, and investment. That logic predicts upstream propagation to suppliers rather than primarily downstream benefits to customers, which is precisely what we observe. Suppliers expand scale immediately, with revenue and material costs rising sharply and persistently, while employment and tangible assets adjust more gradually in a pattern consistent with hiring and investment frictions.

The heterogeneity evidence strengthens this interpretation. Spillovers are more pronounced among financially constrained suppliers and are present throughout the supplier size distribution. This pattern is consistent with a mechanism in which grant-induced demand from treated buyers relaxes upstream liquidity constraints through improved cash flow and utilization. Constrained suppliers expand quickly along variable-input margins, and then adjust labor and capital as financing, investment, and hiring frictions are overcome. Less constrained suppliers plausibly smooth shocks more easily, yielding smaller proportional adjustments.

These results also inform debates on place-based policy and the spatial incidence of industrial interventions. Standard concerns emphasize that local programs may leak, diluting benefits for targeted places (Neumark and Simpson, 2015; Bartik, 2020). Our evidence supports a more nuanced view. Spillovers are stronger for geographically proximate buyer–supplier relationships, but they are not confined to local links. Revenue effects are concentrated among nearby and intermediate-distance suppliers, while material costs and employment adjust persistently even for distant suppliers. Tangible asset accumulation is more geographically concentrated than variable-input responses, consistent with spatial frictions affecting investment more than short-run scale adjustments. This implies that network connectivity, not geography alone, is central for predicting where the benefits of innovation spending accrue. The comparison to Lerche (2025) is instructive. While some place-based interventions generate highly localized spillovers, our results indicate materially less spatial confinement and a larger aggregate upstream footprint relative to the direct effect, which highlights the importance of mapping firm-to-firm linkages when evaluating incidence.

Our cost-benefit calculations are deliberately conservative. They focus on measurable expansions in activity and associated fiscal flows and do not attempt to value longer-run innovation outputs or broader knowledge spillovers that motivate R&D subsidies (Arrow, 1962; Bloom et al., 2019). Nor do they incorporate general equilibrium adjustments that could arise at larger policy scale. Even under this conservative accounting,

the conclusion is clear. Innovation policy operates in networked economies, and production networks can amplify its effects substantially. Evaluations that ignore these linkages risk materially understating program impacts and drawing misleading conclusions about policy returns.

A natural question is whether our magnitudes generalize. Croatia is a high-income, small open EU economy that the European Innovation Scoreboard (European Commission, 2025) classifies as an "Emerging Innovator": average innovative capacity below the EU frontier and meaningful frictions in private credit markets. This profile is shared by a large share of EU member states and by most middle-income economies outside the advanced-economy core, and is more representative of the population of countries that actually deploy R&D grants than the frontier economies that dominate the existing evaluation literature. Within the framework of Section 3, the joint (A_i, λ_i) position of the Croatian economy is favorable for grant additionality: financial constraints are tight enough to bind among recipients, while average innovative capacity is sufficient for a non-trivial share of funded projects to reach the market. Consistent with this, Santoleri et al. (2024) find that grant effects on firm activity are larger in economies and regions with lower economic and financial development, and stronger for younger and smaller firms. Importantly, however, the relevance of R&D grants is not confined to constrained settings: Howell (2017) documents large positive effects of US SBIR grants on recipient firms despite the United States having one of the deepest private capital markets in the world, indicating that even in frontier economies a meaningful share of recipients translate grant funds into real activity rather than absorbing them through crowding-out.

The propagation mechanism is more portable than the precise magnitudes. Pre-existing buyer–supplier relationships are a feature of any economy with a functioning production network, and the logic that grant-induced scale expansion at recipients generates additional intermediate input demand—and that this demand flows through firms' existing procurement relationships—does not depend on the Croatian institutional environment. What our setting contributes is the comprehensive transaction-level data that resolves these flows at firm-to-firm level, which is rare in the cross-country administrative data landscape. The same upstream channel should operate wherever R&D grants generate real expansions at recipient firms, including in advanced economies with deep capital markets where Howell (2017) and others have documented sizable direct effects. Where it does, evaluations restricted to recipient firms will understate the broader returns to R&D policy.

8 Conclusion

This paper shows that R&D grants generate substantial spillovers through production networks. Grants lead to sizable expansions among recipient firms, increasing revenue, material costs, employment, and tangible assets. These direct effects propagate upstream. Suppliers connected to funded applicants experience persistent post-award increases in activity, particularly in revenue and intermediate input use, alongside more modest but durable employment gains and gradual capital accumulation. When these upstream responses are incorporated, a simple aggregation suggests that the total effect of the program is roughly four times larger than what is captured by recipient-only evaluations.

Several features shape how spillovers propagate. First, production linkages matter more than geography alone. Spillovers are strongest for geographically proximate buyer–supplier relationships, but they remain statistically and economically meaningful even for suppliers located far from the recipient firm. This pattern implies that spatial proximity facilitates transmission, yet network connectivity, rather than physical distance per se, ultimately determines exposure. Second, spillovers are amplified among financially constrained suppliers, consistent with the idea that demand shocks from better-funded customers relax binding liquidity constraints upstream. Third, spillovers accrue disproportionately to firms that are smaller than grant recipients on average, highlighting that innovation policy can affect segments of the firm size distribution that are not directly targeted by subsidies.

These findings have direct implications for the evaluation and design of innovation policy. Standard cost–benefit analyses that focus exclusively on recipients substantially understate program impacts when production networks transmit policy-induced expansions beyond treated firms. More broadly, the results caution against evaluating innovation subsidies in isolation from the economic networks in which firms operate. Policies that stimulate activity at one node of the network can generate effects that are both quantitatively large and widely distributed across connected firms.

Methodologically, the paper demonstrates how combining matched-sample designs with administrative transaction data allows evaluators to trace causal spillovers through production networks in settings where clean experimental variation is limited. Regression-discontinuity evidence based on multi-criteria evaluation thresholds provides complementary sensitivity checks, but the core contribution lies in showing how pre-existing buyer–supplier links can be leveraged to study policy propagation beyond direct recipients. This approach is applicable to a broad class of public interventions, particularly in environments where multiple programs operate simultaneously and firms are embedded in dense production networks.

Several directions remain for future research. An important open question concerns the long-run persistence of network spillovers beyond the medium-run horizon we observe, as well as their interaction with knowledge spillovers and reallocation across suppliers. Another avenue relates to optimal policy design, including whether targeting firms with particular network positions or supplier structures can enhance aggregate returns. Future work could also extend the measurement framework in several directions: incorporating product-level transaction data to disentangle price and quantity margins in supplier responses, tracing propagation to recipients' customers and competitors, and extending the analysis across borders to capture indirect effects that our domestic VAT records leave unmeasured. Nonetheless, the central conclusion is clear. Production networks amplify the effects of innovation policy. Ignoring these linkages leads to systematic understatement of policy impacts and risks underinvestment in innovation support.

Literature

Acemoglu, Daron, Ufuk Akcigit, and William Kerr, "Networks and the macroeconomy: An empirical exploration," *Nber macroeconomics annual*, 2016, *30* (1), 273–335.

_ , Vasco M Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi, "The network origins of aggregate fluctuations," *Econometrica*, 2012, *80* (5), 1977–2016.

Aghion, Philippe, Nicholas Bloom, Richard Blundell, Rachel Griffith, and Peter Howitt, "Competition and Innovation: An Inverted-U Relationship," *The Quarterly Journal of Economics*, 2005, *120* (2), 701–728.

Alfaro-Urena, Alonso, Isabela Manelici, and Jose P Vasquez, "The effects of joining multinational supply chains: New evidence from firm-to-firm linkages," *The Quarterly Journal of Economics*, 2022, *137* (3), 1495–1552.

Allen, Treb and Costas Arkolakis, "Trade and the Topography of the Spatial Economy," *Quarterly Journal of Economics*, 2018, *133* (3), 1085–1140.

Altman, Edward I, Marco Balzano, Alessandro Giannozzi, and Stjepan Srhoj, "Revisiting SME default predictors: The omega score," *Journal of Small Business Management*, 2023, *61* (6), 2383–2417.

Arrow, Kenneth J, "Economic welfare and the allocation of resources for invention," *The rate and direction of inventive activity: Economic and social factors*, 1962, pp. 609–626.

Baker, Andrew, Brantly Callaway, Scott Cunningham, Andrew Goodman-Bacon, and Pedro H. C. Sant'Anna, "Difference-in-Differences Designs: A Practitioner's Guide," *Journal of Economic Literature*, 2026, *64* (2), 498–557.

Baqaee, David Rezza and Emmanuel Farhi, "The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten's Theorem," *Econometrica*, 2019, *87* (4), 1155–1203.

Barattieri, Alessandro, Matteo Cacciatore, and Nora Traum, "Estimating the effects of government spending through the production network," Technical Report, National Bureau of Economic Research 2023.

Barrot, Jean-Noël and Julien Sauvagnat, "Input specificity and the propagation of idiosyncratic shocks in production networks," *The Quarterly Journal of Economics*, 2016, *131* (3), 1543–1592.

Bartik, Timothy J, "Using place-based jobs policies to help distressed communities," *Journal of Economic Perspectives*, 2020, *34* (3), 99–127.

Bigio, Saki and Jennifer La'o, "Distortions in production networks," *The quarterly journal of economics*, 2020, *135* (4), 2187–2253.

Bloom, Nicholas, John Van Reenen, and Heidi Williams, "A toolkit of policies to promote innovation," *Journal of Economic Perspectives*, 2019, *33* (3), 163–184.

Boehm, Christoph E, Aaron Flaaen, and Nitya Pandalai-Nayar, "Input linkages and the transmission of shocks: Firm-level evidence from the 2011 Tōhoku earthquake," *Review of Economics and Statistics*, 2019, *101* (1), 60–75.

Bronzini, Raffaello and Eleonora Iachini, "The impact of R&D subsidies on firm innovation," *Research Policy*, 2014, *43* (2), 442–457.

Cai, Jing and Adam Szeidl, "Indirect effects of access to finance," *American Economic Review*, 2024, *114* (8), 2308–2351.

Callaway, Brantly and Pedro H. C. Sant'Anna, "Difference-in-differences with variation in treatment timing," *Journal of Econometrics*, 2021, *225* (2), 254–277.

Carvalho, Vasco M, "From micro to macro via production networks," *Journal of Economic Perspectives*, 2014, *28* (4), 23–48.

_ , Makoto Nirei, Yukiko U Saito, and Alireza Tahbaz-Salehi, "Supply chain disruptions: Evidence from the great east japan earthquake," *The quarterly journal of economics*, 2021, *136* (2), 1255–1321.

Cattaneo, Matias D., Brigham R. Frandsen, and Rocio Titiunik, "Randomization Inference in the Regression Discontinuity Design: An Application to Party Advantages in the U.S. Senate," *Journal of Causal Inference*, 2015, *3* (1), 1–24.

_ , Rocio Titiunik, and Gonzalo Vazquez-Bare, "Inference in Regression Discontinuity Designs under Local Randomization," *The Stata Journal*, 2016, *16* (2), 331–367.

Criscuolo, Chiara, Ralf Martin, Henry G Overman, and John Van Reenen, "Some causal effects of an industrial policy," *American Economic Review*, 2019, *109* (1), 48–85.

David, Paul A., Bronwyn H. Hall, and Andrew A. Toole, "Is Public R&D a Complement or Substitute for Private R&D? A Review of the Econometric Evidence," *Research Policy*, 2000, *29* (4-5), 497–529.

Dechezleprêtre, Antoine, Elias Einiö, Ralf Martin, Kieu-Trang Nguyen, and John Van Reenen, "Do tax incentives increase firm innovation? An RD design for R&D, patents, and spillovers," *American Economic Journal: Economic Policy*, 2023, *15* (4), 486–521.

Demir, Banu, Beata Javorcik, Tomasz K Michalski, and Evren Ors, "Financial constraints and propagation of shocks in production networks," *Review of Economics and Statistics*, 2024, *106* (2), 437–454.

Dhyne, Emmanuel, Ayumu Ken Kikkawa, Magne Mogstad, and Felix Tintelnot, "Trade and domestic production networks," *The Review of Economic Studies*, 2021, *88* (2), 643–668.

Dimos, Christos and Geoff Pugh, "The effectiveness of R&D subsidies: A meta-regression analysis of the evaluation literature," *Research Policy*, 2016, *45* (4), 797–815.

Dvouletỳ, Ondřej, Stjepan Srhoj, and Smaranda Pantea, "Public SME grants and firm performance in European Union: A systematic review of empirical evidence," *Small Business Economics*, 2021, *57* (1), 243–263.

European Commission, "European Innovation Scoreboard 2025," 2025.

Fazzari, Steven M., R. Glenn Hubbard, and Bruce C. Petersen, "Financing Constraints and Corporate Investment," *Brookings Papers on Economic Activity*, 1988, *1988* (1), 141–206.

Furman, Jeffrey L., Michael E. Porter, and Scott Stern, "The Determinants of National Innovative Capacity," *Research Policy*, 2002, *31* (6), 899–933.

Gabaix, Xavier, "The Granular Origins of Aggregate Fluctuations," *Econometrica*, 2011, *79* (3), 733–772.

Golub, Benjamin, "Eigenvalues in microeconomics," *arXiv preprint arXiv:2502.12309*, 2025.

Hall, Bronwyn H and Josh Lerner, "The financing of R&D and innovation," in "Handbook of the Economics of Innovation," Vol. 1, Elsevier, 2010, pp. 609–639.

Holland, Paul W., "Statistics and Causal Inference," *Journal of the American Statistical Association*, 1986, *81* (396), 945–960.

Howell, Sabrina T, "Financing innovation: Evidence from R&D grants," *American Economic Review*, 2017, *107* (4), 1136–64.

Hubbard, R. Glenn, "Capital-Market Imperfections and Investment," *Journal of Economic Literature*, 1998, *36* (1), 193–225.

Hulten, Charles R, "Growth accounting with intermediate inputs," *The Review of Economic Studies*, 1978, *45* (3), 511–518.

Imbens, Guido W. and Donald B. Rubin, *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*, New York: Cambridge University Press, 2015.

Iyoha, Ebehi, "Estimating productivity in the presence of spillovers: Firm-level evidence from the us production network," *Harvard Business School Working Paper 24-033*, 2023.

Jaffe, Adam B, Manuel Trajtenberg, and Rebecca Henderson, "Geographic localization of knowledge spillovers as evidenced by patent citations," *The Quarterly Journal of Economics*, 1993, *108* (3), 577–598.

Lee, David S and Thomas Lemieux, "Regression discontinuity designs in economics," *Journal of economic literature*, 2010, *48* (2), 281–355.

Lerche, Adrian, "Direct and indirect effects of investment tax incentives," *American Economic Review*, 2025, *115* (8), 2781–2818.

Lucking, Brian, Nicholas Bloom, and John Van Reenen, "Have R&D spillovers declined in the 21st century?," *Fiscal Studies*, 2019, *40* (4), 561–590.

Myers, Kyle and Lauren Lanahan, "Estimating spillovers from publicly funded R&D: Evidence from the US Department of Energy," *American Economic Review*, 2022, *112* (7), 2393–2423.

Neumark, David and Helen Simpson, "Place-based policies," *Handbook of Regional and Urban Economics*, 2015, *5*, 1197–1287.

Ramondo, Natalia, Andrés Rodríguez-Clare, and Milagro Saborío-Rodríguez, "Trade, Multinational Production, and the Gains from Openness," *Journal of Political Economy*, 2016, *124* (1), 273–322.

Rosenbaum, Paul R and Donald B Rubin, "The central role of the propensity score in observational studies for causal effects," *Biometrika*, 1983, *70* (1), 41–55.

Rubin, Donald B., "Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies," *Journal of Educational Psychology*, 1974, *66* (5), 688–701.

Santoleri, Pietro, Andrea Mina, Alberto Di Minin, and Irene Martelli, "The causal effects of R&D grants: Evidence from a regression discontinuity," *Review of Economics and Statistics*, 2024, *106* (6), 1495–1510.

Srhoj, Stjepan and Josip Mikulić, "The hidden geography of tourism firm spending: Tracking economic leakages with firm-to-firm transactions," *Annals of Tourism Research*, 2025, *112*, 103933.

Sun, Liyang and Sarah Abraham, "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects," *Journal of Econometrics*, 2021, *225* (2), 175–199.

Takalo, Tuomas, Tanja Tanayama, and Otto Toivanen, "Estimating the Benefits of Targeted R&D Subsidies," *Review of Economics and Statistics*, 2013, *95* (1), 255–272.

– APPENDIX –

[1] Faculty of Economics and Business, University of Dubrovnik, Dubrovnik, Croatia

[2] Harvard Business School, Entrepreneurial Management Unit, Boston, MA

[3] Department of Economics, Faculty of Economics, Business and Tourism, University of Split, Split, Croatia

[4] Department for European Policies, Institute for Development and International Relations, Zagreb, Croatia

[*] Correspondence: Stjepan Srhoj <ssrhoj@efst.hr>

A Production Network Spillovers from EU Funded R&D Grants

Table A1: Pre vs post matching: R&D grant grants

	Pre-match			Post-match		
	Funded	Not funded	Difference	Funded	Not funded	Difference
Young firm (age < 4)	0.09	0.3	-0.21***	0.09	0.09	0
Mature firm (age > 20)	0.49	0.22	0.27***	0.48	0.48	0
Age (in years) (t=-1)	19.03	10.83	8.2***	18.81	19.03	-0.22
Large firm (balance sheet criteria)	0.11	0	0.11***	0.09	0.09	0
Large firm (employee criteria)	0.13	0	0.12***	0.12	0.12	0
Large firm (revenue criteria)	0.12	0	0.11***	0.1	0.1	0
Medium firm (balance sheet criteria)	0.13	0.01	0.12***	0.13	0.13	0
Medium firm (employee criteria)	0.32	0.02	0.3***	0.31	0.31	0
Medium firm (revenue criteria)	0.18	0.01	0.17***	0.18	0.18	0
Small firm (balance sheet criteria)	0.35	0.04	0.31***	0.36	0.36	0
Small firm (employee criteria)	0.4	0.1	0.3***	0.4	0.4	0
Small firm (revenue criteria)	0.31	0.04	0.27***	0.32	0.32	0
Micro firm (balance sheet criteria)	0.4	0.94	-0.54***	0.42	0.42	0
Micro firm (employee criteria)	0.16	0.88	-0.72***	0.16	0.16	0
Micro firm (revenue criteria)	0.39	0.95	-0.56***	0.4	0.4	0
Employees (t=-1)	3.85	1.14	2.7***	3.8	3.73	0.07
Employees (t=-2)	3.67	1.02	2.64***	3.61	3.57	0.05
Employees (t=-3)	3.39	0.93	2.46***	3.33	3.42	-0.09
Employees (t=-4)	3.35	0.85	2.49***	3.29	3.33	-0.05
Tangible Assets (t=-1)	13.04	7.33	5.71***	12.95	13.11	-0.16

Table A1: Pre vs post matching: R&D grant grants

	Pre-match			Post-match		
	Funded	Not funded	Difference	Funded	Not funded	Difference
Tangible Assets (t=-2)	12.73	6.66	6.07***	12.62	12.78	-0.16
Tangible Assets (t=-3)	11.98	5.93	6.05***	11.86	12.15	-0.29
Tangible Assets (t=-4)	11.8	5.44	6.36***	11.68	11.79	-0.11
Revenue (t=-1)	15.22	10.37	4.84***	15.16	15.09	0.07
Revenue (t=-2)	14.84	9.37	5.47***	14.77	14.73	0.04
Revenue (t=-3)	13.99	8.36	5.64***	13.9	14.26	-0.37
Revenue (t=-4)	13.85	7.55	6.3***	13.75	13.95	-0.19
Value Added (IHS) (t=-1)	16.93	11.3	5.64***	16.86	16.6	0.26
Value Added (IHS) (t=-2)	16.48	10.19	6.29***	16.4	16.15	0.25
Value Added (IHS) (t=-3)	15.67	9.11	6.56***	15.57	15.98	-0.41
Value Added (IHS) (t=-4)	15.4	8.2	7.2***	15.28	15.37	-0.09
Profit (IHS) (t=-1)	14.36	7.78	6.58***	14.37	14.3	0.08
Profit (IHS) (t=-2)	13.82	7.23	6.59***	13.83	13.59	0.24
Profit (IHS) (t=-3)	13.22	6.46	6.76***	13.2	13.2	-0.01
Profit (IHS) (t=-4)	13.33	5.76	7.57***	13.2	13.3	-0.1
Material Costs (t=-1)	14.51	9.83	4.68***	14.45	14.37	0.08
Material Costs (t=-2)	14.16	8.86	5.29***	14.09	14.04	0.04
Material Costs (t=-3)	13.28	7.89	5.39***	13.19	13.52	-0.33
Material Costs (t=-4)	13.18	7.13	6.05***	13.08	13.26	-0.18
Labor Productivity (t=-1)	0.13	0.08	0.05***	0.13	0.16	-0.03

Table A1: Pre vs post matching: R&D grant grants

	Pre-match			Post-match		
	Funded	Not funded	Difference	Funded	Not funded	Difference
Labor Productivity (t=-2)	0.12	0.07	0.05***	0.12	0.12	0
Labor Productivity (t=-3)	0.11	0.06	0.05***	0.11	0.11	0
Labor Productivity (t=-4)	0.11	0.05	0.06***	0.11	0.11	0.01
TFP (ACF) (t=-1)	8.12	6.62	1.5***	8.12	8.01	0.12
TFP (ACF) (t=-2)	7.89	5.9	1.99***	7.88	7.73	0.15
TFP (ACF) (t=-3)	7.56	5.23	2.33***	7.55	7.77	-0.23
TFP (ACF) (t=-4)	7.35	4.64	2.71***	7.33	7.32	0.01
Total Exports (t=-1)	11.79	1.67	10.12***	11.65	11.57	0.08
Total Exports (t=-2)	11.31	1.55	9.76***	11.18	10.97	0.21
Total Exports (t=-3)	10.2	1.42	8.78***	10.02	10.24	-0.22
Total Exports (t=-4)	10.05	1.28	8.77***	9.94	9.63	0.31
Exporter Dummy (t=-1)	0.86	0.16	0.7***	0.86	0.87	-0.01
Exporter Dummy (t=-2)	0.84	0.15	0.69***	0.83	0.82	0.01
Exporter Dummy (t=-3)	0.76	0.13	0.63***	0.75	0.78	-0.03
Exporter Dummy (t=-4)	0.76	0.12	0.64***	0.76	0.74	0.02
Total Imports (t=-1)	9.82	1.6	8.22***	9.74	9.56	0.18
Total Imports (t=-2)	9.04	1.45	7.6***	8.93	9.08	-0.15
Total Imports (t=-3)	8.49	1.28	7.21***	8.36	8.84	-0.48
Total Imports (t=-4)	8.08	1.09	6.99***	7.86	7.88	-0.02
Importer Dummy (t=-1)	0.73	0.16	0.57***	0.72	0.72	0

Table A1: Pre vs post matching: R&D grant grants

	Pre-match			Post-match		
	Funded	Not funded	Difference	Funded	Not funded	Difference
Importer Dummy (t=-2)	0.68	0.14	0.53***	0.67	0.7	-0.03
Importer Dummy (t=-3)	0.64	0.13	0.51***	0.64	0.69	-0.05
Importer Dummy (t=-4)	0.63	0.11	0.52***	0.62	0.62	0
Intellectual Assets (value) (t=-1)	8.93	0.99	7.94***	8.78	8.35	0.43
Intellectual Assets (value) (t=-2)	8.11	0.91	7.2***	7.96	7.53	0.43
Intellectual Assets (value) (t=-3)	7.29	0.83	6.46***	7.1	6.9	0.2
Intellectual Assets (value) (t=-4)	6.85	0.78	6.07***	6.64	6.32	0.32
Intellectual Assets (dummy) (t=-1)	0.8	0.13	0.67***	0.79	0.79	0
Intellectual Assets (dummy) (t=-2)	0.74	0.12	0.62***	0.73	0.73	0
Intellectual Assets (dummy) (t=-3)	0.68	0.11	0.58***	0.67	0.69	-0.02
Intellectual Assets (dummy) (t=-4)	0.65	0.1	0.55***	0.64	0.65	-0.01
Gross Patents (t=-1)	0.09	0	0.08**	0.03	0.03	0
Gross Patents (t=-2)	0.08	0	0.08**	0.03	0.03	0.01
Gross Patents (t=-3)	0.07	0	0.07**	0.03	0.02	0.01
Gross Patents (t=-4)	0.06	0	0.06**	0.02	0.02	0
Firms with Patents (t=-1)	0.05	0	0.05***	0.03	0.03	0
Firms with Patents (t=-2)	0.05	0	0.05***	0.03	0.02	0.01
Firms with Patents (t=-3)	0.04	0	0.04***	0.02	0.01	0.01
Firms with Patents (t=-4)	0.04	0	0.04***	0.02	0.01	0.01
Supplier Purchases	10.22	5.95	4.27***	10.13	10.14	-0.01

Table A1: Pre vs post matching: R&D grant grants

	Pre-match			Post-match		
	Funded	Not funded	Difference	Funded	Not funded	Difference
Number of Suppliers	4.17	2.37	1.8***	4.12	4.22	-0.1
Top 1 Supplier Share in Applicant's Total Costs	0.31	0.38	-0.07***	0.31	0.3	0.01
Top 4 Suppliers Share in Applicant's Total Costs	0.58	0.63	-0.04**	0.59	0.58	0.01
Applicant Purchases HHI	1904.47	2734.38	-829.91***	1890.91	1763.2	127.71
Purchases from Q1 Suppliers	9.39	5.26	4.13***	9.29	9.25	0.04
Purchases from Q2 Suppliers	7.52	3.4	4.12***	7.43	7.55	-0.12
Purchases from Q3 Suppliers	5.52	2.15	3.38***	5.42	5.65	-0.23
Purchases from Q4 Suppliers	4.14	1.07	3.07***	4.04	4.21	-0.17
Purchases from Large Suppliers	7.8	4.16	3.63***	7.7	7.66	0.04
Purchases from Medium Suppliers	7.34	3.15	4.18***	7.25	7.23	0.02
Purchases from Small Suppliers	7.87	3.68	4.19***	7.79	7.94	-0.15
Purchases from Micro Suppliers	8.47	4.42	4.05***	8.36	8.33	0.03
Same Municipality Supplier Purchases	8.12	4.13	3.99***	8	7.88	0.12
Other Municipality Supplier Purchases	3.44	2.24	1.2***	3.45	3.81	-0.36
Other County Supplier Purchases	3.65	1.78	1.87***	3.63	3.78	-0.14
Other Region Supplier Purchases	8.61	4.58	4.03***	8.54	8.68	-0.14
Customer Sales	9.85	4.69	5.16***	9.75	9.83	-0.08
Number of Customers	3.21	1.36	1.84***	3.17	3.33	-0.16
Top 1 Buyer Share in Applicant's Total Sales	0.41	0.37	0.04	0.4	0.39	0.02
Top 4 Buyers Share in Applicant's Total Sales	0.64	0.52	0.12***	0.64	0.63	0.01

Table A1: Pre vs post matching: R&D grant grants

	Pre-match			Post-match		
	Funded	Not funded	Difference	Funded	Not funded	Difference
Applicant Sales HHI	3056.09	3170.04	-113.94	3003.75	2829.71	174.04
Sales to Q1 Customers	8.42	3.61	4.81***	8.33	8.55	-0.22
Sales to Q2 Customers	5.93	2.3	3.63***	5.97	6.42	-0.46
Sales to Q3 Customers	3.74	1.43	2.31***	3.72	3.91	-0.19
Sales to Q4 Customers	3.08	1.04	2.04***	2.99	3.33	-0.34
Sales to Large Customers	6.37	1.58	4.79***	6.29	6.59	-0.3
Sales to Medium Customers	6.29	1.81	4.48***	6.2	6.36	-0.16
Sales to Small Customers	6.11	2.48	3.63***	6.06	6.63	-0.57
Sales to Micro Customers	6.32	3.02	3.3***	6.3	6.89	-0.59
Same Municipality Customer Sales	6.97	2.94	4.03***	6.87	7.24	-0.37
Other Municipality Customer Sales	2.95	1.53	1.42***	2.97	3.33	-0.37
Other County Customer Sales	3.2	1.1	2.09***	3.21	3.47	-0.26
Other Region Customer Sales	7.15	2.7	4.45***	7.12	7.6	-0.48
N	164	106027		159	159	

Note: All monetary variables are expressed in logarithms. Profit variables are transformed using the inverse hyperbolic sine transformation. Total factor productivity is estimated using the Ackerberg–Caves–Frazer procedure. Funded denotes firms receiving an R&D grant. Unfunded denotes all other firms in the economy that did not receive R&D grants or other growth-oriented grant schemes (digital transformation grants and innovation commercialization grants).

Table A2: Distribution of firms receiving R&D grants

Category	Firms	Share (%)
Panel A. Top 10 broad sectors		
IT services	32	20.1
Electronics, electrical and telecom	17	10.7
Metals and mineral products	16	10.1
Trade, transport and vehicles	13	8.2
Engineering	12	7.5
Professional and business services	11	6.9
R&D, information and publishing	11	6.9
Construction and installation	10	6.3
Food and beverages	8	5.0
Chemicals, pharmaceuticals and fuels	6	3.8
Other sectors	23	14.5
Panel B. Knowledge-intensity classifications		
High or medium-high digital intensity	105	66.0
Higher-technology industries[a]	95	59.7
High or medium-high R&D intensity	68	42.8
Total firms	159	100.0

Notes: The table reports the distribution of firms receiving R&D grants. Panel A presents the ten most represented broad sectors. Panel B summarizes the share of firms operating in industries classified as high- or medium-high digital intensity, higher-technology industries, and high- or medium-high R&D intensity. Percentages are calculated relative to the full sample of 159 recipient firms. [a] Higher-technology industries comprise higher-technology manufacturing and higher-technology services.

Table A3: Pre vs post matching: suppliers

	Pre-match			Post-match		
	Funded	Not funded	Difference	Funded	Not funded	Difference
Young firm (age < 4)	1.2	1.25	-0.04***	1.2	1.2	0
Mature firm (age > 20)	0.25	0.16	0.09***	0.25	0.25	0
Age (in years) (t=-1)	11.54	8.96	2.59***	11.48	11.43	0.05
Large firm (balance sheet criteria)	0.01	0	0***	0	0	0
Large firm (employee criteria)	0.01	0	0.01***	0	0	0
Large firm (revenue criteria)	0	0	0**	0	0	0
Medium firm (balance sheet criteria)	0.02	0.01	0.01***	0.01	0.01	0
Medium firm (employee criteria)	0.03	0.01	0.03***	0.03	0.03	0
Medium firm (revenue criteria)	0.02	0	0.02***	0.02	0.02	0
Small firm (balance sheet criteria)	0.07	0.03	0.04***	0.07	0.07	0
Small firm (employee criteria)	0.2	0.07	0.14***	0.2	0.2	0
Small firm (revenue criteria)	0.07	0.02	0.05***	0.07	0.07	0
Micro firm (balance sheet criteria)	0.91	0.96	-0.06***	0.91	0.91	0
Micro firm (employee criteria)	0.76	0.92	-0.17***	0.76	0.76	0
Micro firm (revenue criteria)	0.91	0.98	-0.07***	0.91	0.91	0
Employees (t=-1)	1.56	0.83	0.73***	1.54	1.56	-0.02
Employees (t=-2)	1.42	0.74	0.69***	1.4	1.42	-0.02
Employees (t=-3)	1.3	0.66	0.64***	1.29	1.3	-0.01
Employees (t=-4)	1.2	0.6	0.6***	1.18	1.18	0
Tangible Assets (t=-1)	8.1	5.8	2.3***	8.06	8.05	0.01

Table A3: Pre vs post matching: suppliers

	Pre-match			Post-match		
	Funded	Not funded	Difference	Funded	Not funded	Difference
Tangible Assets (t=-2)	7.42	5.26	2.16***	7.38	7.36	0.02
Tangible Assets (t=-3)	6.67	4.65	2.02***	6.64	6.56	0.08
Tangible Assets (t=-4)	6.17	4.25	1.92***	6.13	6.06	0.07
Revenue (t=-1)	11.11	8.49	2.62***	11.08	11.11	-0.04
Revenue (t=-2)	10.3	7.67	2.64***	10.26	10.35	-0.08
Revenue (t=-3)	9.34	6.75	2.59***	9.31	9.38	-0.07
Revenue (t=-4)	8.52	6.03	2.49***	8.48	8.51	-0.03
Value Added (IHS) (t=-1)	12.26	9.25	3.01***	12.23	12.22	0.01
Value Added (IHS) (t=-2)	11.4	8.31	3.09***	11.37	11.45	-0.08
Value Added (IHS) (t=-3)	10.32	7.33	2.99***	10.29	10.37	-0.07
Value Added (IHS) (t=-4)	9.49	6.51	2.98***	9.45	9.45	0.01
Profit (IHS) (t=-1)	8.89	6.18	2.7***	8.87	8.91	-0.04
Profit (IHS) (t=-2)	8.44	5.75	2.69***	8.42	8.53	-0.11
Profit (IHS) (t=-3)	7.72	5.03	2.68***	7.69	7.84	-0.15
Profit (IHS) (t=-4)	7.02	4.43	2.59***	6.99	7.02	-0.03
Material Costs (t=-1)	10.47	8.02	2.45***	10.44	10.45	-0.01
Material Costs (t=-2)	9.71	7.22	2.48***	9.67	9.74	-0.08
Material Costs (t=-3)	8.79	6.36	2.43***	8.76	8.78	-0.03
Material Costs (t=-4)	8	5.68	2.32***	7.96	7.99	-0.03
Labor Productivity (t=-1)	0.09	0.06	0.03***	0.09	0.1	-0.02

Table A3: Pre vs post matching: suppliers

	Pre-match			Post-match		
	Funded	Not funded	Difference	Funded	Not funded	Difference
Labor Productivity (t=-2)	0.09	0.05	0.04***	0.09	0.1	-0.02
Labor Productivity (t=-3)	0.07	0.04	0.03***	0.07	0.08	-0.01
Labor Productivity (t=-4)	0.06	0.04	0.03***	0.06	0.07	-0.01
TFP (ACF) (t=-1)	6.73	5.63	1.1***	6.73	6.72	0
TFP (ACF) (t=-2)	6.2	4.99	1.21***	6.19	6.27	-0.07
TFP (ACF) (t=-3)	5.58	4.36	1.22***	5.58	5.64	-0.06
TFP (ACF) (t=-4)	5.1	3.82	1.28***	5.09	5.11	-0.02
Total Exports (t=-1)	2.14	1.18	0.96***	2.09	2.19	-0.1
Total Exports (t=-2)	2.01	1.08	0.93***	1.95	2.05	-0.1
Total Exports (t=-3)	1.8	0.98	0.83***	1.76	1.9	-0.15
Total Exports (t=-4)	1.69	0.86	0.83***	1.64	1.72	-0.07
Exporter Dummy (t=-1)	0.2	0.11	0.08***	0.19	0.2	-0.01
Exporter Dummy (t=-2)	0.18	0.1	0.08***	0.18	0.18	0
Exporter Dummy (t=-3)	0.16	0.09	0.07***	0.16	0.17	-0.01
Exporter Dummy (t=-4)	0.15	0.08	0.07***	0.15	0.15	0
Total Imports (t=-1)	2.17	0.99	1.18***	2.14	2.14	0.01
Total Imports (t=-2)	1.92	0.88	1.04***	1.9	1.93	-0.04
Total Imports (t=-3)	1.69	0.76	0.92***	1.66	1.66	0
Total Imports (t=-4)	1.44	0.64	0.8***	1.41	1.41	0
Importer Dummy (t=-1)	0.21	0.11	0.11***	0.21	0.21	0

Table A3: Pre vs post matching: suppliers

	Pre-match			Post-match		
	Funded	Not funded	Difference	Funded	Not funded	Difference
Importer Dummy (t=-2)	0.19	0.09	0.09***	0.18	0.19	0
Importer Dummy (t=-3)	0.16	0.08	0.08***	0.16	0.16	0
Importer Dummy (t=-4)	0.14	0.07	0.07***	0.14	0.14	0
Intellectual Assets (value) (t=-1)	1.44	0.66	0.77***	1.41	1.4	0.01
Intellectual Assets (value) (t=-2)	1.33	0.61	0.71***	1.29	1.29	0
Intellectual Assets (value) (t=-3)	1.19	0.55	0.63***	1.16	1.18	-0.02
Intellectual Assets (value) (t=-4)	1.11	0.51	0.6***	1.08	1.07	0.01
Intellectual Assets (dummy) (t=-1)	0.17	0.09	0.09***	0.17	0.17	0
Intellectual Assets (dummy) (t=-2)	0.16	0.08	0.08***	0.16	0.16	0
Intellectual Assets (dummy) (t=-3)	0.15	0.07	0.07***	0.14	0.15	0
Intellectual Assets (dummy) (t=-4)	0.14	0.07	0.07***	0.14	0.14	0
Gross Patents (t=-1)	0	0	0	0	0	0
Gross Patents (t=-2)	0	0	0	0	0	0
Gross Patents (t=-3)	0	0	0	0	0	0
Gross Patents (t=-4)	0	0	0	0	0	0
Firms with Patents (t=-1)	0	0	0**	0	0	0
Firms with Patents (t=-2)	0	0	0**	0	0	0
Firms with Patents (t=-3)	0	0	0**	0	0	0
Firms with Patents (t=-4)	0	0	0**	0	0	0
N	3392	114313		3370	3370	

Table A3: Pre vs post matching: suppliers

	Pre-match			Post-match		
	Funded	Not funded	Difference	Funded	Not funded	Difference

Note: All monetary variables are expressed in logarithms. Profit variables are transformed using the inverse hyperbolic sine transformation. Total factor productivity is estimated using the Ackerberg–Caves–Frazer procedure. Treated (funded) suppliers denote suppliers linked to firms receiving an R&D grant. Control (not funded) suppliers denote all other suppliers in the economy that did not transact with firms receiving R&D grants or other growth-oriented grant schemes (digital transformation grants and innovation commercialization grants).

Table A4: Dynamic upstream spillovers to suppliers after applicant grant awards

	Revenue	Material Costs	Tangible Assets	Employees
-4	0.069	0.087	0.052	0.026
	(0.057)	(0.051)	(0.080)	(0.016)
-3	0.024	0.085	0.025	0.027
	(0.055)	(0.050)	(0.078)	(0.016)
-2	0.046	0.068	0.004	0.003
	(0.054)	(0.048)	(0.076)	(0.015)
0	0.248**	0.260**	0.057	−0.016
	(0.052)	(0.047)	(0.071)	(0.014)
1	0.323**	0.380**	0.159*	0.088**
	(0.052)	(0.047)	(0.071)	(0.014)
2	0.367**	0.412**	0.247**	0.112**
	(0.053)	(0.048)	(0.072)	(0.015)
3	0.376**	0.402**	0.333**	0.105**
	(0.053)	(0.048)	(0.072)	(0.015)
4	0.346**	0.351**	0.302**	0.101**
	(0.054)	(0.048)	(0.073)	(0.015)
Num.Obs.	53 393	53 324	49 995	51 328
R2	0.736	0.785	0.799	0.895
R2 Adj.	0.679	0.739	0.752	0.871
AIC	200 535.6	189 015.1	214 120.8	58 103.6
BIC	284 049.8	272 517.2	297 016.9	141 247.1
RMSE	1.33	1.19	1.71	0.35
Mean dep. var.	0.842	0.511	0.626	11.939
Median dep. var.	0.137	0.064	0.013	4.000
Unique firms	6740	6740	6740	6740

Note: This table reports event-study difference-in-differences estimates of supplier outcomes around the applicant's grant-award year. Treated suppliers are those linked (in the pre-award period) to funded applicants. Control suppliers are those linked to matched non-funded applicants and, by construction, have no sales to any treated firm. Event time is defined relative to the grant decision year, and $t = 1$ is the omitted reference period. Regressions include supplier fixed effects and calendar-year fixed effects. Outcomes are in logs. Pre-award supplier links are defined using baseline transaction relationships and held fixed throughout. $^{*}p < 0.05$; ** $p < 0.01$.

Table A5: Heterogeneous supplier spillovers by applicant-supplier geographic distance

	Revenue			Material Costs		
	(bottom)	(middle)	(top)	(bottom)	(middle)	(top)
-4	0.180	0.048	−0.070	0.109	0.079	0.039
	(0.110)	(0.100)	(0.107)	(0.097)	(0.091)	(0.097)
-3	0.033	0.032	−0.078	0.108	0.086	0.045
	(0.107)	(0.096)	(0.104)	(0.094)	(0.088)	(0.095)
-2	0.049	0.072	−0.025	0.073	0.057	0.063
	(0.104)	(0.093)	(0.102)	(0.092)	(0.086)	(0.092)
0	0.349**	0.227*	0.085	0.231**	0.244**	0.225*
	(0.100)	(0.090)	(0.098)	(0.088)	(0.083)	(0.088)
1	0.480**	0.326**	0.241*	0.436**	0.374**	0.350**
	(0.102)	(0.091)	(0.098)	(0.090)	(0.084)	(0.089)
2	0.431**	0.536**	0.132	0.383**	0.507**	0.359**
	(0.103)	(0.092)	(0.099)	(0.091)	(0.085)	(0.090)
3	0.397**	0.511**	0.174	0.417**	0.497**	0.297**
	(0.104)	(0.093)	(0.100)	(0.091)	(0.085)	(0.090)
4	0.427**	0.486**	0.087	0.371**	0.433**	0.259**
	(0.105)	(0.094)	(0.101)	(0.092)	(0.086)	(0.091)
Num.Obs.	17 791	17 950	17 652	17 770	17 948	17 606
R2	0.747	0.732	0.728	0.796	0.780	0.778
R2 Adj.	0.679	0.656	0.654	0.740	0.718	0.718
AIC	68 311.3	68 277.2	68 118.3	63 675.9	65 199.7	64 338.4
BIC	97 790.8	99 294.9	97 396.9	93 150.9	96 216.9	93 607.2
RMSE	1.33	1.30	1.35	1.17	1.19	1.21
Mean dep. var.	1.033	0.808	0.568	0.601	0.515	0.363
Median dep. var.	0.129	0.150	0.154	0.059	0.072	0.071
Unique firms	2054	2091	2041	2052	2089	2038

Note: This table reports event-study difference-in-differences spillover estimates (as in Table A4 / Figure 1) separately by terciles of geographic distance between the applicant and the observed supplier. Distance is measured as straight-line (air-line) distance between the municipalities of the applicant's headquarters and the supplier's headquarters. Terciles are defined as: (1) bottom: < 23.55 km; (2) middle: 23.55-171.39 km; (3) top: ≥ 171.39 km. Outcomes are in logs. Pre-award supplier links are defined using baseline transaction relationships and held fixed throughout. *$p < 0.05$; ** $p < 0.01$.

Table A6: Heterogeneous supplier spillovers by applicant-supplier geographic distance

	Tangible Assets			Employees		
	(bottom)	(middle)	(top)	(bottom)	(middle)	(top)
-4	0.181	−0.064	−0.022	0.031	0.022	0.021
	(0.161)	(0.139)	(0.148)	(0.030)	(0.028)	(0.031)
-3	0.113	0.013	−0.105	0.043	0.008	0.042
	(0.157)	(0.135)	(0.144)	(0.030)	(0.027)	(0.030)
-2	0.009	0.021	−0.012	0.020	−0.018	0.017
	(0.152)	(0.131)	(0.140)	(0.029)	(0.027)	(0.029)
0	−0.028	0.098	0.226	−0.029	−0.007	−0.008
	(0.140)	(0.122)	(0.131)	(0.027)	(0.025)	(0.027)
1	0.179	0.100	0.282*	0.110**	0.088**	0.080**
	(0.141)	(0.123)	(0.131)	(0.027)	(0.026)	(0.028)
2	0.264	0.283*	0.181	0.135**	0.112**	0.099**
	(0.143)	(0.124)	(0.132)	(0.028)	(0.026)	(0.028)
3	0.430**	0.356**	0.192	0.149**	0.105**	0.076**
	(0.144)	(0.125)	(0.133)	(0.028)	(0.026)	(0.028)
4	0.386**	0.420**	0.136	0.148**	0.089**	0.076**
	(0.145)	(0.126)	(0.135)	(0.028)	(0.026)	(0.028)
Num.Obs.	16 510	16 892	16 593	17 015	17 329	16 984
R2	0.791	0.800	0.804	0.903	0.895	0.886
R2 Adj.	0.728	0.738	0.747	0.875	0.864	0.853
AIC	72 764.4	73 229.8	72 294.6	19 790.5	20 936.6	21 544.6
BIC	101 961.0	104 005.7	101 340.4	49 101.2	51 814.2	50 678.1
RMSE	1.74	1.67	1.70	0.35	0.35	0.37
Mean dep. var.	1.046	0.546	0.419	14.321	12.545	12.520
Median dep. var.	0.016	0.024	0.025	4.000	4.000	5.000
Unique firms	1724	1827	1760	1850	1926	1864

Note: This table reports event-study difference-in-differences spillover estimates (as in Table A4 / Figure 1) separately by terciles of geographic distance between the applicant and the observed supplier. Distance is measured as straight-line (air-line) distance between the municipalities of the applicant's headquarters and the supplier's headquarters. Terciles are defined as: (1) bottom: $<$ 23.55 km; (2) middle: 23.55-171.39 km; (3) top: $\geq$ 171.39 km. Outcomes are in logs. Pre-award supplier links are defined using baseline transaction relationships and held fixed throughout. *$p < 0.05$; ** $p < 0.01$.

Table A7: Heterogeneous supplier spillovers by supplier financial constraints

	Revenue			Material Costs		
	(bottom)	(middle)	(top)	(bottom)	(middle)	(top)
-4	0.023	0.007	0.185	0.034	−0.049	0.182
	(0.054)	(0.096)	(0.141)	(0.057)	(0.096)	(0.128)
-3	0.033	−0.125	0.163	0.080	−0.060	0.102
	(0.054)	(0.091)	(0.135)	(0.056)	(0.092)	(0.123)
-2	0.026	0.049	0.048	0.051	−0.006	0.128
	(0.054)	(0.087)	(0.129)	(0.056)	(0.088)	(0.118)
0	0.019	0.122	0.489**	0.093	0.140	0.522**
	(0.053)	(0.085)	(0.122)	(0.056)	(0.086)	(0.112)
1	0.106*	0.177*	0.677**	0.146**	0.247**	0.723**
	(0.054)	(0.085)	(0.125)	(0.056)	(0.086)	(0.115)
2	0.135*	0.229**	0.752**	0.174**	0.276**	0.761**
	(0.054)	(0.086)	(0.128)	(0.056)	(0.087)	(0.117)
3	0.085	0.295**	0.713**	0.122*	0.321**	0.741**
	(0.054)	(0.087)	(0.130)	(0.056)	(0.087)	(0.119)
4	0.047	0.237**	0.761**	0.083	0.277**	0.693**
	(0.054)	(0.087)	(0.132)	(0.056)	(0.088)	(0.121)
Num.Obs.	18 366	17 005	15 798	18 362	17 029	15 880
R2	0.853	0.676	0.680	0.871	0.723	0.717
R2 Adj.	0.816	0.587	0.585	0.838	0.647	0.633
AIC	49 943.3	62 160.7	69 156.9	51 309.5	62 452.6	66 716.5
BIC	79 230.5	90 377.6	97 021.1	80 595.8	90 674.7	94 599.5
RMSE	0.77	1.21	1.72	0.80	1.22	1.57
Mean dep. var.	1.382	0.611	0.509	0.795	0.376	0.354
Median dep. var.	0.194	0.138	0.094	0.102	0.068	0.042
Unique firms	2081	2083	2022	2072	2075	2032

Note: This table reports event-study difference-in-differences spillover estimates (as in Table A4 / Figure 1) separately by terciles of suppliers' financial constraints, proxied by the Omega Score (Altman et al., 2023), where more negative values indicate tighter financial constraints. Terciles are defined as: (1) bottom: > 24.37; (2) middle: -24.37 to -76.01; (3) top: < -76.01. Outcomes are in logs. Pre-award supplier links are defined using baseline transaction relationships and held fixed throughout. *$p < 0.05$; ** $p < 0.01$.

Table A8: Heterogeneous supplier spillovers by supplier financial constraints

	Tangible Assets			Employees		
	(bottom)	(middle)	(top)	(bottom)	(middle)	(top)
-4	0.108	−0.054	0.096	0.005	0.023	0.044
	(0.111)	(0.143)	(0.179)	(0.017)	(0.029)	(0.039)
-3	0.035	0.009	0.075	0.005	0.021	0.065
	(0.110)	(0.137)	(0.173)	(0.017)	(0.028)	(0.038)
-2	0.010	0.022	0.002	0.008	−0.014	0.014
	(0.109)	(0.130)	(0.164)	(0.017)	(0.027)	(0.036)
0	−0.002	0.149	0.012	0.001	−0.010	−0.016
	(0.106)	(0.121)	(0.150)	(0.017)	(0.025)	(0.033)
1	0.035	0.110	0.268	0.032	0.077**	0.151**
	(0.106)	(0.122)	(0.152)	(0.017)	(0.026)	(0.033)
2	0.104	0.221	0.388*	0.032	0.112**	0.187**
	(0.106)	(0.123)	(0.154)	(0.017)	(0.026)	(0.034)
3	0.281**	0.311*	0.348*	0.047**	0.101**	0.159**
	(0.107)	(0.124)	(0.156)	(0.017)	(0.026)	(0.034)
4	0.154	0.263*	0.451**	0.036*	0.095**	0.168**
	(0.107)	(0.125)	(0.158)	(0.017)	(0.026)	(0.035)
Num.Obs.	17 480	15 945	14 415	18 195	16 673	14 278
R2	0.824	0.771	0.784	0.950	0.877	0.862
R2 Adj.	0.775	0.704	0.712	0.937	0.843	0.815
AIC	70 434.8	68 574.1	66 793.9	7279.5	20 703.1	22 379.5
BIC	99 536.7	96 556.5	94 325.2	36 531.7	48 848.1	49 876.1
RMSE	1.46	1.65	1.91	0.24	0.36	0.41
Mean dep. var.	1.438	0.397	0.236	14.988	13.467	10.655
Median dep. var.	0.027	0.016	0.019	4.000	4.000	4.000
Unique firms	1876	1834	1601	2044	2051	1545

Note: This table reports event-study difference-in-differences spillover estimates (as in Table A4 / Figure 1) separately by terciles of suppliers' financial constraints, proxied by the Omega Score (Altman et al., 2023), where more negative values indicate tighter financial constraints. Terciles are defined as: (1) bottom: $>$ 24.37; (2) middle: -24.37 to -76.01; (3) top: $<$ -76.01. Outcomes are in logs. Pre-award supplier links are defined using baseline transaction relationships and held fixed throughout. *$p < 0.05$; ** $p < 0.01$.

Table A9: Heterogeneous upstream spillovers by supplier size

	Revenue			Material Costs		
	(bottom)	(middle)	(top)	(bottom)	(middle)	(top)
-4	0.274	0.098	0.056	0.173	0.105	0.062
	(0.147)	(0.093)	(0.077)	(0.114)	(0.083)	(0.071)
-3	0.142	0.067	0.013	0.113	0.138	0.057
	(0.140)	(0.090)	(0.076)	(0.109)	(0.080)	(0.070)
-2	0.114	0.059	0.023	0.087	0.071	0.073
	(0.133)	(0.087)	(0.075)	(0.103)	(0.078)	(0.069)
0	0.614**	0.124	0.083	0.443**	0.218**	0.110
	(0.126)	(0.086)	(0.074)	(0.097)	(0.076)	(0.068)
1	0.731**	0.270**	0.178*	0.635**	0.293**	0.211**
	(0.129)	(0.087)	(0.075)	(0.099)	(0.077)	(0.069)
2	0.744**	0.321**	0.272**	0.649**	0.331**	0.258**
	(0.131)	(0.087)	(0.075)	(0.101)	(0.078)	(0.069)
3	0.747**	0.314**	0.250**	0.583**	0.322**	0.305**
	(0.133)	(0.088)	(0.076)	(0.102)	(0.078)	(0.069)
4	0.695**	0.308**	0.259**	0.470**	0.274**	0.321**
	(0.135)	(0.089)	(0.076)	(0.103)	(0.079)	(0.070)
Num.Obs.	16 090	17 447	17 899	15 832	17 396	17 874
R2	0.540	0.404	0.596	0.619	0.517	0.693
R2 Adj.	0.408	0.248	0.489	0.507	0.390	0.611
AIC	71 738.6	64 148.4	60 632.7	61 966.2	59 837.5	57 363.2
BIC	99 315.8	92 326.8	89 839.0	89 485.4	88 005.2	86 564.3
RMSE	1.80	1.24	1.07	1.37	1.10	0.98
Mean dep. var.	0.029	0.155	2.356	0.017	0.082	1.443
Median dep. var.	0.025	0.135	0.725	0.011	0.064	0.411
Unique firms	2088	2087	2087	2009	2084	2086

Note: This table reports event-study difference-in-differences spillover estimates (as in Table A4 / Figure 1) separately by terciles of baseline supplier size, measured by suppliers' pre-award revenue (EUR). Terciles are defined as: (1) bottom: $<$ €89,440; (2) middle: €89,440-€381,046; (3) top: $\geq$ €381,046. Outcomes are in logs. Pre-award supplier links are defined using baseline transaction relationships and held fixed throughout. *$p < 0.05$; ** $p < 0.01$.

Table A10: Heterogeneous upstream spillovers by supplier size

	Tangible Assets			Employees		
	(bottom)	(middle)	(top)	(bottom)	(middle)	(top)
-4	−0.037	0.164	0.086	0.034	0.041	0.033
	(0.213)	(0.131)	(0.098)	(0.029)	(0.026)	(0.027)
-3	−0.038	0.079	0.079	0.036	0.036	0.035
	(0.205)	(0.127)	(0.096)	(0.027)	(0.025)	(0.027)
-2	−0.031	0.017	0.059	−0.014	0.015	0.019
	(0.194)	(0.122)	(0.094)	(0.026)	(0.025)	(0.026)
0	−0.012	0.006	0.125	−0.049*	0.011	0.009
	(0.170)	(0.116)	(0.092)	(0.023)	(0.024)	(0.026)
1	−0.008	0.157	0.232*	0.085**	0.079**	0.092**
	(0.172)	(0.117)	(0.093)	(0.024)	(0.024)	(0.026)
2	0.147	0.227	0.303**	0.102**	0.103**	0.116**
	(0.174)	(0.118)	(0.093)	(0.024)	(0.024)	(0.026)
3	0.226	0.358**	0.327**	0.100**	0.097**	0.104**
	(0.176)	(0.119)	(0.094)	(0.024)	(0.024)	(0.026)
4	0.178	0.331**	0.307**	0.103**	0.078**	0.109**
	(0.177)	(0.120)	(0.094)	(0.025)	(0.025)	(0.026)
Num.Obs.	13 811	16 509	17 520	14 471	17 003	17 672
R2	0.724	0.699	0.783	0.690	0.741	0.884
R2 Adj.	0.627	0.614	0.724	0.588	0.670	0.852
AIC	66 362.5	69 938.4	66 632.0	13 313.1	18 640.8	22 314.8
BIC	93 391.7	97 916.3	95 758.1	40 509.8	46 725.7	51 473.2
RMSE	2.06	1.62	1.31	0.30	0.34	0.37
Mean dep. var.	0.059	0.124	1.804	1.878	4.956	30.167
Median dep. var.	0.004	0.015	0.104	1.000	4.000	13.000
Unique firms	1387	1908	2016	1589	2003	2048

Note: This table reports event-study difference-in-differences spillover estimates (as in Table A4 / Figure 1) separately by terciles of baseline supplier size, measured by suppliers' pre-award revenue (EUR). Terciles are defined as: (1) bottom: < €89,440; (2) middle: €89,440-€381,046; (3) top: ≥ €381,046. Outcomes are in logs. Pre-award supplier links are defined using baseline transaction relationships and held fixed throughout. *$p < 0.05$; ** $p < 0.01$.

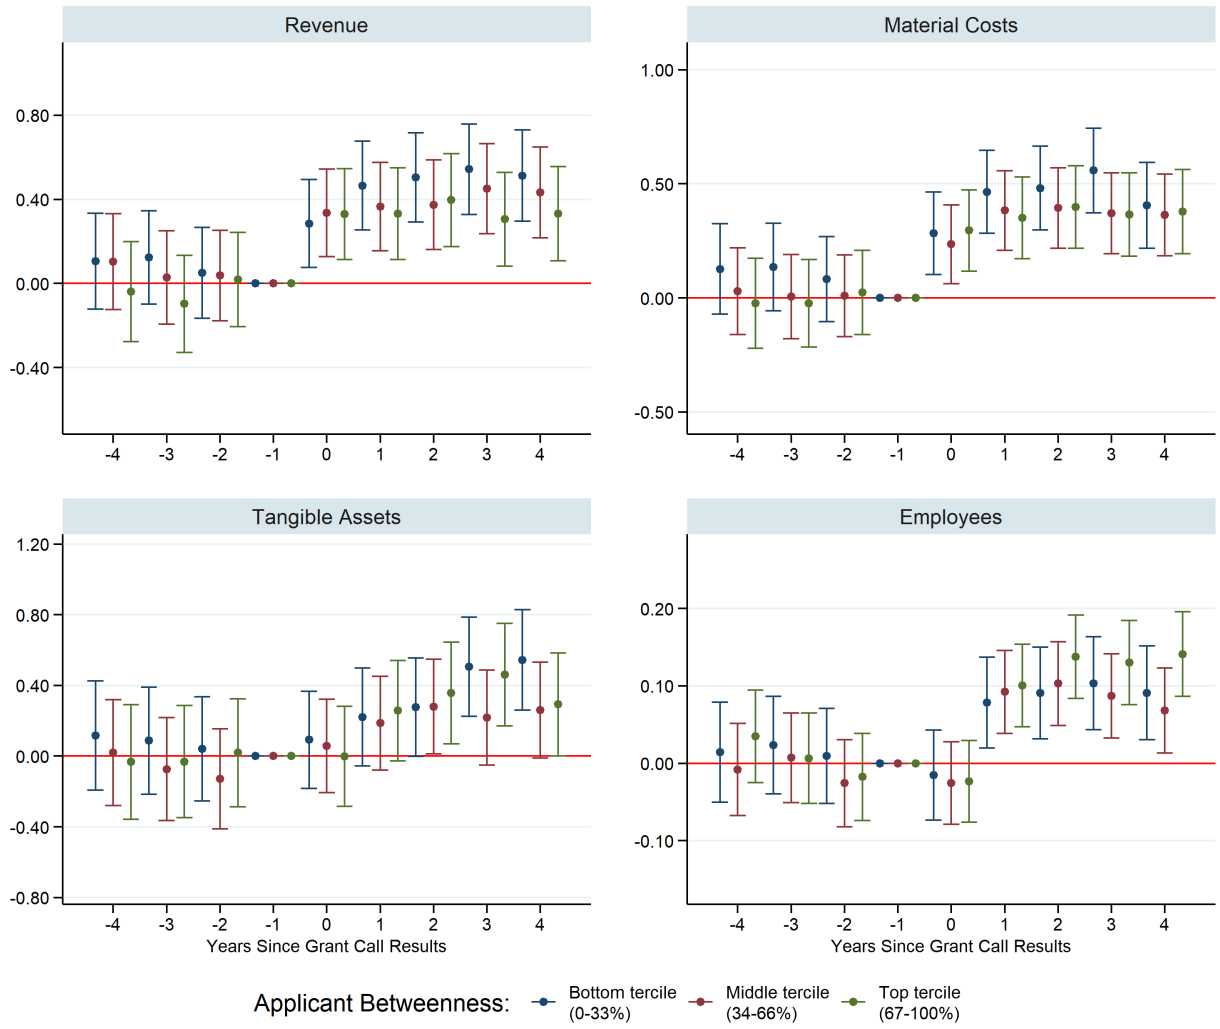


Figure A1: Upstream spillovers on suppliers, split by applicant betweenness centrality

Note: Event-study estimates of supplier outcomes (revenue, material costs, tangible assets, and employment) as in Figure 1, estimated separately by terciles of the grant recipient's betweenness centrality in the pre-award firm-to-firm transaction network. Betweenness centrality measures how often the applicant lies on the shortest path between other firms in the network, capturing its role as a broker in intermediating trade flows. The omitted reference period is year −1. Whiskers denote 95% confidence intervals.

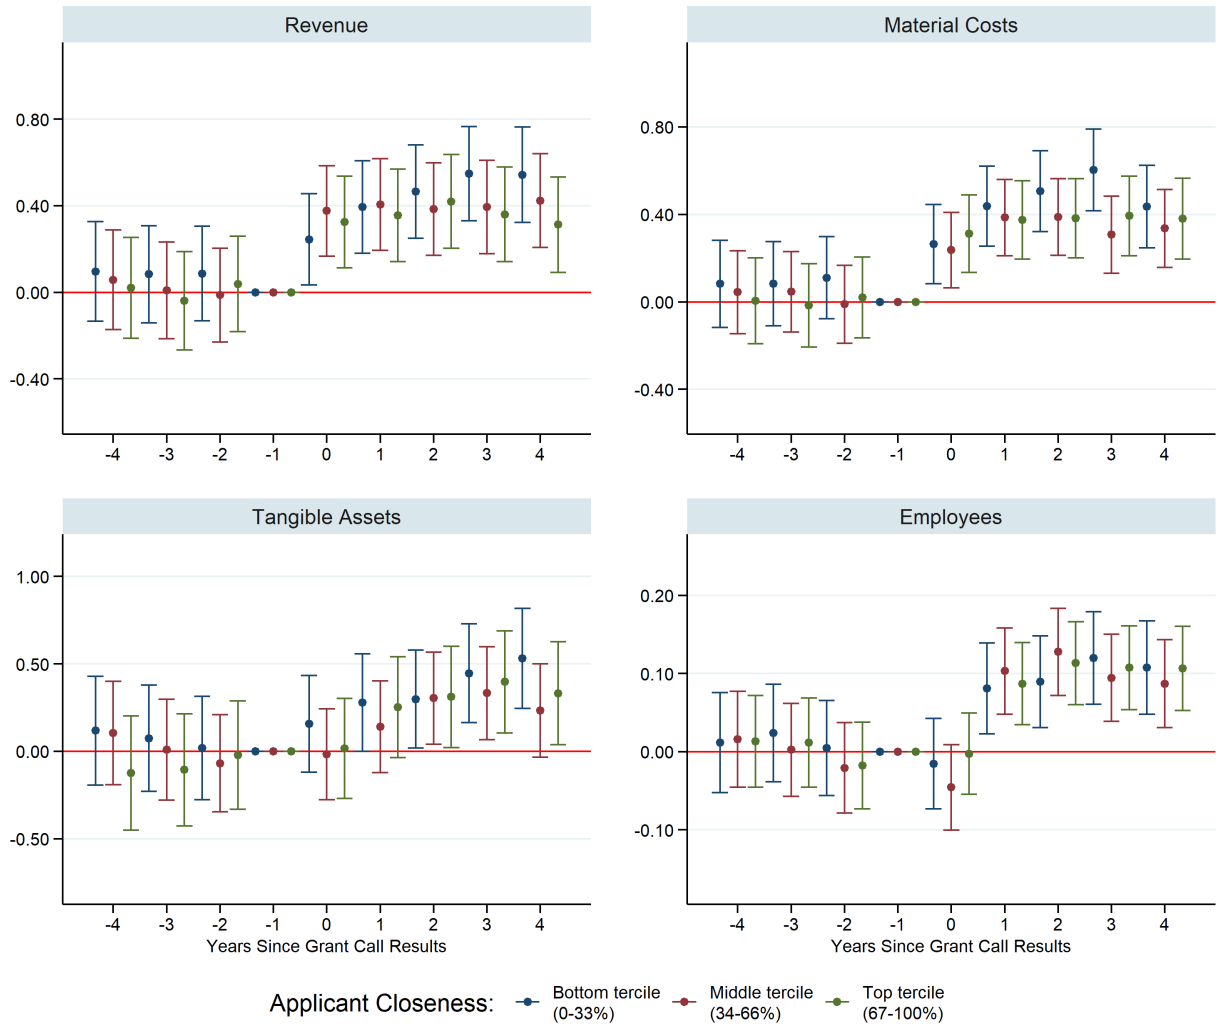


Figure A2: Upstream spillovers on suppliers, split by applicant closeness centrality

Note: Event-study estimates of supplier outcomes (revenue, material costs, tangible assets, and employment) as in Figure 1, estimated separately by terciles of the grant recipient's closeness centrality in the pre-award firm-to-firm transaction network. Closeness centrality is the inverse of the applicant's average distance to all other firms in the network, reflecting how quickly a demand shock originating at the applicant can reach other firms. The omitted reference period is year −1. Whiskers denote 95% confidence intervals.

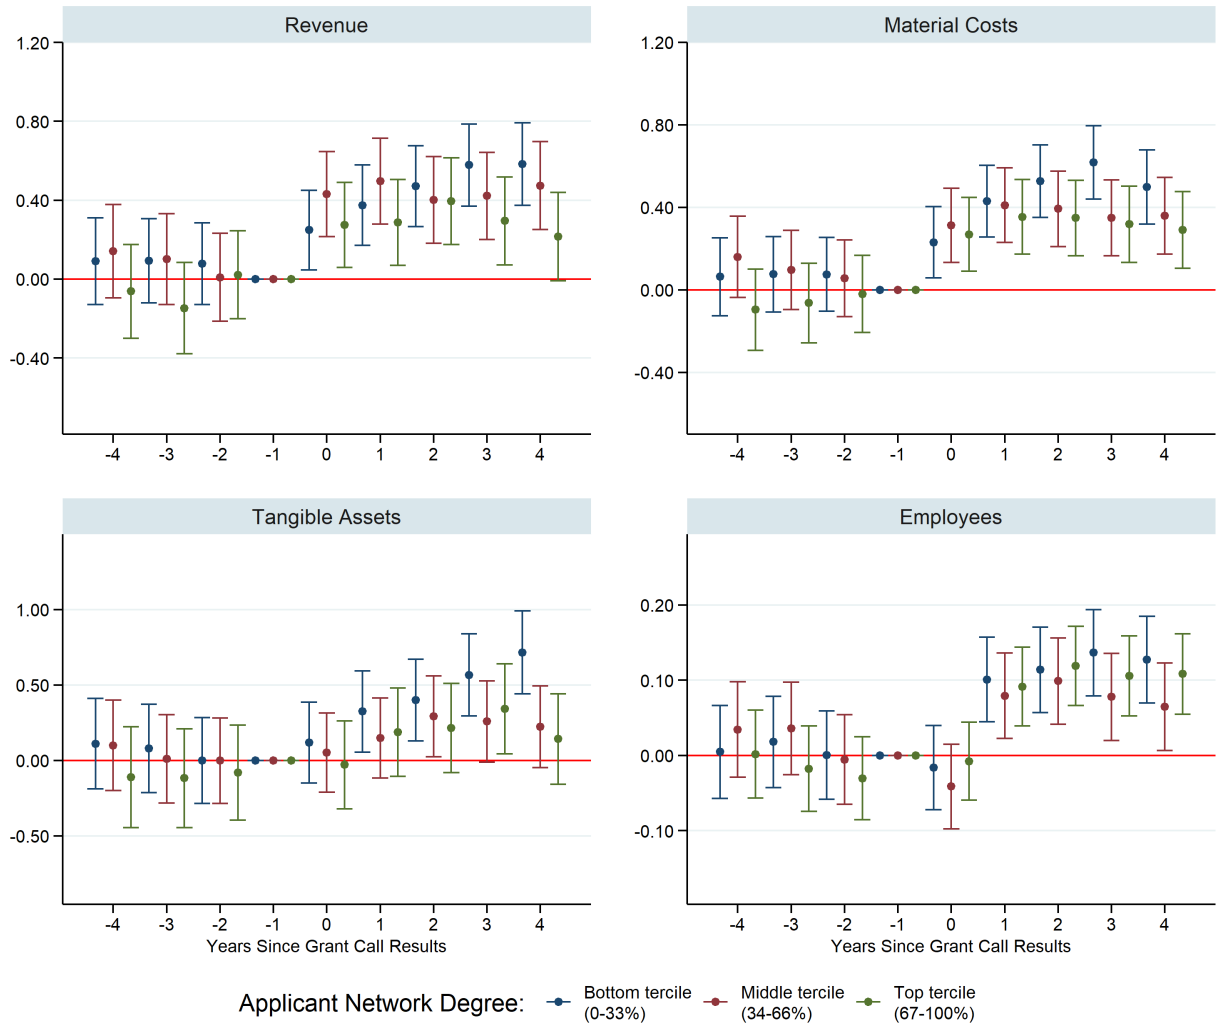


Figure A3: Upstream spillovers on suppliers, split by applicant degree centrality

Note: Event-study estimates of supplier outcomes (revenue, material costs, tangible assets, and employment) as in Figure 1, estimated separately by terciles of the grant recipient's degree centrality in the pre-award firm-to-firm transaction network. Degree centrality counts the applicant's number of direct trading partners, reflecting the breadth of its market relationships and its direct connectivity within the network. The omitted reference period is year −1. Whiskers denote 95% confidence intervals.

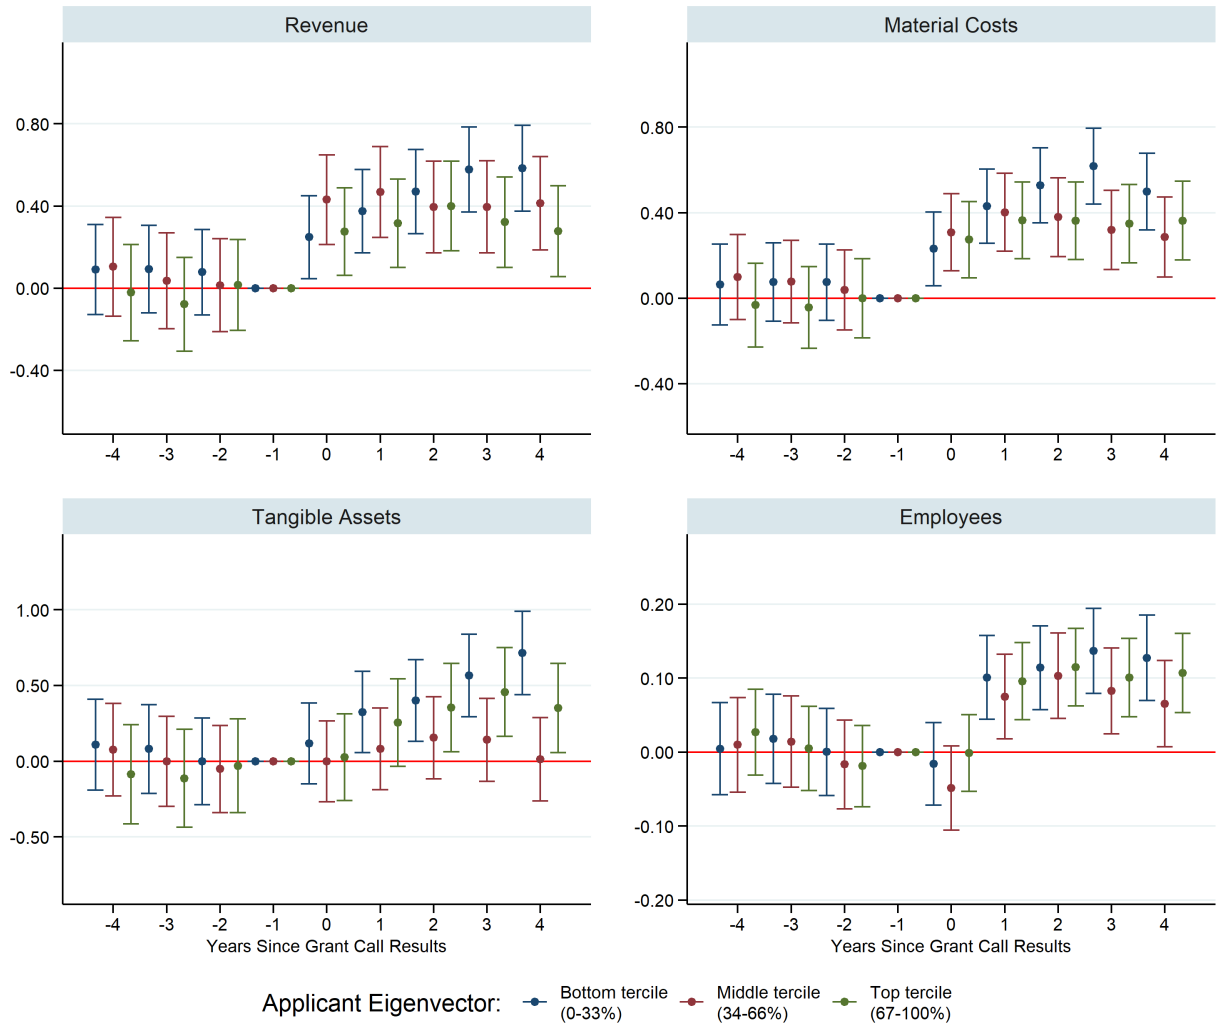


Figure A4: Upstream spillovers on suppliers, split by applicant eigenvector centrality

Note: Event-study estimates of supplier outcomes (revenue, material costs, tangible assets, and employment) as in Figure 1, estimated separately by terciles of the grant recipient's eigenvector centrality in the pre-award firm-to-firm transaction network. Eigenvector centrality weights a firm's connections by the centrality of its neighbors recursively, and corresponds to the limiting case of Katz-Bonacich centrality as indirect network effects receive maximum weight (**?**); it is the firm-level analogue of the influence vector that governs aggregate shock transmission in Acemoglu et al. (2012). The omitted reference period is year −1. Whiskers denote 95% confidence intervals.

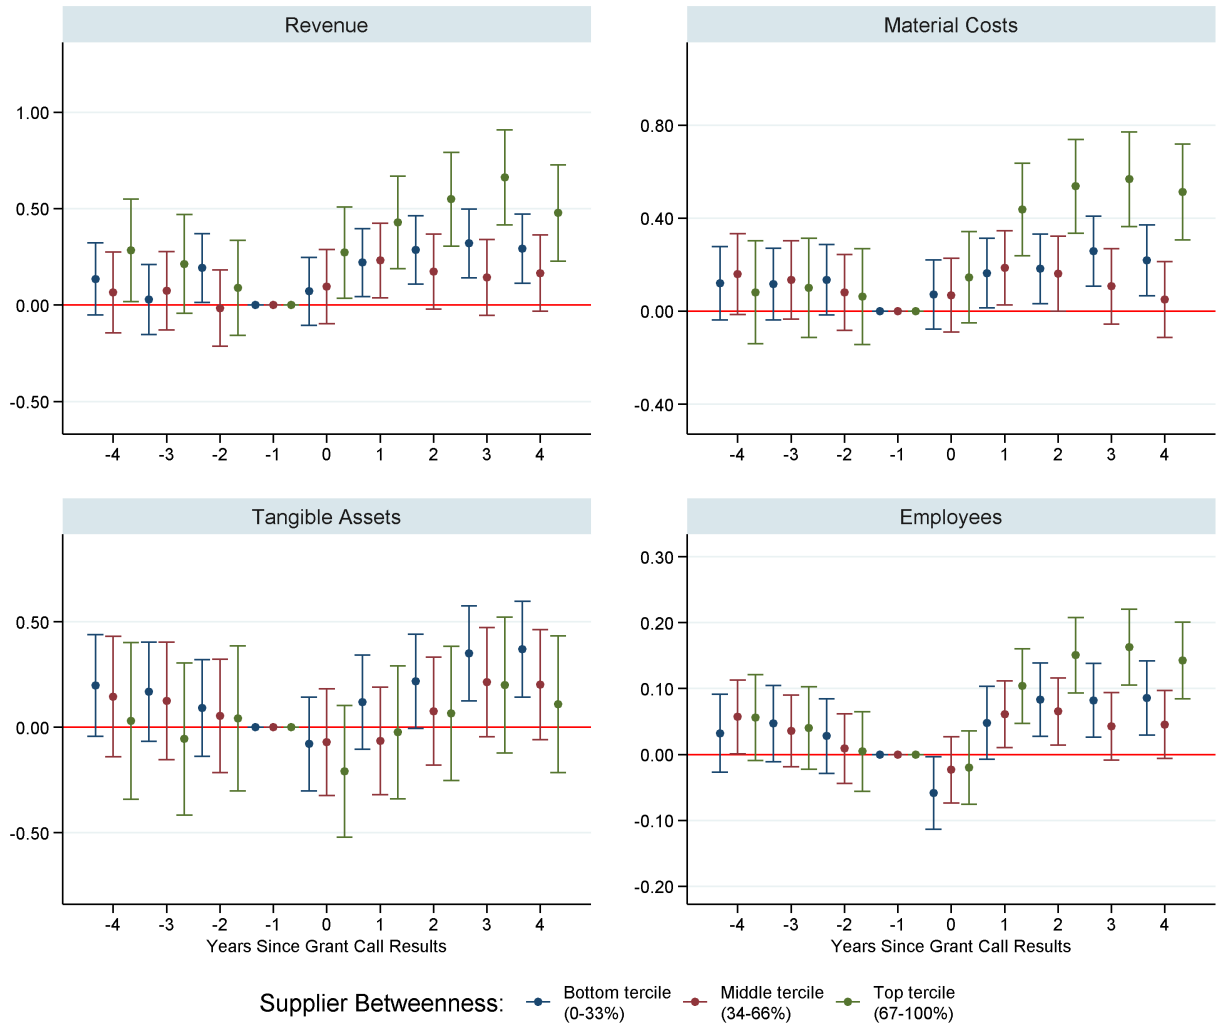


Figure A5: Upstream spillovers on suppliers, split by supplier betweenness centrality

Note: Event-study estimates of supplier outcomes (revenue, material costs, tangible assets, and employment) as in Figure 1, estimated separately by terciles of the supplier's betweenness centrality in the pre-award firm-to-firm transaction network. Betweenness centrality measures how often the supplier lies on the shortest path between other firms in the network, capturing its structural role as a broker and its capacity to re-transmit grant-induced demand shocks to other parts of the economy. The omitted reference period is year −1. Whiskers denote 95% confidence intervals.

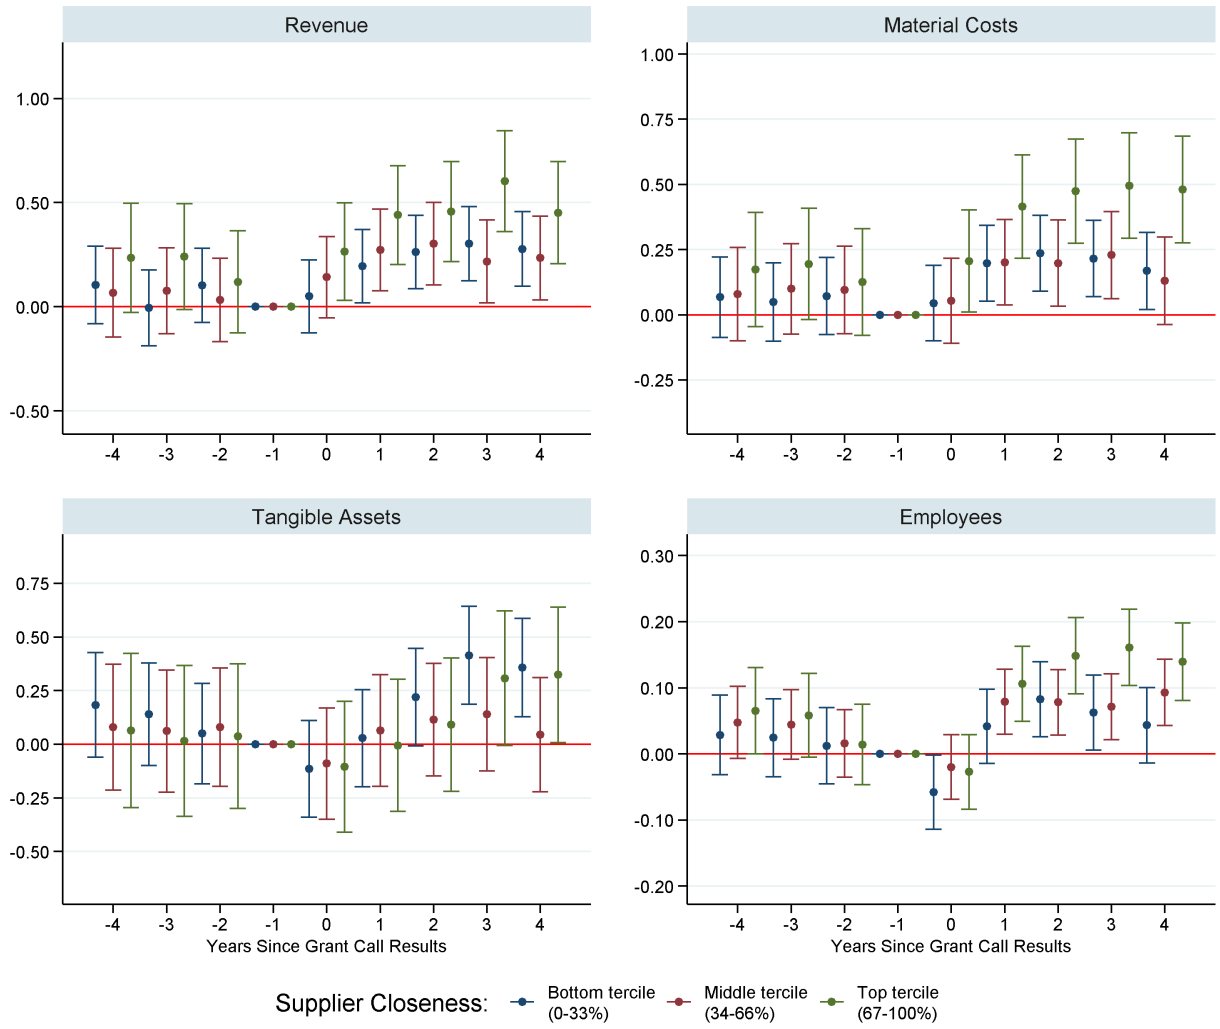


Figure A6: Upstream spillovers on suppliers, split by supplier closeness centrality

Note: Event-study estimates of supplier outcomes (revenue, material costs, tangible assets, and employment) as in Figure 1, estimated separately by terciles of the supplier's closeness centrality in the pre-award firm-to-firm transaction network. Closeness centrality is the inverse of the supplier's average distance to all other firms in the network, reflecting how rapidly grant-induced demand shocks can propagate outward from the supplier to the rest of the economy. The omitted reference period is year −1. Whiskers denote 95% confidence intervals.

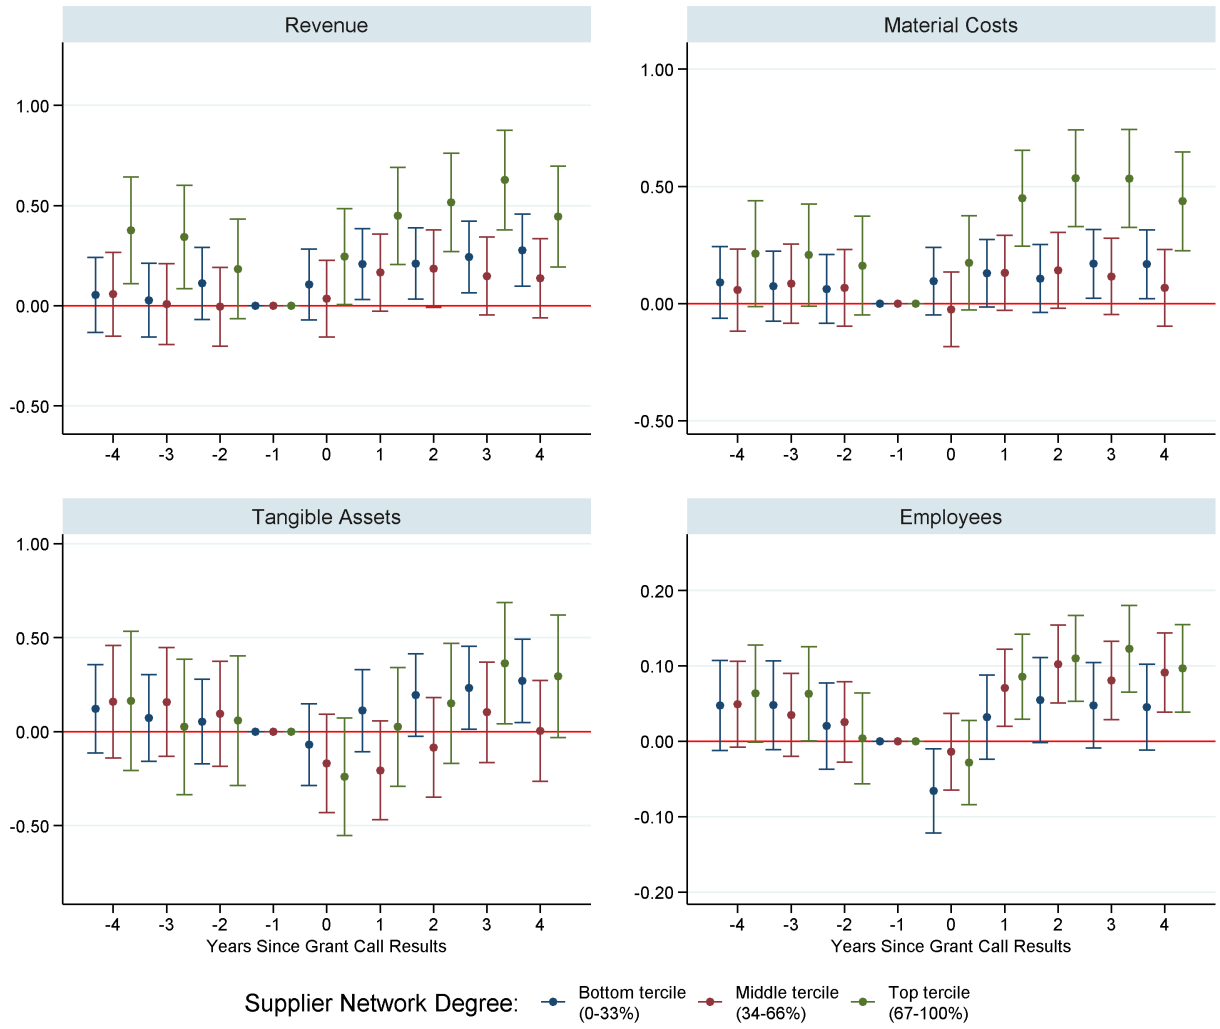


Figure A7: Upstream spillovers on suppliers, split by supplier degree centrality

Note: Event-study estimates of supplier outcomes (revenue, material costs, tangible assets, and employment) as in Figure 1, estimated separately by terciles of the supplier's degree centrality in the pre-award firm-to-firm transaction network. Degree centrality counts the supplier's number of direct trading partners, reflecting the breadth of its market relationships, its capacity to absorb demand expansions across a diversified customer base, and its access to diverse upstream sourcing that facilitates further propagation of grant-induced demand. The omitted reference period is year −1. Whiskers denote 95% confidence intervals.

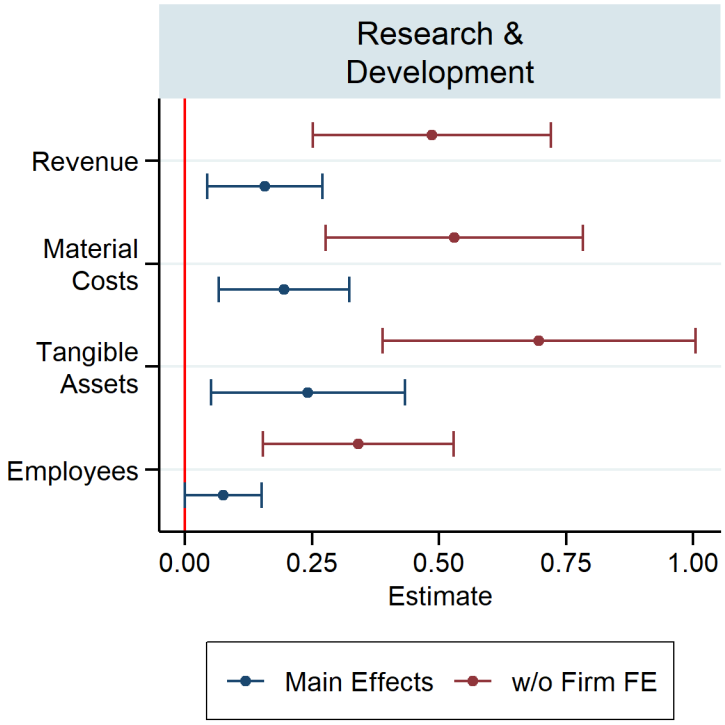


Figure A8: Direct effects of R&D grants: with and without firm fixed effects

Note: This figure reports estimated post-award percent effects of grant receipt on applicant outcomes (revenue, material costs, tangible assets, and employment) in the matched applicant sample. The "Main Effects" specification corresponds to the baseline model that includes firm fixed effects and local industry by region by year fixed effects. The "w/o Firm FE" specification omits firm fixed effects while otherwise maintaining the same outcome definitions, post-award indicator, and inference procedure. Whiskers denote 95% confidence intervals.

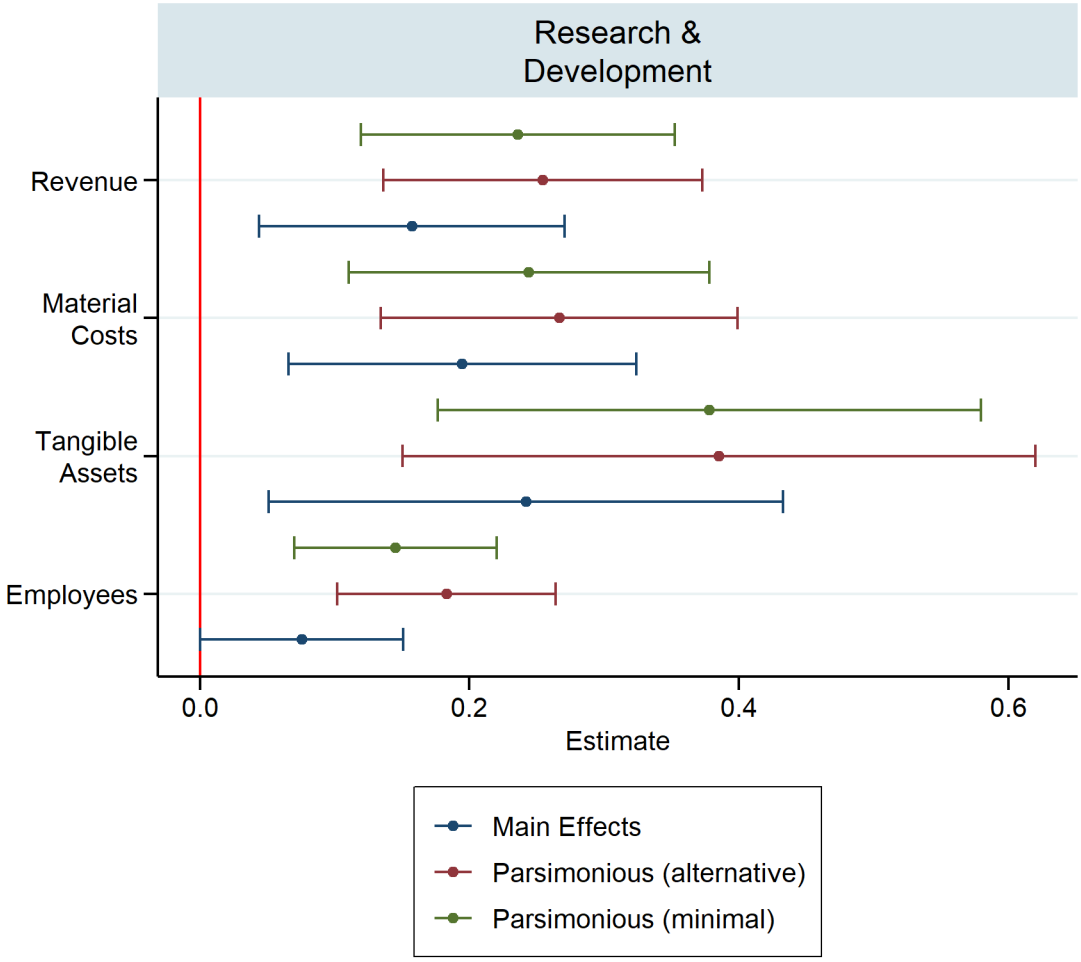


Figure A9: Direct effects of R&D grants: parsimonious matching

Note: This figure compares estimated post-award percent effects of grant receipt across three matched-sample constructions: (i) the baseline, exceptionally detailed matching procedure; (ii) a parsimonious procedure that uses roughly 50% fewer pre-treatment matching variables; and (iii) a minimal procedure that conditions only on eight baseline variables. The figure shows that estimated effects are larger under less detailed matching, consistent with residual positive selection when fewer pre-treatment characteristics are used to construct the counterfactual. Whiskers denote 95% confidence intervals.

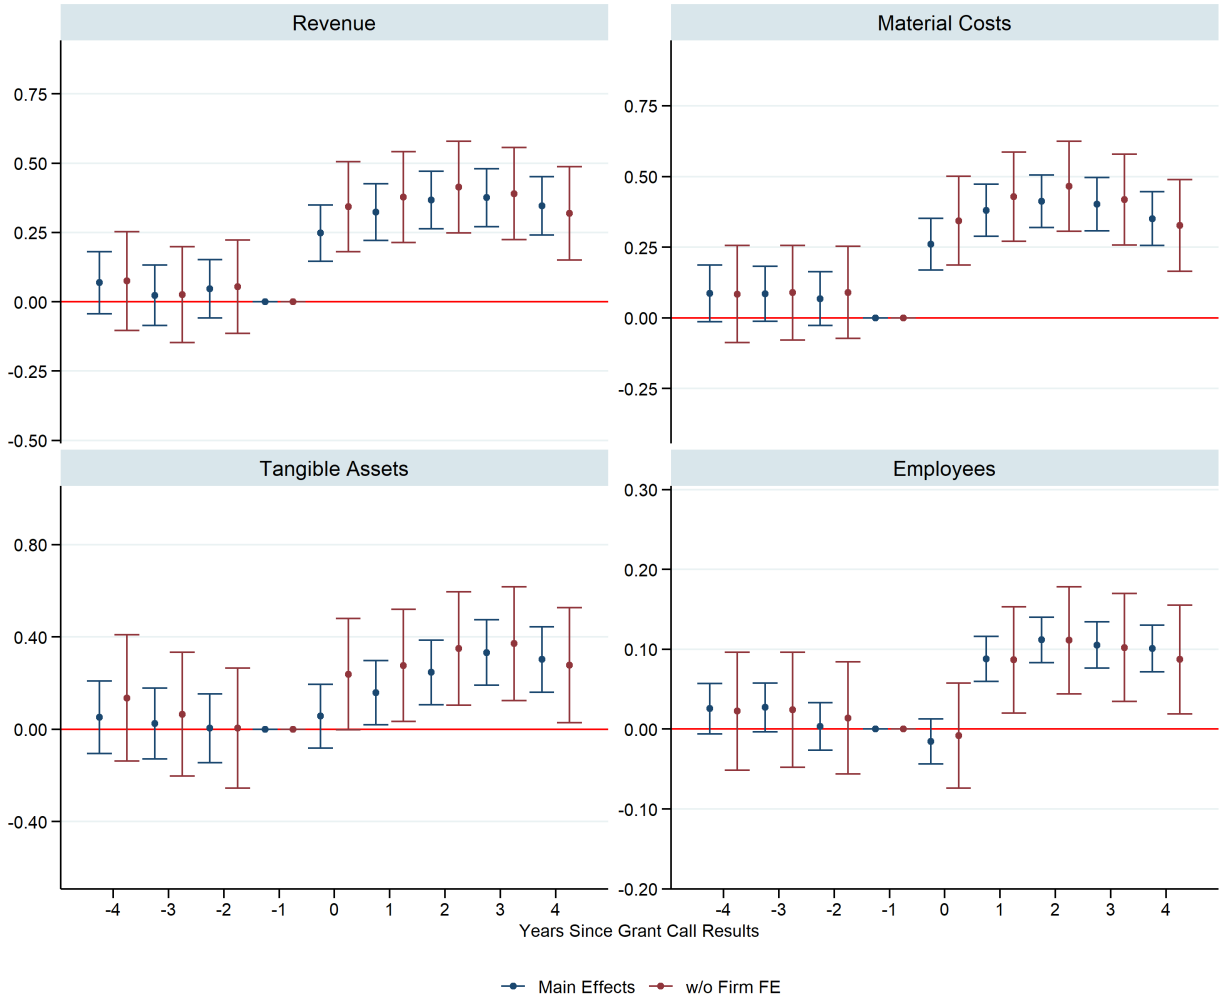


Figure A10: Supplier effects: with and without firm fixed effects

Note: This figure reports supplier event-study estimates of the upstream effects of buyers' R&D grant receipt on supplier outcomes (revenue, material costs, tangible assets, and employment). Points show coefficient estimates for event time relative to the buyer's grant award; whiskers denote 95% confidence intervals. The figure contrasts an otherwise identical specification estimated with supplier fixed effects versus without supplier fixed effects. The two specifications deliver similar point estimates, while confidence intervals are tighter in the fixed-effects specification.

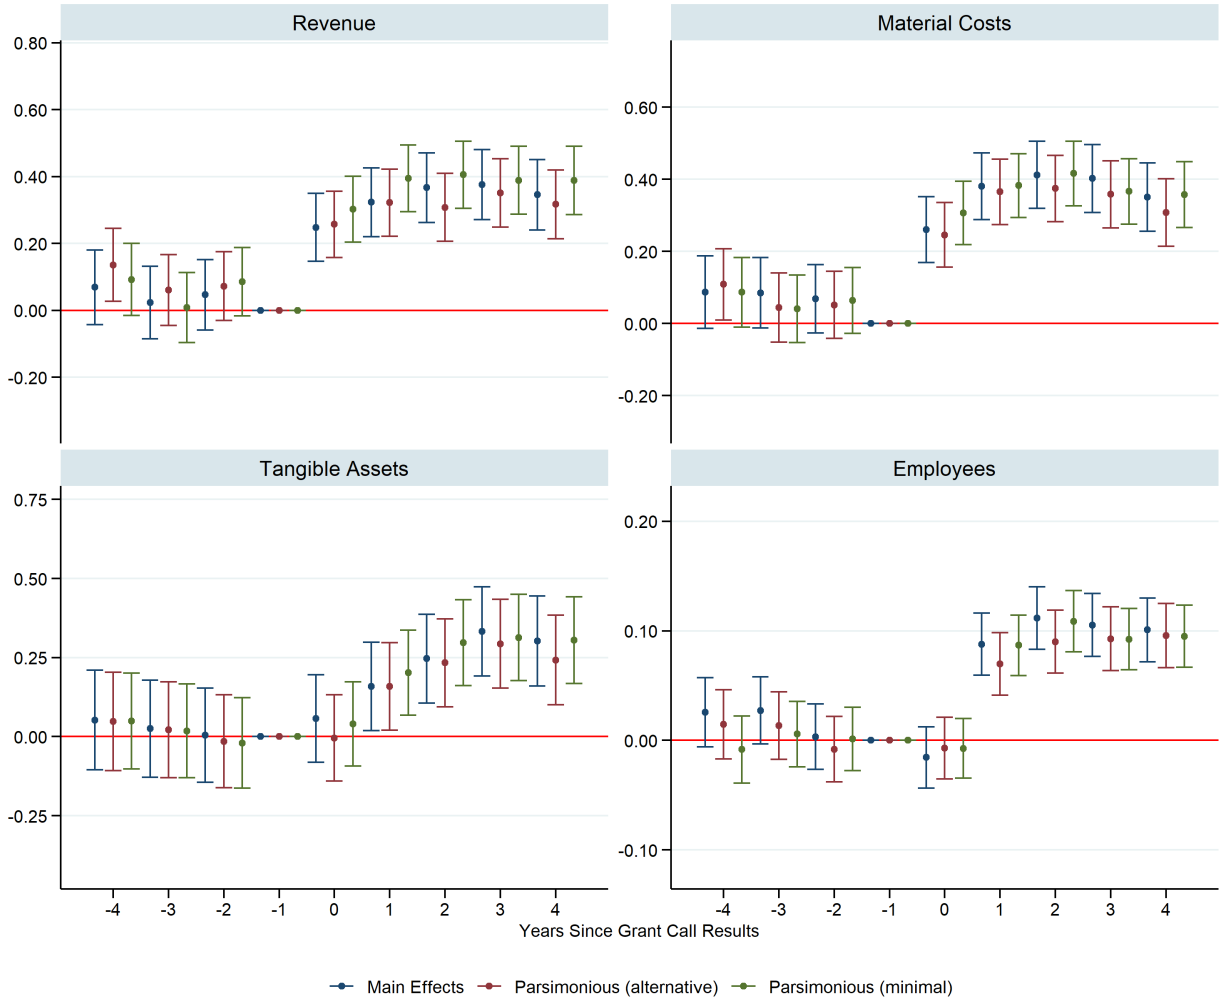


Figure A11: Supplier effects: parsimonious matching

Note: This figure compares supplier event-study estimates obtained under three alternative constructions of the matched applicant sample used to define treated buyers and their counterfactuals. The baseline results use the main (high-dimensional) matching procedure. The two alternatives are (i) a parsimonious matching specification that uses roughly 50% fewer pre-treatment covariates than the baseline, and (ii) a minimal matching specification that conditions only on eight baseline covariates. Points show event-time coefficients and 95% confidence intervals. Estimates are very similar across the three matching approaches, indicating that upstream spillovers are not sensitive to the richness of the matching design.

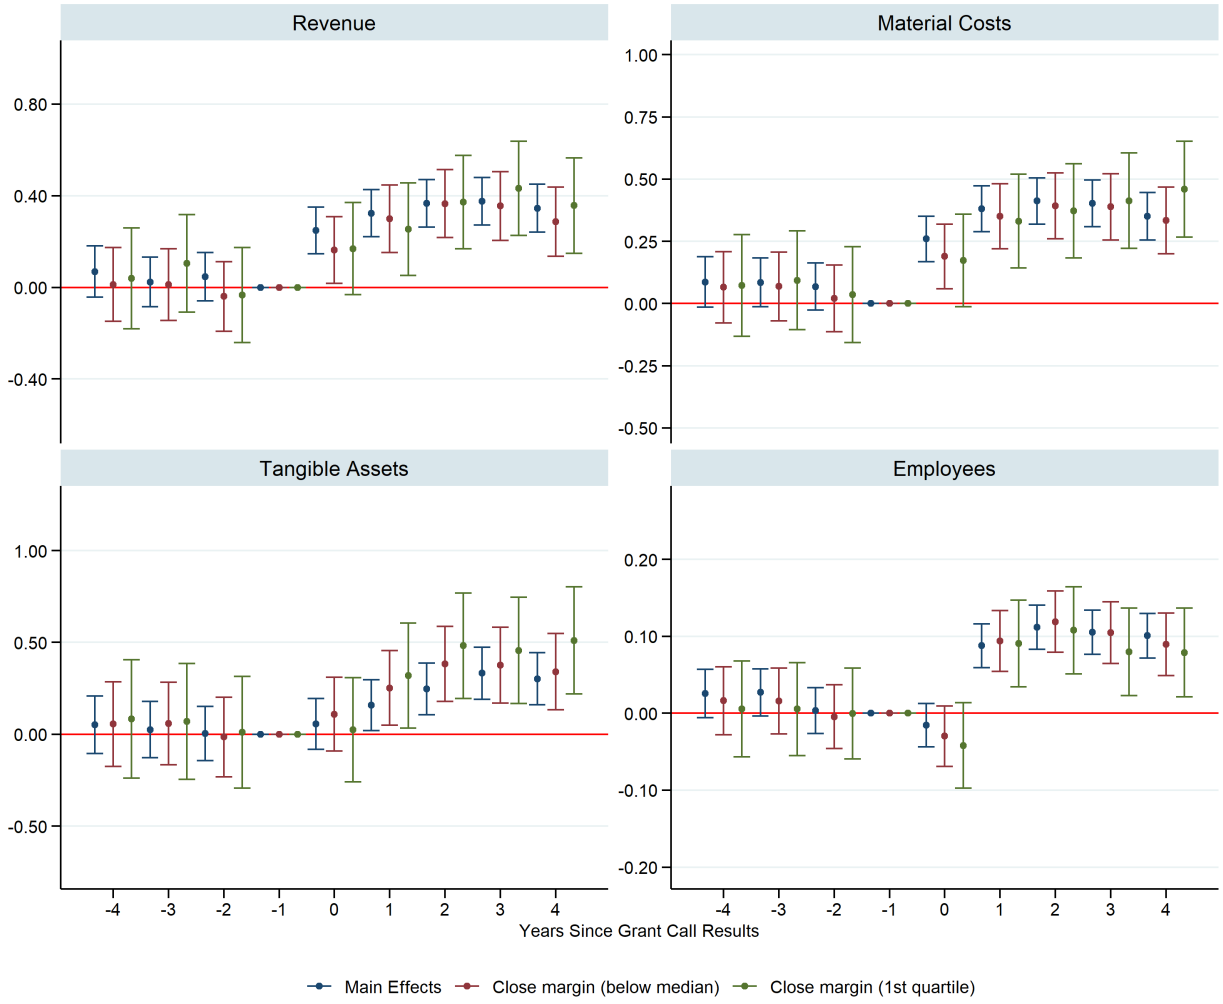


Figure A12: Supplier effects: results for the top 25% and top 50% best graded projects

Note: This figure assesses whether upstream spillovers are concentrated among the highest-quality funded projects. Using the evaluation score assigned to each funded application, we re-estimate supplier event-study effects for suppliers linked to the top 25% best-graded funded projects and, separately, for suppliers linked to the top 50% best-graded funded projects. The figure compares these estimates to the baseline supplier effects. Points show event-time coefficients and 95% confidence intervals. The similarity of estimates across the top-score subsamples and the baseline indicates no evidence of a "quality premium" in supplier spillovers.

B Difference-in-Discontinuities Design

B.1 Multi threshold scoring

The program assigns grants according to a deterministic eligibility rule that generates a sharp discontinuity at the eligibility frontier. Under the standard regression discontinuity (RD) continuity assumption, applicants located sufficiently close to the cutoff are comparable in all respects other than treatment status, allowing discontinuities in outcomes at the threshold to be interpreted causally (Lee and Lemieux, 2010). While the assignment mechanism is cross-sectional, outcomes are observed over multiple periods for each firm. We therefore implement a difference-in-discontinuities framework that exploits within-firm variation over time while preserving the local comparison central to RD designs. Estimation is performed on samples restricted to a neighborhood of the cutoff, reflecting the discrete support of the running variable and the requirement that identification remain local.

A complication is that eligibility depends on multiple scoring thresholds rather than a single scalar index. To implement an RD framework consistent with this institutional structure, we construct a one-dimensional running variable that captures an application's proximity to failing eligibility. Let S_{ik} denote applicant i's score on criterion k, and C_k as the corresponding threshold. We define the normalized margin to each cutoff as

$$M_{ik} = \frac{(S_{ik} \; - \; C_k)}{C_k} \tag{14}$$

and define the running variable as the most binding normalized margin:

$$R_i = \min_k M_{ik} \tag{15}$$

By construction, $R_i \geq 0$ if and only if the applicant passes all thresholds, $R_i < 0$ if at least one threshold is missed, and $R_i = 0$ at the eligibility frontier. Table A11 reports the distribution of binding constraints across applications, where the binding constraint is the evaluation criterion that determines an application's position relative to the eligibility frontier. The pattern differs sharply by funding status. For funded applicants, the binding constraint is always the *Total Score* (100%), indicating that awarded projects satisfy all criterion-specific minimum thresholds and that variation in proximity to the cutoff is governed by the aggregate score requirement. For non-funded applicants, in contrast, the binding constraint is rarely the *Total Score* (0.8%). Instead, ineligibility is driven overwhelmingly by failing criterion-specific minimum standards, most commonly *Project Design and Maturity* (42.9%) and *Implementation Capacity* (38.3%), followed by *Value*

for Money (9.0%), with the remaining criteria accounting for only small shares. Overall, the table indicates that rejected proposals typically fail on specific evaluation dimensions rather than narrowly missing the overall score threshold, a feature that contributes to the discrete support and uneven density of the running variable around the cutoff.

Table A11: Binding Constraint Distribution (Before Filtering)

Funded	Binding Constraint	N	Percentage
Yes	Total Score	178	100.0
No	Project Design and Maturity	57	42.9
No	Implementation Capacity	51	38.3
No	Sustainability	4	3.0
No	Novelty Significance	6	4.5
No	Value for Money	12	9.0
No	Total Score	1	0.8
No	Potential for Commercialization	2	1.5

The binding-constraint distribution clarifies how applications sort around the eligibility frontier and motivates our use of a panel-based RD implementation. Because firm outcomes are observed over multiple periods and are highly persistent, we estimate direct effects using a difference-in-discontinuities specification with firm and time fixed effects:

$$Y_{it} = \eta_i + \lambda_t + \tau \left(D_i \times \text{Post}_t \right) + \varepsilon_{it} \tag{16}$$

Here, η_i are firm fixed effects and λ_t are time fixed effects. The coefficient τ captures the causal effect of grant receipt at the eligibility frontier under the maintained smoothness assumption.

Because RD identification is local, we report estimates for (i) the full estimation sample and (ii) progressively more local neighborhoods around the cutoff. Specifically, we re-estimate the model on subsamples defined by

$$R_i \in [-h_L, h_R], \tag{17}$$

Where $h_L > 0$ and $h_R > 0$ denote the bandwidths to the left and right of the cutoff. Our baseline uses symmetric windows with $h_L = h_R = h$, and we assess robustness to asymmetric windows ($h_L \neq h_R$) motivated by the discrete support of the running variable and the uneven density of observations on either side of the threshold.

Because the running variable is discrete, observations cannot be arbitrarily close to the cutoff. In our data, the closest support points around the eligibility frontier are $R_i = -1$ on the untreated side and $R_i = 0$

on the treated side, implying that one unit is the smallest meaningful distance in R_i for local comparisons at the threshold. We therefore use h=1 as a minimally local bandwidth that compares marginally ineligible applicants $R_i = -1$ to marginally eligible applicants ($R_i \in \{0, 0.5, 1\}$). However, the discrete support and uneven density of observations across the two sides of the cutoff imply that very narrow windows can yield small effective sample sizes and imprecise estimates. To address this, we also report results using a wider symmetric window $R_i \in [-3, 3]$, which retains locality while increasing statistical power. In addition, given that the mass of observations is larger on one side of the cutoff, we consider an asymmetric window, which preserves the minimally local comparison on the sparse side while allowing a broader bandwidth on the high-density side to improve precision.

B.2 Difference-in Discontinuties Design: Results

Figure A13 plots the distribution of the normalized running variable around the grant-award cutoff. The distribution is strongly asymmetric, with substantially more mass on the funded side of the threshold and a relatively long left tail among non-funded applicants. This pattern is consistent with the institutional structure of the evaluation process and the construction of the running variable.

Eligibility requires meeting multiple criterion-specific minimum thresholds as well as an overall requirement. Because the running variable is defined by the weakest normalized margin across criteria, a single failed criterion can place an application well below the eligibility frontier even if other scores are high. As a result, non-funded applications tend to be dispersed further into negative values of the running variable, while relatively few observations appear immediately below the cutoff.

While Figure A13 is informative about the support and density of the running variable, it should not be interpreted as a formal test for sorting or manipulation around the cutoff. Rather, the asymmetry and the scarcity of observations just below zero highlight practical features of the setting—discrete support and uneven density—that motivate treating the regression discontinuity exercise as a local benchmark and complementing it with a higher-powered matched-sample design.

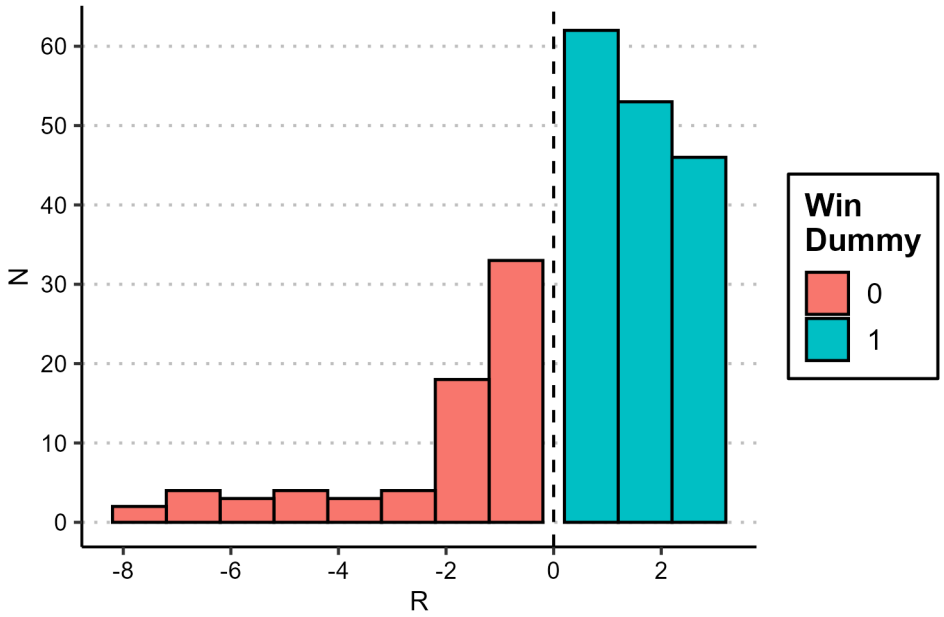


Figure A13: Distribution of the normalized running variable around the grant-award cutoff

Note: This figure plots a histogram of the normalized running variable R. The variable is centered so that $R = 0$ corresponds to the grant-award threshold (vertical dashed line). Bars are shown separately by grant-award status (Win Dummy $= 1$ for awarded applicants; Win Dummy $= 0$ for non-awarded applicants).

Table A12 reports summary statistics for applicants that received evaluation scores, in logs for monetary outcomes. In the full scored sample, the data include 144 funded and 65 non-funded applicants. Funded firms are systematically larger along several observable dimensions. In the pre-award period t-1, recipients exhibit higher revenue (14.864 vs. 13.820, $p < 0.01$), material costs (14.130 vs. 12.814, $p < 0.01$), tangible assets (12.657 vs. 11.424, $p < 0.01$), and employment (3.807 vs. 3.116, $p < 0.01$). Labor productivity is also higher among recipients (11.140 vs. 10.834, $p < 0.05$), while average wages and TFP do not differ significantly. These differences are consistent with selection into grant receipt in the broader scored sample.

Restricting attention to applicants closer to the eligibility frontier substantially changes this comparison. In Close sample 1, defined by the symmetric window ($R \in [-3, 3]$), the sample includes 144 funded and 50 non-funded applicants. Within this window, differences in mean revenue, tangible assets, employment, wages, labor productivity, and TFP are no longer statistically significant; the only remaining difference is a modest gap in material costs (14.130 vs. 13.430, significant at $p < 0.01$). Tightening the restriction further reinforces this pattern. In Close sample 2 ($R \in [-3, 1]$); 56 funded and 50 non-funded) and Close sample 3 ($R \in [-1, 1]$); 56 funded and 26 non-funded), we find no statistically significant differences in baseline firm characteristics across treatment status.

Overall, the attenuation of observable differences as the sample is restricted around the cutoff is consistent with the identifying assumptions of the regression discontinuity design. While recipients and non-recipients differ in the full scored sample, applicants near the eligibility frontier appear comparable in terms of scale, input use, asset base, labor, and wages, supporting the interpretation that variation in grant receipt close to the cutoff is plausibly orthogonal to underlying firm characteristics.

Table A12: Summary statistics by treatment status

Variable	Full Sample		Close Sample 1		Close Sample 2		Close sample 3	
	T	C	T	C	T	C	T	C
Revenue	14.864	13.820***	14.864	13.820	14.850	14.376	14.850	14.214
Material costs	14.130	12.814***	14.130	13.430*	14.181	13.430	14.181	13.170
Tangible assets	12.657	11.424***	12.657	12.153	12.606	12.153	12.606	12.058
Employees	3.807	3.116***	3.807	3.399	3.754	3.399	3.754	3.306
Average wage	9.429	9.333	9.428	9.480	9.410	9.480	9.410	9.497
Labor productivity	11.140	10.834**	11.140	11.053	11.162	11.053	11.162	10.947
TFP	8.866	8.920	8.866	8.633	8.743	8.633	8.743	8.281
N	144	65	144	50	56	50	56	26

Note: monetary values given in logs. Full sample includes all firm applying and receiving a score in the evaluation procedure. Close sample 1 includes firms within $R \in [-3, 3]$. Close sample 2 includes firms within $R \in [-3, 1]$. Close sample 3 includes firms within $R \in [-1, 1]$. Statistically significant difference in means is highlighted with $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Figure A14 provides a descriptive visualization of applicant outcomes around the grant-award cutoff using post-award observations. The figure plots binned averages of (log) revenue, material costs, tangible assets, and employment against the normalized running variable, separately on either side of the eligibility frontier, together with local linear fits. By construction, the horizontal axis measures applicants' distance to the cutoff, while the vertical axis reports realized outcomes in the post-award period.

Across all four outcomes, post-award averages are higher on the funded side of the cutoff than on the non-funded side, a pattern qualitatively consistent with grant receipt increasing firm scale and input use. Interpretation of the figure should nevertheless remain cautious. The running variable has discrete support and there are relatively few observations close to the cutoff—especially immediately below zero—so the plot is not well suited for fine-grained smoothness diagnostics. Moreover, because the plotted outcomes are realized post-award outcomes, a discontinuity at the frontier is consistent with treatment effects by construction rather than providing a direct visual test of the identifying assumption. We therefore treat Figure A14 as an illustrative complement to the formal difference-in-discontinuities estimates in Table A13, and we rely primarily on balance in predetermined characteristics in local windows around the cutoff (Table A12) as

supporting evidence for the RD continuity assumption.

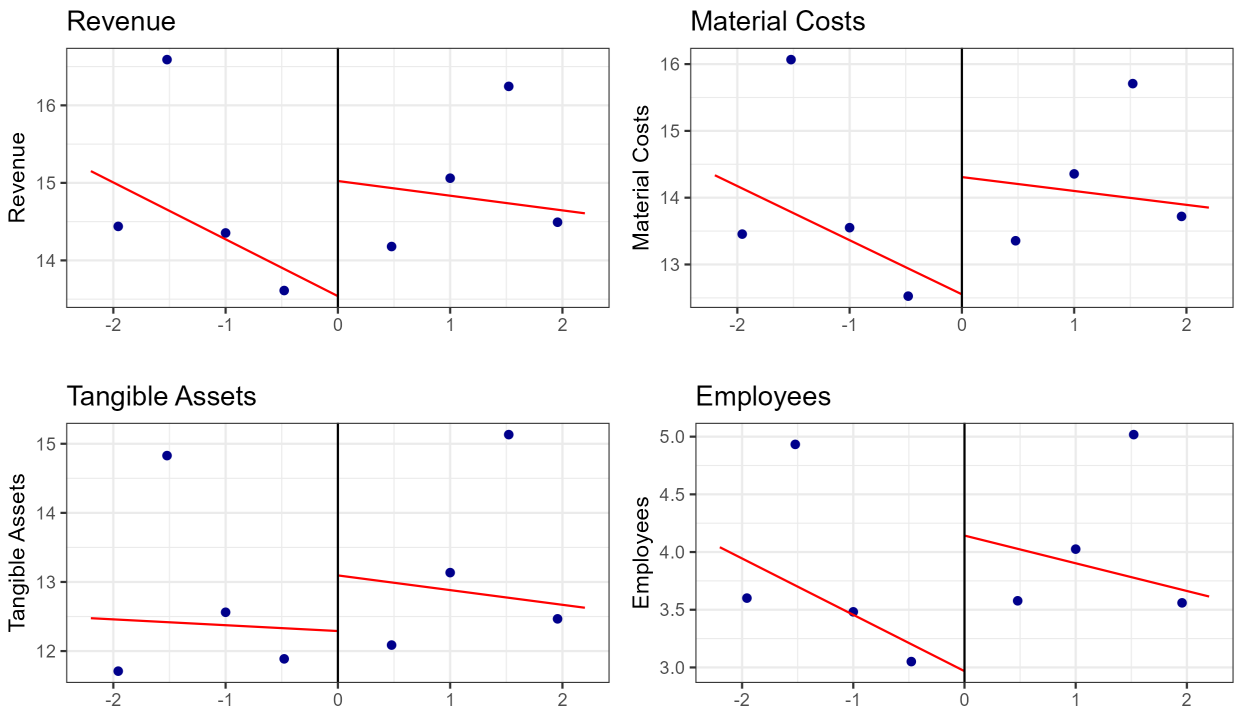


Figure A14: Applicant outcomes around the grant-award cutoff

Note: This figure presents binned scatter plots of applicant outcomes against the normalized running variable R used in the regression discontinuity sensitivity analysis. The vertical line at $R = 0$ denotes the grant-award cutoff. Each dot represents the mean outcome within a bin of the running variable, and the solid lines show separate local linear fits on either side of the cutoff. Panels report (log) Revenue, Material Costs, Tangible Assets, and Employees; log differences can be interpreted approximately as percentage differences.

Table A13 reports difference-in-discontinuities (DiDisc) estimates of the direct effect of R&D grants on applicant outcomes at the grant-award cutoff. Across specifications, the estimated effects are economically meaningful for firm scale and input use, and the patterns are broadly stable as the sample is restricted around the eligibility frontier.

In the full sample, grant receipt increases revenue by 0.193 log points and material costs by 0.147 log points, and raises employment by 0.136 log points. Interpreted as approximate percentage changes, these estimates correspond to about a 21% increase in revenue, a 16% increase in material costs, and a 15% increase in employment. The tangible asset estimate is positive (0.135 log points, about 15%) but not statistically distinguishable from zero. The full-sample specification uses 1,853 observations.

The magnitudes remain similar and, for several outcomes, become larger as the analysis focuses on applicants closer to the cutoff. In Close sample 1, revenue increases by 0.212 log points (about 24%) and material costs by 0.196 log points (about 22%). The tangible asset response rises to 0.285 log points (about 33%), and employment increases by 0.111 log points (about 12%). In Close sample 2, the estimates imply further strengthening along the capital margin: revenue increases by 0.252 log points (about 29%) and material costs by 0.219 log points (about 25%), while tangible assets increase by 0.429 log points (about 54%). Employment increases by 0.099 log points (about 10%). Close sample 1 and Close sample 2 use 1,721 and 936 observations, respectively.

Finally, in Close sample 3, which provides the most local comparison, revenue remains positive at 0.185 log points (about 20%) but is estimated less precisely, while material costs (0.216 log points, about 24%), tangible assets (0.339 log points, about 40%), and employment (0.146 log points, about 16%) are all positive and statistically significant at conventional levels for material costs, assets, and especially employment. This narrow-window specification uses 750 observations.

Table A13: Difference in discontinuities

	Revenue	Material costs	Tangible assets	Employees	N
Full sample					
	0.193**	0.147*	0.135	0.136***	1853
	(0.089)	(0.081)	(0.143)	(0.039)	
Close sample 1					
	0.212**	0.196**	0.285*	0.111**	1721
	(0.095)	(0.087)	(0.146)	(0.040)	
Close sample 2					
	0.252**	0.219*	0.429***	0.099**	936
	(0.115)	(0.112)	(0.160)	(0.049)	
Close sample 3					
	0.185	0.216*	0.339*	0.146***	750
	(0.126)	(0.124)	(0.177)	(0.056)	

Note: This table reports difference-in-discontinuities (DiDisc) estimates of the direct effect of R&D grants on applicant outcomes, using the evaluation threshold as a quasi-experimental cutoff. *$p < 0.1$; **$p < 0.05$; ***$p < 0.01$

Two patterns stand out. First, the joint increase in revenue and material costs indicates that grants translate into higher production and greater intermediate input demand, rather than purely financial or accounting changes. Second, the tangible asset response becomes much larger as the sample is restricted toward the cutoff, suggesting that marginal recipients close to the eligibility frontier adjust particularly strongly along the capital margin, consistent with capacity expansion following grant receipt.

At the same time, the regression discontinuity design is unavoidably local and inherits several limitations

from the institutional setting and the discrete running variable. Because observations cannot be arbitrarily close to the cutoff, the effective sample near the eligibility frontier is small, which limits statistical power and makes it difficult to study dynamics or systematically explore heterogeneous effects. In addition, the density of observations is uneven around the frontier, with substantially more firms on the funded side than on the non-funded side in narrow windows. This imbalance further reduces precision and raises sensitivity to bandwidth choices. Finally, while the attenuation of observable differences near the cutoff is reassuring, we cannot rule out residual sorting around the threshold on unobserved dimensions, for example through application behavior, project selection, or administrative discretion that is not fully captured in the data.

To address these concerns, we extend the analysis using the matched-sample framework described in Section 4. This approach applies the same empirical logic to the full distribution of funded applicants and constructs comparison firms from the universe of non-funded firms in the economy. The administrative data allow us to match on a rich set of pre-treatment characteristics, substantially expanding the sample and enabling more precise estimation of dynamic responses and production-network spillovers while reducing reliance on the narrow local variation at the cutoff.

C Local Randomization Inference: Sensitivity Analysis

This appendix complements the difference-in-discontinuities analysis of Section **??** with local randomization inference following Cattaneo et al. (2015). The local randomization framework provides exact finite-sample inference within a narrow window around the cutoff under the assumption that, locally, treatment assignment is as good as random. Unlike standard regression discontinuity, this approach does not rely on asymptotic approximations that may be unreliable at our sample size, and it is specifically designed for settings with discrete running variables and small effective samples around the cutoff. Implementation uses the `rdlocrand` package of Cattaneo et al. (2016).

C.1 Window Selection

The local randomization framework requires selecting a window around the cutoff in which pre-treatment covariates are balanced between the treated and control sides. We implement the data-driven window selection procedure of `rdwinselect`, which steps outward from the cutoff in increments of 0.25 and reports, at each candidate window, the smallest balance test p-value across a vector of pre-treatment covariates. Our

covariate vector includes eleven pre-treatment characteristics: firm employment, revenue, material costs, capital, total factor productivity, labor productivity, average wage, contracts in non-financial transactions, contracts at high-value and low-value tiers, and the share of high-skilled employment.[7]

Table A14 reports the window-selection diagnostic. The first three windows (± 0.5 to ± 1.0) cannot be evaluated because they contain zero observations on the left side of the cutoff—a direct consequence of the asymmetric density of the running variable noted in Section **??**. From ± 1.25 onward, the joint balance test runs and is initially driven by pre-treatment material costs. The minimum balance p-value crosses the conventional 0.15 threshold at the window $[-1.75, 1.75]$, where it equals 0.184 with 33 control and 59 treated firms. We adopt this as our preferred window: it is the narrowest interval in which joint covariate balance holds and provides a defensible foundation for randomization-based inference.

Table A14: Window selection via covariate balance

Window	Min. p-value	Variable	Binom. test	Obs $< c$	Obs $\geq c$
$[-1.25, 1.25]$	0.100	Material costs (t-1)	0.005	29	56
$[-1.50, 1.50]$	0.108	Material costs (t-1)	0.005	29	56
$[-1.75, 1.75]$	**0.184**	**Material costs (t-1)**	0.009	**33**	**59**
$[-2.00, 2.00]$	0.209	Material costs (t-1)	0.009	33	59
$[-2.25, 2.25]$	0.224	Material costs (t-1)	0.000	43	101
$[-2.50, 2.50]$	0.235	Material costs (t-1)	0.000	43	101
$[-2.75, 2.75]$	0.331	Material costs (t-1)	0.000	46	103
$[-3.00, 3.00]$	0.320	Material costs (t-1)	0.000	46	103

Notes: Window-selection diagnostic from `rdwinselect`. The minimum p-value column reports the smallest balance test p-value across eleven pre-treatment covariates at each candidate window. The chosen window $[-1.75, 1.75]$, shown in bold, is the narrowest interval where the minimum balance p-value exceeds the conventional 0.15 threshold. The first three windows cannot be evaluated due to zero observations on the left side of the cutoff. Number of observations in the full window-selection sample = 209.

C.2 Main Results: First-Differenced Outcomes

Within the chosen window, we estimate the effect of grant receipt on first-differenced outcomes constructed as $\Delta y_i = y_{i,t=+1} - y_{i,t=-1}$. We use first differences rather than post-treatment levels because pre-treatment levels are imbalanced across the cutoff in a manner reflecting selection of larger firms into successful grant applications (see Section C.4). First differencing nets out time-invariant pre-existing gaps and isolates differential growth post-grant.

Table A15 reports estimates at horizon $t = +1$ using three test statistics: difference in means, Wilcoxon

[7] We exclude pre-treatment intellectual assets from the window-selection covariate vector. This variable persistently fails balance across all candidate windows, plausibly because firms with greater intellectual capital are systematically better at writing successful grant applications. We treat this as a consequence of selection into applying rather than as a confounder for post-grant outcomes.

rank-sum, and Kolmogorov-Smirnov. All p-values are exact finite-sample p-values from 5,000 randomization draws.

Table A15: Local randomization inference: main results at $t = +1$

Outcome	Effect (log pts)	Diff. in means p-value	Rank-sum p-value	Kolmogorov-Smirnov p-value
Δ Employment	0.132	**0.072**	**0.089**	0.151
Δ Revenue	0.126	0.309	0.437	0.453
Δ Material costs	0.093	0.643	0.308	0.658
Δ Capital	0.218	0.463	**0.172**	**0.155**

Notes: Exact finite-sample p-values from 5,000 randomization draws within the window $[-1.75, 1.75]$ (33 control firms, 59 treated firms). Outcomes are first differences from $t = -1$ to $t = +1$ relative to the grant decision year. Bold indicates $p < 0.20$. Window selected via `rdwinselect` to satisfy joint covariate balance on eleven pre-treatment characteristics (Table A14).

The point estimates are positive for all four outcomes and broadly consistent in sign and rough magnitude with the matched-sample estimates of Table **??**. The headline employment effect of 0.132 log points (about 14%) is statistically significant under both the difference-in-means test ($p = 0.072$) and the rank-sum test ($p = 0.089$), and borderline under the more demanding Kolmogorov-Smirnov test ($p = 0.151$). The capital effect of 0.218 log points (about 24%) is missed by the difference-in-means test but detected by both the rank-sum ($p = 0.172$) and Kolmogorov-Smirnov ($p = 0.155$) tests, a pattern consistent with heterogeneous capital adjustment across treated firms—some respond strongly while others do not move—that distributional tests are designed to detect. Effects on revenue (+0.126) and material costs (+0.093) are positive in sign but statistically imprecise at this horizon.

C.3 Dynamic Effects Across Post-Treatment Horizons

To assess whether the cutoff-based design captures dynamic adjustment patterns, we re-estimate the main specification at four post-treatment horizons. Table A16 reports the difference-in-means point estimates and p-values across $t = +1, +2, +3, +4$.

Table A16: Local randomization inference: dynamic results across post-treatment horizons

	$t = +1$		$t = +2$		$t = +3$		$t = +4$	
Outcome	Eff.	p	Eff.	p	Eff.	p	Eff.	p
Δ Employment	0.132	**0.072**	0.181	**0.049**	0.241	**0.030**	0.294	**0.039**
Δ Revenue	0.126	0.309	0.090	0.593	0.217	0.282	0.336	0.317
Δ Material costs	0.093	0.643	−0.102	0.753	0.036	0.910	−0.022	0.946
Δ Capital	0.218	0.463	0.090	0.807	0.254	0.532	0.219	0.602

Notes: Difference-in-means estimates with exact finite-sample p-values from 5,000 randomization draws within $[-1.75, 1.75]$. Effective sample is 33 control / 59 treated firms at $t = +1, +3, +4$ and 33 / 58 at $t = +2$. Outcomes are first differences from $t = -1$ to the indicated post-horizon. Bold p-values indicate significance at the 10% level. Rank-sum and Kolmogorov-Smirnov p-values for employment range from 0.077 to 0.224 across horizons.

The employment effect strengthens monotonically across post-treatment horizons. The point estimate rises from 0.132 log points at $t = +1$ to 0.181 at $t = +2$, 0.241 at $t = +3$, and 0.294 at $t = +4$, corresponding to employment increases of approximately 14%, 20%, 27%, and 34% among marginal recipients. The difference-in-means test is statistically significant at the 10% level at every horizon and crosses the 5% threshold from $t = +2$ onward, with p-values of 0.072, 0.049, 0.030, and 0.039 respectively. The convergence of a monotonically growing point estimate with persistent statistical significance across four horizons—at a sample size where exact finite-sample inference is conservative by construction—provides strong cutoff-based evidence of grant effects on recipient employment.

Effects on revenue and capital are positive across horizons. By $t = +4$, the cross-sectional revenue effect reaches 0.336 log points and the capital effect reaches 0.219 log points—magnitudes broadly comparable to the matched-sample estimates of Table **??**. Their statistical imprecision in the cross-section reflects the limited power of randomization inference at $N = 92$ rather than absence of effects, and the gradual emergence of cross-sectional revenue and capital point estimates over time mirrors the lagged dynamics documented in the matched-sample event study (Figure **??**). Material costs remain close to zero in the cross-section throughout, plausibly because the small marginal-applicant sample masks heterogeneity in input substitution that the panel-based design can detect.

C.4 Falsification Tests

We conduct three falsification exercises to support the validity of the local randomization design.

Pre-treatment levels. Table A17 reports balance tests on pre-treatment levels of the four main outcomes within the chosen window. Treated firms have higher pre-period levels of employment, revenue, material

costs, and capital than control firms by between 0.32 and 0.78 log points, but none of these differences is statistically significant at conventional levels. The pattern reflects selection of larger firms into successful grant applications and motivates our use of first-differenced outcomes in the main specification.

Table A17: Placebo: pre-treatment levels

Pre-treatment outcome	Mean (left)	Mean (right)	Difference	p-value
Employment ($t = -1$)	3.485	3.806	0.321	0.347
Revenue ($t = -1$)	14.474	14.918	0.444	0.319
Material costs ($t = -1$)	13.486	14.268	0.782	0.181
Capital ($t = -1$)	12.395	12.721	0.326	0.628

Notes: Difference-in-means p-values from 5,000 randomization draws within the window $[-1.75, 1.75]$ (33 control, 59 treated). All outcomes are pre-treatment levels in logs.

Pre-treatment trends (parallel-trends test). Because the main specification uses first differences, the directly relevant falsification is whether pre-treatment *trends*—not levels—differ across the cutoff. Table A18 reports this parallel-trends test using the change in each outcome from $t = -2$ to $t = -1$.

Table A18: Placebo: pre-treatment trends (parallel-trends test)

Pre-treatment trend	Mean (left)	Mean (right)	Difference	p-value
Δ Employment ($t = -2 \to t = -1$)	0.154	0.166	0.012	0.912
Δ Revenue ($t = -2 \to t = -1$)	0.387	0.166	−0.221	0.420
Δ Material costs ($t = -2 \to t = -1$)	0.131	0.216	0.085	0.546
Δ Capital ($t = -2 \to t = -1$)	0.115	0.263	0.148	0.702

Notes: Difference-in-means p-values from 5,000 randomization draws within the window $[-1.75, 1.75]$ (33 control, 57 treated; two firms drop out due to missing values at $t = -2$). Outcomes are pre-period first differences in logs.

Pre-trend differences are economically small and statistically insignificant for all four outcomes. The contrast with the main results is particularly striking for employment: the post-treatment effect at $t = +1$ of 0.132 log points ($p = 0.072$) is more than ten times larger than the pre-treatment trend differential of 0.012 log points ($p = 0.912$), and only the post-treatment difference is statistically detectable. For revenue, the pre-trend differential is in the *opposite* direction from the post-treatment effect (−0.221 vs. +0.126), implying that any pre-trend bias would attenuate rather than inflate the post-treatment estimate. The combination of differencing in the main specification and the absence of pre-trend differences in this falsification rules out the concern that the post-treatment effects reflect pre-existing trajectories rather than the grant.

Artificial cutoffs. As a third falsification, we re-estimate the main specification at artificial cutoffs $c = -2$ and $c = +2$, restricting attention to firms on one side of the true cutoff so that the "treated" group at the artificial cutoff is genuinely untreated. Table A19 reports the results.

Table A19: Placebo: artificial cutoffs

Fake cutoff	Outcome	Effect	p-value	n_L	n_R
$c = -2$	Δ Employment	-0.110	0.240	18	33
	Δ Revenue	0.147	0.376	18	33
	Δ Material costs	0.085	0.792	18	33
	Δ Capital	0.589	0.173	18	33
$c = +2$	Δ Employment	-0.088	0.101	55	85
	Δ Revenue	-0.071	0.547	55	85
	Δ Material costs	0.057	0.650	55	85
	Δ Capital	0.110	0.693	55	85

Notes: Difference-in-means p-values from 5,000 randomization draws within windows of half-width 1.75 around each artificial cutoff. The estimation sample for each fake cutoff is restricted to one side of the true cutoff so that the artificial "treated" group is genuinely untreated. n_L and n_R denote effective sample sizes below and above the artificial cutoff.

None of the eight placebo estimates reaches the 10% significance level. The closest result is employment at $c = +2$ ($p = 0.101$), but the point estimate is -0.088—*opposite in sign* from the post-treatment effect at the true cutoff. The contrast is sharpest for employment: the true cutoff produces a positive effect of $+0.132$ log points at $t = +1$, while the two fake cutoffs produce small *negative* effects of -0.110 and -0.088. This indicates that the post-treatment employment effect is specific to the true eligibility frontier rather than a generic artifact of dividing the running variable at any arbitrary point.

C.5 Discussion

Local randomization inference provides cutoff-based evidence that complements the matched-sample design under markedly different identifying assumptions. Within a covariate-balanced window of 92 marginal applicants, exact finite-sample inference detects a positive employment effect of grant receipt that strengthens monotonically across four post-treatment horizons (from +14% at $t = +1$ to +34% at $t = +4$) and remains statistically significant at the 10% level under the difference-in-means test at every horizon. Effects on revenue and capital are positive in sign across horizons and reach magnitudes by $t = +4$ that are comparable to the matched-sample estimates, but they remain statistically imprecise in the cross-section. Three independent falsification exercises—pre-treatment levels, pre-treatment trends, and artificial cutoffs—provide no evidence of spurious discontinuities that could explain the post-treatment results.

The convergence of point estimates across the matched-sample design (Section **??**) and the local randomization inference design—each relying on distinct identifying assumptions—supports the causal interpretation of grant effects on recipient activity. At the same time, the weaker statistical precision of the cutoff-based design at $N = 92$ reinforces our reliance on the matched-sample design as the paper's primary identification

strategy: the larger effective sample and panel structure permit precise estimation of dynamic responses and production-network spillovers that the local randomization framework cannot deliver in this setting.

D Network Exposure

D.1 Network Exposure: Empirical Approach

To identify upstream spillovers from grant-funded applicants to their suppliers, we exploit variation in suppliers' pre-treatment exposure to applicant firms combined with the discontinuity-inducing structure of the grant assignment rule. The empirical design is motivated by the observation that suppliers are not directly scored by the program; instead, they are indirectly affected through their trading relationships with applicants. Identification therefore relies on cross-sectional variation in suppliers' connections to firms whose outcomes are shifted by grant receipt.

Supplier exposure is defined using administrative transaction data observed prior to the grant decision. Let W_i denote supplier i's pre-treatment exposure to applicant firms, measured as the share of the supplier's total pre-award sales accounted for by applicants included in the cutoff-based sample. Because W_i is constructed entirely from pre-award transactions, it is predetermined with respect to grant receipt and subsequent outcomes. This construction isolates variation arising from baseline production linkages rather than endogenous post-award adjustments in trading relationships, a standard concern in production-network settings.

We estimate supplier responses using a panel specification that relates supplier outcomes to interactions between exposure and the post-award period. Let Y_{it} denote an outcome of supplier i in calendar year t, and let D_{it} indicate the post-award period relative to the grant decision year of the associated applicants. The baseline specification takes the form:

$$\begin{aligned} Y_{it} = {} & \beta_1 \left(W_i \times D_{it}\right) + \beta_2 \left(W_i \times R_j\right) + \\ & + \beta_3 \left(W_i \times R_j \times D_{it}\right) + \beta_4 \left(W2_i \times D_{it}\right) + \\ & + \rho Y_{i,t-1} + \lambda_t + \varepsilon_{it} \end{aligned} \tag{18}$$

The interaction $W_i \times D_{it}$ captures first-order upstream spillovers to suppliers arising from grant receipt by directly connected applicants. The terms $W_i \times R_j \times D_{it}$ flexibly control for smooth differences in supplier

outcomes associated with applicant quality and allow spillovers to vary with proximity to the cutoff. This is important because suppliers more exposed to applicants with stronger evaluation scores may differ systematically from other suppliers even in the absence of funding.

To examine whether spillovers propagate beyond directly connected suppliers, we construct a second-order exposure measure $W2_i$ defined as supplier i's pre-treatment exposure to other suppliers that are themselves exposed to applicants. The interaction $W2_i \times D_{it}$ therefore captures potential spillovers operating one additional step upstream ("suppliers of suppliers"). All exposure measures are constructed exclusively from pre-award transaction data. We include pre-period ($t = -1$) dependent variables to capture persistence in firm outcomes and year fixed effects λ_t to absorb aggregate shocks.

D.2 Network Exposure: Results

We next provide cutoff-based evidence on upstream spillovers by extending the discontinuity design to firms connected to applicants through pre-existing buyer–supplier relationships. The key identifying variation is that grant receipt changes discontinuously at the evaluation threshold for applicants near the eligibility frontier. We use this quasi-experimental variation to test whether applicant-level shocks transmit to suppliers. Unlike the matched-sample supplier analysis, this exercise conditions directly on the cutoff and therefore does not rely on balancing treated and control applicants through matching.

We estimate supplier-level panel regressions in which outcomes are in logs and the treatment intensity is a supplier's pre-award exposure to cutoff-sample applicants. Let W denote the share of a supplier's pre-award sales accounted for by applicants in the cutoff-based sample. We interact W with an indicator for the post-award period, D, so that $W \times D$ captures first-order spillovers to suppliers from applicants whose funding status changes at the cutoff. To flexibly control for smooth differences related to applicant quality and proximity to the threshold, we include $W \times R$ and $W \times R \times D$, where R is the normalized running variable measuring distance to the cutoff. We also include a second-order exposure measure, $W2$, defined as exposure to other suppliers that are themselves exposed to applicants, and estimate $W2 \times D$ to test for propagation beyond the first tier. The specification includes a lagged dependent variable to absorb persistence and year fixed effects to control for aggregate shocks. Identification is local to the cutoff and hinges on quasi-random assignment at the frontier, but power to detect higher-order and interaction effects is limited because treatment variation is generated only by applicants close to the threshold.

Table A20 reports the estimates. The coefficient on $W \times D$ is positive and statistically significant across

all outcomes, indicating economically meaningful first-order upstream spillovers. For revenue, the estimate is 0.331 log points (s.e. 0.094); for material costs, 0.349 (0.088); for tangible assets, 0.371 (0.152); and for employment, 0.116 (0.031). Because W is a share, these coefficients scale linearly with exposure. A convenient interpretation is per 10 percentage points of pre-award exposure: a supplier whose exposure is higher by 0.10 experiences post-award increases of about 0.033 log points in revenue, 0.035 in material costs, 0.037 in tangible assets, and 0.012 in employment. Interpreted as approximate percentage changes, these correspond to roughly 3 to 4 percent higher revenue, material costs, and assets, and about 1 percent higher employment per 10 percentage points of exposure. The joint rise in supplier revenue and material costs is consistent with a real activity channel in which funded applicants expand production and intermediate input demand, increasing sales and input use among directly connected suppliers.

The terms involving the running variable provide two additional diagnostics. First, $W \times R$ is negative and statistically significant for revenue and material costs (both -0.147 with $p < 0.05$) and for employment (-0.050 with $p < 0.05$), indicating that supplier outcomes covary smoothly with applicant distance to the cutoff in a way that is correlated with exposure. Second, W×R×D is small and statistically indistinguishable from zero across outcomes (for example 0.061 for revenue and 0.037 for material costs), suggesting limited evidence that the spillover slope varies systematically with proximity to the cutoff once the baseline smooth relationship is controlled for. These interaction terms are also estimated less precisely, which is consistent with limited power for detecting heterogeneity in a cutoff-based design.

Finally, there is little evidence of systematic propagation beyond the first tier. The coefficient on $W2 \times D$ is small and statistically indistinguishable from zero in all columns (0.047 for revenue, -0.009 for material costs, 0.088 for tangible assets, -0.019 for employment). Taken together, the cutoff-based network exposure results support the presence of first-order upstream spillovers to directly connected suppliers, while suggesting that effects attenuate rapidly with network distance within the scope of what can be detected in a design anchored to applicants near the eligibility frontier.

Table A20: Supplier-Level Effects

	Revenue	Material Costs	Tangible Assets	Employees
WD	0.331***	0.349***	0.371**	0.116***
	(0.094)	(0.088)	(0.152)	(0.031)
WR	−0.147**	−0.147*	−0.148	−0.050**
	(0.070)	(0.065)	(0.113)	(0.023)
WRD	0.061	0.037	0.072	0.019
	(0.080)	(0.075)	(0.130)	(0.027)
W2D	0.047	−0.009	0.088	−0.019
	(0.043)	(0.040)	(0.069)	(0.014)
Revenue (t-1)	0.796***			
	(0.004)			
Material Costs (t-1)		0.843***		
		(0.003)		
Tangible Assets (t-1)			0.743***	
			(0.003)	
Employees (t-1)				0.877***
				(0.002)
Num.Obs.	32 205	32 205	32 205	32 205
R2	0.613	0.685	0.588	0.861
R2 Adj.	0.613	0.684	0.588	0.861
R2 Within	0.612	0.684	0.587	0.861
R2 Within Adj.	0.612	0.684	0.587	0.861
AIC	119 972.7	115 879.1	150 970.9	48 734.6
BIC	120 048.1	115 954.5	151 046.3	48 810.0
RMSE	1.56	1.46	2.52	0.52
Year FE	Yes	Yes	Yes	Yes
Mean dep. var.	1.532	1.062	1.124	15.409
Median dep. var.	0.124	0.057	0.014	3.000

Note: W denotes a supplier's pre-treatment exposure to applicant firms, measured as the share of the supplier's pre-award sales accounted for by applicants included in the cutoff-based sample. D is an indicator for the post-award period for applicants, equal to one in years after the grant decision and zero otherwise. R is the normalized running variable capturing an applicant's distance to the grant-award threshold, constructed as the most binding normalized margin across the program's evaluation criteria; by construction, $R \geq 0$ indicates that an applicant satisfies all eligibility thresholds, $R < 0$ indicates failure to meet at least one threshold, and $R = 0$ corresponds to the eligibility frontier. $W \times D$ captures first-order upstream spillovers to suppliers arising from grant receipt by directly connected applicants. $W \times R$ and $W \times R \times D$ flexibly control for smooth differences in supplier outcomes associated with applicant quality and allow spillovers to vary with proximity to the cutoff. $W2$ denotes second-order exposure, defined as a supplier's pre-treatment exposure to other suppliers that themselves are exposed to applicants, and $W2 \times D$ captures potential spillovers to suppliers of suppliers. All exposure measures are constructed using pre-award transaction data. All monetary mean and median dependent variable values given are in mil. EUR. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.