Value Capture in Commercial Open Source: Services, Funding Paths, and the FLOSS Value Network (Europe, 2026)

Stefane Fermigier (Abilian) · sf@abilian.com

Draft working paper v0.4, 2026-07-16. Conceptual synthesis with a verified formal core; every proposition carries a proof, a numeric cross-check, and a Monte-Carlo robustness report; no firm-level estimate is reported.

Abstract

Most accounts of the commercial open-source sector are written about its loudest ten percent: venture-backed product vendors relicensing to fend off hyperscalers. This paper starts from the other ninety percent. It builds on the theory of FLOSS projects and open-source business-model dynamics of Jullien, Viseur and Zimmermann (2025), which theorizes that a FLOSS project's value creation lies in managing the software flow (the stream of versions) rather than the stock (the code, free), and maps which business models a project supports at each of its three phases. That theory assembles a value network in the sense of Morgan, Feller and Finnegan (2013) around the 3A services of Jullien and Zimmermann (2009) (assurance, adaptation, assistance), a four-way segmentation of needs (Control, Lock-out, Adaptable, Low-cost), and the capability typology of von Hippel (2001), Kogut and Metiu (2001) and Jullien and Zimmermann (2011) (Innovative, Frontier, No-skill).

Within that frame the paper develops one axis the theory does not treat and European practitioners live by: the firm's **funding path**. Self-funded, service-centric firms build durable businesses on the 3A services across every phase, without needing to control or capture a core. Venture-backed product vendors pursue editor, publisher, and managed-SaaS models that require holding a project's alienation rights, and it is there, and only there, that defensive licensing, hyperscaler competition, and the fixed compliance cost of the Cyber Resilience Act concentrate.

A two-period model turns the axis into theorems, with the funding path reduced to a discount parameter $\beta \in (0, 1]$ and a binding period-one revenue floor acting exactly like a reduction of effective β. Four results follow, each with a proof, a numeric cross-check, and a Monte-Carlo robustness report. Optimal contribution to the underlying project has the closed form $x^* = (\beta s A \eta/2)^2$ and rises with patience: the funding path predicts contribution intensity. Enclosure has a unique patience threshold β^* and requires immediate product revenue to beat immediate service revenue. A community able to fork shrinks the region where enclosure pays, and erosion of the appropriable service margin or slowing project growth widens it, which is when and why the visible relicensings happened. A heterogeneous population bifurcates into an open-service cluster and an enclosed-product cluster.

The paper revisits the 2018 ten-firm blueprint element by element and states, as testable and unmeasured predictions, how the 2026 European context (SaaS delivery, sovereignty procurement, the CRA) should have shifted it, then maps every claim to the survey instrument and program gates that will test it. The empirical test is specified as a design; it is not run here.

1. Introduction

The public account of commercial open source in the 2020s is a story of licensing wars: Elastic, HashiCorp, and Redis abandoning permissive licenses to block cloud resale, and the forks that followed. That story concerns a small, visible population. The economic mass of the sector sits elsewhere, in firms that sell services around software they help maintain but do not own, and that never face a relicensing decision at all. The theory this paper builds on says as much about the industry at large: the 3A services implemented in the peripheral open-source models "account for the largest share of business activity" (Jullien, Viseur & Zimmermann 2025, preprint p. 71; hereafter JVZ). This paper is about how those firms capture value, and about the one strategic choice that most shapes it.

The frame reads a FLOSS project as a value network, an organization of actors who create and capture value in the production of shared code (Morgan, Feller & Finnegan 2013), evolving through phases, and treats open-source companies as the participants who sell services complementary to that code. Section 2 walks that apparatus in detail, because the rest of the paper stands on it: the stock-flow distinction, the 3A services, the user typology, the three phases with their phase-by-model map, and alienation management. Section 2.7 then places the model in the canonical economics of open source, which explains why individuals contribute (Lerner & Tirole 2002), why firms participate (Bonaccorsi & Rossi 2003; von Hippel & von Krogh 2003), which licenses projects carry (Lerner & Tirole 2005), how the public good is underprovided (Johnson 2002), and how a project evolves against commercial competitors (Athey & Ellison 2014), and which prices none of it against the contributor's own cost of patience. Within that frame the paper adds four things. First, an axis the theory does not treat (it classifies companies by their role in the project and is silent on how they are financed): whether the firm is self-funded or venture-backed, which in the European SME reality is the first-order predictor of which business model a firm can run, how it relates to the FLOSS project, and whether it is exposed to the licensing and regulatory dynamics that dominate the discourse. Second, a formal core that makes the axis operative: a two-period model of the choice between the open service model and enclosure, proved, cross-checked, and stress-tested. Third, the 2026 European institutional layer: SaaS delivery, sovereignty-weighted procurement, and the Cyber Resilience Act, each read as a shift in the payoffs of the theory's models. Fourth, a bounding of the licensing and regulatory shocks to the population they actually govern, with a claim-by-claim map to the measurement design that will test the whole structure.

The formal core (Sections 5 to 8) is a two-period model in which the funding path enters as a discount parameter $\beta \in (0, 1]$ and a binding period-one revenue floor acts exactly like a reduction of effective β (Remark 1: a constraint with multiplier μ yields effective discount $\beta/(1 + \mu)$). Four results follow. Optimal contribution to the underlying project has the closed form $x^* = (\beta s A \eta/2)^2$, strictly increasing in patience β, service margin s, adoption base A, and contribution productivity η, so the funding path predicts contribution intensity, with quadratic force: halving effective patience cuts contribution to a quarter (Proposition 1). Under one interpretable assumption (the open second period beats the enclosed second period even at zero contribution), the value difference between staying open and enclosing rises strictly with β, so outside two corner regimes there is a unique threshold β^*: enclosure is chosen exactly when the firm is impatient ($\beta < \beta^*$) and conversion-rich ($p\lambda > s$, immediate product revenue beating immediate service revenue) at once (Proposition 2). The threshold falls with the community's fork capability and with the severity of community degradation, and rises with the product margin and the convertible fraction, which is fork deterrence stated as comparative statics (Proposition 3); it also falls with the service margin and the project's baseline growth, so erosion of the appropriable service layer and project maturity both expand the enclosure region (Proposition 4). A population heterogeneous in margins and patience bifurcates into an open-service cluster and an enclosed-product cluster (Corollary 1), the sorting the program's empirical design tests. Every result carries a written proof (Appendix), a numeric cross-check against brute force, and a Monte-Carlo robustness run (Section 7). Section 8 states which sentences are theorems, which are interpretations, and which await data.

The wider program (`../notes/common/research-program.md`) treats this paper as the supply-side half of a pair; the demand-side companion (`../../oss-adoption/`) reads the same value network from the adopter's side, meeting this paper at the user typology. Section 3 revisits the 2018 blueprint. Section 4 deepens the economics of each of the 3A services.

Section 5 sets out the model, Section 6 its results, Section 7 the verification, and Section 8 the model's boundaries. Sections 9 and 10 treat the service majority and the product minority's two shocks. Section 11 maps claims to measurement. Sections 12 and 13 discuss and conclude.

2. What this paper builds on, and what it adds

Page references to JVZ are to the preprint archived in `../../references/` (HAL hal-04963100); the published article (*Journal of Systems and Software* 224:112383) paginates differently. References to Spijkerman & Jansen (2018; hereafter S&J) use the CEUR volume's pagination (pp. 128–143).

The theory is an assembly, and its parts have distinct provenance, which this section keeps visible because the present paper adds to some of them and not others. Its own contribution, as its authors state it, is a taxonomy of FLOSS and open-source business models that "explains at what point in the life of the FLOSS project each open-source model is possible" (JVZ, p. 5), together with the theorization that value creation lies in the software flow. The value-network frame is Morgan, Feller & Finnegan's (2013), the 3A services are Jullien & Zimmermann's (2009), the capability typology descends from von Hippel (2001), Kogut & Metiu (2001), Ethiraj et al. (2005) and Jullien & Zimmermann (2011), and the stock/flow vocabulary is Shaikh & Cornford's (2011). The four-way need segmentation is JVZ's own construction on axes drawn from that sourcing literature.

2.1 Stock versus flow

The theory's starting point is that "the main value creation of FLOSS projects involves managing the software 'flow'" (JVZ, p. 6): the stream of versions that develops functionality, fixes security bugs, and adapts the software to a changing technical environment. The stock, the code "as is", is free; value creation and capacity requirement both operate on the flow. Benefiting from the flow requires participating in the production, seeing one's needs taken into account, and developing the skills to contribute or at least to understand how the solution can be adapted (JVZ, p. 6), a mechanism the theory reads as absorptive capacity. Users who lack those skills, or who find it cheaper to outsource them, "may pay open-source companies to 'participate' for them" (JVZ, p. 6). The strategic weight a buyer gives to stock and flow dimensions is itself a choice variable in the sourcing literature the theory builds on (Shaikh & Cornford 2011).

This paper treats the firm side of flow management as an investment choice: the model of Section 5 makes it the choice variable x (contribution intensity) with a closed-form optimum, and Proposition 1 shows that the funding path shifts it.

2.2 The 3A services

Open-source companies sell "3A" services, "assurance, adaptation, and assistance or support for use", which are "complementary to the access to the software" (JVZ, p. 1). The term is Jullien & Zimmermann's (2009), and JVZ carry it forward with the caveat that the value creation it names "is not really new" (JVZ, p. 56); what the 2025 theory adds is which services a company can sell at which phase of a project. The theory's synthesis makes each precise (JVZ, p. 56): assistance is help with needs specification and with use; adaptation is the adaptation "of software to the organization, and the organization to the software"; assurance is "access to a maintained/assured technical capacity", resting on the "mastery of a software factory", a combination of technical tools and human skills. In European terms: assurance is the guaranteed, security-maintained edition of an open-source infrastructure component sold under SLA to a public administration; adaptation is the integration of an open-source ERP into a mid-market manufacturer's processes, with the generic modules contributed back upstream; assistance is the specification, training, and support that carry a municipality's move onto an open-source content platform (for a documented public-sector pooling case, see Viseur & Jullien 2023). Providing these services credibly requires participating in the FLOSS project, which in turn supplies the project with the resources to sustain its flow (JVZ, p. 1); value capture and contribution are two sides of one relationship (Jullien & Zimmermann 2009), and the community functions as the complementary asset of Teece (1986) in the specific, participation-built form documented for open source (Dahlander & Wallin 2006).

This paper conditions the 3A mix on the funding path and on the phase of the projects a firm works on, and gives each service a conceptual cost structure (Section 4).

2.3 The user typology

The theory segments users on two dimensions. By need: the requirement's importance on the upgradability axis (must it evolve over time?) and the adaptability axis (must it fit specific needs?) yields four positions, Low-cost (standard and cheap), Adaptable (personalized at best cost), Lock-out (standard, durable, avoiding lock-in), and Control (personalized, perennial, adaptable over time) (JVZ, pp. 10–12, Figure 1). By capability: the capacity to translate needs into technological requirements and the capacity to develop them in code (Ethiraj et al. 2005) split users into Innovative-users (both), Frontier-users (specify but do not implement), and No-skill-users (neither) (JVZ, pp. 13–14, Table 1). This second typology is the older of the two and is not new in JVZ, who name von Hippel (2001) for the innovative user and Kogut & Metiu (2001) for the frontier user, and base the three-way table on Ethiraj et al. (2005), Jullien & Zimmermann (2011) and Viseur & Jullien (2023). Open-source companies are the users who sell these capacities to others: they "have to develop the same participating capacities as an Innovative-user" (JVZ, p. 70), and they "can provide services of specifying the needs and evaluating the (FLOSS) solutions according to these needs" for No-skill users, and develop needs and get them accepted by the project for all users (JVZ, p. 15).

This paper uses the typology as the demand map for the 3A economics of Section 4; the sibling adoption project reads the same typology from the buyer's side.

2.4 The three phases and the phase-by-model map

A FLOSS project's collaboration mechanisms are built in three phases: creation, development (flourishing or decline), and maturity (JVZ, p. 23). In Phase 1, Innovative-users with Control needs initiate a project to co-develop a solution they use themselves; the matching open-source business model is customized software development inside the project, billed as development hours (JVZ, pp. 26–27). "Opening up the project beyond the closed club of initiators marks the transition from Phase 1 to Phase 2" (JVZ, p. 28). In Phase 2 the project structures open collaboration, and the editor role, coordinating and selecting contributions and managing the roadmap, exhibits increasing returns that explain why one actor usually holds it (JVZ, p. 32). The editor-family models are the foundation (participant-neutral management of the project, financed by membership fees; JVZ, pp. 33–34), the ecosystem manager (certification of modules and of the actors able to implement them; JVZ, p. 35), and the publisher (guaranteed, assured editions financed by "assurance subscriptions relating to the evolution of the software"; JVZ, p. 35; Riehle 2021). Peripheral models grow on the modular architecture: the business-solution consultant selling need specification and delivery, and the business-module publisher selling assurance on modules (JVZ, pp. 37–39). The theory condenses this into two capacity families with different scale economics: models "based on human skills and the sale of expertise time, with little economies of scale", and models selling maintained technical solutions with "more room for economies of scale" (JVZ, p. 38). In Phase 3 the intellectual appeal and the contributor inflow decline, and the project's problem becomes organizing the financing of core maintenance (JVZ, pp. 40–42); the matching models are open-source outsourcing (facility management of FLOSS-based infrastructure; JVZ, p. 44), the SaaS publisher (a hosted standard solution earning "significant economies of scale thanks to the extreme standardization of the source code deployed and the machines managed"; JVZ, p. 46), the cloud platform (JVZ, p. 48), and the business-solution publisher (verticalized stand-alone solutions on a stable core; JVZ, pp. 50–

51). The theory also documents the failure mode: owners tempted to stop maintaining the open version at maturity, with the warning that "the decline in investment in community management accelerates the exodus of contributors" (JVZ, p. 41).

The theory maps which models are possible when; this paper adds who can run them: the funding path acts as a selection on the map, and the 2026 European layer (procurement, delivery, the CRA) shifts the payoffs of specific cells (Sections 3 and 10).

2.5 Alienation management

The theory's power table ranks what each role in a project controls (JVZ, p. 22, Table 3): everyone has access to the software stock; contributors also reach the production system; core contributors and package managers control the software flow; owner and core contributors manage the project's evolution; and at the top sits alienation management, held only by the owner, alienation "in the legal sense of the term, i.e. the transfer of intellectual property rights": the software license, contributor contracts, the brand. Alienation rights are disciplined from below: if Innovative-users feel their needs are not being respected by a player seen as too invasive, "they may 'fork' the project or create a new one" (JVZ, p. 65). The theory records real Phase-3 exercises of these rights and their cost, including an ERP publisher that stopped maintaining its FLOSS version in 2010 and the abandonment of CentOS Linux (JVZ, p. 41).

This paper reads the 2020s relicensings as owners exercising alienation rights against their own value network, and the model prices the exercise: the fork option enters as a probability ϕ and shrinks the parameter region where enclosure pays (Sections 5 and 6).

2.6 The two works the theory develops, and where this paper's results meet them

JVZ position their article as developing the "preliminary propositions" of Jullien, Stol & Herbsleb (2019) and the findings of Lin & Maruping (2022). Both sit close enough to this paper's results to be stated explicitly, because parts of what follows have qualitative antecedents there.

Jullien, Stol & Herbsleb (2019) sketch an ecosystem microeconomics as eighteen propositions across the same three phases the 2025 theory later develops. Three of them touch this paper's model. Their Proposition 12 holds that the more central an open-source project is to a company's business, the more that company contributes; the closed form of Proposition 1 below restates that comparative static, since x^* rises in the adoption base A and the service margin s, which are centrality in this model's terms. What the present paper adds on top of it is the patience term β: two firms identical in centrality contribute differently when they are financed differently, and the dependence is quadratic. Their Propositions 5 and 6 treat the decision to open-source a component as a function of its criticality to the core business and its potential for further evolution, with the open-core arrangement (keep the strategic parts closed, open the standard part) as the resolution when a critical component is involved. That is the mirror image of the move analysed here, which is the enclosure of a project already open, and the two meet at the same appropriability tension from opposite sides. Their Proposition 17, that companies stop investing in a project once they no longer perceive business value in it, is the qualitative form of Proposition 4′s comparative static on the appropriable service margin. The value added here is the formal apparatus rather than the direction of any of these effects: a closed form, an exact threshold, and a stated failure region.

Lin & Maruping (2022) supply the closest empirical antecedent. On a matched panel of 17,552 digital startups observed monthly between 2008 and 2017, they find the value effect of open-source engagement contingent on venture maturity and on the mode of engagement: startups in the conception and commercialization stages benefit more from inbound engagement, those in the growth stage from outbound engagement. Two things follow for this paper. Their outbound engagement is this model's contribution intensity, and their result is a stage-contingent version of the claim that contribution pays where the base to earn on is large. And their population is the venture-track slice, which is exactly the minority this paper bounds the licensing and compliance shocks to, so their evidence characterizes the enclosed-product cluster of Corollary 1 rather than the service majority. A test of the sorting prediction on data of their kind would therefore have to sample outside the startup population to see both clusters, which is the design constraint Section 11 records.

2.7 The canonical economics of open source, and where this model sits

The value-network theory is one branch of a literature with a founding frame, and the model of Section 5 sits inside that frame rather than beside it.

Lerner & Tirole (2002) asked why programmers contribute to code they give away, and answered at the individual level with signaling and career concerns paid out in delayed private returns. Their firm-level observation, that commercial participants earn on complements to the free code, is the ancestor of the service margin s here, as is the survey of firm motivations in Bonaccorsi & Rossi (2003). What this paper adds at that level is the intertemporal price of participating: the return on contribution arrives a period after its cost, so the firm's discount factor decides how much of it is bought.

Lerner & Tirole (2005) is the closest antecedent to the enclosure result. Studying license choice across nearly 40,000 SourceForge projects, they find corporate-developed and end-user-oriented projects carrying more restrictive licenses, and permissively licensed projects attracting more contributors. Two differences place the present model. Their object is the license chosen at the outset, while this one is an incumbent owner's decision to relicense a project already open, where what stands to be lost is a community built by past contribution. And their empirical regularity, that permissive licensing attracts contributors, is the reduced form of what Proposition 3 prices: fork capability deters enclosure, so a licensor's control is bounded by the community's exit option.

von Hippel & von Krogh (2003) name the mechanism the 3A services instantiate, contributors giving away the innovation while capturing private benefits from the process of producing it. Here the private benefit is commercial and the contributor is a firm, so the private-collective logic runs through an appropriable service margin rather than through learning and reputation.

Johnson (2002) supplies the underprovision benchmark: with a public good and private contributors, free-riding leaves the flow undersupplied. That is the background condition for commercial participation to matter at all, and it is the benchmark the companion buyers'-club paper works against directly. This paper's model has a single firm choosing its own contribution and aggregates nothing, so Johnson's result bounds the reading of Proposition 1 rather than entering it.

Athey & Ellison (2014) is the canonical dynamic treatment, with reciprocal altruism driving the joint evolution of quality, contributors, and user support, and with commercial firms competing against the project on price. This paper's dynamics run on the other side of that boundary: the firm participates in the project instead of competing with it, its motive is appropriable rather than altruistic, and its two periods carry a financing constraint rather than a stock of accumulated quality.

Three strategy papers sit closer still to specific results here. Casadesus-Masanell & Llanes (2011) model the firm's choice of how much of a product to open, which is the open-core arrangement as a design problem; this paper's object is an owner reversing a commitment already made, where what is at stake is a community built under the old one. Zhu & Liu (2018) document a platform owner entering its complementors' spaces on Amazon data, which is the hyperscaler squeeze of Section 10.1 arriving from above rather than from the licensor. Nagle (2018) is the closest empirical antecedent to Proposition 1′s motive, finding that firms contributing to crowdsourced public goods gain competitive advantage from doing so; he measures the return to contributing, and this model prices which firms can afford to wait for it.

Stated once, the delta against this literature is the same as the delta against the value-network theory. It explains why individuals contribute, why firms participate, and which licenses projects carry. It does not price the commercial contributor's intertemporal problem with the firm's own cost of patience as an argument, and the funding path is absent from it.

3. The 2018 blueprint, revisited

Spijkerman and Jansen (2018) built the reference picture of B2B open-source business models from ten comparative case studies, interviewing practitioners on the elements of Schief's Software Business Model Framework and modeling the stakeholder value exchanges in e3-value (S&J, pp. 130–133). The sample, anonymized, spans foundings from 1992 to 2016 and sizes from under fifty to over ten thousand employees, from project-management tools and a CMS to integration platforms and two Linux distributions (S&J, p. 132, Table 1). The authors flag the convenience sample and the external-validity limits themselves (S&J, p. 131); the blueprint is a status picture of successful firms.

Two of its findings anchor this program. First, the sample is service-centric: "for all companies the revenue source is directly paying customers" (S&J, p. 135), the majority offer recurring subscription-style payment and layered support contracts, and since open-source companies cannot charge for the software itself, "business models of OSS companies are primarily built around software related services" (S&J, p. 136). Second, the funding-path signal is already there: none of the interviewees named investment as a prerequisite for success (S&J, p. 137), venture capital is described as "often not interested in an OSS product" (S&J, p. 138), and interviewee J states the irreversibility that the model of Section 5 formalizes as community-asset degradation: "you can make things open source only one time" (S&J, p. 138). The study's headline conclusion is that "it is not harder to establish a business around OSS than around a proprietary product" (S&J, p. 141).

The table below sets the blueprint's observed patterns against what the 2026 European context predicts has shifted. Every entry in the two right-hand columns is a hypothesis: stated to be testable against the 2018 baseline as a distributional shift, and unmeasured here. Instrument items refer to `../notes/value-capture/instrument.md`.

Blueprint element	2018 pattern (S&J)	2026 European prediction (unmeasured)	Instrument
Value proposition	Quality: "high consistency and dependability" (p. 134); quality heads the blueprint (p. 140)	Quality persists; sovereignty and auditability enter as stated value propositions for public and regulated buyers	Q4, Q11, Q12
Cost structure	R&D dominant: eight of ten firms (p. 135)	R&D still leads; cloud operations gain share where delivery is managed SaaS; CRA compliance appears as a fixed cost line for product vendors	Q5, Q10
Sales channel	Sales agents (p. 140), with heavy online marketing (p. 136)	Partner and integrator channels and cloud marketplaces gain share; direct outbound concentrates in the venture-backed segment	Q6
Delivery	Hybrid of on-premise and on-demand most chosen (pp. 136, 140)	Fully managed SaaS gains share overall; sovereignty procurement sustains self-hosted and hybrid delivery in public and regulated segments	Q3, Q11
Growth model	Growth investment horizon for most firms (p. 134)	Bifurcation by funding path: growth targets in the venture-	Q2, Q13

Blueprint element	2018 pattern (S&J)	2026 European prediction (unmeasured)	Instrument
		backed cluster, margin and durability objectives in the self-funded cluster	
Customer segment	Large organizations, above 250 employees (pp. 136, 139)	Persists; public-sector weight higher in the European frame	Q11
Licensing	Copyleft-dominant: six of ten GPL, four dual-licensed (p. 135)	Stable for the service majority; source-available transitions concentrated in the venture-backed product subsample	Q7, Q8

Read as a whole, the 2018 baseline already looks like the service-centric majority of this paper, and the predicted shifts are of two kinds. Delivery and channel shifts (SaaS, marketplaces, partners) are industry-wide and should show up across both populations. The licensing and compliance shifts should be concentrated: if the funding-path reading is right, the license column moves only in the venture-backed product subsample, and the CRA cost line appears asymmetrically (Section 10). A uniform shift across both populations on those two rows would count against the axis.

4. The economics of the 3A services

The theory's two capacity families (JVZ, p. 38) give the 3A services their cost structures: expertise sold as time, with little scale economy, and maintained technical capacity sold as subscription, with much more. The treatment is conceptual throughout; no magnitudes are stated.

4.1 Assurance: selling the guaranteed flow

Assurance sells access to a maintained/assured technical capacity (JVZ, p. 56): a guaranteed edition, a security-maintenance commitment, an SLA, a certification. The revenue form is the assurance subscription tied to the software's evolution (JVZ, p. 35), which prices the flow itself. The cost structure is dominated by fixed capacity: the software factory (build and test infrastructure, non-regression bases, vulnerability handling, release management) costs roughly the same whether it serves ten clients or a thousand, so the marginal cost of an additional subscriber is low and the economics reward scale, through the automated integration of large module sets in the publisher case (JVZ, p. 35) and through extreme standardization in the SaaS case (JVZ, p. 46). Accumulated private assets sharpen the position: the brand, and knowledge bases such as bug and needs databases that are kept up to date and confidential (JVZ, p. 56). The natural buyers are Frontier-users and No-skill-users with Lock-out, Adaptable, or Low-cost needs, who want the flow's benefits without building flow-management capacity (JVZ, p. 36, Table 7). The phase fit is flourishing and maturity: assurance presupposes a stable core worth guaranteeing, and the SaaS variant is a Phase-3 model (JVZ, pp. 46–47).

4.2 Adaptation: selling need translation and integration

The adaptation service sells the work of fitting software "to the organization, and the organization to the software" (JVZ, p. 56): customization, integration into an existing information system, and module development. The cost structure is the mirror image of assurance: mostly marginal, since delivery consumes expert time client by client, with little economies of scale and some learning effects (JVZ, pp. 37–38); the compensating property is that revenue arrives with the first engagement and capacity grows with headcount, so the model is financeable out of operating cash flow. The buyers are Control-need Innovative-users who outsource implementation while keeping direction, and Adaptable-need Frontier-users who can specify but cannot build (JVZ, pp. 37–39); module adaptation extends the offer to No-skill users through the modular architecture. The phase fit is the widest of the three: customization is the Phase-1 model itself (JVZ, pp. 26–27), continues into Phase 2 for complex accounts, and reappears at maturity as verticalized business solutions on a stable core (JVZ, pp. 50–51).

4.3 Assistance: selling use
Assistance sells help with needs specification and support for use (JVZ, p. 56): helping a buyer state what it needs, select and evaluate solutions, train users, and resolve operating problems. For No-skill users this includes the full service of specifying needs and evaluating FLOSS solutions against them (JVZ, p. 15); for Frontier-users it is help at the border of the project. The cost structure is human time plus a mostly fixed knowledge base that accumulates as a joint product of participation: the same engineers who follow the project's flow answer support questions, which is why credible assistance is hard to sell without participation. Scale economies are mild; the service travels well in subscription form when bundled with assurance. The phase fit is broad, thickening as the user base widens beyond the initiating Innovative-users, in late flourishing and maturity, where Adaptable and Low-cost segments dominate the clientele (JVZ, pp. 54–55, Table 14).

The split matters for the funding-path axis because the fixed and marginal structures select different financing. Adaptation and assistance produce cash from the first client and scale with hires, so they are feasible on retained earnings; that is the mechanism behind the self-funded service majority. Assurance at scale requires the factory to exist before the subscribers do, so it is financed by capital, by accumulated surplus, or by mutualizing the fixed cost across a project's ecosystem through a foundation. This is a conceptual mapping from cost structure to capital structure; the 3A revenue-mix item in the instrument measures whether the sector actually sorts this way.

5. The model

The axis of the preceding sections is verbal. This section and the next reduce it to one behavioral primitive and derive its consequences inside a model small enough to prove things about. The community enters as the complementary asset of Section 2.2, and enclosure severs the participation relationship that sustains it.

Two periods, $t = 1, 2$. A firm commercializes around an open-source project with initial adoption base $A > 0$ and discounts period 2 by $\beta \in (0, 1]$. The parameter encodes the funding path: a firm financing itself from revenue has β near one, while a venture-backed firm facing fund horizons and growth targets behaves as a low-β firm, in the following exact sense.

Remark 1 (revenue deadlines as impatience). *(i) On the contribution margin the equivalence is exact. Within OPEN, a binding period-one revenue floor $R_1 \geq \bar{R}$ yields $x^c = sA - \bar{R}$, which is the unconstrained optimum of a firm with effective discount $\beta_{\text{eff}} = 2\sqrt{sA - \bar{R}}/(sA\eta) < \beta$. (ii) On the strategy margin it is directional. Against no floor, a floor that binds under OPEN and not under ENCLOSE weakly enlarges the enclosure region, since it lowers the open branch's value and leaves the enclosed branch's untouched. Against the as-if firm at β_{eff}, the constrained firm still discounts the enclosed branch at β, so under A1 it weights the open second period more and is strictly the more open-leaning of the two: substituting β_{eff} for the constraint overstates the pull toward enclosure. And a floor above the open branch's feasible period-one revenue ($\bar{R} > sA$) forces ENCLOSE at every β, which no reduction of β reproduces.* Growth targets and runway deadlines therefore enter the model as reductions of effective patience on the margin the propositions below use, and as a weak push toward enclosure on the margin they do not.

At $t = 1$ the firm commits to one of two strategies.

OPEN (service model). The firm sells the 3A services at margin $s > 0$ per unit of base and period, and chooses a contribution intensity $x \geq 0$ at cost x: code, maintenance, and community work invested in the project. Contribution grows the base, $A_2 = (1 + g_0 +$

$\eta\sqrt{x})A$, where $g_0 \geq 0$ is the project's baseline growth and $\eta > 0$ the productivity of the firm's contribution; the square root captures decreasing returns. The firm's value is

$$V_{\text{open}}(\beta) = \max_{x \geq 0} sA - x + \beta\, sA\,(1 + g_0 + \eta\sqrt{x}).$$

ENCLOSE (product capture). The firm relicenses and converts a fraction $\lambda \in (0,1]$ of the base into product customers at margin $p > s$. The move breaks the reciprocity that sustains the community, so the complementary asset degrades: the retained base is $A_2 = \delta A$ with $\delta < 1$, and the contribution and growth channel is dead ($x = 0$). With probability $\phi \in [0,1)$ the community is capable of forking, and the fork erodes the locked base in period 2. Expected value is

$$V_{\text{enc}}(\beta, \phi) = p\lambda A + \beta\, p\lambda\, \delta(1 - \phi)A.$$

The difference $D(\beta) = V_{\text{open}}(\beta) - V_{\text{enc}}(\beta, \phi)$ determines the choice: the firm encloses exactly when $D(\beta) < 0$.

Assumption A1 (the community asset matters). $s(1 + g_0) > p\lambda\delta(1 - \phi)$.

Per unit of initial base, the left side is the open strategy's period-2 revenue at zero contribution and the right side is the enclosed strategy's expected period-2 revenue. A1 says the open second period beats the enclosed second period even when the firm invests nothing in the project: enclosure buys period-1 conversion and sacrifices community growth, an undegraded base, and immunity to forks. The sacrificed assets are worth more later. A1 fails only where conversion is broad and rich, degradation mild, and the community unable to fork; Section 7 samples that region too and reports what still holds there.

6. Results

Proposition 1 (contribution). *The open firm's optimal contribution is interior, with closed form*

$$x^*(\beta) = \left(\frac{\beta\, sA\eta}{2}\right)^2$$

from the first-order condition $\beta sA\eta/(2\sqrt{x}) = 1$, and x^* is strictly increasing in β, s, A, and η.

This is the funding-path axis made formal: the funding path predicts contribution intensity. Contribution is an investment whose return, a larger base earning the service margin, arrives one period later, so its optimal level scales with the square of patience; halving effective β cuts contribution to a quarter. The same closed form says who contributes most at a given patience: firms with more base (high A), a better service margin on that base (high s), and more productive engineering inside the project (high η). Self-funded service firms, the patient population of Section 9, are the model's natural maintainers; a venture-backed firm running the same service model under a revenue deadline contributes less on the identical asset. Value capture and contribution move together, the coupling documented by Jullien and Zimmermann (2009).

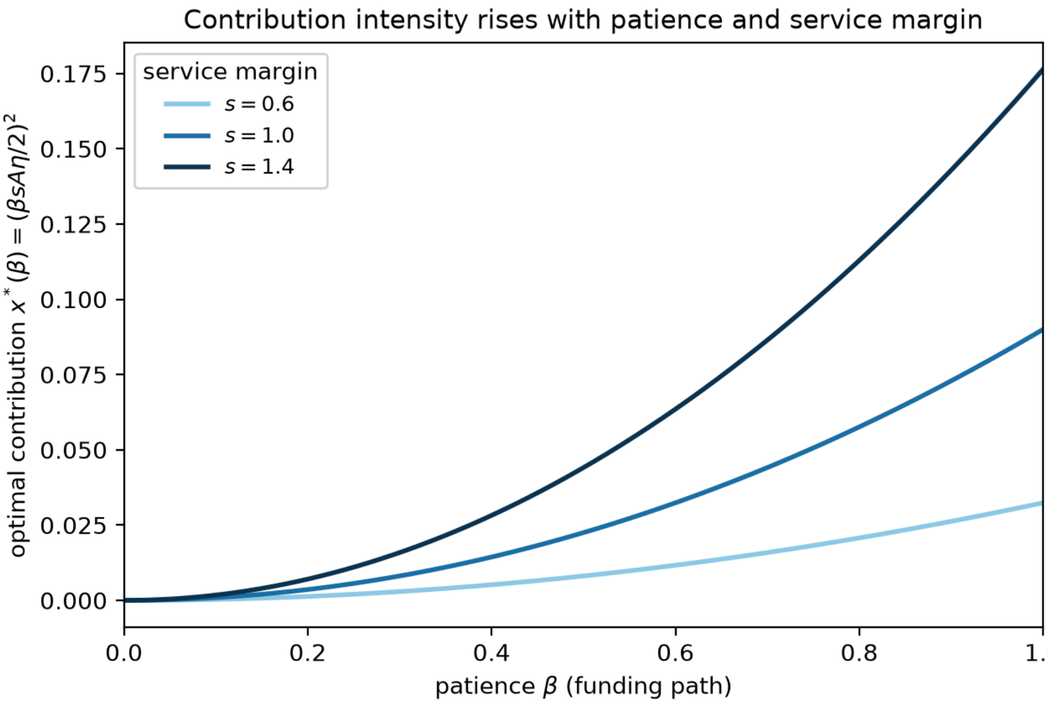

Proposition 2 (enclosure threshold). *Evaluate V_{open} at $x = x^*(\beta)$. Under A1, D is strictly increasing on $(0, 1]$. Consequently: (i) if $s \geq p\lambda$, then $D > 0$ on $(0, 1]$ and OPEN is optimal for every β. (ii) If $D(1) \leq 0$, then ENCLOSE is optimal for every $\beta \in (0, 1)$ and weakly optimal at $\beta = 1$. (iii) If $D(0) < 0 < D(1)$, there is a unique $\beta^* \in (0, 1)$ with $D(\beta^*) = 0$, and ENCLOSE is optimal exactly when $\beta < \beta^*$. Monotonicity of D uses the envelope theorem on $x^*(\beta)$.*

The threshold organizes the funding-path axis. Enclosure requires two things at once: $p\lambda > s$, so that converting the base to a product beats serving it today, and $\beta < \beta^*$, so that the firm discounts what enclosure destroys tomorrow. Impatience is the operative condition: the venture-backed, deadline-driven firm weighs the period-1 conversion heavily and the degraded, fork-exposed period 2 lightly, while a patient firm internalizes the community asset's future return and stays open even when the product margin is rich. Outside the corner regime (ii), where conversion economics dominate at every patience level, patient service firms never enclose. The Monte Carlo of Section 7 reports how often each regime arises in a broad parameter box.

Remark 2 (single crossing without A1). *D is a convex quadratic in β: $D = c_0 + c_1\beta + c_2\beta^2$ with $c_0 = A(s - p\lambda)$, $c_1 = A(s(1 + g_0) - p\lambda\delta(1 - \phi))$, $c_2 = (sA\eta)^2/4 > 0$, and A1 is $c_1 > 0$. When A1 fails, $p\lambda > s$ necessarily, so $c_0 < 0$ and the sublevel set $\{D \leq 0\}$ is an interval containing 0: the enclosure region keeps the form $[0, \beta^*)$ even where D dips before rising. A1 buys the monotonicity of D and the signed comparative statics of Proposition 3; single crossing is global in this specification (Appendix).*

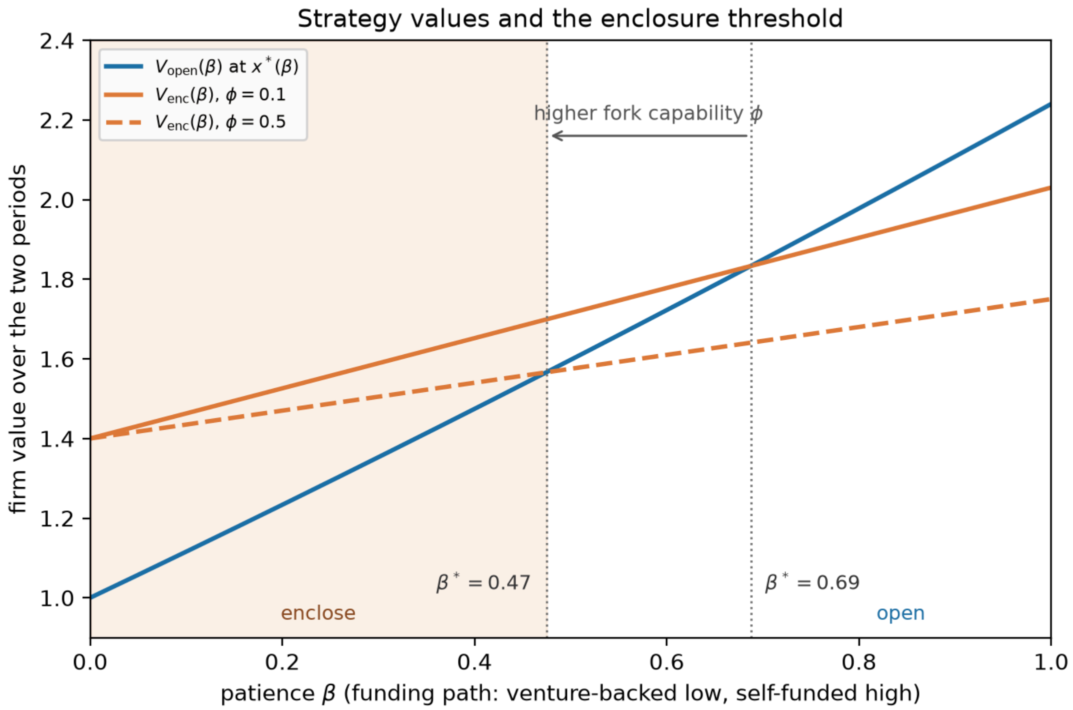

Proposition 3 (fork deterrence). *At an interior threshold, β^* is strictly decreasing in ϕ, strictly increasing in δ, and strictly increasing in p and λ. Equivalently, the enclosure region $[0, \beta^*)$ shrinks with the community's fork capability and with the severity of degradation $1 - \delta$, and expands with the product margin and the convertible fraction. The proof is the implicit function theorem on $D(\beta^*) = 0$, using $D' > 0$ from Proposition 2.*

The comparative statics are the alienation-management story of Section 2.5 run in reverse: relicensing is the owner exercising alienation rights, and the community's capacity to reconstitute the project elsewhere is a tax on that exercise, paid in expected period-2 revenue. A capable community (high ϕ) deters enclosure among the marginal firms, exactly those with β just below the old threshold; the 2020s relicensings drew forks where communities were capable, which the model reads as realized draws of ϕ among firms impatient enough to enclose despite it. Degradation works from the other side: the more the complementary asset collapses on enclosure (low δ), the smaller the enclosure region. Rich conversion economics (p, λ) pull the other way and expand it. The model states deterrence as comparative statics; it assigns no probability to any particular fork.

Proposition 4 (margin erosion and maturity). *At an interior threshold, β^* is strictly decreasing in the service margin s and in the baseline growth g_0. Equivalently, the enclosure region $[0, \beta^*)$ expands when the appropriable service margin erodes and when the project's autonomous growth slows.*

The proposition supplies the timing the cases exhibit, as comparative statics. The trigger of the relicensing wave was a competitor capturing the service layer: when a hyperscaler's managed offering absorbs the revenue that the firm's 3A services earned on the base, s falls, both periods of the open value fall with it, and firms whose patience sat just above the old threshold cross into enclosure. Maturity works the same way from the growth side: a project whose autonomous growth g_0 has flattened offers less second-period reward to staying open, so the enclosure region widens exactly at the life-cycle phase where the visible relicensings occurred. Neither force touches the enclosure payoff; both work by cheapening what enclosure destroys.

Corollary 2 (endogenous fork capability). *Let fork capability increase with the innovative share of the project's user base, $\phi = \phi(\sigma_I)$ with $\phi' > 0$. Then β^* is decreasing in σ_I: enclosure is deterred on projects whose users can develop, and undeterred on projects whose users cannot.*

This connects the supply side to the user typology of Section 2.3: the same Innovative users who self-serve on the adoption side are the community that can reconstitute a relicensed project, and their presence disciplines the owner. Developer-facing infrastructure carries its deterrent inside its own user base; end-user-facing applications with No-skill populations do not.

Corollary 1 (bifurcation). *Let firms be heterogeneous in (s, p, β), and possibly in the remaining parameters, each choosing by Proposition 2. The population splits into an open-service cluster (high β, or s/p high enough that the enclosure region is empty) and an enclosed-product cluster (low β together with high $p\lambda$). The statement is Proposition 2 applied firm by firm; no further proof is needed.*

This is the model's testable sorting prediction, and the shape the empirical program's cluster analysis and fsQCA are designed to detect (Section 11): funding path and the service-versus-product margin structure should jointly separate the two populations, with contribution intensity tracking β inside the open cluster (Proposition 1). A continuum of intermediate strategies would falsify the model's discrete-choice structure; a split without the β gradient would falsify the funding-path reading.

7. Verification

Every proposition carries a proof (Appendix), a numeric cross-check, and a Monte-Carlo robustness run. The suite is `../verify/funding_path.py` (seed 20260716); it regenerates both figures and writes `../verify/results/funding_path.json`. All numbers below are from that run.

- **Contribution (Prop. 1).** The closed form x^* against brute-force maximization of V_{open} over x (grid bracket plus bounded refinement) on 108 (β, s, A, η) grid cases: maximum relative error 2.7×10^{-8}. In the Monte Carlo, the same brute force at a random β on each A1-satisfying draw agrees to 3.7×10^{-8}, and all 1468 upward bumps of β, s, A, η raise the brute-forced optimum.
- **Threshold (Prop. 2).** β^* computed three ways (bisection on D, the stable quadratic root, a 20001-point scan of the argmax strategy over β) on the 360 cases of a 384-case $(s, p, \lambda, \delta, g_0, \eta, \phi, A)$ grid that satisfy A1: zero disagreements, and the scan boundary sits within 2.5×10^{-5} of the bisection value (half the scan step). Regimes on the grid: OPEN dominant on 192 cases, ENCLOSE dominant on 37, interior threshold on 131. In the Monte Carlo (400 draws; 379 satisfy A1), the enclosure set is an initial segment $[0, \beta^*)$, empty, or everything on all 379 draws (173 open-dominant, 60 enclose-dominant, 146 interior); no draw shows more than one sign change of D.
- **Fork deterrence (Prop. 3).** On the 146 interior draws, bumping each parameter while keeping A1: β^* falls in ϕ on 146 of 146 checks, and rises in δ on 145 of 145, in p on 145 of 145, and in λ on 145 of 145 (a bumped case is skipped when the bump breaks A1).
- **Margin erosion and maturity (Prop. 4).** On the same interior draws, β^* falls under an upward bump of s on 146 of 146 checks and under an upward bump of g_0 on 146 of 146 (bumps that break A1 or $p > s$ are skipped). Corollary 2 inherits the ϕ checks through the chain rule and receives no separate run.
- **Revenue floors (Remark 1), in both directions.** On 18,225 A1-satisfying draws with a binding, feasible floor: part (i) is exact, the constrained contribution x^c matching $x^*(\beta_{\text{eff}})$ to 2.7×10^{-15}. Part (ii) is measured rather than assumed. The constrained firm at true β and the as-if firm at β_{eff} rank OPEN against ENCLOSE differently on 6.6 percent of those draws, which is the size of the gap the old Lagrangian argument papered over, and the constrained firm is the more open-leaning of the two on every draw, as A1 requires. The check is `check_remark1`.
- **A1 violations (Remark 2).** 21 of 400 draws violate A1. On all 21, $p\lambda > s$ holds, so $D(0) < 0$ as Remark 2 requires. D is non-monotone (dips before rising) on 14 draws

and monotone decreasing on 7, so the assumption does real work; the sign pattern is nonetheless single crossing on all 21 (20 enclose-everywhere, 1 interior upcrossing, none open-everywhere, none with two crossings). In the violating region the model keeps the interval structure of the enclosure region and loses the monotonicity of D.

- **Figures.** Parameters chosen for legibility of the crossing: at $s = 1$, $p = 1.75$, $\lambda = 0.8$, $\delta = 0.5$, $g_0 = 0.15$, $\eta = 0.6$, $A = 1$, the threshold is $\beta^* = 0.687$ at $\phi = 0.1$ and $\beta^* = 0.475$ at $\phi = 0.5$.

8. What the model does and does not say

The model speaks to incentives: who contributes and how much, who encloses, and how community capability moves the boundary. Enclosure is modeled as revenue capture only: relicensing converts base into product customers at a higher margin, and nothing is said about product quality, pricing against a hyperscaler, or the legal design of the license itself. The fork arrives as an exogenous probability ϕ read as community capability; fork timing, coordination, and governance are outside the model, as is any claim about how a particular community would respond. Regulation is absent by design: compliance-cost incidence lives in Section 10.2, and a fixed per-period cost common to both strategies cancels out of D without touching the threshold. The setting is one firm in partial equilibrium, with no competition among service firms and no strategic play by the community, and the two-period horizon compresses fund structure into a single β, with Remark 1 as the reduction's defense. No efficacy claim is made for openness, for forks, or for any policy instrument: the outputs are incentive statements at given parameters, and the parameters await the measurement program of Section 11.

The status ledger is as follows. The propositions and corollaries of Section 6 are theorems conditional on the model, each carried by a written proof (Appendix), a numeric cross-check against brute force, and a Monte-Carlo robustness run (Section 7). Three mappings are interpretations, defended but outside the theorems: reading β as the funding path (via Remark 1), reading ϕ as community capability, and reading enclosure as relicensing to source-available terms. Everything quantitative about the population awaits data: the parameter magnitudes, the existence and sharpness of the sorting, and the contribution gradient inside the open cluster. Section 11 maps each to its test.

9. The service-centric majority: where the value is

For the self-funded majority the durable questions are operational: which projects to build capacity in, at which phases, and with which 3A mix. Revenue follows the client typology: need specification and full delivery for No-skill buyers, adaptation and module development for Frontier buyers who can specify, assurance and maintenance for buyers who want a guaranteed version of the flow. None of this requires controlling a core, so hyperscaler resale and upstream relicensing arrive as supplier-selection events rather than existential threats: when an upstream vendor relicenses and a fork reconstitutes the project, a service firm's continuity runs through the fork, and its client relationships, integration knowledge, and participation capacities carry over.

The 2018 blueprint already pointed here. Its ten successful firms earned from directly paying customers, sold layered support and subscriptions, and built "primarily... around software related services" (S&J, pp. 135–136); its interviewees reported that establishing an open-source business is no harder than a proprietary one (S&J, p. 141) and did not treat outside investment as a prerequisite (S&J, p. 137). The European specifics of 2026 sharpen the picture along the lines of the predictions in Section 3: regional integration expertise as the scarce asset, public-sector and regulated buyers who weight auditability and sovereignty in procurement, and a partner-and-integrator channel in place of direct outbound sales. The macroeconomic weight of open source in the EU economy is documented at the aggregate level (Blind et al. 2021; Blind & Schubert 2023), which makes the composition question, who inside the sector carries the employment and revenue, the open one.

Two public-sector deployments give that demand a name and a size. Germany's openDesk, developed by the Zentrum für Digitale Souveränität, is being rolled out across federal and regional administrations "with a planned deployment of approximately 160,000 licences by the end of 2025", and "over 60,000 teacher workstations in Baden-Württemberg are migrating to openDesk"; France's La Suite numérique, run by DINUM and ANCT, is "already used daily by several hundred thousand public servants" (European Commission 2026b). No deployment, migration, or switching cost is reported for either programme, so the figures count seats and say nothing about the economics of the service relationships that would generate the 3A revenue of Section 4.

The bifurcation claim is a hypothesis: the self-funded service majority carries most of the sector's employment and revenue while receiving little of the discourse's attention. It is motivated by the visible-vendor selection bias of the trade press, and it is exactly what the empirical program's cluster analysis and sampling frame are designed to test (Section 11). No share is claimed here.

10. The product-vendor minority and its two shocks

The licensing and regulatory dynamics that dominate the public account belong mostly to the venture-backed product minority, and this paper treats them as bounded results about that population.

10.1 The appropriability tension

A product vendor captures value from a core it controls under a weak appropriability regime: a permissive license chosen to maximize adoption. When a hyperscaler turns that adoption into a managed service and captures the value, the vendor's defensive move is to relicense to source-available terms. The move strengthens legal protection and, in the same act, breaks the reciprocity that recruited the community, prompting a foundation-backed fork that reconstitutes the value network without the vendor. In the theory's language, relicensing is the owner exercising alienation rights in a way that regresses the project's value network (Section 2.5); in the model's language, it is the enclosure branch of Section 5, chosen where impatience and conversion economics coincide, and taxed by the community's fork capability (Proposition 3).

The indicators assembled for the Commission size the position under contest. Cloud ranks first or joint first for adoption in each of the twelve industrial ecosystems of the critical-capacities study, at 31 to 53 percent of firms; private investment concentrates on cloud, big data, and IoT, with cloud the largest at approximately €137 bn against €90 bn for big data and €85 bn for IoT; and in the study's ranking of profit margins across its technology areas, "firms in cloud-edge-IoT have higher profit margins, companies in quantum report overall negative values. AI is also still at the end of the ranking" (European Commission 2026a). The layer with the widest adoption is also the layer with the largest spend and the highest margins, which is the position a managed-service operator occupies and the one a relicensing vendor is trying to reach.

The public record of the 2020s supplies the episodes as timeline sentences; the primary sources are listed in the References section "Documented relicensing & fork episodes". In 2018, MongoDB relicensed to the Server Side Public License; no fork of comparable weight followed. In 2021, Elastic relicensed Elasticsearch to the SSPL and the Elastic License, and AWS launched the OpenSearch fork. In 2023, HashiCorp relicensed Terraform and other core tools to the Business Source License, and the OpenTofu fork followed under the Linux Foundation. In 2024, Redis moved from BSD to dual SSPL/RSALv2 licensing, and the Valkey fork followed. The pattern is consistent with the model's comparative statics: relicensing into a dense, corporate-backed ecosystem, where fork capability is high, drew forks; the early MongoDB move, whose value had already migrated to a managed-cloud asset the firm owned, did not. The sustainability of communities beyond a fork has a documented precedent in LibreOffice (Gamalielsson & Lundell 2014). These are consistency observations on four public episodes; the model assigns no episode-level probabilities, and the mediation claim (owned complementary assets make relicensing survivable) is tested on the product-vendor subsample, per Section 11.

10.2 The compliance-cost tension

The Cyber Resilience Act attaches obligations to placing products with digital elements on the EU market: a secure development process, vulnerability handling, an SBOM, technical documentation, and conformity assessment, falling on the category the act defines as manufacturers, with a distinct and lighter regime for the entities it defines as open-source software stewards (Regulation (EU) 2024/2847; sector-facing guidance in CNLL & Inno^3 2024; early sector positions in Eclipse Foundation et al. 2023). For this paper the act is a cost-incidence object, and two hypotheses with named mechanisms follow, both unmeasured.

The burden hypothesis: the compliance cost is largely fixed, and fixed cost is regressive in firm size, so a security mandate has a size-regressive incidence and can behave as an artificial appropriability regime that advantages incumbents. Where the manufacturer category catches a small product vendor, the fixed line lands on a small revenue base; the steward regime bounds the exposure of foundations and non-commercial maintainers by design. How national enforcement draws the boundary between the categories is itself endogenous, so realized incidence will differ from statutory incidence.

The offset hypothesis: the same documentation and process artifacts that the act requires are, on the demand side, saleable assurance. Auditability and a documented security posture are procurement advantages with public and regulated buyers, and the service majority can package them as an assurance service (Section 4.1). The net of burden and offset is a firm-level empirical question, conditioned on funding path and on which side of the manufacturer/steward boundary a firm sits.

No claim is made that the CRA delivers security or protection to anyone; the paper prices cost incidence and demand composition only. One structural remark carries over from the model (Section 8): a per-period compliance cost common to both strategies cancels out of the enclosure comparison entirely, so the CRA moves the funding-path margins only through its asymmetries, fixed cost against firm size, and manufacturer against steward.

11. From conceptual claims to measurement

The measurement vehicle is the executive survey instrument and analysis plan in `../notes/value-capture/instrument.md`: fourteen closed questions plus two profile items, under eight minutes, with the funding path as the primary conditioning variable and the 3A revenue mix as the model descriptor. The licensing and CRA items (Q7–Q10) are analyzed on the venture-backed product subsample where they apply. Self-reports are triangulated against license histories, GitHub participation series, and Crunchbase/Dealroom funding records, per the plan's common-method-bias standards, and the sampling frame runs through the European trade associations (CNLL, APELL, OW2) with pre-registration and documented non-response. The program's phase-and-gate structure (`../notes/common/research-program.md`) sequences the work: Phase 1 is this paper's referee-legible conceptual gate, Phase 2 the pre-registered instrument, Phase 3 the secondary-data triangulation, Phase 4 the survey, Phase 5 the empirical paper, held to the same verification standard as this paper's propositions (Section 7).

The claim-by-claim map:

Claim	Status	Where it gets tested
Contribution intensity rises with patience, $x^* = (\beta s A \eta/2)^2$	Theorem (Prop. 1), conditional on the model	Funding-path item crossed with GitHub contribution series (Phase 3)
Revenue deadlines act as reduced effective patience	Theorem on the contribution margin, directional on the strategy margin (Remark 1)	Q13 board stance and funding records as patience proxies (Phases 3–4)
Enclosure requires $p\lambda > s$ and $\beta < \beta^*$	Theorem (Prop. 2)	Q7–Q8 license trajectories on the venture-backed subsample, with funding data (Phases 3–5)
Fork capability shrinks the enclosure region	Theorem (Prop. 3)	Q9 fork exposure; contributor series around relicensing dates (Phase 3)
Population bifurcates into two clusters	Corollary (Cor. 1) plus hypothesis	Cluster analysis and fsQCA on the full sample (Phase 5)
Service majority carries most employment and revenue	Hypothesis	Sampling frame plus survey distributions (Phases 2–4)
3A mix varies with phase profile and client capability	Interpretation of JVZ	3A revenue-mix item with Q4–Q6 (Phases 4–5)
Delivery shifted toward managed SaaS since 2018	Hypothesis	Q3 against the S&J baseline (Phase 4)
Sovereignty procurement is a demand channel	Hypothesis	Q11 (Phase 4)
CRA fixed cost is size-regressive	Hypothesis	Q10 against firm size, product subsample (Phase 5)
Auditability sells as assurance (CRA offset)	Hypothesis	Q12 on the service majority (Phases 4–5)

Two features of the map carry the program's epistemics. The theorems are falsifiable only through their interpretations: a failed contribution gradient impugns the reading of β as funding path before it impugns Proposition 1, and the design separates the two by measuring patience proxies (board stance, funding structure, growth targets) independently of contribution. And every number that eventually fills the right-hand column must clear the same standard before it enters a paper: a written proof or computation, a numeric cross-check, and a robustness report. Nothing in this table reports a result.

12. Discussion

For theory. The paper extends the FLOSS phase model with a funding-path axis and a European institutional layer, and bounds the licensing and regulatory dynamics to the product-vendor population they govern. The axis carries the formal weight of the paper: the model of Sections 5 and 6 turns it into a closed-form contribution schedule, a threshold characterization of enclosure, and deterrence and timing comparative statics, all conditional on a small model whose assumptions are stated and stress-tested. The theoretical claim carried over from the value-network account is that appropriability and the complementary asset are coupled in open source: for vendors who capture a core, moving the license moves the community, with opposite sign. The extension offers the theory a candidate answer to a question it leaves open (why some editors enclose and others never do), in one parameter the theory does not carry: the patience induced by the firm's capital structure.

For managers. The first decision is the funding path, because it selects the game. A founder on the self-funded path should build 3A-service depth, treat the core as shared infrastructure, and read Proposition 1 as a return schedule: patience raises the optimal investment in the project quadratically, and that investment makes the service offer credible.

A founder on the venture path should decide early which complementary asset the firm will own before it needs to defend value, because relicensing is survivable only for a firm that already owns one. The same founder should prefer instruments that raise appropriability without breaking reciprocity before reaching for a non-open license; the fork episodes of Section 10 are the cost side of that ledger.

For policy. The population a European open-source policy most wants to sustain is the self-funded service majority, and a security mandate with a size-regressive fixed cost selects against exactly the small firms and foundations at its base. If the goal is a resilient European open-source sector, the mandate should be paired with instruments that flatten the fixed-cost curve for small firms and foundations: mutualized compliance tooling, shared conformity infrastructure, steward-regime clarity. This is a claim about instrument design; it reports no measured outcome, and it credits no instrument with protection it has not been shown to deliver.

European policy analysis holds open source in two positions at once, and the model says what reconciles them. The Commission's critical-capacities study records in its SWOT that "Open source is seen as a strength and a potential strategy to close gaps where Europe lags behind global competitors, particularly the US and China. In areas like software, open source could provide a more agile and collaborative route to innovation". The same study's dependency inventory lists "Opensource standards (Kubernetes, Spark)", "Orchestration Software (e.g. Kubernetes, OpenStack) for Cloud", "AI Frameworks (TensorFlow, PyTorch)", and "Key Blockchain Protocols (Ethereum)" among the EU's technological dependencies, notes that core protocols "are governed outside the EU, limiting influence over updates and standards", and observes that "In NGI, although much software is open source, many initiatives fail to reach widespread adoption, and public institutions do not subsequently deploy or endorse these home-grown alternatives" (European Commission 2026a). Both readings hold under a condition this paper's machinery states: an open license confers control on an adopter only where the surrounding community can fork the project and sustain it afterwards, which is the same capability ϕ that deters enclosure in Proposition 3. The license is the legal permission to fork; the capacity to exercise it is an asset that has to exist and be paid for. The study names open source as a strategy without attaching an allocation to it: open source appears in none of the four investment options' budget grids and in no funding condition, and the proposed programme design carries no open-licensing requirement, no open-standards clause, and no reuse obligation. Whether such conditions would move the fork capability of any project is an empirical question that the design as described leaves untested.

13. Conclusion

Open-source businesses capture value by selling assurance, adaptation, and assistance around code they help maintain, and the shape of that business is set first by the firm's funding path. The theory supplies the machinery: value lives in the flow, the 3A services price access to it, the typology says who buys which, the phases say when each model is possible, and alienation rights say who can enclose. This paper's addition is the axis that selects among the possibilities, now backed by theorems: contribution rising with patience in closed form, enclosure confined to impatient and conversion-rich firms below a unique threshold, forks deterring enclosure at the margin, margin erosion and maturity widening the region where enclosure pays, and a population that should split into two clusters if the reading is right. The 2018 blueprint gives the baseline, the 2026 European layer gives the predicted shifts, and the instrument gives every claim its test. The drama of the licensing wars belongs at the edge of this picture, and the economics of the service majority at its center. The claims are framed to be testable and are left untested; the contribution of this draft is the structure and its measurement map.

Appendix: proofs

Remark 1. Technique: constrained optimization on an interval, then a signed comparison of value gaps. Under OPEN, $R_1 = sA - x$ and $R_2 = sA(1 + g_0 + \eta\sqrt{x})$, so the floor is the interval constraint $x \le sA - \bar{R}$, feasible iff $\bar{R} \le sA$. (i) $V_{\text{open}}(\beta, \cdot)$ is strictly concave in $\sqrt{x}$ with unconstrained maximizer $x^*(\beta) = (\beta sA\eta/2)^2$, so the constrained maximizer is $x^c = \min\{x^*(\beta), sA - \bar{R}\}$, equal to $sA - \bar{R}$ when the floor binds. Solving $x^*(\beta_{\text{eff}}) = sA - \bar{R}$ gives $\beta_{\text{eff}} = 2\sqrt{sA - \bar{R}}/(sA\eta)$, and $\beta_{\text{eff}} < \beta$ exactly when the floor binds. (ii) First claim: the constraint shrinks OPEN's feasible set, so $V^c_{\text{open}}(\beta) \le V^*_{\text{open}}(\beta)$, while $V_{\text{enc}}(\beta)$ is unchanged whenever $\bar{R} \le p\lambda A$, hence $D^c(\beta) \le D(\beta)$ pointwise and the enclosure set weakly expands. Second claim: $V^c_{\text{open}}(\beta) = V_{\text{open}}(\beta, x^c)$ and $V^*_{\text{open}}(\beta_{\text{eff}}) = V_{\text{open}}(\beta_{\text{eff}}, x^c)$ share the same x^c, so their difference is $(\beta - \beta_{\text{eff}})\, sA(1 + g_0 + \eta\sqrt{x^c})$, while $V_{\text{enc}}(\beta) - V_{\text{enc}}(\beta_{\text{eff}}) = (\beta - \beta_{\text{eff}})\, p\lambda A\delta(1 - \varphi)$. A1 states $s(1 + g_0) > p\lambda\delta(1 - \varphi)$, and $\eta\sqrt{x^c} \ge 0$, so the first difference strictly exceeds the second and $D^c(\beta) > D(\beta_{\text{eff}})$. Third claim: $\bar{R} > sA$ makes $\{x \ge 0 : sA - x \ge \bar{R}\}$ empty. The earlier version of this remark asserted the equivalence on both margins by a Lagrangian sufficiency argument. That argument needs a convex choice set, which the plan space is not, since it carries the discrete OPEN/ENCLOSE choice; part (ii) is what survives, and nothing downstream uses more, Propositions 2 to 4 never invoking the remark. □

Proposition 1. Technique: change of variables and strict concavity. Write $y = \sqrt{x}$. The open objective becomes $sA + \beta sA(1 + g_0) + \beta sA\eta\, y - y^2$, strictly concave in y on $[0, \infty)$ with unique maximizer $y^* = \beta sA\eta/2 > 0$, which is interior for all $\beta, s, A, \eta > 0$. Hence $x^* = (y^*)^2 = (\beta sA\eta/2)^2$, at which the first-order condition in x, $\beta sA\eta/(2\sqrt{x}) = 1$, holds. As the square of a product of positive, strictly increasing functions of each argument, x^* is strictly increasing in β, s, A, and η. The optimized value is $V_{\text{open}}(\beta) = sA(1 + \beta(1 + g_0)) + (\beta sA\eta)^2/4$. □

Proposition 2. Technique: envelope theorem and the intermediate value theorem. The maximizer $x^*(\beta)$ is unique and interior, so by the envelope theorem $V'_{\text{open}}(\beta) = sA(1 + g_0 + \eta\sqrt{x^*(\beta)}) \ge sA(1 + g_0)$, while $V'_{\text{enc}}(\beta) = p\lambda\delta(1 - \phi)A$. Under A1, $D'(\beta) \ge A\,[s(1 + g_0) - p\lambda\delta(1 - \phi)] > 0$, so D is strictly increasing and continuous on $[0, 1]$ with $D(0) = A(s - p\lambda)$ (the continuous extension at $\beta = 0$, where $x^* = 0$). (i) If $s \ge p\lambda$ then $D(0) \ge 0$ and $D(\beta) > D(0) \ge 0$ for $\beta > 0$. (ii) If $D(1) \le 0$ then $D(\beta) < D(1) \le 0$ for $\beta < 1$, so ENCLOSE is strictly optimal on $(0, 1)$ and weakly optimal at $\beta = 1$. (iii) If $D(0) < 0 < D(1)$, strict monotonicity and the intermediate value theorem give a unique $\beta^* \in (0, 1)$ with $D < 0$ on $(0, \beta^*)$ and $D > 0$ on $(\beta^*, 1]$. □

Remark 2. Technique: convexity of quadratics and interval sublevel sets. Substituting $x^*(\beta)$ gives $D(\beta) = c_0 + c_1\beta + c_2\beta^2$ with the coefficients of the statement and $c_2 > 0$, a convex quadratic. If A1 fails, then $s \le s(1 + g_0) \le p\lambda\delta(1 - \phi) < p\lambda$, using $\delta(1 - \phi) < 1 \le 1 + g_0$. Hence $c_0 = A(s - p\lambda) < 0$. Sublevel sets of a convex function are intervals, so $\{\beta \ge 0 : D(\beta) \le 0\}$ is an interval containing 0, and the enclosure region on $(0, 1]$ is an initial segment: empty, all of $(0, 1]$, or $(0, \beta^*)$ for a unique upcrossing β^*. □

Proposition 3. Technique: implicit function theorem with the envelope theorem. At an interior threshold, $D(\beta^*; \theta) = 0$ with $\partial D/\partial\beta\,|_{\beta^*} > 0$ by Proposition 2, and D is continuously differentiable in (β, θ) for $\theta \in \{\phi, \delta, p, \lambda\}$, so $\beta^*(\theta)$ is differentiable with $d\beta^*/d\theta = -(\partial D/\partial\theta)/(\partial D/\partial\beta)$. The maximizer x^* does not depend on these parameters, so by the envelope theorem the partials are those of the direct payoffs: $\partial D/\partial\phi = \beta Ap\lambda\delta > 0$, while $\partial D/\partial\delta = -\beta Ap\lambda(1 - \phi) < 0$, $\partial D/\partial p = -\lambda A\,(1 + \beta\delta(1 - \phi)) < 0$, and $\partial D/\partial\lambda = -pA\,(1 + \beta\delta(1 - \phi)) < 0$. The signs of $d\beta^*/d\theta$ follow: negative in ϕ, positive in δ, p, and λ. □

Proposition 4. Technique: the same implicit-function argument, with the envelope theorem applied on the s margin where x^* does depend on the parameter. For g_0: x^* does not depend on g_0, so $\partial D/\partial g_0 = \beta sA > 0$ directly. For s: by the envelope theorem, differentiate $V_{\text{open}}(\beta) = sA(1 + \beta(1 + g_0)) + (\beta sA\eta)^2/4$ at the optimizer, giving $\partial D/\partial s = A(1 + \beta(1 + g_0)) + \beta^2 sA^2\eta^2/2 > 0$, while V_{enc} does not involve s. With $\partial D/\partial\beta > 0$ at the threshold, $d\beta^*/ds < 0$ and $d\beta^*/dg_0 < 0$. □

Corollary 2. Chain rule: $d\beta^*/d\sigma_I = (d\beta^*/d\phi)\,\phi'(\sigma_I) < 0$, the first factor negative by Proposition 3 and the second positive by assumption. □

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 two anchor papers and the 2026-07-15 re-anchoring decisions.
- `../notes/common/research-program.md`: research questions, the two-contribution structure, phases and gates, target venues.
- `../notes/common/provenance-and-salvage.md`: provenance of the source material and what was kept.
- `../notes/value-capture/instrument.md`: the executive survey instrument and analysis plan of Section 11; the cluster-analysis and fsQCA design that tests Corollary 1.
- `notes/value-capture-wp.md`: internal notes for this paper (version history, venue thoughts, bibliography debt).
- `notes/archived-funding-path-theory-wp-v0.2.md`: the formal core's standalone pre-merge version (archived 2026-07-16); its internal notes remain in `notes/funding-path-theory-wp.md`.

Annex B. Reproduction: scripts, data, and figures

- `../verify/funding_path.py`: the verification suite for the formal core (fixed seed 20260716). One run recomputes every number in Section 7, regenerates both figures, and rewrites the results file; dependencies are declared inline, so `uv run ../verify/funding_path.py` from `papers/` (or `uv run verify/funding_path.py` from the project root) reproduces everything.
- `../verify/results/funding_path.json`: machine-readable results of the run (seed, grid and Monte-Carlo configurations, per-check outputs).
- `figures/funding-path-contribution.png` and `figures/funding-path-value.png`: the two figures of Section 6, written by the suite.

The paper reports no other computed number and uses no dataset.

References

Page references to Jullien, Viseur & Zimmermann (2025), abbreviated JVZ in the text, are to the preprint archived in `../../references/` (HAL hal-04963100); the published article (*Journal of Systems and Software* 224:112383) paginates differently. Page references to Spijkerman & Jansen (2018), abbreviated S&J, are to CEUR Vol. 2305 (pp. 128–143). The formal results of Sections 5 to 8 are verified by `../verify/funding_path.py` (seed 20260716).

- Athey, S., & Ellison, G. (2014). Dynamics of open source movements. *Journal of Economics & Management Strategy*, *23*(2), 294–316.
- Blind, K., Böhm, M., Grzegorzewska, P., Katz, A., Muto, S., Pätsch, S., & Schubert, T. (2021). *Study about the impact of open source software and hardware on technological independence, competitiveness and innovation in the EU economy*. European Commission, Directorate-General for Communications Networks, Content and Technology. https://op.europa.eu/en/publication-detail/-/publication/d38e0892-7eb6-11eb-9ac9-01aa75ed71a1/language-en/format-PDF
- Blind, K., & Schubert, T. (2023). Estimating the GDP effect of Open Source Software and its complementarities with R&D and patents: Evidence and policy implications. *The Journal of Technology Transfer*. https://doi.org/10.1007/s10961-023-10033-1
- Bonaccorsi, A., & Rossi, C. (2003). Why open source software can succeed. *Research Policy*, *32*(7), 1243–1258.

- Casadesus-Masanell, R., & Llanes, G. (2011). Mixed source. *Management Science*, *57*(7), 1212–1230.
- CNLL & Inno³. (2024). *Cyber Resilience Act Compliance Guide for Open Source Players*.
- Dahlander, L., & Wallin, M. W. (2006). A man on the inside: Unlocking communities as complementary assets. *Research Policy*, *35*, 1243–1259.
- Eclipse Foundation, APELL, CNLL, OSI, OW2, OpenForum Europe, OSBA, Linux Foundation Europe, ESOP, COSS, The Document Foundation, & Software Heritage (2023). *Open letter to the European Commission on the Cyber Resilience Act*, 17 April 2023.
- 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.
- European Commission, Directorate-General for Communications Networks, Content and Technology. (2026a). *Study on the EU's critical digital capacities deployment beyond 2027: final report.* Prepared by Technopolis Group; K. Izsak, P. Simmonds, M. Ploeg, S. Viscido, A. Lotito, C. Moreno & R. Alasghar (eds.). Luxembourg: Publications Office of the European Union. ISBN 978-92-68-40079-1, doi:10.2759/0673677.
- European Commission, Directorate-General for Communications Networks, Content and Technology. (2026b). *Study on the EU's critical digital capacities deployment beyond 2027, Annex 15: Interoperability and GovTech.* Prepared by A. Murray, A. Adewusi, & T. Marshall, Technopolis Group. Luxembourg: Publications Office of the European Union. ISBN 978-92-68-40089-0, doi:10.2759/5257213.
- European Union. (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, 20 November 2024.
- Gamalielsson, J., & Lundell, B. (2014). Sustainability of open source software communities beyond a fork: How and why has the LibreOffice project evolved? *Journal of Systems and Software*, *89*, 128–145.
- 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 (HAL: hal-04963100).
- 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.
- Kogut, B. M., & Metiu, A. (2001). Open source software development and distributed innovation. *Oxford Review of Economic Policy*, *17*(2), 248–264.
- Lerner, J., & Tirole, J. (2002). Some simple economics of open source. *The Journal of Industrial Economics*, *50*(2), 197–234.
- Lerner, J., & Tirole, J. (2005). The scope of open source licensing. *The Journal of Law, Economics, and Organization*, *21*(1), 20–56. https://doi.org/10.1093/jleo/ewi002
- Lin, Y.-K., & Maruping, L. M. (2022). Open source collaboration in digital entrepreneurship. *Organization Science*, *33*(1), 212–230. https://doi.org/10.1287/orsc.2021.1538
- Morgan, L., Feller, J., & Finnegan, P. (2013). Exploring value networks: Theorising the creation and capture of value with open source software. *European Journal of Information Systems*, *22*, 569–588.

- Nagle, F. (2018). Learning by contributing: Gaining competitive advantage through contribution to crowdsourced public goods. *Organization Science*, *29*(4), 569–587.
- Riehle, D. (2021). The open source distributor business model. *Computer*, *54*(12), 99–103.
- Shaikh, M., & Cornford, T. (2011). *Total cost of ownership of open source software: A report for the UK Cabinet Office supported by OpenForum Europe*. Technical report, UK Cabinet Office.
- Spijkerman, Z., & Jansen, S. (2018). The Open Source Software Business Model Blueprint: A Comparative Analysis of 10 Open Source Companies. In *Software Business Workshop (SiBW 2018)*, CEUR Workshop Proceedings, Vol. 2305, pp. 128–143. http://ceur-ws.org/Vol-2305/paper10.pdf
- Teece, D. J. (1986). Profiting from technological innovation: Implications for integration, collaboration, licensing and public policy. *Research Policy*, *15*(6), 285–305.
- 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.
- Zhu, F., & Liu, Q. (2018). Competing with complementors: An empirical look at Amazon.com. *Strategic Management Journal*, *39*(10), 2618–2642.

Documented relicensing & fork episodes

Primary sources for the four episodes of Section 10.1 are the company license announcements and the foundation launch posts; the URLs remain to be added.

- MongoDB relicensing to the Server Side Public License (SSPL) (2018).
- Elastic relicensing of Elasticsearch to the Server Side Public License (SSPL) and the Elastic License (2021); AWS OpenSearch fork.
- HashiCorp relicensing of Terraform, Vault, and other core tools to the Business Source License (BSL) (2023); OpenTofu fork (Linux Foundation).
- Redis relicensing from BSD to SSPL and RSALv2 (2024); Valkey fork.