

Who Pays for the Battery? Why Energy-Only Settlement Cannot Reliably Finance Community Storage

Nicolas J. Greber^{a,b}

^a*University of Zürich, Zürich, Switzerland*

^b*Swiss Economics, Zürich, Switzerland*

Abstract

Local energy communities are expected to attract investment in shared battery storage through internal energy trading. We show that this expectation is not generally warranted when storage operators recover costs solely from the internal price spread. Under energy-only settlement, two opposing forces, declining utilization with scale and endogenous spread compression, determine whether arbitrage revenue covers cost. A scarcity-clearing internal price balances them for every mismatch distribution, as in competitive wholesale models; the canonical local-market rules do not. Each fails for a different reason: no revenue, a constant spread that only coincidentally equals the required value, or compression ahead of scarcity. Calibrating to quarter-hourly smart-meter data from a Swiss collective self-consumption community, we find that at the benchmark cost supply–demand-ratio settlement finances a tenth to a fifth of the surplus-optimal capacity, depending on the convention, and under-invests in every scenario. Community ownership, where a balanced-budget membership fee replaces the energy spread as the cost-recovery instrument, restores surplus-optimal investment in tariff-value terms and undistorted participation. Where the price corridor makes shared storage efficient, regulatory frameworks, including Swiss local electricity communities and analogous EU schemes, should therefore enable non-energy financing instruments for community-level storage.

Keywords: Energy communities, Collective self-consumption, Battery storage, Settlement rules, Financing gap, Electricity market regulation

JEL classification: D47, L22, L51, L94, Q41

1. Introduction

Energy communities are expanding across Europe, and their legal basis now extends to shared assets. The Renewable Energy Directive (RED II) gives jointly acting renewables self-consumers the right to share self-generated electricity, including through peer-to-peer arrangements, and defines renewable energy communities that may own generation and storage (European Parliament and Council of the European Union, 2018); the Electricity Directive adds citizen energy communities (European Parliament and Council of the European Union, 2019); and the 2024 reform of the electricity market design introduces an explicit right to energy sharing, under which a third party may own or operate a storage facility of up to 6 MW without thereby becoming an active customer (European Parliament and Council of the European Union, 2024, Art. 15a). National transpositions differ, but they share a structure. Members within a defined perimeter share locally produced electricity *on top of* regulated retail supply: each household retains the option to import at the regulated retail price and to export surplus at the regulated feed-in tariff.¹ The difference between these two prices is a *price corridor* within which internal exchange creates value (Gautier et al., 2018), and which tariff components apply to internal exchange determines how wide it is.

Switzerland is our empirical setting and a clean example of this structure. The *Zusammenschluss zum Eigenverbrauch* (ZEV) lets households behind a common connection point self-consume jointly, netted behind the meter, so that internal flows avoid the full volumetric retail tariff. Since January 2026, local electricity communities (LECs; *Lokale Elektrizitätsgemeinschaften*) extend this to households sharing electricity over the existing distribution grid as a virtual microgrid, within a single municipality and network area (Romano and Trutnevyte, 2026), with a reduced network charge on internal flows, residual imports and exports settled with the basic supplier at standard terms, and third-party storage operators admitted as participants (Der Bundesrat, 2025). Swiss rules also limit systematic grid-to-community arbitrage by storage, so the value of a community battery lies mainly in shifting local production within the day. The implementing ordinances fix the metering and the corridor and leave the internal price to the members; they do not provide for non-energy charges such as membership fees or capacity payments. Under the ZEV, such charges are a matter of private contracting and arise only where governance supports collective investment. Neither framework prohibits a capacity instrument; neither creates one.

The stakes are not small: Romano and Trutnevyte (2026) estimate that LECs in rural Swiss municipalities alone could supply 23% of the 2035 renewable target—but only under a social objective that adds generation and storage together; under a private objective that maximizes return on investment the contribution falls below 12%, and storage enters only in the social stage, where it raises self-sufficiency from 14% to 45%. Behind these frameworks lies a common expectation: that internal trading arrangements designed to share PV generation will extend naturally to community-level batteries. If households can trade at prices that reflect local scarcity, the resulting intertemporal price signal should attract storage investment. This paper asks whether that expectation is warranted. *Can internal energy settlement alone finance efficient community-level storage, and if not, which organizational form can?*

We use a few terms throughout. *Energy-only settlement* means that a storage operator is paid only per kWh it buys and sells inside the community, with no separate payment for capacity. The *settlement rule* is the formula that sets internal prices; a *scarcity-reflective* rule, such as the supply–demand–ratio (SDR) rule widely used in local-market pilots, raises the internal price when local supply is short and lowers it when supply is abundant. *Spread compression* is the effect of a battery on that rule: as the battery shifts more energy into the scarce period, it relieves the scarcity and narrows the very buy–sell spread that pays for it. A *balanced-budget membership fee* is a fixed charge per member set exactly to cover the battery’s cost.

Our answer depends on the community’s organizational form. Under decentralized energy-only settlement—the institutional default in most current implementations—two forces shape storage incentives. First, utilization declines with scale: a larger battery is fully used on fewer days, so its average return per unit of capacity exceeds the marginal value of the last unit, a force that on its own would invite *over*-investment.

¹In Switzerland the feed-in tariff is the distribution system operator’s purchase price for surplus PV (*Rückliefervergütung*), set periodically against the federal reference market price; we use “feed-in tariff” throughout, as in Gautier et al. (2025) and Romano and Trutnevyte (2026).

Second, spread compression pushes revenue below cost as capacity grows, a force toward *under*-investment. Whether revenue covers cost depends on how the settlement rule balances the two. A scarcity-clearing internal price—one that stays at the retail ceiling until the matchable energy is exhausted—balances them exactly, so that energy-only revenue equals the marginal value of capacity, as in competitive wholesale models (Schmalensee, 2022; Korpås and Botterud, 2020). The canonical local-market rules do not do so robustly: time-invariant prices give no intertemporal signal, an administrative spread hits the balance only if set to a value that depends on the mismatch distribution and the cost-corridor ratio, and scarcity-reflective formulas such as SDR let the discharge price fall before scarcity is relieved. Auction-based and bilateral rules, which we do not analyze formally, are subject to the same condition.

We develop a parsimonious analytical framework in which storage is supplied by a competitive fringe, prosumers decide whether to join, and internal prices follow a settlement rule, and we calibrate it to quarter-hourly smart-meter data from a Swiss ZEV. Three results follow. First, we characterize when energy-only settlement can finance efficient storage: revenue adequacy requires the settlement rule’s compression of the corridor at the efficient scale to equal the curvature of the shifting technology (eq. (15)). Scarcity clearing satisfies this condition for every mismatch distribution; each canonical rule violates it, for the reasons just given, and under SDR applied to the matchable market the shortfall is positive for every mismatch distribution, with a closed form under uniform mismatch. Second, decoupling storage cost recovery from the internal spread—through a platform’s capacity subsidy or a community-owned battery financed by a balanced-budget fee—is necessary for efficient investment under time-invariant prices, under matched-market SDR for every mismatch distribution, under the full-deficit base in our calibration, and under any administrative spread not calibrated to the community; Gautier et al. (2025) establish that it is sufficient. Together, among the canonical rules, the organizational form matters at least as much as the settlement rule. Community ownership achieves surplus-optimal investment and undistorted participation; a profit-maximizing platform achieves efficient capacity per member but serves half the efficient membership, so aggregate capacity is half the efficient level.

Third, the magnitudes are large. At the benchmark cost, SDR settlement finances about a tenth of the surplus-optimal capacity when the supply–demand ratio is measured against the matched internal market (Table 5), and about a fifth when it is measured against the full period-2 deficit (Table 6). Section 5 reports how far these shares travel: they move with the retail–feed-in price ratio, with the mismatch distribution and with seasonal pooling, and stay below one in every scenario: from 10% to 51% of the surplus-optimal capacity across the cost scenarios, and as low as 6.5% when the retail–feed-in price ratio is raised at a fixed corridor.

The mechanism does not depend on Swiss tariffs. The corridor, the settlement rule and the absence of a capacity instrument are features of every collective self-consumption framework we are aware of; what varies across countries and communities is the width of the corridor—set by which network charges internal flows avoid—and the mismatch distribution, which set the levels of the efficient and the financed battery. The direction is under-investment in every scenario we compute (Section 5.4). Under the matched-market convention under-investment is proven for every mismatch distribution (Corollary 1); under the full-deficit convention it is a property of the community’s deficit-to-matchable ratio and can reverse where that ratio is large (Section 3.3), so the direction travels across communities only under the former. The policy implication is correspondingly general: allowing trade within a corridor is not sufficient to deliver efficient flexibility investment. What matters is whether the regulatory framework permits organizational forms and cost-recovery instruments beyond the energy spread, such as community ownership with membership fees or regulated capacity payments.

Related literature

Our paper sits at the intersection of two literatures. The economics of energy communities studies how internal prices and sharing rules distribute a community’s value; the economics of storage studies how arbitrage value responds to installed capacity. We join them. Dinh et al. (2022) size a community battery as the solution to a Stackelberg game in which a profit-maximizing aggregator sets both the internal prices and the battery’s charging schedule while prosumers adjust consumption, and report payback periods of 5.1 and 8.7 years under two pricing schemes; their aggregator is the analogue of the platform regime of Section 4.1.

We ask the prior question—whether *competitive entry* under settlement rules the community takes as given recovers cost at the efficient scale—which neither literature has addressed.

Energy communities. Two contributions provide our point of departure. [Gautier et al. \(2025\)](#) establish that total community value depends only on external regulated prices and the physical self-consumption function, not on internal prices or sharing rules, and show in an appendix that community-owned storage financed through the community budget inherits this independence. Their storage extension models the battery as a device that converts exports into self-consumption at fixed internal prices; they leave arbitrage under time-varying internal prices, which they note would further increase the scope for arbitrage, for future research. [Balegamire and Gautier \(2025\)](#) characterize which combinations of time-invariant internal prices and sharing rules satisfy members’ participation constraints. We take up the deferred case: storage as a temporal-arbitrage asset remunerated under dynamic, settlement-rule-dependent internal prices. Their value-independence result implies that decoupling cost recovery from settlement is sufficient for efficient investment; we show it is necessary under time-invariant prices, under matched-market SDR settlement, and under any administrative spread not set to the one value that matches the curvature of the shifting technology at the efficient scale—a value the rule-setter cannot know without the mismatch distribution. Together the two results say that, among the canonical rules, no settlement rule finances the efficient battery robustly, whereas community ownership does so under every rule. On the stability and formation of communities, [Abada et al. \(2020a\)](#) use cooperative game theory to characterize when a community is stable against member exit, and [Abada et al. \(2020b\)](#) trace how community formation feeds back into grid tariffs; [Andreolli et al. \(2025\)](#) study how peer-to-peer trading affects household PV–battery adoption under uncertainty, and [Bravo-Melgarejo and Haritchabalet \(2025\)](#) compare grid- and individual-storage regulation for prosumers. [Romano and Trutnevskyte \(2026\)](#) model LEC formation in all 730 rural Swiss municipalities with a two-stage algorithm: a private stage that sizes generation to maximize members’ return and excludes storage by construction, and a social stage that adds batteries by minimizing the community’s total bill—the planner of our Section 4.2. Our paper supplies a rationale for that modeling choice: under the settlement rules communities actually use, a competitive storage supplier remunerated through the internal spread finances only a fraction of the efficient battery in our calibration, so the efficient battery is reliably built only where the community finances it collectively.

Settlement rules. The local-market literature proposes and pilots the rules we study. [Liu et al. \(2017\)](#) introduce SDR settlement, in which internal prices respond to the ratio of local supply to local demand; [Long et al. \(2017\)](#) propose the mid-market rate; [Morstyn et al. \(2018\)](#) propose federated power plants, in which peer-to-peer transactions form a coordinated prosumer coalition; [Mengelkamp et al. \(2018\)](#) derive seven required components for microgrid market design and evaluate the Brooklyn Microgrid against them; [Sorin et al. \(2018\)](#) develop consensus-based peer-to-peer clearing with product differentiation; and pilot implementations in Switzerland and Germany clear at 15-minute resolution ([Ableitner et al., 2020](#); [Meeuw et al., 2020](#); [Weinhardt et al., 2019](#)), while an Austrian commercial platform lets participants set their own buying and selling prices ([Energie Steiermark, 2024](#)). Surveys by [Soto et al. \(2021\)](#), [Reis et al. \(2021\)](#) and [Islam \(2024\)](#) map the resulting design space of trading mechanisms and business models. [Greber et al. \(2026\)](#) study such platforms computationally, letting prosumers learn trading policies against an endogenously cleared price; capacity withholding emerges there from price-aware learning, and the findings are reproduced on measured demand from twelve Swiss communities. These rules are designed to allocate surplus fairly and to give a scarcity signal; they take the asset base as given. We ask whether the signal they generate can finance the asset.

Storage economics. Storage value is tied to temporal price variation and the institutional environment governing remuneration ([Sioshansi, 2010, 2014](#); [Borenstein, 2012](#)), and that arbitrage value declines as more storage is added is established in wholesale markets ([Sioshansi et al., 2009](#)); the arbitrage value available to a price-taking device has been measured for European day-ahead markets ([Zafirakis et al., 2016](#); [Mercier et al., 2023](#)). In competitive wholesale models with uncapped scarcity prices this decline is nonetheless consistent with efficient investment: energy-only revenue equals the marginal value of storage

at the cost-minimizing capacity (Schmalensee, 2022; Korpås and Botterud, 2020; Tarel et al., 2024). Our point is that the canonical settlement rules used inside communities are not scarcity-clearing prices, so this benchmark does not carry over: within the institutional constraints of a local community—a bounded corridor, energy-only settlement under an administrative or formula-based rule, and no capacity mechanism—the same declining-value mechanism leaves the efficient battery unfinanced. The resulting problem resembles the wholesale missing-money problem (Joskow, 2008; Cramton et al., 2013; Newbery, 2016), but arises endogenously without a price cap. Market power is the other source of storage distortion studied in this literature: Andrés-Cerezo and Fabra (2023) show that storage investment depends on market structure, with over-investment relative to the first best when a dominant generator withholds output and lower capacity under vertical integration, Williams and Green (2022) that concentrated storage ownership reduces the welfare gains from storage though by less than market power in generation does, and Karaduman (2023) finds that consumer-owned storage roughly doubles the consumer-surplus gain relative to monopoly ownership, a close analogue of our community-versus-platform comparison. Our competitive-fringe assumption removes market power deliberately, to isolate the revenue-adequacy friction.

Ownership and intermediation. From an engineering perspective, Rodrigues et al. (2020) size a shared battery under user ownership and third-party (energy-sharing-provider) ownership but take remuneration as given; we supply the economic reason ownership matters, namely that it determines whether a non-energy financing instrument exists. When a platform mediates access, its fee and its transfers to storage affect both investment and participation, as in two-sided markets (Rochet and Tirole, 2003, 2006; Armstrong, 2006). The membership fee is the community-level analogue of a fixed connection charge that finances capacity which volumetric revenues alone cannot sustain, connecting to the case for multi-part tariffs for lumpy infrastructure (Cambini and Soroush, 2019) and to the grid-tariff design in Gautier et al. (2025).

2. The Model

We use a deliberately parsimonious framework—the same two-periods-per-day aggregation that Romano and Trutnevyte (2026) use to size community storage at national scale—to compare how different organizational forms perform on storage investment and participation, and to provide structure for the calibration in Section 5. The model captures the essential structure of a local energy community with a regulated price corridor, inflexible prosumers, and competitive storage entry, while abstracting from within-period spatial matching, multi-period dispatch, and seasonal variation.

2.1. Environment and the price corridor

We study an energy community of rooftop-PV households. We model a representative day with two periods: Period 1, in which all participants have a PV surplus, and Period 2, in which all face a consumption deficit. The structure is that of a single-cycle, PV-dominated community; communities that add wind, hydropower or baseload biomass, as the municipal communities of Romano and Trutnevyte (2026) do, have mismatch distributions with several cycles per day and are outside the scope of the two-period model (Appendix D.2; Section 6 returns to it among the limitations). In the baseline, prosumers are *synchronous*: there is no contemporaneous within-period matching between surplus and deficit households, so internal coordination arises only through *intertemporal* shifting via community-level storage.

Absent internal coordination, Period-1 surplus is exported at the feed-in tariff $\underline{p}$ and Period-2 deficit is imported at the retail price $\bar{p} > \underline{p}$. These regulated prices define an external *price corridor*:

$$\Delta_0 \equiv \bar{p} - \underline{p} > 0. \quad (1)$$

Shifting one kWh from Period 1 to Period 2 creates a gross value of Δ_0 per kWh shifted. The corridor is determined by the regulated tariff structure at the community’s connection point; in the Swiss ZEV framework, internal flows are netted behind the meter and avoid volumetric network charges, so Δ_0 equals

the full corridor. Under the LEC regime in force since January 2026, a reduced network tariff may apply to internal exchanges and the corridor adjusts accordingly.²

Throughout, we treat Δ_0 as exogenous: the community is a price-taker vis-à-vis the grid, and the surplus we compute measures *bill savings within the price corridor*—the *tariff-value optimum*—not resource-cost welfare. The distinction matters because Δ_0 embeds volumetric network charges and levies, which are cost-recovery instruments rather than marginal resource costs. A system-value benchmark that strips those charges out would use a narrower corridor and a (weakly) smaller optimal battery, provided the resource value of exported energy and any network benefit are held at their tariff values. If internal settlement were built on that corridor, the financing mechanism would be identical, and the efficient and the financed battery would both move—through the cost-corridor ratio and, under SDR, through the price ratio $\bar{p}/\underline{p}$ as well; if settlement remains tariff-based while the benchmark is system value, the two corridors differ and the comparison is a separate one (Section 5.6).

2.2. Agents

Prosumers. There is a unit mass of potential prosumers. Prosumers are operationally inflexible within the day: they do not adjust consumption or generation in response to internal prices. Their only choice is whether to join the community. Joining entails an idiosyncratic adoption cost $\phi \sim U[0, \bar{\phi}]$, drawn independently across households. Let $N \in [0, 1]$ denote the mass of participating prosumers.

Community-level storage. A competitive fringe of storage operators can install battery capacity. We model storage as provided by competitive entrants rather than as a community-owned asset. Many prosumers may already own behind-the-meter batteries; in our setting these are absorbed into the distribution of net surplus and deficit positions (Section 2.6). Our object of interest is investment in *community-level* capacity that arbitrages intraday mismatches and is remunerated through internal settlement. Free entry is deliberately the assumption most favorable to the energy-only market: it rules out market power and isolates a pure revenue-adequacy friction.³ Let $K \geq 0$ denote total installed usable energy capacity (kWh). Storage entails a per-unit (levelized) resource cost $c > 0$ per day, so total resource cost is cK .⁴

2.3. Mismatch and the time-shifting technology

In day ω , a participating prosumer has a Period-1 surplus $s(\omega) \geq 0$ that would be exported absent local trade and a Period-2 deficit $d(\omega) \geq 0$ that would be imported. The intertemporally matchable amount is the short side,

$$m(\omega) \equiv \min\{s(\omega), d(\omega)\} \in [0, \bar{m}], \quad (2)$$

the energy (kWh) that could be shifted from Period 1 to Period 2 if sufficient community-level storage is available. Across days, $m(\omega)$ is drawn from a stationary distribution with cdf F . We interpret $m(\omega)$ as the *residual intertemporal mismatch* after any contemporaneous within-period matching that may occur due to cross-sectional heterogeneity.

Given participation $N > 0$, define capacity per member as $k \equiv K/N$. In state ω , the battery shifts

$$q(\omega, k) \equiv \min\{k, m(\omega)\}. \quad (3)$$

²The corridor can be decomposed into grid-cost components following the tariff structure in [Gautier et al. \(2025\)](#): avoided import cost minus export value, which simplifies to the injection fee net of any self-consumption grid charge. For our purposes, the specific decomposition is immaterial; what matters is that $\Delta_0 > 0$ and is exogenous to the community.

³This institutional assumption is consequential. In Section 4.2, we show that community ownership of storage—where costs are covered by a membership fee rather than arbitrage revenue—eliminates the financing problem we identify here. The competitive-fringe setup isolates the case in which no non-energy financing instrument is available.

⁴A mapping from capital expenditure, lifetime, and discount rate to the daily levelized cost c is provided in [Appendix A](#).

We assume perfect roundtrip efficiency.⁵ The expected shifted energy per participant is

$$v(k) \equiv \mathbb{E}_\omega[\min\{k, m(\omega)\}], \quad (4)$$

and aggregate shifted energy is $Q \equiv N v(k)$.

By standard properties of the newsvendor expectation, $v(\cdot)$ is increasing, concave, and continuous, with $v(0) = 0$ and $v'(0) = \Pr(m > 0) > 0$. Concavity captures a physical regularity: as battery capacity expands, it is fully utilized on fewer days, so each additional kWh of capacity yields a smaller expected increase in shifted energy. The utilization rate $v(k)/k$ is consequently non-increasing in k .

2.4. Internal settlement and dynamic local prices

The central design question for a local energy community is how to price internal exchanges. Three objects need separating. The *price corridor* Δ_0 is fixed by the regulated outside options and by which charges internal exchange avoids; it determines the total value of shifting a kWh and is exogenous to the community. The *matching rule* determines how much energy counts as internal in each interval, and enters the model through $m(\omega)$ and $v(k)$. The *settlement rule* determines the internal prices p^1 and p^2 , and hence how the corridor is divided between storage and prosumers; the matching and settlement rules together are what [Gautier et al. \(2025\)](#) call the community's internal organization. Only the settlement rule is a pricing instrument, and it is the object we vary. Among the rules proposed in the local-market literature, scarcity-reflective ones—supply–demand ratio (SDR) settlement ([Liu et al., 2017](#)) chief among them—share one feature: internal prices are high when local supply is scarce and low when it is abundant, generating the intertemporal signal that can remunerate storage.

We formalize this by letting the internal settlement rule determine how the corridor Δ_0 is split between storage operators and prosumers. In day ω , let $p^1(\omega, k)$ denote the price paid by storage when charging in Period 1, and $p^2(\omega, k)$ the price received when discharging in Period 2. The within-day internal margin is

$$\Delta(\omega, k) \equiv p^2(\omega, k) - p^1(\omega, k) \in [0, \Delta_0]. \quad (5)$$

Storage earns $\Delta(\omega, k)$ per shifted kWh; the residual $\Delta_0 - \Delta(\omega, k)$ accrues to prosumers through more favorable internal terms of trade.

Under scarcity-reflective settlement, such as SDR, the margin responds endogenously to the local balance. When the battery is small relative to the mismatch ($k \ll m(\omega)$), the discharge period remains scarce and prices stay close to the retail ceiling $\bar{p}$, yielding a large margin. As capacity expands, scarcity is relieved and the margin compresses toward zero. This endogenous spread compression is the defining feature of scarcity-reflective settlement from the perspective of storage investment incentives. It is the local, settlement-driven counterpart of the well-documented decline of storage arbitrage value with capacity in wholesale markets—an *arbitrage limit* ([Sioshansi et al., 2009](#); [Sioshansi, 2010](#); [Antweiler, 2021](#)) that has been documented empirically for utility-scale batteries ([Lamp and Samano, 2022](#)).

Since participation and investment are chosen ex ante, we summarize the settlement mechanism by the expected margin per shifted kWh:

$$\Delta(k) \equiv \frac{\mathbb{E}_\omega[\Delta(\omega, k) q(\omega, k)]}{\mathbb{E}_\omega[q(\omega, k)]} = \frac{\mathbb{E}_\omega[\Delta(\omega, k) \min\{k, m(\omega)\}]}{v(k)}. \quad (6)$$

We impose a regularity condition on this induced margin.

Assumption 1 (Settlement margin). The function $\Delta: (0, \infty) \rightarrow [0, \Delta_0]$ is continuous, admits a finite right limit $\Delta(0^+) \equiv \lim_{k \downarrow 0} \Delta(k) \in (0, \Delta_0]$, and is weakly decreasing in k .

⁵With roundtrip efficiency $\eta \in (0, 1)$, charging q kWh yields ηq kWh at discharge, so the effective corridor becomes $\Delta_0^\eta = \eta \bar{p} - p < \eta \Delta_0$ and the cost–corridor ratio rises to $\alpha^\eta = c/\Delta_0^\eta$. Losses also change the matchable energy to $m^\eta(\omega) = \min\{s(\omega), d(\omega)/\eta\}$ (eq. (2)), which in charge units is weakly larger than m . The planner's and community-ownership results carry over exactly with Δ_0 replaced by Δ_0^η and m by m^η ; under SDR the period-2 supply becomes ηq , so ρ and the margin change as well, and the replacement is approximate. [Appendix D.5](#) re-runs the dispatch and the entry condition at $\eta = 0.90$.




Figure 1: Community-level storage compresses the internal margin $\Delta(k)$ within the external corridor $\Delta_0 = \bar{p} - \underline{p}$. Generic illustration; under SDR settlement p^1 remains at $\underline{p}$ and only p^2 falls (Appendix B.2).

Scarcity clearing (Appendix B.4) has a continuous margin when F is atomless, but that margin need not be weakly decreasing: across a gap in the support of F , or where its density is low, the throughput weights shift toward scarce days and the volume-weighted $\Delta(k)$ can rise even though every state's spread weakly falls. What the analysis below uses is monotonicity of unit revenue g , which scarcity clearing satisfies for every F because $g^{Wal}(k) = \Delta_0 \Pr(m > k)$; we therefore treat it separately in Appendix B.4. The assumption that $\Delta(k)$ is weakly decreasing captures spread compression. For SDR settlement under a uniform mismatch distribution we derive $\Delta(k)$ in closed form in Appendix B.5 and verify numerically that it is decreasing for the prices of Section 5; under the empirical mismatch distribution it is verified numerically on the relevant range of capacities. For other scarcity-reflective mechanisms—double auctions or bilateral negotiation—monotonicity is economically natural (more capacity relieves scarcity and compresses the margin) but has not, to our knowledge, been formally derived. We therefore treat Assumption 1 as a maintained regularity condition for the general analysis; the calibration rests on the SDR microfoundation.

Special case: time-invariant internal prices. Our framework nests the simpler case of a fixed internal price as a boundary case. If the community sets a single price p^s for all internal exchanges—the arrangement analyzed in much of the energy-community literature—then a storage operator charging in Period 1 and discharging in Period 2 faces the same price in both periods, so $\Delta(\omega, k) = 0$ for all (ω, k) and energy-only revenue for storage is identically zero. A uniform price is an instrument for allocating *contemporaneous* self-consumption, not for providing an intertemporal signal. Intermediate cases are also possible: if the community administratively sets distinct buy and sell prices with a constant margin δ , then $\Delta(k) = \delta$ for all k and the framework applies with no spread compression. The spectrum from zero margin (fixed price) through constant margin (administrative rule) to endogenous, declining margin (scarcity-reflective settlement) is ordered by the responsiveness of internal prices to local conditions. It is not ordered by the revenue available to storage: a constant spread can pay storage more or less than a scarcity-reflective rule at a given capacity, and a spread equal to Δ_0 pays weakly more than any compressed margin; only the zero-margin rule is unambiguously lowest. Section 3 shows that what matters for investment is not the level of the margin but its relation to the curvature of the shifting technology.

Scope of the rules analyzed. We evaluate four rules: a time-invariant internal price, which covers bill-sharing and the mid-market rate; a constant administrative spread; SDR settlement; and scarcity clearing (Section 3, Appendix B.4). The first three are the canonical mechanisms of the local-market literature: bill sharing and the mid-market rate are proposed by Long et al. (2017), SDR settlement by Liu et al. (2017), and Zhou et al.

(2018) evaluate all three. The fourth is the competitive benchmark. Auction-based rules, such as the double auction with participant-set limit prices used in the Quartierstrom pilot (Ableitner et al., 2020), and bilateral negotiation, which Morstyn et al. (2018) discuss among the mechanisms for forming prosumer coalitions, generate state-dependent prices; they fall under Assumption 1 when the induced volume-weighted margin $\Delta(k)$ is weakly decreasing, which a declining realized spread in every state makes plausible but does not guarantee, because the weights $q(\omega, k)$ shift with capacity. Where it holds, the revenue-adequacy condition of Section 3 applies to them; whether a given auction design satisfies it depends on bidding behavior, which we do not model. A double auction in which all participants bid their outside-option prices would reproduce scarcity clearing; one in which participants set limit prices inside the corridor, as observed in the pilots, would not. Our claims about “the rules in use” are therefore claims about the canonical rules.

2.5. Community value and the financing problem

The total community surplus from time-shifting is

$$W = \Delta_0 Q - cK - \int_0^N \phi(x) dx = N(\Delta_0 v(k) - ck) - \frac{\bar{\phi}}{2} N^2. \quad (7)$$

Define the *tariff-value surplus* created by one participant at capacity per member k :

$$S(k) \equiv \Delta_0 v(k) - ck. \quad (8)$$

A community planner who internalizes the full corridor chooses (N, k) to maximize W . At an interior optimum (N^W, k^W) , the first-order conditions are

$$\Delta_0 v'(k^W) = c, \quad (9)$$

$$N^W = \frac{S(k^W)}{\bar{\phi}}. \quad (10)$$

We call k^W the *efficient*, or *surplus-optimal*, capacity; it is the tariff-value optimum of Section 2.1. Condition (9) equates the marginal value of an additional unit of storage to its resource cost; since $v'(k) = \Pr(m > k) \leq 1$, a positive battery is efficient only if the cost–corridor ratio

$$\alpha \equiv \frac{c}{\Delta_0} \quad (11)$$

is below $v'(0) = \Pr(m > 0)$, hence a fortiori below one, and k^W is then the $(1 - \alpha)$ -quantile of the mismatch distribution.⁶ Condition (10) sets participation so that the marginal adopter’s cost equals the per-participant surplus.

Two properties of $S(k)$ deserve emphasis. First, it depends only on the corridor Δ_0 , the shifting function $v(k)$, and the resource cost c —not on any internal settlement rule. *The total value of time-shifting is independent of how that value is divided; whether the community can finance the asset that creates it depends entirely on the settlement rule.* Second, the surplus-maximizing capacity per member k^W can be achieved only if the community can *finance* that level of storage. Under decentralized energy-only settlement, storage is financed through the internal margin $\Delta(k)$, which is set by the settlement rule and compressed by the very capacity it is meant to finance. The instrument that distributes value among members is simultaneously the instrument that must finance the asset. We call the resulting tension the *financing problem* for community storage and analyze it formally in Section 3.

This financing problem is a consequence of the institutional arrangement, not of storage technology per se. If storage costs are covered by a non-energy instrument—a membership fee, a capacity payment, or a community budget—then the investment decision is decoupled from internal settlement and the community can evaluate storage at the full corridor. We analyze these organizational alternatives in Section 4.

⁶At an atom of F the derivative in (9) is a subgradient and the condition reads $\Pr(m > k^W) \leq \alpha \leq \Pr(m \geq k^W)$, the form used in Section 5 and Appendix B.4. Throughout we assume the $(1 - \alpha)$ -quantile is unique, i.e. that $\Pr(m > k) = \alpha$ does not hold on an interval of k ; otherwise S is flat on that interval, the efficient capacity is a set, and strict statements of the form $S(k^D) < S(k^W)$ for $k^D \neq k^W$ become weak.

2.6. Organizational regimes

We compare three regimes for financing community-level storage. Each corresponds to an arrangement observed or provided for in practice. Third-party storage operators remunerated through internal settlement are explicitly foreseen as participants in Swiss LECs ([Der Bundesrat, 2025](#)) and correspond to the merchant model of community batteries; commercial aggregators and battery-as-a-service providers that contract capacity to a community for a fixed rental fee resemble the platform regime on the financing margin; and cooperative energy communities that procure a shared asset and recover its cost through member contributions correspond to community ownership. The battery-as-a-service case matches the platform regime only in part: the participation result of Section 4.1 additionally requires control over per-member access pricing and endogenous membership, which a rental supplier need not have. Where it does not, competitive rental procurement run by the community reproduces the cost-recovery outcome without the participation distortion. The regimes are ordered by how far storage cost recovery is decoupled from the internal spread: not at all, partly through a capacity subsidy financed by a fee, and fully through a fee; the decentralized regime is the reference case against which the revenue-adequacy condition of Section 3 is evaluated.

Decentralized energy-only settlement. There is no intermediary. Storage operators recover costs solely from the internal margin $\Delta(k)$. A prosumer of type ϕ who joins obtains expected surplus $(\Delta_0 - \Delta(k))v(k) - \phi$ from improved internal terms of trade. This regime corresponds to most current energy-community implementations, including Swiss ZEVs and the majority of European collective self-consumption schemes, where internal exchange is settled through energy prices alone and no separate capacity remuneration mechanism exists.

Platform intermediation. A monopoly platform charges a membership fee f and a per-unit capacity charge τ levied on installed storage (a subsidy to storage if $\tau < 0$). A prosumer obtains

$$u(\phi) = (\Delta_0 - \Delta(k))v(k) - f - \phi. \quad (12)$$

Platform profit is $\Pi = fN + \tau K$. The platform can use τ to decouple storage cost recovery from the endogenous margin, creating a non-energy revenue stream for battery investment.

Community-owned storage. A non-profit community procures storage and finances it through a balanced-budget membership fee $f = ck$. Cost recovery is separated from internal energy settlement: whatever sharing rule or settlement rule the community adopts for day-to-day transactions, the *investment cost* of the battery is financed through the fee. As we show in Section 4.2, this separation allows the community to evaluate storage at the full corridor Δ_0 and restores unanimity among members on the efficient investment level.

Individual versus shared storage. Behind-the-meter household batteries exist in all three regimes and are absorbed into the mismatch distribution. A standalone household battery earns the full corridor without any internal settlement and so faces no financing problem of the kind we study, but it forgoes the gains from pooling that [Appendix C](#) sets out. In the multi-family ZEV we calibrate, individual units cannot install standalone storage that settles against the tariff corridor without the common connection point, so the relevant alternative to the shared battery is no battery; this is why the participation conditions below benchmark joining against zero storage surplus.

3. Energy-Only Settlement and the Financing Gap

We now analyze the institutional baseline in most current energy-community implementations: decentralized local trade with no fixed fees and no side transfers ($f = 0, \tau = 0$), where storage operators recover investment costs solely from local arbitrage revenues. This is the regime Swiss LECs provide for by admitting third-party storage operators as participants, and it is the default wherever a collective self-consumption scheme settles internal exchange through energy prices alone. The calibrated ZEV, which nets at the connection point and has no internal market, is the physical setting on which we overlay it (Section 5.4). The question is whether this arrangement can finance the surplus-optimal battery.

3.1. The free-entry condition

Let $N > 0$ and recall that capacity per member is $k \equiv K/N$. Expected trading revenue for storage is $\Delta(k) N v(k)$, and the total resource cost is cNk . Under competitive entry, any equilibrium with positive storage satisfies

$$\Delta(k) v(k) = ck. \quad (13)$$

This is an *average revenue equals average cost* condition: storage enters until its expected arbitrage income per unit of capacity,

$$g(k) \equiv \Delta(k) \frac{v(k)}{k}, \quad (14)$$

equals the levelized cost c . The condition is an *average-revenue* condition because daily dispatch is rationed across installed units in proportion to capacity: on a day with $m(\omega) < k$ each unit of capacity shifts $m(\omega)/k$ kWh and earns $\Delta(\omega, k) m(\omega)/k$, even though a marginal unit adds nothing to shifted energy that day. Price-taking alone would not deliver this payoff; pro-rata dispatch does, and it is the natural convention for a fleet of identical units behind one connection point. It is also what makes the utilization-decline force of Section 3.3 operate under a constant spread or the full-deficit SDR base, where the margin stays positive on those days; scarcity clearing and matched-market SDR shut that force down by driving the margin to zero whenever $m(\omega) \leq k$. Under Assumption 1 and concavity of v , g is weakly decreasing: $v(k)/k$ is weakly decreasing because v is concave with $v(0) = 0$, and $\Delta(k)$ is weakly decreasing by assumption. Throughout the equilibrium analysis we use the slightly stronger property that g is *strictly* decreasing wherever it is positive; this is proven for matched-market SDR under any mismatch distribution (Corollary 1) and in closed form for the uniform benchmark (Corollary 3); for the full-deficit base it is verified numerically under the empirical distribution.

The planner's optimality condition (9), by contrast, is a *marginal* condition: an additional unit of capacity is worthwhile if $\Delta_0 v'(k) = c$. The planner evaluates capacity at the full corridor Δ_0 times the marginal increase in shifted energy $v'(k)$; the market evaluates it at the endogenous margin $\Delta(k)$ times average shifted energy per unit of capacity $v(k)/k$. These two conditions coincide only under restrictive circumstances, for reasons we now develop.

Lemma 1 (Decentralized equilibrium). *Let g be strictly decreasing where positive. If $\lim_{k \downarrow 0} g(k) = \Delta(0^+) v'(0) > c$, the free-entry condition (13) has a unique solution $k^D > 0$; otherwise no positive storage investment breaks even and $K^D = 0$.*

Proof. g is continuous on $(0, \infty)$ under Assumption 1, and $g(k) \leq \Delta_0 \mathbb{E}[m]/k \rightarrow 0$ as $k \rightarrow \infty$; existence follows from the intermediate value theorem and uniqueness from strict monotonicity. If $\Delta(0^+) v'(0) \leq c$, then $g(k) < c$ for all $k > 0$. $\square$

3.2. The revenue-adequacy condition

The question is whether the decentralized capacity per member k^D coincides with the planner's k^W . The planner's condition (9) is $c = \Delta_0 v'(k^W)$; the free-entry condition (13) is $c = \Delta(k^D) v(k^D)/k^D$. For $k^D = k^W$ these require

$$\frac{\Delta(k^W)}{\Delta_0} = \frac{v'(k^W)}{v(k^W)/k^W}. \quad (15)$$

We call this the *revenue-adequacy condition*. The right-hand side is determined by the curvature of the shifting function—a physical property of the mismatch distribution. The left-hand side is determined by how much the settlement rule compresses the corridor at the efficient capacity—a design property of the internal market. When the condition holds, arbitrage revenue at the efficient capacity exactly covers cost and competitive entry delivers $k^D = k^W$; when it fails, the decentralized outcome diverges from the planner's.

One rule satisfies (15) for every mismatch distribution: scarcity clearing.⁷ If the matchable market cleared at Walrasian prices, storage would pay $\underline{p}$ and receive $\bar{p}$ on every day with $m(\omega) > k$ —there is excess demand

⁷For distributions with atoms the rule must be completed with a price convention on days where supply exactly equals demand; see the discussion of equality states in Appendix B.4.

for stored energy at any price below $\bar{p}$ —and the margin would be competed to zero on days with $m(\omega) \leq k$, so that $\Delta(\omega, k) = \Delta_0 \mathbf{1}\{m(\omega) > k\}$. Then $\Delta(k) v(k)/k = \Delta_0 \Pr(m > k) = \Delta_0 v'(k)$: average revenue equals the planner’s marginal value and free entry delivers $k^D = k^W$ (Appendix B.4). This is the community-scale counterpart of the cost-recovery results for competitive storage in wholesale markets (Schmalensee, 2022; Korpås and Botterud, 2020). Scarcity clearing achieves this while paying storage strictly less than the full corridor: its compression ratio is $\Delta(k)/\Delta_0 = v'(k)/(v(k)/k)$ at every k , which equals $1 - \mathbb{E}[m \mathbf{1}\{m \leq k\}]/v(k)$ and falls below one exactly when $\Pr(0 < m \leq k) > 0$, i.e. as soon as the battery is only partly filled on a positive-probability set of days; days with no matchable energy contribute neither revenue nor throughput and do not compress. Being paid less than Δ_0 is therefore not by itself the source of the financing gap; what matters is whether the compression a rule imposes matches the curvature of the shifting technology.

The canonical rules are not scarcity-clearing prices. A time-invariant price sets the left-hand side to zero. An administrative spread fixes it at δ/Δ_0 irrespective of curvature: it satisfies (15) if and only if $\delta = \delta^* \equiv \Delta_0 v'(k^W)/(v(k^W)/k^W)$, which requires knowing F and α and ceases to hold when either moves (Corollary 2). With uniform mismatch on $[0, \bar{m}]$ and $\alpha = \frac{1}{2}$ this is $\delta^* = 2\Delta_0/3$, and at the calibration of Section 5 it is 12.32 cents/kWh (Table 7). SDR settlement maps the supply–demand ratio smoothly into prices, so the discharge price falls below $\bar{p}$ as soon as any capacity is installed, rather than only once the matchable energy is exhausted; the margin is compressed before scarcity is relieved. Section 5 quantifies the gap under SDR and shows that it is not a marginal friction.

Proposition 1 (Energy-only settlement and the financing gap). *The decentralized energy-only equilibrium implements the efficient capacity per member, $k^D = k^W$, if and only if the revenue-adequacy condition (15) holds—that is, if and only if the settlement rule’s compression ratio $\Delta(k^W)/\Delta_0$ equals the curvature ratio $v'(k^W)/(v(k^W)/k^W)$. Scarcity-clearing settlement satisfies the condition for every mismatch distribution. Among the canonical rules: under a time-invariant internal price $k^D = 0$; under a constant administrative spread the market over-invests, under-invests, or implements k^W exactly according as $\delta \gtrless \delta^*$ (Corollary 2); and under SDR settlement applied to the matchable market it strictly under-invests for every mismatch distribution (Corollary 1).*

Corollary 1 (Matched-market SDR under-invests for every mismatch distribution). *Let SDR settlement be applied to the matchable market (Appendix B.2) with $0 < \underline{p} < \bar{p}$, let F be any mismatch distribution with $\Pr(m > 0) > \alpha > 0$ —so that $k^W > 0$. Then g is strictly decreasing wherever positive, $g(k^W) < c$, and $k^D < k^W$.*

Proof. On $\{m(\omega) > k\}$ the period-2 ratio is $\rho = k/m \in (0, 1)$ and, by (B.15), $\Delta^{\text{SDR}}(\omega, k) = \underline{p} \Delta_0 (1 - \rho)/(p + \Delta_0 \rho) < \Delta_0$, since the inequality reduces to $-\underline{p} \rho < \Delta_0 \rho$. On $\{m(\omega) \leq k\}$ we have $\rho = 1$ and $\Delta^{\text{SDR}} = 0$. Since $q(\omega, k) = k$ on the first event and the margin vanishes on the second,

$$\mathbb{E}[\Delta^{\text{SDR}}(\omega, k) q(\omega, k)] = k \mathbb{E}[\Delta^{\text{SDR}}(\omega, k) \mathbf{1}\{m > k\}] < \Delta_0 k \Pr(m > k),$$

strictly whenever $\Pr(m > k) > 0$. Dividing by k gives $g(k) < \Delta_0 \Pr(m > k) = \Delta_0 v'(k)$. Evaluating at k^W and using the planner’s condition (9), read at atoms as $\Pr(m > k^W) \leq \alpha \leq \Pr(m \geq k^W)$, yields $g(k^W) < \Delta_0 \alpha = c$; if $\Pr(m > k^W) = 0$ then $g(k^W) = 0 < c$ directly. For monotonicity, note that the margin vanishes on $\{m \leq k\}$ and $q = k$ on $\{m > k\}$, so $g(k) = \mathbb{E}[\Delta^{\text{SDR}}(\omega, k) \mathbf{1}\{m > k\}]$. Take $k_1 < k_2$. On $\{m > k_2\}$ the ratio $\rho = k/m$ is larger at k_2 and (B.15) is strictly decreasing in ρ , so the k_2 term is strictly smaller; on $\{k_1 < m \leq k_2\}$ the k_2 term is zero while the k_1 term is positive; elsewhere both vanish. Hence $g(k_2) < g(k_1)$ whenever $\Pr(m > k_1) > 0$ —strict monotonicity where positive, for every F and without appeal to Assumption 1. Since $g(k^D) = c$ at any positive k^D , it follows that $k^D < k^W$. $\square$

Equivalently, the compression ratio $\Delta^{\text{SDR}}(k)/\Delta_0$ lies strictly below the curvature ratio $v'(k)/(v(k)/k)$ at every k with $\Pr(m > k) > 0$, which is the form that connects the result to (15) and Corollary 2. Corollary 4 remains the uniform case, supplying the closed-form magnitude of the shortfall and its scaling in $\bar{p}/\underline{p}$, which Corollary 1 does not.

3.3. Two forces, and who receives the corridor

Why do the canonical rules fail the condition? Two forces create a gap between the market’s average-revenue condition and the planner’s marginal condition—but they work in opposite directions, which is why the distortion varies across settlement regimes.

The first is *utilization decline*. Concavity of v with $v(0) = 0$ implies that average utilization $v(k)/k$ strictly exceeds marginal utilization $v'(k)$ for all $k > 0$: as the battery grows, it reaches full utilization on fewer days. This gap between average and marginal product is familiar from any decreasing-returns technology. In the storage context, it means that a competitive entrant earning average revenue per unit of capacity finds the market *more* attractive at k^W than the planner’s marginal condition would warrant. Taken alone, this force pushes toward *over*-investment: the market keeps entering past the point where the last unit of capacity is socially worthwhile, because it earns the average—not the marginal—return.

The second is *spread compression*. Under Assumption 1, the expected margin $\Delta(k)$ is weakly decreasing in capacity. As the battery expands, it smooths local surplus-deficit imbalances, improving internal terms of trade for prosumers but eroding the per-kWh arbitrage rent that storage earns. This force pushes toward *under*-investment: the market earns the compressed margin $\Delta(k) < \Delta_0$ rather than the full corridor that defines social returns.

Spread compression is a pecuniary externality among storage units. Because the margin $\Delta(k)$ depends on total capacity, each additional unit relieves local scarcity and lowers the rent earned by every other unit; that rent does not disappear but is transferred to prosumers, who obtain better internal terms of trade. The social value of the marginal unit, $\Delta_0 v'(k)$, is unaffected by this transfer. What matters for investment is who receives it: under energy-only settlement the marginal unit of capacity creates value Δ_0 per shifted kWh, of which storage is paid only $\Delta(k)$, while the remainder $\Delta_0 - \Delta(k)$ accrues to prosumers as an uncompensated benefit. Storage is thus remunerated at a fraction of the value it creates. Under the canonical scarcity-reflective rules the fraction shrinks faster than the curvature of the shifting technology warrants, so that entry halts once the compressed margin no longer covers cost. This is the sense in which centralization restores efficiency: a planner or a community-owned battery (Section 4) is the residual claimant of the whole corridor and, by financing storage through an instrument that does not depend on the margin, evaluates capacity at its full value Δ_0 .

Formally, let $R(k) \equiv \Delta(k) v(k)$ denote expected energy-only revenue of the fringe per participant. Wherever Δ is differentiable,

$$R'(k) = \underbrace{\Delta(k) v'(k)}_{\text{marginal unit's own margin}} + \underbrace{\Delta'(k) v(k)}_{\text{change in the inframarginal margin} \leq 0}. \quad (16)$$

The second term is, to first order, the revenue the marginal unit removes from all inframarginal throughput and hands to prosumers. Because $\Delta(k)$ is volume-weighted, it mixes a price effect with a change in the state weights $q(\omega, k)$; the decomposition is exact either way, and neither the free-entry nor the revenue-adequacy condition depends on how it is labelled. No free entrant faces it: each operator earns the average $g(k) = R(k)/k$. The decomposition also shows what the under-investment does *not* stem from: the fringe’s failure to internalize this term. A single storage owner who did internalize it would maximize $R(k) - ck$ and stop where $R'(k) = c$; since $R'(k) = g(k) + k g'(k) < g(k)$, it would install *less* than the free-entry fringe, which stops where $g(k) = c$.⁸ The distortion lies in the gap between what the fringe is paid per unit of capacity, $g(k) = \Delta(k) v(k)/k$, and the marginal value of that unit, $\Delta_0 v'(k)$ —not in the compression term as such, and not in the fact that $\Delta(k) < \Delta_0$, which scarcity clearing shares (Section 3.2, Table 7).

When the alignment fails, the direction of the distortion follows from which side of (15) is larger.

⁸Under SDR settlement with a uniform mismatch distribution and the prices of Section 5, the single owner installs 43–47% of the free-entry capacity across $\alpha \in [0.5, 0.9]$. Market power in storage (Williams and Green, 2022) therefore compounds the financing gap rather than correcting it. How market power in generation would interact with it lies outside a model in which generation is fixed; Andrés-Cerezo and Fabra (2023) show that its effect on storage investment depends on market structure.

Table 1: Settlement rules and the financing gap for community-level storage. The revenue-adequacy condition (15) holds under scarcity-clearing settlement for every mismatch distribution, and under a constant administrative spread only at $\delta = \delta^*$.

Settlement rule	$\Delta(k)$	Direction	Financing gap at k^W
Time-invariant price	0	$k^D = 0$	100% of cost; no storage revenue whatsoever
Constant spread buy-sell	δ	$k^D \gtrless k^W$	Depends on δ : over- or under-investment, zero at $\delta = \delta^*$
Scarcity-reflective (e.g. SDR settlement)	Decreasing	$k^D \leq k^W$ in general; $k^D < k^W$ under matched-market SDR (Cor. 1)	Potentially large; $\ll k^W$ under the matched-market base in the calibration
Scarcity-clearing (Walrasian; not used in practice)	Δ_0 if $m > k$, 0 otherwise	$k^D = k^W$	Zero (Appendix B.4)

Corollary 2 (Direction of the distortion). *Suppose g is strictly decreasing where positive and $k^D > 0$. Then*

$$k^D \leq k^W \iff \frac{\Delta(k^W)}{\Delta_0} \leq \frac{v'(k^W)}{v(k^W)/k^W}.$$

The market under-invests exactly when the settlement rule compresses the margin at the efficient scale by more than the curvature discount of the shifting technology, and over-invests in the opposite case.

Proof. Since g is strictly decreasing where positive and $g(k^D) = c$, we have $k^D < k^W$ iff $g(k^W) < c$ (and symmetrically for $>$). Substituting $c = \Delta_0 v'(k^W)$ from (9) and $g(k^W) = \Delta(k^W) v(k^W)/k^W$ and rearranging yields the display. $\square$

Under SDR applied to the matchable market, compression more than offsets the utilization effect and the market under-invests for every cost level—a claim we prove for every mismatch distribution (Corollary 1) and quantify under the uniform-mismatch benchmark (Corollary 4), where the capacity ratio k^D/k^W falls from about 20% at $\alpha = 0.5$ to 9% at $\alpha = 0.9$ and approaches one only as $\alpha \rightarrow 0$ —and the shortfall is large across the cost range we calibrate in Section 5.

The qualification of the SDR result by convention in Proposition 1 matters. Under SDR the strength of compression depends on the denominator of the period-2 ratio. If that denominator is the full deficit $d(\omega)$ rather than the matchable energy $m(\omega)$, the ratio is q/d , and as d/m grows the discharge price stays near $\bar{p}$ for any installed capacity: compression vanishes, the rule approaches a constant spread of Δ_0 , and Corollary 2 then implies *over*-investment. With $m \sim U[0, \bar{m}]$ and $d = \lambda m$, revenue at k^W falls short of cost for λ below a threshold that rises with the cost-corridor ratio. That comparison does not aggregate to the heterogeneous case, however, because $d(\omega)/m(\omega)$ varies widely across days; the sign under the full-deficit base is therefore an empirical matter, settled by direct computation in Section 5, where the market under-invests in every scenario we compute.

3.4. The financing gap across settlement regimes

Table 1 summarizes how the financing gap varies across the settlement rules introduced in Section 2.4.

The one rule that finances efficient storage for every mismatch distribution—scarcity clearing—is not used by any community implementation we are aware of. The reason is not member welfare within the model: with free entry, prosumers receive $S(k^D)$ in expectation under any rule (Section 3.5), so scarcity clearing, which delivers $k^D = k^W$, gives members the same expected surplus as community ownership. It lies in what the symmetric, risk-neutral payoffs omit. On days when the battery cannot absorb the matchable surplus the rule assigns the entire corridor to storage and nothing to prosumers, so members see no benefit precisely when the battery is most valuable, bear the resulting bill volatility, and—where some are net buyers and others net sellers—divide the corridor unequally; the canonical rules are adopted precisely to share the corridor state by state. Sharing it that way is compatible with financing the asset only at one

calibrated spread. Matched-market SDR, which shares it in every state, over-compresses for every mismatch distribution; whether another scarcity-reflective rule does so depends on the comparison in (15).

The per-unit financing gap at the efficient capacity per member is $c - g(k^W)$, positive whenever the market under-invests; under a time-invariant price it equals the full cost c , and under SDR settlement—the scarcity-reflective rule we use in the calibration—it is quantified in Section 5. The revenue shortfall has a family resemblance to the “missing money” problem in wholesale electricity markets, where energy-only prices may not cover peaker fixed costs (Joskow, 2008; Cramton et al., 2013). The mechanism here is different. In wholesale markets, the missing-money problem arises because administrative price caps (set below the value of lost load) truncate scarcity rents that would otherwise remunerate peaking capacity. In our setting, the external corridor $[\underline{p}, \bar{p}]$ does bound internal prices, but this bound is not the source of the revenue shortfall: even if internal prices could freely reach $\bar{p}$ in every deficit period, the endogenous spread compression under scarcity-reflective settlement would still erode storage revenue as capacity expands. There is also no reliability externality. But the policy implication is similar: a non-energy revenue stream is needed to finance capacity that the energy market alone cannot sustain.

Environment dependence. The financing gap varies across communities. Both $v(k)$ and $\Delta(k)$ depend on the mismatch distribution F , which reflects community composition: adding non-PV consumers introduces within-period matching that changes the residual intertemporal component—reducing it where the surplus side binds, which in our data is about half the days, but not necessarily elsewhere—and with it both k^W and $g(k)$. Seasonal variation has a similar effect. Accordingly, no administrative or formula-based settlement rule delivers robust storage incentives across all configurations, and only scarcity clearing does so by construction. The same holds across regulatory frameworks: the corridor width Δ_0 depends on which network charges internal flows avoid—the full volumetric tariff in a Swiss ZEV, a reduced network charge in a LEC, and country-specific reductions under RED II transpositions—and it scales the levels of k^W and k^D through $\alpha = c/\Delta_0$ and, under SDR, also through the price ratio $\bar{p}/p$, which enters the margin separately (Section 5 reports both sensitivities); under the matched-market SDR convention it leaves the sign of the gap unchanged, since $g(k) < \Delta_0 \Pr(m > k)$ whenever that tail is positive (Corollary 1).

3.5. Consequences for participation

The distortion in capacity per member propagates to participation. Under energy-only settlement at capacity per member k^D , a prosumer joins if the surplus from improved internal terms of trade covers the adoption cost:

$$\phi \leq (\Delta_0 - \Delta(k^D)) v(k^D). \quad (17)$$

Using the free-entry condition (13), this cutoff equals $\Delta_0 v(k^D) - ck^D = S(k^D)$. With adoption costs uniformly distributed on $[0, \bar{\phi}]$, participation is

$$N^D = \min \left\{ 1, \frac{S(k^D)}{\bar{\phi}} \right\}. \quad (18)$$

Since k^W is the unique maximizer of $S(k)$, any departure $k^D \neq k^W$ implies $S(k^D) < S(k^W)$ and therefore $N^D \leq N^W$ (with strict inequality whenever the cap at 1 does not bind); under a revenue-adequate rule $k^D = k^W$ and participation is undistorted.

The mechanism is straightforward: an undersized battery generates less surplus per participant, which makes community membership less attractive, which reduces the number of households that join. This is a second-order consequence—the primary distortion is in investment—but it amplifies the welfare loss. With $N^D = S(k^D)/\bar{\phi}$ (interior case), total welfare (7) is $W^D = S(k^D)^2/(2\bar{\phi})$, so relative to the planner’s $W^W = S(k^W)^2/(2\bar{\phi})$ the decentralized community captures the *square* of the per-participant surplus ratio, $W^D/W^W = (S(k^D)/S(k^W))^2$. Section 5 reports this ratio for the Swiss community.

4. Organizational Responses

The preceding section showed that the decentralized energy-only regime—the institutional default in most current implementations—cannot reliably finance efficient community-level storage. We now ask whether alternative organizational forms perform better.

We analyze two alternatives: platform intermediation and community ownership. Both introduce a *non-energy* financing instrument—a payment to storage not tied to the per-kWh internal margin. They differ in who controls the instrument, how participation is affected, and therefore in the share of community surplus that is ultimately realized.

4.1. Platform intermediation

At the efficient capacity per member k^W , storage earns $g(k^W)$ per unit of capacity from energy-only settlement while its cost is c ; the per-unit financing gap is

$$c - g(k^W). \quad (19)$$

This is positive under every rule that under-invests at k^W —the decentralized case of Section 3 and of the calibration—negative under an over-investing administrative spread, and zero under a revenue-adequate one. Closing it requires a transfer to storage that is not tied to the per-kWh margin, and the two responses below differ in whether that transfer is administered for profit or at cost.

Consider a monopoly platform that mediates between prosumers and storage operators. This is the closest analogue in the literature to the regime [Dinh et al. \(2022\)](#) study computationally: their market aggregator is a Stackelberg leader that sets internal prices and the battery’s schedule and sizes the asset to maximize its own profit. The platform has two instruments: a membership fee f charged to each participating prosumer, and a per-unit capacity charge τ levied on storage operators, so that $\tau < 0$ is a subsidy paid to them. Both are set ex ante and do not vary with the daily state.

Investment. Under the platform regime, a storage operator’s profit per unit of capacity is $g(k) - c - \tau$. Competitive entry requires

$$g(k) = c + \tau. \quad (20)$$

If the platform sets $\tau < 0$, it lowers the break-even threshold. The transfer selects a unique capacity provided $g(k^W) > 0$ and g is strictly decreasing where positive. Under a time-invariant internal price $g \equiv 0$, so $\tau = -c$ satisfies (20) at *every* capacity and the two price instruments do not pin down k ; the platform must then contract for capacity directly, which this model does not represent.⁹ To implement the efficient capacity per member k^W , the platform chooses

$$\tau^* = g(k^W) - c. \quad (21)$$

The transfer is a subsidy when (19) is positive and a capacity charge when it is negative; under every rule we calibrate it is a subsidy equal to the financing gap. It bridges the shortfall between energy-only revenue and full cost, allowing competitive storage operators to break even at the efficient scale. The platform finances this subsidy from prosumer membership fees.

Participation and platform optimization. Let k^P denote the equilibrium capacity per member under the platform regime. A prosumer who joins obtains $(\Delta_0 - \Delta(k^P))v(k^P) - f - \phi$, so the marginal participant has $\phi = (\Delta_0 - \Delta(k^P))v(k^P) - f$ and participation is $N^P = [(\Delta_0 - \Delta(k^P))v(k^P) - f]/\phi$. The platform chooses (f, τ) to maximize

$$\Pi = f N^P + \tau k^P N^P, \quad (22)$$

⁹The hypothesis holds throughout the calibration: at K^W the compression ratio is 0.240 under the matched-market base and 0.380 under the full-deficit base (Table 7), giving $g(K^W) = 2.67$ and 4.22 cents per unit of capacity respectively—both strictly positive, and both below $c = 9.79$.

subject to competitive storage entry (20) and this participation constraint. Since the platform's payoff is increasing in per-participant surplus $S(k)$, it chooses τ^* to implement $k^P = k^W$ (as in (21)), then sets f to maximize Π . Substituting and optimizing over f yields

$$f^* = \frac{S(k^W)}{2} + (c - g(k^W)) k^W, \quad (23)$$

where the first term is the standard monopoly markup—half the per-participant surplus at the efficient capacity per member—and the second term recovers the storage subsidy from prosumers. The resulting participation is

$$N^P = \frac{S(k^W)}{2\bar{\phi}} = \frac{1}{2} N^W. \quad (24)$$

The efficient capacity per member is the platform's *optimum*, not an assumption: at the optimal fee the platform's reduced-form profit is $\Pi^*(k) = S(k)^2/(4\bar{\phi})$, which inherits the single peak of $S(\cdot)$ at k^W , so the platform sets τ^* exactly to implement $k^P = k^W$.

Proposition 2 (Platform intermediation). *Let $g(k^W) > 0$, let g be strictly decreasing where positive, and consider the interior case $S(k^W) \leq \bar{\phi}$. A profit-maximizing platform implements the efficient capacity per member $k^P = k^W$ through the capacity transfer $\tau^* = g(k^W) - c$ (a subsidy when $g(k^W) < c$) and the fee $f^* = S(k^W)/2 + (c - g(k^W)) k^W$. Participation is $N^P = \frac{1}{2} N^W$: the platform delivers efficient capacity per member but serves only half the efficient number of members, so aggregate capacity is $K^P = \frac{1}{2} K^W$.*

For $S(k^W) > \bar{\phi}$ the cap $N \leq 1$ binds for the planner; the platform then serves $N^P = \min\{1, S(k^W)/(2\bar{\phi})\}$, and once $S(k^W) > 2\bar{\phi}$ both regimes serve everyone. Capacity per member remains k^W throughout.

Assessment. The platform achieves efficient capacity per member by using the capacity transfer to decouple cost recovery from the endogenous margin. But it serves only half the efficient number of participants. This is the standard monopoly-pricing distortion on a uniform distribution and is not specific to the energy-community setting; we include it to complete the institutional comparison. The platform pays storage the value it creates through the subsidy but extracts the resulting surplus through the fee, excluding prosumers whose adoption costs exceed the net-of-fee benefit.

The tradeoff is sharp. The platform solves the investment problem completely— $k^P = k^W$ —but creates a participation problem that the decentralized regime does not have. Under energy-only settlement, participation is distorted only indirectly, through the suppressed surplus caused by under-investment. Under the platform, participation is distorted directly, through the fee. With $N^P = N^W/2$ and $k^P = k^W$, total welfare (7) is

$$W^P = N^P S(k^W) - \frac{\bar{\phi}}{2} (N^P)^2 = \frac{3}{4} W^W, \quad (25)$$

of which the platform's profit $\Pi^* = S(k^W)^2/(4\bar{\phi})$ is two-thirds and prosumer surplus one-third. The platform regime thus achieves three-quarters of the potential welfare, against the squared surplus ratio $(S(k^D)/S(k^W))^2$ of the decentralized regime (Section 3.5). The distributional consequences differ, however: the platform concentrates the remaining surplus in its own hands rather than leaving it with prosumers.

A regulated platform facing a price cap on f could in principle approach the efficient participation level. We do not model regulation explicitly, but the implication is clear: the platform regime is most attractive when the regulator can constrain the fee without undermining the capacity subsidy. This is a non-trivial regulatory design problem, since the fee finances the subsidy and capping one constrains the other.

4.2. Community-owned storage

Now consider a non-profit community that procures and owns storage collectively, financing it through a balanced-budget membership fee. There is no platform; the community itself decides on the capacity per member k and sets the fee to cover costs.

Budget balance and fee structure. The community's cost for capacity per member k is ck per participant per day. Under a balanced-budget rule, the membership fee is

$$f^C = ck. \quad (26)$$

Internal settlement of energy flows—whatever pricing or sharing mechanism the community adopts for day-to-day transactions—determines how prosumers transact with one another and with the battery on a day-to-day basis. But the *investment cost* of the battery is financed through the fee, not through the energy margin. This is the key structural distinction from the decentralized regime.

Investment. A prosumer of type ϕ who joins obtains expected surplus

$$u^C(\phi) = \Delta_0 v(k) - ck - \phi = S(k) - \phi. \quad (27)$$

This expression is independent of the internal settlement rule. Whatever operating revenue the community-owned battery earns from internal trades— $\Delta(k)v(k)$ per participant under any rule with margin $\Delta(k)$ —accrues to the community itself and is rebated to members, or offsets the fee, in the same proportions in which the fee is levied; one member's payment is another's receipt, and the transfers net to zero within the cooperative. What remains is the gross value of shifting $\Delta_0 v(k)$ minus the resource cost ck , which is the value-independence property of [Gautier et al. \(2025\)](#).¹⁰ Because every member's payoff is $S(k) - \phi$, every member's preferred capacity is the one that maximizes S , which is the planner's k^W : with cost recovery separated from the energy margin, the community evaluates storage at Δ_0 rather than at $\Delta(k)$, the average-marginal gap of Section 3 does not arise, and members are unanimous whatever rule—per-capita sharing, pro-rata allocation, or scarcity-reflective settlement—governs day-to-day exchange. Let k^C denote the community's choice; then $k^C = k^W$. Section 5 reports the fee and the net benefit per member for the Swiss community.

Participation. At the efficient capacity per member, a prosumer joins if $S(k^W) \geq \phi$, giving

$$N^C = \frac{S(k^W)}{\bar{\phi}} = N^W. \quad (28)$$

Community ownership achieves both efficient investment and undistorted participation: $N^C = N^W$, the planner's benchmark level. (This is generally less than full uptake, since prosumers with high adoption costs optimally remain outside the community.) There is no surplus extraction, no participation distortion, and no need for regulatory oversight of fees.

Proposition 3 (Community ownership). *Under community ownership financed by the balanced-budget fee $f^C = ck$, every member unanimously prefers the efficient capacity per member $k^C = k^W$, regardless of the internal settlement rule, and participation is undistorted, $N^C = N^W$.*

Heterogeneous members and fee design. Proposition 3 treats members as symmetric. With heterogeneous mismatch profiles, member i 's gross benefit from the community battery is $\Delta_0 v_i(k)$, and a *uniform* fee ck can violate individual rationality for low-mismatch members— $\Delta_0 v_i(k^W) < ck^W$ —even though k^W maximizes aggregate surplus: equal cost contributions sit poorly with unequal benefits. Two observations follow. First, heterogeneity affects fee design, not the investment result: k^W maximizes aggregate surplus for any benefit profile. Second, benefit-proportional cost shares restore participation and, under a natural condition, unanimity. If profiles are scaled versions of a common shape and capacity is allocated in proportion to them, so that $v_i(k) = \theta_i v(k)$ with mean θ normalized to one, then the benefit-proportional fee is $f_i = \theta_i ck$ and, with operating revenue rebated in the same proportions, member i 's net payoff is $\theta_i S(k)$ (a per-capita rebate

¹⁰The settlement rule still determines how *individual* members' benefits depend on their consumption profiles and sharing coefficients; only the *aggregate* benefit per participant is rule-independent.

Table 2: Organizational regimes for community-level storage. k^W and N^W denote the planner’s efficient capacity per member and participation. The decentralized row reports the SDR case used in the calibration; under a time-invariant price $k^D = 0$, and under a revenue-adequate rule (scarcity clearing, or $\delta = \delta^*$) $k^D = k^W$ and $N^D = N^W$. The internal settlement rule governs day-to-day exchange in all three regimes and is not what distinguishes them.

Regime	Who installs and owns storage	Internal settlement rule	Financing instrument	Investment	Participation
Decentralized	Third-party operators, competitive entry	Any; efficient only if (15) holds (SDR in calibration)	Internal only	spread $k^D < k^W$ under matched-market SDR ($\ll$ in the calibration); $k^D = k^W$ if (15) holds	$N^D \leq N^W$
Platform	Third-party operators, contracted by a for-profit platform	Any with $g(k^W) > 0$	Membership fee + capacity subsidy	k^W	$\frac{1}{2}N^W$
Community-owned	The community (non-profit)	Any	Balanced-budget membership fee	k^W	N^W

would instead give $\theta_i S(k) + (1 - \theta_i) R(k)$, which need not peak at k^W : every member’s payoff is a positive multiple of the same function, so all members unanimously prefer k^W and all gain, gross of adoption costs, whenever $S(k^W) > 0$. For genuinely heterogeneous shapes, the choice among participation-compatible fee structures is the community-level analogue of the pricing-and-sharing problem studied by [Balegamire and Gautier \(2025\)](#); it determines who joins and how surplus is divided; for a given membership it does not change the efficient scale of the asset, although with voluntary entry and exit the membership a fee structure induces—and hence the pooled mismatch distribution and the surplus-maximizing capacity—can differ across structures. ZEV membership is fixed by building ownership (Section 5), so this margin is absent from the calibration.

4.3. Comparing the three regimes

Table 2 summarizes outcomes across organizational regimes. Three patterns emerge.

First, there is a hierarchy of financing instruments. The decentralized regime relies solely on the endogenous energy margin; the platform adds a capacity subsidy financed by a membership fee; the community replaces the margin entirely with a budget instrument. Each step further decouples storage cost recovery from the internal spread. The platform corrects capacity per member but adds a participation distortion, so in the interior case of Proposition 2 it dominates the decentralized regime only when $(S(k^D)/S(k^W))^2 < 3/4$, i.e. when the decentralized regime captures less than 87% of the per-member surplus—which holds in every matched-market scenario of Table 5 but not at $\alpha = 0.9$ under the full-deficit base (Table 6). Community ownership dominates both.

Second, among the regimes that rely on a non-energy instrument, the investment–participation tradeoff is resolved only under community ownership. The platform achieves efficient capacity per member but serves only half the efficient membership, so aggregate capacity is $K^W/2$. This exclusion is not driven by battery physics or settlement design; it is driven by market power. The relevant institutional margin is therefore not “platform versus community” per se, but whether the non-energy financing instrument is administered for profit or on a cost-recovery basis. A competitive intermediary—or a regulated platform constrained to charge at cost—would replicate the community-ownership outcome. What matters is the governance of the fee, not its existence.

Third, the magnitudes are substantial. Under community ownership, per-participant surplus equals $S(k^W)$ and participation reaches N^W , so the community captures the full potential welfare. The platform achieves the same per-participant surplus (since $k^P = k^W$), but monopoly pricing admits only half the efficient number of participants; total welfare is three-quarters of the community-ownership benchmark (25). The decentralized regime loses on both margins: per-participant surplus falls because of under-investment, participation contracts proportionally, and total welfare falls to the square of the surplus ratio (Section 3.5; Section 5 gives the number). These are not marginal differences in allocative efficiency; they represent qualitatively different levels of community functionality under the matched-market base. The choice of

organizational form determines whether the community can capture the value that local time-shifting could create.

The logic extends beyond energy communities. Partial capture of the value an asset creates does not by itself imply under-remuneration: scarcity clearing pays storage less than Δ_0 and still finances it exactly (Appendix B.4). What does is that per-unit energy revenue at the efficient scale falls short of marginal value, $g(k^W) < c$ —that the rule’s compression exceeds the curvature of the technology, condition (15). Whenever that holds for a durable asset, a capacity-type instrument is needed to close the gap. Wholesale markets furnish the familiar analogue, where energy-only prices are argued to leave a “missing money” shortfall for peaking capacity (Joskow, 2008; Cramton et al., 2013; Newbery, 2016); there the shortfall arises from administrative price caps, whereas here the compression is endogenous. Community storage is a particularly clean instance, because the corrective instrument—a membership fee—is internal to the community rather than a regulated market-wide payment.

4.4. Qualifications and extensions

The theoretical ranking in Table 2 abstracts from governance and transaction costs. In practice, community ownership requires collective decision-making on investment scale, technology, and cost allocation—decisions that can be costly in heterogeneous communities. Platform intermediation economizes on governance by centralizing decisions but transfers surplus to the operator. The choice between forms therefore involves a tradeoff between governance costs and allocative efficiency that our model does not capture but that is central to policy design.¹¹

Several extensions deserve note. First, platform competition drives fees toward cost recovery and replicates the community-ownership outcome (Appendix B.6), though local energy communities may constitute natural monopolies at the neighborhood scale, limiting the scope for entry. Second, the community could finance storage through instruments other than a flat fee—usage-based levies, investment shares, or blended financing—provided the non-energy component covers the gap (19). For symmetric members the specific instrument matters for distributional equity but not for investment efficiency; with heterogeneous members different cost shares change individual preferred capacities (Section 4.2), so the instrument can matter for whether collective choice lands on k^W even though k^W itself is unchanged. Third, when multiple communities share a distribution network, their storage investments interact through grid-level externalities not captured by $S(k)$, potentially justifying additional public support. Fourth, we model the storage fringe as serving a single community. In practice a supplier may hold contracts with several communities or other flexibility buyers. Such an outside option does not overturn the financing result (Appendix B.7 formalizes the argument): it raises the operator’s reservation revenue, so the community must cover the shortfall $\max\{r, c\} - g(k)$ through a non-energy instrument—the outside return sets the threshold, and internal arbitrage revenue still counts toward it—and a higher threshold lowers the capacity any one community attracts at a given internal margin. Where the rule under-invests at every cost level, as matched-market SDR does, the energy-only margin therefore remains insufficient to finance the efficient community-scale battery against the benchmark set by the opportunity cost of capacity; under the full-deficit convention the sign at that benchmark is again an empirical matter. The assumption is also institutionally grounded here: Swiss implementation ties community storage to time-shifting within the perimeter and constrains systematic grid-to-community arbitrage (Section 1). A full model of capacity allocation across markets is a natural extension, and we expect the need for a non-energy financing instrument to be robust to it.

Notwithstanding these qualifications, the theoretical ranking is unambiguous: among the organizational forms we consider, only community ownership (or its competitive equivalent) implements the efficient outcome on both the investment and participation margins under every settlement rule; the decentralized regime does so only under a rule that satisfies (15), and the platform on the investment margin alone.

¹¹Gui and MacGill (2018) discuss intermediary organizations that broker knowledge across communities and, in some cases, provide commercial, technical and financial support.

Table 3: The price corridor by regulatory framework, at the 2026 tariff components of [Appendix D.1](#). “Still paid” is the volumetric charge an internally exchanged kWh continues to bear; Δ_0 is the resulting corridor; α is evaluated at the benchmark resource cost $c = 9.79$ cents/(kWh·day). A positive battery requires $\alpha < \Pr(m > 0)$.

Framework	Still paid	Δ_0	α at $c = 9.79$
Behind a common connection point (ZEV)	0.00	13.98	0.70
Over the grid, one voltage level (−40%)	7.58	6.40	1.53
Over the grid, with transformation (−20%)	8.78	5.20	1.88

5. Calibration

We calibrate the framework to a Swiss ZEV using quarter-hourly smart-meter data to assess whether the institutional differences identified above are quantitatively important. The calibration has two goals. First, we discipline the physical side of the model by measuring the distribution of daily matchable energy that a community-level battery could physically shift from surplus to deficit periods given temporal generation constraints. Second, taking the regulated price corridor as given, we quantify the financing gap by comparing the surplus-optimal battery capacity against the capacity that energy-only settlement can sustain.

Because ZEV membership is determined by building ownership and cannot be adjusted at the storage-investment margin, we treat the number of participants N as fixed throughout this section. We therefore work directly with aggregate capacity K rather than capacity per member $k = K/N$; all results can be rescaled by N if desired.

5.1. ZEV setting and prices

In a ZEV, the community appears as a single end-consumer vis-à-vis the distribution network operator and is settled at the common connection point; internal exchanges are netted out behind the meter (see [EnergieSchweiz, 2021](#); [Verband Schweizerischer Elektrizitätsunternehmen \(VSE\), 2025](#)). Consequently, any internally self-consumed kWh reduces net imports and avoids the *volumetric* import tariff components (energy, volumetric network charges, and levies).

The value of shifting one kWh from a surplus to a deficit period is therefore the corridor $\Delta_0 = \bar{p} - \underline{p}$ of eq. (1), with $\bar{p}$ the retail price at the connection point and $\underline{p}$ the feed-in tariff.

In the benchmark calibration we set $\bar{p} = 21.98$ cents/kWh and $\underline{p} = 8.00$ cents/kWh, implying a corridor of $\Delta_0 = 13.98$ cents/kWh. Here and throughout, one cent is one Swiss centime (CHF 0.01). The retail figure is the *volumetric* stack a shifted kWh actually avoids: energy at 12.00 plus the 9.98 of network charges and levies that make up the distribution tariff. It excludes the capacity charge (*Leistungstarif*) and the fixed metering charge (*Messentgelt*), which a shifted kWh does not avoid at the margin unless it also lowers the monthly peak—a separate value channel that our model, whose only value is energy time-shifting, does not represent. We take 8.00 cents/kWh as the export price: the statutory minimum for installations below 30 kW plus the certificate-of-origin top-up. [Appendix D.1](#) documents the tariff components and the basis for the export price, and Section 5.5 reports the sensitivity to $\bar{p}/\underline{p}$.

The corridor is not a property of the technology but of the regulatory treatment of internal exchange: it is whatever volumetric charges an internally consumed kWh avoids, net of the export revenue forgone. Table 3 contrasts the two Swiss frameworks on that basis. Behind a common connection point the internal kWh avoids the whole volumetric stack. Over the public grid it avoids only a reduction in the network tariff—40% on one voltage level, 20% where a transformation step is used—and continues to pay system services, the reserve levy, solidarised costs, the renewable promotion levy and the concession fee in full.

Only the first framework supports a positive battery at the benchmark cost. In the other two the corridor is too narrow for any shared storage to be efficient, and the financing question of Section 3 does not yet arise; it becomes binding only once installed costs fall below roughly 260 CHF/kWh, where α reaches $\Pr(m > 0)$ on the one-voltage-level corridor. The calibration below therefore uses the behind-the-meter corridor, and we return to the distinction in Section 6.

Table 4: Distribution of daily matchable energy $m(\omega)$ (ZEV data, 2024–2025).

Statistic	Value
Mean	61.63 kWh/day
Median	60.37 kWh/day
$\Pr(m > 0)$	0.928
Days	656

Because Swiss LECs entered into force only from 2026, we calibrate the mismatch distribution $m(\omega)$ using quarter-hourly data from an existing Swiss ZEV. The ZEV data discipline only the *physical* shifting potential; LEC-relevant institutional differences enter through Δ_0 and the settlement rule; narrowing the corridor raises α , which Table 5 sweeps, and also changes the price ratio $\bar{p}/p$ that enters SDR revenue separately (Section 5.5), so transferring the financed shares to a storage-viable LEC requires both.

5.2. Data and matchable energy

The model requires the distribution of daily matchable energy $m(\omega)$. We compute $m(\omega)$ from quarter-hourly production and consumption measured in a ZEV¹² over 656 consecutive days (22 January 2024 to 7 November 2025, with no missing days; the first and last months are partial). The ZEV comprises 24 units (21 apartments and 3 hobby rooms) with a combined PV capacity of 99.4 kWp, rounded to 100 kWp below. Table 4 summarizes the resulting distribution.

To ensure physical feasibility, $m(\omega)$ is obtained from a chronological dispatch simulation over the 96 quarter-hourly intervals of each day: an effectively unconstrained battery charges from community-level PV surplus and discharges into the subsequent deficit, subject to a 1C power limit, and starts each day empty. Appendix D.2 describes the simulation and validates the one-cycle form used below: a second surplus–deficit cycle accounts for less than 1.5% of matchable energy, carrying energy across midnight changes shifted energy by 0.3% at the surplus-optimal capacity and by less at smaller ones, and the simulated throughput matches $\min\{K, m(\omega)\}$ to within 1% on more than 99% of days at every capacity up to 150 kWh.

Given $\{m(\omega)\}$, the expected shifted energy for a battery of capacity K is

$$V(K) \equiv \mathbb{E}[\min\{K, m(\omega)\}], \quad (29)$$

computed by sample averaging. The function $V(K)$ is plotted in Figure 2.

5.3. Surplus-optimal battery size

With ZEV membership fixed at N participants, a planner chooses aggregate capacity K to maximize net surplus

$$S(K) \equiv \Delta_0 V(K) - cK, \quad (30)$$

where c is the daily levelized cost per kWh of usable capacity (cents per kWh-capacity per day). Appendix A details the derivation of c from capital cost, lifetime, and discount rate assumptions.

Since the right derivative is $V'(K) = \Pr(m(\omega) > K)$ (and the left derivative $\Pr(m(\omega) \geq K)$; the two coincide except at atoms of the empirical distribution), the first-order condition for an interior optimum is

$$\Pr(m(\omega) > K^W) \leq \alpha \leq \Pr(m(\omega) \geq K^W), \quad \text{where } \alpha \equiv \frac{c}{\Delta_0}. \quad (31)$$

Thus, the surplus-optimal capacity K^W corresponds to the $(1 - \alpha)$ -quantile of the empirical distribution of $m(\omega)$.

¹²Data from a ZEV in the Canton of Luzern, Switzerland, in the supply area of CKW; the load profiles in Romano and Trutnevyte (2026) are drawn from the same utility’s smart-meter base in the same canton (over 100,000 CKW meters in Canton Luzern).

Benchmark scenario. For the main results, we set $\alpha = 0.70$, which implies $c = 9.79$ cents/(kWh·day) under the benchmark corridor $\Delta_0 = 13.98$ cents/kWh. This corresponds to an installed battery cost of approximately 426 CHF/kWh (assuming a 15-year lifetime and 3% real discount rate; see [Appendix A](#)). The average end-customer price in Switzerland had fallen to around 800 CHF per kWh of installed capacity by December 2024, while the installers surveyed regarded 300–600 CHF/kWh as the range at which storage becomes economically sensible; the benchmark lies inside that range, and the same survey attributes the continuing decline in storage prices to technological progress ([Swissolar, 2025](#)). Note also that a positive battery requires $\alpha < \Pr(m > 0) = 0.93$ for this community, i.e. an installed cost below roughly 565 CHF/kWh—a tighter threshold than $\alpha < 1$ alone implies. At $\alpha = 0.70$, the surplus-optimal capacity is $K^W = 37.66$ kWh.

5.4. Energy-only equilibrium under SDR

We model internal pricing using the supply–demand-ratio (SDR) rule of [Liu et al. \(2017\)](#), applied as in [Appendix B.2](#).

Denote by $\Delta^{SDR}(K)$ the expected internal margin per shifted kWh. Revenue per unit of installed capacity is then

$$g(K) \equiv \Delta^{SDR}(K) \frac{V(K)}{K}. \quad (32)$$

The energy-only equilibrium capacity K^D solves the zero-profit condition

$$g(K^D) = c. \quad (33)$$

Figure 3 plots unit revenue $g(K)$ against the cost c . The revenue curve slopes downward because both utilization ($V(K)/K$) and the SDR spread ($\Delta^{SDR}(K)$) decline with capacity. Under SDR applied to the matched internal market, spread compression dominates utilization decline, producing severe under-investment relative to K^W ; the full-deficit paragraph below sets out how the picture changes when the ratio is measured against the full deficit instead.

Two-period vs. sub-daily SDR computation. The SDR margin is computed on the model’s two-period aggregation, applying the rule once per period to each day’s consolidated surplus and deficit. In a 15-minute implementation the margin would also depend on the within-day profile of surplus and deficit and on the dispatch rule; [Appendix D.3](#) explains why the direction of the difference cannot be signed without specifying how interval-level internal demand is allocated.

Matched-market vs. full-deficit SDR base. A second question concerns the demand side of the Period-2 ratio, and it is a question about the *matching rule* rather than a rival way of computing one quantity: the SDR price formula is fixed, but which quantity it is applied to is a design choice. [Liu et al. \(2017\)](#) define the ratio as total internal supply over total internal demand, with demand the energy-sharing zone’s full net deficit, so the literal reading of the published rule uses $\text{SDR}'_2 = q/d(\omega)$. Applied to the internal market of the two-period model—the matchable market of size $m(\omega)$ stipulated in [Appendix B](#)—the rule uses $\text{SDR}_2 = q/m(\omega)$: only matchable energy is traded internally and the residual deficit $d - m$ is imported at retail. The calibrated ZEV has no internal market of either kind: it nets at the common connection point and imports the residual, and its netting does not distinguish the unmet matchable deficit $m - q$ from the non-matchable deficit $d - m$. Both conventions are therefore counterfactual settlement overlays on the same physical data. We report the matched-market base in the main tables because it is the model-consistent one, and the full-deficit base in [Table 6](#) because it is the rule as published. Since $d(\omega) \geq m(\omega)$, the matched base gives a weakly higher ratio, a weakly lower discharge price and hence weakly lower storage revenue: it is the stronger form of compression, because the battery is deemed to relieve scarcity as soon as it absorbs the matchable surplus, and the period-2 ratio reaches one whenever $K \geq m(\omega)$. Under the full-deficit base, compression is weaker because stored energy remains scarce relative to total evening demand. In our data the deficit is large relative to the matchable energy— $d(\omega)$, the community deficit over the intervals that follow the day’s first surplus, averages 119 kWh against 62 kWh for $m(\omega)$ —so the two bases differ materially; [Appendix D.4](#) sets out why the period-2 rather than the whole-day deficit is used and reports the whole-day alternative, which is an




Figure 2: Expected shifted energy $V(K) = \mathbb{E}[\min\{K, m\}]$ as a function of battery capacity. The dotted line indicates the 45-degree benchmark where $V(K) = K$ (full utilization). Vertical lines mark the surplus-optimal capacity K^W (dashed) and the SDR-financed capacity K^D (dash-dot) for the benchmark scenario $\alpha = 0.70$. Data: 656 days from a Swiss ZEV in the CKW supply area.

upper bound on the financed share. Table 6 recomputes the equilibrium under SDR₂: at the benchmark the SDR-financed capacity rises from 3.7 to 7.9 kWh (21% of K^W) and the surplus loss falls from 80% to 62%; across cost scenarios the financed share is 21–51% and unit revenue at K^W is below cost in every cost scenario, most narrowly at $\alpha = 0.9$, where it falls short by 2.7% ($\tilde{g}(K^W) = 12.24$ against $c = 12.58$) and stays below cost in 97.2% of day-level resamples (Appendix D.4). The two bases do not bracket compression in general—the weakest case is no compression at all, a constant spread of Δ_0 , under which the market over-invests—but they bracket it here.

5.5. Sensitivity to storage costs

To test the robustness of the financing gap, we vary c through the cost–corridor ratio $\alpha = c/\Delta_0$ and recompute K^W and K^D . Table 5 reports the results for $\alpha \in \{0.50, 0.60, 0.70, 0.80, 0.90\}$.

In the benchmark scenario ($\alpha = 0.70$), we obtain $K^D = 3.73$ kWh—about a tenth of the surplus-optimal capacity. Net surplus is $S(K^W) = 49.7$ cents/day, CHF 182 per year, against $S(K^D) = 10.0$ cents/day, CHF 37 per year: energy-only settlement under SDR achieves only 20% of surplus-optimal community surplus. With endogenous participation this ratio would enter squared (Section 3.5): a community that could also lose members would retain about $0.20^2 \approx 4\%$ of the achievable welfare.

Table 7 places the settlement rules side by side at the surplus-optimal capacity. Scarcity clearing meets the revenue-adequacy condition (15) while paying storage 12.32 of the 13.98 cents/kWh corridor; both SDR conventions compress the margin far below the curvature ratio of 0.881.

Sensitivity and robustness. Across the cost scenarios in Table 5 the SDR-financed capacity stays at 10–19% of the surplus-optimal level under the matched-market base and 21–51% under the full-deficit base (Table 6), with unit revenue at K^W below cost in every scenario; the achieved surplus is 20–31% and 38–88% respectively. The gap persists when storage is cheap because compression intensifies as the battery expands.




Figure 3: Revenue per unit of installed capacity under energy-only settlement with SDR settlement. The solid curve shows $g(K) = \Delta^{SDR}(K) V(K)/K$; the horizontal dashed line is the daily capacity cost $c = 9.79$ cents/(kWh-day). The SDR equilibrium K^D occurs where $g(K) = c$. At K^W , unit revenue is far below cost ($g(K^W) < c$). This gap quantifies the financing shortfall defined in Section 3: investors would incur a loss of $c - g(K^W)$ per unit of capacity if they attempted to install the efficient size.

How far does the magnitude travel beyond this community? We vary the three things it depends on. First, sampling: resampling the 656 days with replacement, i.i.d., with K^W and K^D recomputed on each resample, puts the benchmark share at 9.9% with a 95% interval of [8.9, 10.9]%. This measures sampling variability under the empirical distribution; it does not account for serial dependence in weather or for interannual variation in PV output, so it supports the stability of the ratio within this sample rather than insensitivity to sample length. Second, the price level: the SDR margin depends on $\underline{p}$ and $\bar{p}$ separately and not only on their difference, and Corollary 4 shows the gap scaling with $\bar{p}/\underline{p}$. Holding Δ_0 fixed and moving that ratio from 2.0 to 4.5 moves the financed share from 12.9% to 6.5%: a community facing a less generous feed-in tariff relative to retail faces a wider gap. The two move together when the export price itself changes. At the Q4 2024 BFE reference price of $\underline{p} = 8.80$ instead of the benchmark 8.00, the corridor narrows to $\Delta_0 = 13.18$ and α rises to 0.742, so K^W falls to 31.20 kWh and K^D to 2.83 kWh—a financed share of 9.1% against 9.9% at the benchmark. Third, the mismatch distribution: the uniform benchmark of Appendix B.5 gives 14.0% at the benchmark cost against 9.9% here, so the shape of F matters for the level while leaving the order unchanged.

Seasonality matters for the level. Splitting the sample gives 15.0% in summer and 15.0% in winter, against 9.9% pooled. The within-season figures describe a battery sized for one season alone, which moves both $g(K)$ and K^W , so the comparison does not isolate a single mechanism; the pooled figure is the one relevant to a battery installed once and operated year-round, and days on which it absorbs the entire matchable energy ($\rho = 1$, zero SDR margin) are genuine operating outcomes rather than an artifact of aggregation. Evidence from more than one community is correspondingly valuable: Greber et al. (2026) reproduce the endogenous-price mechanism on measured demand from twelve Swiss communities.

Engineering assumptions enter through the same two channels. Round-trip losses, degradation and shorter lifetimes raise the effective cost or narrow the effective corridor, i.e. they raise α , and Table 5 shows the ratio staying between a tenth and a fifth over the range of α it reports. Re-running the dispatch and the entry condition with a round-trip efficiency of $\eta = 0.90$ (Appendix D.5) bears this out: the surplus-optimal

Table 5: Surplus-optimal vs. SDR-financed capacity across cost/corridor ratios (matched-market SDR base). $\alpha = c/\Delta_0$ is the cost–corridor ratio; c is in cents/(kWh·day); capacities are in kWh; $S(K^W)$ is the net community surplus at the surplus-optimal capacity, in cents/day.

α	c	K^W	$S(K^W)$	K^D	K^D/K^W	$S(K^D)/S(K^W)$
0.50	6.99	60.37	188.5	11.46	19%	31%
0.60	8.39	49.04	111.4	7.05	14%	24%
0.70	9.79	37.66	49.7	3.73	10%	20%
0.80	11.18	14.62	10.0	1.40	10%	22%
0.90	12.58	1.36	0.2	0.13	10%	24%

Table 6: Robustness: SDR-financed capacity under the full-deficit base $\text{SDR}'_2 = q/d(\omega)$, with $d(\omega)$ the period-2 deficit (Section 5.4, Appendix D.4). $\tilde{g}(K^W)$ is unit revenue at the surplus-optimal capacity under this convention, to be compared with c ; capacities in kWh, revenue and cost in cents/(kWh·day). A whole-day denominator raises every entry in the last four columns—at the benchmark to 7.47, 16.97, 45% and 68%—and is an upper bound.

α	c	K^W	$\tilde{g}(K^W)$	$K^{D'}$	$K^{D'}/K^W$	$S(K^{D'})/S(K^W)$
0.50	6.99	60.37	2.35	19.28	32%	48%
0.60	8.39	49.04	3.12	12.96	26%	41%
0.70	9.79	37.66	4.22	7.86	21%	38%
0.80	11.18	14.62	7.99	3.80	26%	49%
0.90	12.58	1.36	12.24	0.69	51%	88%

capacity falls by three quarters, to 10.22 kWh, while the financed share falls only from 9.9% to 8.1%. Losses do absorb most of the net community surplus at the benchmark cost (CHF 14 per year against CHF 182), so the *level* of the community-ownership gain reported below does not survive realistic losses at benchmark costs, although the ranking of the regimes does.

Community ownership in the calibrated ZEV. Under community ownership (Proposition 3) the community installs $K^W = 37.66$ kWh and finances it through the balanced-budget fee $f^C = cK^W = 369$ cents/day, i.e. CHF 1,345 per year for the community or about CHF 56 per unit per year across the 24 units. The battery’s gross value is $\Delta_0 V(K^W) = 418$ cents/day (CHF 1,527 per year), so the net community surplus of CHF 182 per year is about CHF 8 per unit. The net benefit is positive at every α below $\Pr(m > 0)$ in Table 5 by construction ($S(K^W) > 0$), but it is small at the benchmark because $\alpha = 0.70$ lies close to the viability threshold: the fee absorbs 88% of the gross value. The level grows quickly as batteries get cheaper—at $\alpha = 0.60$ ($I \approx 365$ CHF/kWh) the net community surplus at K^W is CHF 407 per year (CHF 17 per unit), and at $\alpha = 0.50$ ($I \approx 304$ CHF/kWh) it is CHF 688 per year (CHF 29 per unit) against a fee of CHF 64 per unit—whereas the SDR-financed share of the efficient capacity stays between 10% and 19% over $\alpha \in [0.5, 0.9]$ (Table 5). The object of interest is therefore the ratio K^D/K^W and the mechanism behind it, not the level of the net gain at any one cost. With a uniform fee and benefits allocated in proportion to matchable energy, a unit whose share of the gross battery value falls below $cK^W/(24 \Delta_0 V(K^W)) \approx 0.88/24 \approx 3.7\%$ —88% of an equal share, the curvature ratio of Table 7 again—pays more than it gains; the benefit-proportional fee of Section 4.2 removes this, and its implementation requires the per-unit metering that a ZEV’s billing already provides.

5.6. Who captures the value? A functional decomposition

Community ownership closes the financing gap and delivers the surplus-optimal net benefit $S(K^W)$, CHF 182 per year at the benchmark cost. The remaining question is how that value is *distributed*. Our data record community-level totals rather than individual meters, so we cannot trace the distribution across members. We can, however, decompose it along the three functional roles the model distinguishes: the

Table 7: Settlement rules at the surplus-optimal capacity, benchmark scenario ($\alpha = 0.70$). The revenue-adequacy condition (15) holds when the compression ratio $\Delta(K^W)/\Delta_0$ equals the curvature ratio $v'(K^W)/(v(K^W)/K^W) = 0.881$, evaluated with $v'(K^W) = \alpha$, the subgradient selected by the planner’s condition. Scarcity clearing attains it while still paying storage less than the full corridor, so underpayment relative to Δ_0 is not by itself the source of the gap.

Settlement rule	$\Delta(K^W)/\Delta_0$	vs. 0.881	Outcome
Constant spread $\delta = \Delta_0$	1.000	above	over-invests
Constant spread $\delta = \delta^*$	0.881	equal	efficient
Scarcity clearing (Walrasian)	0.881	equal	efficient
SDR, full-deficit base	0.380	below	under-invests
SDR, matched-market base	0.240	below	under-invests
Time-invariant internal price	0.000	below	$K^D = 0$

Table 8: Functional decomposition of the gross value of shifting $G = \Delta_0 V(K)$ under SDR settlement, benchmark scenario ($\alpha = 0.70$). Values in cents/day; shares of G in parentheses. “Battery net” = $\Delta(K)V(K) - cK$; a negative value is the financing gap.

	G	PV	Battery / net	Consumers
At K^W	418.3	0 (0%)	100.4 (24%) / -268.2	317.9 (76%)
At K^D	46.5	0 (0%)	36.5 (78%) / ≈ 0	10.0 (22%)

surplus (PV) side, the battery, and the deficit (consumer) side. Per shifted kWh, the corridor Δ_0 splits into the PV premium $p^1 - \underline{p}$, the battery margin $\Delta(K) = p^2 - p^1$, and the consumer discount $\bar{p} - p^2$ (Section 2.4).

Under SDR settlement this split has a sharp form. Because period-1 supply is abundant, storage charges at the feed-in floor, $p^1 = \underline{p}$ (Appendix B.2), so the PV side earns *no* premium and the entire non-battery value accrues to consumers. Table 8 reports the split at the surplus-optimal and SDR-financed capacities. At K^W , the battery’s *gross* arbitrage revenue $\Delta(K^W)V(K^W)$ recovers only a fraction of its cost cK^W ; the remainder—the negative “battery net” entry—is exactly the financing gap. Consumers capture the bulk of the gross value of shifting while the asset that creates it runs a loss under energy-only settlement.

The decomposition clarifies what the membership fee does distributionally: it transfers the shortfall from the members who capture the value—consumers—to the battery, converting a privately unprofitable asset into a collectively financed one. It also shows that the distribution is a property of the settlement rule, not a primitive. Under the mid-market rate (MMR), $p^1 = p^2$ is the corridor midpoint, so the battery earns nothing and the PV and consumer sides share Δ_0 equally (Appendix B.3); the distribution shifts entirely even though the physical value $\Delta_0 V(K)$ is unchanged.

Tariff-value vs. system-value interpretation. The surplus figures above measure private bill savings within the price corridor $\Delta_0 = 13.98$ cents/kWh. The corridor embeds substantial volumetric charges that are cost-recovery instruments rather than marginal resource costs (Appendix D.1). To gauge the system-value analogue, consider a system corridor of $\Delta_0^{\text{sys}} \approx 4$ cents/kWh—the residual after stripping out the 9.98 cents/kWh of network charges and policy levies. At the benchmark resource cost $c = 9.79$ cents/(kWh-day), the system-value cost–corridor ratio is $\alpha^{\text{sys}} = c/\Delta_0^{\text{sys}} \approx 2.4 > 1$, so the system-value optimum is the corner $K^{W,\text{sys}} = 0$: storage is not yet efficient at current costs when evaluated against marginal resource values alone. In this case the financing gap is moot from a system perspective—the energy-only market’s failure to finance a large battery is not a social loss. The 4-cent figure is a tariff-component proxy—the retail energy component net of the feed-in payment—not an estimate of the marginal resource value of shifting, which depends on the timing of surplus and deficit; what follows is conditional on it. As battery costs decline and α^{sys} falls below one, two cases arise. If internal settlement were based on the same corridor the planner values, the analysis of Section 3 applies with system-value prices in place of $(\bar{p}, \underline{p})$ and the financing gap recurs under system-value accounting exactly as under tariff-value accounting. If instead settlement stays tariff-based while the benchmark is system value, storage revenue is unchanged by the change of benchmark:

the market still installs the capacity the tariff schedule $g(K)$ supports, and the comparison with $K^{W,\text{sys}}$ can go either way. At $c = 3.6$ cents/(kWh·day)—an installed cost of about 157 CHF/kWh, well below the range of Table A.1 and so a future-cost scenario— $\alpha^{\text{sys}} = 0.90$ gives $K^{W,\text{sys}} = 1.36$ kWh while tariff-based SDR entry ($\alpha = 0.26$) sustains 29.45 kWh under the matched-market base—and more under the full-deficit base, where the same cost supports 44 kWh: the levies embedded in the tariff corridor then act as an implicit subsidy to storage relative to its resource value, and the market over-invests by more than a factor of twenty. The *mechanism* of Section 3 is corridor-invariant only when settlement and benchmark use the same corridor. At the benchmark cost the system net value of $K^W = 37.66$ kWh is about -249 cents/day, so the battery the fee finances is efficient in tariff-value terms and not in system-value terms.

6. Conclusion and Policy Implications

Swiss LECs and similar European collective self-consumption frameworks are designed on the assumption that internal energy trading will attract community-level storage investment. Our analysis shows that this assumption does not hold under energy-only settlement. Calibrating to quarter-hourly smart-meter data from a Swiss ZEV, we find that, at the benchmark cost, scarcity-reflective SDR settlement finances about a tenth of the surplus-optimal battery capacity when the supply–demand ratio is measured against the matched internal market, and about a fifth when it is measured against the full period-2 deficit; the community forgoes between roughly three fifths and four fifths of the potential value from intraday shifting. The direction is under-investment in every scenario. This value is measured in tariff-value terms—bill savings within the price corridor. At the December 2024 average end-customer price no shared battery is justified even in tariff-value terms (Appendix A); at the benchmark cost, under the tariff-component proxy for the system-value corridor in Section 5.6, none is justified in system-value terms either, and the financing gap is then a private coordination failure rather than a social efficiency loss; whether the resource value of intraday shifting exceeds the roughly 10.5 cents/kWh that would reverse this at the benchmark cost is a question our data cannot answer. As battery costs decline the same gap recurs under system-value accounting if internal settlement comes to reflect system values; if settlement stays tariff-based, the tariff corridor can instead over-reward storage relative to its resource value (Section 5.6).

Community ownership with a balanced-budget membership fee closes the gap entirely in tariff-value terms, achieving the surplus-optimal investment and undistorted participation; it implements the community’s private optimum, which coincides with the system optimum only when the corridor reflects resource values. Platform intermediation achieves efficient capacity per member but serves half the efficient membership through monopoly pricing, so aggregate capacity is half the efficient level. Among the canonical rules, organizational form is the lever whose efficiency does not hinge on a parameter set once and left alone: sizing the battery requires the community’s mismatch distribution under either lever, but under community ownership a misestimate costs a misestimated battery and nothing more, whereas an administrative spread suffices only if calibrated to the distribution and only until it or the cost moves. A scarcity-clearing price, which no community adopts, is the only rule that suffices for every distribution.

Three policy implications follow.

First, regulatory frameworks for energy communities should explicitly enable non-energy financing instruments for community-level assets. Current frameworks are designed around energy settlement, which is appropriate for allocating the benefits of shared PV but not for financing storage. Allowing communities to levy transparent, earmarked membership fees—under balanced-budget or cooperative governance rules—closes the financing gap without requiring public subsidy. This is a capacity-side instrument and complements the generation-side instruments—surplus buy-back, power purchase agreements, contracts for difference—that Romano and Trutnevyte (2026) recommend for the same communities: theirs address what the community builds, ours whether the storage that its own self-consumption logic justifies can be paid for. A prior condition deserves emphasis, because it orders the two levers. A financing instrument only matters where the corridor is wide enough to make a shared battery efficient in the first place, and that is a regulatory choice rather than a technological one: behind a common connection point, where internal exchange avoids the whole volumetric tariff, the corridor supports a positive battery at the benchmark cost; over the public grid at a reduced network tariff it does not, at present costs (Table 3). Corridor width and financing instrument are

therefore levers in sequence, and a framework that grants only a partial network-charge reduction has not yet reached the question this paper answers.

Second, where platforms intermediate community storage, regulators face a tension between investment incentives and consumer protection. A profit-maximizing platform delivers efficient capacity per member but extracts rents that exclude half the efficient membership. Fee caps and competitive procurement obligations can narrow this gap, but a cap set too low undermines the capacity subsidy. Platform regulation for energy communities deserves attention comparable to that devoted to distribution network tariffs.

Third, settlement-rule design and organizational-form design are distinct policy levers. Settlement rules allocate surplus given a fixed asset base; organizational forms determine whether the asset gets built. The financing gap is a property of energy-only settlement under the canonical rules together with the absence of a capacity instrument, not of storage technology or of any single tariff. Its practical consequence, where households can install individually, is a bias toward standalone behind-the-meter batteries, which face no settlement and no financing gap but forgo the capacity-pooling gain of a shared asset and the netting a community provides without storage ([Appendix C](#)); how large those gains are in a given community is a question for household-level data.

Several limitations qualify these conclusions. The two-period structure abstracts from within-period spatial matching, multi-period dispatch, and seasonal variation, and does not cover communities whose generation mix produces several surplus–deficit cycles per day. The SDR margin is computed on daily aggregates rather than interval-by-interval, and its magnitude depends on whether the period-2 ratio is measured against the matchable energy or the full deficit (Section 5), neither of which is dictated by current ZEV metering, so the precise magnitude under a 15-minute implementation remains to be validated. We evaluate the canonical settlement rules and the scarcity-clearing benchmark, not auction-based or bilateral rules, whose outcome depends on bidding behavior. The competitive-fringe assumption rules out strategic behavior, and the welfare benchmark treats the full price corridor as the value of shifting, whereas a portion reflects cost-recovery charges rather than marginal resource costs.

Future work should embed the financing problem in richer operational models with realistic dispatch, extend the analysis to heterogeneous communities where distributional concerns interact with investment incentives, and assess whether non-energy financing instruments can maintain broad participation across income groups—or whether they risk replicating, at the community scale, the access inequities that energy-community policy is intended to address.

Funding

This work was supported by the Swiss Federal Office of Energy (SFOE) under grant number SI/502525. The funding source had no involvement in the study design; in the collection, analysis, and interpretation of data; in the writing of the report; or in the decision to submit the article for publication.

Declaration of competing interest

The author is employed part-time at Swiss Economics SE AG, a consultancy that provides economic advice to public and private clients in the energy sector. This research was conducted independently of any client mandates. The views expressed in this paper are those of the author and do not necessarily reflect the position of Swiss Economics or its clients. The author declares that he has no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

CRedit authorship contribution statement

Nicolas J. Greber: Conceptualization, Methodology, Software, Formal analysis, Data curation, Investigation, Writing – original draft, Writing – review & editing.

Data availability

The replication package is provided as supplementary material. It contains the derived daily series for the 656 days of the sample—for each day the date, the daily surplus $s(\omega)$, the matchable energy $m(\omega)$ without losses and at $\eta = 0.90$, and the community deficit $d(\omega)$ on both the period-2 and the whole-day convention—and two scripts. The first reproduces from this series alone Tables 4–8, Table A.1, Figures 2–3 and every other calibrated number in Section 5 and Appendix D that does not require the quarter-hourly series, together with the uniform-benchmark figures of Section 3.3 and Appendix B.5. The second computes Table 3, the volumetric components of Appendix D.1 and the installed-cost thresholds quoted with them from the published tariff sheet and $\Pr(m > 0)$. Three checks of Appendix D.2—the second-cycle share, the capacity-by-capacity throughput check and the carry-over across midnight—require the underlying quarter-hourly series, which is not published: it records the aggregate production and consumption of a single identifiable building of 24 units (21 apartments and 3 hobby rooms) and is held under the terms on which it was obtained from the community. It is available from the author on reasonable request, subject to those terms.

Declaration of generative AI and AI-assisted technologies in the writing process

During the preparation of this work the author used Claude (Anthropic) to draft and edit the text of the manuscript on the basis of the author’s own model, derivations, calibration and results and to draft the revisions and the response to the referee reports. After using this tool, the author reviewed and edited the content as needed and takes full responsibility for the content of the published article. The research question, the analytical framework, the proofs, the data work, the calibration and the interpretation of the results are the author’s own.

References

- Abada, I., Ehrenmann, A., Lambin, X., 2020a. On the viability of energy communities. *The Energy Journal* 41, 113–150. doi:[10.5547/01956574.41.1.iaba](https://doi.org/10.5547/01956574.41.1.iaba).
- Abada, I., Ehrenmann, A., Lambin, X., 2020b. Unintended consequences: The snowball effect of energy communities. *Energy Policy* 143, 111597. doi:[10.1016/j.enpol.2020.111597](https://doi.org/10.1016/j.enpol.2020.111597).
- Ableitner, L., Tiefenbeck, V., Meeuw, A., Wörner, A., Fleisch, E., Wortmann, F., 2020. User behavior in a real-world peer-to-peer electricity market. *Applied Energy* 270, 115061. URL: <https://doi.org/10.1016/j.apenergy.2020.115061>, doi:[10.1016/j.apenergy.2020.115061](https://doi.org/10.1016/j.apenergy.2020.115061).
- Andreolli, F., D’Alpaos, C., Kort, P., 2025. Does p2p trading favor investments in pv–battery systems? *Energy Economics* 145, 108418.
- Andrés-Cerezo, D., Fabra, N., 2023. Storing power: Market structure matters. *The RAND Journal of Economics* 54, 3–53.
- Antweiler, W., 2021. Microeconomic models of electricity storage: Price forecasting, arbitrage limits, curtailment insurance, and transmission line utilization. *Energy Economics* 101, 105390. doi:[10.1016/j.eneco.2021.105390](https://doi.org/10.1016/j.eneco.2021.105390).
- Armstrong, M., 2006. Competition in two-sided markets. *The RAND Journal of Economics* 37, 668–691.
- Balegamire, R., Gautier, A., 2025. Pricing and sharing rules for energy communities. *Utilities Policy* 96, 101995.
- Borenstein, S., 2012. The private and public economics of renewable electricity generation. *Journal of Economic Perspectives* 26, 67–92. doi:[10.1257/jep.26.1.67](https://doi.org/10.1257/jep.26.1.67).

- Bravo-Melgarejo, S., Haritchabalet, C., 2025. Prosumers: Grid vs. individual storage. *Energy Economics* 141, 108095. doi:[10.1016/j.eneco.2024.108095](https://doi.org/10.1016/j.eneco.2024.108095).
- Cambini, C., Soroush, G., 2019. Designing grid tariffs in the presence of distributed generation. *Utilities Policy* 61, 100979.
- CKW, 2025a. CKW führt dynamischen wahltarif ein und passt stromtarif an. URL: <https://www.ckw.ch/ueber-ckw/medienstelle/medienmitteilungen/2025/ckw-fuehrt-dynamischen-wahltarif-ein-und-passt-stromtarif-an>.
- CKW, 2025b. Preisinformation stromprodukte 2026: Privat- und gewerbekunden (unter 50'000 kwh/jahr). Tariff sheet, Basiskundengruppe.
- CKW, 2025c. Rückliefervergütung: überschüssigen strom verkaufen. URL: <https://www.ckw.ch/energie/strom/ruecklieferverguetung>.
- Cramton, P., Ockenfels, A., Stoft, S., 2013. Capacity market fundamentals. *Economics of Energy & Environmental Policy* 2, 27–46.
- Der Bundesrat, 2025. Bundesrat setzt zweites paket des bundesgesetzes für eine sichere stromversorgung in kraft. Medienmitteilung. URL: <https://www.news.admin.ch/de/nsb?id=104172>.
- Dinh, N.T., Pourmousavi, S.A., Karimi-Arpanahi, S., Kumar, Y.P.S., Guo, M., Abbott, D., Liisberg, J.A.R., 2022. Optimal sizing and scheduling of community battery storage within a local market, in: *Proceedings of the Thirteenth ACM International Conference on Future Energy Systems*, ACM. pp. 34–46. doi:[10.1145/3538637.3538837](https://doi.org/10.1145/3538637.3538837).
- Energie Steiermark, 2024. Energiegemeinschaft mit smartcommunity (product page). URL: <https://www.e-steiermark.com/privat/produkte/energiegemeinschaft>.
- EnergieSchweiz, 2021. Leitfaden Eigenverbrauch. Technical Report. EnergieSchweiz/Bundesamt für Energie (BFE). URL: https://www.ewz.ch/dam/ewz/Geschaeftskunden/Solarenergie/Solarenergie_fuer_Geschaeftskunden/leitfaden_eigenverbrauch.pdf.
- European Parliament, Council of the European Union, 2018. Directive (EU) 2018/2001 of 11 December 2018 on the promotion of the use of energy from renewable sources (recast). *EUR-Lex*. URL: <https://eur-lex.europa.eu/eli/dir/2018/2001/oj>.
- European Parliament and Council of the European Union, 2019. Directive (EU) 2019/944 of 5 June 2019 on common rules for the internal market for electricity and amending directive 2012/27/EU (recast). *Official Journal of the European Union*, L 158, 14.6.2019, pp. 125–199. URL: <https://eur-lex.europa.eu/eli/dir/2019/944/oj>.
- European Parliament and Council of the European Union, 2024. Directive (EU) 2024/1711 of 13 June 2024 amending directives (EU) 2018/2001 and (EU) 2019/944 as regards improving the union’s electricity market design. *Official Journal of the European Union*, L, 2024/1711, 26.6.2024. URL: <https://eur-lex.europa.eu/eli/dir/2024/1711/oj>.
- Gautier, A., Jacqmin, J., Poudou, J.C., 2018. The prosumers and the grid. *Journal of Regulatory Economics* 53, 100–126. doi:[10.1007/s11149-018-9350-5](https://doi.org/10.1007/s11149-018-9350-5).
- Gautier, A., Jacqmin, J., Poudou, J.C., 2025. The energy community and the grid. *Resource and Energy Economics* 82, 101480.
- Greber, N., Eschenbaum, N., Szehr, O., 2026. Peer-to-peer electricity platforms with endogenous prices: A multi-agent neural network control approach. *Energy and AI* 25, 100809. doi:[10.1016/j.egyai.2026.100809](https://doi.org/10.1016/j.egyai.2026.100809).

- Gui, E.M., MacGill, I., 2018. Typology of future clean energy communities: An exploratory structure, opportunities, and challenges. *Energy Research & Social Science* 35, 94–107. doi:[10.1016/j.erss.2017.10.019](https://doi.org/10.1016/j.erss.2017.10.019).
- Islam, S.N., 2024. A review of peer-to-peer energy trading markets: Enabling models and technologies. *Energies* 17, 1702.
- Joskow, P.L., 2008. Capacity payments in imperfect electricity markets: Need and design. *Utilities Policy* 16, 159–170.
- Karaduman, Ö., 2023. Economics of Grid-Scale Energy Storage in Wholesale Electricity Markets. Working Paper. Stanford University.
- Korpås, M., Botterud, A., 2020. Optimality conditions and cost recovery in electricity markets with variable renewable energy and energy storage. CEEPR Working Paper 2020-005. MIT Center for Energy and Environmental Policy Research.
- Lamp, S., Samano, M., 2022. Large-scale battery storage, short-term market outcomes, and arbitrage. *Energy Economics* 107, 105786. doi:[10.1016/j.eneco.2021.105786](https://doi.org/10.1016/j.eneco.2021.105786).
- Liu, N., Yu, X., Wang, C., Li, C., Ma, L., Lei, J., 2017. Energy-sharing model with price-based demand response for microgrids of peer-to-peer prosumers. *IEEE Transactions on Power Systems* 32, 3569–3583.
- Long, C., Wu, J., Zhang, C., Thomas, L., Cheng, M., Jenkins, N., 2017. Peer-to-peer energy trading in a community microgrid, in: 2017 IEEE Power & Energy Society General Meeting, IEEE. pp. 1–5. doi:[10.1109/PESGM.2017.8274546](https://doi.org/10.1109/PESGM.2017.8274546).
- Meeuw, A., Schopfer, S., Wörner, A., Tiefenbeck, V., Ableitner, L., Fleisch, E., Wortmann, F., 2020. Implementing a blockchain-based local energy market: Insights on communication and scalability. *Computer Communications* 160, 158–171.
- Mengelkamp, E., Gärttner, J., Rock, K., Kessler, S., Orsini, L., Weinhardt, C., 2018. Designing microgrid energy markets: A case study: The Brooklyn Microgrid. *Applied Energy* 210, 870–880.
- Mercier, T., Olivier, M., De Jaeger, E., 2023. The value of electricity storage arbitrage on day-ahead markets across Europe. *Energy Economics* 123, 106721.
- Morstyn, T., Farrell, N., Darby, S.J., McCulloch, M.D., 2018. Using peer-to-peer energy-trading platforms to incentivize prosumers to form federated power plants. *Nature Energy* 3, 94–101.
- Newbery, D., 2016. Missing money and missing markets: Reliability, capacity auctions and interconnectors. *Energy Policy* 94, 401–410.
- Reis, I.F.G., Gonçalves, I., Lopes, M.A.R., Henggeler Antunes, C., 2021. Business models for energy communities: A review of key issues and trends. *Renewable and Sustainable Energy Reviews* 144, 111013. doi:[10.1016/j.rser.2021.111013](https://doi.org/10.1016/j.rser.2021.111013).
- Rochet, J.C., Tirole, J., 2003. Platform competition in two-sided markets. *Journal of the European Economic Association* 1, 990–1029.
- Rochet, J.C., Tirole, J., 2006. Two-sided markets: a progress report. *The RAND Journal of Economics* 37, 645–667.
- Rodrigues, D.L., Ye, X., Xia, X., Zhu, B., 2020. Battery energy storage sizing optimisation for different ownership structures in a peer-to-peer energy sharing community. *Applied Energy* 262, 114498.

- Romano, E., Trutnevyte, E., 2026. Local energy communities in rural Switzerland: National scalability under different incentives schemes and economic scenarios. *Renewable Energy* 256, 124569. doi:[10.1016/j.renene.2025.124569](https://doi.org/10.1016/j.renene.2025.124569).
- Schmalensee, R., 2022. Competitive energy storage and the duck curve. *The Energy Journal* 43, 1–16. doi:[10.5547/01956574.43.2.rsch](https://doi.org/10.5547/01956574.43.2.rsch).
- Sioshansi, R., 2010. Welfare impacts of electricity storage and the implications of ownership structure. *The Energy Journal* 31, 173–198.
- Sioshansi, R., 2014. When energy storage reduces social welfare. *Energy Economics* 41, 106–116.
- Sioshansi, R., Denholm, P., Jenkin, T., Weiss, J., 2009. Estimating the value of electricity storage in PJM: Arbitrage and some welfare effects. *Energy Economics* 31, 269–277. doi:[10.1016/j.eneco.2008.10.005](https://doi.org/10.1016/j.eneco.2008.10.005).
- Sorin, E., Bobo, L., Pinson, P., 2018. Consensus-based approach to peer-to-peer electricity markets with product differentiation. *IEEE Transactions on Power Systems* 34, 994–1004.
- Soto, E.A., Bosman, L.B., Wollega, E., Leon-Salas, W.D., 2021. Peer-to-peer energy trading: A review of the literature. *Applied Energy* 283, 116268.
- Swissolar, 2025. Batteriespeicher mit Photovoltaik 2025. Technical Report. Swissolar. URL: https://www.swissolar.ch/02_markt-politik/batteriesmonitor/130525_batteriespeicher_bericht_sws.pdf.
- Tarel, G., Korpås, M., Botterud, A., 2024. Long-term equilibrium in electricity markets with renewables and energy storage only. *Energy Systems* doi:[10.1007/s12667-024-00654-y](https://doi.org/10.1007/s12667-024-00654-y).
- Verband Schweizerischer Elektrizitätsunternehmen (VSE), 2025. Handbuch Eigenverbrauchsregelung (HER). Technical Report. VSE. URL: <https://www.strom.ch/de/media/15256/download>.
- Weinhardt, C., Mengelkamp, E., Cramer, W., Hambridge, S., Hobert, A., Kremers, E., Otter, W., Pinson, P., Tiefenbeck, V., Zade, M., 2019. How far along are local energy markets in the DACH+ region? A comparative market engineering approach, in: *Proceedings of the Tenth ACM International Conference on Future Energy Systems*, pp. 544–549.
- Williams, O., Green, R., 2022. Electricity storage and market power. *Energy Policy* 164, 112872.
- Zafirakis, D., Chalvatzis, K.J., Baiocchi, G., Daskalakis, G., 2016. The value of arbitrage for energy storage: Evidence from European electricity markets. *Applied Energy* 184, 971–986.
- Zhou, Y., Wu, J., Long, C., 2018. Evaluation of peer-to-peer energy sharing mechanisms based on a multiagent simulation framework. *Applied Energy* 222, 993–1022. doi:[10.1016/j.apenergy.2018.02.089](https://doi.org/10.1016/j.apenergy.2018.02.089).

Appendix A. Levelized Cost of Storage Capacity

We interpret c as the daily levelized resource cost of one unit of community-level battery capacity (1 kWh of usable energy). Let I denote the installed cost (CHF/kWh), L the economic lifetime (years), and ϱ the real discount rate. Here I is measured per kWh of *usable* capacity and excludes VAT, consistent with the tariff components of Section 5; end-customer figures quoted for comparison may be stated per nominal kWh and including VAT, which would move the implied α for a given quoted price but not any result expressed in α . Abstracting from annual O&M (typically negligible for residential lithium-ion systems under warranty), the capital recovery factor is

$$CRF(\varrho, L) = \frac{\varrho}{1 - (1 + \varrho)^{-L}}. \quad (\text{A.1})$$

The daily levelized cost in cents/(kWh·day) is then

$$c = \frac{CRF(\varrho, L) \times I \times 100}{365}. \quad (\text{A.2})$$

Mapping α to installed cost. Table A.1 shows the implied installed cost I for selected values of α , assuming $\Delta_0 = 13.98$ cents/kWh, $L = 15$ years, and $\varrho = 3\%$ (hence $CRF \approx 0.0838$).

Table A.1: Cost–corridor ratio α and implied battery CAPEX.

α	c (cents/kWh·day)	Implied I (CHF/kWh)
0.50	6.99	304
0.60	8.39	365
0.70	9.79	426
0.80	11.18	487
0.90	12.58	548

The average end-customer price in Switzerland was around 800 CHF/kWh in December 2024, which corresponds to $\alpha > 1$ given current tariffs; a positive battery for this community requires an installed cost below roughly 565 CHF/kWh, where α reaches $\Pr(m > 0)$. The scenarios in Table 5 therefore represent projected cost trajectories as prices decline or support mechanisms are introduced.

Appendix B. Microfoundation of internal margins under settlement rules

Appendix B.1. Two-period corridor: objects used in the model

A day ω has a surplus period (1) and a deficit period (2). Working *per participating prosumer*, let $s(\omega) \geq 0$ and $d(\omega) \geq 0$ denote the period-1 surplus and period-2 deficit of eq. (2) in the main text (lower-case s, d avoid a clash with the surplus function $S(\cdot)$). The *matchable market* is the short side:

$$m(\omega) \equiv \min\{s(\omega), d(\omega)\} \in [0, \bar{m}]. \quad (\text{B.1})$$

Let $k \geq 0$ denote capacity per participating prosumer (so in the main text $k = K/N$ when $N > 0$). Given k , the battery can time-shift

$$q(\omega, k) \equiv \min\{k, m(\omega)\}. \quad (\text{B.2})$$

For $m(\omega) > 0$, define the utilization ratio

$$\rho(\omega, k) \equiv \frac{q(\omega, k)}{m(\omega)} \in [0, 1], \quad (\text{B.3})$$

and set $\rho(\omega, k) = 0$ when $m(\omega) = 0$ (then $q(\omega, k) = 0$).

External prices are the feed-in tariff $\underline{p}$ and the retail price $\bar{p} > \underline{p}$, with corridor $\Delta_0 \equiv \bar{p} - \underline{p} > 0$. Each settlement rule induces state-dependent internal prices in each period, which we map into the main-text objects:

$$p^1(\omega, k) \quad (\text{price paid by storage when charging in period 1}), \quad (\text{B.4})$$

$$p^2(\omega, k) \quad (\text{price received by storage when discharging in period 2}). \quad (\text{B.5})$$

The within-day internal margin is

$$\Delta(\omega, k) \equiv p^2(\omega, k) - p^1(\omega, k) \in [0, \Delta_0]. \quad (\text{B.6})$$

Finally, consistent with the main text, define the *expected internal margin per shifted kWh* as the volume-weighted average

$$\Delta(k) \equiv \frac{\mathbb{E}_\omega[\Delta(\omega, k) q(\omega, k)]}{\mathbb{E}_\omega[q(\omega, k)]}, \quad k > 0, \quad (\text{B.7})$$

so that expected storage trading revenue equals $\Delta(k) Q$ with $Q = N \mathbb{E}[q(\omega, k)]$.

The mechanisms below are applied to the *matchable* market of size $m(\omega)$. Any residual (non-matchable) surplus $s(\omega) - m(\omega)$ or deficit $d(\omega) - m(\omega)$ is settled externally at $\underline{p}$ or $\bar{p}$ and does not affect $p^1(\omega, k)$ and $p^2(\omega, k)$.

Appendix B.2. SDR (Supply–Demand Ratio) settlement

We follow the SDR (supply–demand ratio) rule in [Liu et al. \(2017\)](#). In any period, define

$$\text{SDR} \equiv \frac{\text{local (internal) supply}}{\text{local (internal) demand}}. \quad (\text{B.8})$$

SDR computes period-specific sell and buy prices as functions of SDR:

$$p^{\text{sell}}(\text{SDR}) = \begin{cases} \underline{p}, & \text{SDR} > 1, \\ \frac{\underline{p} \bar{p}}{(\bar{p} - \underline{p}) \text{SDR} + \underline{p}}, & \text{SDR} \leq 1, \end{cases} \quad (\text{B.9})$$

$$p^{\text{buy}}(\text{SDR}) = \begin{cases} \underline{p}, & \text{SDR} > 1, \\ p^{\text{sell}}(\text{SDR}) \cdot \text{SDR} + \bar{p} \cdot (1 - \text{SDR}), & \text{SDR} \leq 1. \end{cases} \quad (\text{B.10})$$

Applying SDR to the corridor. In period 1 (surplus), internal matchable supply is $m(\omega)$ and storage charging demand is $q(\omega, k) \leq m(\omega)$. Thus

$$\text{SDR}_1(\omega, k) = \begin{cases} \frac{m(\omega)}{q(\omega, k)}, & q(\omega, k) > 0, \\ +\infty, & q(\omega, k) = 0, \end{cases} \Rightarrow \text{SDR}_1(\omega, k) \geq 1. \quad (\text{B.11})$$

From (B.9)–(B.10), storage charges at

$$p^1(\omega, k) = p^{\text{buy}}(\text{SDR}_1(\omega, k)) = \underline{p}. \quad (\text{B.12})$$

In period 2 (deficit), internal matchable supply is the storage discharge $q(\omega, k)$ and internal matchable demand is $m(\omega)$. Hence

$$\text{SDR}_2(\omega, k) = \frac{q(\omega, k)}{m(\omega)} = \rho(\omega, k) \in [0, 1] \quad (m(\omega) > 0). \quad (\text{B.13})$$

By (B.9), storage receives

$$p^2(\omega, k) = p^{\text{sell}}(\text{SDR}_2(\omega, k)) = \frac{\underline{p} \bar{p}}{(\bar{p} - \underline{p}) \rho(\omega, k) + \underline{p}} = \frac{\underline{p} \bar{p}}{\underline{p} + \Delta_0 \rho(\omega, k)}. \quad (\text{B.14})$$

Induced margin. For $m(\omega) > 0$,

$$\Delta^{SDR}(\omega, k) \equiv p^2(\omega, k) - p^1(\omega, k) = \frac{\underline{p}\bar{p}}{\underline{p} + \Delta_0 \rho(\omega, k)} - \underline{p} = \frac{\underline{p}\Delta_0(1 - \rho(\omega, k))}{\underline{p} + \Delta_0 \rho(\omega, k)}. \quad (\text{B.15})$$

Moreover, if $k \geq m(\omega)$ then $\rho(\omega, k) = 1$ and $\Delta^{SDR}(\omega, k) = 0$ (full arbitrage). If $0 < k < m(\omega)$ then $\rho(\omega, k) = k/m(\omega)$ and (B.15) is strictly decreasing in k .

Expected margin per shifted kWh. For $k > 0$,

$$\Delta^{SDR}(k) \equiv \frac{\mathbb{E}_\omega[\Delta^{SDR}(\omega, k) q(\omega, k)]}{\mathbb{E}_\omega[q(\omega, k)]}. \quad (\text{B.16})$$

Appendix B.3. Fixed internal tariffs, including the mid-market rate

Suppose internal prices are fixed constants, independent of (ω, k) :

$$p^1(\omega, k) \equiv p^{\text{buy}}, \quad p^2(\omega, k) \equiv p^{\text{sell}}, \quad \text{with } \underline{p} \leq p^{\text{buy}} \leq p^{\text{sell}} \leq \bar{p}. \quad (\text{B.17})$$

Then the margin is constant,

$$\Delta(\omega, k) = p^{\text{sell}} - p^{\text{buy}} \equiv \delta \in [0, \Delta_0], \quad (\text{B.18})$$

and the volume-weighted expected margin equals

$$\Delta(k) = \frac{\mathbb{E}_\omega[\delta q(\omega, k)]}{\mathbb{E}_\omega[q(\omega, k)]} = \delta, \quad k > 0. \quad (\text{B.19})$$

Hence there is no spread compression: $\Delta(k)$ does not vary with capacity per member. Fixed tariffs can be viewed as administratively set internal prices. Any net imbalance between physical flows and fixed-price internal settlements must be absorbed elsewhere (e.g. by the default supplier or an operator), so this rule does not clear internal markets endogenously.

Mid-market rate. The mid-market rate of Long et al. (2017) is the special case $p^{\text{buy}} = p^{\text{sell}} = (\underline{p} + \bar{p})/2$, so $\delta = 0$: storage charges and discharges at the corridor midpoint, earns no energy-only margin, and the financing gap equals the full cost. The corridor is delivered through the *level* of the transfer price, split equally between the surplus and the deficit side, rather than through an intraday spread.

Appendix B.4. Scarcity-clearing (Walrasian) settlement

Suppose the matchable market in each period clears at a Walrasian price, with prosumers' outside options at $\underline{p}$ (export) and $\bar{p}$ (import). In period 1, matchable supply is $m(\omega)$ at any price weakly above $\underline{p}$ and storage demand is $q(\omega, k) = \min\{k, m(\omega)\}$. If $k < m(\omega)$ there is excess supply at every price above $\underline{p}$, so $p^1(\omega, k) = \underline{p}$. In period 2, matchable demand is $m(\omega)$ at any price weakly below $\bar{p}$ and supply is $q(\omega, k)$. If $k < m(\omega)$ there is excess demand at every price below $\bar{p}$, so $p^2(\omega, k) = \bar{p}$. If $k \geq m(\omega)$, storage capacity is in excess supply in both periods and competition among storage units drives the margin to zero. Hence

$$\Delta^{\text{Wal}}(\omega, k) = \Delta_0 \mathbf{1}\{m(\omega) > k\}, \quad \Delta^{\text{Wal}}(k) = \frac{\Delta_0 k \Pr(m > k)}{v(k)}, \quad g^{\text{Wal}}(k) = \Delta_0 \Pr(m > k) = \Delta_0 v'(k). \quad (\text{B.20})$$

Free entry, $g^{\text{Wal}}(k) = c$, coincides with the planner's condition (9), so $k^D = k^W$ for every mismatch distribution F and every corridor.

Equality states. On days with $m(\omega) = k$ period-2 supply $q = \min\{k, m\} = k$ equals demand m , so every price in $[\underline{p}, \bar{p}]$ clears and the indicator in (B.20) picks one extreme. Writing the general selection $\Delta^{Wal}(\omega, k) = \Delta_0 \mathbf{1}\{m(\omega) > k\} + \vartheta(k) \Delta_0 \mathbf{1}\{m(\omega) = k\}$ with $\vartheta(k) \in [0, 1]$ gives $g^{Wal}(k) = \Delta_0 [\Pr(m > k) + \vartheta(k) \Pr(m = k)]$, which ranges over $[\Delta_0 \Pr(m > k), \Delta_0 \Pr(m \geq k)]$ —exactly the interval of subgradients the planner’s condition admits at k . When F is atomless the selection is immaterial and $k^D = k^W$ as stated. When F is discrete, g^{Wal} is a step function that can cross c without attaining it, and free entry holds at the planner’s optimum only under the selection ϑ solving $\Pr(m > k) + \vartheta \Pr(m = k) = \alpha$. The empirical distribution of Section 5 is of this kind: the survival function takes values in multiples of $1/656$ and $\alpha = 0.70$ is not among them, so the optimum is the observation at which the subgradient interval contains α . The reported K^W is the linear-interpolation quantile, which differs from that observation by 0.04 kWh and leaves every reported figure unchanged at the precision shown. The distributional counterpart is that on days with $m(\omega) > k$ prosumers obtain none of the gross value of shifting; the entire margin accrues to storage. The rules in Appendix B.2–Appendix B.3 depart from (B.20) in order to share the corridor with prosumers on those days. Sharing as such does not generate the gap—a constant spread with δ/Δ_0 above the curvature ratio shares value with prosumers on scarce days yet over-invests, because it also pays storage on days when scarcity is already relieved. By Corollary 2 the gap opens when expected unit revenue at k^W falls below the planner’s marginal value, which is what SDR’s state-by-state compression produces. Relative to (B.20), SDR settlement lowers p^2 below $\bar{p}$ as soon as $\rho(\omega, k) > 0$ (eq. (B.15)), rather than only when $\rho = 1$.

Appendix B.5. Closed-form SDR expressions and monotonicity of the revenue function

Assume $m(\omega) \sim U[0, \bar{m}]$ i.i.d. across days. Then for $0 \leq k \leq \bar{m}$,

$$v(k) = \mathbb{E}[\min\{k, m\}] = k - \frac{k^2}{2\bar{m}} \quad \Rightarrow \quad \frac{v(k)}{k} = 1 - \frac{k}{2\bar{m}}, \quad (\text{B.21})$$

so that the planner’s condition (9) gives $k^W = \bar{m}(1 - \alpha)$ and $S(k^W) = \Delta_0 \bar{m}(1 - \alpha)^2/2$ for $\alpha \in (0, 1)$. Moreover, for $k \geq \bar{m}$ we have $q(\omega, k) = m(\omega)$ a.s., so SDR fully arbitrages the matchable market in every state and $\Delta^{SDR}(k) = 0$.

A closed form for $\Delta^{SDR}(k)$ on $(0, \bar{m})$. For $k \in (0, \bar{m})$,

$$\Delta^{SDR}(k) = \frac{\Delta_0}{1 - \frac{k}{2\bar{m}}} \left[\frac{\bar{m} - k}{\bar{m}} + \frac{k\bar{p}}{\bar{m}\underline{p}} \ln\left(\frac{k\bar{p}}{\underline{p}\bar{m} + \Delta_0 k}\right) \right]. \quad (\text{B.22})$$

Monotonicity of g on the relevant range. Define, as in the main text,

$$g(k) \equiv \Delta^{SDR}(k) \frac{v(k)}{k}. \quad (\text{B.23})$$

Combining (B.21) and (B.22) yields (for $k \in (0, \bar{m})$)

$$g(k) = \Delta_0 h(k), \quad h(k) \equiv \frac{\bar{m} - k}{\bar{m}} + \frac{k\bar{p}}{\bar{m}\underline{p}} \ln\left(\frac{k\bar{p}}{\underline{p}\bar{m} + \Delta_0 k}\right). \quad (\text{B.24})$$

Corollary 3. *Under SDR with $m \sim U[0, \bar{m}]$, the function $g(k)$ is strictly decreasing on $(0, \bar{m})$.*

Proof. Since $\Delta_0 > 0$, it suffices to show $h'(k) < 0$ on $(0, \bar{m})$. Let

$$L(k) \equiv \ln\left(\frac{k\bar{p}}{\underline{p}\bar{m} + \Delta_0 k}\right), \quad k \in (0, \bar{m}). \quad (\text{B.25})$$

Then $L(k) < 0$ for $k \in (0, \bar{m})$ and $L(\bar{m}) = 0$. Differentiating (B.24) gives

$$h'(k) = -\frac{1}{\bar{m}} + \frac{\bar{p}}{\bar{m}\underline{p}} L(k) + \frac{\bar{p}}{\underline{p}\bar{m} + \Delta_0 k}. \quad (\text{B.26})$$

Set $x \equiv k/\bar{m} \in (0, 1)$ and define

$$y(x) \equiv \frac{x\bar{p}}{\underline{p} + \Delta_0 x} \in (0, 1), \quad \text{so that} \quad L(k) = \ln(y(x)). \quad (\text{B.27})$$

Using the strict inequality $\ln y < y - 1$ for $y \in (0, 1)$, we obtain

$$\bar{m} h'(k) = -1 + \frac{\bar{p}}{\underline{p}} \ln(y(x)) + \frac{\bar{p}}{\underline{p} + \Delta_0 x} < -1 + \frac{\bar{p}}{\underline{p}} (y(x) - 1) + \frac{\bar{p}}{\underline{p} + \Delta_0 x}. \quad (\text{B.28})$$

A direct simplification of the right-hand side yields

$$-1 + \frac{\bar{p}}{\underline{p}} (y(x) - 1) + \frac{\bar{p}}{\underline{p} + \Delta_0 x} = \frac{\bar{p}^2(x - 1)}{\underline{p}(\underline{p} + \Delta_0 x)} < 0 \quad \text{for } x \in (0, 1). \quad (\text{B.29})$$

Hence $h'(k) < 0$ on $(0, \bar{m})$, so $g(k)$ is strictly decreasing on $(0, \bar{m})$. $\square$

For $k \geq \bar{m}$, SDR yields $\Delta^{SDR}(k) = 0$ and thus $g(k) = 0$. The strict monotonicity in Corollary 3 is the relevant property on the interior range where the matchable market is not fully arbitrated in every state.

SDR under-investment holds for every cost level. The closed form (B.24) yields a sharper result than monotonicity: under SDR settlement with a uniform mismatch distribution, energy-only revenue falls short of cost at the efficient scale for *every* cost–corridor ratio and *every* price corridor, and the financing gap has a closed form.

Corollary 4 (Generic under-investment under SDR, matched-market base). *Let SDR be applied to the matchable market as in Appendix B.2, let $m \sim U[0, \bar{m}]$ and $\alpha \equiv c/\Delta_0 \in (0, 1)$, so that $k^W = \bar{m}(1 - \alpha)$. Then*

$$c - g(k^W) = \Delta_0 \frac{\bar{p}}{\underline{p}} (1 - \alpha) \ln \frac{1}{y(1 - \alpha)} > 0, \quad \text{where } y(x) \equiv \frac{x\bar{p}}{\underline{p} + \Delta_0 x} \in (0, 1), \quad (\text{B.30})$$

so $g(k^W) < c$ and hence $k^D < k^W$ for all $\alpha \in (0, 1)$ and all corridors $0 < \underline{p} < \bar{p}$.

Proof. Evaluating h in (B.24) at $k^W = \bar{m}(1 - \alpha)$, i.e. at $x = 1 - \alpha$, gives

$$h(k^W) = \alpha + \frac{\bar{p}}{\underline{p}} (1 - \alpha) \ln y(1 - \alpha).$$

Since $y(x) \in (0, 1)$ for $x \in (0, 1)$, the log term is strictly negative, so $g(k^W) = \Delta_0 h(k^W) < \Delta_0 \alpha = c$, and the gap $c - g(k^W)$ equals the display. Because g is strictly decreasing on $(0, \bar{m})$ (Corollary 3) and $g(k^D) = c$ whenever $k^D > 0$, $g(k^W) < c$ implies $k^D < k^W$. $\square$

Three features of (B.30) are worth noting. First, the per-unit gap vanishes only at the boundaries $\alpha \rightarrow 0$ (where the efficient battery fully arbitrages the matchable market and the margin, cost, and gap all collapse together) and $\alpha \rightarrow 1$ (where the efficient battery shrinks to zero); it is strictly positive for all interior α . (The ratio k^D/k^W is not constant in α : under the uniform benchmark it falls from about 20% at $\alpha = 0.5$ to about 9% at $\alpha = 0.9$, close to the calibrated range in Table 5.) Second, the gap scales with the corridor ratio $\bar{p}/\underline{p}$: the lower the feed-in floor relative to the retail price, the faster the SDR discharge price collapses as local scarcity is relieved, and the larger the shortfall at the efficient scale. Third, the result is specific to the matched-market base, under which the period-2 ratio is q/m . If the ratio is computed against the full deficit, q/d with $d = \lambda m$, the margin becomes $\underline{p} \Delta_0 (1 - \rho/\lambda) / (\underline{p} + \Delta_0 \rho/\lambda)$, which converges to the constant Δ_0 as $\lambda \rightarrow \infty$; by Corollary 2 the market then over-invests. The sign of the distortion under the full-deficit base depends on λ and is settled empirically in Section 5.

Appendix B.6. Platform competition

Section 4.4 notes that competition among platforms replicates the community-ownership outcome. We prove this here.

Game. Throughout we maintain $g(k^W) > 0$ and g strictly decreasing where positive, as in Proposition 2. Two identical platforms simultaneously announce contracts (f_j, τ_j) , $j \in \{1, 2\}$, where f_j is a membership fee and τ_j is a per-unit capacity charge on storage (a subsidy if negative). Given (f_j, τ_j) , competitive storage entry on platform j yields capacity per member k_j satisfying $g(k_j) = c + \tau_j$. A prosumer who joins platform j obtains expected surplus

$$u_j(\phi) = (\Delta_0 - \Delta(k_j)) v(k_j) - f_j - \phi. \quad (\text{B.31})$$

Each prosumer joins the platform offering higher surplus $u_j(\phi)$; in case of a tie, prosumers split equally. Prosumers who prefer not to join either platform remain independent. Platform j 's profit is $\Pi_j = (f_j + \tau_j k_j) N_j$, where N_j is the mass of prosumers who join platform j .

Equilibrium. We show that in any Nash equilibrium both platforms earn zero profit and implement k^W . The argument proceeds in two steps.

Step 1: Zero profit. Let $M_j \equiv f_j + \tau_j k_j$ denote platform j 's margin per member, so $\Pi_j = M_j N_j$. Suppose some platform earns $\Pi_j > 0$; then $M_j > 0$ and $N_j > 0$. Let $i \neq j$ be the platform with $\Pi_i \leq \Pi_j$. Since j has members, i 's offer gives prosumers weakly lower surplus than j 's. Let i deviate to $(f_j - \varepsilon, \tau_j)$ with $0 < \varepsilon < M_j/2$. The induced capacity is unchanged— $g(k) = c + \tau_j$ still gives k_j —and i 's offer now yields strictly higher surplus than j 's, so every prosumer who joins either platform joins i : $N'_i \geq N_i + N_j$. If $N_i = 0$, the deviation earns $(M_j - \varepsilon)N_j > 0 = \Pi_i$. If $N_i > 0$, both platforms had members, so their offers tied and the tie-splitting rule gives $N_i = N_j = N/2$ with N the total mass served; then $\Pi_i \leq \Pi_j = M_j N/2 < (M_j - \varepsilon)N$ for $\varepsilon < M_j/2$. Either way the deviation is strictly profitable, so no candidate with positive profit for either platform is a Nash equilibrium. In any equilibrium $\Pi_j = 0$ for both; a negative margin is ruled out because a platform can always deviate to an offer no one accepts and earn zero, and for any platform with members zero profit requires

$$f_j = -\tau_j k_j. \quad (\text{B.32})$$

Step 2: Efficient capacity per member. Under zero profit, substitute $f_j = -\tau_j k_j$ into prosumer surplus. Using $\tau_j = g(k_j) - c$ (from competitive storage entry) and $g(k_j) = \Delta(k_j) v(k_j)/k_j$:

$$f_j = -(g(k_j) - c) k_j = c k_j - \Delta(k_j) v(k_j). \quad (\text{B.33})$$

Prosumer surplus becomes

$$\begin{aligned} u_j(\phi) &= (\Delta_0 - \Delta(k_j)) v(k_j) - (c k_j - \Delta(k_j) v(k_j)) - \phi \\ &= \Delta_0 v(k_j) - c k_j - \phi \\ &= S(k_j) - \phi. \end{aligned} \quad (\text{B.34})$$

The $\Delta(k_j)$ terms cancel: one enters positively through improved terms of trade, the other negatively through the fee that recovers the storage subsidy. Under zero profit, prosumers receive the full corridor net of resource costs, regardless of the settlement rule.

Each platform attracts prosumers by offering higher surplus $S(k_j)$. Suppose a zero-profit candidate has $k_1 \neq k^W$, so $S(k_1) < S(k^W)$ by uniqueness of the maximizer (Section 2.5). Platform 2 deviates to $k_2 = k^W$, implemented by $\tau_2 = g(k_2) - c$ and $f_2 = -\tau_2 k_2 + \varepsilon'$ with $0 < \varepsilon' < S(k^W) - \max\{S(k_1), 0\}$. Prosumer surplus on platform 2 is then $S(k^W) - \varepsilon' - \phi > S(k_1) - \phi$, so every prosumer who joins either platform joins 2, and $N'_2 = \min\{1, (S(k^W) - \varepsilon')/\bar{\phi}\} > 0$. Platform 2's profit is $\varepsilon' N'_2 > 0$, contradicting $\Pi_2 = 0$. Hence both platforms set $k_j = k^W$ in any equilibrium; at $k_j = k^W$ with zero margins no deviation attracts members at a positive margin, so the profile is an equilibrium.

Outcome. At $k_j = k^W$, both platforms offer surplus $S(k^W) - \phi$. Participation is $N^* = \min\{1, S(k^W)/\bar{\phi}\} = N^W$, split equally between platforms. The outcome—efficient storage investment k^W , undistorted participation N^W , and 100% of the potential surplus captured—is identical to community ownership (Section 4.2).

The result assumes homogeneous platforms and at least two viable entrants; with differentiated platforms or a natural monopoly at the neighborhood scale, the outcome lies between the monopoly and community-ownership benchmarks (Section 4.4).

Appendix B.7. Outside options for the storage fringe

The main text models the competitive fringe as serving a single community. Suppose instead that each unit of storage capacity can alternatively be deployed outside the community—serving other communities, a distribution system operator, or wholesale arbitrage—earning an exogenous reservation return r per unit of capacity per day, with elastic outside demand. We take r as the return available in a large outside market; were entry into that market also free, r would be competed down to c and the single-community case of the main text is recovered. This presumes that deployable capacity is drawn from a common pool, so that a unit installed in the community displaces one outside it; if instead capacity can be added to the community at cost c without displacing any outside deployment, the community condition remains $g(k) = c$ and the outside option is irrelevant. Entry into the community then requires the internal return to match the best alternative, so any equilibrium with positive community storage satisfies

$$g(k) = \max\{r, c\}. \quad (\text{B.35})$$

Lemma 2 (Outside options reduce community storage and preserve the sign of the distortion). *Let g be strictly decreasing where positive, let $k^D(z)$ denote the solution to $g(k) = z$ with $k^D(z) = 0$ if $\Delta(0^+)v'(0) \leq z$, and let $k^W(z)$ solve $\Delta_0 v'(k) = z$, with $k^W(z) = 0$ if $z \geq \Delta_0 \Pr(m > 0)$. For any $r \geq c$, community storage satisfies $k^D(r) \leq k^D(c)$, strictly whenever $r > c$ and $k^D(c) > 0$; and, whenever the settlement rule under-invests at cost c for every cost-corridor ratio (as matched-market SDR does for every mismatch distribution, Corollary 1), under-investment persists against the benchmark set by the opportunity cost of capacity, $k^D(r) < k^W(r)$, for every r with $0 < r/\Delta_0 < \Pr(m > 0)$; for $r \geq \Delta_0 \Pr(m > 0)$ both capacities are zero.*

Proof. Immediate from the monotonicity of g : a higher break-even threshold is reached at a smaller capacity per member. The comparison to k^W is Corollary 1. $\square$

Three remarks. First, the equilibrium *margin* is higher with the outside option— $\Delta(k^D(r)) \geq \Delta(k^D(c))$ because less capacity is installed locally—which is the sense in which the ability to divert capacity elsewhere “mitigates” spread compression. But the mitigation operates by shrinking the community battery. Whether the *gap* widens is a separate matter: when r is a genuine opportunity cost the efficient benchmark $k^W(r)$ falls as well, so neither the absolute shortfall nor the ratio k^D/k^W is necessarily worse; what persists is the sign (Lemma 2). The margin schedule $\Delta(\cdot)$ is itself unaffected: under SDR the supply-demand ratio is computed on quantities internal to the community, so compression depends on the capacity deployed *in* the community and not on the operator’s wider portfolio. Second, if the outside return $r > c$ reflects genuine value of capacity elsewhere, the community’s efficient benchmark changes as well: the relevant opportunity cost of capacity becomes r , and the efficient capacity per member solves $\Delta_0 v'(k) = r$. The structure of the financing problem is unchanged—Corollary 4 applies verbatim with $\alpha = r/\Delta_0$ for $r/\Delta_0 \in (0, 1)$ —so a non-energy instrument, now a fee financing capacity procurement at the market return r , remains necessary for efficiency, and the community-ownership result of Proposition 3 carries over with c replaced by r . Third, if instead the same capacity earns an outside return r^{out} per unit *alongside* the internal margin (value stacking, as in battery-as-a-service arrangements), the effective cost of capacity to the community falls to $c - r^{\text{out}}$, provided $r^{\text{out}} < c$; Corollary 4 then applies with $\alpha = (c - r^{\text{out}})/\Delta_0$, and Table 5 shows that K^D/K^W stays between a tenth and a fifth over the cost range it reports.

Table B.2: Nomenclature of variables.

Variable	Description
$\underline{p}$	Feed-in tariff (price received for exporting 1 kWh).
$\bar{p}$	Retail price (all-in marginal cost of importing 1 kWh).
$\Delta_0 \equiv \bar{p} - \underline{p}$	External price corridor (retail minus feed-in).
$\alpha \equiv c/\Delta_0$	Cost-corridor ratio.
δ^*	Constant administrative spread satisfying (15): $\delta^* = \Delta_0 v'(k^W)/(v(k^W)/k^W)$.
ω	Daily state (realization of PV generation and consumption conditions).
$m(\omega)$	Matchable (time-shiftable) energy per participant in day ω .
$s(\omega), d(\omega)$	Period-1 surplus and Period-2 deficit per participant; $m = \min\{s, d\}$ (Appendix B.1).
$\bar{m}$	Upper bound of matchable energy.
F	CDF of $m(\omega)$ across days.
N	Mass of participating prosumers (uptake), $N \in [0, 1]$.
ϕ	Idiosyncratic adoption cost of joining the community.
$\bar{\phi}$	Upper bound of adoption costs; $\phi \sim U[0, \bar{\phi}]$.
K	Total installed community-level battery capacity.
$k \equiv K/N$	Capacity per member (installed capacity per participating prosumer), for $N > 0$.
c	Per-period (levelized) resource cost per unit of storage; total resource cost cK .
$q(\omega, k) \equiv \min\{k, m(\omega)\}$	Shifted energy per participant in day ω .
$v(k) \equiv \mathbb{E}_\omega[q(\omega, k)]$	Expected shifted energy per participant.
$p^1(\omega, k)$	Internal price in the surplus period (state ω , capacity per member k).
$p^2(\omega, k)$	Internal price in the deficit period (state ω , capacity per member k).
$\Delta(\omega, k) \equiv p^2(\omega, k) - p^1(\omega, k)$	Internal within-day spread in state ω .
$\Delta(k)$	Expected internal margin per shifted kWh (volume-weighted average).
$\Delta(0^+) \equiv \lim_{k \downarrow 0} \Delta(k)$	Limiting internal margin as capacity becomes negligible.
$g(k) \equiv \Delta(k) v(k)/k$	Arbitrage revenue per unit of capacity.
$R(k) \equiv \Delta(k) v(k)$	Expected energy-only revenue of the fringe per participant.
$\rho(\omega, k)$	Supply-demand ratio in the deficit period (Appendix B.2).
$\text{SDR}_2 \equiv q/m(\omega)$	Period-2 ratio on the matched-market base.
$\text{SDR}'_2 \equiv q/d(\omega)$	Period-2 ratio on the full-deficit base.
$\tilde{g}(K)$	Unit revenue under the full-deficit base.
$K^{D'}$	Energy-only equilibrium capacity under the full-deficit base.
$\lambda \equiv d/m$	Deficit-to-matchable ratio.
$\vartheta(k) \in [0, 1]$	Equality-state price selection under scarcity clearing (Appendix B.4).
θ_i	Member i 's benefit scale factor, $v_i(k) = \theta_i v(k)$, mean one (Section 4.2).
η	Round-trip efficiency; $\Delta_0^\eta = \eta \bar{p} - \underline{p}$ (Section 2.3, Appendix D.5).
f	Membership fee charged to participating prosumers (platform regime).
τ	Per-unit capacity charge levied on storage (a subsidy to storage if $\tau < 0$).
r	Return per unit of capacity available to the fringe outside the community (Appendix B.7).
r^{out}	Outside return earned <i>alongside</i> the internal margin (value stacking).
ϱ, L, I	Real discount rate, economic lifetime and installed cost (Appendix A).
$u(\phi)$	Prosumer expected utility: $u(\phi) = (\Delta_0 - \Delta(k))v(k) - f - \phi$.
Π	Platform profit: $\Pi = fN + \tau K$.
W	Community surplus: $W = \Delta_0 Q - cK - \frac{\bar{\phi}}{2} N^2$ (with $Q = 0$ if $N = 0$).
$S(k) \equiv \Delta_0 v(k) - ck$	Net surplus created by one participant at capacity per member k (planner object).

Appendix C. Individual versus shared storage

A standalone household battery outside any community has no financing problem of the kind we study: it shifts the household’s own surplus into its own deficit and earns the full corridor Δ_0 per kWh without any internal settlement, so it is the community-owned regime of Section 4.2 at $N = 1$. A battery owned by a community *member* is different: when it diverts energy from a contemporaneous internal match its private return depends on the settlement rule and can fall short of Δ_0 . What a standalone battery forgoes is pooling, which has two distinct components. First, within the synchronous structure of the model—every member has a period-1 surplus $s_i \geq 0$ and a period-2 deficit $d_i \geq 0$ —pooled matchable energy weakly exceeds the sum of the households’: $m = \min\{\sum_i s_i, \sum_i d_i\} \geq \sum_i \min\{s_i, d_i\}$, strictly when households sit on different short sides. Two households with $(s, d) = (10, 1)$ and $(1, 10)$ have $\sum_i m_i = 2$ but $m = 11$. Moreover a shared battery of capacity K shifts weakly more than the same capacity split across households: $\min\{K, m\} \geq \sum_i \min\{k_i, m_i\}$ for any split with $\sum_i k_i = K$, since the right-hand side is bounded by both $\sum_i k_i = K$ and $\sum_i m_i \leq m$. The same resource cost therefore buys weakly more throughput at higher utilization. Second, where some households are in deficit while others have surplus in the same period, contemporaneous netting serves part of the load without any storage—with period-1 positions +10 and −10 the pooled surplus is zero and $m = 0$, although $\sum_i m_i = 10$ —which removes shifting *need* rather than adding throughput, and reverses the first inequality once the synchronous hypothesis fails. Whether matchable energy per member rises or falls with community size therefore depends on which component dominates in the profile mix, and is an empirical question; quantifying either requires household-level meter data, which our community-level dataset does not contain (Section 5). The financing gap we identify bears on the first component: it determines whether the pooled asset is built or whether households fall back on individual batteries. Where households could install individually at the same unit cost, the reservation payoff would be S^{indiv} rather than zero, and the pooling components above are what make $S^{\text{indiv}} \leq S(k^W)$, strictly when households sit on different short sides.

Appendix D. Calibration details

Appendix D.1. Tariff components and prices

Based on the CKW 2026 tariff structure for residential customers (*Basiskundengruppe*, excl. MWST): the volumetric network charge (*Netz-Einheitstarif*) is 6.00 cents/kWh, Swissgrid system services (*SDL*) 0.27 cents/kWh, the federal strategic reserve levy (*Stromreserve*) 0.41 cents/kWh, solidarised costs 0.05 cents/kWh, the federal renewable energy promotion levy (*Netzzuschlag*) 2.30 cents/kWh, and a municipal concession fee (*Konzessionsabgabe*) of 0.95 cents/kWh. The first five are published on the 2026 tariff sheet, which lists them under *CKW Netz E* but states no total and refers the concession to a municipality-specific schedule; we take the concession at 0.95 cents/kWh, so that the volumetric network charges and levies avoided by internal self-consumption total 9.98 cents/kWh. These charges are avoided by internal self-consumption in a ZEV but primarily reflect fixed infrastructure cost recovery, policy levies, and transfers (see tariff details in CKW, 2025b). We take the residential basic-group tariff as the reference stack; the community’s own tariff group is not reported to us, and a demand-metered tariff would recover part of the network cost through a capacity charge rather than volumetrically, narrowing Δ_0 and raising $\alpha = c/\Delta_0$ —the margin along which Table 5 reports the financed share. These components are year-specific—the 2027 sheet already differs—so Δ_0 is a 2026 figure. The energy procurement component is the element that most closely tracks marginal wholesale costs; it was 8.50 cents/kWh in 2025 and 12.00 cents/kWh in 2026, and it is the 2026 figure that enters $\bar{p}$ in Section 5 (CKW, 2025a). The feed-in tariff tracks the BFE quarterly reference market price, with a statutory minimum of 6 cents/kWh for PV installations under 30 kW plus a voluntary CKW top-up of 2 cents/kWh for certificates of origin (*Herkunftsnachweise*); the Q4 2024 reference price was 8.8 cents/kWh (CKW, 2025c). We take 8.00 cents/kWh as the export price on this basis; the installation-level configuration behind the community’s 99.4 kWp is not reported to us, so for a configuration to which the statutory minimum did not apply the operative tariff would follow the quarterly reference price instead. Section 5.5 reports how the financed share moves with $\bar{p}/\underline{p}$ at a fixed corridor, and separately the case $\underline{p} = 8.80$, in which the corridor narrows and the price ratio falls together.

Appendix D.2. Dispatch simulation and validation of the one-cycle form

To ensure physical feasibility, we simulate a battery dispatch for each day. The simulation proceeds chronologically through the 96 quarter-hourly intervals of each day. In each interval with a community-level PV surplus (generation exceeding consumption), the battery charges at the minimum of the available surplus and the remaining capacity, subject to a 1C power limit (i.e., maximum charge/discharge power equals the capacity in kW). In each subsequent interval with a community-level deficit, the battery discharges at the minimum of the deficit and the stored energy, again subject to the 1C limit. The battery starts each day empty. This simulation and the validations below are computed on the quarter-hourly series; see the Data availability statement for what the replication package contains. Resetting at midnight could in principle forgo shifting of late-evening surplus into post-midnight deficit. Re-running the dispatch with carry-over across days changes shifted energy by 0.0% at $K = 10\text{--}20$ kWh and by 0.3% at $K^W = 37.66$ kWh; the effect becomes material only far above the capacities studied here. We use a large reference capacity (effectively unconstrained) so that the daily total $m(\omega)$ captures the full physically matchable energy.

The 1C limit is conservative for modern lithium-ion systems and matches typical residential battery specifications. Because the reference battery is unconstrained in capacity, the power limit—not capacity—is the binding constraint on days with very concentrated surplus peaks. In the data, however, PV surplus is spread over 4–6 midday hours for most days, so the power constraint rarely binds before the surplus window ends.

One-cycle-per-day validation. The functional form $V(K) = \mathbb{E}[\min\{K, m(\omega)\}]$ treats K as the effective daily throughput limit, which implicitly assumes one full charge–discharge cycle per day. We validate this by examining whether secondary surplus windows (which would enable a second cycle) are present in the data. Across all 656 days, the share of total matchable energy attributable to a second or subsequent surplus–deficit cycle is less than 1.5%. The generation profile of this ZEV is structurally unimodal: a single consolidated midday surplus block followed by a consolidated evening/night deficit block. The one-cycle approximation is therefore accurate for this community; it may be less appropriate for communities with significant midday demand dips or multiple generation technologies.

Dispatch feasibility at smaller capacities. The functional form $V(K) = \mathbb{E}[\min\{K, m(\omega)\}]$ implicitly assumes that a battery of capacity K can fully charge and discharge $\min\{K, m(\omega)\}$ kWh within each day’s temporal structure. Under the 1C power limit, charging K kWh requires at least one hour (four quarter-hourly intervals). Since the surplus window in this ZEV spans 4–6 midday hours on typical production days, this constraint is not binding for the capacity range relevant to the calibration ($K^W \leq 50$ kWh with 100 kWp installed PV). We verify this by running the chronological dispatch simulation at each candidate capacity $K \in \{1, 2, \dots, 150\}$ kWh and confirming that the simulated shifted energy matches $\min\{K, m(\omega)\}$ to within 1% on more than 99% of days. For communities with more concentrated surplus peaks or lower PV-to-storage ratios, the power constraint could bind at lower capacity levels, and a dispatch-based $V(K)$ should replace the newsvendor form.

The reference battery is “large” only in the sense that capacity does not bind when measuring $m(\omega)$; dispatch feasibility still respects the imposed power limit (1C) and the chronological charge–discharge constraint within each day.

Appendix D.3. Two-period versus sub-daily SDR computation

The SDR margin $\Delta^{SDR}(K)$ is computed using the model’s two-period aggregation: each day’s surplus and deficit are consolidated into single blocks, and the SDR algorithm is applied once per period. In an actual 15-minute implementation, the supply–demand ratio would be computed interval by interval, and the resulting spread would depend on the within-day temporal profile of surplus and deficit—not only on their daily totals. Two effects could cause the 15-minute and two-period margins to diverge. First, within-day variation in the surplus–deficit ratio would generate heterogeneous SDR values across intervals, and the volume-weighted average margin could differ from the margin computed on aggregated quantities. Second, the margin would depend on the dispatch rule. We conjecture that the two-period aggregation understates the financing gap, for the following reason, although the comparison cannot be settled without specifying

how interval-level internal demand is allocated—which the matched-market computation leaves open. The discharge price is decreasing in the period’s supply–demand ratio. Under the chronological dispatch we simulate, the battery discharges $\min\{\text{deficit}_t, \text{stored energy}\}$ in each deficit interval, so in every interval in which it covers the entire deficit the interval ratio equals one and the discharge price equals the feed-in tariff $\underline{p}$: the margin on that energy is zero. Energy discharged in those intervals earns a positive margin only in the two-period computation, where it is pooled with the rest of the evening deficit. Interval-level settlement with this dispatch would therefore lower the margin on energy discharged in fully served intervals; whether the volume-weighted margin over all intervals is lower depends on the margin earned in partially served intervals under the 1C limit, which can exceed the daily matched-market margin, so the direction is not established here. A storage operator could mitigate this by spreading discharge across intervals so that no interval ratio reaches one, which makes the energy-only margin dispatch-dependent in a way the two-period model abstracts from. Confirming the direction and magnitude through a full 15-minute SDR simulation with alternative dispatch rules is beyond the scope of this paper but would be a valuable robustness check in future work.

Appendix D.4. The period-2 deficit under the full-deficit base

The full-deficit base of Section 5.4 measures the period-2 ratio against $d(\omega)$, constructed as $\sum_t \max\{0, \text{consumption}_t - \text{production}_t\}$ over the intervals that follow the day’s first surplus. We take the period-2 deficit rather than the whole day’s for two reasons. The two-period model places period 2 *after* period 1, so morning deficit that precedes any PV surplus belongs to neither period; and the published rule is defined per time slot, so in a sub-daily implementation those intervals are separate slots in which internal supply is zero and nothing is traded, which makes folding them into the period-2 denominator an aggregation artifact rather than a reading of the rule. Summing the deficit over all 96 intervals instead gives a mean of 319 kWh and a benchmark financed share of 45.1% in place of 20.9%, so a whole-day denominator is an upper bound on the financed share.

The highest-cost scenario of Table 6 is the narrowest case: at $\alpha = 0.9$ unit revenue at K^W falls short of cost by 2.7% ($\tilde{g}(K^W) = 12.24$ against $c = 12.58$), and resampling the 656 days 4,000 times with K^W recomputed on each resample puts $\tilde{g}(K^W) < c$ in 97.2% of resamples, with a 95% interval for $\tilde{g}(K^W) - c$ of $[-0.80, 0.00]$ cents. Under-investment is therefore established in every cost scenario under this base as well, most narrowly at the highest cost.

Appendix D.5. Round-trip losses

This appendix re-runs the dispatch of Section 5.2 and the entry condition of Section 5.4 with a round-trip efficiency of $\eta = 0.90$. Measuring capacity in charged kWh, so that charging q delivers ηq at discharge and the cost c per unit of rated capacity is unchanged, the effective corridor per charged kWh is $\Delta_0^\eta = \eta\bar{p} - \underline{p} = 11.78$ cents/kWh and $\alpha^\eta = 0.831$. Matchable energy rises in charge units, because a given deficit now requires $1/\eta$ as much stored energy ($\mathbb{E}[m^\eta] = 66.1$ against $\mathbb{E}[m] = 61.6$ kWh), and $K^W = 10.22$ kWh. The financed capacity solves $g^\eta(K) = c$ with the operator’s margin per charged kWh equal to $\eta p^2(\rho) - \underline{p}$ and $\rho = q/m^\eta$, giving $K^D = 0.83$ kWh and $K^D/K^W = 8.1\%$. The surplus-optimal capacity falls by three quarters because α^η moves close to $\Pr(m > 0) = 0.93$, while the financed share falls modestly, from 9.9% to 8.1%: the level is sensitive to losses, the ratio much less so. Two qualifications follow. First, the net community surplus at the re-optimized capacity is $S^\eta(K^W) = 3.9$ cents/day, CHF 14 per year against CHF 182 at $\eta = 1$, with the balanced-budget fee absorbing 96% of the gross value: the *level* of the community-ownership gain reported in Section 5.5 does not survive realistic losses at benchmark costs, although the ranking of the regimes does; as with the lossless case the level recovers quickly as batteries get cheaper, to CHF 136 per year at $\alpha = 0.60$ and CHF 374 at $\alpha = 0.50$. Second, the margin per charged kWh turns negative once $p^2 < \underline{p}/\eta = 8.89$ cents, i.e. once $\rho > (\eta\bar{p} - \underline{p})/\Delta_0 = 0.843$, so a merchant operator facing losses stops charging before scarcity is relieved rather than shifting at a loss. At K^D that threshold is crossed on 17 of the 609 operating days and on 1.5% of shifted volume, and at K^W on 74 days and 5.5% of volume, so the truncation does not bind where the entry condition is evaluated; it would matter for a community operating well above its financed capacity.