Buyers' Clubs for the Open-Source Flow: Collective Provision in the FLOSS Value Network

Stefane Fermigier (Abilian) · sf@abilian.com

Draft working paper v0.5, 2026-07-17. Theory and empirical design; every proposition carries a proof, a numeric cross-check, and a Monte-Carlo robustness report; the design is specified and no result is estimated.

Abstract

European supervisors ask firms to pool procurement and co-fund open-source alternatives to dominant providers. This paper models the institution that proposal implies: a buyers' club for the open-source flow. The code stock is free and inherited; funding is needed for the flow of maintenance, evolution, and compliance evidence, plus the excludable assurance, adaptation, and assistance services that hold a club together. Firms differ in capability (the cost of self-providing those services) and in need (the value of autonomy), the user typology of the FLOSS value network.

The model adds two open-source-specific margins to the collective-provision problem. First, membership can be paid in code: a firm can deliver fee-equivalent flow work at a private cost that falls with its capability, subject to a conversion friction. Riding (self-serving on the open flow without joining) then persists only when the friction is high relative to the fee, and since the fee falls as the club grows, scale converts riders into contributing members: the free-rider problem attenuates itself along the in-kind margin. Clubs sort into code-paying capable members and cash-paying incapable ones, the pattern observed in real consortia. With any compliance burden, near-costless code dues also make arbitrarily small clubs weakly viable, so the anchor buyer's role sharpens into jumping the critical mass for the cash population.

Second, the license regime enters: under effective reciprocity (copyleft reaching the deployment mode in use), riders' adaptations return upstream and defray the flow, so the club's fee and its under-participation wedge both fall with reciprocity strength, and free-riding becomes partially self-financing. The generic collective-provision results hold under the extension: participation is a game of strategic complements with a zero equilibrium and a critical mass, an anchor buyer of computable size eliminates the zero equilibrium of the cash benchmark, compliance burdens recruit the capable while pricing out marginal need, and equilibrium membership is inefficiently low, with the same wedge on both margins under a permissive license. Under reciprocity the need-margin wedge shrinks, and converting riders into members pays only while the per-rider return stays below the wedge.

Two further results locate the club against the voluntary-provision reading of open source: when visible contribution carries signaling returns that rise with the club's audience, the conversion friction falls with club size and rider conversion arrives earlier; and the club weakly dominates voluntary provision firm by firm, with its strict contribution exactly on the population priced out of self-provision and nothing to add below its critical mass. The club is also a natural open-source steward in the Cyber Resilience Act's sense, and a club large enough to fork disciplines its suppliers, connecting this model to the enclosure threshold of the supply-side companion project.

The second half of the paper specifies the measurement on the open-source poll: five hypotheses covering the typology prior on who joins, firm-level regulatory recruitment, conditional cooperation, the binding constraint on cooperation, and a within-firm criticality contrast in the accepted sovereignty premium, elicited through commitment-style outcomes with behavioral anchors, a conditional-cooperation ladder that doubles as a provision-point elicitation, a randomized binding-constraint item, and a mirrored supply block. The mapping from elicited constraints to policy instruments is fixed before fieldwork. Identification is limited to contrasts within the stated-preference boundary; the design is specified and no result is estimated.

1. Introduction

The policy call is on record: five Dutch supervisory authorities jointly ask firms and public buyers to pool demand, set joint requirements, and invest together in European alternatives, including open-source ones, arguing that first movers bear costs that later adopters do not, so that market forces alone will not build alternatives at scale (ACM et al. 2026). This paper supplies the theory of the institution that proposal implies and the design of its measurement, and it takes the open-source setting seriously enough to let it change the model. Three questions organize the empirical part: which firms are willing to pool procurement or co-fund development, under what conditions their willingness becomes commitment, and what constraint binds for the rest. The population is the one where such clubs would recruit first, European open-source users and providers, reached through the open-source poll.

Open source changes the collective-provision problem in three ways that a generic model of pooled IT procurement misses. First, the good has two layers with opposite economics: the code stock is free to all by construction, while the flow that keeps it alive (maintenance, security work, evolution, compliance evidence) is costly and undersupplied, the stock-flow distinction (Shaikh & Cornford 2011) that Jullien, Viseur & Zimmermann (2025) place at the base of their theory of FLOSS projects and open-source business-model dynamics, where a project's value creation is located in the flow. A buyers' club funds the flow and contracts the excludable 3A services (assurance, adaptation, assistance; the term is Jullien & Zimmermann's 2009) around it; it cannot and need not own the stock. This paper takes from the same source the two-way segmentation used from Section 6 onward: adopters by capability into Innovative, Frontier, and No-skill users (von Hippel 2001; Kogut & Metiu 2001; Ethiraj et al. 2005; Jullien & Zimmermann 2011), and needs into Control, Lock-out, Adaptable, and Low-cost positions (Jullien, Viseur & Zimmermann 2025). Second, capable firms can pay in code. Real consortia run on exactly this: vendors and capable users contribute engineering while others pay cash, and whether a capable firm free-rides or joins as a contributing member is a question of the conversion friction between its engineering and the club's needs. Third, the license regime is a policy lever inside the model: under effective reciprocity, the adaptations that self-serving riders build return upstream and defray the very flow the club funds.

The model delivers eight results. A distribution-free participation partition now has two membership currencies: code-paying members at low capability cost, cash-paying members at high capability cost and high need, riders in a band that exists only when the in-kind friction is high relative to the fee, and the rest staying with the incumbent (Proposition 1). Participation remains a game of strategic complements with a zero equilibrium and a critical mass (Proposition 2), and an anchor buyer of computable size eliminates the zero equilibrium of the cash benchmark, while the in-kind margin makes arbitrarily small clubs weakly viable whenever a compliance burden exists (Proposition 3). Scale converts riders: the rider band shrinks as the club grows and vanishes at a finite threshold, so free-riding is a small-club disease on the in-kind margin (Proposition 4). Compliance burdens raise every firm's willingness to join at a given fee and recruit riders into membership while pricing out marginal-need members, an ambiguity that is real in the Monte Carlo (Proposition 5). Equilibrium membership is inefficiently low, with an identical wedge on both margins under a permissive license, and the wedge falls with reciprocity: copyleft is a tax on self-service that partially self-finances the club, and where the return channel is strong enough, converting riders into members stops paying (Proposition 6). When visible contribution carries the signaling returns of Lerner and Tirole (2002), the conversion friction itself falls with club size, every monotone conclusion still holds, and rider conversion arrives earlier (Proposition 7). Against the voluntary-provision benchmark of Johnson (2002), the club weakly dominates firm by firm, its strict contribution is exactly the population priced out of self-provision, and below the critical mass it adds nothing (Proposition 8). A section of institutional readings sets the club against the Cyber Resilience Act's open-source steward category, reads the operating steward bodies of the policy record through the model's objects, and sets the club against the enclosure threshold of the supply-side companion, where a club's fork capability disciplines project owners.

The pooled object here is open source: the paper models and measures the open-source buyers' club. The three commitments the survey elicits are joint requirements-setting toward providers, pooled purchase contracts for supported open solutions, and co-funded development and maintenance of open-source components the members depend on; the club's economics run on the split between the free code stock and the costly service flow around it (§4). Open code changes joint investment in both directions: the co-funded asset cannot be withheld from any member, which removes the holdup risk that makes firms wary of co-funding a proprietary product. The same openness hands the code to non-members, which sharpens free-riding and moves the excludable services to the center of club design. The eventual data come from the open-source poll shared with the companion adoption study (`stratified-adoption-wp.md`): a sample of open-source-using and open-source-providing firms reached through European associations. Because the sampling frame includes open-source providers, who are themselves users, a mirrored block asks the supply side whether it would serve buyer collectives or pool with other vendors.

The measurement is the first behavioral test of the supervisors' proposal. Four elements of its design have no precedent in the literatures of Section 3: willingness measured as commitment under stated conditions in place of slogan agreement, a conditional-cooperation elicitation imported from public-goods experiments into a firm survey, the binding constraint elicited with the instrument mapping fixed before fieldwork, and both sides of the pooling market observed in one instrument through the sector's provider-user duality.

One boundary is set at the outset: the model and the design speak to willingness and structure. Whether pooled demand in fact scales European suppliers, lowers their costs, or delivers autonomy is a causal question that neither the theory nor stated willingness can answer, and the supervisors' efficacy claims remain assumptions here.

Section 2 lays out the policy context, Section 3 the related literatures, Section 4 the model, Section 5 the results, Section 6 the institutional readings, Section 7 the hypotheses and the empirical design with its identification limits, Section 8 the mapping from elicited constraints to policy instruments, Section 9 the verification suite, and Section 10 the boundaries of both parts; Appendix A carries the proofs and Appendix B the item wordings.

2. Policy context

2.1 The supervisors' proposal

Towards digital autonomy (ACM et al. 2026) is a joint position paper by five Dutch supervisory authorities: competition, financial markets, data protection, the central bank, and digital infrastructure. The signature matters, since these are the bodies that supervise the firms being addressed, speaking together across their mandates. The paper is normative and carries no data, survey or otherwise. Its central argument is the coordination failure of §1, and its remedy is collective action on the demand side, private alongside public.

The proposed mechanisms are five, and Table 1 arranges them against the assertions behind them and what remains untested. Around the mechanisms, the authorities assert, without measurement: that firms are reluctant to cooperate and that competition-law uncertainty is the cause; that pooled demand would let subscale European providers scale; that first movers face a cost premium; that autonomy justifies temporary cost and functionality concessions, more for vital processes; and that open standards and open source extend switching options. Every one of these is a testable firm-level claim, and none has been tested.

Table 1. The proposed mechanisms (ACM et al. 2026), the assertions behind them, and what is untested.

Mechanism	Who acts	What the authorities assert	What is untested
Buyer collectives: pooled demand, joint requirements	Firms, jointly	Pooled demand would let subscale European providers scale; reluctance to cooperate stems from competition-law uncertainty	Whether firms would join, which firms, and whether the legal diagnosis names the binding constraint (P1, P4)
Joint investment in European services, including open-source alternatives	Firms, co-funding	First movers face a cost premium that joint investment spreads	Whether firms would commit funds or staff time to co-funded development, and which firms (P1, P2)
Public authorities as long-term launch customers	Public buyers	First movers face a cost premium	Whether private willingness is in fact conditional on a public anchor (P3, the anchor rung of the ladder)
Procurement collectives and joint framework agreements	Public and private buyers	Reluctance to cooperate stems from competition-law uncertainty	Whether requirements are poolable and coordination costs bearable (P4)
A legally required European fallback for vital processes	The legislator, then firms	Autonomy justifies temporary cost and functionality concessions, more for vital processes	Whether firms accept larger concessions for critical workloads (P5)

A sixth assertion, that open standards and open source extend switching options, runs under all five rows and is the premise this paper's open-source framing takes as its object. The authorities emphasize the competition-law point: they state that competition law permits far more cooperation than firms assume, and they announce guidance as the remedy. That announcement makes the diagnosis testable. If elicited constraints put legal uncertainty first, the remedy is well targeted; if coordination costs, free-riding, heterogeneity, or a missing organizer dominate, guidance treats a cause that does not bind (§8).

2.2 The attitude-action gap

The representative benchmark is the Bitkom Cloud Report 2026 (Bitkom Research; CATI; n = 603 German firms with 20 or more employees; field weeks 14 to 20 of 2026; weighted to the federal business register; a repeated annual cross-section). Its attitude numbers are strong and rising: 85 percent of firms say Germany is too dependent on US cloud providers, up from 78 percent in 2025; 91 percent would prefer a German provider while 53 percent use one and 71 percent use a US one; 64 percent of cloud users say current US policy forces a strategy rethink, up from 50 percent in 2025; and 43 percent see no equivalent European alternative for their requirements. The willingness numbers are thin: 37 percent would accept at least one disadvantage in exchange for a Germany-only, foreign-access-protected cloud, 58 percent refuse every disadvantage offered, and 12 percent accept a price premium of 10 to 20 percent at the single price point asked. Realized behavior tells the same story from the other side: lock-in is the top realized switching barrier (59 percent), while easy switching is a must-have selection criterion for only 37 percent.

The gap between the two sets of numbers is the phenomenon this study measures around. Concern without individual willingness is exactly the gap collective action claims to close: if each firm's unwillingness reflects the first-mover problem (bearing switching costs alone while the benefits of a broader supply base accrue to everyone), then a coordination device changes the individual calculus, and the stated unwillingness is compatible with high participation in a collective. If the unwillingness reflects absent valuation instead, no coordination device helps. The two readings are observationally equivalent in an attitude survey and separable in a conditional design, and §7 builds exactly that design.

Two absences locate the open-source angle. The Bitkom report contains no open-source item and nothing on joint procurement or buyer cooperation, so both the instrument the supervisors propose and the asset class they name are missing from the representative record.

The institutional landscape the proposal points toward is documented at scale by Linåker and Muto (2026), a sixteen-country study of public-sector open-source policies and support structures, and their evidence brackets the model of this paper from both sides. The steward bodies they document are operating instances of the institution modeled here: OS2 in Denmark, a municipal association counting roughly 80 of the country's 98 municipalities, procures development and maintenance from an ecosystem of more than 60 vendors under standardized memoranda of understanding (Frey 2023); IMIO in Wallonia is a service provider co-owned by more than 140 municipalities that develops and maintains common components against its owners' requirements (Viseur & Jullien 2023); ADULLACT in France and VNG in the Netherlands play the same role. The dissolved initiatives are the predicted failure mode: Malta and Iceland both launched adoption-and-reuse programs in the early 2010s that dissipated once support lapsed, which is the fate the model assigns to a club that never reaches, or falls below, its critical mass (Proposition 2).

3. Related work

3.1 Cooperative purchasing

The cooperative-purchasing literature studies buying groups directly. Nollet and Beaulieu (2005) frame the join decision from the member's side: a purchasing group offers volume discounts, shared transaction costs, and pooled expertise, and it charges for them in reduced flexibility, diluted supplier relationships, and governance overhead. Whether joining pays therefore depends on how standardizable the requirement is and how much the member must adapt. Schotanus and Telgen (2007) supply the institutional map: a typology of organisational forms of cooperative purchasing, distinguished by how intensively members cooperate and which activities they share, from loose piggy-backing on another buyer's contract to jointly staffed purchasing organizations. Two regularities carry over to the present setting: groups form around standardizable requirements and stall on heterogeneous ones, and disputes over governance and gain sharing sink groups that scale. The evidence base is health-care and public-sector procurement. One documented case brings the literature to open source: Viseur and Jullien (2023) describe Communesplone, in which municipalities pooled development resources around a shared open-source platform. It is a public-sector case, and it observes a pool that formed; it says nothing about which private firms would join one or what stops the rest. Industrial organization studies buyer groups as instruments of bargaining power and commitment: Dana (2012) shows buyers gaining by committing jointly to a single seller before prices are set, which intensifies seller competition, and Inderst and Wey (2007) analyze how buyer power feeds back on suppliers' incentives to invest. In both, the pooled object is a rival good sold by strategic sellers; the free code stock, the in-kind margin, and a flow that non-members can take lie outside their scope.

3.2 Collective action and conditional cooperation

The collective-action literature explains why appeals to cooperate fail. Olson (1965) locates the failure in group structure: when the benefit of an action is non-excludable, members of a large group rationally withhold contribution unless selective incentives, small numbers, or coercion change the calculus, and latent groups stay latent however large the common gain. The experimental literature sharpened the picture at the individual level: in the strategy-method public-goods design of Fischbacher, Gächter and Fehr (2001), most contributors are conditional cooperators whose contribution rises with the contributions of others, unconditional cooperators are rare, and a minority free-rides throughout. Two implications matter for measurement. First, a stated willingness share is uninformative on its own: the same share is consistent with a group that assembles once a coordination device exists and with a dead letter, depending on how much of the willingness is conditional. Second, participation in a club with membership complementarities is a coordination problem with critical mass (Rohlfs 1974; Katz & Shapiro 1985), so beliefs about others are part of the demand structure, and an elicitation that omits them measures the wrong object.

3.3 Club goods and the private provision of open source

The theory of clubs and private provision says what can hold a group together around a non-excludable asset. Buchanan (1965) defines the club good: excludable, shared, non-rival up to congestion, with membership size and provision level determined together. The literature that grew from that definition treats membership and provision as jointly chosen design variables: Sandler and Tschirhart (1980) survey its central questions, among them member heterogeneity, mixed clubs, and discriminatory fee schedules, which reappear here as the two-currency membership of Proposition 1. Cornes and Sandler (1996) place club goods on the spectrum between private and public goods and treat joint products, where one activity yields a private and a collective output at once, the static form of the in-kind margin: a member's code contribution is simultaneously its dues and an input to the shared flow. Bergstrom, Blume and Varian (1986) characterize the private provision of a pure public good: provision concentrates on the highest-valuation contributors, aggregate provision falls short of the efficient level, and redistribution among contributors is largely neutralized. Johnson (2002) applies the private-provision logic to open source: development happens when some user's private value of an improvement exceeds its cost, and the equilibrium carries both free-riding and duplicated effort; von Hippel and von Krogh (2003) describe the same phenomenon as a private-collective model in which contributors earn private returns from producing a public good. Around open code, the excludability line is sharp: the code is a public good the moment it ships, while the services around it (assurance, adaptation, assistance, and the influence that comes with membership) are excludable and priceable. A buyers' club for open source is therefore a Buchanan club wrapped around a Bergstrom-Blume-Varian public good, and its design problem is which layer to make excludable. Research consortia are the nearest cost-sharing institution: Katz (1986) models cooperative R&D in which members share costs and results, with welfare turning on the spillovers to outsiders. The buyers' club differs in that its shared output is open to outsiders by construction, which moves the design problem from controlling spillovers to pricing the excludable services around them.

Johnson (2002) also supplies the benchmark provision mode against which the club must be judged: the flow provisioned by uncoordinated individual contributions, each made because some user's private value of an improvement exceeds its cost. Most of open-source history was written in that mode, and the club adds nothing where it works: where Innovative users are dense, private values cover the flow and coordination buys little. The club of this paper dominates the voluntary mode exactly where the typology says self-provision is priced out: firms with Control or Lock-out needs and Frontier or No-skill capability value the flow, cannot produce it, and are reached by no spontaneous contribution equilibrium. The two modes are complements across the type space, and the design's segmentation measures which mode a given population is in. The cost-sharing motive itself has an antecedent in the ecosystem-microeconomics propositions of Jullien, Stol & Herbsleb (2019), whose Proposition 5 holds that a firm favours an open-source strategy for a component that is not critical to its core business and still expected to evolve, precisely in order to share the cost of development. That proposition names the motive; what this paper adds is the institution that acts on it, with a membership rule, a fee schedule, an in-kind margin, and an equilibrium in which the free-riding region is characterized rather than assumed. Proposition 8 formalizes the boundary: the club weakly dominates the voluntary mode firm by firm, its strict contribution has positive measure exactly when a positive equilibrium exists, and that contribution lies entirely in the population self-provision does not reach.

3.4 Governance and steward institutions

A fourth literature describes the institutions that administer shared code. O'Mahony (2003) documents how community-managed projects guard the commons: foundations hold trademarks and copyrights, licenses and norms block private appropriation, and the legal shell is built ahead of the commercial partners it is meant to constrain. O'Mahony and Ferraro (2007) trace the emergence of governance inside one such community, a blend of bureaucratic rules and democratic selection in which leadership is granted conditionally and authority stays limited. Both study the supply side's self-organization, and the devices they document, membranes that keep the shared asset open while gating what surrounds it, are the design space of Section 4. On the buyer side the institutional record is newer: Linåker, Nummelin Carlberg and O'Riordan (2024) map open source programme offices across the European public sector into national-government, institution-centric, and association-based archetypes, the last pooling the resources of members individually too small to sustain their own open-source work. Linåker and Muto (2026) document these structures operating across sixteen countries. The association-based archetype is the institutional shell of the club modeled here, and Section 6 reads its operating instances through the model's objects.

Across these literatures, three gaps remain after the wider sweep. Buyer-group models in industrial organization and cooperative purchasing price the pooling of rival goods, and research consortia share excludable results; none models a club whose pooled asset is open to non-members by construction, with membership payable in kind. The conditional-cooperation evidence comes from individuals in laboratory games, while the units the policy addresses are firms deciding through committees. The club-theoretic and private-provision work on open source models the supply of code, and the governance and steward literature documents institutional forms; neither carries a participation model for demand-side pooling or a firm-level measurement of who would join. The policy documents that do address digital sovereignty carry no behavioral data (§2.1). The measurement specified here addresses those gaps: firm-level, commitment-based, conditionality-aware, and specific to the open-source object.

4. Model

A sector contains a continuum of firms of total mass $N > 1$. A firm's type is (k, v), drawn from an atomless joint distribution with bounded support, normalized to $[0,1]^2$, with positive density on its interior. Here k is the firm's cost of self-providing the 3A service layer on the open flow: low k is the Innovative user who can specify and develop, intermediate k the Frontier user, high k the No-skill user. The value v is what autonomy is worth to the firm; high v corresponds to the Control and Lock-out need types of the value network, low v to Adaptable and Low-cost.

The status quo is an incumbent (proprietary, typically non-European) solution with payoff normalized to zero. The open alternative has a free stock and a costly flow: keeping the software maintained, evolving, and compliant costs $C_0 > 0$ per period, and a regulation adds a compliance-evidence cost $R \geq 0$ per operator, borne in full by a firm operating alone and produced once inside a club.

A **club** of mass m (plus an anchor buyer of mass $M \geq 0$) funds the flow and compliance and delivers the 3A layer at per-member cost c, with per-member coordination and legal cost τ: a Buchanan (1965) club wrapped around a privately provided flow (Bergstrom, Blume & Varian 1986; for open source, Johnson 2002). Members pay dues in one of two currencies. **Cash**: the balanced-budget fee φ. **Code**: fee-equivalent flow work, which a firm of capability cost k produces at private cost $\psi k\varphi$, where $\psi > 0$ is the in-kind conversion friction (contribution pipelines, code review, the distance between firm-specific work and the project's needs; ψ is low when the firm's own adaptations double as flow work). A member's dues cost is therefore $\mu(k) = \varphi \min(1, \psi k)$: firms with $\psi k \leq 1$ pay in code, the rest in cash.

Non-members with the capability to do so can **ride**: build on the open flow privately, paying $k + R$ (their own service layer, their own compliance evidence). Under a license regime with **effective reciprocity** $\rho_\kappa \geq 0$, a rider's adaptations return upstream with value ρ_κ per rider, defraying the club's funding requirement. The balanced-budget fee at membership m and rider mass r is

$$\varphi(m, r) = \frac{\max(C_0 + R - \rho_\kappa r, 0)}{m + M} + c + \tau.$$

Effective reciprocity is a joint property of the license and the deployment mode: internal use escapes GPL-style obligations, network use triggers AGPL-style ones, and permissive licenses set ρ_κ near zero; the model treats ρ_κ as the reduced-form strength of the return channel.

Given a club, each firm chooses among **stay** (payoff 0), **ride** (payoff $v - k - R$), and **join** (payoff $v - \mu(k)$, in the cheaper currency). Firms take aggregates as given.

Assumption A1 (no solo provision). $C_0 + R > 1$ (the top of the v support): no single firm funds the flow alone. The coordination-failure premise, made explicit.

5. Results

5.1 Sorting: two membership currencies, and who still rides

Proposition 1 (participation partition; distribution-free). *Given a fee $\varphi > 0$, optimal choices partition the type space, up to ties of measure zero: firms join iff $\mu(k) \leq k + R$ and $v \geq \mu(k)$, paying in code iff $\psi k \leq 1$ and in cash otherwise; firms ride iff $k + R < \mu(k)$ and $v \geq k + R$, and all others stay. The rider set is nonempty only if $\psi(\varphi - R) > 1$, and it is the band $k \in (R/(\psi\varphi - 1), \varphi - R)$ intersected with $v \geq k + R$.*

The partition is the club composition observed in practice, derived: capable firms are code members, incapable high-need firms are cash members, and the need threshold for code members is $\psi k\varphi$, rising in k, so even modest-need capable firms join when their conversion friction is low. Riding, the traditional self-service of capable open-source users, occurs only in a capability band where the firm is too capable to pay cash happily and too friction-bound to pay in code. In the cash benchmark ($\psi \to \infty$) the generic partition reappears: members are the high-need, high-k corner and every capable high-need firm rides.

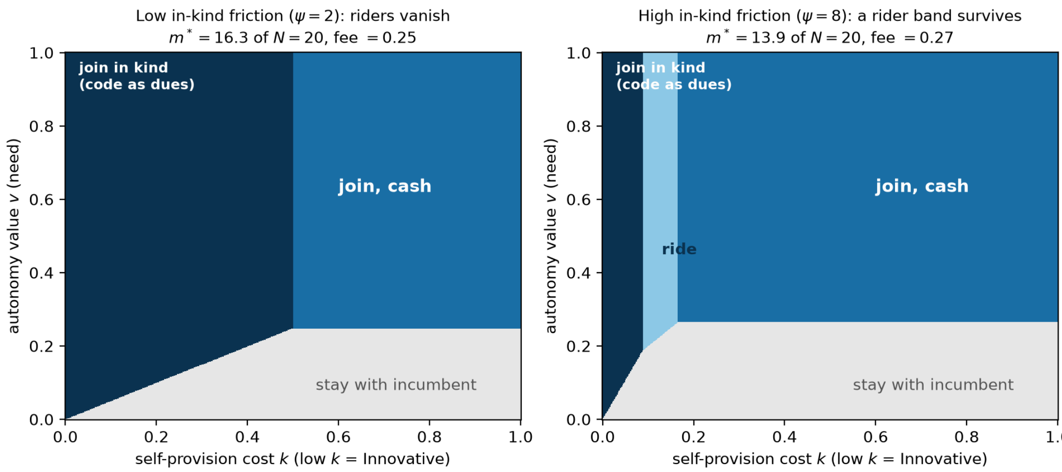

5.2 Critical mass, the anchor, and in-kind seeding

Let $B(m)$ denote the mass of firms preferring membership when the club has mass m, with the fee solving the balanced-budget identity at the induced rider mass; participation then has the interdependent-demand structure of Rohlfs (1974) and Katz & Shapiro (1985), joining paying more when more have joined.

Proposition 2 (equilibrium structure). *B is nondecreasing in m for $\rho_\kappa = 0$ (the dues cost $\mu(k)$ rises pointwise with the fee, so the join set shrinks as the fee rises), $B(0^+) = 0$ under A1, and the equilibrium set is a nonempty complete lattice containing $m = 0$. A positive equilibrium exists iff $B(m) \geq m$ for some $m > 0$, and when B is S-shaped, an unstable critical mass separates the basins. Any change that lowers the fee pointwise (lower C_0, c, or τ, or higher M) weakly raises the largest equilibrium $\bar{m}$.*

Proposition 3 (the anchor, and what in-kind dues change). *(i) Cash benchmark ($\psi \to \infty$): with full-support types, an anchor of mass $M > M^* = (C_0 + R)/(1 - c - \tau)$ makes $B_M(0) > 0$ and eliminates the zero-participation equilibrium; $\bar{m}$ is nondecreasing in M. (ii) With a finite conversion friction and $R > 0$, the join set has positive measure at every finite fee (firms with k near zero pay near-zero code dues and prefer them to duplicating compliance), so $B_M(0) > 0$ for every $M > 0$: arbitrarily small clubs are weakly viable, and the anchor's role is jumping the critical mass for the cash population rather than creating the first member.*

Part (ii) is an open-source-specific fact with a policy face. Compliance burdens plus code dues seed micro-clubs: a handful of capable firms sharing compliance evidence for a component they all patch is a club the model predicts at any scale. Such seeds lack the cash population, whose need threshold is the full fee; the anchor and the critical mass govern that population exactly as in the cash benchmark.

Remark (why neutrality does not apply). The neutrality theorem of Bergstrom, Blume and Varian (1986) says lump-sum public contributions to a voluntarily provided pure public good crowd out private contributions one for one, and would appear to gut the anchor result. It does not apply here, for two reasons the model makes exact. First, members do not make voluntary contributions that they could reduce: they pay a price for an excludable service layer, set by the balanced-budget fee schedule. Second, the anchor operates through that schedule: $\partial\varphi/\partial M < 0$ raises every type's private return to joining, at every membership level including the empty club where the voluntary-provision game has no margin at all. Crowding out would require the anchor to lower some firm's return to participation, and the fee identity shows the opposite. The neutrality logic lives where Johnson's voluntary-provision reading of open source lives, outside the club's membrane.

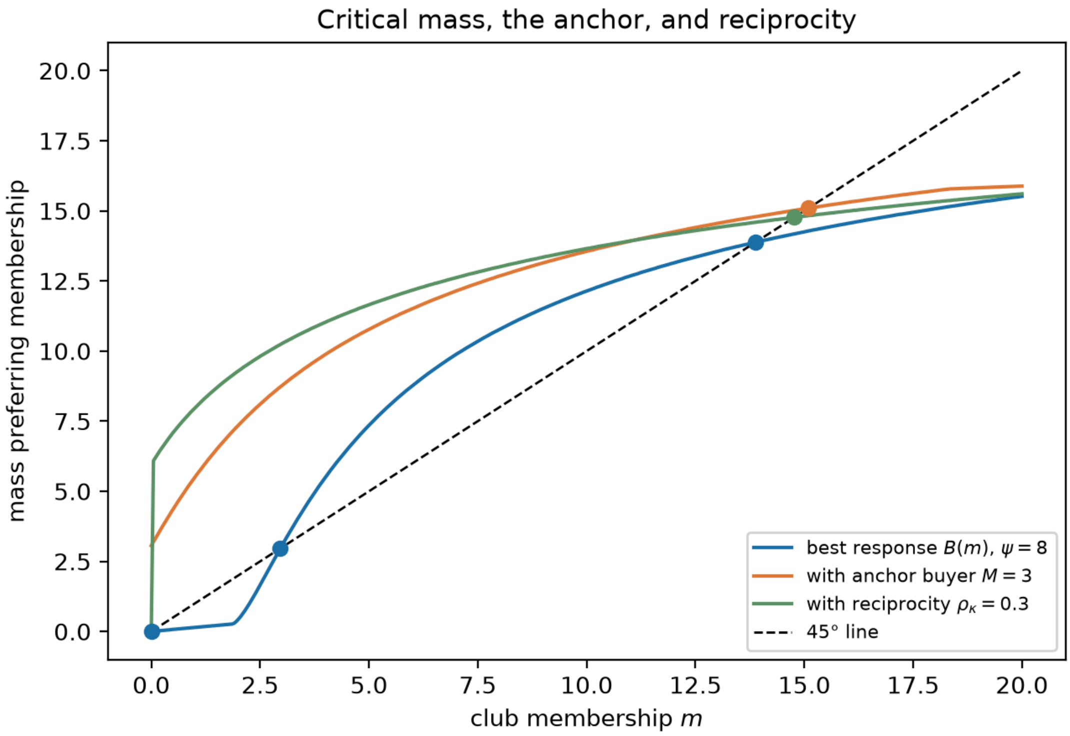

5.3 Scale converts riders; regulation recruits and excludes

Proposition 4 (free-riding is a small-club disease on the in-kind margin). *The rider band of Proposition 1 shrinks pointwise as the fee falls, hence as m grows; it is empty whenever $\psi(\varphi - R) \leq 1$. If the friction and the limiting fee satisfy $\psi(c + \tau + (C_0 + R)/(N + M) - R) \leq 1$, there is a club size beyond which no firm rides.*

As the club grows, the code-dues cost $\psi k \varphi$ falls with the fee while the ride cost $k + R$ does not, so the capable convert. A cash-only model has free-riding as a permanent feature of openness; the in-kind margin makes it a disequilibrium phenomenon of small clubs and high-friction contribution channels. The comparative institution is sharp: clubs that lower ψ (clean contribution pipelines, upstream-friendly governance, dues denominated in engineering) recruit exactly the population that cash-only clubs lose to riding. Section 5.4 endogenizes the friction: when visible contribution carries signaling returns, ψ itself falls with club size and the conversion arrives earlier (Proposition 7).

Proposition 5 (compliance costs: recruitment and exclusion). *(i) Firm level: at any fixed fee, raising R weakly expands the join set (the ride option deteriorates one-for-one while dues are unchanged) and never converts a member to a rider. (ii) Equilibrium: the fee rises with R at rate $1/(m + M)$, so marginal-need members are priced out while riders are recruited; the net effect on $\bar{m}$ depends on the type densities at the two margins, and both signs occur in the Monte Carlo.*

Proposition 6 (the wedge, and what copyleft does to it). *At any equilibrium with membership $\bar{m} > 0$, rider mass r, funding gap $G = \max(C_0 + R - \rho_\kappa r, 0)$, and fee $\varphi = G/(\bar{m} + M) + c + \tau$, define the wedge $w = G/(\bar{m} + M)$, the in-kind surplus mass $S = N \int_I (1 - \psi k)\, dF < \bar{m}$ over the in-kind member set I, and the factor $\Phi = 1 - S/(\bar{m} + M) \in (0, 1]$. Holding the other firms' choices fixed and rebalancing the budget: (i) the marginal social value of membership is $w\Phi$ on the need margin, on both dues branches, and $(w - \rho_\kappa)\Phi$ on the capability margin (converting a marginal rider) when $G > 0$, and both margins carry zero marginal value when the fee floor is slack. (ii) At $\rho_\kappa = 0$ the two margins carry the identical positive value $w\Phi$ and equilibrium membership is inefficiently low. (iii) For $\rho_\kappa > 0$, need-margin entry remains weakly beneficial, while converting riders raises welfare iff $\rho_\kappa < w$: under strong reciprocity riding self-finances the flow and conversion weakly lowers welfare. (iv) w is decreasing in $\rho_\kappa r$: effective* *reciprocity defrays the flow and shrinks both the fee and the need-margin under-participation gap. Under a permissive regime ($\rho_\kappa = 0$) the full wedge stands on both margins.*

Reciprocity is the license regime as club policy. A copyleft-covered flow taxes the self-servers in kind: their adaptations return, the funding requirement falls, and the club's fee with it. The model identifies a smaller under-provision problem where the return channel is strong, and it prices rider conversion: each convert's upstream return is lost, so recruiting riders into membership pays only while $\rho_\kappa < w$, and under strong reciprocity the club's interest is to let the capable ride. The choice of license for a co-funded component is therefore part of the club's design space, along with the friction ψ and the anchor. France's government instruction Circulaire 5608, recommending that five to ten percent of the savings from an OSS-based acquisition be contributed back to the concerned projects (Linåker & Muto 2026), is the reciprocity lever ρ_κ as an existing policy instrument.

5.4 Endogenous friction: signaling and scale

The friction ψ has so far been a technological constant. Lerner and Tirole (2002) identify a private return to visible open-source contribution, career concerns and reputation, and the return rises with the size of the audience that observes the work. In the club, the audience grows with membership, so the effective conversion friction, the net private cost of delivering one fee-unit of flow work, falls as the club grows. Proposition 7 endogenizes the friction.

Proposition 7 (signaling makes the friction fall with scale). *Let the conversion friction be a positive nonincreasing function $\psi(m)$ of club mass, and let $\rho_\kappa = 0$. (i) Dues $\mu(k; m) = \varphi(m)\min(1, \psi(m)k)$ are nonincreasing in m for every k, B is nondecreasing in m, the equilibrium set is a nonempty complete lattice containing $m = 0$ under A1, and the conclusions of Propositions 2, 3(ii), and 4 hold with $\psi(m)$ in place of ψ. (ii) If $\psi_A(m) \le \psi_B(m)$ for every m, then $B_A \ge B_B$ pointwise, the largest equilibria satisfy $\bar{m}_A \ge \bar{m}_B$, the rider set under ψ_A is contained in the rider set under ψ_B at every m, and wherever the rider band is empty under ψ_B it is empty under ψ_A. In particular, relative to the constant benchmark $\psi \equiv \psi(0)$, a signaling club has weakly larger membership, a pointwise smaller rider band, and weakly earlier rider conversion.*

Dues now fall with m through two channels, the fee and the friction, so the join set expands at least as fast as under constant friction and the conversion threshold of Proposition 4 arrives no later. The comparative institution of Section 5.3 gains a scale reading: governance choices that make contribution visible (public attribution, maintainer seats, credited releases) are ψ-reducers whose bite grows with the club, and a small club that cannot yet offer an audience is exactly the club that most needs clean contribution pipelines.

5.5 The Johnson boundary

Johnson (2002) supplies the benchmark provision mode of Section 3.3: uncoordinated self-provision, each firm acting alone when its private value covers its private cost. In the model that mode is the no-club outcome: each firm solves $\max(0, v - k - R)$, the served set is $V = \{(k, v) : v \ge k + R\}$, and the flow receives only the reciprocity returns ρ_κ per self-provider. Proposition 8 locates the boundary between the two institutions.

Proposition 8 (the Johnson boundary). *(i) At any equilibrium of the club game, every firm's payoff weakly exceeds its voluntary-mode payoff, strictly for members up to ties of measure zero, and the served set weakly contains V: the club never displaces voluntary provision. (ii) At any positive equilibrium, the flow requirement $C_0 + R$ is funded in full, and the club's strict addition to the served set is $\{(k, v) : \mu(k) \le v < k + R\}$, up to ties of measure zero; the addition lies entirely outside V and has positive measure whenever $R > 0$. (iii) If $B(m) < m$ for every $m > 0$, the zero equilibrium is unique and the club outcome coincides with the voluntary mode. The club therefore strictly dominates voluntary provision exactly when a positive equilibrium exists, and its entire strict contribution is on firms priced out of self-provision.*

The boundary organizes the institutional comparison. Where Innovative users are dense, most of the population sits inside V, voluntary provision already serves it, and the club's contribution to those firms is at most a cost saving, zero for the ones that keep riding. Where need concentrates in the Control and Lock-out rows and capability in the Frontier and No-skill columns, the population sits outside V, no spontaneous contribution equilibrium reaches it, and everything the club serves there is net new. The funding side splits the same way: above the critical mass the balanced budget covers the flow in full, while the voluntary mode contributes ρ_κ per self-provider and nothing under a permissive regime. The segmentation of the empirical design (Section 7) measures which side of the boundary a surveyed population is on.

6. Institutional readings

The club as steward. The Cyber Resilience Act defines, alongside manufacturers, the lighter category of open-source software stewards: entities providing sustained support for open-source products (Regulation (EU) 2024/2847; practitioner mapping in CNLL & Inno^3 2024). The model's club, which funds the flow and produces compliance evidence once for a shared component, is a candidate steward, and the sharing arrangement behind Proposition 5(i) is the economic content of that legal form. Nothing here claims the regulation achieves its security aims; R is cost incidence, and the steward category is the institutional container that makes sharing it lawful and legible.

The club in the wild. The steward bodies of Section 2 instantiate the model's objects with different membrane choices. OS2 collects membership from roughly 80 of Denmark's 98 municipalities and procures the flow and the 3A layer from an ecosystem of more than 60 vendors, with governance, procurement, and development models standardized and signed by the vendors through memoranda of understanding (Frey 2023; Linåker & Muto 2026). Read through the model, the standardized MoU is the membrane made contractual: it fixes the terms once for every member-vendor pair, which lowers the coordination cost τ. The members that initiate and collaborate on projects (Linåker & Muto 2026) pay part of their dues in code, the in-kind margin of Proposition 1 operating inside a mostly cash club. IMIO occupies the opposite corner, the club as its own supplier: more than 140 Walloon municipalities co-own the provider that develops and maintains their common components against their requirements (Viseur & Jullien 2023). The delivery cost c is therefore internalized, the fee schedule becomes an owner's budget contribution, and the in-kind margin takes the form of requirements work and co-development rather than arm's-length code dues. ADULLACT and VNG pool the support layer itself in association form (Linåker & Muto 2026), clubs whose excludable product is knowledge, coordination, and shared procurement capacity around the free stock.

The dissolved cases sit on the other side of Proposition 2′s threshold. Malta and Iceland launched adoption-and-reuse programs in the early 2010s that dissipated once support lapsed (Linåker & Muto 2026). In the model's terms the public program was the anchor: it made small-scale participation viable while it lasted, and its withdrawal returned the system to the basin of the zero equilibrium before a self-sustaining membership mass had formed. The bodies still operating hold large fractions of their eligible populations, which is where the S-shaped best response of Proposition 2 locates a stable club; the dissolved ones had no membership mass of their own to hold when the anchor left. Six cases fitting the model's partition of outcomes is consistency, and the measurement of Section 7 tests the participation structure these readings assume.

The club as fork capability. In the supply-side companion project's enclosure model (in `../../oss-business-models/papers/value-capture-wp.md`), relicensing pays only below a patience threshold, and the threshold falls with the community's capability to fork. A buyers' club is organized, funded fork capability: it holds the component's compliance evidence, employs or contracts maintainers, and can redirect its co-funding to a fork at low coordination cost. Membership mass therefore raises the fork probability in the companion's sense and shrinks the enclosure region. A club does more than buy services: it disciplines its upstream. This cross-model complementarity is stated as an interpretation; neither model prices the other's equilibrium.

7. Hypotheses and empirical design

The hypotheses derive from the model of Sections 4 and 5: P1 restates the participation partition (Proposition 1), P2 the firm-level half of the regulation result (Proposition 5), P3 the belief structure that the critical-mass and anchor results make decision-relevant (Propositions 2 and 3), P4 the map from constraints to the parameters the policy instruments act on (organized by the under-participation wedge of Proposition 6), and P5 the monotonicity of choice in the autonomy value.

7.1 Hypotheses

- **P1 (typology prior).** Willingness to enter the cash-margin commitments, joint requirements-setting (B2a) and pooled purchase (B2b), is highest for firms combining a Control or Lock-out need with Frontier or No-skill capability. For co-funding (B2c), which admits payment in staff time, the corner ordering is predicted only under high contribution friction, and no direction is registered (pre-registration amendment of 2026-07-17).
- **P2 (fixed-cost sharing).** Exposure to regulatory compliance obligations raises pooling willingness.
- **P3 (conditional cooperation).** Willing firms are mostly conditional cooperators: stated participation rises with the number and identity of committed peers, and unconditional joiners are a minority.
- **P4 (the binding constraint).** Among firms unwilling or hesitant to pool, the elicited binding constraint adjudicates between competition-law uncertainty (the supervisors' diagnosis), coordination cost, free-riding concern, requirement heterogeneity, and the absence of a credible organizer. The test is symmetric: either outcome is informative, and the announced remedy is well targeted only if the legal cause dominates.
- **P5 (criticality).** Within firms, the accepted sovereignty premium (price, feature lag, or reduced functionality) is higher for critical workloads than for non-critical ones.

Each hypothesis carries an explicit mechanism.

P1, mechanism. The value of membership to a firm is the gap between what it needs and what it can self-provide. Control and Lock-out are the needs the incumbent market leaves unserved: a firm that wants independence from a vendor cannot buy that independence from the vendor. Frontier and No-skill firms cannot close the gap internally; Innovative firms can, at the cost of their own engineering time (direct participation in the projects a firm depends on is the classic behavior of the capable user; Jullien & Zimmermann 2009; von Hippel & von Krogh 2003); Adaptable and Low-cost needs are met by the existing market, open or proprietary. Willingness therefore concentrates where the need is sovereignty-shaped and the capability is insufficient, and Proposition 1 derives the corner from payoff primitives: the Control or Lock-out rows of the need classifier crossed with the Frontier or No-skill columns of the capability classifier. The in-kind margin of Proposition 1 is the design lever that can convert the capable from riders into members. The corner prediction carries a scope condition that Proposition 1 makes exact. Cash membership requires $\psi k > 1$ and $\varphi \leq k + R$, so cash members occupy the high-k, high-v corner at every friction level, and the corner contrast on the cash-margin commitments is distribution-free. Membership as a whole is not cornered: the ordering holds only in the high-friction regime, where $\psi\varphi$ is large enough that the cheap-code region $k < R/(\psi\varphi - 1)$ lies below the Innovative capability range and the capable firm's best response is riding. When friction is low, code dues undercut self-provision across the Innovative column, the join region reaches down to low need levels there, and the overall corner ordering inverts: at the partition figure's parameters, the Innovative column's join rate is 0.92 against the other columns' 0.76 when $\psi = 2$, and 0.61 against 0.73 when $\psi =$

8, while the high-need cells saturate at low friction, so the corner cell is matched instead of dominant (Section 9). B2c admits staff time and therefore elicits both currencies: its cell ordering is diagnostic of the friction regime and carries no registered direction. The classifier cells come from Module A, Section 0, measured before any pooling item is shown, so the segmentation cannot be contaminated by the outcome.

P2, mechanism. Compliance obligations of the CRA type are fixed costs, largely invariant to firm size and duplicated across every user of the same component. A club converts many duplicated fixed costs into one shared cost, so the private return to membership rises with exposure, which is Proposition 5(i) at the firm level. The exposure measure is the regulatory perceived-risk composite from Module A, Section 3 (items RR1 to RR4), which captures the burden as the firm perceives it; the prediction is a positive coefficient on that composite in the commitment models. The claim is firm-level by design: whether regulation grows or shrinks the club in equilibrium depends on forces the survey cannot see, and Proposition 5(ii) shows the equilibrium sign can go either way.

P3, mechanism. A club with membership complementarities has multiple equilibria: joining pays if enough others join, and staying out is self-fulfilling when everyone expects everyone else to stay out. In such a game a firm's participation decision is a function of its beliefs about others, and the experimental evidence on individuals (Fischbacher, Gächter & Fehr 2001) finds exactly the profile the theory predicts: conditional cooperation as the norm, unconditional cooperation as the exception. Propositions 2 and 3 give the firm-level analogue: an anchor buyer of sufficient committed size removes the empty equilibrium, which is the formal content of the supervisors' launch-customer instrument. P3 predicts that the profile carries over from individuals to firms: stated participation should rise along the ladder of committed peers, the public-anchor rung should attract a distinct mass, and unconditional joiners should be a minority of the willing.

P4, mechanism. The five candidate constraints are five different failure models with five different remedies. Competition-law uncertainty is a perceived legal cost that guidance can remove; coordination cost is a transaction cost that an organizer can amortize; free-riding concern is the Olson problem that membrane design (§4) addresses; requirement heterogeneity is the stall condition of the cooperative-purchasing literature and bounds the poolable scope; a missing organizer is the coordination failure that anchors and credible institutions resolve. The supervisors assert the first and announce guidance as remedy. Eliciting the constraint adjudicates the diagnosis, and the test is symmetric because a distribution dominated by any other constraint redirects the instrument choice (§8). Proposition 6 organizes the same instruments around the need-margin under-participation wedge, so the elicited distribution also selects which parameter of the model policy should move.

P5, mechanism. The supervisors' proportionality premise holds that autonomy justifies larger concessions for vital processes. In the model's terms, the autonomy value of an alternative scales with the cost of losing, or being coerced through, the incumbent service, and that cost is larger for critical workloads. If choice is monotone in the autonomy value, accepted premia are ordered by criticality within each firm. The within-firm design carries the identification: firm-level heterogeneity (sovereignty attitudes, budget slack, identity-driven responding) is differenced out when the same firm answers the same battery for a critical and a non-critical workload, so the criticality gradient is estimated inside firms.

Table 2. Hypotheses, measurement, and tests (item IDs per the survey instrument; Annex A).

Hypothesis	Prediction	Measured by	Test
P1	Cash-margin commitment willingness (B2a, B2b) highest in the Control/ Lock-out × Frontier/No-	B2a–B2c; capability and need classifier (Module A, Section 0)	Binary and ordered models on B2a and B2b; registered directional corner contrast; B2c and the two-propor-

Hypothesis	Prediction	Measured by	Test
	skill corner; B2c direction unregistered, friction-dependent		tion cell contrasts reported alongside
P2	Regulatory exposure raises willingness	B2a–B2c; RR composite (Module A, Section 3)	Same models with the RR composite as predictor; firm-level claim only
P3	Participation rises along the ladder; unconditional joiners a minority of willing firms	B3, asked for B2b and B2c	Profile classification; one-sided binomial test that unconditional joiners are fewer than half of willing firms; threshold distribution reported
P4	The constraint distribution adjudicates the supervisors' diagnosis	B4 (buyers), S3 (providers)	Multinomial confidence intervals; modal constraint named with its margin over the runner-up
P5	Larger accepted premium for critical workloads	B6, asked twice (critical, non-critical)	Paired within-firm McNemar-type test on accepting any disadvantage; premium item reported separately

7.2 Design elements

The study is Module B of a three-module survey shared with the companion acceptance study; the poll architecture, the full item set with budgets, the assembled questionnaire, and the pre-registration draft are working documents listed in Annex A. Module order is fixed: the acceptance module (A) precedes the pooling module (B) and the context items (C), so the acceptance measurement is never primed by pooling or sovereignty content. Module B costs about four minutes of a fifteen-minute instrument, and the item budget is contractual: every added item must displace another. Appendix B carries the wordings.

Commitment-style outcomes. Willingness is measured as concrete commitments. The three B2 items name concrete acts under a common stem ("Consider software or cloud services your organization depends on. Would your organization be willing to…"): setting requirements jointly with other buyers and buying against them (B2a), entering a pooled purchase contract with other organizations (B2b), and contributing funding or staff time to the shared development or maintenance of an open-source component the firm depends on (B2c). Each is answered on a three-point scale (yes; yes, under conditions; no), and every "yes, under conditions" opens a free-text field so that the conditions become data. A slogan item ("businesses should join forces") appears nowhere in the module: agreement with it would measure acquiescence and produce an editorial. Behavioral anchoring comes first: B1 asks, before any willingness item, whether the organization has in the past five years actually set requirements jointly, bought through a pooled or framework contract, co-funded a component, or held membership in a procuring consortium, which gives every stated answer a realized-behavior benchmark.

The conditional-cooperation ladder. For the two heaviest commitments, the pooled purchase (B2b) and the co-funding of a component (B2c), item B3 asks under which condition the organization would actually commit, with five rungs: among the first; if a few comparable organizations (2–5) had committed; only if many organizations in the sector had committed; only if a public-sector anchor buyer participated; at no number. The ladder is the strategy method of Fischbacher, Gächter and Fehr (2001) translated from contribution schedules to membership thresholds. It classifies each firm as an unconditional joiner (first rung), a conditional cooperator with a stated threshold (middle rungs), or a non-joiner (last rung), which is the structure P3 needs. The public-anchor rung doubles as an instrument test: it measures demand for the launch-customer policy directly, firm by firm.

Binding-constraint elicitation. Every respondent who answers no or yes-under-conditions anywhere in B2 receives B4: the main obstacle, as a single choice with randomized option order and an open "other". The six options operationalize the candidate diagnoses of §2.1: uncertainty about what competition law allows; the time and cost of coordinating with other organizations; that others would benefit without contributing their share; requirements too specific to pool; no credible organizer; no need. Randomizing the order protects the P4 comparison from primacy effects, and the routing includes the hesitant, whose stated conditions are exactly the constraints policy could relax. B5 follows with organizer trust (an industry association, a public agency or buyer, a large peer organization, a vendor consortium, a European institution, none), so that if the missing organizer wins in B4, B5 names the candidate that could fill the role.

The within-firm criticality contrast. The sovereignty trade-off battery B6 is asked twice, once for a workload the organization considers critical and once for a non-critical one: would the organization accept, to obtain a European or sovereignty-protected alternative for this workload, a price premium of roughly 10 to 20 percent, a delay in receiving new features, a reduced feature set, or none of these (multi-select). The paired difference between the two batteries is the P5 variable, estimated within firm; a cross-firm comparison of critical against non-critical users would confound the criticality effect with everything else that differs across firms. The wording is EU-level; Bitkom's Germany-specific trade-off item (Bitkom 2026, p. 24) is the contextual reference, without a replication claim, and the premium point matches the single price point the benchmark asked.

The supply-side mirror. Provider respondents, screened in Module A's opening section, answer four mirrored items: whether they would adapt their offer to a buyer collective's joint requirements in exchange for pooled demand (S1); whether they would team up with other vendors through consortium bids, shared platforms, or joint compliance work (S2); the supply-side binding constraint, with margin pressure from pooled buyers added to the option set (S3); and a behavioral anchor on past joint bidding and shared compliance work (S4). Three firm descriptors (funding path F1, primary revenue model F2, project position F3) sit with the demographics, before any pooling item, so they cannot be colored by the pooling content. The sample's provider-user duality thus yields both sides of the pooling market from one instrument: whether the buyers who would join can find sellers who would serve, and whether the two sides are blocked by the same constraint.

7.3 The membership drive as a provision-point mechanism

The conditional-cooperation ladder is, mechanism-theoretically, an assurance-contract elicitation: Bagnoli and Lipman (1989) show that provision-point mechanisms, where commitments are refunded unless a stated threshold is met, implement efficient provision of a discrete public good in undominated equilibria. Tabarrok (1998) extends them to dominant assurance contracts, where the organizer additionally pays committed contributors if the threshold fails, making commitment a dominant strategy. Read through the model, the organizer's instrument is exactly that family: refundable letters of intent conditional on reaching the critical mass of Proposition 2 convert conditional cooperators into an assured coalition, and the elicited thresholds of the ladder measure where the provision point must sit. This turns the binding-constraint result into design guidance: where the elicited constraint is the missing organizer, the remedy has a known mechanism, and the survey measures its feasibility.

7.4 Identification: what stated commitments deliver, and what they cannot

Everything Module B elicits is stated. Stated commitments can identify, under explicit assumptions, structure: which segment is more willing than which (P1, P2), how willingness moves along the conditionality ladder (P3), which constraint a firm names when it holds back (P4), and how acceptance differs between two workloads inside one firm (P5). Every one of these estimands is a contrast, within respondent or across cells measured identically. They cannot identify levels: the share of firms that would actually sign, the realized membership of any club, or take-up under any instrument. Stated willingness overstates realized take-up; the size of the overstatement is unknown for this population, and nothing in this design measures it. The pre-registration therefore commits to contrasts as the registered quantities and bars any reading of a stated share as a market forecast.

Three devices buy internal validity within the stated-preference boundary. The conditional ladder converts an uninterpretable level into a profile: a willingness share means one thing if it is unconditional and another if it evaporates without an anchor, and the ladder observes the difference. The within-firm contrast (B6 asked twice) removes firm-level heterogeneity by construction, so any hypothetical bias that is constant within firm cancels from the P5 estimate. The behavioral anchors (B1, S4) attach every stated answer to a record of realized cooperation, so the analysis can condition on demonstrated behavior and report whether stated willingness concentrates among firms that have already cooperated once.

One caveat remains: if hypothetical bias varies across segments, conditions, or workload criticality, the contrasts themselves are contaminated, and no internal device fully rules that out. The design's defenses are the identical commitment wording across cells, the anchoring above, and disclosure: the estimates are stated-preference estimates, and every output carries that label.

7.5 Estimation and analysis shells

Estimation follows the pre-registration draft (Annex A). For P1 and P2, binary or ordered response models of the B2 outcomes on capability segment, need position, and regulatory exposure, with organization type, sector, size, and country as controls; the registered P1 contrast is the theory corner (Control or Lock-out crossed with Frontier or No-skill) against the rest, directional, on the cash-margin outcomes B2a and B2b, with B2c and the raw two-proportion cell contrasts reported alongside (B2c carries no registered direction and its cell ordering is read as a friction diagnostic per §7.1). For P3, classification of the B3 profiles and a one-sided binomial test that unconditional joiners are fewer than half of willing firms. For P4, the full B4 and S3 distributions with multinomial confidence intervals; the modal constraint is named with its margin over the runner-up, and no single-winner claim is made without a non-overlapping interval. For P5, a paired McNemar-type test on accepting any disadvantage, with the premium item reported separately. The supply-demand match (B2 against S1 and S2) is descriptive and labeled as such. The tables below are shells: every cell is to be estimated after fieldwork, and none carries a number now.

Table 3. Commitment willingness by segment cell (shell; to be estimated).

Segment cell	B2a joint requirements	B2b pooled purchase	B2c co-funding	n
Control/Lock-out × Frontier/No-skill	·	·	·	·
Control/Lock-out × Innovative	·	·	·	·
Other needs × Frontier/No-skill	·	·	·	·
Other needs × Innovative	·	·	·	·
All respondents	·	·	·	·

Table 4. Conditional-cooperation profiles (shell; to be estimated).

Profile (from B3)	Pooled purchase (B2b)	Co-funding (B2c)
Unconditional (among the first)	·	·
Conditional: a few comparable organizations (2–5)	·	·
Conditional: many in the sector	·	·
Conditional: public-sector anchor buyer	·	·
Non-joiner (at no number)	·	·

Table 5. Binding-constraint distribution (shell; to be estimated).

Constraint	Buyers (B4)	Providers (S3)
Competition-law uncertainty	·	·
Coordination time and cost	·	·
Free-riding concern	·	·
Requirements too specific or heterogeneous	·	·
No credible organizer	·	·
No need	·	·
Margin pressure from pooled buyers	n/a	·
Other (open)	·	·

7.6 Statistical power

Power is computed: the numbers come from a simulation with 2,000 replications per cell at alpha .05 (scripts, seeds, and result files in Annex B; the power-analysis note is listed in Annex A). The P1-style two-proportion contrast needs about 87 respondents per cell for 80% power under strong sorting (25% against 45%, a 20-point difference) and about 330 per cell under moderate sorting (25% against 35%, a 10-point difference). The analytic approximation and the simulation agree to within 0.019. Association distribution is expected to yield 150 to 300 respondents in total. The design consequences are drawn in advance: the 20-point scenario is within reach of a decent fieldwork outcome, the 10-point scenario is out of reach and no claim about it will be made, and the pre-registration commits to reporting achieved sensitivity (the minimum detectable difference at the realized N) with the results. The proportion tests do not depend on measurement models, which is why Module B delivers value at any realistic N. P5′s power was not simulated, and no number is claimed for it.

7.7 Sampling, disclosure, and gates

The frame is the companion study's: distribution through CNLL, APELL, OW2, and Eclipse networks, in French, English, and German. This is a specialty sample of open-source-using and open-source-providing firms. It supports within-sample structural claims, and population shares are out of its reach; the selection is also informative in its own right, since association-reachable open-source firms are the population where a buyers' club would recruit first. Linåker and Muto (2026) close their sixteen-country study by calling for prototyped and evaluated questionnaires for open-source enablement. The instrument specified here is such a questionnaire, aimed at the firm population their public-sector evidence does not cover (their clubs are municipal; this design samples firms, and the model covers both). Descriptive results feed a trilingual policy barometer for the distributing associations (skeleton listed in Annex A), with the sample's selection disclosed on every page and representative benchmarks (Bitkom 2026) cited as approximate context. The barometer's disclosure box states that percentages describe the sample and that willingness figures are stated intentions. No estimate is reported before the instrument clears the pre-registration and its validation gates; the design is pre-registered on Module B alone, and its hypothesis set does not grow because the shared survey carries other modules.

8. Policy mapping: what each constraint would indict

The mapping below is stated before fieldwork so that the reading of P4 is fixed in advance. It says which instrument each elicited constraint would indict; it advocates none of them, and it claims nothing about whether an indicted instrument would deliver, which is an efficacy question outside this design (§10). The same mapping applies to the supply side through S3, with margin pressure as the added row.

Table 6. From elicited constraint to indicated instrument (what the data will adjudicate).

Elicited constraint	Failure model	Instrument it would indict
Competition-law uncertainty	Perceived legal risk deters lawful cooperation	Guidance and legal clarity, the supervisors' announced remedy
Coordination time and cost	Transaction costs exceed the member's expected surplus	An organizer that amortizes setup: standard membership terms, an existing association or public agency as host (B5 names the trusted candidate)
Free-riding concern	The non-excludable flow invites riding (§4)	Membership design: pricing the excludable 3A layer, in-kind contribution routes, and the license levers of Proposition 6
Requirements too specific	Heterogeneity, the stall condition of cooperative purchasing	Scope discipline: pool the standardizable layer (shared components, compliance evidence) and leave differentiated needs out
No credible organizer	Coordination failure among conditional cooperators	Anchor buyers and credible organizers; Proposition 3 is the anchor's formal version, and a provision-point membership drive with refundable commitments (§7.3) is the organizer's implementable mechanism
No need	Absent demand	None
Margin pressure (providers, S3)	Monopsony concern on the supply side	Governance limits on the collective's buyer power; a heavy showing here qualifies the pooling case itself

The table is read jointly with P3: even if no constraint dominates B4, a heavy conditional-cooperation profile in B3 indicts the organizer and anchor instruments on its own, since conditionality is a demand for a coordination device. The symmetric warning stands as well: if no need dominates, the data will have adjudicated against the premise of the policy, and reporting that outcome is part of the design.

9. Verification

Every proposition carries a proof (Appendix A), a numeric cross-check, and a Monte-Carlo robustness run. The scripts, seeds, and result files that reproduce every number and figure are listed in Annex B.

- **Partition (Prop. 1).** Brute-force best responses over the four options for 20,000 sampled firms per case against the analytic partition: zero mismatches across all 95 positive-equilibrium Monte-Carlo draws.
- **Closed form (numeric cross-check).** A piecewise closed form of the uniform-type membership mass against two-dimensional brute-force decision integration on 48 parameter cases: zero mismatches. The cash benchmark ($\psi = 10^9$, $R = 0$) reproduces the

interior equilibria of the generic cash model's closed-form cubic on 27 grid cases with zero mismatches.

- **Equilibria and comparative statics (Props. 2, 3).** On 200 Monte-Carlo draws (random C_0, c, τ, R, N, ψ and independent Beta type marginals): B nondecreasing in m on all draws; the largest equilibrium weakly falls with τ and C_0 and rises with M on all draws; the cash-benchmark anchor threshold formula matches bisection on 200 of 200 tested draws.
- **Rider conversion (Prop. 4).** At every positive equilibrium tested away from the boundary, riders exist exactly when $\psi(\varphi - R) > 1$ (93 checks, zero failures); the mean in-kind share of membership across positive-equilibrium draws is 31 percent.
- **Regulation (Prop. 5).** The firm-level expansion is monotone by construction and verified inside the partition check. In equilibrium, raising R by 0.1 raises the largest equilibrium on 23 of the 200 draws, lowers it on 38, and leaves it within tolerance on 139: both signs occur, as the proposition states.
- **Welfare margins (Prop. 6).** On 60 independently seeded draws (30 positive equilibria tested; every fourth draw permissive, the rest with ρ_κ up to 0.4), total surplus is integrated from primitives and a marginal atom is injected at each testable margin with the continuum's choices frozen and the budget rebalanced: the Richardson-extrapolated numerical derivative matches $w\Phi$ on the need margin (58 tests, both dues branches) and $(w - \rho_\kappa)\Phi$ on the capability margin (6 tests, the rider-band edges where a band exists at equilibrium) with zero mismatches and worst relative error below 10^{-9}. Conversion lowers welfare on 4 of the 6 tested capability margins and raises it on 2: both reciprocity regimes occur. Raising ρ_κ from zero to 0.15 raises the largest equilibrium on 24 draws, lowers it on none, and leaves it unchanged within tolerance on 176. The inner fee-rider fixed point is unique (the map is monotone in the fee), solved by a damped iteration with a bisection fallback on a rider mass computed as a continuous band integral; the returned fee-rider pairs are consistent by construction, and the worst budget residual across the suite is below 10^{-14} of the fee range.
- **Signaling friction (Prop. 7).** On 40 independently seeded Monte-Carlo draws with $\psi(m) = \psi_0/(1 + \lambda m)$, random λ, and $\rho_\kappa = 0$: B remains nondecreasing in m, dominates the constant-ψ_0 benchmark pointwise, rider mass is pointwise weakly smaller, the rider-vanish condition holds under $\psi(m)$ wherever it holds under ψ_0, and the largest equilibrium is weakly larger, each on 40 of 40 draws.
- **Johnson boundary (Prop. 8).** On the 32 positive equilibria among 60 independently seeded draws, with brute-force best responses over 20,000 sampled firms per case: the club payoff is weakly above the voluntary payoff $\max(0, v - k - R)$ for every sampled firm, every voluntarily served firm stays served, and the strict addition matches $\{\mu(k) \le v < k + R\}$ away from ties with zero mismatches; the addition's mass is positive in all 32 cases and averages 60 percent of equilibrium membership.
- **P1 scope condition (§7.1).** At the partition figure's parameters with uniform types, the join rate by capability column is 0.92 (Innovative) against 0.76 (Frontier and No-skill) at $\psi = 2$, and 0.61 against 0.73 at $\psi = 8$: the column ordering inverts with the friction regime, while the high-need cells saturate at low friction (both at 1.00), so the registered corner contrast is confined to the cash margin.

10. What the model and the design do and do not say

The model speaks to willingness and structure on the demand side of open source: who joins, in which currency, who rides and until what scale, what an anchor does, how compliance burdens and license regimes move the margins, and why membership is undersupplied. Delivery at cost c, club governance, and the valuation of in-kind contributions at par are assumed; a club that cannot review and integrate contributed code has, in the model's terms, a high ψ, and at that point its governance problem reappears as economics. Effective reciprocity is a reduced form over license and deployment mode, and the model assigns no value to any particular license. Multiple competing clubs, dynamic membership, and the supply side's response are outside the model; the last of these is treated in the supply-side companion project. Efficacy is out of scope throughout: nothing here shows that pooling scales European suppliers, the design of Section 7 measures willingness, never outcomes, and the mapping of Section 8 indicts instruments without certifying that any of them works.

Stated preference. All willingness measures are stated, anchored by past behavior. Levels are never headline claims; contrasts are the registered quantities; and hypothetical bias that varies with the contrast dimensions escapes every internal defense (§7.4) and is disclosed wherever results appear.

Specialty sample. The frame is association-recruited open-source users and providers. Within-sample structure is estimable; population shares are out of reach, and none will be reported. The skew also bounds the external reading: low pooling willingness in the friendliest reachable population would be strong evidence against the policy's demand premise, while high willingness here licenses no extrapolation to the general firm population.

Committee decisions. Pooling decisions are committee decisions, which makes Module B less exposed to individual careerism than solo technology choices; the taxonomy and the auxiliary robustness covariates are recorded in the behavioral-agency note (Annex A).

Benchmark comparability. The only representative benchmark is German (Bitkom 2026), and its trade-off stem is Germany-specific; Module C uses EU-level wording of its own, and every comparison to the benchmark is labeled approximate context, never replication.

Cross-section. One wave. Conditionality is measured as stated structure at one point in time; the dynamics it points to (clubs forming, anchors arriving, thresholds being crossed) are outside the observation window.

11. Conclusion

Taking open source seriously turns the buyers' club from a generic pooling institution into something more specific and more hopeful. The free stock and costly flow define what a club funds; the in-kind margin lets the capable pay in code, converts free-riding into a small-club, high-friction phenomenon, and produces the two-tier composition real consortia display; compliance burdens seed micro-clubs and recruit the capable while the anchor jumps the critical mass for everyone else; the license regime is a design lever that shrinks the under-provision wedge where reciprocity reaches the deployment mode; and the club's strict contribution relative to voluntary provision lies exactly where the typology says self-provision is priced out. The club is a natural steward in the Cyber Resilience Act's sense and a standing fork threat in the supply-side model's sense. Whether firms will accept the supervisors' invitation is an empirical question with exactly this structure: collective action predicts that participation is conditional and organizer-dependent, club theory locates the excludable layer that can hold a group together around free code, and the value network predicts which firms need the collective at all. Sections 7 and 8 specify the hypotheses with their mechanisms, the instrument, the identification limits, and the mapping from elicited constraints to the instruments they would indict; the estimation waits for the poll, and until it is done the paper reports no coefficient.

Appendix A: proofs

Proposition 1. Fix $\varphi > 0$. Dues are paid in the cheaper currency, so membership costs $\mu(k) = \varphi \min(1, \psi k)$, and code is the cheaper currency iff $\psi k \leq 1$. Join dominates ride iff $v - \mu(k) \geq v - k - R$, that is iff $\mu(k) \leq k + R$. Join dominates stay iff $v \geq \mu(k)$, and ride dominates stay iff $v \geq k + R$. If $\mu(k) \leq k + R$, the choice is join versus stay, decided by $v \geq \mu(k)$, and otherwise ride versus stay, decided by $v \geq k + R$. For the rider set: on the code branch ($\psi k \leq 1$), $k + R < \psi k\varphi$ iff $k(\psi\varphi - 1) > R$, requiring $\psi\varphi > 1$ and $k > R/(\psi\varphi - 1)$, while on the cash branch, $k + R < \varphi$ iff $k < \varphi - R$. The band is nonempty iff $R/(\psi\varphi - 1) < \varphi - R$, which rearranges to $\psi(\varphi - R) > 1$. Ties lie on null boundary sets. $\square$

Proposition 2. For $\rho_\kappa = 0$ the fee $\varphi(m) = (C_0 + R)/(m + M) + c + \tau$ is strictly decreasing in m, and $\mu(k) = \varphi \min(1, \psi k)$ is strictly increasing in φ for every $k > 0$. The join event $\{\mu(k) \leq k + R, v \geq \mu(k)\}$ is nonincreasing in φ as a set, so B is nondecreasing in m for any joint type distribution (monotone set inclusion; Topkis 1998; Milgrom & Shannon 1994). As $m \to 0^+$ with $M = 0$, $\varphi \to \infty$, and the join set is contained in $\{k \leq R/(\psi\varphi - 1)\} \cup \{k \geq \varphi - R\}$, whose probability tends to zero (the second component leaves the support, the first shrinks to a null set), so $B(0^+) = 0$, and A1 rules out unilateral provision, making $m = 0$ an equilibrium. Tarski's fixed-point theorem on the complete lattice $[0, N]$ gives the equilibrium lattice; $\bar{m} = \sup\{m : B(m) \geq m\}$, and pointwise increases in B raise $\bar{m}$ by the standard argument. $\square$

Proposition 3. (i) At $\psi = \infty$, $\mu \equiv \varphi$ and the cash-benchmark argument is direct: with anchor M, $\varphi_M(0) = (C_0 + R)/M + c + \tau$ is finite, $B_M(0) > 0$ iff $\varphi_M(0) < 1$ under full support (both thresholds interior with $\min(1, 1 + R) = 1$), i.e. iff $M > M^*$, and every fixed point then satisfies $m \geq B_M(0) > 0$. Monotonicity in M is pointwise. (ii) For finite ψ and $R > 0$, at any finite fee φ the set $\{k \in (0, R/(\psi\varphi - 1))\}$ (or all low k when $\psi\varphi \leq 1$) has positive measure, and on it $\mu(k) = \psi k\varphi < k + R$ with $\mu(k)$ below the interior of the v support for k small, so the join set has positive measure. Any anchor $M > 0$ makes the fee finite at $m = 0$, hence $B_M(0) > 0$. $\square$

Proposition 4. The rider band $(R/(\psi\varphi - 1), \varphi - R)$ has both endpoints moving toward each other as φ falls: the left endpoint rises, the right endpoint falls, and the band is empty once $\psi(\varphi - R) \leq 1$ (Proposition 1). The fee is decreasing in m, bounded below by $c + \tau + (C_0 + R)/(N + M)$, so if the stated condition holds at that bound, there is a finite $\hat{m}$ with $\psi(\varphi(\hat{m}) - R) \leq 1$, beyond which no firm rides. $\square$

Proposition 5. (i) At fixed φ, dues $\mu(k)$ do not involve R while the ride payoff falls one-for-one, so the join set $\{\mu \leq k + R, v \geq \mu\}$ weakly expands and no member's dues change. (ii) In equilibrium $\partial\varphi/\partial R = 1/(m + M)$, so the need margin $v \geq \mu(k)$ tightens at rate $\min(1, \psi k)/(m + M)$ while the join-versus-ride margin $\mu(k) \leq k + R$ relaxes at rate $1 - \min(1, \psi k)/(m + M) > 0$ for $m + M > 1$, and the net movement of B depends on the densities at the two margins, and Section 9 reports both signs. $\square$

Proposition 6. The technique is direct resource accounting on a balanced allocation, one margin perturbed at a time. Fix the equilibrium classification: join set J with in-kind subset I, rider set D, and fee φ satisfying the budget identity $\varphi(\bar{m} + M) = G + (c + \tau)(\bar{m} + M)$ with $G = \max(C_0 + R - \rho_\kappa r, 0)$. Total surplus, up to the anchor's unmodeled gross value, is $W = N\int_J v\, dF + N\int_D (v - k - R)\, dF + \varphi S - G - (c + \tau)(\bar{m} + M)$: members' values enter gross and their cash dues are transfers, riders bear their own $k + R$, the club spends $G + (c + \tau)(\bar{m} + M)$ with cash-funded production purchased at cost equal to value, and each in-kind member produces work worth φ at private cost $\psi k\varphi$, so club production is cheaper than its cash funding by φS with $S = N\int_I (1 - \psi k)\, dF$. Substituting the budget identity gives the equivalent form $W = \Sigma - \varphi M$ with Σ the sum of private payoffs, which cross-checks the accounting. $S < \bar{m}$ since $1 - \psi k < 1$ on I up to a null set, so $\Phi = 1 - S/(\bar{m} + M) > 0$. Need margin: injecting mass ε at $(k_0, v_0 = \mu(k_0))$ leaves r and G unchanged and rebalances the fee to $\varphi_\varepsilon = G/(\bar{m} + \varepsilon + M) + c + \tau$, so $dW/d\varepsilon = v_0 + (1 - \psi k_0)\varphi\, \mathbf{1}\{k_0 \in I\} - (c + \tau) + S\, d\varphi/d\varepsilon$. On the cash branch $v_0 = \varphi$, on the code branch $v_0 + (1 - \psi k_0)\varphi = \varphi$, and $d\varphi/d\varepsilon = -w/(\bar{m} + M)$, so both branches give $w - Sw/(\bar{m} + M) = w\Phi$. Capability margin with $G > 0$: converting mass ε of riders at k^* with $\mu(k^*) = k^* + R$ saves their private cost $k^* + R$, adds delivery $c + \tau$, raises the gap to $G + \rho_\kappa\varepsilon$ (the lost upstream return must now be produced), rebalances the fee with $d\varphi/d\varepsilon = (\rho_\kappa - w)/(\bar{m} + M)$, and adds the production surplus $(1 - \psi k^*)\varphi$ on the code branch, where $\psi k^*\varphi = k^* + R$. Collecting terms on either rider-band edge, using $k^* + R = \varphi$ on the cash edge, $dW/d\varepsilon = (k^* + R) - (c + \tau) - \rho_\kappa + (\varphi - k^* - R)\, \mathbf{1}\{\text{code}\} + S(\rho_\kappa -$

$w)/(\bar{m}+M) = (w-\rho_\kappa)\Phi$. When the floor is slack ($\rho_\kappa r > C_0 + R$), G stays at zero under either perturbation, $\varphi = c+\tau$, $w = 0$, and the same computations give zero on both margins. The remaining claims read off these expressions: at $\rho_\kappa = 0$, A1 gives $G = C_0 + R > 0$, both margins equal $w\Phi > 0$, marginal expansion on either margin raises total surplus, and equilibrium membership is inefficiently low; for $\rho_\kappa > 0$ the capability margin is positive iff $\rho_\kappa < w$, and $w = G/(\bar{m}+M)$ is decreasing in $\rho_\kappa r$ directly. □

Proposition 7. The technique is monotone comparative statics by pointwise set inclusion (Topkis 1998; Milgrom & Shannon 1994), as in Proposition 2. (i) With $\rho_\kappa = 0$, the fee $\varphi(m) = (C_0+R)/(m+M)+c+\tau$ is positive and strictly decreasing in m, and $\min(1, \psi(m)k)$ is nonnegative and nonincreasing in m (as ψ is), so the product $\mu(k;m)$ is nonincreasing in m for every k. The join event $\{\mu(k;m) \le k+R, v \ge \mu(k;m)\}$ is then nondecreasing in m as a set, so B is nondecreasing for any joint type distribution. As $m \to 0^+$ with $M = 0$, $\varphi \to \infty$ while $\psi(m) \ge \psi(N) > 0$, so the join set is contained in $\{k \le R/(\psi(N)\varphi - 1)\} \cup \{k \ge \varphi - R\}$, whose probability tends to zero; A1 rules out unilateral provision, so $m = 0$ is an equilibrium, and Tarski's fixed-point theorem on $[0, N]$ gives the lattice. Proposition 3(ii) needs only a finite friction at the evaluated mass. Proposition 1′s band formula holds at each m with $\psi(m)$ in place of ψ, so the rider band at mass m is empty whenever $\psi(m)(\varphi(m)-R) \le 1$; both factors are nonincreasing in m (and $\varphi(m) - R < 0$ empties the band outright), so the finite conversion threshold of Proposition 4 exists whenever $\psi(N)(c+\tau+(C_0+R)/(N+M)-R) \le 1$. (ii) $\psi_A \le \psi_B$ gives $\mu_A(k;m) \le \mu_B(k;m)$ pointwise, hence the join set under ψ_B is contained in the join set under ψ_A at every m, so $B_A \ge B_B$ pointwise and $\bar{m}_A = \sup\{m : B_A(m) \ge m\} \ge \bar{m}_B$. The rider set $\{k+R < \mu\}$ shrinks as μ falls, so the rider set under ψ_A is contained in that under ψ_B at every m. Where $\psi_B(m)(\varphi(m)-R) \le 1$ with $\varphi(m) \ge R$, also $\psi_A(m)(\varphi(m)-R) \le 1$, while $\varphi(m) < R$ empties the band under both. The benchmark comparison takes $\psi_B \equiv \psi(0) \ge \psi(m) = \psi_A(m)$. □

Proposition 8. The technique is a revealed-preference payoff comparison with set inclusion. (i) The voluntary-mode payoff is $\max(0, v-k-R)$ and the club payoff is $\max(0, v-k-R, v-\mu(k))$, a maximum over a superset of options, hence weakly larger for every firm. A member has $v - \mu(k) \ge \max(0, v-k-R)$, strictly off the tie sets $\{v = \mu(k)\}$ and $\{\mu(k) = k+R\}$, which are null. For the served sets: $v \ge k+R$ gives service in the voluntary mode; under the club, if $\mu(k) \le k+R$ the firm joins (since $v \ge k+R \ge \mu(k)$), and otherwise it rides (since $v \ge k+R$), so it stays served either way. (ii) The balanced-budget identity of Proposition 6 funds $C_0 + R - \rho_\kappa r$ from dues at any positive equilibrium, so dues and rider returns together meet the flow requirement $C_0 + R$ in full. A firm is served under the club and unserved voluntarily iff it joins with $v < k+R$. Joining is then optimal iff $\mu(k) \le v$, since the ride payoff $v-k-R$ is negative while $v-\mu(k) \ge 0$, and $\mu(k) \le v < k+R$ also delivers the join condition $\mu(k) \le k+R$. The addition is therefore $\{\mu(k) \le v < k+R\}$, up to the null tie sets, and it lies outside V by construction. Positive measure for $R > 0$: for k in a right neighborhood of zero, $\mu(k) = \psi\varphi k \to 0$ while $k+R \to R > 0$, so the strip $\{(k,v) : k \in (0,\varepsilon), \mu(k) \le v < \min(k+R, 1)\}$ has positive area, and the type density is positive on the interior of the support. (With $R = 0$ the addition can vanish only on the knife-edge $\psi\varphi = 1$ with $\varphi \ge 1$, where code dues replicate the self-provision cost exactly.) (iii) If $B(m) < m$ for every $m > 0$, the only fixed point is $m = 0$, an equilibrium under A1 by Proposition 2, and at $m = 0$ every firm solves $\max(0, v-k-R)$: the voluntary mode. □

Appendix B. Item wordings: Module B and the supply-side mirror

Wordings are verbatim from the survey instrument; the instrument and the assembled questionnaire, including Module A and the context items, are listed in Annex A. Response formats are summarized with complete option lists.

Table A1. Module B (all respondents).

Item	Wording	Response format
B1	"In the past five years, has your organization taken part in any of the following?"	Multi-select: joint requirements-setting with other buyers; a pooled or framework purchase with other organizations; co-funding development or maintenance of a software component with other organizations; membership in a consortium or industry body that procures or maintains software; none of these
B2 stem	"Consider software or cloud services your organization depends on. Would your organization be willing to..."	Each sub-item: yes / yes-under-conditions / no; "yes-under-conditions" opens a short free-text condition field
B2a	"...set requirements jointly with other buyers (for example on interoperability, exit options, or data location) and buy against them?"	Per B2 stem
B2b	"...enter a pooled purchase contract with other organizations for such a service?"	Per B2 stem
B2c	"...contribute funding or staff time to the shared development or maintenance of an open-source component you depend on?"	Per B2 stem
B3 (asked for B2b and B2c separately)	"Under which condition would your organization actually commit?"	Single choice: we would commit even if we were among the first / if a few comparable organizations (2–5) had committed / only if many organizations in our sector had committed / only if a public-sector anchor buyer participated / we would not commit at any number
B4 (asked of every respondent answering no or yes-under-conditions anywhere in B2)	"What is the main obstacle to your organization joining such an effort?"	Single choice plus "other", option order randomized: uncertainty about what competition law allows / the time and cost of coordinating with other organizations / others would benefit without contributing their share / our requirements are too specific to pool / there is no credible organizer for such an effort / we see no need
B5	"Who would your organization most trust to organize a joint procurement or co-funding effort?"	Single choice: an industry association / a public agency or public buyer / a large peer organization / a vendor or vendor consortium / a European institution / none
B6 (asked twice: once for "a workload or application your	"To obtain a European or sovereignty-protected alternative for	Multi-select: a price premium of roughly 10 to 20 percent / a de-

Item	Wording	Response format
organization considers critical", once for "a non-critical one")	this workload, would your organization accept…"	lay in receiving new features / a reduced feature set / none of these

Table A2. Supply-side mirror and descriptors (provider respondents only).

Item	Wording	Response format
S1	"Would your firm adapt its offer to the joint requirements of a buyer collective (for example on interoperability, exit options, or compliance evidence) in exchange for pooled demand?"	Yes / yes-under-conditions / no
S2	"Would your firm team up with other vendors (a consortium bid, a shared platform, joint compliance work) to serve demand at a scale you cannot reach alone?"	Yes / yes-under-conditions / no
S3	"What is the main obstacle?"	Single choice plus "other": pooled buyers would press margins too hard / uncertainty about what competition law allows / coordinating with competing vendors is too costly / buyer requirements are too heterogeneous / there is no credible organizer / we see no need
S4	"Has your firm ever bid jointly with other vendors for a contract, or contributed to shared compliance work (for example CRA documentation for a common component)?"	Yes, several times / yes, once / no
F1 (funding path; descriptor, placed with the demographics)	Funding path	Single choice: self-funded or bootstrapped / venture- or private-equity-backed / subsidiary of a larger group / public or mixed ownership
F2 (primary revenue model; descriptor)	Primary revenue model	One primary, optionally one secondary: support and assurance subscriptions / custom development and adaptation / consulting, integration, and training / product licenses or dual licensing / SaaS or hosting of the product / other
F3 (project position; descriptor)	Project position	Single choice: we maintain or lead an open-source project / we contribute regularly to projects we do not lead / we build on open source without regular upstream contribution

References

- ACM, AFM, AP, DNB, & RDI. (2026). *Towards digital autonomy: Increased control over digital dependencies*. Joint position paper of five Dutch supervisory authorities.
- Bagnoli, M., & Lipman, B. L. (1989). Provision of public goods: Fully implementing the core through private contributions. *The Review of Economic Studies*, *56*(4), 583–601.
- Bergstrom, T., Blume, L., & Varian, H. (1986). On the private provision of public goods. *Journal of Public Economics*, *29*(1), 25–49.
- Bitkom. (2026). *Cloud Report 2026: Status Quo & Trends in Wirtschaft und Politik*. Bitkom Research. doi:10.64022/2026-cloud-report-2026.
- Buchanan, J. M. (1965). An economic theory of clubs. *Economica*, *32*(125), 1–14.
- CNLL & Inno³. (2024). *Cyber Resilience Act Compliance Guide for Open Source Players*.
- Cornes, R., & Sandler, T. (1996). *The Theory of Externalities, Public Goods, and Club Goods* (2nd ed.). Cambridge University Press.
- Dana, J. D. (2012). Buyer groups as strategic commitments. *Games and Economic Behavior*, *74*(2), 470–485.
- European Parliament & Council. (2024). Regulation (EU) 2024/2847 of the European Parliament and of the Council of 23 October 2024 on horizontal cybersecurity requirements for products with digital elements (Cyber Resilience Act). *Official Journal of the European Union*, L series, 2024/2847, 20 November 2024.
- Fischbacher, U., Gächter, S., & Fehr, E. (2001). Are people conditionally cooperative? Evidence from a public goods experiment. *Economics Letters*, *71*(3), 397–404.
- Ethiraj, S. K., Kale, P., Krishnan, M. S., & Singh, J. V. (2005). Where do capabilities come from and how do they matter? A study in the software services industry. *Strategic Management Journal*, *26*(1), 25–45.
- Frey, F. (2023). How we lead successful open source collaborations in the Danish public sector. *IEEE Software*, *40*(4), 19–24.
- Inderst, R., & Wey, C. (2007). Buyer power and supplier incentives. *European Economic Review*, *51*(3), 647–667.
- Johnson, J. P. (2002). Open source software: Private provision of a public good. *Journal of Economics & Management Strategy*, *11*(4), 637–662.
- Jullien, N., Stol, K.-J., & Herbsleb, J. D. (2019). A preliminary theory for open-source ecosystem microeconomics. In B. Fitzgerald, A. Mockus, & M. Zhou (Eds.), *Towards Engineering Free/Libre Open Source Software (FLOSS) Ecosystems for Impact and Sustainability* (pp. 49–68). Springer. (Preprint: arXiv:1905.05985.)
- Jullien, N., Viseur, R., & Zimmermann, J.-B. (2025). A theory of FLOSS projects and Open Source business models dynamics. *Journal of Systems and Software*, *224*, 112383. https://doi.org/10.1016/j.jss.2025.112383
- Jullien, N., & Zimmermann, J.-B. (2009). Firms' contribution to open-source software and the dominant user's skill. *European Management Review*, *6*, 130–139.
- Jullien, N., & Zimmermann, J.-B. (2011). FLOSS firms, users and communities: A viable match? *Journal of Innovation Economics*, *1*(7), 31–53.
- Katz, M. L. (1986). An analysis of cooperative research and development. *RAND Journal of Economics*, *17*(4), 527–543.
- Katz, M. L., & Shapiro, C. (1985). Network externalities, competition, and compatibility. *American Economic Review*, *75*(3), 424–440.
- Lerner, J., & Tirole, J. (2002). Some simple economics of open source. *The Journal of Industrial Economics*, *50*(2), 197–234.
- Linåker, J., & Muto, S. (2026). Advancing digital government through open source indicators. *Journal of Systems and Software*, *242*, 113031. https://doi.org/10.1016/j.jss.2026.113031

- Linåker, J., Nummelin Carlberg, A., & O'Riordan, C. (2024). *Institutionalising Support for Open Source Software in the European Public Sector*. European Commission, Directorate-General Informatics, Brussels.
- Milgrom, P., & Shannon, C. (1994). Monotone comparative statics. *Econometrica*, *62*(1), 157–180.
- Kogut, B. M., & Metiu, A. (2001). Open source software development and distributed innovation. *Oxford Review of Economic Policy*, *17*(2), 248–264.
- Nollet, J., & Beaulieu, M. (2005). Should an organisation join a purchasing group? *Supply Chain Management: An International Journal*, *10*(1), 11–17.
- Olson, M. (1965). *The Logic of Collective Action: Public Goods and the Theory of Groups*. Harvard University Press.
- O'Mahony, S. (2003). Guarding the commons: How community managed software projects protect their work. *Research Policy*, *32*(7), 1179–1198.
- O'Mahony, S., & Ferraro, F. (2007). The emergence of governance in an open source community. *Academy of Management Journal*, *50*(5), 1079–1106.
- Rohlfs, J. (1974). A theory of interdependent demand for a communications service. *Bell Journal of Economics and Management Science*, *5*(1), 16–37.
- Sandler, T., & Tschirhart, J. T. (1980). The economic theory of clubs: An evaluative survey. *Journal of Economic Literature*, *18*(4), 1481–1521.
- Schotanus, F., & Telgen, J. (2007). Developing a typology of organisational forms of cooperative purchasing. *Journal of Purchasing and Supply Management*, *13*(1), 53–68.
- Shaikh, M., & Cornford, T. (2011). *Total cost of ownership of open source software: A report for the UK Cabinet Office supported by OpenForum Europe*. London School of Economics.
- Tabarrok, A. (1998). The private provision of public goods via dominant assurance contracts. *Public Choice*, *96*(3–4), 345–362.
- Topkis, D. M. (1998). *Supermodularity and Complementarity*. Princeton University Press.
- Viseur, R., & Jullien, N. (2023). Communesplone, an original open-source model of resource pooling in the public sector. *IEEE Software*, *40*, 46–54.
- von Hippel, E. (2001). Innovation by user communities: Learning from open-source software. *MIT Sloan Management Review*, *42*(4), 82–86.
- von Hippel, E., & von Krogh, G. (2003). Open source software and the "private-collective" innovation model: Issues for organization science. *Organization Science*, *14*(2), 209–223.

Annex A: Working notes (temporary annex)

This annex lists internal working documents and will be removed from the final version.

- `../notes/common/theory-foundation.md`: the shared theoretical foundation of the project and its re-centering decisions.
- `../notes/common/poll-design.md`: the one-poll, three-module survey architecture, from which the design derives.
- `../notes/joining-forces/instrument.md`: Modules B and C, the full item set with budgets, and the supply-side mirror (the source of the Appendix B wordings).
- `../notes/common/questionnaire-en.md`: the assembled English master questionnaire, including Module A and the context items.
- `../notes/joining-forces/prereg-paper2.md`: the pre-registration draft, to be frozen after the pilot.
- `../notes/common/power-analysis.md`: computed power and quota guidance.
- `../notes/common/behavioral-agency.md`: the agency taxonomy and the auxiliary robustness covariates of Section 10.
- `../drafts/barometer-skeleton.md`: the trilingual policy-barometer skeleton of Section 7.

- `../../oss-business-models/papers/value-capture-wp.md`: the supply-side companion project's working paper behind the fork-capability reading of Section 6.
- `../notes/common/research-program.md`: the wider research program.

Annex B: Reproduction: scripts, data, and figures

- `../verify/collective_provision.py` (seed 20260716): all numeric cross-checks and Monte-Carlo runs of Section 9; writes `../verify/results/collective_provision.json` and generates the figures `figures/collective-provision-partition.png` and `figures/collective-provision-equilibria.png`.
- `../verify/power_analysis.py` (seed 20260716, 2,000 replications per cell, alpha .05): the proportion-contrast power simulation of Section 7; writes `../verify/results/power_analysis.json` and generates the figure `figures/power-curves.png`.
- `../verify/run_all.py`: reproduces every computed number and figure of the project in one command.
- The design artifacts (questionnaire and pre-registration drafts) are listed in Annex A.