Public Capital, Private Selection: Additionality in European Venture Capital

[Preliminary Version: Please Do Not Circulate]

Sevket Camoglu*

July 31, 2026

Abstract. Government venture capital (GovVC) is a leading instrument of European industrial policy, yet a first-order question remains open: does it finance firms that would otherwise go unfunded, or does it merely substitute for private capital? Because venture capital is a two-sided matching market, the firms a public investor backs reflect both its own preferences and its competition with private investors—two forces that reduced-form designs cannot separate. I estimate a two-sided matching model of European venture capital (1989–2024) and re-solve the equilibrium with government investors removed. Only 24.9% of GovVC-backed firms would have gone unfunded in their absence; the remaining three quarters are re-absorbed by private capital. Recovered preferences reverse the reduced-form targeting evidence: GovVC favors high-tech firms but loses them to private competition, and later-stage ventures over early stages. The *additional* firms exit less often than identical private-backed firms (4.88 pp.), but the shortfall vanishes once the geographic mandates are relaxed: the apparent quality gap partly reflects single-market fragmentation, not adverse selection. Around the round, these firms show higher *employment* but no change in *profitability*. Together, these results suggest an employment orientation rather than the selection of high-exit frontier firms.

JEL: G24 · G28 · O38 · L26.
Keywords: Government Venture Capital · Additionality · Two-sided Matching · Industrial Policy · Entrepreneurial Finance.

* I am indebted to Stefano Colonnello for his guidance, support, and encouragement throughout this project. I also thank Loriana Pelizzon, Matthias Efing, and Johan Hombert for their valuable comments and feedback. Finally, I am grateful to the participants of the S&O Research Day at HEC Paris and the FMA Europe 2026 conference for their helpful comments and suggestions. Address: Ca' Foscari University of Venice, Cannaregio 873, Fondamenta San Giobbe, 30121 Venezia, Italy. E-mail: sevket.camoglu@unive.it.

1. Introduction

No EU company founded in the last fifty years has reached a market capitalization of €100 billion, while six U.S. firms founded over the same period exceed €1 trillion (Draghi 2024). This divergence has put the financing of European innovation at the center of the policy debate. The shortfall is usually traced to three margins: a supply-side margin—whether Europe lacks a strong entrepreneurial pipeline (Aghion and Howitt 1990; Michelacci and Schivardi 2020); an institutional margin—fragmented legal and tax frameworks that impede cross-border scaling (Draghi 2024; Enriques et al. 2025, 2026); and a financial margin—whether risk capital is deployed efficiently, especially at early stages (Da Rin et al. 2006).

This paper studies the financial margin, focusing on government-backed venture capital (GovVC). GovVC is a primary instrument of national and EU innovation policy and supplies a substantial share of early-stage capital in Europe (Figure A1); in my sample it participates in roughly one investment in four (Berger et al. 2024). Because these funds are publicly financed, the social return to the intervention turns on a single question: *does GovVC create financing that the private market would not have supplied, or does it substitute for capital that would have arrived anyway?*

This question is hard to answer with reduced-form tools, and the difficulty is structural rather than statistical. Venture capital is a matching market: the observed composition of GovVC's portfolio is the joint product of the public investor's *preferences* and the *competition* it faces from private capital (Cumming and Johan 2019). A negative cross-sectional correlation between GovVC participation and, say, high-technology status can arise either because public investors dislike high-tech firms or because private investors outcompete them for the high-tech segment, leaving GovVC the residual. No regression of outcomes on a GovVC indicator can tell these apart, yet the policy question—what happens when GovVC is withdrawn?—depends entirely on which is true.

I therefore estimate a two-sided matching model in the tradition of Sørensen (2007) that recovers, for every feasible firm–VC pair, the latent value of the match. The model is identified from the revealed-preference inequalities that a stable matching implies, and it separates

each investor type's preferences from the competitive equilibrium in which those preferences are exercised. With the latent match values in hand, I measure additionality directly: I re-solve the equilibrium with public investors removed and partition GovVC-backed firms into those that find a private match (substitution) and those left unfunded (additionality). I run the same machinery on a parallel U.S. sample as an external benchmark.

The paper delivers four results.

First, substitution dominates additionality. Re-solving the European equilibrium without GovVC, 24.9% of its portfolio firms go unfunded, while the remaining 75.1% are re-absorbed by private investors operating within their observed capacity. The parallel U.S. additionality rate is lower still, 13.5%. In both economies, then, the large majority of public-VC activity backs firms that the private market would have financed.

Second, the structural preferences reverse the sign of the reduced-form coefficients on the two characteristics that matter most for policy. In the cross-section, GovVC participation is 5.6 percentage points *less* likely in high-tech firms and 1.5 points *more* likely in early-stage firms (Table A1). The structural preferences point the other way: at the average firm, switching to high-tech *raises* the probability of a GovVC match by 6.8 points, while switching to early-stage *lowers* it by 2.8 points (Table 2). The reversal is exactly the preference-versus-competition confound the model is built to resolve: private VC's appetite for high-tech is so much stronger that it absorbs the segment in equilibrium, leaving GovVC the residual even though GovVC's own preference for high-tech is positive.

Third, the additionality cohort underperforms on exits and shows an employment—but not a profitability—orientation. A coarsened exact-matching comparison against observationally identical private-backed firms reveals a 4.88-percentage-point shortfall in successful-exit rates for additionality firms. On the real-side outcomes that public mandates invoke, a staggered difference-in-differences design shows that these firms are associated with markedly higher *employment* around the round (0.75–0.86 log points relative to both private-backed and general-population controls), while their operating profitability and asset turnover show no significant change. I read the employment result with caution: the pre-round trend is not flat, so parallel trends is not satisfied and the estimate is descriptive rather than causal. Its

direction is nonetheless consistent with an employment orientation rather than the selection of high-exit "frontier" firms in the sense of Draghi (2024).

Fourth, the exit shortfall is fragile and partly a single-market friction, not clear evidence of adverse selection. National GovVC funds operate under a hard geographic mandate—Italy's development bank cannot freely back a German firm—which I model as a feasibility constraint. Counterfactually relaxing it raises additionality only marginally (+1.0 point). More tellingly, when selection is measured on the model's *own* additionality set rather than the raw cohort, the private-VC exit shortfall is small and statistically insignificant, and relaxing the constraint shrinks it further. The apparent low quality of the additionality cohort is therefore weaker evidence of adverse selection than the raw comparison suggests, and to the extent a gap remains it is in part a product of the cross-border legal fragmentation emphasized by Draghi (2024), which pushes national public capital into worse local matches.

Related literature. The paper contributes to two strands.

The first is the empirical literature on government participation in venture capital. Lerner (2000) provides the foundational U.S. evaluation, showing that SBIR awardees grow faster and attract more follow-on private financing. International evidence is heterogeneous: Brander et al. (2015) report positive aggregate effects across countries, while Grilli and Murtinu (2014) and Bertoni and Tykvová (2015) find mixed results in European samples. Lerner (2002) stresses the political constraints that distort public selection; Howell (2017) and Cumming et al. (2017) document certification and crowding effects in specific settings. A growing literature extends these themes to European institutions and development finance (Bai et al. 2021; Andonov et al. 2025; Colonnelli et al. 2024; De Haas and González-Uribe 2024; Köppl et al. 2025; Berger et al. 2024, 2025). A common limitation is that selection is read off reduced-form correlations that conflate preference with competition—the difficulty Cumming and Johan (2019) highlight and that I address through structural matching.

The second is the structural matching literature in venture finance. Sørensen (2007) develops the two-sided matching estimator to separate sorting from value-added between top- and bottom-tier private investors. Ewens et al. (2022) use another framework to study contract

design, and Aragoneses and Saxena (2025) examine industrial-policy bias in European VC matching. My contribution is to bring this toolkit to government additionality: I adapt the estimator to a two-type market—government versus private—and use the recovered preferences to quantify how much of GovVC's footprint is new financing rather than substitution, and to characterize what the additional financing achieves.

2. Institutional Background

Government participation is a defining feature of the European venture-capital market. In my sample GovVC is involved in roughly 25% of investments over the past three decades. Relative to the United States, public investors play a larger role in many European markets, consistent with a thinner private VC sector and a longer tradition of public involvement in entrepreneurial finance (Bottazzi and Da Rin 2002; Da Rin et al. 2006; Leleux and Surlemont 2003).

In venture capital, limited partners (LPs) supply capital while general partners (GPs) select, monitor, and support portfolio firms (Gompers and Lerner 2004; Sahlman 2022). Public actors enter on either margin: governments may invest indirectly as LPs in venture funds, or directly as GPs through publicly sponsored vehicles. I classify both channels as GovVC, because each represents public capital shaping firm-level financing.[1]

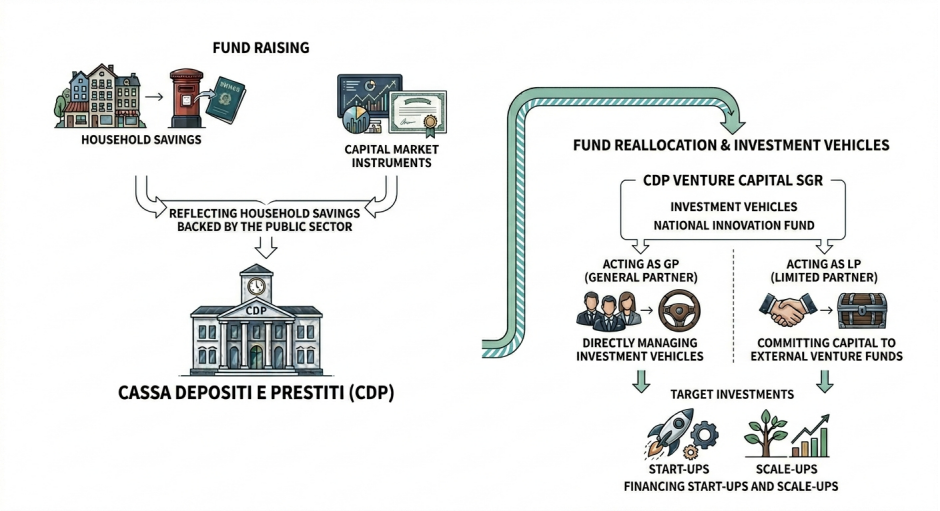
A useful illustration is Cassa Depositi e Prestiti (CDP), Italy's national development bank. CDP raises resources through state-guaranteed postal savings and capital-market instruments and channels them into investment vehicles, including the National Innovation Fund managed by CDP Venture Capital SGR.[2] Within this structure CDP acts both as an LP (committing capital to external funds) and as a GP (directly managing investment programs). At the supranational level, the European Investment Fund (EIF) and the European Investment Bank (EIB) play an analogous role under an explicitly cross-border mandate (Kraemer-Eis et al. 2015).

GovVC and private VC differ not only in ownership but in objective. Private VC is dis-

[1] Governments frequently act on both margins at once—committing capital as an LP to publicly backed funds while managing affiliated GP vehicles.

[2] CDP, "Financial Report," and https://www.cdpventurecapital.it.

FIGURE 1. *Government VC Example from Italy*



ciplined primarily by financial returns; GovVC is also guided by policy goals—innovation support, market completion, regional cohesion, and long-run development (Lerner 2002). Even where aggregate financing gaps are small, governments may intervene because entrepreneurial activity is spatially concentrated and some sectors and regions remain persistently underserved, or because they prioritize strategic domains—health, green technology, deep tech—in which private incentives are weaker or more short-term (Köppl et al. 2025). Whether these mandates direct capital toward firms that are genuinely constrained yet socially valuable is the empirical question this paper takes up.

3. Data and Descriptive Statistics

I compile a dataset of VC/PE investments in the 30 European countries (the 27 EU member states plus the United Kingdom, Norway, and Switzerland) from LSEG Refinitiv Deal Screener over 1989–2024. The data capture investments by both EU and non-EU investors into firms located in Europe. I merge with Orbis for firm financials and with Preqin for selected fund-level information.

GovVC investors are flagged using the lists in Berger et al. (2024) together with extensive manual verification. They include European-level institutions (EIF, EIB, EBRD, the European Commission), national development banks (KfW, Bpifrance, CDP), national and local funds

(Norfund, Friulia, Bayern Kapital), and public university funds (GU Ventures AB).[3] Foreign government investors[4] are reclassified as private VC, since they reflect a different policy regime and must remain in the control market for the additionality counterfactual to isolate the effect of *European* public-VC supply.

Descriptive statistics. Figure A1 panel (A) shows a rising trend in GovVC investment, reaching roughly 25% of European investments in recent years. Public participation is broad: GovVC invests across financing stages and across industry sectors, as panel (B) makes clear. Across the structural estimation sample, GovVC-backed firms are slightly older at first financing, somewhat less concentrated in high-tech, and far more likely to be matched to a same-country investor (96% versus 54% for private VC), reflecting the geographic mandates discussed below.

4. Reduced-Form Evidence

I begin with the cross-sectional and dynamic correlations between GovVC participation and firm outcomes. These motivate the structural analysis; as I argue in Section 4.4.

4.1. Selection patterns of portfolio companies

I relate first-round investment outcomes Y_d to the composition of the founding syndicate, absorbing time-invariant heterogeneity through fixed effects. The baseline (Panel A of Table A1) captures the intensive margin of government presence:

$$Y_d = \beta\, \text{GovVC}_d + X_d'\delta + \mu_c + \lambda_j + \theta_y + \varepsilon_d, \tag{1}$$

where GovVC_d equals one if at least one government investor participates in the firm's first round. To distinguish the mode of intervention, Panel B replaces the indicator with two

[3] Because the classification relies heavily on manual verification, the figures below capture, at most, a *lower bound* of GovVC activity, in both Europe and the parallel U.S. sample.

[4] Seventy-two unique non-European GovVC investors participate in European investments during the sample, 21% of all unique VC investors; there are 343 GovVC investors in total, European and non-European.

mutually exclusive treatments,

$$Y_d = \beta_1 \,\text{Pure}_d + \beta_2 \,\text{Mixed}_d + X_d'\delta + \mu_c + \lambda_j + \theta_y + \varepsilon_d, \tag{2}$$

where Pure_d denotes a syndicate composed exclusively of government investors and Mixed_d one combining government and private investors; purely private syndicates are the omitted category. The vector X_d collects investment-level controls (the number of investors, firm age at financing, and—where it is not itself the outcome—an early-stage indicator); μ_c, λ_j, and θ_y are country, industry, and year fixed effects, so coefficients are identified within country–industry–year cells. I estimate (1) and (2) by OLS; binary outcomes are linear probability models, and I omit the industry effect in the high-tech regression since that outcome is defined at the industry level.

European GovVC invests substantially smaller amounts: government participation is associated with investment volume about 24% lower on average (−0.276 log points), and the gap is far larger for purely public syndicates (−0.490, about 39%) than for mixed ones (−0.158, about 15%). It tilts modestly toward early-stage targets—raising the probability of an early-stage round by 1.5 percentage points overall and 4.0 points for pure-government syndicates—yet is markedly *less* likely to back high-technology firms (−5.6 points overall, −9.4 for pure-government). I find no detectable difference in successful-exit rates between GovVC- and privately backed firms.

Taken at face value, this pattern suggests that European GovVC fills an early-stage capital gap rather than identifying frontier firms: smaller checks into earlier, less technology-intensive ventures, with no exit-performance premium. But "at face value" is precisely the problem. The economically relevant questions—how much of this financing is incremental, and what it achieves—require separating GovVC's preferences from its competitive position, which I take up in Sections 5–7.

4.2. Crowding patterns: an event study

I next ask whether early GovVC backing crowds private capital in or out over the firm's subsequent rounds. I estimate a panel PPML model of follow-on syndicate size and investment volume:

$$E[Y_{i,r} \mid \mathbf{X}] = \exp\left(\alpha_i + \mu_t + \sum_{k=2}^{5} \beta_k (Pure_i \times I_r^k) + \sum_{k=2}^{5} \gamma_k (Mixed_i \times I_r^k) + \delta' \mathbf{X}_{i,r}\right), \tag{3}$$

where $Y_{i,r}$ is syndicate size or volume for firm i in round r, α_i are firm fixed effects, μ_t are month-year fixed effects, and $Pure_i$ ($Mixed_i$) indicates that firm i's first round was funded purely (or in a mixed syndicate) by GovVC.

Figure 2 reports the estimates. The aggregate effect of any GovVC presence (Panels A–B) is muted to slightly negative. This masks sharp heterogeneity. Firms whose first round was *pure* GovVC subsequently raise capital in increasingly larger syndicates (Panels C–D): the β_k are positive and widen with firm age, consistent with a certification channel (Lerner 2000) in which a credible public anchor reduces the asymmetric-information friction faced by later co-investors. Firms whose first round was a *mixed* public–private syndicate, by contrast, attract *fewer* additional investors later, with no volume differential—plausibly because the private partner's round-one participation already conveys the relevant signal. The cross-pattern—positive at the extensive margin of pure-GovVC presence, negative on the syndicate-composition margin—is difficult to reconcile with a simple "GovVC crowds out private capital" story; the public–private interaction operates differently depending on whether GovVC anchors alone or co-invests.

4.3. Scale-up to follow-on rounds

Because the crowding patterns are conditional on raising follow-on capital, I next ask whether GovVC- or private-backed firms are more likely to scale up. I define $ScaleUp_{i,r}$ as an indicator equal to one if firm i raises target round $r \in \{1, \ldots, 5\}$, and estimate

$$ScaleUp_{i,r} = \alpha_i + \sum_{k=2}^{5} \lambda_k I_r^k + \sum_{k=2}^{5} \beta_k \left(Pure_i \times I_r^k\right) + \sum_{k=2}^{5} \gamma_k \left(Mixed_i \times I_r^k\right) + \epsilon_{i,r}, \tag{4}$$

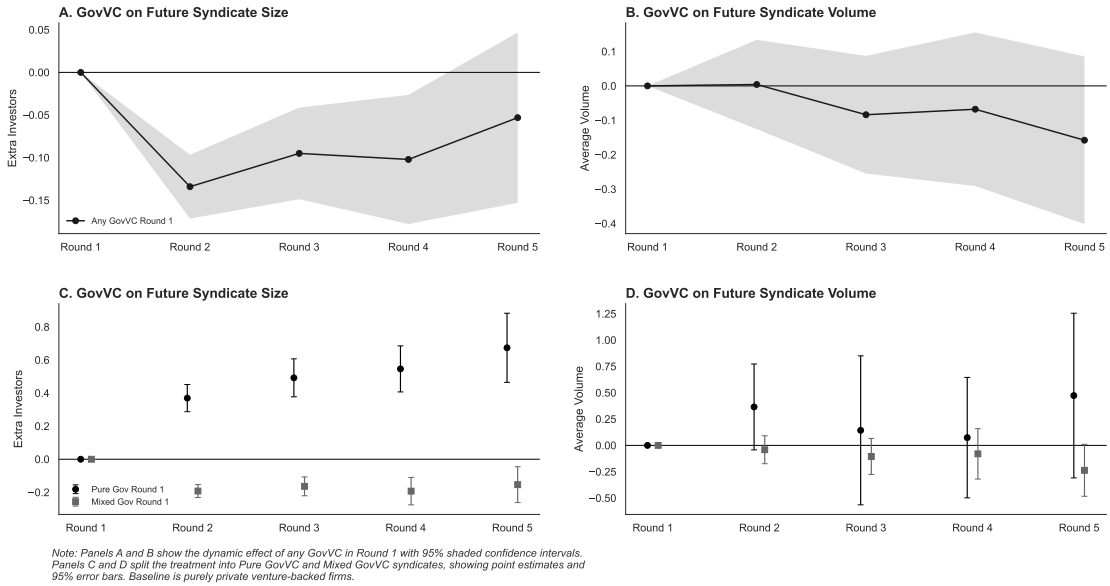


FIGURE 2. Dynamic effects of GovVC presence on follow-on syndicate size and volume. PPML event-study coefficients from Equation (3). Panels A–B: the aggregate effect of any GovVC presence in Round 1 on future syndicate size and volume; Panels C–D split the Round 1 treatment into Pure and Mixed GovVC syndicates. Baseline: pure-private-backed firms. All models include firm and month-year fixed effects and investment controls. 95% CIs clustered at the firm level.

where α_i are firm fixed effects, I_r^k is a target-round dummy, and β_k, γ_k measure differential scale-up relative to private-backed firms.

Pooling all government backing (Panel A of Figure 3) yields a small, mostly insignificant negative effect. Splitting the first-round treatment (Panel B) again reveals opposing patterns. Pure-GovVC firms are substantially *less* likely to raise follow-on rounds—about −6 points in Round 2, narrowing but remaining negative and significant through Round 5. Mixed firms are *more* likely to scale up in the near term (+3 points in Round 2, +2 in Round 3, both significant), with the premium fading by Rounds 4–5. The contrast reinforces a selection reading: a mixed first round reveals that at least one private investor anchored the firm at entry, and such firms clear continuation hurdles at higher rates; a purely public first round identifies firms no private investor anchored, and these struggle to attract continuation financing. Because pure-GovVC firms are the least likely to survive to later rounds, the larger follow-on syndicates observed among the survivors (previous subsection) reflect a positively selected subset, not a representative effect. Whether pure GovVC props up firms that would otherwise fail—a lemons concern—or sustains constrained-but-viable ventures cannot be settled by scale-up

rates alone. It is precisely the question the structural model is built to answer.

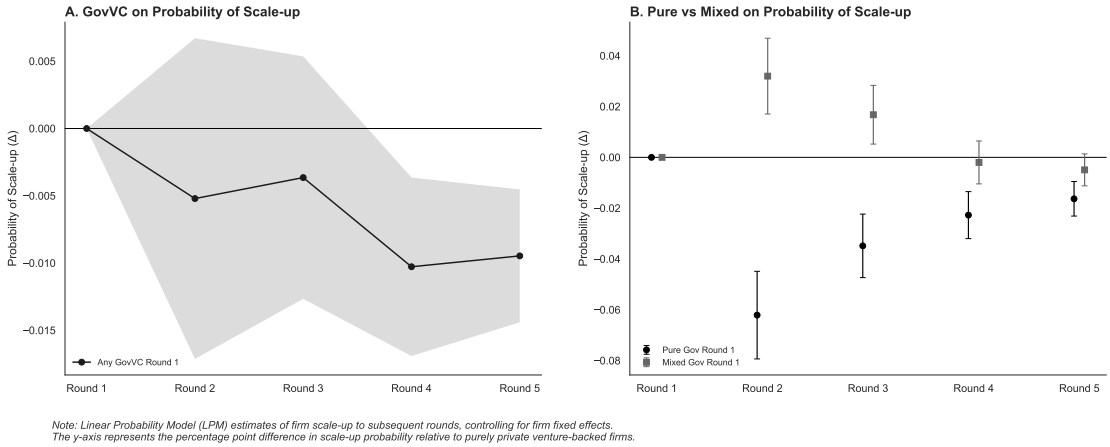


FIGURE 3. Dynamic effects of GovVC on firm scale-up. LPM event-study estimates with firm fixed effects. Panel A: any GovVC in Round 1; Panel B: Pure and Mixed GovVC separately. 95% CIs clustered at the firm level.

4.4. From correlations to preferences: the case for structural estimation

The reduced-form patterns measure the *composition* of observed allocations—which firm characteristics are over-represented in GovVC's portfolio. In a two-sided matching market, composition is the joint outcome of two distinct objects: the preferences of each investor type, and the competitive equilibrium under which those preferences are exercised (Sørensen 2007). The policy questions in this paper concern the first, not the second.

A simple example fixes ideas. Suppose both private VC and GovVC value high-tech firms, but private VC's appetite is slightly stronger and it is less capacity-constrained—more funds, no geographic restrictions. In equilibrium, the best high-tech firms match with private VC, and what remains for GovVC is lower-quality high-tech firms plus non-high-tech residuals. The OLS coefficient of GovVC participation on the high-tech indicator can then be *negative* even though GovVC's preference parameter strictly *favors* high-tech. The same logic runs in reverse for early-stage firms: GovVC can show a higher observed early-stage share—even when private VC weights early-stage more—if private VC is selective enough within the segment that GovVC absorbs the residual. The observed correlation does not, in general, identify which type weights a characteristic more.

What I recover instead are the preference parameters themselves. They are identified from the matching equilibrium: a stable matching implies a set of inequalities—each matched firm prefers its actual VC to feasible alternatives, and each VC prefers its matched firm to alternatives within its capacity—and the structural model below recovers the type-specific weights γ_g and γ_p from cross-VC and cross-firm variation in these inequalities, holding the competitive structure fixed. As Section 6.1 shows, the reduced-form coefficient and the structural preference can have *opposite* signs in this market. Reading the OLS coefficients as preferences would lead to qualitatively wrong policy conclusions.

5. A Structural Empirical Model of the VC Matching Market

I implement a two-sided matching model in the spirit of Sørensen (2007). The model recovers the latent match value of every feasible firm–VC pair, which I use in Section 7 to compute equilibrium counterfactuals.

5.1. Environment

The economy consists of firms $j \in \mathcal{J}$ and VCs $i \in \mathcal{I}$. Each firm j is located in nation $n(j) \in \mathcal{N}$, with $|\mathcal{N}| \leq 30$. Each VC has a type $\tau(i) \in \{g, p\}$. National government VCs face a hard same-region feasibility constraint, $\mathcal{F}_i = \{j : \text{region}(j) = \text{region}(i)\}$; pan-European institutions (EIF, EIB, EBRD, the European Commission) and all private VCs face no geographic restriction. Foreign government investors are reclassified as private VC.

Time is discrete, with each period t a calendar year. A firm receives funding from at most one VC per year, and both types compete simultaneously in every yearly market. Each VC's annual capacity is bounded by an exogenous quota q_i^t, set equal to the number of matches the VC executes in year t, following Sørensen (2007). The quota is a central modeling object: it fixes private VC's aggregate annual matching volume and therefore governs how counterfactual GovVC removal propagates through the equilibrium (Section 7). The maintained interpretation is that a VC's scarce resource is partner time, not money, and that deals absorb roughly equal time; Kaplan and Strömberg (2004) document that VCs ration deals on a time

rather than a capital constraint, so the number of investments per year is effectively capped[5].

5.2. Assumptions

The estimator rests on five assumptions, adapted from Sørensen (2007) and Sørensen (2005, chs. 2–3).

(A1) *One-period matching.* A firm receives at most one VC per year, so the empirical unit is a first-round (firm, VC, year) triple.

(A2) *Exogenous quotas.* q_i^t is taken as given, equal to the observed number of matches by VC i in year t. In particular, private VCs cannot expand year-t capacity in response to changes in the GovVC contract space.

(A3) *Hard feasibility.* The contract space $\mathcal{X}_t \subset \mathcal{I}_t \times \mathcal{J}_t$ excludes contracts ruled out by the geographic mandate (a national German GovVC cannot match an Italian firm).

(A4) *Linear-Normal match value.* The latent value of a feasible contract is the linear index in Equation (6), with i.i.d. Gaussian noise.

(A5) *Stability at the observed allocation.* The data X_t are the outcome of a stable matching given the realized $\{\varepsilon_{ij}\}$, so the inequalities defining stability are valid revealed-preference moments.

Assumptions (A1)–(A3) define the matching environment, (A4) imposes the parametric structure, and (A5) is the equilibrium-selection assumption that converts the matching problem into an estimable model.

5.3. Payoffs, the sharing rule, and stability

Each feasible match (i, j) has a valuation M_{ij}, interpreted as the expected net present value of the relationship at the time of the investment decision. Better firms and higher-value-added investors generate higher valuations; because I distinguish government and private investors, the two types may add different value to the same firm—a private VC may bring deeper commercial networks, while a public investor may provide access to public procurement or strategic legitimacy.

[5] Sørensen (2007) uses exaclty same modelling with the same reasoning.

The same valuation M_{ij} determines the preferences of *both* sides. Investors receive a fixed fraction λ of the surplus and firms the complement $1 - \lambda$, with λ exogenous and common. This is a *no-transfer* assumption: a weaker firm cannot attract a stronger investor by ceding a larger equity share. Its economic justification is the pledgeability constraint of Holmstrom and Tirole (1997): liquidity-constrained entrepreneurs facing moral hazard have limited pledgeable income, so over-pledging equity mutes incentives. The fixed sharing rule also delivers a unique equilibrium, without which the likelihood is not well defined (Bresnahan and Reiss 1991). Payoffs are therefore

$$\Pi_i\big(\mu(i)\big) = \lambda \sum_{j \in \mu(i)} M_{ij}, \qquad \Pi_j = (1 - \lambda)\, M_{\mu(j)j}. \tag{5}$$

The equilibrium concept is two-sided stability (Hatfield and Milgrom 2005; Roth and Sotomayor 1992). A matching X_t is stable if no unmatched firm–VC pair would prefer each other to their current partners and no VC could improve by reallocating its quota. Concretely, for each matched pair $(i, j) \in X_t$:

(i) $M_{ij} \geq M_{i'j}$ for every alternative feasible VC $i' \neq i$ that firm j could match (firm side); and

(ii) $M_{ij} \geq M_{ij'}$ for every unmatched feasible firm j' that VC i could take within quota (VC side).

The two sides are jointly necessary for identification: firm-side inequalities identify the type differential $\gamma_g - \gamma_p$ from cross-VC variation within firm, and VC-side inequalities identify the level γ_τ from cross-firm variation within VC. Because the sharing rule is fixed, stability reduces to this collection of inequalities and the equilibrium is unique (Gale and Shapley 1962). Under stability the parameters are identified up to a scale normalization on the latent value (I set $\mathrm{Var}(\varepsilon_{ij}) = 1$). The equilibrium is constructed via the (firm-proposing) cumulative-offer algorithm of Hatfield and Milgrom (2005), which under responsive preferences coincides with firm-proposing deferred acceptance; it is presented in Appendix B.

5.4. The match-value equation

For a feasible contract (i, j, τ), the latent match value is

$$M_{ij} = X_j'\gamma_\tau + Z_i'\eta_\tau + X_j'\Lambda_\tau Z_i + \delta_\tau D_{ij} + \varepsilon_{ij}, \qquad \varepsilon_{ij} \sim \mathcal{N}(0,1). \tag{6}$$

Firm j is summarized by $X_j \in \mathbb{R}^3$: log firm age, an early-stage indicator, and a high-tech indicator. Investor i is summarized by $Z_i \in \mathbb{R}^2$: log cumulative deal experience and the log of the investor's average deal size. The type-specific vectors $\gamma_\tau \in \mathbb{R}^3$ and $\eta_\tau \in \mathbb{R}^2$ capture how government and private investors weight firm and investor characteristics. The matrix $\Lambda_\tau \in \mathbb{R}^{3\times 2}$ governs firm × investor complementarities, so the marginal value of a firm attribute can depend on the matched investor's attributes.[6] The indicator $D_{ij} \in \{0,1\}$ equals one when firm and investor share a country, and δ_τ is the type-specific home-bias coefficient. The error variance is normalized to one, the surplus being identified only up to scale under stability.

The intercept μ_τ is omitted: stability inequalities are invariant to $M_{ij} \mapsto M_{ij} + c$, so any constant shift is not identified and would random-walk in the sampler. Consequently only cross-type differences ($\gamma_g - \gamma_p$, etc.) and the implied matching probabilities are interpretable; level coefficients absorb the unidentified location. I therefore report differentials and marginal probabilities as the primary quantities.

5.5. Estimation sample

The estimation sample covers 36 yearly markets, 1989–2024. I drop investments with missing deal amount, missing average investor deal experience, or missing firm age at financing. The final sample contains 22,957 firms receiving a first-round investment, matched with 4,194 unique VCs; GovVC accounts for 20.8% of investments and same-country matches for 63%. Continuous covariates (log firm age, log deal experience, log average deal amount) are z-standardized; the binary early-stage and high-tech indicators are kept on $\{0,1\}$. The high-tech indicator equals one when the firm's Refinitiv VentureXpert sector[7] is Biotechnol-

[6] Unlike Sørensen (2007), who omits these interactions for data reasons, I include the full 3×2 set; they also aid identification of the firm-side preferences.

[7] The "Investee Company VE Industry Sub-Group 1" field in VentureXpert.

ogy, Communications, Computer Hardware, Computer Other, Computer Software, Internet Specific, Medical/Health, or Semiconductor/Electronics.

5.6. Estimation strategy

I estimate $\Theta = \{\gamma_\tau, \eta_\tau, \Lambda_\tau, \delta_\tau\}_{\tau \in \{g,p\}}$ by Bayesian data-augmented Gibbs sampling, following Sørensen (2007) and Sørensen (2005, chs. 3–4).[8] The Bayesian route is not incidental. Maximizing the stability-implied likelihood is infeasible at scale: the likelihood contains a high-dimensional latent object—the unobserved match value of every feasible pair—and the stability inequalities carve this object into a non-rectangular region. Sørensen (2007) shows the posterior can instead be sampled by alternating two closed-form conditionals: a truncated-normal draw for the latent values (subject to stability) and a conjugate Gaussian draw for the coefficients. Each block is closed-form, so each sweep is linear in the number of feasible alternatives.

Two observations clarify the logic. Conditional on the augmented latent values, Equation (6) is a Gaussian-prior linear regression, with exact Gaussian posterior. Conditional on the coefficients, each M_{ij} is a univariate truncated normal whose truncation points are set by the rankings of alternatives in the firm's and VC's comparison sets. Iterating these draws yields a Markov chain whose stationary distribution is the posterior of Θ given the observed allocation (Gelman et al. 2013).

For tractability with 22,957 investments, I subsample comparison sets within each sweep: each matched firm is paired with $K_{\text{firm}} = 15$ randomly drawn feasible alternative VCs (firm side) and each matched VC with $K_{\text{vc}} = 15$ randomly drawn unmatched feasible firms (VC side). Random subsampling preserves posterior consistency under McFadden (1978)'s alternative-sampling theorem; standard Sørensen-style applications use $K = 5$–15 (Train 2009; Akkus et al. 2016). I run 200,000 iterations, discard 40,000 as burn-in, thin every 20, and retain 8,000 posterior draws. The prior on each coefficient is independent $\mathcal{N}(0, 10)$, weakly informative so the posterior is data-dominated even in small type-by-characteristic cells. The full algorithm

[8] See Chib and Greenberg (1995) on data augmentation and Logan et al. (2008) on two-sided matching with Gibbs.

and conditional posteriors are in Appendix C.

6. Preferences and Marginal Probabilities

The posterior coefficients are reported in Table 1 (firm and VC main effects, by type and as a differential). Because levels absorb the unidentified intercept, I read the estimates through the type *differential* $\gamma_g - \gamma_p$ and through the marginal matching probabilities below.

TABLE 1. Type-specific match-value coefficients: GovVC vs Private VC

	GovVC		**Private VC**		**Differential**	
Variable	Mean	SD	Mean	SD	$(\gamma_g - \gamma_p)$	SD
(Firm) log firm age	0.009	(0.011)	-0.045^{***}	(0.005)	0.053^{***}	(0.012)
(Firm) early stage	0.032^{***}	(0.012)	0.177^{***}	(0.006)	-0.145^{***}	(0.013)
(Firm) high-tech	-0.850^{***}	(0.021)	-2.371^{***}	(0.011)	1.522^{***}	(0.022)
(VC) log deal experience	0.015	(0.032)	0.191^{***}	(0.007)	-0.175^{***}	(0.033)
(VC) log average deal amount	-0.070^{*}	(0.038)	-0.034^{***}	(0.007)	−0.036	(0.038)
Same-geography (δ)	-0.157^{***}	(0.025)	1.056^{***}	(0.009)	-1.212^{***}	(0.027)

Notes: Posterior means with standard deviations in parentheses. "Differential" is the posterior of the GovVC-minus-Private-VC coefficient gap, signed against zero significantly positive (negative) values indicate that GovVC investors weight that characteristic more (less) heavily than Private VC in the match-value equation. Stars: * p<0.10, ** p<0.05, *** p<0.01 (posterior pseudo-*t*).

Because only cross-type differences are identified, Table 1 is read through its final column, the differential $\gamma_g - \gamma_p$. Four patterns stand out. First, GovVC weights *high-tech* far more heavily than private VC relative to its own baseline—the differential +1.522 ($p < 0.01$) is the largest in the table. Both types load negatively on high-tech in levels, but private VC's aversion is nearly three times as strong (−2.371 versus −0.850); the level coefficients are not separately interpretable (they absorb the unidentified intercept), whereas the differential says unambiguously that, holding competition fixed, GovVC is the *relatively* more high-tech-inclined type. Second, GovVC weights *early-stage less* than private VC (−0.145, $p < 0.01$): private VC's early-stage loading (+0.177) is more than five times GovVC's (+0.032), so at the preference level it is private capital, not public capital, that leans hardest into the seed stage. Third, GovVC places essentially no weight on investor *deal experience* ($\eta_g = 0.015$, insignificant) while private VC rewards it strongly (+0.191), giving a negative differential (−0.175)—consistent with private syndicates sorting on investor quality in a way public vehicles do not. Fourth, the

home-bias parameter δ diverges sharply: private VC's same-country coefficient is large and positive (+1.056) while GovVC's is *negative* (−0.157). This does not mean public investors shun local firms—the hard feasibility constraint (A3) already forces national GovVCs to stay within region, so the binding geography is absorbed by the constraint rather than by δ; conditional on feasibility, the residual same-country pull is simply not a positive selection margin for public capital the way it is for private capital.

6.1. From coefficients to marginal probabilities

The coefficients are most readable as *marginal probabilities* of matching with one type rather than the other. Consider an average firm (covariates at the sample mean) that receives one offer each from a representative GovVC and a representative private VC (each with characteristics at its within-type mean). The firm matches with whichever yields the higher latent value. Under $\mathrm{Var}(\varepsilon_{ij}) = 1$, the difference $M_{ig} - M_{ip}$ is Normal with mean Δ from Equation (6) and variance 2, so the probability of a GovVC match is $\Phi(\Delta/\sqrt{2})$. Shifting one firm characteristic at a time—one standard deviation for continuous variables, $0 \to 1$ for indicators—and recomputing Δ yields the per-characteristic marginal effects in Table 2.

TABLE 2. Marginal effect on the probability of matching with a GovVC (per 1-SD change in firm characteristic)

Firm Characteristic	$\Delta P_{\mathbf{GovVC}}$ **(pp)**	**SD**
log firm age	$+0.93^{***}$	(0.25)
early stage	-2.79^{***}	(0.30)
high-tech	$+6.79^{***}$	(0.08)
Baseline: $P(\text{matched to GovVC} \mid \text{average firm}) = 80.2\%$ (SD = 0.41 pp)		

Notes: Sørensen-style marginal-probability presentation. I construct an 'average firm' with X_F at the sample mean and an 'average GovVC'/'average Private VC' with X_V at the within-type sample mean, then compute the implied differential match value $m_g - m_p$. Because $\mathrm{Var}(\varepsilon) = 1$ in each match-value equation (Sørensen p. 2734 normalization), the differential is approximately $\mathcal{N}(\Delta, 2)$ and I report $\Phi(\Delta/\sqrt{2})$. Each row shifts one firm covariate by +1 SD (industry dummies set to 1 vs. baseline share) and reports the implied change in $P(\text{matched to GovVC})$ in percentage points. The baseline geography contribution is held at zero (firm and VC in different country/state) so the table isolates the role of firm characteristics from geographic homophily δ_τ. Posterior means over MCMC draws, standard deviations in parentheses. Stars: * p<0.10, ** p<0.05, *** p<0.01 (posterior pseudo-t).

At the average European firm, the model implies an 80.2% probability of matching with a government rather than a private investor in this one-against-one tournament at within-type

average characteristics. This is a *within-firm preference share at average characteristics*, and it should not be confused with the unconditional GovVC share in the data (20.8%), which is jointly determined by quotas, feasibility, and equilibrium competition across the full investor set. A one-standard-deviation increase in log firm age raises the implied probability by 0.93 points; switching to early-stage *lowers* it by 2.79 points; and switching to high-tech *raises* it by 6.79 points. High-tech is the single largest driver, consistent with European GovVC mandates that emphasize strategic and high-technology sectors (Lerner 2002). The early-stage result deserves emphasis. At the preference level—stripped of the competitive composition that drives the reduced-form coefficient the other way—European GovVC tilts *away* from the earliest-stage firms and toward more mature, later-stage rounds. This echoes Aragoneses and Saxena (2025), who show that industrial-policy considerations pull European public and public-influenced venture capital toward later-stage, larger-ticket deals rather than toward the seed-stage financing gap that public programs are nominally designed to close. The Λ complementarities in Appendix D rationalize the same tilt from the investor side: large-check investors of both types load negatively on early-stage, and Europe's public capital is concentrated in exactly such deep-pocketed, later-stage vehicles.

6.2. Two sign reversals: preference versus competition

Two of these marginal probabilities are informative precisely because their *sign* diverges from the reduced-form OLS coefficient on the same variable (Table A1). The OLS coefficient on high-tech is negative (−0.056, $p < 0.01$): in the cross-section, GovVC participation is 5.6 points less likely in high-tech firms. The structural marginal probability is positive (+6.79 points): switching a firm to high-tech raises the probability of a GovVC match by nearly seven points. A symmetric reversal occurs for early-stage status (OLS +0.015, $p < 0.05$; structural −2.79 points).

These estimates are not contradictory—they answer different questions, as anticipated in Section 4.4. OLS reports the *composition* of observed investments: the high-tech share among realized GovVC investments, jointly produced by GovVC's preference for high-tech *and* private VC's competition for the same firms. Private VC has more aggregate capacity,

no geographic restrictions, and competes hardest in high-tech; in equilibrium it absorbs a disproportionate share of high-tech firms, pushing GovVC's observed high-tech share below private VC's. The structural marginal probability strips out this competitive composition and reports the *preference*: at fixed capacity, an average GovVC outbids an average private VC for high-tech firms more often than not. The two agree on the underlying ranking only if private VC's competitive footprint in the segment is small; here it is large—private VC's high-tech coefficient is roughly three times GovVC's in magnitude (Table 1)—so private VC wins the segment in equilibrium even though GovVC's own preference for high-tech is positive. The same logic explains the early-stage reversal: private VC weights early-stage more at the preference level ($\gamma_g - \gamma_p = -0.145,\ p < 0.01$) but is so selective *within* the segment that GovVC ends up with the early-stage residual, producing a positive OLS coefficient.

The reversal is the substantive content of Section 4.4, and it is not an artefact. The policy interpretation of the additionality counterfactual rests on the structural preferences—"what happens if GovVC is removed?" is a preference-and-equilibrium question—not on observed composition. Reading the OLS coefficients as preferences would yield qualitatively wrong conclusions in this market.

7. Counterfactual Analysis

The structural model recovers the latent values $\{M_{ij}\}$ for every feasible pair. At the posterior-mean coefficients, I compute counterfactual matchings by re-solving the equilibrium under modified contract spaces. In each yearly market, the cumulative-offer algorithm (Appendix B) returns the unique stable matching consistent with the modified contract space and the observed quotas.

7.1. CF1: removing all GovVC

In CF1 I restrict the contract space to private contracts and re-solve. The GovVC-funded firms partition into two groups. Firms still matched to a private VC under CF1 represent a *substitution* channel: private VC was a viable backup. Firms now unmatched represent *additionality*:

GovVC was indispensable, and absent its supply these firms would have remained unfunded.

A feature of the design that the reader should keep in mind is that the fixed quota rules out genuine crowding-*in*. Each private VC's annual capacity is fixed at its observed volume, so removing GovVC cannot expand aggregate private investment—it can only reshuffle existing private matches. When a private VC absorbs a firm originally matched to GovVC, that absorption displaces a different firm the same VC funded in the data. The substitution I measure is therefore crowding-*out* within a fixed private envelope: GovVC presence displaces a fraction of private matches, and removing GovVC reverses that displacement at the margin. This makes the additionality estimate *conservative*: it counts as "substitutable" any firm a private VC would take by bumping another firm, which is the most demanding bar for declaring GovVC additional.

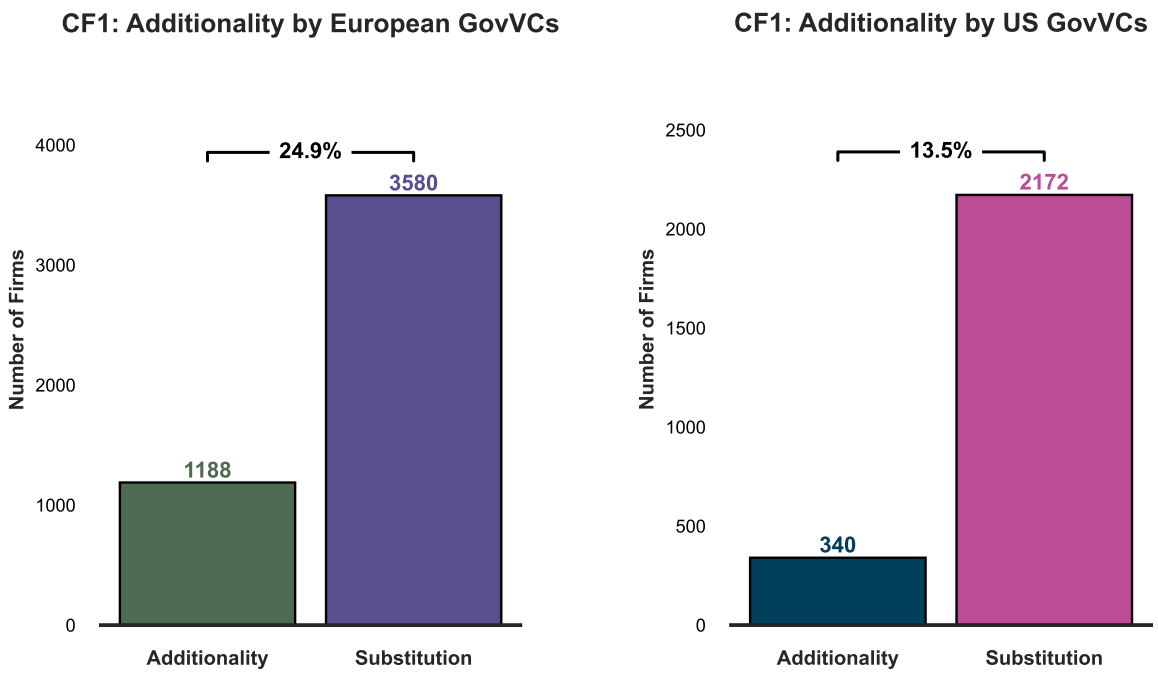
7.2. Results

Figure 4 reports CF1 for the European baseline and the parallel U.S. specification. In Europe, additionality is 24.9% (1,188 of 4,768 GovVC-backed firms): roughly one GovVC-backed firm in four would have gone unfunded absent public capital, while the remaining 75.1% are re-absorbed by private VC. The U.S. additionality rate is lower, 13.5%. In both economies, substitution—not additionality—is the dominant margin.

The Europe–U.S. contrast is informative about market structure. The U.S. private VC market is deeper but more geographically concentrated, and its public investors (SBIC and federal-program funds) target segments private VC bypasses; the lower U.S. additionality is consistent with this. European private VC is shallower in volume. The higher European additionality reflects this asymmetry: a non-trivial fraction of European GovVC matches occurs in firms within reach of some private VC, leaving a smaller additionality core whose financing genuinely depends on the public channel.

7.3. Are the additionality firms lemons?

If additionality firms are simply too risky for private VC to fund profitably, the additionality measured above could reflect a lemons problem rather than a corrected market failure



(a) European baseline. *(b) U.S. baseline.*

FIGURE 4. CF1: additionality versus substitution. Bars decompose GovVC-funded firms into Additionality (firms left unmatched when GovVC is removed) and Substitution / crowding-out (firms re-absorbed by private VC at the cost of another private match). Computed by re-solving the cumulative-offer algorithm at the posterior-mean parameters with GovVC removed. The European additionality rate is 24.9%; the U.S. rate is 13.5%.

(Akerlof 1978). I test this by comparing post-investment outcomes of additionality firms with those of observationally identical private-backed firms.

I use Coarsened Exact Matching (Iacus et al. 2012). Each additionality firm—a GovVC-matched firm that CF1 classifies as unmatched under GovVC removal—is matched to private-backed firms with the identical combination of firm-age bracket (0–1, 1–3, 3–7, 7–15, 15+ years), high-tech indicator, country, and three-year investment-year window. To equalize the exit window, I restrict the comparison to investments at least six years old. A successful exit is an IPO, acquisition, or merger; bankrupt or defunct firms are failures; all other statuses are no-exit.

Figure 5 reports the result. Successful exit is 14.1% among matched additionality firms versus 19.0% among matched private controls—a Private VC edge of 4.88 percentage points ($t = -3.39$, $p = 0.001$). Three qualifications discipline the reading. First, CEM controls for ob-

servables only; if private VCs also select on unobservable quality, the matched gap overstates the lemon component. Second, GovVC mandates explicitly target firms private capital is unlikely to fund, so a lower exit rate is partly mechanical—it is the population governments are designed to reach. Finally, the exit-rate metric is silent on the social returns—employment, R&D, regional development—that public mandates invoke. Section 8 turns to these outcomes.

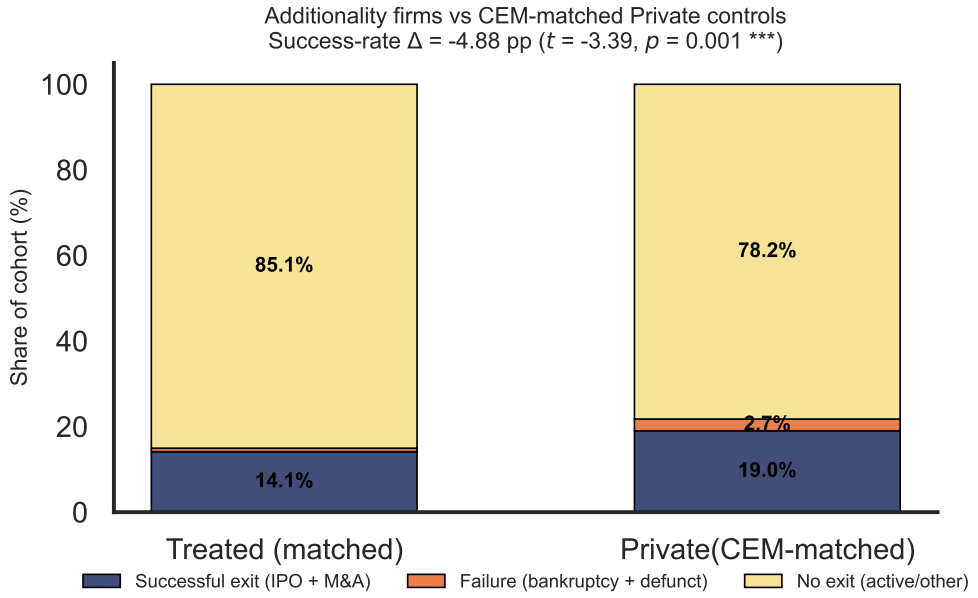


FIGURE 5. The Private VC exit edge. Outcome distribution for CEM-matched additionality firms (treated) and pure-private firms (control). Coarsened Exact Matching on (firm-age bracket, high-tech indicator, country, three-year investment-year window). Sample restricted to investments at least six years old. Successful exit is IPO, acquisition, or merger; failure is bankruptcy or defunct; no exit is everything else. The success-rate gap is −4.88 points ($t = -3.39$, $p = 0.001$).

7.4. CF-GEO: relaxing the single-market friction

European industrial policy is exercised within a fragmented single market. A start-up in Italy that expands to Germany confronts a different legal, tax, and compliance regime—an institutional friction absent in the United States, where an Ohio firm scales across states freely (Draghi 2024; Enriques et al. 2025). The friction binds investors as well, and most tightly the national GovVCs: Cassa Depositi e Prestiti, by mandate, cannot freely select portfolio firms across European borders. My model encodes this as the same-region feasibility constraint (A3). CF-GEO asks what additionality and selection would look like if that constraint were lifted—if

national public capital could move across European regions as freely as it moves within them.

A model-consistent difference-in-differences. The comparison has to be made with care, because the observed data already embed the constraint. Simply contrasting the observed additionality rate with a constraint-free re-solution would confuse two things: the effect of geography and the model's imperfect fit to the observed matching. I therefore run the counterfactual as a *within-model* difference. I re-solve the equilibrium twice on the same simulated base—once reproducing reality's geography (a_0) and once with the national-GovVC region constraint removed (a_1)—and read the geography effect as the difference $a_1 - a_0$. The intuition is a difference-in-differences: both re-solutions start from the same simulated firms and identical estimated preferences, and differ only in the geography "switch," so $a_1 - a_0$ nets out anything common to the two worlds (including model fit) and leaves the pure effect of the constraint. The remaining quantity, the gap between a_0 and the observed CF1 headline, is not noise to be discarded but a useful diagnostic: it is a *model-validation* check on how faithfully the estimated match values reproduce the additionality actually observed in the data. I apply the identical decomposition to the CEM exit edge of Section 7.3.

TABLE 3. Geographic counterfactual: GovVC additionality (CF1 kept; geography effect = $a_1 - a_0$)

	GovVC-funded	Additionality	Add. rate (%)
CF1 (observed)	4,768	1,188	24.9
a_0 model baseline (simulated)	4,692	1,404	29.9
a_1 CF-GEO (simulated)	4,768	1,475	30.9
Model-fit gap (CF1 − a_0) = −5.01 pp			
Geography effect ($a_1 - a_0$) = +1.01 pp (same simulated base)			
Step 1: 13,192 firms re-matched (baseline→CF); cross-area national-GovVC matches 0→4,531			

Notes: CF1 is the observed-factual headline and is kept. a_0 re-simulates the SAME geography as reality (its gap to CF1 is model fit); a_1 re-simulates under the changed geography. The geography effect $a_1 - a_0$ uses the same simulated base and the same firms, so it isolates geography without losing observations or conflating model fit.

Table 3 reports the additionality decomposition, and the model validates well. The model-consistent baseline is $a_0 = 29.9\%$ against an observed CF1 headline of 24.9%: the re-simulation modestly over-states additionality (a model-fit gap of −5.0 points), but it tracks the observed rate closely and preserves its economic magnitude, which is reassuring given that the match

values are estimated, not fitted to the additionality target.[9] Removing the cross-border constraint raises additionality to $a_1 = 30.9\%$, a geography effect of $a_1 - a_0 = +1.0$ point: freed from the region mandate, national GovVC reaches a modest set of additional firms it could not otherwise feasibly serve. The effect is small because the re-allocation is large in gross but nearly offsetting in net—lifting the constraint re-routes 4,531 national-GovVC matches (from zero) across region lines and re-shuffles 13,192 firms relative to the constrained baseline, but most of the newly reachable firms displace others that were themselves additional.

TABLE 4. Geographic counterfactual: Private VC exit edge over the additionality cohort

Additionality cohort	Add. exit (%)	Private exit (%)	Edge (pp)
a_0 model baseline (sim.)	17.26	19.00	+1.75
a_1 CF-GEO (sim.)	18.28	19.65	+1.37
Geography effect on edge ($a_1 - a_0$) = −0.38 pp (same simulated base)			
CF1 observed cohort edge = +4.88 pp			

Notes: "Edge" is the Private-minus-additionality gap in successful-exit rate among CEM-matched firms. The geography effect compares the a0 (baseline) and a1 (CF) cohorts, both simulated on the same base. A wider edge means the geography change pushes GovVC further into lower-quality (lemon) firms.

The selection result is the more informative one (Table 4), and it substantially qualifies the lemons reading. The large, statistically significant exit edge in the *observed* CF1 cohort (+4.88 points, $p = 0.001$; Section 7.3) does not survive the model-consistent re-simulation. On the common simulated base, the additionality cohort's exit edge is only +1.75 points under the region constraint (a_0) and +1.37 points once it is lifted (a_1), a geography effect on the edge of −0.38 points. Crucially, *neither simulated edge is statistically distinguishable from zero*: the underlying Welch tests on the success-rate gap give $t = -1.18$ ($p = 0.24$) for a_0 and $t = -0.89$ ($p = 0.37$) for a_1.[10] In other words, among the firms the structural model itself classifies as additional, there is *no significant Private-VC exit edge*—no measurable lemon. Two forces are consistent with this. First, the observed CEM cohort and the model's additionality set do not perfectly coincide, and the raw edge is driven partly by firms the CEM procedure treats

[9] The model-fit gap, CF1 − $a_0 = -5.01$ points, is the distance between the observed-factual headline and its model-consistent re-simulation; it is reported for transparency and is netted out of the geography effect by construction.

[10] These p-values are not shown in Table 4, which reports only the point edges; I state them here because they are what discipline the interpretation.

as additional but the model does not; the 4.88-point figure is thus an upper bound on any true lemons component. Second, to the extent a small residual gap remains, relaxing the single-market constraint shrinks it, so part of it is attributable to the friction rather than to selection—the region mandate pushes national public capital into worse *local* matches, and freeing it lets the same funds reach firms no worse than private-backed ones. Either way, the evidence for a genuine adverse-selection "lemons" problem is considerably weaker than the raw CF1 comparison suggests, a direct if quantitatively modest instance of the Draghi (2024) thesis that cross-border fragmentation, not only a financing gap, degrades the allocation of European risk capital.

8. What Does GovVC Deliver? Real Effects on the Additionality Cohort

The exit-rate evidence shows that additionality firms are less likely to reach an IPO or acquisition. But exit is a financial-return metric, and the stated mandates of European GovVC are social—employment, innovation, regional development (Lerner 2002; Howell 2017). A firm that grows payroll without producing a lucrative exit can still discharge a public mandate. This section asks what the additionality firms actually deliver on the real-side outcomes the mandates invoke, and in doing so recovers the *revealed* objective function of European GovVC.

8.1. Design

I take the additionality cohort from CF1—the GovVC-backed firms the structural model identifies as genuinely additional—and trace their Orbis financial histories around the first GovVC round. I study three outcomes available in Orbis: log employment, operating revenue scaled by total assets, and operating profitability (EBIT scaled by total assets, similar to Novy-Marx (2013)'s specification). Employment is the headcount metric most directly tied to the job-creation mandate; the two asset-scaled flows capture commercial traction and operating efficiency.

Identification is the staggered difference-in-differences estimator of Callaway and Sant'Anna

(2021). A firm's cohort g_i is the calendar year of its first GovVC round—the only round used—and event time is $e \equiv t - g_i$. I do *not* report two-way fixed-effects estimates: under staggered timing with dynamic effects, TWFE contaminates the estimand with already-treated comparisons (Goodman-Bacon 2021).

Because GovVC is selected rather than randomly assigned, I form the comparison group by matching before estimating, and I use *two* control groups so the result does not hinge on a single counterfactual population:

a. **Private-matched.** Privately backed firms, matched to the additionality firms exactly on country and broad sector and, within each cell, by propensity score on pre-round characteristics (firm age, log total assets, solvency, current ratio, return on assets, and vintage), $k = 5$ nearest neighbors without replacement within a 0.2-SD caliper.
b. **Orbis-matched.** The general Orbis firm population in the same countries, matched within country × cohort-year by propensity score on the same pre-round characteristics, calendar-aligned to the treated firm's round.

Matching exactly on the categorical confounders and by propensity within the cell keeps country (and, for the private comparison, sector) as hard constraints while balancing the continuous dimensions finely; coarsened exact matching on all covariates would fragment 30 countries × sectors × ratios into near-empty cells. The private comparison balances closely on the matched covariates. The general-population comparison is harder: because the additionality startups are drawn from a very different firm universe than the average Orbis firm, residual imbalance on firm size is difficult to eliminate, and I therefore treat the Orbis-matched estimates as a robustness check rather than the primary comparison. Appendix E reports the matching and estimation details.

Estimator. For a treated cohort g and calendar year $t \geq g$, the group-time average treatment effect on the treated is

$$ATT(g,t) = \underbrace{\mathbb{E}\left[Y_{it} - Y_{i,g-1} \mid G_i = g\right]}_{\text{additionality firms}} - \underbrace{\mathbb{E}\left[Y_{it} - Y_{i,g-1} \mid C = 1\right]}_{\text{matched, never-treated controls}}, \tag{7}$$

where the within-firm difference $Y_{it} - Y_{i,g-1}$ absorbs the firm fixed effect and differencing against same-year controls absorbs the year effect; identification rests on parallel trends between treated and matched controls, made credible by the pre-round matching, so no further covariates enter. Stacking cohorts on event time and normalizing to the year before the round gives the event-study coefficients

$$\theta(e) = \sum_{g \in \mathcal{G}} \omega_g \, ATT(g, g+e), \qquad \omega_g = \Pr(G_i = g \mid G_i \in \mathcal{G}, g+e \le T), \qquad \theta(-1) \equiv 0, \tag{8}$$

with $\mathcal{G}$ the set of treated cohorts and ω_g the cohort-share weights. Coefficients at $e \in \{-4, \ldots, -2\}$ are pre-trend placebos; $e \ge 0$ are the dynamic treatment effects. Inference uses the Callaway and Sant'Anna (2021) influence-function variance, clustered at the firm level, and the overall ATT is the post-round average $\bar{\theta} = \sum_{e \ge 0} \omega^e \theta(e)$. The model is estimated separately for each outcome and each control group.

8.2. Results

Figure 6 reports the event studies; each panel overlays the two matched control groups.

Employment. Post-round, additionality firms display sharply higher employment: the overall post-round ATT is +0.748 log points relative to private-matched controls and +0.862 relative to the general Orbis population (both $p < 0.01$), building close to monotonically through four or five years after the round. Taken alone this looks like a large job-creation effect. But the *pre-round* coefficients are not flat: they sit visibly below zero at $e = -4, \ldots, -2$ and are individually significant at $e = -3$ (panel (a) of Figure 6). Treated and control firms were therefore already on *diverging* employment paths before the GovVC round, so parallel trends fails and the post-round estimate cannot be given a causal interpretation. I therefore present the employment finding as a *descriptive association*: firms that receive additional public capital are on a markedly steeper employment trajectory around the round: not as evidence that GovVC *causes* that trajectory.

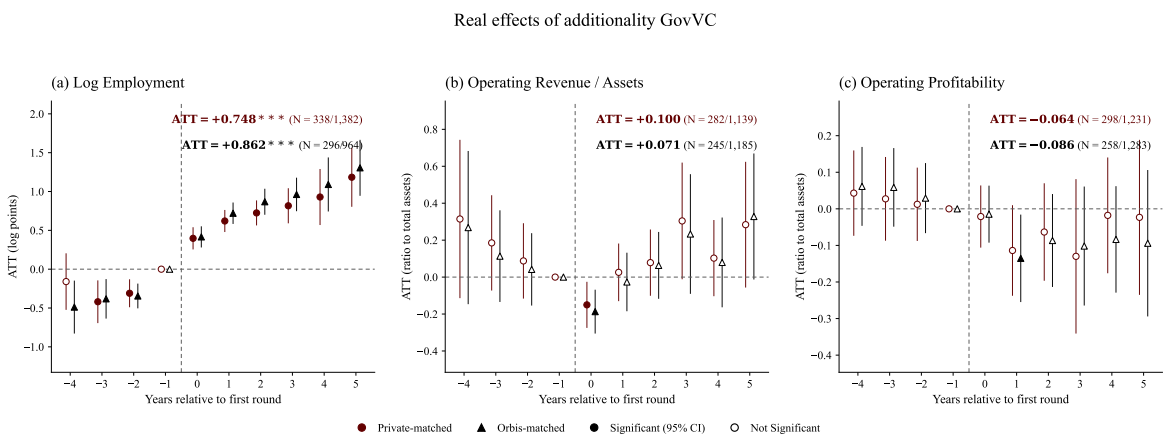


FIGURE 6. Real effects of GovVC on the additionality cohort. Callaway–Sant'Anna (2021) event-study coefficients $\theta(e)$ for the CF1 additionality firms, by outcome: (a) log employment, (b) operating revenue over total assets, (c) operating profitability (operating profit/loss over total assets). Each panel overlays two matched control groups—privately backed firms (circles) and the general Orbis population (triangles)—each matched within country on pre-round characteristics. Event time is normalized to $e = -1$. Filled markers denote point estimates whose 95% confidence interval excludes zero; hollow markers do not. Annotations report the overall post-round ATT with significance stars and matched sample sizes. Standard errors clustered at the firm level. Note that the employment series (panel a) exhibits a significant *negative* pre-trend at $e = -3$, so its post-round path is descriptive rather than causal (Section 8).

Financial performance. Operating revenue over assets is small and statistically insignificant (+0.100 and +0.071 for the private- and Orbis-matched comparisons). Operating profitability is, if anything, slightly *negative* and likewise insignificant (−0.064 and −0.086). On measured financial performance, then, additionality firms are indistinguishable from their matched counterparts: there is no profitability premium, and the point estimates lean the other way—a direction at least consistent with the lower-exit-rate evidence rather than against it.

Interpretation. Read together with the exit-rate shortfall, the real-effects evidence does not support a story in which GovVC selects high-performing frontier firms: additionality firms exit less often and are no more profitable per unit of assets. What the data do show—subject to the parallel-trends caveat—is an employment orientation, the pattern one would expect if European public capital is channelled toward job-supporting ventures that the private market, disciplined by financial returns, declines to fund.

9. Conclusion

I formalize selection in the European venture-capital market through a two-sided matching estimator in the spirit of Sørensen (2007), adapted to the government-versus-private application, and use the recovered preferences to re-solve the equilibrium without public capital.

Four findings are central. First, substitution dominates additionality: only 24.9% of European GovVC-backed firms (and 13.5% in the United States) would have gone unfunded absent public capital, while three quarters are re-absorbed by private VC within its existing capacity. Second, the structural preferences overturn the reduced-form narrative on targeting: at average characteristics, switching a firm to high-tech raises its probability of a GovVC match by 6.8 points and switching to early-stage lowers it by 2.8 points—both opposite in sign to the OLS coefficients—because the observed composition reflects private VC's competition, not GovVC's preferences; the early-stage tilt toward later, larger rounds echoes Aragoneses and Saxena (2025). Third, on real outcomes the additionality cohort shows an employment—but not a profitability—orientation: these firms follow a markedly steeper employment path around the round, while operating profitability and asset turnover do not move. I read the employment pattern cautiously, because its pre-round trend is not flat—consistent with an employment orientation rather than frontier-firm selection, but not a clean causal effect. Fourth, the exit shortfall itself: the 4.88-point Private-VC edge in the raw cohort is important once selection is measured on the model's own additionality set, and however relaxing the cross-border feasibility constraint shrinks it further, so the raw "lemon" signature is weaker evidence of adverse selection than it first appears and is partly a fragmented-market artefact.

These results recast the policy question. The relevant debate is not whether GovVC "works" in the aggregate—most of its portfolio would have been privately funded—but whether the employment associated with its additional quarter, and the efficiency lost to cross-border frictions, justify the public outlay.

References

Aghion, P. and Howitt, P. (1990). A model of growth through creative destruction.

Akerlof, G. A. (1978). The market for "lemons": Quality uncertainty and the market mechanism. In *Uncertainty in economics*, pages 235–251. Elsevier.

Akkus, O., Cookson, J. A., and Hortaçsu, A. (2016). The determinants of bank mergers: A revealed preference analysis. *Management Science*, 62(8):2241–2258.

Andonov, A., Li, A., and Smeets, P. (2025). Do development financial institutions create impact through venture capital investments? *Available at SSRN 5101408*.

Aragoneses, M. and Saxena, S. (2025). Industrial policy via venture capital. *Available at SSRN 5316925*.

Bai, J., Bernstein, S., Dev, A., and Lerner, J. (2021). The dance between government and private investors: Public entrepreneurial finance around the globe. Technical report, National Bureau of Economic Research.

Berger, M., Criscuolo, C., and Dechezleprêtre, A. (2025). Understanding government venture capital: A primer and a taxonomy. *OECD Science, Technology and Industry Working Papers*, 2025(06).

Berger, M., Dechezleprêtre, A., and Fadic, M. (2024). What is the role of government venture capital for innovation-driven entrepreneurship? *Documents de travail de l'OCDE sur la science, la technologie et l'industrie*.

Bertoni, F. and Tykvová, T. (2015). Does governmental venture capital spur invention and innovation? evidence from young european biotech companies. *Research Policy*, 44(4):925–935.

Bottazzi, L. and Da Rin, M. (2002). Venture capital in europe and the financing of innovative companies. *Economic policy*, 17(34):229–270.

Brander, J. A., Du, Q., and Hellmann, T. (2015). The effects of government-sponsored venture capital: international evidence. *Review of Finance*, 19(2):571–618.

Bresnahan, T. F. and Reiss, P. C. (1991). Empirical models of discrete games. *Journal of Econometrics*, 48(1-2):57–81.

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

Chib, S. and Greenberg, E. (1995). Understanding the metropolis-hastings algorithm. *The american statistician*, 49(4):327–335.

Colonnelli, E., Li, B., and Liu, E. (2024). Investing with the government: A field experiment in china. *Journal of Political Economy*, 132(1):248–294.

Cumming, D. and Johan, S. (2019). Government venture capital research: fake science and bad public

policy. *Venture Capital*, 21(1):121–131.

Cumming, D. J., Grilli, L., and Murtinu, S. (2017). Governmental and independent venture capital investments in europe: A firm-level performance analysis. *Journal of corporate finance*, 42:439–459.

Da Rin, M., Nicodano, G., and Sembenelli, A. (2006). Public policy and the creation of active venture capital markets. *Journal of Public Economics*, 90(8-9):1699–1723.

De Haas, R. and González-Uribe, J. (2024). Public policies for private finance. *Annual Review of Financial Economics*, 17.

Draghi, M. (2024). The future of european competitiveness part a: A competitiveness strategy for europe.

Enriques, L., Nigro, C. A., and Troeger, T. H. (2026). The eu inc.: A half-naked 28th regime. *A Half-Naked 28th Regime(June 06, 2026). European Corporate Governance Institute-Law Working Paper*, (942).

Enriques, L., Nigro, C. A., and Tröger, T. (2025). Can us venture capital contracts be transplanted into europe? systematic evidence from germany and italy. Technical report, SAFE Working Paper.

Ewens, M., Gorbenko, A., and Korteweg, A. (2022). Venture capital contracts. *Journal of Financial Economics*, 143(1):131–158.

Gale, D. and Shapley, L. S. (1962). College admissions and the stability of marriage. *The American mathematical monthly*, 69(1):9–15.

Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., and Rubin, D. B. (2013). *Bayesian Data Analysis*. CRC Press, 3 edition.

Gompers, P. A. and Lerner, J. (2004). *The venture capital cycle*. MIT press.

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

Grilli, L. and Murtinu, S. (2014). Government, venture capital and the growth of european high-tech entrepreneurial firms. *Research Policy*, 43(9):1523–1543.

Hatfield, J. W. and Milgrom, P. R. (2005). Matching with contracts. *American Economic Review*, 95(4):913–935.

Holmstrom, B. and Tirole, J. (1997). Financial intermediation, loanable funds, and the real sector. *the Quarterly Journal of economics*, 112(3):663–691.

Howell, S. T. (2017). Financing innovation: Evidence from r&d grants. *American economic review*, 107(4):1136–1164.

Iacus, S. M., King, G., and Porro, G. (2012). Causal inference without balance checking: Coarsened exact matching. *Political Analysis*, 20(1):1–24.

Kaplan, S. N. and Strömberg, P. E. (2004). Characteristics, contracts, and actions: Evidence from venture capitalist analyses. *The journal of finance*, 59(5):2177–2210.

Köppl, S., Köppl-Turyna, M., and Christopoulos, D. (2025). The performance of government-backed venture capital investments. *Research policy*, 54(8):105270.

Kraemer-Eis, H., Lang, F., Torfs, W., and Gvetadze, S. (2015). European small business finance outlook. *EIF Research & Market Analysis*.

Leleux, B. and Surlemont, B. (2003). Public versus private venture capital: seeding or crowding out? a pan-european analysis. *Journal of Business Venturing*, 18(1):81–104.

Lerner, J. (2000). The government as venture capitalist: the long-run impact of the sbir program. *The Journal of Private Equity*, pages 55–78.

Lerner, J. (2002). When bureaucrats meet entrepreneurs: the design of effectivepublic venture capital'programmes. *The economic journal*, 112(477):F73–F84.

Logan, J. A., Hoff, P. D., and Newton, M. A. (2008). Two-sided estimation of mate preferences for similarities in age, education, and religion. *Journal of the American Statistical Association*, 103(482):559–569.

McFadden, D. (1978). Modelling the choice of residential location. In Karlqvist, A., Lundqvist, L., Snickars, F., and Weibull, J. W., editors, *Spatial Interaction Theory and Planning Models*, pages 75–96. North-Holland, Amsterdam.

Michelacci, C. and Schivardi, F. (2020). Are they all like bill, mark, and steve? the education premium for entrepreneurs. *Labour Economics*, 67:101933.

Novy-Marx, R. (2013). The other side of value: The gross profitability premium. *Journal of financial economics*, 108(1):1–28.

Roth, A. E. and Sotomayor, M. (1992). Two-sided matching. *Handbook of game theory with economic applications*, 1:485–541.

Sahlman, W. A. (2022). The structure and governance of venture-capital organizations. In *Venture capital*, pages 3–51. Routledge.

Sørensen, M. (2005). *Essays on Venture Capital*. PhD thesis, Graduate School of Business, Stanford University.

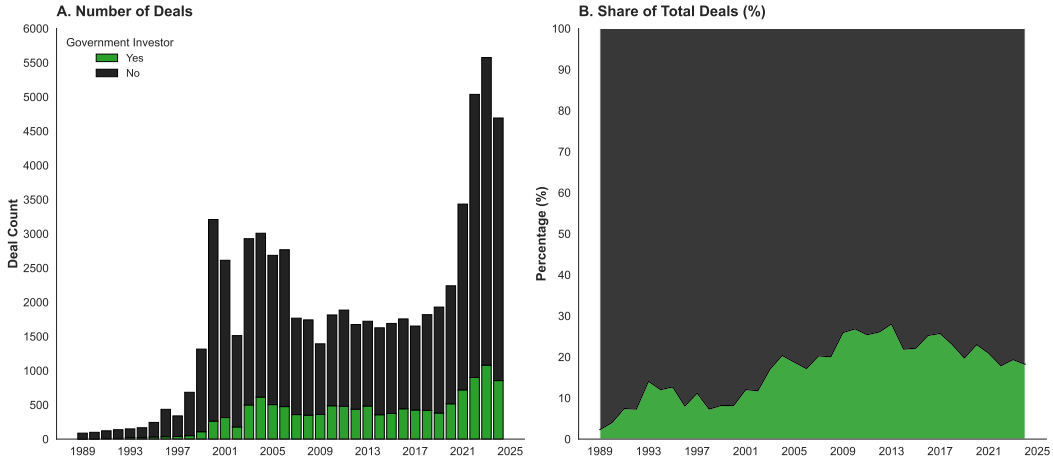
Sørensen, M. (2007). How smart is smart money? a two-sided matching model of venture capital. *The Journal of Finance*, 62(6):2725–2762.

Train, K. E. (2009). *Discrete Choice Methods with Simulation*. Cambridge University Press, 2 edition.

Appendix A. Additional figures and tables

A.1. Trends in Europe

(a) Government VC participation in Europe.



(b) Distribution by financing stage and industry segment.

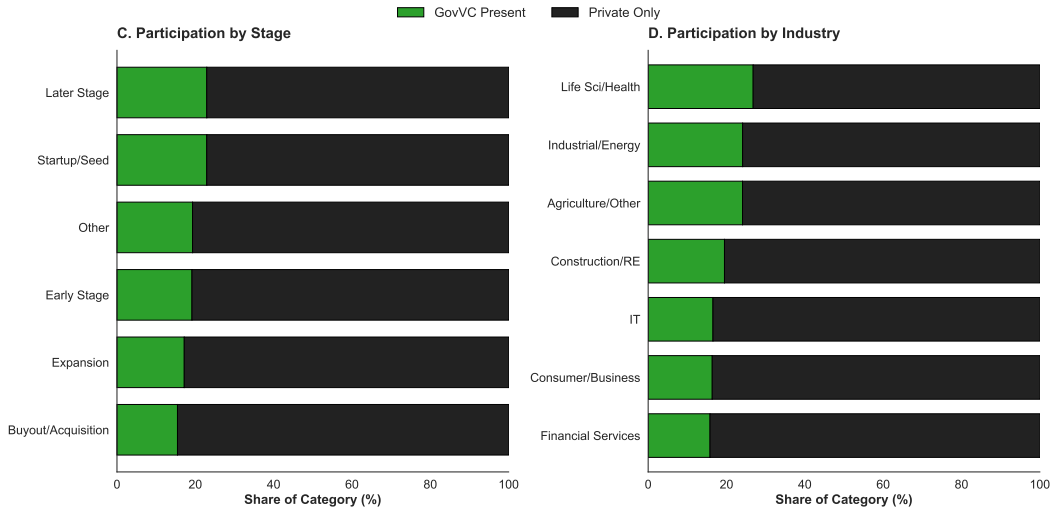


FIGURE A1. The landscape of government venture capital in Europe. Panel (a) plots the total count and relative share of GovVC investments over time. Panel (b) shows the composition of investor types across financing stages and industry sectors.

A.2. Selection of portfolio companies and GovVC participation—OLS results

TABLE A1. The Effect of Government VC on (*only first round*) Investment Outcomes

Panel A: Baseline Government Presence					
	Ln(Deal Vol)	*Success Exit*	*Early Stage*	*High Tech*	*Sustainable*
	(1)	(2)	(3)	(4)	(5)
Mean:	0.627	0.078	0.539	0.691	0.014
GovVC	-0.276***	-0.003	0.015**	-0.056***	-0.001
	(0.027)	(0.004)	(0.007)	(0.007)	(0.002)
Controls	Yes	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	No	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Observations	24,093	37,139	37,139	37,139	37,139
R^2	0.224	0.054	0.135	0.126	0.036
Panel B: Pure vs. Mixed Syndication					
	Ln(Deal Vol)	*Success Exit*	*Early Stage*	*High Tech*	*Sustainable*
	(1)	(2)	(3)	(4)	(5)
Pure GovVC	-0.490***	-0.008	0.040***	-0.094***	-0.003
	(0.043)	(0.006)	(0.010)	(0.009)	(0.003)
Mixed GovVC	-0.158***	0.002	-0.006	-0.027***	0.000
	(0.032)	(0.005)	(0.009)	(0.009)	(0.002)
Controls	Yes	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	No	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Observations	24,093	37,139	37,139	37,139	37,139
R^2	0.226	0.054	0.135	0.126	0.036

A.3. Full posterior of the structural parameters

The complete posterior of the type-specific coefficients and the Λ complementarity block is reported in Tables 1 and A2 in the main text; the analogous U.S. estimates underlie the U.S. counterfactual in Figure 4(b).

Appendix B. Cumulative-Offer Algorithm

The equilibrium matching is computed using the cumulative-offer algorithm (COA) of Hatfield and Milgrom (2005). In each yearly market the algorithm produces the firm-optimal stable allocation consistent with the observed quotas $\{q_i^t\}$ and the contract space $\mathcal{X}_t$. Under the responsive (rank-by-M_{ij}) preferences implied by the match-value equation, COA coincides with the many-to-one firm-proposing deferred-acceptance algorithm of Gale and Shapley (1962), which the estimation code implements directly.

Algorithm 1 Cumulative-offer / firm-proposing deferred acceptance in market $\mathcal{M}_t$.

Require: Firms $\mathcal{J}_t$, VCs $\mathcal{I}_t$, contract space $\mathcal{X}_t$, quotas $\{q_i^t\}_{i\in\mathcal{I}_t}$, posterior-mean match values $\{M_{ij}\}$.

Initialize: each firm's pointer at the top of its preference list; held set $H_i \leftarrow \varnothing$ for every VC i.
repeat
  for each unmatched firm j **do**
    Let $i^\star$ be the highest-ranked feasible VC in firm j's preference list not yet rejected.
    Firm j proposes contract $(i^\star, j)$ to VC $i^\star$.
  end for
  for each VC i **do**
    Add new proposals to the cumulative held set H_i.
    Choose $C_i \subseteq H_i$ to be the top q_i^t contracts in H_i ranked by M_{ij}.
    Reject all contracts in $H_i \setminus C_i$.
  end for
until no firm has a new proposal to make
return $\bigcup_{i\in\mathcal{I}_t} C_i$ as the stable matching.

For counterfactuals (Section 7) I run Algorithm 1 on the modified contract space (CF1: $\mathcal{X}_t \setminus \{(i,j) : \tau(i) = g\}$; CF-GEO: $\mathcal{X}_t$ augmented with the cross-region national-GovVC contracts excluded by (A3)). Firms unmatched at termination form the additionality cohort; firms

initially matched to GovVC but matched to private VC at termination form the substitution cohort.

Appendix C. Bayesian Gibbs Sampler in Detail

I estimate $\Theta = \{\gamma_\tau, \eta_\tau, \Lambda_\tau, \delta_\tau\}_{\tau \in \{g,p\}}$ by data-augmented Gibbs sampling, iterating two blocks with closed-form conditional posteriors.

Block 1: latent match-value augmentation. Conditional on Θ and the observed allocation X_t, each augmented value M_{ij} is drawn from a truncated normal whose truncation points are set by the firm-side and VC-side stability inequalities. For each matched pair $(i^\star, j) \in X_t$ I draw a comparison set of $K_{\text{firm}} = 15$ alternative feasible VCs i' (firm side) and $K_{\text{vc}} = 15$ alternative unmatched feasible firms j' (VC side), and draw

$$M_{i^\star j}^{(s)} \sim \mathcal{TN}\Big(\bar{m}_{i^\star j},\ 1;\ \big[\max\{M_{i' j}^{(s-1)}, M_{i^\star j'}^{(s-1)}\},\ +\infty\big)\Big),$$

truncated below by the largest recently drawn alternative; alternatives are drawn truncated above by $M_{i^\star j}^{(s)}$. Here $\bar{m}_{ij}$ is the structural mean from (6) at $\Theta^{(s-1)}$. Random subsampling preserves consistency under McFadden (1978).

Block 2: conjugate Gaussian posterior for the coefficients. Conditional on the augmented values, (6) is a linear regression with response $\mathbf{M}$ and design $\mathbf{W}$ whose rows are the type-interacted regressors

$$\big[\mathbb{1}_{\tau=g} X_j',\ \mathbb{1}_{\tau=p} X_j',\ \mathbb{1}_{\tau=g} Z_i',\ \mathbb{1}_{\tau=p} Z_i',\ \mathbb{1}_{\tau=g}\,\mathrm{vec}(X_j Z_i')',\ \mathbb{1}_{\tau=p}\,\mathrm{vec}(X_j Z_i')',\ \mathbb{1}_{\tau=g} D_{ij},\ \mathbb{1}_{\tau=p} D_{ij}\big].$$

With $\mathrm{Var}(\varepsilon_{ij}) = 1$ fixed and prior $\Theta \sim \mathcal{N}(0, 10 \cdot I)$ (Sørensen 2007, App. eq. (A1)), the posterior is conjugate Gaussian:

$$\Theta^{(s)} \sim \mathcal{N}\Big(\hat{\Theta}^{(s)},\ (\mathbf{W}'\mathbf{W} + 0.1\, I)^{-1}\Big), \qquad \hat{\Theta}^{(s)} = (\mathbf{W}'\mathbf{W} + 0.1\, I)^{-1}\, \mathbf{W}'\mathbf{M}^{(s)}.$$

Because $\mathbf{W}$ depends only on observables, $\mathbf{W}'\mathbf{W}$ and its Cholesky factor are pre-computed once, so each sweep requires one matrix–vector product and one multivariate-normal draw. I run 200,000 iterations, discard 40,000 as burn-in, thin every 20, and retain 8,000 posterior draws.

Appendix D. The Λ Complementarity Block

TABLE A2. Complementarity Matrix (Λ): Firm × VC interactions in the match-value equation

	log deal experience		log average deal amount	
Firm Characteristic	**GovVC**	**Priv VC**	**GovVC**	**Priv VC**
log firm age	0.006	−0.006	0.014	0.028***
	(0.013)	(0.004)	(0.014)	(0.003)
early stage	0.031**	0.025***	−0.072***	−0.044***
	(0.013)	(0.004)	(0.015)	(0.004)
high-tech	0.023	−0.100***	−0.170***	0.054***
	(0.026)	(0.007)	(0.032)	(0.008)

Notes: Posterior means with standard deviations in parentheses. $\Lambda_\tau[k, l]$ captures the interaction in the match-value equation between firm characteristic k (rows) and VC characteristic l (columns) for VC type τ. Following Sørensen (2007, fn. 9), firm × VC interactions span (log firm age, early stage, high-tech) × (log deal experience, log average deal amount), i.e. a 3×2 sub-block per VC type. With a single high-tech indicator (collapsed from the 8 Refinitiv high-tech sector codes following Sørensen §IV.B), the type-specific interaction with deal experience and deal-size proxies is economically interpretable end-to-end. Stars: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ (posterior pseudo-t).

The complementarity matrix Λ (Table A2) sharpens the picture by letting a firm attribute's weight depend on the matched investor's characteristics. The economically largest entry is GovVC's high-tech × deal-size interaction, −0.170 ($p < 0.01$): a GovVC's preference for high-tech firms *weakens* as the investor writes larger checks, whereas for private VC the same interaction is positive (+0.054). Public high-tech support is thus concentrated among smaller-ticket public vehicles rather than the large development banks. Both types show a negative early-stage × deal-size interaction (GovVC −0.072, private −0.044), so large-check investors of either type tilt away from the earliest-stage firms toward later, larger rounds—a pattern I return to below. Finally, private VC exhibits a strong negative high-tech × experience interaction (−0.100) that is absent for GovVC, indicating that the most experienced private investors are

precisely those that pull back from high-tech, while public investors' high-tech stance is invariant to experience. Taken together, the Λ block says that the "public high-tech" preference recovered above is a small-fund, low-experience phenomenon, not a property of the deep-pocketed public banks.

Appendix E. Real-Effects Matching and Estimation Details

The treated group is the CF1 additionality cohort. I link each firm to Orbis through an exact-match name waterfall (accent-folded, suffix-stripped name keys, blocked by website domain, country, and city, with a final globally-unique-name pass; only unambiguous one-to-one matches retained), and pull annual financials around the first GovVC round, collapsing to one record per firm-year. Outcomes are log employment, operating revenue over total assets, and operating profitability (operating profit/loss over total assets); the two ratio outcomes are winsorized at the 1st and 99th percentiles.

I match treated firms to controls measured at the pre-round year $g - 1$. For the private comparison, matching is exact on country $\times$ broad sector with a within-cell propensity score on $\{$firm age, log total assets, solvency, current ratio, return on assets, vintage$\}$; for the Orbis general-population comparison, matching is exact on country $\times$ cohort-year with a pooled propensity score on the same pre-round characteristics, and each matched control is anchored at the treated firm's cohort year and treated as never-GovVC-treated. In both cases I take the $k = 5$ nearest neighbors on the propensity logit, without replacement, within a 0.2-SD caliper. Group-time effects are estimated with a never-treated comparison group and aggregated to event time per Equations (7)–(8); the event window is $e \in [-4, 5]$, normalized to $e = -1$.