Certification as a Signal, Not a Guarantee: How Innovation Status Recapitalises Italian SMEs through Equity

Daniel Ruiz Romera* Filippo di Pietro† Belén Martín Castro‡ Vincenzo Gentile§

August 17, 2026

Abstract

Does public innovation certification expand firms' debt capacity, or change how they finance themselves? We study Italy's *PMI innovativa* programme, a policy mix that bundles certification with a broad set of fiscal, financial, and governance benefits, three of which bear on capital structure: legally binding public certification, access to a sovereign-backed credit guarantee, and a 30% tax relief for equity investors. Because the status confers these jointly, we estimate the effect of certification itself and ask, from the margins it moves, which instrument the response points to, using a near-universal panel of 320,599 SMEs (2015–2024), treatment assigned by tax identifier, matched controls, and a staggered difference-in-differences estimator. The robust effect is on the equity side: certification raises paid-in share premium (1.48 pp, p =< 0.001), the precise instrument the decree subsidises. This is capital paid in above par by outside investors rather than founders' nominal subscriptions, and total equity rises by a comparable amount (1.66 pp). The response accumulates through the fourth post-certification year (1.75 pp) rather than spiking at registration, is near-identical across five matching and weighting designs, and survives relative-magnitude parallel-trends bounds. Exploiting the register's record of *why* each firm qualifies, the equity response concentrates where innovation is most *fundable*, among firms holding registered intellectual property (patents or registered software) rather than the most opaque, R&D-qualifying firms, a direct

*Universidad de Sevilla. Corresponding author: druiz8@us.es.
†JRC European Commission.
‡Universidad de Sevilla.
§Università del Salento.

test of the information mechanism. On the debt side the robust finding is what does *not* happen: certification does not raise leverage, against the premise that a guarantee should lever firms up. Point estimates under the primary design in fact turn negative (concentrated in short-term debt), but this contraction is design-sensitive and, unlike the equity result, does not survive the most demanding trend bounds, so we rest no causal claim on the deleveraging itself. Because the outcomes are ratios and the asset base expands, the robust reading is that the financing *mix* does not tilt toward debt, not that firms borrow less in absolute terms. The recapitalisation carries a cost: return on assets falls (-1.14 pp, $p = 0.003$), reflecting both the larger asset base the new equity funds and a compression in operating margins as firms front-load investment ahead of longer-horizon returns. Within this policy mix, the financing response runs through the equity incentive, not the credit guarantee: an evaluation scored on loan volumes would miss it.

Keywords: innovation certification, recapitalisation, equity financing, capital structure, innovation policy, policy mix, staggered difference-in-differences, SME finance, Italy, *PMI innovativa*.

JEL Classification: O31, G32, L25, C21.

1 Introduction

Innovative small and medium enterprises are the firms public policy most wants to finance and the ones private lenders find hardest to fund. Their value lies in intangible assets and, by the same token, in a scarcity of the tangible assets a lender can pledge—so the quality of their projects is opaque to outside investors and their collateral base is thin (Beck et al., 2008; Hall and Lerner, 2009). This tangible-asset scarcity is what a collateral-relaxing guarantee is meant to address. Governments increasingly respond with public certification schemes that vouch for a firm's innovative status and often attach subsidised credit guarantees. Whether these schemes actually loosen the financing constraint—and, if so, on which margin—is a first-order question for evaluating them. The implicit premise of a guarantee-based scheme is that it should lever firms up: relax the collateral constraint and more debt will flow. But a certification that above all conveys information can do something different—move firms toward the equity that suits their intangible, long-horizon assets, and away from the debt, and especially the short-term debt, on which opaque innovators usually rely (Meuleman and De Maeseneire, 2012; Bellucci et al., 2023b). An evaluation that tracks only the volume of credit will miss exactly this kind of effect.

Italy's *PMI innovativa* certification, introduced by Decree-Law 3/2015, offers a clean setting in which to separate volume from structure. The status is a broad *policy mix*; the three instruments most relevant to capital structure each target a distinct friction: *information validation*, through legally binding registration on a public register maintained by the Ministry of Enterprise; *collateral substitution*, through preferential access to the *Fondo di Garanzia*, a sovereign-backed guarantee fund; and a direct *equity* incentive, a 30% tax relief (since 2017) for investors who subscribe capital in certified firms (Section 2). The programme now certifies thousands of firms, yet its consequences for capital structure are essentially unstudied; Aiello et al. (2024) examine certified firms' profitability but leave the financing channel open. We ask whether *PMI innovativa* expands firms' debt capacity, as a guarantee scheme is meant to, or instead changes the composition of their financing—raising the subscribed equity the law subsidises—and at what cost to current performance.

Using a near-universal panel of 320,599 Italian SMEs over 2015–2024, treatment assigned by tax identifier, matched controls, and the staggered difference-in-differences estimator of Callaway and Sant'Anna (2021), we reach three conclusions. First, the certification works as a *signal*, not a credit guarantee: its robust effect is on the equity side, where paid-in share premium rises (1.48 pp, $p =< 0.001$)—the precise margin the 30% investor tax relief targets—with total shareholders' funds rising by a comparable amount (1.66 pp). This is fresh outside capital subscribed *above par*, not retained earnings or founders' par-value top-ups, and it accumulates through the fourth post-certification year (1.75 pp) rather than coinciding with registration. Second, the guarantee does not fire: certification does not raise leverage—if anything the ratio falls under the primary design (ATT = -2.23 pp total, -3.24 pp short-term), but this contraction does not survive the most demanding identification checks, so we treat the recapitalisation, not the deleveraging, as the causal core. Since the asset base also expands, this says the financing *mix* does not tilt toward debt, not that firms borrow less outright. Third, the response concentrates where innovation is most *fundable*—among firms holding registered intellectual property, not the most opaque, R&D-qualifying firms. The recomposition is not free: return on assets falls (-1.14 pp, $p = 0.003$), reflecting both the asset-base expansion the new equity funds and a genuine compression in operating margins as firms front-load investment ahead of returns that accrue over a longer horizon.

Because the effects are of modest size and certification is voluntary, we place unusual weight on identification and are explicit about its limits. A regression-discontinuity design at the eligibility thresholds is not available: meeting a criterion—say, the 3% R&D threshold—makes a firm *eligible* but not certified, since it must still choose to apply and file a binding self-certification, so crossing a threshold produces no jump in the probability of treatment to exploit. We therefore rely on matching and staggered difference-in-differences (Section 3). The equity response is the sturdier of the two margins: it is near-identical across all five matching and weighting designs, remains bounded away from zero under relative-magnitude parallel-trends violations, and lies well outside a randomization-inference timing placebo (RI $p = 0.010$). Its one natural threat is timing—a firm might register precisely when raising capital, to capture the investor tax relief. We address this head-on. Most directly, the effect persists at +1.13 pp ($p < 0.001$) among the 1,864 certified firms that did *not* raise paid-in capital in the two years before registration, so it cannot be a mechanical product of firms certifying mid-raise (it is larger, +3.39 pp, among those that did). The share-premium path also keeps climbing through the fourth post-treatment year rather than jumping once at registration, a dynamic a contemporaneous, one-off capital increase would not produce, and the rise runs through subscribed capital while retained earnings fall, ruling out a mechanical bookkeeping effect. The deleveraging is weaker. It is stable across the matching designs that retain explicit control firms (coarsened exact matching, our primary design, and Mahalanobis matching), the anticipation window, and a leave-one-cohort-out exercise; but it is not separable from a randomization-inference timing placebo (RI $p = 0.239$ for short-term leverage), it attenuates under reweighting designs, and under the most demanding trend bounds its confidence intervals admit zero. We report this asymmetry openly and rest the paper's conclusions on the margin that survives it.

Two pieces of mechanism evidence, associational by construction, indicate how certification reshapes financing. First, after certification the sensitivity of long-term leverage to a firm's own tangibility falls, consistent with the public guarantee standing in for the collateral that intangible-intensive firms lack. Second, because the register records which eligibility criterion each certified firm satisfies, we test how the response varies by qualification channel. The equity response is concentrated not among the most opaque firms—those qualifying through R&D spending—but among those holding registered intellectual property, the asset an outside investor can most readily value and underwrite. Certification activates equity where the signal is most *fundable*, a registry-based test of the information mechanism more direct than the sectoral proxies used in prior work (Section 4.5). This concentration is itself evidence against a purely mechanical reading of the equity result: were the rise simply certified firms booking subsidised capital, it would not single out the firms whose innovation an outside investor can most readily underwrite.

This paper makes three contributions. *First*, it shows how an innovation-certification policy mix reshapes SME capital structure: by estimating leverage and equity jointly, we find the programme mobilises external equity on the precise margin it targets (paid-in share premium) without raising leverage, which points to the equity incentive rather than the credit guarantee as the instrument in the mix that moves financing. *Second*, it exploits the register's disclosure of each firm's eligibility criterion as a direct test of the information mechanism, indicating the equity response concentrates where innovation is most fundable—firms holding registered intellectual property—a sharper test than the sectoral proxies of prior work and one whose logic travels beyond this scheme. *Third*, it provides the first causal evidence, to our knowledge, on the capital-structure effects of *PMI innovativa* (Bellucci et al., 2023b; Biancalani et al., 2022) and draws a lesson for programme evaluation: an assessment scored on the volume or maturity of *debt* would misjudge the programme, whereas the financing response it robustly produces runs through capital, at a measurable cost to current profitability.

The remainder of the paper proceeds as follows. Section 2 describes the *PMI innovativa* programme, reviews the related literature, and develops the hypotheses. Section 3 presents the data and the empirical strategy. Section 4 presents the causal results and mechanism evidence. Section 5 reports the multi-design robustness analysis. Section 6 concludes.

2 Institutional Background, Literature, and Hypotheses

2.1 The *PMI innovativa* Scheme

Decree-Law 3/2015 (Governo Italiano, 2015) created the "innovative SME" (*PMI innovativa*) status to stimulate innovation among Italian SMEs. Eligible firms must meet the EU SME definition (European Commission, 2003), have an Italian production centre, be unlisted, and not already be registered as an innovative startup; they must satisfy at least two of three criteria—R&D of at least 3% of the larger of production costs or revenues, a qualified workforce (doctoral or master's-level), or registered intellectual property. Certified firms are listed in a dedicated section of the *Registro delle Imprese*, making their status publicly visible, and must renew an annual binding self-certification (*autocertificazione*) checked administratively against filed accounts. The status bundles a broad set of regulatory and financial advantages (Ministero dello Sviluppo Economico, 2015, 2022), which fall into four groups. On the fiscal side, qualifying equity investors—individuals and corporations, including venture-capital funds—receive, since 2017, a 30% tax relief on the capital they subscribe (an IRPEF deduction for individuals, up to €1 million, and an IRES deduction for corporations, up to €1.8 million), conditional on a three-year holding period.[1] Employee stock-option grants are tax-exempt, and various administrative levies (such as stamp duties and Chamber of Commerce fees) are waived; certified firms also enjoy relaxed crisis-management rules, including deferral of the mandatory recapitalisation that ordinarily follows a capital loss exceeding one-third of equity, postponed to the second subsequent fiscal year. On the corporate-governance side, the regime permits tailored shareholder rights that facilitate complex venture-capital contracting and

[1]Since 2020 a parallel *de minimis* scheme has offered individuals a larger 50% deduction (up to €300,000 per year); we refer throughout to the general 30% incentive, available since the 2017 Budget Law and built on the framework of Decree-Law 3/2015.

the issuance of equity instruments without diluting founders' control. On the financing side, certified firms obtain simplified, fee-free access to the *Fondo di Garanzia per le PMI* and may raise capital through authorised equity-crowdfunding platforms. Finally, on the growth and internationalisation side, they receive subsidised support from the Italian Trade Agency and access to international innovation events. The 30% equity tax relief and the guarantee-fund access are the two instruments most directly relevant to capital structure, and they map onto the financing responses documented below—a shift toward equity and a contraction in debt.

These two instruments work through distinct frictions. The first, an information-validation channel, operates through public registration: because the *autocertificazione* is legally binding and false declarations expose the legal representative to civil and criminal liability, registration converts a firm's innovation profile from a private claim into a credible, third-party-verifiable signal available to lenders and investors at no search cost. The second, a collateral-substitution channel, operates through simplified, fee-free access to the *Fondo di Garanzia per le PMI*, which provides an 80% sovereign guarantee on bank loans up to €5 million without requiring the firm to post collateral for the covered tranche. Unlike innovative *startups*, for which the guarantee is granted automatically, certified PMI obtain only *streamlined* access: the Fund still assesses each loan's credit risk on a five-band rating and denies the guarantee to applicants in the lowest band (Ministero dello Sviluppo Economico, 2022). Because the Fund substitutes for pledgeable assets, its value is largest for firms with the tightest collateral constraints—though this creditworthiness gating already hints at why the credit channel may not bind. These two channels motivate the dual-mechanism predictions developed in Section 2.3.

2.2 Related Literature

The literature speaks to four themes. Public interventions in SME finance address two market failures: information asymmetry in credit markets and underinvestment in high-social-return innovation (Beck et al., 2008; Hall and Lerner, 2009). Asymmetric information can ration credit even when borrowers are willing to pay higher rates (Stiglitz and Weiss, 1981; Bester, 1985), and in bank-based systems these frictions bind hardest for small firms lacking credit histories (Magri, 2009). Policies therefore pair financial instruments (subsidies, guarantees, tax relief) with informational ones (certification, registry inclusion); evidence on government-backed guarantees in Italy in particular shows measurable effects on SME outcomes (Gai et al., 2023). The additionality of such guarantees is, however, contested: a regression-discontinuity evaluation of Italy's own *Fondo di Garanzia* finds that eligibility raises bank lending but not its price and can even raise default risk (de Blasio et al., 2018), and where a scheme bundles a guarantee with a quality certification the certification component, rather than the guarantee, carries the effect outside crises (Bonfim et al., 2023), while mutual-guarantee arrangements cut borrowing costs for opaque Italian firms (Columba et al., 2009). Because take-up of these schemes is voluntary, separating selection from treatment is the central identification problem, which we confront by design (Section 3). The evidence is mixed on whether such support expands debt *volume* (Minard, 2016; Martí and Quas, 2018; Ullah, 2020) or instead improves maturity and pricing without raising aggregate leverage (Meuleman and De Maeseneire, 2012; Bellucci et al., 2023b). This distinction is central to how a programme should be evaluated. If the policy-relevant outcome is the quantity of credit, a programme that leaves total debt unchanged—or reduces it—looks like a failure; but if firms are moving toward the equity that their intangible, long-horizon assets call for, the same debt figures conceal a welfare-improving recapitalisation, a better match between asset and liability horizons. Most existing evaluations report a single headline coefficient on total or bank debt and therefore cannot separate these cases. Our design is built around the distinction: we estimate debt—total, short-term, and long-term—and equity jointly, and treat the debt–equity mix, rather than the volume of credit, as the object of interest.

Capital-structure theory offers two organising frameworks. The trade-off view balances debt's tax advantage against expected bankruptcy and agency costs, implying an interior optimum that rises with tangibility and profitability (Kraus and Litzenberger, 1973; DeAngelo and Masulis, 1980); the pecking-order view, grounded in asymmetric information, predicts instead a financing hierarchy—internal funds, then debt, then external equity—so that opaque firms lean on whatever external finance is least information-sensitive (Myers and Majluf, 1984; Leland and Pyle, 1977; Ross, 1977). The two are not mutually exclusive, and the empirical determinants of leverage—size, age, tangibility, profitability, growth—are robust across countries (Antoniou et al., 2008; De Jong et al., 2008; Frank and Goyal, 2009). For SMEs the evidence adds heavier reliance on short-term, bank-intermediated debt, a strong negative profitability–leverage relation, and pronounced collateral sensitivity (Michaelas et al., 1999; Hall et al., 2004; Chittenden et al., 1996; Degryse et al., 2012; Palacín-Sánchez et al., 2013; Daskalakis et al., 2017). These regularities frame our control set.

Against this backdrop, the capital structure of innovative firms differs systematically because of a *scarcity of tangible, pledgeable assets*: their value sits in intangibles, and their returns are harder to verify. It is this shortage of collateral—not intangible intensity *per se*—that defines the binding constraint a guarantee scheme targets and bounds what such a guarantee can do. Pecking-order logic predicts reliance on internal funds under asymmetric information (Myers and Majluf, 1984; Frank and Goyal, 2009), while collateral-based lending makes tangibility a robust predictor of debt capacity, especially at long maturities (Rajan and Zingales, 1995; Titman and Wessels, 1988; Rampini and Viswanathan, 2013). Since the irrelevance benchmark of Modigliani and Miller (1958), the literature has located the determinants of leverage precisely in such frictions, and for high-technology SMEs these produce a distinctive, intangible-heavy capital structure (Serrasqueiro et al., 2016; Ali et al., 2024). The maturity margin is especially salient for small firms, whose debt-maturity structure responds to asset composition, information quality, and monitoring (Scherr and Hulburt, 2001; Feito-Ruiz et al., 2023). For R&D-intensive firms, long-horizon and weakly collateralizable projects fit poorly with short-term facilities and face rationing in long-term credit (Hall and Lerner, 2009), leaving such firms dependent on short-tenor instruments even when aggregate ratios look similar (Brown et al., 2009; D'Amato, 2020). This literature motivates a focus on debt composition and maturity, not only total leverage—the analytical choice at the heart of this paper.

Turning to certification, such schemes institutionalize signalling: a signal is informative when mimicry by low-quality firms is costly (Spence, 1973), a cost created here by eligibility thresholds and legal accountability, and innovation in particular is hard to communicate credibly to outside financiers (Bhattacharya and Ritter, 1983). Empirically, certification improves credit-access probabilities (Minard, 2016; Ullah, 2020) and can lower the cost of bank debt by substituting public for private information (Bellucci et al., 2023a); when bundled with financial instruments, long-term debt often responds more than short-term debt (Meuleman and De Maeseneire, 2012; Bellucci et al., 2023b); and policy-linked equity incentives can activate targeted equity components rather than broad equity growth. Evidence on investor tax incentives elsewhere is sobering: U.S. state angel tax credits raise angel investment but often draw inexperienced or insider investors toward lower-growth firms, with limited effects on real outcomes (Denes et al., 2023)—a caveat we weigh in reading the equity response and its cost. Related work on innovative firms' access to public funds and venture capital documents how status and public support reshape financing (Colombo et al., 2007; Lee et al., 2015). For Italy, evidence on *PMI innovativa* has addressed performance and regional outcomes (Aiello et al., 2024), leaving the capital-structure channel—and the separation of information validation from collateral substitution—unexplored.

On identification, difference-in-differences inference must contend both with serial correlation (Bertrand et al., 2004) and, under staggered timing and heterogeneous effects, with the bias of the conventional two-way fixed-effects estimator (Goodman-Bacon, 2021; Sun and Abraham, 2021; Wooldridge, 2021). The Callaway and Sant'Anna (2021) estimator addresses the latter by comparing cohort-time effects against explicit control groups under transparent weighting, and has become the preferred design for level effects under progressive adoption. Matching complements DiD by improving covariate overlap before estimation (Rosenbaum and Rubin, 1983; Caliendo and Kopeinig, 2008); because propensity-score matching collapses the covariate vector into a single index and can achieve weaker balance than distance- or moment-based methods (King and Nielsen, 2019; Hainmueller, 2012), we verify our results across several matching and weighting designs and report, for the debt margins, how sensitive each conclusion is to violations of parallel trends (Rambachan and Roth, 2023). This triangulation reflects a broader lesson from the recent panel-data literature: because no single estimator dominates across data-generating processes, robustness is best assessed across estimators and identifying assumptions rather than asserted for one (Athey et al., 2026).

2.3 Hypotheses

Our predictions concern the *level* effects on the two sides of the balance sheet and the *mechanism* behind them. We order them by the strength of their *a priori* institutional rationale—which, as it turns out, the evidence mirrors. The central prediction is on the equity side and follows directly from the design of the 2017 investor tax relief, which subsidises newly subscribed equity. The debt-side prediction is the question a guarantee scheme invites—whether it levers firms up—which our evidence supports only as a null. The mechanism is read most directly off the eligibility criterion each firm satisfies. Associational mechanism evidence (how borrowing responds to a firm's own collateral) and additional reduced-form patterns are reported as supplementary evidence in Section 4 and the Appendix.

H 1 (Targeted equity recapitalisation)**.** *The 30% tax relief activates equity through* paid-in capital*—share premium subscribed above par—rather than through retained earnings or broad-based equity growth. The causal effect on share premium is positive, with aggregate shareholders' funds rising by a comparable amount, signalling fresh outside capital priced above nominal value rather than internally generated funds or founders' par-value top-ups.*

H 2 (No credit lever)**.** *Certification does not expand credit. The guarantee-scheme premise makes the sharp, falsifiable prediction $ATT(\textit{Total leverage}) > 0$; if certification instead works through signalling and a subsidised equity channel rather than by relaxing a binding credit-supply constraint, that alternative is rejected and total leverage does not rise. Any movement is, if anything, a contraction concentrated in short-term liabilities rather than an expansion—a pattern we treat as exploratory, since the data cannot separate financial from trade payables and the contraction is design-sensitive (Section 5).*

H 3 (Qualification channel: fundability)**.** *The register records which eligibility criterion each firm satisfies—R&D spending of at least 3% of cost or production value, a qualified workforce, or registered intellectual property (a patent or registered software)—which lets us test the information mechanism directly rather than through a sectoral proxy. Under an* opacity *view* *the response concentrates among the least verifiable input, R&D ($|ATT_{R\&D}| > |ATT_{IP}|$); under a* fundability *view it concentrates instead among holders of registered intellectual property—the asset an equity investor can most readily value and underwrite, so that a verifiable signal complements the equity subsidy ($|ATT_{IP}| > |ATT_{R\&D}|$). The data adjudicate. Because a firm may satisfy several criteria, the test compares the effect conditional on* whether *a firm qualifies through a given channel, not a partition of the treated.*

3 Data and Empirical Strategy

This section describes the data and the empirical strategy. The strategy separates three roles: pooled OLS and two-way fixed-effects (FE) regressions describe baseline leverage correlates in the full panel (an associational benchmark, Appendix B); the Callaway and Sant'Anna (2021) staggered DiD estimator identifies causal effects on leverage *levels* (ATTs) in a matched sample; and two-way fixed-effects interaction models assess post-certification changes in *sensitivities* (mechanism-consistent, associational slope evidence). Causal claims are reserved for the DiD levels; mechanism evidence is explicitly associational.

3.1 Data and sample

We analyse Italian SMEs over 2015–2024 at the firm-year level. Financial data come from the *Analisi Informatizzata delle Aziende Italiane* (AIDA) database (Bureau van Dijk / Moody's), which provides harmonised annual accounts for virtually all Italian incorporated firms; the list of certified *PMI innovative* firms is from the Italian Business Register (Unioncamere – Infocamere, 2026). Certified firms are statutory SMEs under the eligibility regime of Decree-Law 3/2015, and the control universe is drawn from the AIDA extraction under the EU SME size criteria (European Commission, 2003); we impose no further accounting size filter on certified firms, so as to retain the full certified population, and restore treated–control size comparability at the matching stage rather than by sample truncation. We exclude firms in liquidation or with missing incorporation year, financial intermediaries (NACE Rev.2 64–66), and real estate (NACE Rev.2 68), whose mortgage-financed balance sheets reflect a structurally distinct collateral regime. The resulting analytical panel comprises **3,201,469** firm-year observations on **320,599** firms, of which **3,065** ever obtained certification. Table 1 summarises the construction; the treatment-cohort distribution underlying the staggered design is reported in Table 11 (Appendix A).

Table 1: Sample Construction Flow

Step	Stage	N Obs.	N Firms	N Treated	N Controls
1	AIDA raw (Excel exports)	322,173	322,169	—	—
2	Clean panel (excl. liquidation, missing incorp.)	3,201,851	320,651	—	—
3	Linked to PMI registry	3,201,851	320,651	3,117	317,534
4	Excl. financial (ATECO 64-66) and real estate (68)	3,201,469	320,599	3,065	317,534
4b	R&D reporters (Spec 2, 58.3% of row 4)	1,866,841	277,736	2,058	275,678
5	Matching pool at t-1 (unique firms)	—	320,403	2,869	317,534
6	Post-matching: coarsened exact matching (CEM)	—	223,300	2,308	220,992
7	GOLD panel for DiD estimation	2,229,994	223,300	2,308	220,992

[a] The raw AIDA export provides one record per firm; 4 record(s) removed due to missing or duplicate tax identifiers. Step 1b expands each firm across available fiscal years (2015–2024); actual observation counts vary with entry and exit dates.

[b] 196 firms in Step 4 lack a valid $t-1$ observation and are excluded from the matching pool.

Other notes: Steps 1–4 describe the full analytical panel used for OLS and FE estimation. Steps 5–7 describe the matched subsample for causal identification via staggered DiD; the doubly-robust DiD estimates (Table 4) are computed on the firm-years in this panel with complete covariate information. R&D expenditure (58.3% coverage) and cash holdings (42% coverage) enter only specifications where available. Reproducible via `scripts/99_Table_SampleFlow.R`. *Source:* Authors' elaboration based on AIDA (Bureau van Dijk (Moody's Analytics), 2024).

Our dependent variables are three book-value leverage measures. Our AIDA extraction records liabilities only as *total payables*, with a short-term/long-term maturity split but *no* breakdown into financial (bank) and trade items—a feature of the harmonised *bilancio abbreviato* schedule under which most non-listed Italian SMEs file, and one that the extraction does not resolve for *any* subset of firms, including those filing the full *bilancio ordinario*. We therefore follow the standard Italian convention (D'Amato, 2020) and measure leverage from total payables: total (D/TA), short-term (STD/TA), and long-term (LTD/TA) liabilities over total assets, at book value. A direct implication is that STD/TA includes trade payables and cannot be decomposed into bank credit and supplier financing for any subsample; we therefore read movements in STD/TA as changes in short-term *liabilities* rather than bank credit specifically, and return to this in Section 6. Where a financing-cost margin is needed, we instead use total financial charges (interest expense), which the extraction *does* report

(Section 5).

Following standard capital-structure practice we control for size (ln total assets), age (ln years since incorporation), asset growth (Δ ln TA), tangibility (tangible fixed assets/TA), profitability (EBITDA/TA), and non-debt tax shield (depreciation/TA). Two knowledge-capital variables are added where available: R&D intensity (R&D/TA, reported by 58.3% of firm-years) and cash holdings (cash/TA, 42.0%), both limited by abbreviated-format reporting; intangibility (intangible assets/TA) has full coverage, with zeros retained as genuinely asset-light firms. All ratios are winsorised at the 1st/99th percentiles of the pooled panel (Dixon and Massey, 1960), so boundaries are set by the large control group, not the small treated group.

Two features of the sample shape the analysis. First, certification is rare and staggered: the 3,065 ever-certified firms enter the register across multiple annual cohorts over 2015–2024 (Table 11), which is what makes a timing-robust estimator necessary and rules out a single-break design.

Second, certified firms are not a random draw from the population. They are markedly asset-light: intangible assets average 19.66% of total assets against 3.71% among controls, collateralisable tangible assets are correspondingly thinner (8.30% versus 19.94%), and R&D intensity is more than an order of magnitude higher (1.24% versus 0.05%). They also carry less debt and more equity: total leverage of 52.67% versus 59.42%, short-term liabilities of 38.06% versus 45.02%, and an equity ratio of 38.81% versus 32.10%. Certification is also highly concentrated by activity: 35.3% of certified firms operate in computer programming and IT services (NACE 62) and 14.6% in scientific R&D (NACE 72), rising to 54.8% once electronics and instruments (NACE 26) are added, so a handful of knowledge-intensive divisions account for the bulk of the certified population.

These gaps are descriptive and measured before matching: they reflect *who* certifies rather than any effect *of* certifying, and they are precisely the opacity and collateral frictions the certification is meant to relax. They are also why the raw treated–control comparison is uninformative and a matched design is required. The profile survives dropping the software sector: excluding NACE 62, the remaining certified firms are still far more intangible-intensive (19.14% intangibility), less levered (54.22%), and more equity-financed (38.21%) than controls, so the asset-light pattern is general to *PMI innovative* rather than a feature of their digital core. The matching stage (Section 3) restores comparability on size, age, leverage, tangibility, profitability, growth, and employment before any causal contrast is drawn.

3.2 Matching

Certification is voluntary and staggered, so treated and never-treated firms differ sharply on observables. Certified PMI are much younger and more asset-light than the SME control universe (Section 3.1), which creates a genuine overlap problem. Within exact year × two-digit NACE Rev.2 cells there are too few young, intangible-intensive controls for nearest-neighbour matching to balance firm age without discarding most treated firms. We confront this directly by pre-processing on seven pre-treatment covariates measured at $t-1$—ln assets, ln age, total leverage, tangibility, EBITDA/TA, asset growth, and ln employees—under several designs and by selecting as primary the one that achieves balance without the pathologies the overlap problem induces. Including pre-treatment total leverage among the covariates strengthens parallel trends for the primary outcome by controlling for the documented persistence of capital structure (Lemmon et al., 2008); conditioning on a pre-treatment outcome level is standard in DiD and is not a bad-control problem (Angrist and Pischke, 2009). Matching improves overlap and balance *before* estimation; it is not itself the source of the causal estimand.

Our primary design is coarsened exact matching (CEM): it bins each covariate—using fine age bins to resolve the age gap—and matches treated to controls within those bins (Iacus et al., 2012); calendar time is absorbed by the staggered estimator, and the exact year × sector cells used by the distance- and propensity-based designs serve as a cross-check rather than a feature of the primary match. Because the share-premium effect is near-identical under Mahalanobis and propensity-score matching, both of which *do* impose exact year × two-digit-sector strata (Section 5), residual industry composition in the primary match does not drive the equity result. CEM brings every covariate below the conventional 0.10 standardised-mean-difference threshold (Rosenbaum and Rubin, 1983; Austin, 2011) while keeping explicit control firms (not weights) and avoiding the extreme-weight fragility that the poor overlap would otherwise produce. The resulting matched "GOLD" panel links 2,308 treated firms to 220,992 controls across 2,229,994 firm-years.[2]

Because no single scheme is uniformly best, we re-estimate the entire design under four further matchers and treat the choice as an object of robustness. Entropy balancing reweights all controls so that their weighted covariate moments equal the treated moments *exactly* (Hainmueller, 2012); it attains the tightest balance and discards no firms, but under limited

[2]Since ln employees has the lowest coverage of the seven covariates, we verified that requiring it does not distort the matched sample. Re-running the entire primary design without it retains 2,032 estimable treated firms (against 1,930; +102) at essentially unchanged balance, and the robust equity effect is unaffected (share-premium ATT +1.52 versus 1.48 pp). We retain ln employees because excluding it weakens the share-premium parallel-trends test (joint pre-trend p moving from 0.093 to 0.037) while adding little—it is itself well balanced when included. Appendix Table 18 reports the full comparison.

overlap it concentrates weight on a few controls, so its effective sample is far smaller than its nominal one. Mahalanobis distance matching selects nearest neighbours within exact year $\times$ sector cells, here with a calliper on firm age to contain the age gap. The covariate-balancing propensity score (Imai and Ratkovic, 2014) estimates the score so as to balance covariates directly, the principled propensity-family answer to the failure of plain propensity-score matching, which collapses the covariates into a single index and, under poor overlap, can leave them more imbalanced than before (King and Nielsen, 2019). Table 2 summarises how the designs differ; Section 5 shows that the substantive conclusions hold across all of them, which is what makes the choice of primary design immaterial to the findings.

Table 2: Matching and weighting designs: how they differ

Design	**How it forms comparisons**	**Behaviour under limited overlap**	**Role**
Coarsened exact matching (CEM)	Coarsens each covariate into bins (fine bins for age) and matches treated to controls within those bins.	Brings every covariate below the 0.10 threshold while keeping explicit control firms and near-uniform weights; no extreme-weight fragility.	**Primary**
Entropy balancing (EB)	Reweights all never-treated controls so their weighted covariate means equal the treated means exactly (Hainmueller, 2012).	Achieves exact moment balance and discards no firm, but concentrates weight on a few controls, so the effective sample is far below the nominal one.	Robustness (exact-balance cross-check)
Mahalanobis distance matching (MDM)	Selects nearest neighbours in multivariate covariate space within exact year $\times$ sector cells, with a calliper on firm age.	Balances most covariates with explicit matched controls, but a residual age gap and sample loss remain because young, asset-light controls are scarce.	Robustness
Covariate-balancing propensity score (CBPS)	Estimates the propensity score so as to balance the covariates directly, then reweights (Imai and Ratkovic, 2014).	The propensity-family method that does balance under poor overlap, unlike plain propensity-score matching.	Robustness (propensity family)
Propensity-score matching (PSM)	Collapses the covariates into a single estimated treatment probability and matches nearest neighbours on it.	Can leave individual covariates more imbalanced than before matching when overlap is poor (King and Nielsen, 2019).	Cautionary case (appendix)

Notes: All designs use the same seven pre-treatment covariates measured at $t-1$ (ln assets, ln age, total leverage, tangibility, EBITDA/TA, asset growth, ln employees) and, where applicable, exact year $\times$ two-digit NACE Rev.2 strata. The 0.10 figure is the conventional standardised-mean-difference threshold (Austin, 2011). Realised balance and effective sample sizes are reported in Section 5.

3.3 Staggered difference-in-differences

For cohort g (first-certification year) and calendar year t, the CS2021 estimator identifies

$$\mathrm{ATT}(g,t) = \mathbb{E}[Y_{it}(g) - Y_{it}(0) \mid G_i = g]\,, \tag{1}$$

comparing each treated cohort to never-treated controls under conditional parallel trends. By using only never-treated firms as the comparison group, the estimator avoids the already-treated-as-control comparisons that bias the static two-way fixed-effects estimator under staggered timing and heterogeneous effects (Goodman-Bacon, 2021; Sun and Abraham, 2021); the headline estimates are thus immune to that contamination by construction, and the only two-way fixed-effects quantity we report—the associational collateral slope of Section 4.5—is flagged accordingly. We use the doubly-robust implementation (combining outcome regression and propensity weighting on the seven covariates, evaluated at the pre-treatment base period), primary-design (CEM) weights, an anticipation window $\kappa = 2$ with reported sensitivity to $\kappa \in \{0, 1, 2\}$, and $B = 999$ bootstrap replications (`set.seed = 2026`). We report the aggregate ATT and the dynamic event-study, with a joint Wald pre-trend test on the design-consistent pre-period $e \leq -(\kappa + 1)$ using the bootstrap influence-function covariance. For the debt outcomes we report, in addition to the headline estimates, relative-magnitude sensitivity bounds (Rambachan and Roth, 2023) that ask how large a parallel-trends violation would have to be to overturn each conclusion; we carry these through transparently, including where they qualify the debt results (Section 5).

The identifying assumption is conditional parallel trends: absent certification, treated and matched control firms would have followed common leverage paths after conditioning on the matching covariates. Two threats deserve comment. The first is selection on time-varying unobservables—firms may certify precisely when they anticipate a financing shift. We address this on three fronts: matching on the pre-treatment level of the outcome and on the standard leverage determinants; testing for and failing to reject differential pre-trends on the genuine pre-period; and reporting formal sensitivity bounds that ask how large a parallel-trends violation would have to be to overturn each conclusion (Section 5). A regression-discontinuity design around the statutory eligibility thresholds—a natural alternative—is not available here, for three reasons. First and most fundamentally, eligibility is not treatment: certification is voluntary, so crossing a threshold confers only the *option* to register, not registration itself; treatment is self-selected rather than assigned at the cutoff, and there is no discontinuity in treatment status to exploit. Second, qualification requires at least two of three criteria (R&D spending of at least 3% of the larger of cost or production value, a qualified workforce, or registered intellectual property), so no single criterion is a binding threshold—a firm below the R&D cutoff can still qualify through the other two. Third, although the register discloses *which* criteria each firm satisfies—information we exploit in the mechanism analysis below—it does not report the underlying continuous magnitudes (for instance, the R&D-to-cost ratio relative to the 3% threshold), so no running variable with a known discontinuity is observed. A sharp RD design is therefore infeasible; a fuzzy design that instrumented registration with crossing an eligibility threshold would fare no better, because the endogenous margin here is the voluntary decision to register, not the meeting of a criterion, so the first stage—threshold crossing predicting registration—is precisely where selection enters. The second threat is anticipation: because investors and lenders may react before the registration date, we treat the two pre-registration years as a potentially contaminated window ($\kappa = 2$) rather than as clean controls, and we show that the conclusions are stable to this choice. Neither device can rule out every confounder, but together they make a purely selection-driven explanation difficult to sustain, especially given the heterogeneity of the response across qualification channels documented below, which a single aggregate financing shock would not reproduce. The mechanism analysis that probes *how* certification reshapes financing is described together with its results in Section 4.5.

4 Results

We organise the evidence in three layers: the structural baseline from the full panel, the causal level effects from the matched staggered DiD, and the mechanism-consistent slope evidence. The causal effects are the paper's core; the baseline establishes the financing frictions that the certification is hypothesised to relax, and the mechanism evidence speaks to how it does so.

4.1 The structural baseline

Before turning to certification, the full-panel fixed-effects estimates document the financing environment that innovative SMEs operate in, and they motivate the maturity focus of the causal analysis (full tables in Appendix B). Three regularities stand out (exact coefficients in Tables 12 and 13, Appendix B). First, profitability enters with a large, precisely estimated negative coefficient on total leverage, the canonical pecking-order pattern in which internally generated funds substitute for external debt (Frank and Goyal, 2009). Second, asset tangibility acts on the *maturity* of debt far more than on its quantity: it loads *negatively* on short-term leverage and *positively*, by a similar amount, on long-term leverage, so that its net effect on total leverage is small. The same maturity-matching signature holds for intangibility. Collateral, in other words, governs chiefly the maturity at which firms borrow—the empirical fact that makes debt composition, alongside its level, a natural margin on which to look for a certification effect. Third, R&D intensity is associated with more short-term but no additional long-term borrowing: the short-term loading is large and significant while the long-term coefficient is statistically indistinguishable from zero. Unable to credibly communicate the long-horizon value of intangible investment, R&D-intensive firms are confined to short-tenor instruments—precisely the baseline that a credible innovation signal, or a shift towards equity, should relax.

4.2 Capital structure: deleveraging and recapitalisation

Table 3 reports covariate balance across all five designs. The primary CEM brings every one of the seven covariates below the conventional $|\text{SMD}| < 0.10$ threshold (Austin, 2011)—the best balance among the designs that retain explicit control firms. Entropy balancing and the covariate-balancing propensity score attain exact moment balance, but only by reweighting the controls down to a small effective sample; Mahalanobis matching leaves a residual gap on size, profitability, and growth; and propensity-score matching leaves the size and age gaps largely intact—the overlap pathology that confines it to a cautionary role (King and Nielsen, 2019). Table 4 reports the CS2021 doubly-robust ATTs.

Table 3: Covariate balance across matching and weighting designs

	Standardised mean difference $(t-1)$					
Covariate	Unmatched	CEM	MDM	EB	PSM	CBPS
ln assets	+0.198	+0.096	+0.154	-0.000	-0.281	+0.000
ln age	-1.157	-0.044	-0.143	+0.000	-0.662	-0.000
Total leverage	-0.257	-0.030	-0.075	-0.000	+0.411	-0.000
Tangibility	-0.854	-0.081	+0.004	+0.000	+0.117	-0.000
EBITDA/TA	-0.112	-0.026	-0.102	-0.000	-0.067	-0.000
Asset growth	+0.451	+0.032	+0.121	-0.000	+0.193	+0.000
ln employees	+0.194	+0.021	+0.057	+0.000	-0.262	+0.000
Max |SMD|		0.096	0.154	0.000	0.662	0.000
Treated (n)		2,308	1,784	2,410	2,382	2,410
Controls		220,992	6,497	3,363 (ESS)	7,956	3,364 (ESS)

Notes: Standardised mean differences on the seven pre-treatment $(t-1)$ covariates, computed by `cobalt` with the treated standard deviation as denominator, before matching (*Unmatched*) and after each design with its weights. CEM (the primary design) and MDM retain explicit matched control firms; EB and CBPS reweight the full control pool (control count reported as the effective sample size, ESS). $|\text{SMD}| < 0.10$ indicates satisfactory balance (Austin, 2011). *Source:* Authors' elaboration based on Bureau van Dijk (Moody's Analytics) (2024).

Table 4: CS2021 Average Treatment Effects on the Treated — Debt Neutrality and Falsification Tests

Outcome	ATT	SE	p-value	Sig	N_{obs}	N_{firms}
Panel A: H_{null} (Debt Neutrality)						
Total leverage (D/TA)	-0.0223	0.0075	0.003	***	2,229,479	223,300
Short-term leverage (STD/TA)	-0.0324	0.0077	0.000	***	2,229,479	223,300
Long-term leverage (LTD/TA)	0.0091	0.0070	0.194		2,229,479	223,300
Panel B: $H_{debt\text{-}spec}$ (Equity Falsification)						
Share premium / TA	0.0148	0.0016	0.000	***	2,229,479	223,300
Shareholders funds / TA	0.0166	0.0080	0.037	**	2,229,479	223,300

Notes: Callaway and Sant'Anna (2021) doubly-robust estimator. Control group: never-treated firms from GOLD matched panel (2,308 treated firms; 220,992 matched controls; 2,229,479 firm-year observations with complete covariate information used by the doubly-robust estimator). Anticipation window: $\kappa = 2$. Bootstrap replications: $B = 999$ (`set.seed = 2026`). Primary-design (coarsened exact matching) weights applied. ATT: overall average treatment effect on the treated (aggregated across cohorts). Pre-trends assessed with a joint Wald test on the design-consistent pre-period $e \leq -(\kappa + 1)$, using the bootstrap influence-function covariance across pre-period event-time coefficients. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. *Source:* Authors' elaboration based on Bureau van Dijk (Moody's Analytics) (2024).

On the debt side, the robust finding is what does *not* happen: certification does not raise leverage, against the guarantee-scheme premise that it should lever firms up. This null—the falsification of the credit-lever channel—is the debt result we stand behind. Under the primary design the point estimates in fact turn negative: total leverage -2.23 pp ($p = 0.003$), concentrated in short-term liabilities (-3.24 pp, $p =< 0.001$), with long-term leverage flat (0.91 pp, $p = 0.194$); the three satisfy the adding-up identity up to estimation noise ($-3.24 + 0.91 \approx -2.23$ pp), and the joint pre-trend test passes ($p = 0.980$ total, $p = 0.424$ short-term). We treat this *contraction* —its sign and magnitude—as exploratory rather than confirmatory: Section 5 shows it is design-sensitive and does not survive the most demanding parallel-trends bounds, so we rest no causal claim on the deleveraging itself.

Two clarifications guard against over-reading the debt result. First, all our outcomes are ratios to total assets, and the asset base itself expands sharply after certification (Section 4.4): a flat or falling debt-to-assets ratio alongside a much larger balance sheet implies that debt in *levels* broadly keeps pace with assets rather than contracting. The robust statement is therefore about the financing *mix*—certification does not tilt it toward debt, and if anything tilts it toward equity—not that firms borrow less in absolute terms. Second, the short-term measure combines revolving credit with trade payables and cannot be attributed to bank borrowing alone. The resulting pattern—an equity-tilted balance sheet rather than a debt-led expansion—is what a signalling, not a credit-lever, interpretation predicts. Its counterpart on the equity side is developed in Section 4.3; because the two margins are estimated separately, we describe their co-movement as a recapitalisation around certification rather than a firm-by-firm one-for-one swap.

Figure 1 traces the full dynamics and shows that the debt response is concentrated in the first post-certification years rather than permanent. Before certification, all three series are flat and statistically indistinguishable from zero (to the left of the dashed line at $e = -0.5$), so matched treated and control firms were on common leverage paths—the parallel-trends premise of the design. After certification, the short-term panel drops sharply: the contraction is already significant at $e = 0$ and deepens to its trough at $e = 1$ (-4.50 pp), then attenuates over the following years, standing at -1.76 pp by $e = 4$, where it is no longer statistically distinguishable from zero. The total-leverage panel mirrors this profile at a smaller scale. Long-term leverage moves little: a mild, short-lived uptick around $e = 1$ (2.42 pp) that fades by $e = 3$. The debt response is thus concentrated in the first two to three post-certification years rather than permanent—a transitory re-profiling of the liability side, and, as Section 5 shows, the margin most sensitive to the identifying assumptions. This temporary debt path stands in contrast to the equity response documented next, which builds and persists.

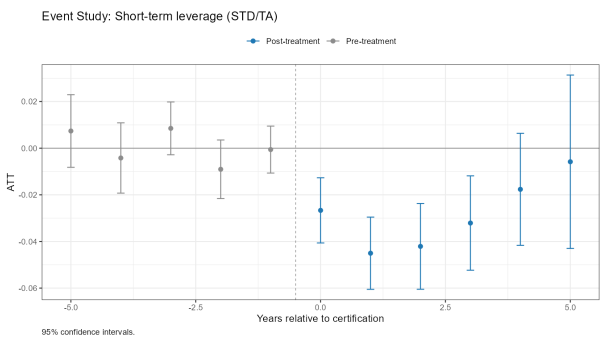

(a) Short-term leverage (STD/TA)

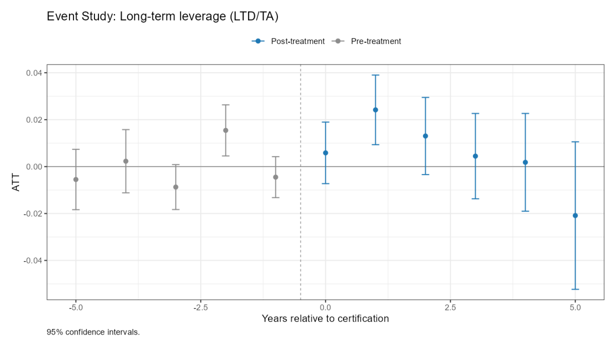

(b) Long-term leverage (LTD/TA)

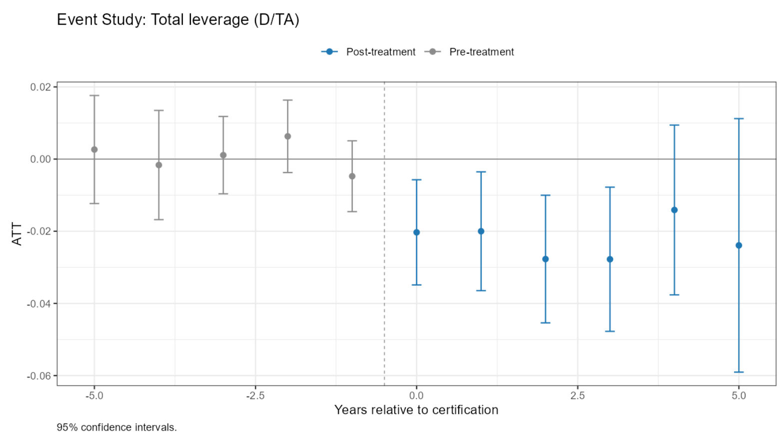

(c) Total leverage (D/TA) — design-sensitive contraction

Figure 1: Event-study ATT: maturity recomposition and the total-leverage contraction.
Notes: CS2021 doubly-robust estimator; never-treated controls; anticipation $\kappa = 2$; $B = 999$ bootstrap (`set.seed=2026`). Bands: 95%. *Source:* Authors' elaboration based on Bureau van Dijk (Moody's Analytics) (2024).

The identifying assumption—conditional parallel trends—is supported on the design-consistent pre-period ($e \leq -3$, excluding the $\kappa = 2$ anticipation window). A joint Wald test fails to reject the null of zero pre-treatment effects for the debt outcomes that carry the deleveraging result—total leverage ($p = 0.980$), short-term leverage ($p = 0.424$), and long-term leverage ($p = 0.287$)—and for shareholders' funds ($p = 0.710$); the corresponding event-study pre-periods in Figure 1 are flat. The one outcome with a borderline pre-trend is the share premium ($p = 0.093$), where a small positive movement in the anticipation years is consistent with investors subscribing equity ahead of formal registration. We treat this honestly rather than explain it away: in Section 5 we report relative-magnitude sensitivity bounds (Rambachan and Roth, 2023) for every outcome. Those bounds are reassuring for the share-premium effect, which remains informative under deviations up to half the maximum pre-trend, but they are demanding for the debt effects, whose confidence intervals admit ezero once even modest parallel-trends violations are allowed. The deleveraging is thus robustly estimated under exact parallel trends, passes the joint pre-trend test, and is stable across matching designs and cohorts, but it is not immune to the most conservative trend-violation bounds—a caveat we carry through to the conclusions.

4.3 The equity channel

Certification also activates a targeted equity response (H1), and its specificity is informative about the channel. Since 2017, investors subscribing equity in certified firms receive a 30% tax relief (Section 2); if the programme works through this instrument, the response should appear in *paid-in capital*—the share premium recorded when new equity is issued above par—rather than in equity components that accumulate mechanically with profitability. This is what we find. The ATT on share premium is 1.48 pp ($p =< 0.001$), and aggregate shareholders' funds rise by a comparable 1.66 pp ($p = 0.037$). Three features tie the equity increase to new subscriptions rather than to internally generated funds. First, the two magnitudes are similar, so most of the rise in total equity is accounted for by paid-in share premium—the line item that records equity issued above par. Second, profitability falls after certification (return on assets, -1.14 pp, $p = 0.003$; Section 4.4), so retained earnings cannot be the source of the additional equity: the recapitalisation must come from fresh outside capital. Third, and most directly, the response is confined to the premium itself: nominal capital stock barely moves (+0.30 pp, $p = 0.544$, not significant), so the new equity enters as capital paid *above* par rather than as par-value subscriptions—the accounting signature of valuation-based rounds priced by outside investors, not of founders topping up nominal capital (Appendix J). This is exactly the margin the 30% investor tax relief is designed to mobilise. This equity activation, the margin on which we rest the paper's causal claims, moves the specific target of the policy; alongside the (design-sensitive) debt contraction documented above, it describes a recapitalisation rather than a balance-sheet expansion. Consistent with this reading, the share-premium effect is the most stable result in the paper, near-identical across the matching designs and in its full dynamic path: it rises from 1.23 pp at impact to 1.75 pp by the fourth post-treatment year, with no reversal. This building path is hard to reconcile with a timing-driven reading: a one-off capital increase at registration would jump and then plateau, not keep accumulating. Figure 2 overlays the two equity event studies under each design; for share premium the four estimators are visually indistinguishable, with near-coincident confidence bands, so the effect does not depend on the matching scheme. The figure does show a mild pre-certification rise in share premium, which we read as investors subscribing ahead of formal registration (and why we adopt the $\kappa = 2$ window); the converse—firms timing certification to an equity round that captures the relief—cannot be ruled out from this alone, which is the equity channel's one real exposure, so we test it directly in Section 5, where the effect survives excluding firms that raised equity before registration. The shareholders'-funds path (lower panel) is the same in sign but noisier, as expected for an aggregate that mixes the subscription channel with retained earnings; the contrast reinforces that the response concentrates in the precise instrument the policy targets, not equity at large.

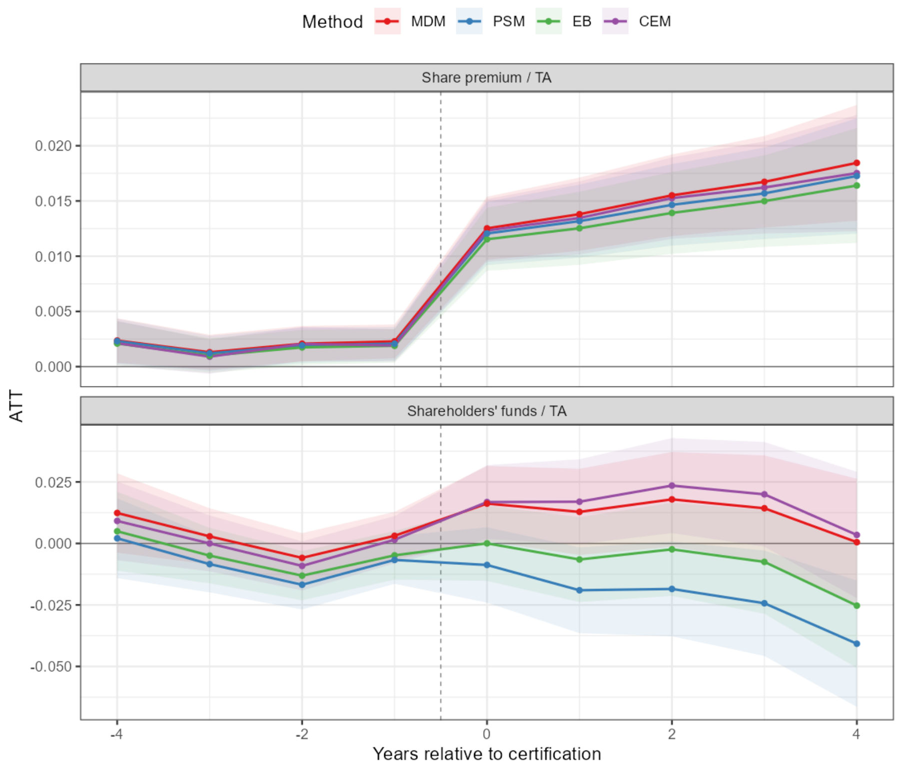

Figure 2: Equity event-study ATTs across matching designs (MDM, PSM, entropy balancing, CEM).

Notes: CS2021 doubly-robust event-study estimates for share premium and shareholders' funds under four matching designs; never-treated controls; anticipation $\kappa = 2$; analytic standard errors. Bands: 95%. *Source:* Authors' elaboration based on Bureau van Dijk (Moody's Analytics) (2024).

4.4 Real consequences: profitability and the cost of debt

If certification reshapes the liability side, does it leave a mark on financial performance? We extend the staggered design to the cost of debt (financial charges over total payables), interest coverage (EBIT over charges), debt servicing (charges over EBITDA), and return on assets. The price and servicing of debt are unaffected (cost of debt $p = 0.929$; interest coverage $p = 0.872$; debt servicing $p = 0.343$). Because the payables stock cannot be split into financial and trade components, these interest-expense flows are the only window on the *financial* margin the guarantee is meant to move—and on it we find no price effect, consistent with an effect on the *quantity and composition* of financing rather than its price. We read these as the absence of a detectable price effect, not precise zeros: scaled by total payables, financial charges are a noisy proxy for the interest rate (Section 3.1).

The one performance margin that does move is profitability: return on assets falls (-1.14 pp, $p = 0.003$) after certification. The decline is modest but precisely estimated, and decomposing it shows that it is *not*, as a simple reading might suggest, the administrative burden of documenting R&D—the causal rise in R&D intensity is far too small to carry it. Two distinct forces do (Appendix Table 21). First, a denominator effect: the asset base expands steeply after certification (ln total assets rise by $+43.4$ log points, $p < 0.001$) as firms draw in the new equity and scale up, and because return on assets divides profit by assets this expansion depresses the ratio even at unchanged profitability. Second, a genuine numerator effect: the operating margin itself compresses (-2.67 pp on profit over revenues, $p = 0.022$), consistent with young firms front-loading investment and qualified personnel ahead of returns that accrue over a longer horizon. Both forces point the same way—certified firms are recapitalising and growing into long-horizon, intangible-intensive investment whose payoff is deferred—and this sits naturally with the rest of the evidence: firms whose value rests on long-horizon, low-collateral investment are precisely those for which equity, rather than debt, is the appropriate financing instrument, which is what the recapitalisation documented above achieves. Read together, this is the paradox at the centre of the paper: certification helps firms recapitalise and tilt their financing towards equity suited to long-horizon innovation, but it does so at a measurable cost to current profitability rather than by lowering the price of credit.

4.5 Mechanism

The deleveraging under the primary design—the design-sensitive margin we do not certify as causal (Section 5)—would, taken at face value, puzzle the instrument most readers associate with the programme: a sovereign credit guarantee should expand borrowing, yet here it does not. We read what follows as an interpretation of *why*, conditional on the debt estimates, not a second causal claim: the guarantee acts on the *margin* of collateral dependence without raising the *volume* of credit. On the matched panel we estimate a firm- and year-fixed-effects regression of long-term leverage on certification interacted with a covariate X (tangibility or intangibility); the interaction coefficient $\hat{\delta}_{X \times PMI}$ measures the post-certification change in the within-firm sensitivity of long-term leverage to X (equation (2) and Table 14 in Appendix D). The post-certification sensitivity of long-term leverage to a firm's own tangibility falls ($\hat{\delta}_{\text{Tang} \times PMI} = -0.047$, $p = 0.030$), consistent with collateral substitution, and the sensitivity to intangibility falls alongside it ($\hat{\delta}_{\text{Intang} \times PMI} = -0.032$, $p = 0.050$). Once the public guarantee can stand in for collateral, term borrowing depends less on the firm's own asset base—tangible or intangible—and not only on its pledgeable tangibles. Because the firm fixed effect absorbs the post-certification level shift, this interaction is identified from within-firm changes in slope and is associational, not a causal contrast; being a two-way fixed-effects estimate under staggered adoption it is also exposed to the heterogeneity-weighting concerns of Goodman-Bacon (2021). The Callaway and Sant'Anna (2021) estimator we use for the headline effects identifies ATTs on *levels* and has no direct analogue for a within-firm *slope* on a continuous covariate, which is precisely why the causal claims rest on the level estimates and this interaction is read only as corroborating texture, not as independent evidence.

Yet the loosened collateral constraint does not coincide with more debt—total and short-term leverage do not rise. The constraint that binds these intangible, long-horizon firms is informational rather than collateral, so relaxing collateral does not make debt the right instrument; meanwhile the programme cheapens the right one—the 30% relief plus the quality signal certification sends—which firms use to retire roll-over-prone short-term debt. The net effect is one of demand, not supply: a guarantee shifts credit supply, but recapitalising firms simply do not draw on it. The guarantee, the programme's headline credit instrument, does not lever firms up; the equity incentive does the work. We turn next to a more direct test of this information channel, exploiting the eligibility criteria each firm satisfies.

Eligibility rests on at least two of three statutory criteria: R&D spending of at least 3% of the larger of cost or production value; a qualified workforce; or registered intellectual property. These differ in how legible they are to outside financiers, and that difference lets us adjudicate between two readings of the mechanism (H3). Under an *opacity* reading, certification works by attesting to hard-to-verify quality, so its effect should be largest for firms qualifying through R&D—the least verifiable, least pledgeable input—and smallest for those holding registered intellectual property, which already carries a verifiable signal: $|\mathrm{ATT}_{\mathrm{R\&D}}| > |\mathrm{ATT}_{\mathrm{IP}}|$. Under a *fundability* reading, the verifiable, pledgeable signal complements the equity subsidy instead of making it redundant, so the equity response should be largest precisely for the IP holders—the firms whose innovation an outside investor can most readily price and underwrite: $|\mathrm{ATT}_{\mathrm{IP}}| > |\mathrm{ATT}_{\mathrm{R\&D}}|$. We estimate the ATT separately for firms that do and do not qualify through each criterion, each against the common never-treated pool, with a two-way fixed-effects interaction as the formal difference test (Table 5). Because a firm may satisfy several criteria, the subgroups overlap: the comparison is conditional on *whether* a firm qualifies through a given channel, not a partition of the treated.

The data favour the fundability reading. The equity response is concentrated among the IP holders: firms qualifying through registered intellectual property raise share premium by 1.77 pp, against 0.82 pp for those that do not, a difference of 0.95 pp that is the only

Table 5: Equity and debt response by eligibility criterion (CS2021 ATT).

Outcome	Criterion	Qualifies ($r = 1$) ATT (pp)	(s.e.)	ATT if not ($r = 0$, pp)	Difference (r=1 vs r=0), p
Total leverage (D/TA)	R&D (≥3%)	−1.90	(0.85)	−3.80	+1.90 (0.310)
	Qualified team	−1.58	(1.03)	−3.19	+1.60 (0.287)
	Patent / software	−2.86	(0.88)	−0.84	−2.02 (0.237)
Short-term leverage (STD/TA)	R&D (≥3%)	−2.70	(0.83)	−5.69	+2.99 (0.145)
	Qualified team	−2.78	(1.06)	−3.94	+1.16 (0.440)
	Patent / software	−3.44	(0.90)	−2.91	−0.53 (0.756)
Share premium / TA	R&D (≥3%)	+1.39	(0.18)	+1.91	−0.52 (0.242)
	Qualified team	+1.36	(0.22)	+1.66	−0.31 (0.366)
	Patent / software	+1.77	(0.20)	+0.82	+0.95*** (0.007)
Shareholders' funds / TA	R&D (≥3%)	+1.31	(0.90)	+3.34	−2.03 (0.328)
	Qualified team	+1.08	(1.10)	+2.53	−1.45 (0.371)
	Patent / software	+2.42	(0.96)	+0.01	+2.41 (0.179)

Notes: Each treated subgroup is matched against the common never-treated pool; ATTs are CS2021 doubly-robust estimates (anticipation $\kappa = 2$) on the primary (CEM) matched sample. Subgroups overlap, since a firm may meet several criteria, so the comparison is conditional on whether a firm qualifies through a given channel. The difference column reports the doubly-robust ATT difference (criterion r=1 minus r=0). Its standard error is covariance-aware where the Callaway and Sant'Anna (2021) influence functions align across the two subgroups and otherwise the conservative $\sqrt{\mathrm{se}_1^2 + \mathrm{se}_0^2}$, which *overstates* it because the shared never-treated controls induce positive covariance between the two ATTs. $^*p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

statistically significant split in the table ($p = 0.007$). The R&D and workforce criteria produce no significant difference on the equity margin ($p = 0.242$ and $p = 0.366$), and no criterion generates a significant difference on total leverage ($p = 0.237$ for the IP split). The opacity prediction—that the least verifiable input drives the response—is not borne out; if anything the equity difference runs the other way. Equity responds where innovation is most fundable: the firms that recapitalise most are those whose value rests on a registered, externally verifiable asset an equity investor can price. The difference is estimated with the conservative standard error, so it does not rest on a favourable covariance adjustment; within the hypothesis-relevant family of share-premium comparisons it survives a Holm multiple-testing correction ($p = 0.021$), though across the full battery of criterion-by-outcome tests it is only marginal ($p = 0.084$). Because the subgroups overlap and we test several criteria against several outcomes, we therefore read this split as suggestive rather than decisive: it is the contrast the fundability account predicts and the opacity account does not, and it refines the headline result rather than carrying it. Certification activates equity through the policy-subsidised channel, and most where the signal is most legible to the investors the 30% tax relief is designed to attract.

5 Robustness and Causal Credibility

The credibility of the two findings—the deleveraging and the equity recapitalisation—rests on more than a single design choice. We subject them to five checks: a multi-design triangulation, anticipation-window sensitivity, a leave-one-cohort-out exercise, a placebo timing test, and formal parallel-trends bounds. We state plainly which conclusions hold throughout and which are design-sensitive: the equity response is robust on every check, the deleveraging on the designs that retain explicit control firms and on the design-based falsifications, while it attenuates under reweighting and under the most conservative trend bounds.

5.1 Multi-design triangulation

We re-estimate the entire design under five matching and weighting schemes: coarsened exact matching (CEM, our primary design), entropy balancing (EB), Mahalanobis distance matching (MDM), the covariate-balancing propensity score (CBPS), and propensity-score matching (PSM) (Table 6). The equity response is near-identical across all five (share premium between 1.38 and 1.53 pp, all $p < 0.001$). The debt results split along a clear matching-versus-weighting line. The designs that retain explicit, balanced control firms—CEM and MDM—show significant contractions in total and short-term leverage. The reweighting designs—EB and CBPS—attenuate the debt effects to statistically zero, with the total-leverage point estimate even turning marginally positive while remaining insignificant. To attain exact moment balance under limited overlap, both concentrate weight on a small set of controls, so the effective sample behind the weighted moments is far smaller than the nominal one. Table 7 makes this concrete: the Kish effective sample of the reweighting designs collapses to a few thousand controls, well below that of the primary CEM design, and a single control carries more than twice the maximum weight it does under CEM—the weight concentration that drives the attenuation. Their nominal precision is comparable to the matched designs—the standard errors are within a few percent of one another across all five methods—so the divergence is one of point estimate, not noise: the exact-balance solution leans on few controls and pulls the counterfactual toward the treated firms' own path, biasing the debt contrast toward zero. PSM, which collapses the covariates into a single index, can leave them more imbalanced than before matching (King and Nielsen, 2019); here it returns sign-flipped, anomalous debt estimates, so we report it as a cautionary case rather than a decisive check. The dynamic event study (Figure 3) makes the same point: the short-term path lies below zero after certification under every design and the equity path steps up and stays elevated under every design, while the total-leverage path declines only under the explicit-matching designs.

The cross-design disagreement is confined to total leverage, and it has a transparent source. Total leverage is the sum of two components that move in opposite directions: short-term debt contracts under all five designs, while long-term debt is positive under all five (Table 6). Total leverage is therefore the small net of two offsetting margins, and its sign turns on which one dominates—the short-term contraction under the explicit-matching designs, the long-term increase under PSM. The directionally robust facts are the components themselves: a short-term contraction and an equity expansion present in every design. We read the net deleveraging as the verdict of the designs that best handle the limited overlap (CEM and MDM), and we treat the sign of total leverage as the most design-sensitive figure in the paper. This is why the recapitalisation interpretation is anchored on the short-term and equity margins, which are directionally stable, rather than on the headline total.

Table 6: Causal Robustness Across Matching Designs: CS2021 ATT

Outcome	CEM	EB	MDM	CBPS	PSM
Total leverage (D/TA)	-0.0223***	0.0072	-0.0175**	0.0069	0.0210***
Short-term leverage (STD/TA)	-0.0324***	-0.0072	-0.0287***	-0.0074	-0.0042
Long-term leverage (LTD/TA)	0.0091	0.0135*	0.0116*	0.0134*	0.0249***
Share premium / TA	0.0148***	0.0138***	0.0153***	0.0138***	0.0145***
Shareholders' funds / TA	0.0166**	-0.0067	0.0134	-0.0064	-0.0204**

Notes: Callaway and Sant'Anna (2021) doubly-robust ATT under four designs: Mahalanobis distance matching (MDM, baseline), propensity-score matching (PSM), entropy balancing (EB, Hainmueller, 2012), and coarsened exact matching (CEM, Iacus et al., 2012). All columns use analytic CS2021 standard errors (`bstrap=FALSE`) for strict comparability; point estimates match the bootstrap runs in the main text. Never-treated controls, anticipation $\kappa = 2$, same covariate vector throughout. EB and CEM reweight the full never-treated pool. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. *Source:* Authors' elaboration based on Bureau van Dijk (Moody's Analytics) (2024).

Table 7: Control-sample weight concentration across designs

Design	Controls	Effective N	ESS/N (%)	Max wt. (%)
CEM	220,992	6,087	2.8	0.395
EB	273,467	3,356	1.2	0.896
MDM	6,497	6,163	94.9	0.076
CBPS	311,597	3,364	1.1	0.895
PSM	7,956	7,956	100.0	0.013

Notes: Kish effective sample size ESS $= (\sum w_c)^2 / \sum w_c^2$ for the control weights of each design, with the share of total control weight carried by the single largest-weighted control. The relevant comparison is the *absolute* effective N and the maximum weight, not the ESS/N ratio: CEM's ratio is low only because it retains the full never-treated pool. The exact-moment reweighting designs (EB, CBPS) concentrate weight on few controls under limited overlap—roughly half CEM's effective N and over twice its maximum weight—collapsing the effective sample, whereas the explicit-matching designs (CEM, MDM) keep near-uniform weights. *Source:* Authors' elaboration based on Bureau van Dijk (Moody's Analytics) (2024).

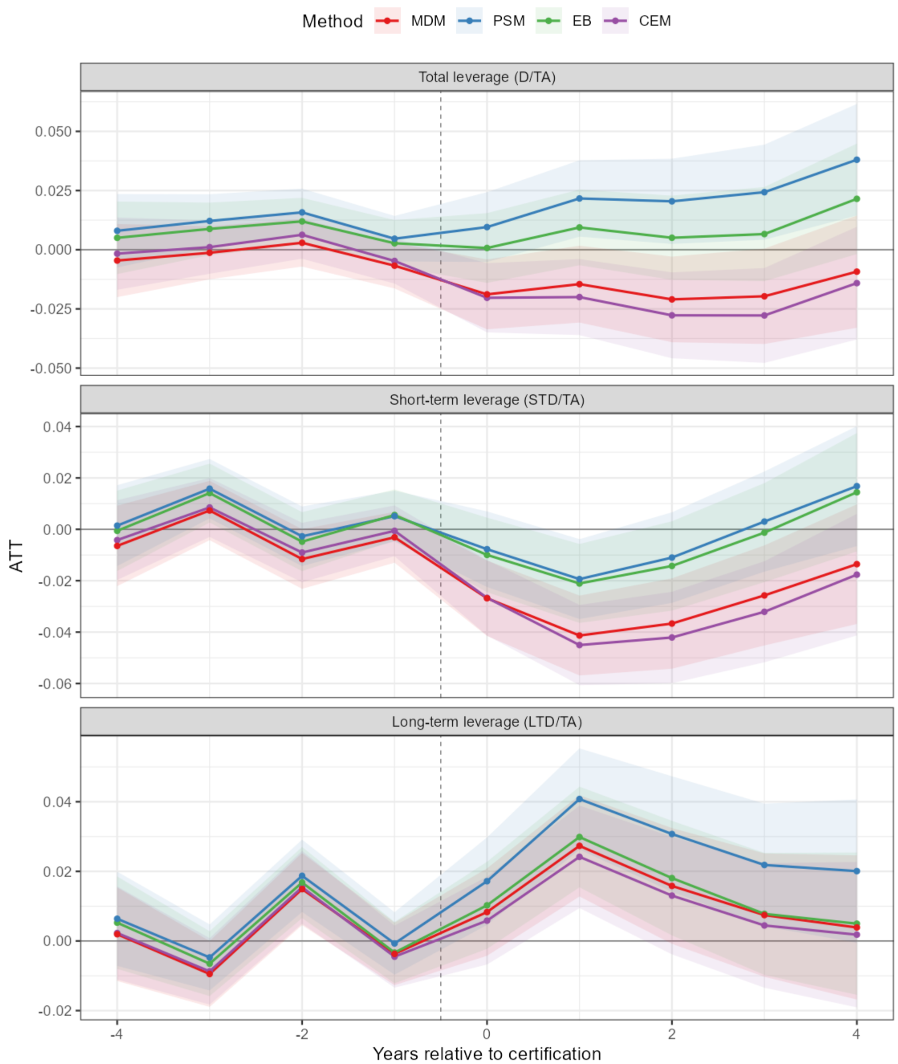

Figure 3: Event-study ATTs across matching and weighting designs.

Notes: CS2021 doubly-robust event-study estimates under five designs; never-treated controls; anticipation $\kappa = 2$; analytic standard errors for cross-design comparability. Bands: 95%. *Source:* Authors' elaboration based on Bureau van Dijk (Moody's Analytics) (2024).

5.2 Anticipation window

Because investors and lenders may react before the registration date, we treat the two pre-registration years as a potentially contaminated window. Table 8 re-estimates the ATTs under $\kappa \in \{0, 1, 2\}$. The deleveraging is stable across all three windows—the total and short-term contractions retain their sign, magnitude, and significance—so the conclusion does not hinge on the choice of anticipation horizon.

Table 8: CS2021 DiD — Anticipation Window Sensitivity ($\kappa = 0, 1, 2$)

Panel A: Anticipation window sensitivity ($\kappa = 0, 1, 2$)

Outcome	κ	ATT	SE	Sig	Pre-trend p
Total leverage (D/TA)	0	-0.0199	0.0058	***	0.668
Short-term leverage (STD/TA)	0	-0.0149	0.0057	***	0.461
Long-term leverage (LTD/TA)	0	-0.0056	0.0047		0.036
Total leverage (D/TA)	1	-0.0324	0.0069	***	0.783
Short-term leverage (STD/TA)	1	-0.0186	0.0062	***	0.337
Long-term leverage (LTD/TA)	1	-0.0145	0.0067	**	0.023
Total leverage (D/TA)	2	-0.0223	0.0075	***	0.979
Short-term leverage (STD/TA)	2	-0.0324	0.0077	***	0.424
Long-term leverage (LTD/TA)	2	0.0091	0.0070		0.287

Panel B:

Primary-design weights vs. no weights ($\kappa = 2$)

Outcome	Specification	ATT	SE	Sig
Long-term leverage (LTD/TA)	No weights, kappa=2	0.0038	0.0071	
Long-term leverage (LTD/TA)	Primary-design weights, kappa=2 (baseline)	0.0091	0.0070	
Short-term leverage (STD/TA)	No weights, kappa=2	-0.0331	0.0075	***
Short-term leverage (STD/TA)	Primary-design weights, kappa=2 (baseline)	-0.0324	0.0077	***
Total leverage (D/TA)	No weights, kappa=2	-0.0292	0.0076	***
Total leverage (D/TA)	Primary-design weights, kappa=2 (baseline)	-0.0223	0.0075	***

Notes: Callaway and Sant'Anna (2021) doubly-robust estimator; never-treated controls; Primary-design (CEM) matching weights; $B = 999$ bootstrap (`set.seed = 2026`). ATT: overall average treatment effect on the treated. Pre-trends: joint χ^2 Wald test on $e \leq -(\kappa + 1)$ using the influence-function covariance. The LTD sign reversal from $\kappa < 2$ to $\kappa = 2$ confirms anticipatory LTD accumulation in the two years before registration; $\kappa = 2$ is the appropriate window. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. *Source:* Authors' elaboration based on Bureau van Dijk (Moody's Analytics) (2024).

5.3 Cohort robustness

Because the aggregate ATT is a cohort-weighted average of group-time effects, a natural concern is that one certification cohort drives the result. A leave-one-cohort-out exercise (Table 9) re-estimates the overall ATT after dropping each cohort's treated firms in turn, keeping the never-treated controls and the remaining treated cohorts. The total and short-term contractions are sign-stable and significant across all leave-out estimates, so no single cohort is responsible for the deleveraging.

Table 9: Leave-One-Cohort-Out (LOCO) Robustness of the CS2021 ATT

Outcome	ATT (baseline)	LOCO min	LOCO max	Sign stable	Sig. stable
Long-term leverage (LTD/TA)	0.0091	0.0042	0.0129	Yes	Yes
Short-term leverage (STD/TA)	-0.0324***	-0.0384	-0.0273	Yes	Yes
Total leverage (D/TA)	-0.0223***	-0.0249	-0.0177	Yes	Yes

Notes: Each model re-estimates the overall ATT after dropping one treated cohort at a time (8 leave-out estimates per outcome); never-treated controls and the remaining treated cohorts are retained. Headline specification: Callaway and Sant'Anna (2021) doubly-robust estimator, never-treated control group, anticipation $\kappa = 2$, $B = 999$ bootstrap replications (`set.seed = 2026`), primary-design (CEM) matching weights. "LOCO min/max" report the range of leave-out ATT estimates. "Sign stable" indicates all leave-out estimates share the baseline sign; "Sig. stable" indicates all leave-out estimates match the baseline's 10% significance status. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. *Source:* Authors' elaboration based on Bureau van Dijk (Moody's Analytics) (2024).

5.4 Placebo timing

A complementary falsification varies treatment *timing* rather than assignment. Restricting the sample to never-treated firms, we assign fake certification years drawn to reproduce the real cohort-size distribution, re-estimate the ATT, and repeat the exercise to trace the randomization-inference null (Table 10; Figure 4). The placebo distributions are centred on zero, so the staggered estimator does not manufacture effects. The test sharpens the asymmetry between the two margins. The share-premium effect lies well outside the reshuffled-timing null (RI $p = 0.010$): paid-in equity is a stable quantity, so random certification dates almost never reproduce a rise of its size. The debt contractions, by contrast, are not separable from the null (total leverage RI $p = 0.383$; short-term RI $p = 0.239$): leverage is volatile enough that reshuffled timing generates movements of comparable magnitude. This is a further reason we anchor the paper's causal claims on the equity margin and read the deleveraging as suggestive.

Table 10: Placebo Timing Test — Randomization Inference on Fake Treatment Dates

Outcome	Real ATT	Placebo mean	Placebo SD	95% placebo interval	RI p
Total leverage (D/TA)	-0.0223	0.0012	0.0314	[-0.0700, 0.0622]	0.383
Short-term leverage (STD/TA)	-0.0324	0.0010	0.0341	[-0.0784, 0.0602]	0.239
Share premium / TA	0.0148	-0.0004	0.0031	[-0.0058, 0.0069]	0.010

Notes: Randomization-inference placebo test on the full never-treated sample. In each of 200 replications a random subset is assigned fake certification years drawn to reproduce the real treated cohort-size distribution, and the overall ATT is re-estimated with the identical headline Callaway and Sant'Anna (2021) doubly-robust estimator (never-treated controls, anticipation $\kappa = 2$, primary-design CEM weights), aggregated to the overall ATT; the real ATT is computed by the same routine. No inner bootstrap is used per draw—inference comes from the across-draw randomization distribution. "Placebo mean/SD" summarise the null distribution of the ATT; the 95% placebo interval is the 2.5–97.5 percentile range. The randomization p-value is the share of placebo estimates whose absolute value weakly exceeds the real |ATT|. *Source:* Authors' elaboration based on Bureau van Dijk (Moody's Analytics) (2024).

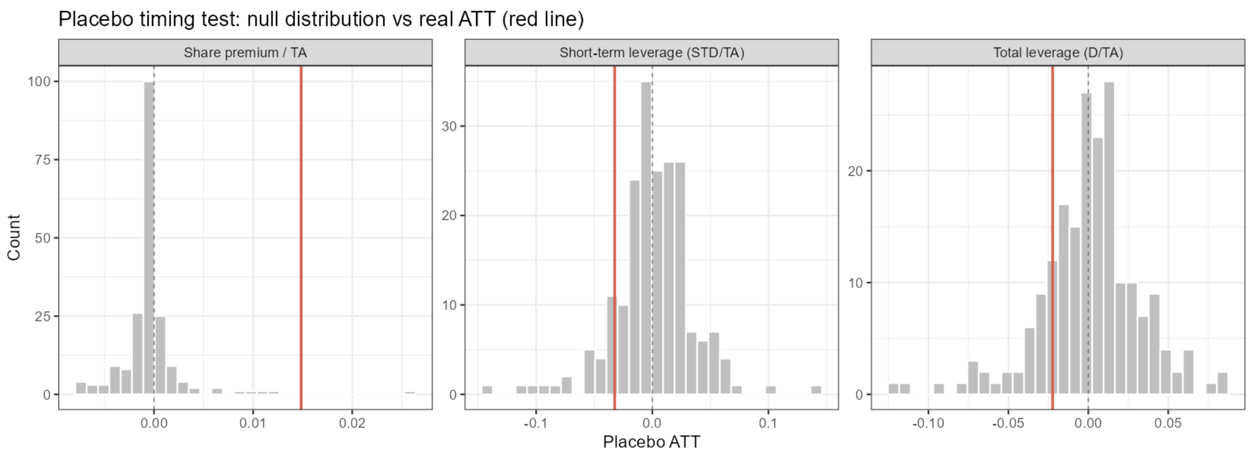

Figure 4: Placebo timing test: randomization-inference null versus the real estimate.
Notes: Each panel plots the distribution of placebo ATTs from fake certification dates assigned to never-treated firms (drawn to match the real cohort-size distribution); the dashed line marks zero and the solid line the real ATT. *Source:* Authors' elaboration based on Bureau van Dijk (Moody's Analytics) (2024).

5.5 Parallel-trends sensitivity

Finally, we ask how large a deviation from parallel trends it would take to overturn each conclusion, using the relative-magnitude and smoothness bounds of Rambachan and Roth (2023) (Table 17). For this test the event study is re-estimated without anticipation, so that the pre-period coefficients that bound the post-treatment violation are genuine and contiguous; the resulting overall ATTs differ slightly from the headline $\kappa = 2$ estimates but are the appropriate baseline here. Because this baseline treats the $e = -1, -2$ anticipation years—the very window the $\kappa = 2$ design quarantines—as ordinary pre-periods, any movement in them feeds into the bound, so the share-premium result must survive deviations constructed from the same pre-trend the headline design sets aside: this makes its survival a more demanding test, not a more lenient one.

The two sides of the balance sheet behave very differently, and that contrast is the centre of our interpretation. The equity response is durable: the robust confidence interval for paid-in share premium continues to exclude zero under post-treatment violations up to $\bar{M} = 0.50$ times the worst observed pre-trend, and under the smoothness restriction it survives deviations of up to one standard error. It also survives a direct reverse-causality check: the share-premium effect persists (+1.13 pp, $p < 0.001$) among the 1,864 certified firms that raised no paid-in capital in the two years before registration (Appendix G), so it is not an artefact of firms certifying mid-recapitalisation. The rise in subscribed equity is therefore not an artefact of differential trends; it is a genuine causal response of the firm's capital base to certification, the empirical signature of the programme acting as a catalyst for equity—a positive, externally legible signal that draws in outside investors and prompts existing owners to capitalise the firm.

The debt response does not clear the same bar. The contraction in total and short-term leverage is significant under exact parallel trends, stable across matching designs and cohorts, and unchanged when pre-treatment leverage is removed from the doubly-robust adjustment so that the lagged outcome is not used twice (Appendix H)—but it does not separate from reshuffled certification timing (Section 5.4): the volatility of leverage lets random dates reproduce contractions of similar size. And its robust intervals admit zero at the smallest violation we consider, which means we cannot rule out that part or all of the measured decline reflects pre-existing differences in leverage trajectories rather than the certification itself. The honest reading is that the debt channel is, at best, weakly identified: we cannot certify that certification causally lowers firms' reliance on bank or trade debt.

This asymmetry lets us say precisely what the programme does and does not do. Certification raises treated firms' equity base, and that claim withstands demanding violations of the identifying assumption; the accompanying reduction in debt does not survive the same scrutiny. We therefore rest our conclusions on the equity margin, where the certification demonstrably bites, and treat the debt margin as a transparently reported but design-sensitive companion.

6 Conclusions

Using a near-universal panel of 320,599 Italian SMEs and a staggered difference-in-differences design with matched controls, we ask whether public innovation certification expands debt capacity or reshapes the financing mix. The robust answer is that *PMI innovativa* recapitalises firms: it raises paid-in share premium (1.48 pp, $p =< 0.001$), with total equity rising by a comparable amount (1.66 pp), through capital paid in *above par* by outside investors rather than retained earnings or founders' par-value top-ups—return on assets falls (-1.14 pp, $p =$ 0.003), reflecting both the asset-base expansion the recapitalisation funds and a compression in operating margins as firms invest ahead of longer-horizon returns. This equity response is the targeted effect of the instrument the decree subsidises, it accumulates through the fourth post-certification year, and it concentrates among firms holding registered intellectual property—those whose innovation an outside investor can most readily value—rather than the most opaque, R&D-qualifying firms. On the debt side the guarantee does not lever firms up: leverage does not rise, and if anything the ratio falls under the primary design (-2.23 pp total, -3.24 pp short-term), but this contraction does not survive our most demanding identification checks, so we rest no causal claim on it. Because the outcomes are ratios and the asset base expands, this concerns the financing *mix* (which tilts to equity), not the absolute volume of borrowing. Read through capital-structure theory, certification moves firms toward the equity their intangible, long-horizon assets call for rather than toward more borrowing.

We hold these conclusions to a demanding identification standard and report it even-handedly. The equity response is robust across all five matching and weighting designs and remains bounded away from zero under relative-magnitude parallel-trends bounds; its one exposure—that firms might time registration to an equity round—is mitigated by its continued accumulation well beyond the registration window and by its running through subscribed capital rather than retained earnings. The deleveraging is weaker. It is stable across the matching designs that retain explicit control firms (coarsened exact matching and Mahalanobis matching), the anticipation window, and a leave-one-cohort-out exercise; but it is not separable from a randomization-inference timing placebo, it attenuates under reweighting designs, where the extreme weights exact balance requires under limited overlap shrink the effective control sample, and it admits zero under the most demanding trend bounds. Mechanism evidence sharpens the equity finding: the response concentrates among firms holding registered intellectual property—the asset an outside investor can most readily value—rather than among the most opaque, R&D-qualifying firms, indicating that what certification mobilises is investable equity, not a mechanical take-up of the subsidy.

For policy, the lesson sharpens once robustness is taken into account. The programme demonstrably works on one margin: it strengthens the equity base of treated firms, a causal effect that survives every design and every parallel-trends bound we impose. What it does *not* do is causally change how much these firms borrow. The apparent deleveraging is precisely the headline an evaluation attentive only to the volume or maturity of *debt* would record as a success, yet it does not withstand demanding trend-sensitivity bounds; we therefore decline to read it as a robust causal effect. The instrument that carries the durable result is the equity incentive (the 30% investor tax relief together with the certification signal), while the credit guarantee, the instrument designed to expand borrowing, leaves no effect we can certify. The broader lesson is for the design of innovation *policy mixes*: when a programme bundles many instruments, the one that moves financing need not be the headline credit lever. Here, where a credible certification and an equity subsidy are jointly present, the financing response runs through capital rather than credit, and an evaluator who scores such schemes on loan volumes would misread them.

We have tried to report which results are robust and which are sensitive as plainly as the estimates themselves: the equity response holds across every design and trend bound, whereas the debt-side effects, significant under our primary design, attenuate under reweighting and under the most conservative parallel-trends bounds.

Several limitations bound these conclusions. Identification rests on conditional parallel trends strengthened by matching; although pre-trend tests pass and results are stable across the explicit-control designs, time-varying unobserved confounders cannot be fully ruled out, and certification is voluntary. Leverage is measured from total payables, so the short-term measure includes trade credit and cannot isolate bank borrowing from supplier financing. A regression-discontinuity design around the eligibility thresholds is not feasible—not because the qualifying criteria are unobserved (the register now discloses which criteria each firm satisfies, and we exploit them) but because eligibility is necessary rather than sufficient for the voluntary treatment, qualification requires at least two of three criteria, and the underlying continuous magnitudes that define the thresholds are not reported. Future work could link AIDA to credit-register data to observe loan-level maturities directly and compare *PMI innovativa* with the more selective *startup innovativa* programme to trace the dose–response between certification stringency and capital-structure effects. A further avenue, which the register's ownership information opens but which our sample is too small to settle, is whether the financing response differs for firms led by under-represented entrepreneurs—women, young, or foreign founders—for whom ex-ante credit frictions are tighter; richer founder-level data would let that distributional question be addressed in its own right.

Data Availability Statement

The firm-level financial data analysed in this study were obtained from the AIDA database (Bureau van Dijk) under an institutional licence held by the Università del Salento. Under the terms of that licence the authors are not authorised to redistribute the data. AIDA is accessible to subscribing institutions; the *PMI innovativa* registry is publicly available from the Italian Business Register (*Registro delle Imprese*).

References

Aiello, F., Errico, L., and Rondinella, S. (2024). Innovative SMEs in Italy. Explaining profitability patterns in inner areas. *Journal of Economic Studies*, 51(9):306–322.

Ali, S., Muhammad, H., and Migliori, S. (2024). R&D investment and SMEs performance: the role of capital structure decisions. *EuroMed Journal of Business*, 20(4):1029–1051.

Angrist, J. D. and Pischke, J.-S. (2009). *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press, Princeton, NJ.

Antoniou, A., Guney, Y., and Paudyal, K. (2008). The Determinants of Capital Structure: Capital Market-Oriented versus Bank-Oriented Institutions. *Journal of Financial and Quantitative Analysis*, 43(1):59–92.

Athey, S., Imbens, G., Qu, Z., and Viviano, D. (2026). Triply robust panel estimators. *Journal of Applied Econometrics*.

Austin, P. C. (2011). An Introduction to Propensity Score Methods for Reducing the Effects of Confounding in Observational Studies. *Multivariate Behavioral Research*, 46(3):399–424.

Beck, T., Demirgüç-Kunt, A., and Maksimovic, V. (2008). Financing patterns around the world: Are small firms different? *Journal of Financial Economics*, 89(3):467–487.

Bellucci, A., Borisov, A., Giombini, G., and Zazzaro, A. (2023a). Information asymmetry, external certification, and the cost of bank debt. *Journal of Corporate Finance*, 78:102336.

Bellucci, A., Pennacchio, L., and Zazzaro, A. (2023b). Debt financing of SMEs: The certification role of R&D Subsidies. *International Review of Financial Analysis*, 90:102903.

Bertrand, M., Duflo, E., and Mullainathan, S. (2004). How much should we trust differences-in-differences estimates? *The Quarterly Journal of Economics*, 119(1):249–275.

Bester, H. (1985). Screening vs. Rationing in Credit Markets with Imperfect Information. *The American Economic Review*, 75(4):850–855.

Bhattacharya, S. and Ritter, J. R. (1983). Innovation and Communication: Signalling with Partial Disclosure. *The Review of Economic Studies*.

Biancalani, F., Czarnitzki, D., and Riccaboni, M. (2022). The Italian Start Up Act: a microeconometric program evaluation. *Small Business Economics*, 58(3):1699–1720.

Bonfim, D., Custódio, C., and Raposo, C. (2023). Supporting small firms through recessions and recoveries. *Journal of Financial Economics*, 147(3):658–688.

Brown, J. R., Fazzari, S. M., and Petersen, B. C. (2009). Financing Innovation and Growth: Cash Flow, External Equity, and the 1990s R&D Boom. *The Journal of Finance*, 64(1):151–185. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1540-6261.2008.01431.x.

Bureau van Dijk (Moody's Analytics) (2024). AIDA – analisi informatizzata delle aziende italiane.

Caliendo, M. and Kopeinig, S. (2008). Some practical guidance for the implementation of propensity score matching. *Journal of Economic Surveys*, 22(1):31–72.

Callaway, B. and Sant'Anna, P. H. C. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230.

Chittenden, F., Hall, G., and Hutchinson, P. (1996). Small firm growth, access to capital markets and financial structure: Review of issues and an empirical investigation. *Small Business Economics*, 8(1):59–67.

Colombo, M. G., Grilli, L., and Verga, C. (2007). High-tech Start-up Access to Public Funds and Venture Capital: Evidence from Italy. *International Review of Applied Economics*, 21(3):381–402. _eprint: https://doi.org/10.1080/02692170701390361.

Columba, F., Gambacorta, L., and Mistrulli, P. E. (2009). Mutual Guarantee Institutions and Small Business Finance.

Daskalakis, N., Balios, D., and Dalla, V. (2017). The behaviour of SMEs' capital structure determinants in different macroeconomic states. *Journal of Corporate Finance*, 46:248–260.

de Blasio, G., De Mitri, S., D'Ignazio, A., Finaldi Russo, P., and Stoppani, L. (2018). Public guarantees to SME borrowing. A RDD evaluation. *Journal of Banking & Finance*, 96:73–86.

De Jong, A., Kabir, R., and Nguyen, T. T. (2008). Capital structure around the world: The roles of firm- and country-specific determinants. *Journal of Banking & Finance*, 32(9):1954–1969.

DeAngelo, H. and Masulis, R. W. (1980). Optimal capital structure under corporate and personal taxation. *Journal of Financial Economics*, 8(1):3–29.

Degryse, H., de Goeij, P., and Kappert, P. (2012). The impact of firm and industry characteristics on small firms' capital structure. *Small Business Economics*, 38(4):431–447.

Denes, M., Howell, S. T., Mezzanotti, F., Wang, X., and Xu, T. (2023). Investor tax credits and entrepreneurship: Evidence from U.S. states. *The Journal of Finance*, 78(5):2621–2671.

Dixon, W. J. and Massey, F. J. (1960). *Introduction to statistical analysis*. McGraw-Hill.

D'Amato, A. (2020). Capital structure, debt maturity, and financial crisis: empirical evidence from SMEs. *Small Business Economics*, 55(4):919–941.

European Commission (2003). Commission Recommendation 2003/361/EC concerning the definition of micro, small and medium-sized enterprises.

Feito-Ruiz, I., Cardone-Riportella, C., and Ughetto, E. (2023). Debt maturity and SMEs: Do auditor's quality and ownership structure matter? *Journal of Small Business Management*, 61(4):1736–1772.

Frank, M. Z. and Goyal, V. K. (2009). Capital Structure Decisions: Which Factors Are Reliably Important? *Financial Management*, 38(1):1–37.

Gai, L., Arcuri, M. C., and Ielasi, F. (2023). How does government-backed finance affect SMEs' crisis predictors? *Small Business Economics*, 61(3):1205–1229.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2):254–277.

Governo Italiano (2015). Decreto-Legge 3/2015: Misure urgenti per il sistema bancario e gli investimenti.

Hainmueller, J. (2012). Entropy Balancing for Causal Effects: A Multivariate Reweighting Method to Produce Balanced Samples in Observational Studies. *Political Analysis*, 20(1):25–46.

Hall, B. H. and Lerner, J. (2009). The Financing of R&D and Innovation.

Hall, G. C., Hutchinson, P. J., and Michaelas, N. (2004). Determinants of the Capital Structures of European SMEs. *Journal of Business Finance & Accounting*, 31(5-6):711–728. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.0306-686X.2004.00554.x.

Iacus, S. M., King, G., and Porro, G. (2012). Causal Inference without Balance Checking: Coarsened Exact Matching. *Political Analysis*, 20(1):1–24.

Imai, K. and Ratkovic, M. (2014). Covariate balancing propensity score. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 76(1):243–263.

King, G. and Nielsen, R. (2019). Why Propensity Scores Should Not Be Used for Matching. *Political Analysis*, 27(4):435–454.

Kraus, A. and Litzenberger, R. H. (1973). A State-Preference Model of Optimal Financial Leverage. *The Journal of Finance*, 28(4):911–922. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1540-6261.1973.tb01415.x.

Lee, N., Sameen, H., and Cowling, M. (2015). Access to finance for innovative SMEs since the financial crisis. *Research Policy*, 44(2):370–380.

Leland, H. E. and Pyle, D. H. (1977). Informational Asymmetries, Financial Structure, and Financial Intermediation. *The Journal of Finance*, 32(2):371–387.

Lemmon, M. L., Roberts, M. R., and Zender, J. F. (2008). Back to the Beginning: Persistence and the Cross-Section of Corporate Capital Structure. *The Journal of Finance*, 63(4):1575–1608.

Magri, S. (2009). The financing of small innovative firms: the Italian case. *Economics of Innovation and New Technology*, 18(2):181–204. _eprint: https://doi.org/10.1080/10438590701738016.

Martí, J. and Quas, A. (2018). A beacon in the night: government certification of SMEs towards banks. *Small Business Economics*, 50(2):397–413.

Meuleman, M. and De Maeseneire, W. (2012). Do R&D subsidies affect SMEs' access to external financing? *Research Policy*, 41(3):580–591.

Michaelas, N., Chittenden, F., and Poutziouris, P. (1999). Financial Policy and Capital Structure Choice in U.K. SMEs: Empirical Evidence from Company Panel Data. *Small Business Economics*, 12(2):113–130.

Minard, P. (2016). Signalling through the noise: private certification, information asymmetry and Chinese SMEs' access to finance. *Journal of Asian Public Policy*, 9(3):243–256. _eprint: https://doi.org/10.1080/17516234.2015.1083412.

Ministero dello Sviluppo Economico (2015). Circolare applicativa PMI innovative.

Ministero dello Sviluppo Economico (2022). Agevolazioni a favore delle PMI innovative. https://www.mise.gov.it/index.php/it/impresa/piccole-e-medie-imprese/pmi-innovative. Direzione Generale per la Politica Industriale, l'Innovazione e le PMI.

Modigliani, F. and Miller, M. H. (1958). The Cost of Capital, Corporation Finance and the Theory of Investment. *The American Economic Review*, 48(3):261–297.

Myers, S. C. and Majluf, N. S. (1984). Corporate financing and investment decisions when firms have information that investors do not have. *Journal of Financial Economics*, 13(2):187–221.

Palacín-Sánchez, M. J., Ramírez-Herrera, L. M., and di Pietro, F. (2013). Capital structure of SMEs in Spanish regions. *Small Business Economics*, 41(2):503–519.

Rajan, R. G. and Zingales, L. (1995). What Do We Know about Capital Structure? Some Evidence from International Data. *Journal of Finance*, 50(5):1421–1460.

Rambachan, A. and Roth, J. (2023). A More Credible Approach to Parallel Trends. *Review of Economic Studies*, 90(5):2555–2591.

Rampini, A. A. and Viswanathan, S. (2013). Collateral and capital structure. *Journal of Financial Economics*, 109(2):466–492.

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

Ross, S. A. (1977). The Determination of Financial Structure: The Incentive-Signalling Approach. *The Bell Journal of Economics*, 8(1):23–40.

Scherr, F. C. and Hulburt, H. M. (2001). The Debt Maturity Structure of Small Firms. *Financial Management*.

Serrasqueiro, Z., Nunes, P. M., and da Rocha Armada, M. (2016). Capital structure decisions: old issues, new insights from high-tech small- and medium-sized enterprises. *The European Journal of Finance*, 22(1):59–79.

Spence, M. (1973). Job market signaling. *The Quarterly Journal of Economics*, 87(3):355–374.

Stiglitz, J. E. and Weiss, A. (1981). Credit Rationing in Markets with Imperfect Information. *The American Economic Review*, 71(3):393–410.

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

Titman, S. and Wessels, R. (1988). The Determinants of Capital Structure Choice. *The Journal of Finance*, 43(1):1–19.

Ullah, B. (2020). Signaling value of quality certification: Financing under asymmetric information. *Journal of Multinational Financial Management*, 55:100629.

Unioncamere – Infocamere (2026). Registro delle imprese – PMI innovative.

Wooldridge, J. M. (2021). Two-way fixed effects, the two-way Mundlak regression, and difference-in-differences estimators. *SSRN Working Paper 3906345*.

A Treatment Cohorts

Table 11 reports the distribution of treated firms by first certification year, the cohort structure that motivates the staggered estimator of Section 3.

Table 11: Treatment cohort distribution: year of first *PMI innovativa* certification

Year first treated	N firms	Cumulative N	% of treated	Cumulative %
2017	99	99	5.1%	5.1%
2018	106	205	5.5%	10.6%
2019	220	425	11.4%	22.0%
2020	241	666	12.5%	34.5%
2021	239	905	12.4%	46.9%
2022	237	1142	12.3%	59.2%
2023	382	1524	19.8%	79.0%
2024	406	1930	21.0%	100.0%
Total	1930	—	100.0%	—

Notes: Each treated firm is assigned to the cohort of its first *PMI innovativa* certification, on the matched (GOLD) sample used in the staggered DiD. The 378 matched treated firms first certified after 2024 fall outside the 2015–2024 panel window, have no post-treatment observation, and so do not contribute treated periods to the estimator. Cohort sizes determine the precision of cohort-specific ATT estimates in the Callaway–Sant'Anna (2021) staggered DiD design. Source: Authors' elaboration based on Unioncamere – Infocamere (2026).

B Full-Panel OLS/FE Determinants

Tables 12–13 report panel fixed-effects estimates of leverage determinants on the full SME panel. The EBITDA ratio enters strongly negative (pecking order); tangibility and intangibility load negatively on short-term and positively on long-term leverage (maturity matching) with near-zero total-leverage effects; R&D intensity loads positively on short-term but not long-term leverage, establishing the baseline that certification is hypothesised to shift. Within the certified subsample the size, age, and profitability sensitivities attenuate or reverse—consistent with certification substituting a direct quality signal for the soft creditworthiness proxies lenders otherwise rely on—a pattern we summarise here and report in full on request.

Table 12: Panel Fixed-Effects Estimates — Full Italian SME Sample

	Dep. variable: Total leverage (Payables/TA)			
	(1) Base	(2) +Intang.	(3) +R&D	(4) +Cash
Tangibility	0.015***	0.016***	0.002	−0.012***
	(0.002)	(0.002)	(0.003)	(0.003)
Profitability	−0.410***	−0.410***	−0.392***	−0.400***
	(0.002)	(0.002)	(0.002)	(0.003)
Size (ln TA)	0.017***	0.017***	0.017***	0.014***
	(0.001)	(0.001)	(0.001)	(0.001)
Age (ln years)	−0.010***	−0.010***	−0.024***	−0.020***
	(0.001)	(0.001)	(0.002)	(0.002)
Asset growth	0.083***	0.083***	0.086***	0.080***
	(0.001)	(0.001)	(0.001)	(0.001)
NDTS	0.504***	0.497***	0.499***	0.478***
	(0.008)	(0.008)	(0.012)	(0.013)
Intangibility		0.018***	0.003	0.014**
		(0.005)	(0.008)	(0.007)
R&D intensity			0.407***	
			(0.070)	
Cash ratio				−0.144***
				(0.011)
Firm FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Observations	2,824,722	2,824,722	1,625,050	1,129,923
Within R^2	0.818	0.818	0.839	0.845

Notes: Panel fixed-effects estimates; firm and year fixed effects included in all specifications. Sample period: 2015–2022. Specification (1) covers 2,824,722 firm-year observations after listwise deletion from the 3,201,469-observation analytical panel (Step 4, Table 1); Specifications (3)–(4) use subsamples of firms with R&D and cash disclosures. Dependent variable is total payables divided by total assets (D'Amato, 2020). Standard errors (in parentheses) are clustered at the firm level. NDTS = non-debt tax shield (depreciation/TA). $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. *Source:* Authors' elaboration based on Bureau van Dijk (Moody's Analytics) (2024).

Table 13: Panel Fixed-Effects Estimates — Full SME Sample, STD and LTD Decomposition

	STD/TA		*LTD/TA*	
	(1) Base	(2) +R&D	(1) Base	(2) +R&D
Tangibility	−0.148***	−0.149***	+0.164***	+0.151***
	(0.002)	(0.003)	(0.002)	(0.003)
Profitability	−0.267***	−0.266***	−0.137***	−0.118***
	(0.002)	(0.002)	(0.001)	(0.002)
Size (ln TA)	−0.007***	−0.005***	+0.023***	+0.020***
	(0.001)	(0.001)	(0.001)	(0.001)
Age (ln years)	−0.035***	−0.054***	+0.029***	+0.033***
	(0.001)	(0.002)	(0.001)	(0.002)
Asset growth	+0.082***	+0.084***	+0.002***	+0.002***
	(0.001)	(0.001)	(0.000)	(0.001)
NDTS	+0.465***	+0.502***	+0.038***	−0.002
	(0.008)	(0.011)	(0.007)	(0.009)
R&D intensity		+0.312***		+0.096
		(0.070)		(0.061)
Firm FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Observations	2,824,722	1,625,050	2,824,722	1,625,050
Within R^2	0.758	0.784	0.720	0.756

Notes: Panel fixed-effects estimates; firm and year FE in all specifications. Standard errors (in parentheses) clustered at the firm level. NDTS = depreciation/TA. The R&D column uses the subsample of firms disclosing R&D expenditure (*bilancio ordinario*, approximately 59% of observations). Financial sector (NACE Rev.2 64–66) and real estate (NACE Rev.2 68) excluded. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. *Source:* Authors' elaboration based on Bureau van Dijk (Moody's Analytics) (2024).

C Full-Panel OLS Interaction Battery

The full-panel OLS interaction estimates are the reduced-form analogue of the matched-sample TWFE estimates in Table 14 and deliver the same qualitative pattern—a negative post-certification interaction between *innovative_pmi* and tangibility on long-term leverage ($\hat{\delta}_{\text{Tang}\times PMI} < 0$), the collateral-substitution signature. Because they replicate the matched-sample evidence on a much larger but non-control-comparable sample, the full battery of interactions is omitted here for brevity and is available from the authors on request.

D Matched-Sample TWFE Detail

The mechanism interactions in Section 4.5 are estimated on the matched panel by

$$\text{LTD}_{it} = \alpha_i + \delta_t + \beta\, X_{it} + \theta\,(\text{PMI}_{it} \times X_{it}) + \gamma' \mathbf{Z}_{it} + \varepsilon_{it}, \tag{2}$$

where X_{it} is the firm's tangibility (or intangibility), $\text{PMI}_{it} = 1$ in its certified years, α_i and δ_t are firm and year fixed effects, $\mathbf{Z}_{it}$ collects the remaining controls, standard errors cluster by firm, and $\theta \equiv \hat{\delta}_{X\times PMI}$ is the post-certification change in the within-firm sensitivity of long-term leverage to X. Table 15 reports the real-activity channel (asset composition as the outcome) under the same two-way fixed-effects specification. We do not tabulate a TWFE replication of the level ATTs: those are estimated by the heterogeneity-robust Callaway and Sant'Anna (2021) design in the main text, and a static TWFE on the same outcomes would be the very estimator Section 3.3 flags as biased under staggered timing. Table 14 reports the matched-sample collateral-substitution interactions discussed in Section 4.5.

Table 14: TWFE Causal Interaction Battery: Sensitivity of LTD/TA to Firm Characteristics

Hyp.	**Interaction term**	**Coef.**	**SE**	p	**Sig.**	$N_{\mathbf{obs}}$
H??	PMI × Tangibility	−0.047	0.022	0.030	**	2,006,162
H??	PMI × Intangibility	−0.032	0.016	0.050	**	2,006,162
H??	PMI × R&D†	+0.495	0.326	0.128		1,122,134

Notes: TWFE estimator with firm and year fixed effects, primary-design (CEM) matching weights. GOLD matched panel: 2,308 treated firms; 220,992 matched never-treated controls. Outcome in all panels: long-term leverage (LTD/TA). Each row reports the PMI × asset-variable interaction coefficient (`innovative_pmi` main effect absorbed by firm FE). Standard errors clustered at the firm level. † R&D subsample: firms filing *bilancio ordinario* with separate R&D disclosure (59% of observations). n.s. = not significant at $p < 0.10$. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. *Source:* Authors' elaboration based on Bureau van Dijk (Moody's Analytics) (2024).

Table 15: TWFE Real-Activity Channel: Asset Composition as Outcome (Panel D)

Outcome	**Coef.**	**SE**	p	**Sig.**	$N_{\mathbf{obs}}$
Tangibility (TFA/TA)	−0.019	0.002	0.000	***	2,006,162
Intangibility (IA/TA)	+0.026	0.004	0.000	***	2,006,162
R&D intensity (R&D/TA)	+0.002	0.000	0.000	***	1,122,134
Cash ratio (Cash/TA)†	+0.004	0.001	0.003	***	817,293

Notes: TWFE estimator, firm and year FE, primary-design (CEM) matching weights, full GOLD matched sample (2,308 treated; 220,992 controls; 223,300 firms). Each row is a separate regression with the listed variable as dependent variable. `innovative_pmi` main effect; controls: EBITDA ratio, ln assets, ln age, asset growth, NDTS. † R&D intensity and cash ratio estimated on firms with non-missing values (59% and 41% coverage respectively). Standard errors clustered by firm. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$. *Source:* Authors' elaboration based on Bureau van Dijk (Moody's Analytics) (2024).

E PSM Estimates and Parallel-Trends Bounds

Table 16 reports the full PSM ATTs underlying the comparison in Section 5. Table 17 reports HonestDiD relative-magnitude sensitivity bounds (Rambachan and Roth, 2023): the share-premium ATT remains informative under moderate parallel-trends violations, whereas the debt-side effects, significant under exact parallel trends and the joint pre-trend test, do not survive the relative-magnitude bounds—a limit we report openly.

Table 16: CS2021 ATT — PSM Robustness (Main Outcomes)

Outcome	**ATT**	**SE**	**p-value**	**Pre-trend p**	**Sig**
Panel A: Debt outcomes (H_{null})					
Total leverage (D/TA)	0.0210	0.0078	0.007	0.072	***
Short-term leverage (STD/TA)	-0.0042	0.0078	0.585	0.024	
Long-term leverage (LTD/TA)	0.0249	0.0069	0.000	0.572	***
Panel B: Equity falsification ($H_{debt\text{-}spec}$)					
Share premium / TA	0.0145	0.0017	0.000	0.064	***
Shareholders funds / TA	-0.0204	0.0081	0.012	0.365	**

Notes: Callaway and Sant'Anna (2021) doubly-robust estimator. Control group: never-treated firms from PSM-matched panel (2382 treated firms; 7956 matched controls). Anticipation window: $\kappa = 2$. $B = 999$ bootstrap replications (`set.seed = 2026`). PSM 1:5 matching weights applied. Pre-trend: approximate joint Wald χ^2 on $e \leq -3$. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. *Source:* Authors' elaboration based on Bureau van Dijk (Moody's Analytics) (2024).

Table 17: HonestDiD Sensitivity Analysis: Robust CIs under Parallel-Trends Violations (Rambachan & Roth, 2023)

Panel A — Relative-magnitude restriction $\Delta^{\text{RM}}(\bar{M})$: violation $\leq \bar{M}\times$ max pre-trend						
Outcome	**ATT**	**SE**	**Robust 95% CI under $\bar{M}$**			
			$\bar{M} = 0$	$\bar{M} = 0.5$	$\bar{M} = 1$	$\bar{M} = 2$
Total leverage (D/TA)	-0.0199	0.0057	—	[-0.062, 0.019]	[-0.101, 0.057]	[-0.122, 0.122]
Short-term leverage (STD/TA)	-0.0149	0.0052	—	[-0.073, 0.045]	[-0.119, 0.101]	[-0.119, 0.119]
Long-term leverage (LTD/TA)	-0.0056	0.0042	—	[-0.075, 0.058]	[-0.103, 0.103]	[-0.103, 0.103]
Share premium / TA	0.0118	0.0013	—	**[0.003, 0.021]**	[-0.005, 0.027]	[-0.021, 0.027]
Panel B — Smoothness restriction $\Delta^{\text{SD}}(M)$: second-difference bound (M = fraction of SE)						
Outcome	**ATT**	**SE**	**Robust 95% CI under M**			
			$M = 0.5\,\text{SE}$	$M = 1.0\,\text{SE}$	$M = 1.5\,\text{SE}$	$M = 2.0\,\text{SE}$
Total leverage (D/TA)	-0.0199	0.0057	[-0.062, 0.030]	[-0.089, 0.064]	[-0.116, 0.095]	[-0.142, 0.126]
Short-term leverage (STD/TA)	-0.0149	0.0052	[-0.058, 0.027]	[-0.088, 0.054]	[-0.117, 0.080]	[-0.146, 0.104]
Long-term leverage (LTD/TA)	-0.0056	0.0042	[-0.038, 0.032]	[-0.058, 0.059]	[-0.078, 0.083]	[-0.097, 0.109]
Share premium / TA	0.0118	0.0013	**[0.009, 0.030]**	**[0.003, 0.036]**	[-0.003, 0.042]	[-0.009, 0.048]

Notes: Sensitivity analysis following Rambachan and Roth (2023). ATT and SE are the overall (cohort-averaged) estimates from the Callaway and Sant'Anna (2021) simple aggregation. Because the relative-magnitude test requires at least two contiguous, genuine pre-periods to bound the post-treatment violation, the sensitivity event study is re-estimated with no anticipation ($\kappa = 0$); the resulting overall ATT therefore differs slightly from the headline $\kappa = 2$ estimate in Table 4 and is the appropriate baseline for this test. *Panel A* reports robust 95% CIs under the relative-magnitude restriction $\Delta^{\text{RM}}(\bar{M})$: post-treatment violations bounded by $\bar{M}$ times the largest observed pre-trend. *Panel B* reports robust 95% CIs under the smoothness restriction $\Delta^{\text{SD}}(M)$: second differences bounded by M (multiples of the overall SE). $\bar{M} = 0$ (Panel A) and $M = 0$ (Panel B) correspond to exact parallel trends. Bold intervals exclude zero. All specifications: Callaway and Sant'Anna (2021) doubly-robust estimator, never-treated controls, primary (CEM) matching weights, $B = 999$ bootstrap (`set.seed = 2026`). *Source:* Authors' elaboration based on Bureau van Dijk (Moody's Analytics) (2024).

F Robustness to the ln employees covariate

Table 18 re-estimates the primary CEM design without ln employees—dropping it from both the match and the doubly-robust adjustment—against the seven-covariate primary, on the same pool and seed. Both columns are re-estimated within this comparison, so the seven-covariate point estimates coincide with Table 4 while the bootstrap standard errors differ marginally. The equity results are unchanged and the debt results are, if anything, slightly stronger; the gain in estimable treated firms is modest and the share-premium parallel-trends test deteriorates, which is why the seven-covariate specification is retained (Section 3.2).

Table 18: Primary design robustness: CEM with vs. without ln employees

	With ln emp. (7 cov.)		Without ln emp. (6 cov.)	
Outcome	ATT	p	ATT	p
Total leverage (D/TA)	-0.0223***	0.006	-0.0295***	0.000
Short-term leverage (STD/TA)	-0.0324***	0.000	-0.0394***	0.000
Long-term leverage (LTD/TA)	+0.0091	0.202	+0.0086	0.263
Share premium / TA	+0.0148***	0.000	+0.0152***	0.000
Shareholders' funds / TA	+0.0166**	0.046	+0.0217**	0.010
Return on assets	-0.0114***	0.004	-0.0079**	0.036

Notes: CS2021 doubly-robust ATT (never-treated controls, $\kappa = 2$, $B = 999$). Both columns use coarsened exact matching on the same pool; the right panel drops ln employees from both the match and the DR adjustment. Estimable treated (first treat ≤ 2024) 1,930 with ln emp. and 2,032 without. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

G Reverse causality: excluding pre-registration equity raisers

A natural threat to the equity result is that firms certify precisely when raising capital, to capture the 30% investor tax relief, so that the share-premium response reflects timing rather than treatment. Table 19 addresses it directly: we flag the 444 certified firms whose paid-in capital (capital stock or share-premium reserve) rose between $e = -2$ and $e = -1$—an equity round in the two years before registration—and re-estimate the headline CS2021 ATT on the 1,864 firms that did not. The share-premium effect remains positive and highly significant (+1.13 pp, $p < 0.001$) among these non-raisers, so it is not a mechanical product of contemporaneous capital increases; the effect is larger (+3.39 pp) among the raisers, as expected for firms in the midst of a recapitalisation, but its survival among non-raisers is what rules out a purely selection-driven reading.

Table 19: Reverse-causality check: ATT excluding pre-registration equity raisers

Sample	Outcome	ATT	SE	p
Non-raisers	Share premium / TA	+0.0113***	0.0018	0.000
Non-raisers	Shareholders' funds / TA	+0.0051	0.0085	0.547
Non-raisers	Total leverage (D/TA)	-0.0127	0.0082	0.123
Raisers	Share premium / TA	+0.0339***	0.0054	0.000
Raisers	Shareholders' funds / TA	+0.0808***	0.0228	0.000

Notes: CS2021 doubly-robust ATT (never-treated controls, $\kappa = 2$, $B = 999$, CEM weights) re-estimated after dropping the 444 treated firms whose paid-in capital (capital stock or share-premium reserve) rose between $e = -2$ and $e = -1$, i.e. those with an equity round in the two years before registration; 1,864 treated firms are retained. Controls are unchanged. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. *Source:* Authors' elaboration based on Bureau van Dijk (Moody's Analytics) (2024).

H Debt estimates without the lagged outcome in the doubly-robust adjustment

The primary specification includes pre-treatment total leverage in both the coarsened exact match and the doubly-robust covariate vector. Because the staggered estimator also differences out the base period, this conditions on the lagged outcome twice, which could in principle pull the debt counterfactual toward the treated pre-path. Table 20 re-estimates the three debt ATTs with pre-treatment leverage retained in the match but removed from the doubly-robust adjustment. The point estimates are essentially unchanged (total leverage −2.38 pp versus −2.23 pp; short-term −3.36 pp versus −3.24 pp), so the contraction is not an artefact of double-conditioning, and the cross-design fragility documented in Section 5 reflects overlap and reweighting rather than this specification choice.

Table 20: Debt ATT excluding the lagged outcome from the doubly-robust adjustment

Outcome	ATT	SE	p
Total leverage (D/TA)	-0.0238**	0.0098	0.015
Short-term leverage (STD/TA)	-0.0336***	0.0088	0.000
Long-term leverage (LTD/TA)	+0.0089	0.0070	0.201

Notes: CS2021 doubly-robust ATT (never-treated controls, $\kappa = 2$, $B = 999$, CEM weights), with pre-treatment total leverage retained in the coarsened exact match but *removed* from the doubly-robust covariate vector, so the lagged outcome is not used twice. Compare with Table 4. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. *Source:* Authors' elaboration based on Bureau van Dijk (Moody's Analytics) (2024).

I Decomposing the return-on-assets decline

The headline return on assets falls after certification, but ROA divides profit by total assets, so the decline can reflect a larger asset base, lower operating profitability, or both. Table 21 separates the two by estimating the ATT on ln total assets (the denominator) and on a denominator-free margin, profit over revenues (the numerator), with the headline CS2021 specification. Total assets rise steeply (+43.4 log points, $p < 0.001$)—the asset-base expansion of firms that draw in the new equity and scale up—while the operating margin compresses by -2.67 pp ($p = 0.022$). The ROA decline therefore combines a mechanical denominator effect from the recapitalisation with a genuine, if modest, fall in operating profitability, rather than reflecting the small causal increase in documented R&D intensity.

Table 21: Decomposing the return-on-assets decline: denominator versus numerator

Outcome	ATT	SE	p
ln total assets (denominator)	+0.4341***	0.0342	0.000
Profit / revenues (margin)	-0.0267**	0.0116	0.022

Notes: CS2021 doubly-robust ATT (never-treated controls, $\kappa = 2$, $B = 999$, CEM weights). "ln total assets" is the log of book total assets (the ROA denominator); "Profit / revenues" is net income over turnover (a denominator-free profitability margin), winsorised at 1%. A positive asset response alongside a flat or smaller margin response indicates that the headline ROA decline (Section 4.4) is driven substantially by the asset-base expansion from the equity recapitalisation rather than by a fall in operating profitability. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. *Source:* Authors' elaboration based on Bureau van Dijk (Moody's Analytics) (2024).

J Origin of the equity response: paid-in premium versus nominal capital

Whether the equity response reflects new outside capital or existing owners capitalising the firm matters for interpretation. Share premium is, by accounting definition, capital paid *above* nominal par value—the signature of valuation-based subscription, in which shares are priced above book—whereas founders injecting par-value capital would move nominal capital stock. Table 22 contrasts the two ATTs: the response is concentrated in the premium (1.48 pp, $p < 0.001$) while nominal capital stock is flat and insignificant (+0.30 pp, $p = 0.544$). The new equity therefore enters as premium-priced subscription consistent with outside investors paying for a now-credible, externally legible signal, rather than as par-value top-ups by incumbent owners.

Table 22: Origin of the equity response: paid-in premium versus nominal capital

Outcome	ATT	SE	p
Share premium / TA (paid above par)	+0.0148***	0.0016	0.000
Nominal capital stock / TA	+0.0030	0.0049	0.544

Notes: CS2021 doubly-robust ATT (never-treated controls, $\kappa = 2$, $B = 999$, CEM weights). Share premium is paid-in capital recorded *above* nominal par value; capital stock is nominal subscribed capital. A response concentrated in the premium rather than in nominal capital indicates valuation-based subscription (shares issued above par) rather than founders injecting par-value capital. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. *Source:* Authors' elaboration based on Bureau van Dijk (Moody's Analytics) (2024).