Leverage Without Additionality: What Co-Funding Ratios Identify in Public Technology Programmes

Stéfane Fermigier

Draft working paper v0.2, 2026-08-04. Four propositions, each with a proof and a numeric cross-check against brute force; the ranking claim carries a Monte-Carlo report. The application is an accounting exercise on published programme data; no causal effect is estimated and none is claimed.

Abstract

European technology programmes report a "leverage factor": the co-funding declared by grant beneficiaries per euro of public contribution. The statistic is presented as evidence that public money mobilizes private money.

We show what it identifies. Under a published co-funding schedule, the ratio is a deterministic function of that schedule, so a programme funding half of each project reports a leverage of exactly one however its beneficiaries behave. The reported ratio bounds the crowding-in ratio from above, with equality only under the assumption that no beneficiary would have invested anything absent the grant, which reverses the direction of the caveat funders usually attach to it. Compliance with the schedule cannot produce a ratio above a computable ceiling, so only observations above that ceiling carry behavioural information. And because declared co-funding pools public and private sources, a rationale resting on private capital-market frictions is addressed only by the private component, which must be reported separately or the statistic is silent on the stated market failure.

We apply the results to the Digital Europe Programme's published figures. Its headline leverage of 0.79 lies below the ceiling its own schedule implies, so it identifies the share of the envelope disbursed at full funding rates (0.21) and nothing about behaviour; one technology area of thirteen exceeds the ceiling. The private component is a minority of reported leverage in nine of the twelve areas where both are published, standing at 9.1% in artificial intelligence and 0.0% in advanced connectivity. In Monte-Carlo, ordering programmes by the reported ratio inverts the ordering by true additionality on 25.2% of sampled pairs, and programmes sharing one schedule report an identical ratio while true additionality spans the whole admissible range.

Keywords: additionality, crowding in, state aid, subsidy design, programme evaluation, technology policy.
JEL: H50, O38, D04, L52.

1. Introduction

Public technology programmes are asked to show that their money moves private money. The instrument they reach for is a ratio: total co-funding declared by beneficiaries divided by the public contribution. For the European Union's flagship digital deployment programme the statistic is reported at 79% in the study assembled to inform the 2028-2034 multiannual financial framework, whose text describes the co-funding as what public money "managed to generate" (European Commission 2026).

This paper asks what such a ratio identifies. The answer is narrow and can be stated exactly. Grant programmes publish co-funding rates: the share of eligible project cost the funder pays. A beneficiary that complies with a 50% rate contributes one euro for every euro received, and the programme records a leverage of one. The number is a restatement of the rate schedule. It is compatible with complete additionality, with complete non-additionality, and with everything between.

Four results follow. Proposition 1 gives the schedule identity and its corollary: under a single rate, the ratio is pinned regardless of behaviour. Proposition 2 shows the ratio is an upper bound on the crowding-in ratio, attained only when counterfactual private investment is zero. This inverts the caveat funders attach to the statistic, which is that exclusions make it an understatement. Proposition 3 characterizes the region in which the ratio carries information: compliance cannot push it above the largest value the schedule admits, so observations at or below that ceiling are uninformative about behaviour. Proposition 4 decomposes the ratio into public and private components and locates which of them a market-failure rationale requires.

Section 6 applies the results to the Digital Europe Programme. The exercise is arithmetic on published tables. It establishes what the programme's own numbers can and cannot support, and it identifies the small set of cells that do carry behavioural content.

The contribution is a measurement result. Nothing here estimates the programme's additionality, which would require a counterfactual the published data do not contain. The claim is that the reported statistic does not estimate it either. The gap between the two is a design choice.

2. Related literature

The econometrics of subsidy additionality is settled in outline and unsettled in magnitude. David, Hall and Toole (2000) survey the evidence on whether public research and development spending complements or substitutes for private spending and find the sign varies with the level of aggregation and the identification strategy. Klette, Møen and Griliches (2000) set out what microeconometric evaluation of such programmes requires, and Lach (2002) and Görg and Strobl (2007) supply firm-level estimates under explicit selection corrections. The common feature is a counterfactual: additionality is a difference between observed investment and investment that would have occurred absent the grant, and recovering it needs either an experiment, a discontinuity, or a defensible selection model.

The co-funding ratio studied here belongs to a different family. It is an accounting statistic computed from grant applications and it appears in programme reporting. The state-aid literature supplies the normative frame the funders themselves invoke: Haucap and Schwalbe (2011) derive the appropriability logic under which capability-building actions warrant high funding rates and deployment actions lower ones, a logic the source programme cites when setting its schedule. That same connection makes the ratio uninformative, since the schedule is chosen on appropriability grounds and the ratio then reports the schedule back.

Audit bodies have criticized leverage and multiplier reporting in European financial instruments on definitional grounds (European Court of Auditors, citation to be verified). This paper's contribution is to state the identification content formally and to separate the two distinct failures: a ratio that restates a schedule, and a ratio that pools public with private co-funding.

3. The statistic, in the record

The figures analysed here appear in a study commissioned by the Commission's Directorate-General for Communications Networks, Content and Technology and published by the Publications Office of the European Union. Its legal notice records that the document reflects the views of its authors and should not be considered representative of the Commission's official position. What follows therefore treats the statistic as a published reporting practice on which programming decisions draw.

The leverage factor for the Digital Europe Programme is reported as follows (European Commission 2026): "The table below presents the direct leverage of the Digital Europe Programme, based on currently granted projects (excluding procurements, financial instruments, contribution agreements), presenting the level of other public and private invest-

ments that the EU public co-funding managed to generate. Direct leverage is measured as the amount of funding committed by beneficiaries as part of the project financial application and reporting procedures."

Three features of this definition matter. The measurement is the beneficiary's own declared commitment at application. The population is other public *and* private investment pooled. And the causal verb ("managed to generate") is attached to a quantity computed without any counterfactual.

The same source attaches a caveat in one direction only: "It is important to note that this does not include leverage of investments at end-user level, as DIGITAL is set up mostly as an indirect delivery programme, with beneficiaries supporting private companies (in particular SMEs) and public authorities. End users also invest (in kind or in cash) in digital transformation as a response to the original EU investment." The exclusions are read as making the figure a lower bound. Proposition 2 shows the construction supports the opposite bound.

4. What a co-funding ratio identifies

4.1 Setup

A programme awards grants indexed by i. Project i has eligible cost B_i, public contribution G_i, and beneficiary-declared co-funding $C_i = B_i - G_i$. The programme's co-funding rate for the project is $s_i = G_i/B_i \in (0, 1]$, fixed by the action type before any behaviour is observed. The reported leverage factor is

$$L = \frac{\sum_i C_i}{\sum_i G_i}.$$

Let $I_i^0 \geq 0$ denote the investment the beneficiary would have made in the same activity absent the grant. Additionality is $A = \sum_i (C_i - I_i^0)$ and the crowding-in ratio is $\rho = A/\sum_i G_i$. A programme is said to crowd in private money when $\rho > 0$.

Call a beneficiary *compliant* when it declares exactly the co-funding the schedule requires, $C_i = G_i(1 - s_i)/s_i$, and *over-contributing* when it declares more.

4.2 Results

Proposition 1 (schedule identity). Let the schedule use rates $s_1, \ldots, s_J$ and let ω_j be the share of total public contribution disbursed at rate s_j, with $\sum_j \omega_j = 1$. If every beneficiary is compliant, then

$$L = \sum_{j=1}^{J} \omega_j \frac{1 - s_j}{s_j}.$$

In particular, under a single rate s the ratio is $L = (1 - s)/s$ identically, so a programme funding half of each project reports $L = 1$ whatever its beneficiaries do.

Proposition 2 (additionality bound). For any programme,

$$\rho = L - \frac{\sum_i I_i^0}{\sum_i G_i} \leq L,$$

with equality if and only if $I_i^0 = 0$ for every i. The reported leverage factor is therefore an upper bound on the crowding-in ratio, and it equals that ratio only under the assumption that no beneficiary would have invested anything in the absence of the grant.

Proposition 3 (informative region). Define the schedule ceiling $\bar{L} = \max_j (1 - s_j)/s_j$. Under compliance, $L \in [\min_j (1 - s_j)/s_j, \bar{L}]$, and every value in that interval is attainable by some allocation ω. Hence an observation $L \leq \bar{L}$ is consistent with $\rho = 0$ and carries no behavioural information; only $L > \bar{L}$ requires over-contribution by some beneficiary.

Proposition 4 (public/private decomposition). Write $C_i = C_i^{\text{pub}} + C_i^{\text{priv}}$ and correspondingly $L = L^{\text{pub}} + L^{\text{priv}}$. If the programme's stated rationale is a friction in private capital

markets, the quantity responsive to that rationale is L^{priv}. A reported L that pools the two is uninformative about the stated failure whenever L^{pub} is an appreciable share of it.

Proofs are in Appendix A.

4.3 Reading

Proposition 1 is an identity. The reported statistic varies across programmes and technology areas for a reason that has nothing to do with beneficiary behaviour. Two programmes with identical additionality report different leverage when they fund different action types, and a programme can raise its reported leverage by shifting the mix toward lower funding rates while changing no one's investment.

Proposition 2 fixes the direction of the bias. Funders present exclusions from the numerator as reasons the figure understates true leverage. The construction implies the reverse: the numerator counts all declared co-funding, including whatever the beneficiary would have spent anyway, so the ratio overstates crowding in by exactly the counterfactual investment per euro of grant. That quantity is unobserved and weakly positive.

Proposition 3 says where to look. Values inside the schedule's range are accounting; values above the ceiling are behaviour, because no compliant portfolio can reach them. This turns the reported table from a performance measure into a screen: the exceptional cells are the ones to examine.

Proposition 4 addresses a substitution that occurs in programme reporting. Public-to-public cost sharing between a union and its member states is a real fiscal fact. Where the stated market failure concerns private financing, only the private component speaks to it.

5. Verification

Every proposition carries a numeric cross-check against brute-force computation, and the ranking claim of §6.3 carries a Monte-Carlo report. The suite is `verify/leverage.py`; results are written to `verify/results/leverage_summary.json`.

`check_L1` builds explicit project portfolios at each schedule rate, sums declared co-funding and public contribution, and compares the realized ratio to the closed form over 2,000 random allocations of the simplex. `check_L2` draws counterfactual investment levels and verifies both the identity and the inequality over 20,000 draws. `check_L3` searches 200,000 random compliant allocations for a portfolio exceeding the ceiling and finds none, confirming the ceiling of 1.00 under the observed two-rate schedule. `check_L4` reconstructs the public residual from the published totals and private components and confirms the decomposition.

6. Application: the Digital Europe Programme

6.1 What the headline identifies

The programme reports $L = 0.79$ across granted projects. Its schedule combines a 50% simple grant with action types funded at 100% (European Commission 2026), so $\bar{L} = 1.00$. The observed value sits below the ceiling, and by Proposition 3 it is consistent with zero crowding in.

It identifies the mix. Under compliance with a two-rate schedule, $L = 1 - \pi$ where π is the share of public contribution disbursed at the full rate. The reported 0.79 therefore implies $\pi = 0.21$: about a fifth of the granted envelope went to fully funded actions. That is a useful accounting fact about programme composition, and it is the only quantity the headline recovers.

Coverage compounds the point. The figure covers granted projects only, excluding procurements, financial instruments and contribution agreements, which corresponds to roughly a fifth of the programme's total envelope.

6.2 The cells that carry information

Of the thirteen technology areas with published totals, one exceeds the ceiling: robotics, at 1.08. Within-area breakdowns give three further cells above 100% (robotics deployment at 128% and robotics ecosystem-building at 139%, cybersecurity digital infrastructure at 119%, blockchain deployment at 101%). By Proposition 3 these are the observations that require over-contribution and therefore carry behavioural content. Every other published cell is consistent with compliance.

6.3 The private component

The programme publishes a second series covering private co-funding alone. Applying Proposition 4, the private share of reported leverage is:

Technology area	**Reported** L	**Private** L^{priv}	**Private share**
Advanced connectivity	0.80	0.00	0.0%
Artificial intelligence	0.77	0.07	9.1%
Robotics	1.08	0.16	14.8%
High-performance computing	0.53	0.10	18.9%
Microelectronics	0.99	0.21	21.2%
Quantum	0.90	0.23	25.6%
Data	0.72	0.19	26.4%
Interoperability	0.90	0.29	32.2%
Cybersecurity	0.81	0.32	39.5%
Blockchain	0.85	0.50	58.8%
Cloud-Edge-IoT	0.84	0.58	69.0%
Next-generation internet and XR	0.43	0.40	93.0%

The private component is a minority of reported leverage in nine of the twelve areas. In artificial intelligence, 70 of the 77 points of reported leverage are other public money; in advanced connectivity, all 80 are. Where the programme's rationale rests on private financing frictions, these are the numbers that address it.

Two cautions apply to this table. The two series are published for different windows, the totals to September 2025 and the private component to 2024, so the shares are indicative. And photonics is recorded as unavailable in the private series and is omitted.

6.4 Does the statistic rank programmes correctly?

Even granting that the ratio is not a level estimate, it might still order programmes by additionality. It does not: in the Monte-Carlo (`mc_rank`, seed 20260718, 20,000 sampled pairs), programmes draw a rate schedule and a counterfactual-investment intensity; ordering by the reported ratio inverts the ordering by true crowding in on 25.2% of pairs. The degenerate case is sharper: programmes sharing a single schedule report an identical ratio by construction while their true crowding-in ratios span the entire admissible interval.

6.5 The class where the statistic is inapplicable by design

The degenerate case runs in the wild, at the corner of the schedule. One class of public technology spending funds the maintenance of open-source commons: infrastructure whose licence prevents its maintainer from appropriating the value it creates, and whose public funders screen applicants for the absence of alternative financing. Germany's Sovereign Tech Agency is the reference programme. It funds critical open-source components through service contracts at full funding rates, and it publishes per-technology figures (EUR 405,888 for OpenSSL, EUR 596,160 for Log4j, EUR 1,056,673 for PyPI, EUR 37.3 million across the portfolio from October 2022 to July 2026; rough, with these figures and the selection-criteria description taken from the agency's published pages, citation to be verified) [Sovereign Tech Agency 2026]. By the mechanics of Proposition 1 a full-funding-rate programme reports a

leverage of exactly zero, so the statistic ranks this programme last in any comparison set. It does so for the very feature that makes the programme's additionality plausibly the highest in the set: selection conditions on the absence of counterfactual funding. Private provision of the class exists and is documented thin (the post-Heartbleed Core Infrastructure Initiative and its successor, the OpenSSF, are consortium funding of exactly the OpenSSL class). Additionality here is therefore high by selection design, a property the co-funding ratio cannot see from either end. For this class the informative statistic is coverage: the fraction of the critical substrate under funded maintenance, the bus factor of each funded component, and funding per unit of criticality. The class also closes the loop with the companion work: maintained commons are the option-maintenance public good of the companion theory paper (eurostack-path-sovereignty/papers/theory-wp.md, §6.2), and the agency is the state variant of the demand-pooling club analyzed in the companion collective-provision paper (oss-adoption/papers/collective-provision-wp.md).

7. What would have to be measured instead

The results imply a short reporting specification. Publish the rate schedule and the distribution of public contribution across rates, which makes the mechanical component of any reported ratio computable by the reader. Publish the private component separately from the public one. Flag observations above the schedule ceiling, since those are the ones carrying behavioural information. And keep the causal vocabulary for quantities estimated against a counterfactual. For the commons-maintenance class of §6.5 the specification differs: publish coverage of the critical substrate, per-component bus factors, and funding per unit of criticality, since a co-funding ratio is uninformative at the full-funding corner where that class correctly operates.

Recovering additionality itself requires a design the current reporting does not contain: variation in funding rates or in award thresholds that is plausibly unrelated to beneficiary investment intentions. Programmes that score and rank applications already generate such variation at the funding cutoff, and the application records that produce the co-funding statistic contain the ingredients for a discontinuity design. That requires a research programme; the evaluation literature cited in §2 has pursued it in other settings.

A separate consideration concerns instrument comparability. A guarantee instrument reports total investment mobilized per euro of budget guarantee, a grant reports third-party contribution per euro disbursed. The two differ in numerator, denominator, and timing, and one of them is typically an ex ante target while the other is partly realized. Ranking instruments by these figures compares different objects. Putting them on a common footing requires expressing both as expected additional investment per euro of expected fiscal cost, which for a guarantee means pricing the contingent liability.

8. Discussion and limits

The analysis takes the funder's own definition of additionality, which counts co-funding declared on the supported project. A broader definition would count investment induced elsewhere, including at end users, and the source programme is right that its reporting excludes that channel. Proposition 2 is unaffected: whatever population the numerator covers, it includes counterfactual investment within that population, and the bound holds on the corresponding ratio.

Compliance is treated as exact. In practice beneficiaries may over-declare co-funding at application without realizing it, which would push reported ratios above the ceiling for reasons unrelated to investment. The cells identified in §6.2 should therefore be read as candidates for behavioural content, pending reconciliation of application figures with final reporting.

The application is arithmetic on published aggregates, and it inherits their limits. The two series differ in window and coverage, the totals are provisional while implementation continues, and the schedule is reconstructed from the programme's stated funding rates. Proposition 1 requires programmes to publish the distribution; its absence is part of the finding.

Nothing here bears on whether the programme is worthwhile. A programme with zero crowding in can be justified on other grounds, including the appropriability logic its own design invokes. The claim is about what one statistic supports.

9. Conclusion

A co-funding ratio computed from grant applications restates the funder's rate schedule. It bounds crowding in from above, it is uninformative about behaviour at any value the schedule can produce, and it conflates public cost sharing with private mobilization. On the Digital Europe Programme's published figures, the headline leverage of 0.79 identifies a programme composition (about a fifth of granted funds at full funding rates) and nothing about behaviour; one technology area of thirteen carries behavioural content; and the private component, which a private-financing rationale requires, is a minority of reported leverage in nine of twelve areas and near zero in two of the largest. The reporting changes that would fix this are small.

Appendix A. Proofs

Proposition 1. For a compliant project at rate s_i, $C_i = B_i - G_i = G_i/s_i - G_i = G_i(1 - s_i)/s_i$. Summing within the group of projects funded at rate s_j, whose total public contribution is $\omega_j \sum_i G_i$, gives group co-funding $\omega_j(1 - s_j)/s_j \sum_i G_i$. Summing across groups and dividing by $\sum_i G_i$ yields the stated expression. With $J = 1$ the weight is unity and $L = (1 - s)/s$, which at $s = 1/2$ equals 1. □

Proposition 2. By definition $\rho = \left(\sum_i C_i - \sum_i I_i^0\right) / \sum_i G_i = L - \sum_i I_i^0 / \sum_i G_i$. Since $I_i^0 \geq 0$ for every i and $\sum_i G_i > 0$, the subtracted term is weakly positive, so $\rho \leq L$. It vanishes if and only if every I_i^0 is zero. □

Proposition 3. By Proposition 1, L is a convex combination of the values $(1 - s_j)/s_j$ with weights ω_j, hence lies in the interval spanned by them and attains any interior value at the appropriate ω. An observation in that interval is thus reproduced by compliance alone, which is consistent with $I_i^0 = C_i$ for every i and therefore with $\rho = 0$ by Proposition 2. Exceeding $\bar{L}$ requires some $C_i > G_i(1 - s_i)/s_i$, which is over-contribution. □

Proposition 4. Additivity of the decomposition is immediate from $C_i = C_i^{\text{pub}} + C_i^{\text{priv}}$ and linearity of the sum. The second clause follows by noting that a friction in private capital markets constrains C_i^{priv} and leaves C_i^{pub} determined by intergovernmental allocation, so variation in L driven by L^{pub} is uninformative about that friction. □

References

- David, P. A., Hall, B. H. & Toole, A. A. (2000). Is public R&D a complement or substitute for private R&D? A review of the econometric evidence. *Research Policy* 29(4–5):497–529.
- European Commission, Directorate-General for Communications Networks, Content and Technology (2026). *Study on the EU's critical digital capacities deployment beyond 2027: final report.* Prepared by Technopolis Group; K. Izsak, P. Simmonds, M. Ploeg, S. Viscido, A. Lotito, C. Moreno & R. Alasghar (eds.). Luxembourg: Publications Office of the European Union. ISBN 978-92-68-40079-1, doi:10.2759/0673677.
- European Court of Auditors (2016). *Implementing the EU budget through financial instruments: lessons to be learnt from the 2007-2013 programme period.* Special Report 19/2016. (Citation to be verified.)

- Görg, H. & Strobl, E. (2007). The effect of R&D subsidies on private R&D. *Economica* 74(294):215–234.
- Haucap, J. & Schwalbe, U. (2011). *Economic principles of state aid control.* DICE Discussion Paper No. 17, Düsseldorf Institute for Competition Economics.
- Klette, T. J., Møen, J. & Griliches, Z. (2000). Do subsidies to commercial R&D reduce market failures? Microeconometric evaluation studies. *Research Policy* 29(4–5):471–495.
- Lach, S. (2002). Do R&D subsidies stimulate or displace private R&D? Evidence from Israel. *Journal of Industrial Economics* 50(4):369–390.
- Sovereign Tech Agency (2026). Published per-technology investment figures, portfolio October 2022 to July 2026. https://www.sovereign.tech. (Citation to be verified: figures taken from the agency's project pages via the 2026-08-04 read-through; the site could not be re-fetched at edit time.)

Annex A. Working notes (temporary annex)

This annex lists internal working documents and will be removed from the final version. Paths are relative to the project root.

- `notes/leverage/source-tables.md`: the published co-funding tables transcribed verbatim, with the extraction anchors.
- `papers/notes/leverage-without-additionality-wp.md`: internal notes for this paper.

Annex B. Reproduction: scripts, data, and figures

One command reproduces every in-text number: `python verify/leverage.py` (seconds; python 3, numpy; deterministic, no network).

- `verify/leverage.py`: the published data (`L_TOTAL`, `L_PRIVATE`, `RATES`), the closed forms, gate-2 checks `check_L1`…`check_L4` (seeds 0–2), and gate-3 `mc_rank` (seed 20260718).
- `verify/results/leverage_summary.json`: the schedule ceiling, the implied full-rate share, the areas above the ceiling, the private-share table of §6.3, and the Monte-Carlo configuration and outcome.

No figures are used in this paper.