I'll Move If You Move: Assurance Campaigns for Platform Migration

Matthew Cashman*

July 16, 2026 (latest version here)

Abstract

A new platform can attract accounts without attracting a community. People may want to move yet remain on an incumbent because the friends and audiences they value are still there. Migration from X to Bluesky illustrates this cold-start problem. I ask whether an entrant can help a community move together. In an assurance campaign, people pledge to make the new platform their main venue if enough others pledge. If the campaign succeeds, everyone is told that it is time to move. Following through remains voluntary, and no one must abandon the incumbent. I compare this campaign with ordinary migration backed by the same practical support. It is worthwhile when the added chance of a successful group launch outweighs its cost and delay. I then ask what people should see as pledges accumulate. Names can attract participants when people they value have joined, but discourage them when the early group looks less promising. A count can therefore sometimes work better. An entrant may overcome the cold start by helping users coordinate, but revealing more does not always produce a better outcome.

Keywords: Platform migration; assurance campaigns; network effects; disclosure design; sequential participation

1 Introduction

A firm entering a market with strong network effects faces a cold-start problem: its product is worth little to any one customer until many adopt it together, so the demand it needs arrives in groups or not at all. For the operator of a destination-side platform, service, or client, the customers who matter are often stuck on an incumbent in this way, and recent migration from X to Bluesky makes the resulting acquisition problem visible: opening a destination account is not the same as moving meaningful activity there. X-to-Bluesky measurement defines transition as account creation, documents continued X use among many transitioners, and measures first-day reconnection to available X ties (Quelle et al. 2025). Bluesky starter packs make named accounts easier to discover and follow, but membership is not an agreement to

*Cashman: MIT Sloan School of Management, cashman@mit.edu.

move (Balduf et al. 2025). Earlier Twitter-to-Mastodon measurement is context about linked users, accounts, instances, and social connections, not a model calibration (He et al. 2023). These studies neither implement nor validate the pledge mechanism, the activity model, or the three service policies proposed below. No cited evidence establishes how prevalent the targeted segment is, whether its members will incur consequential durable preparation costs, whether invitees comprehend the proposed interface, or how often a live count leaves invitees uncertain who has already pledged. The product question is whether the operator of a particular destination-side migration service or client can help a bounded professional or social community originate enough activity and active ties for Bluesky to become its default venue for a stated context.

I study a hypothetical *assurance campaign*: a bounded, managed, expiring, one-shot invitation wave run by that operator. A participant prepares a destination account and records the same nonpersonalized pledge to move the default venue for the stated context if enough eligible accounts pledge before a deadline. Enrollment obtains informed consent to publication of the final signer roster on success (the *successful roster*) and, under the live-roster policy defined below, to immediate identity display after acceptance. Pledges arrive sequentially. When the threshold is first reached, the operator publishes the roster and sends a common notice. The operator controls only its wave's eligibility, outreach, interface, record, publication, and notices. It cannot compel later action or govern users, protocol components, or third-party services outside its authority. Participants may continue using X, and every destination-side launch decision remains voluntary, without a bond, penalty, legal enforcement, or technological lock.

The operator of this service or client is also the *provider* whose expected value the analysis ranks; I write operator for the campaign mechanics and provider for that objective. It faces two nested choices. Policy N, ordinary migration with no campaign, retains the same account creation, matching, lawful import, partial tie reconstruction, starter packs, public activity, announcements, and onboarding as the campaign policies, but adds no contingent record or common threshold notice. This is a deliberately strong benchmark: research on friend-finding tools shows that copying prior ties can quickly produce a connected, reciprocal, clustered destination network and that copied connections see greater interaction (Zhong et al. 2014). Policy C adds a live authenticated count while withholding signer identities from the campaign interface during accrual. Policy R differs from C only by showing the complete live roster. Both campaign policies reveal the same successful roster and use the same notice, reconstruction tools, and voluntary activity stage. The first question is whether either assurance policy adds expected activity-and-tie value over N. The second, conditional question is whether live names improve on a live count.

The first comparison translates established adoption and assurance logic into a platform-acquisition decision. Interactive media can stall below critical mass because each user's benefit depends on reciprocal participation (Markus 1987). Expectation-dependent switching between standards can produce excess inertia (Farrell and Saloner 1985). Expectations also shape competition when products are incompatible (Katz and Shapiro 1985), and product preannouncements can alter adoption (Farrell and Saloner 1986) and diffusion (Le Nagard-Assayag and Manceau 2001). Marketing research shows why multihoming and complement overlap cannot be reduced to raw platform size (Landsman and Stremersch 2011). These forces explain why the provider's acquisition target is realized destination activity and active ties rather than account creation alone.

An assurance campaign also shares the preconsumption timing of advance selling and the threshold structure of group buying. Advance selling turns early demand into purchases before consumption and can expand sales when buyers are uncertain about their future valuations (Xie and Shugan 2001). Group buying can pair a required group size with price to induce informed customers to recruit others (Jing and Xie 2011); real-time sign-up data show a surge around the threshold (Wu et al. 2015), and sequential group-buying research shows that visible progress can change later participation (Hu et al. 2013). A migration pledge is different: it is neither a purchase nor a binding reservation to use the destination, and reaching the threshold does not produce destination activity. This distinction is why the model separates campaign formation from voluntary execution.

The underlying assurance logic is established. Count-contingent adoption insurance addresses positive adoption externalities (Dybvig and Spatt 1983). Provision-point mechanisms develop the same broad logic (Palfrey and Rosenthal 1984; Bagnoli and Lipman 1989), and dominant assurance is another established variant (Tabarrok 1998). Proposition 1 accordingly makes no novelty claim for assurance itself. It gives the provider a net assurance-lift formula relative to otherwise identical ordinary migration support.

The narrower contribution concerns cohort composition. Marketing work on network diffusion shows why composition matters: influentials and imitators play asymmetric roles (Van den Bulte and Joshi 2007), innovator and follower hubs affect adoption speed and eventual market size differently (Goldenberg et al. 2009), and seed choice changes diffusion (Hinz et al. 2011). Network position makes equal-sized sets economically different (Ballester et al. 2006). The information shown during recruitment can also change later behavior: public communication affects higher-order beliefs (Chwe 2001); crowdfunding research studies predecessor histories (Lin et al. 2024) and identity concealment (Burtch et al. 2016); and related work studies threshold-conditioned identity disclosure (Braghieri et al. 2026). These literatures do not make a roster uniformly desirable. A live count pools early identities; a live roster reveals which people and prospective ties are represented. That information changes recruitment and can therefore change the value of the migration cohort.

Proposition 1 separates whether the campaign fills (formation) from whether signers then move (voluntary execution) and gives the net assurance lift; its formula, the probability the campaign fills times the conditional probability that voluntary launches tip destination activity into its high basin ($\alpha_d q_d$), uses a coarse two-level approximation to horizon value, while the general comparison uses expected full-path value. Proposition 2 is an uncalibrated existence result in a five-account example built around an *anchor*: one prominent early signer whose identity matters to a later pivotal invitee. Averaged over a random arrival order common to both policies, and holding execution, outside behavior, provider branch values, and costs equal across them, neither display always recruits better. The roster gives better information: a count can induce preparation after the weaker anchor that disclosing that identity would deter, while a roster can recruit after the stronger anchor that a pooled count would not. The count advantage is strategic opacity in recruitment, not total product value or welfare. Section 2 describes the product, Section 3 gives the common environment, Section 4 presents the results, and Section 5 develops the decision rules and scope conditions.

2 The Assurance Campaign

Table 1 defines this conditional, untested Bluesky product hypothesis for a screened community: a bounded, managed wave run by the operator of a particular destination-side service or client. Uptake and trust are pilot estimands: quantities a pilot must measure, not assumptions the model supplies. Within the functions it controls, the operator keeps the baseline stack identical across policies: N retains portability, discovery, public communication, and ordinary activity opportunities. Assurance adds the prepared pledge, policy-required identity-display consent, authenticated live state, threshold, deadline, and first-hit notice; launch remains voluntary. Policies C and R differ in whether identities are displayed live with consent and in the information that display conveys.

Table 1: Matched service policies

Feature	N: ordinary	C: count	R: roster
Account creation, lawful import, matching, partial tie reconstruction, starter packs	Same baseline	Same baseline	Same baseline
Public activity, announcements, and ordinary onboarding	Available	Available	Available
Prepared contingent pledge	None	Common pledge	Common pledge
Identity-publication consent	None	Successful roster	Live and successful roster
Live campaign display	None	Authenticated count	Count and identities
Operator threshold, deadline, and first-hit notice	No assurance campaign	Same rule	Same rule
Successful roster	Not applicable	Released	Released
Destination action and incumbent use	Voluntary; multihoming allowed	Voluntary; multihoming allowed	Voluntary; multihoming allowed

Before outreach, the operator precommits to and logs a treatment-invariant rule for the bounded community, finite eligible set, threshold K, short expiration time D, and activity context across N, C, and R. Account-control authentication prevents one qualifying linked account from occupying two record positions. It does not establish unique people, eliminate duplicate-person accounts or matching error, or say anything about sincerity, readiness, or performance. Acceptance requires the common preparation bundle: creating or linking the destination account, completing authentication, setting up the profile and lawful import, and staging a launch. It also requires one nonpersonalized migration pledge. The pledge is an agreement to move the context's default venue after success. Enrollment separately obtains informed consent to successful-roster publication under C and R and to immediate live identity display under R; no pledge term depends on a future named signer. The eligibility rules, exclusions, overlap, matches, and handling of duplicates are auditable.

Each eligible account receives one effective opportunity in a policy-invariant sequence and either accepts or declines. Later invitees observe the state produced by earlier acceptances: under C, the clock and authenticated count; under R, those objects and signer identities. A decline is not publicly labeled. The short-window strategic benchmark is append-only: invitees cannot wait, revisit a decision, or condition acceptance on a future signer. Operationally, if an account withdraws consent required by its assigned display policy before the trigger, the operator deletes the pledge, removes an R identity prospectively, excludes the account from the successful roster, and decrements the count; copies already made under R cannot be erased. Announced removals for fraud, lost authentication, law, harassment, or safety do the same. Neither decrement is available as a strategic action in the append-only witness.

The formal result assumes that events do not tie and that the campaign closes at the exact first hit. As soon as the count reaches K before D, accrual closes; publication uses consent already obtained, the same successful roster appears under C and R, and the operator sends a common timestamped notice. An implementation that batches nearly simultaneous events must announce its clock tolerance and deterministic within-batch order in advance; the formal roster remains the exact first K. If the wave expires below K, no coordinated launch occurs. Count disclosure is temporary interface nondisclosure rather than permanent anonymity: a successful C campaign reveals its roster, and outside announcements or timing may reveal identities earlier.

The public states are credible only if the operator commits in advance not to fabricate counts, suppress acceptances, alter authenticated identities, selectively reveal names, or change eligibility or invitation timing after observing the state. An audit log records eligibility, authentication, acceptances, withdrawals and removals, displays, first hit, roster publication, and notice attempts. Under R, removing an identity cannot erase copies that are already public. Network reconstruction remains partial, lawful, authorized, and opt-in where required. The operator holds its own support fixed across N, C, and R, while independently operated protocol components and third-party services remain outside its control.

Success still does not move an account. After notice, each signer decides whether to undertake the stated destination-side launch action. Notice delivery, exposure, and timing must be logged; missed or staggered notices enter the execution distribution rather than being treated as receipt. The economic interpretation is making Bluesky the default venue for the campaign's context, not deleting X or ending all multihoming. The provider observes actual destination actions, activity, and ties at a common horizon H. Pledges and threshold events are process measures; they have no direct value in the provider objective.

3 Model

3.1 Sequential Recruitment

Before policy assignment, the operator's logged treatment-invariant rule fixes a finite eligible set E for one bounded, managed, expiring wave. Account i receives one effective campaign opportunity at time $\tau_i < D$, with ties occurring with probability zero. The joint law of opportunity times and E is common across C and R. At its opportunity, the account observes its private preparation cost κ_i and the policy-specific public history. A C history reports the clock and authenticated count; an R history reports those objects and the ordered set of accepted identities. In the append-only strategic benchmark, the account accepts or declines once: acceptance appends it to $S(h)$ and increments $c(h)$, while a decline creates no named public event. Operational withdrawals and administrative decrements are outside this short-window game.

The exact game stops at the first history satisfying $c(h) = K$ (ties occur with probability zero), or at D if none occurs; announced batching tolerances are implementation rules, not theorem primitives. Let S_d denote the successful roster under policy $d \in \{C, R\}$ and set $S_d = \emptyset$ after expiration. The eligibility and opportunity process, authentication, threshold, deadline, successful-roster consent and publication, and notice are common; R additionally requires consent to its live identity message. The operator's mechanical rule and audit commitment are common knowledge. The formal recruitment comparison holds direct utility effects of display and consent outside the witness. The model treats the notice as common, while application records delivery and exposure and absorbs failures into Q below. Under N, the same accounts receive the same ordinary outreach and destination tools, but there is no recorded contingent action or common threshold event. This finite game does not cover true open enrollment, strategic waiting, repeat visits, endogenous sharing or reach, or timing-equilibrium selection; those questions are deferred rather than claimed as robustness.

The model focuses on accounts for which an active destination is attractive but the current low-activity destination is not: $U_i(B_H) > U_i(X) > U_i(B_L)$. This ordering is a conditional market segment, not a claim about every X user and not something the provider observes perfectly. It locates the coordination problem in meaningful destination activity rather than low-cost registration. Ordinary movers and the existing active seed are held fixed across all three policies.

3.2 Voluntary Execution and Provider Value

For accounts i and j, let $\omega_{ij} \geq 0$ measure how much i values j's destination activity after feasible tie reconstruction. In the focal witness, this nonnegative affinity applies only to screened identities that i weakly values; antagonistic, unsafe, or negatively valued identities require signed utility or a separate disutility channel. Directed affinities allow many users to value one prominent account or a user to value several local contacts. They are distinct from the provider's nonnegative weights in the narrow positive-spillover specialization: a persuasive roster need not be the provider's most valuable cohort, and vice versa.

For a successful roster S, let $Q(X \mid S)$ be the joint law of the subset $X \subseteq S$ that voluntarily undertakes the destination-side launch action after notice. The law integrates over signer types and may contain idiosyncratic, common, and clustered nonperformance. I hold it common to C and R conditional on the same roster. This excludes any separate effect of the earlier display history or trigger time on post-notice execution. A pledge is therefore neither an action nor a guarantee; it changes the distribution of rosters presented to the voluntary execution stage.

For each $d \in \{N, C, R\}$, let $\mathcal{W}_d$ be the random gross provider value at horizon H from realized destination activity and active ties on the full path. For a successful assurance roster S and either $d \in \{C, R\}$, the maintained common post-notice law gives $W(S) = \mathbb{E}_d[\mathcal{W}_d \mid S_d = S]$, integrating over $Q(X \mid S)$ and subsequent network activity. If $S_d = \emptyset$, the campaign expired without a provider-coordinated success notice, but $\mathcal{W}_d$ still measures the realized ordinary, failed-path, and transient outcomes. Let $A_d \geq 0$ and $D_d \geq 0$ be incremental administration and campaign-induced delay, measured relative to N in provider-value units and excluding effects already in $\mathcal{W}_d$. Then

$$J_d = \mathbb{E}_d[\mathcal{W}_d] - A_d - D_d, \quad d \in \{C, R\}, \qquad J_N = \mathbb{E}_N[\mathcal{W}_N].$$

The nested platform quantities are

$$\Delta_A = \max\{J_C, J_R\} - J_N, \qquad \Delta_R = J_R - J_C.$$

The first governs whether to add assurance; the second governs the live message only after assurance is worthwhile. This is narrow expected provider value in commensurable units. The objective excludes pledges, triggers, direct publicity value, symbolic value, and data-acquisition value; it does not assert that any is absent. Empirical use requires validated endpoints and prespecified horizons. Risk, ambiguity, hard safety constraints, participant payoffs, and social welfare remain outside the ranking.

For basin language, consider the uncalibrated scalar benchmark $y_{t+1} = G(y_t)$, not an asserted sufficient statistic for Bluesky. The same continuous increasing map governs ordinary migration and post-launch response, with fixed points $y^- < y^u < y^+$. A campaign changes only the post-notice initial condition through actual voluntary launchers. Paths below and above y^u remain in their respective basins and converge asymptotically to y^- and y^+. Convergence motivates the labels but does not imply equal finite-H activity or value; $\mathcal{W}_d$ retains same-basin transient differences. The basin structure and common map are maintained, uncalibrated assumptions. The online appendix gives the sign conditions and proof.

An additional binary-value approximation sets $\mathcal{W}_d = W_L$ for every noncrossing outcome, including expiry and successful but subcritical execution, and $\mathcal{W}_d = W_H > W_L$ for every crossing outcome. These are coarse basin-contingent horizon values, not exact finite-H activity levels implied by convergence. Let α_d be the probability of a successful roster and q_d the conditional probability that its voluntary execution crosses y^u, with $q_d = 0$ if $\alpha_d = 0$; then $\alpha_d q_d$ is their joint probability. Crossing is strict. Proposition 1 assumes no boundary atom; for a discrete atom at $y^u - y^-$, equality belongs to the low class and q_d uses the strict event.

3.3 Five-Account Witness

The disclosure result uses an uncalibrated five-account existence witness with $K = 3$. The eligible accounts are two possible anchors a and b whose authenticated identities p recognizes under R, a pivotal account p, a reliable signer s, and a reserve signer r; anchor recognition is an uncalibrated condition and pilot estimand. Exactly one anchor is ready in the campaign window. Write $Z \in \{a, b\}$ for that anchor, with $\Pr(Z = a) = \pi \in (0, 1)$. The two anchors and p receive opportunities in a policy-invariant random early block whose order is independent of Z and symmetric under exchanging the anchor labels. Invitation identities and declines are not public under C. Let $\beta \in (0, 1)$ be the probability that the ready anchor precedes p; a uniform random order of the three early accounts gives $\beta = 1/2$. The analysis integrates over the full early-block order rather than selecting a favorable sequence.

1. In the random early block, the ready anchor Z accepts and the unready anchor declines. If p arrives before Z, it sees an empty campaign and uses the same pooled decision under either display. If Z arrives first, p sees one pledge and makes the identity-sensitive decision below.

2. Account s receives the next opportunity and accepts. If p accepted, s supplies the third pledge and the campaign closes immediately with roster $\{Z, p, s\}$.

3. If the campaign remains open because p declined, r receives the next opportunity. It is ready with probability $\rho \in (0, 1)$; if ready it accepts and completes roster $\{Z, s, r\}$, and otherwise the campaign expires with two pledges.

This is an ordinary first-hit sequence: the reserve is reached only if earlier opportunities leave the campaign below threshold. On identity-sensitive histories, C pools the authenticated anchor identity, leaving posterior π, whereas R reveals Z. Relative to a common baseline that may include continued X use, normalize refusal of destination launch after prepared success to zero. Let $L_z > 0$ be incremental destination-launch value and $O_z \geq 0$ the optimized ordinary continuation after declining, including later migration after successful-roster publication. By assumption, O_z is common across C and R conditional on z, parallel to policy-common $Q(\cdot \mid S)$ conditional on successful roster S; only recruitment and cohort formation differ. Define the incremental value of accepting, before the preparation cost κ, by

$$u_z = L_z - O_z, \qquad u_a > u_b > 0, \qquad \bar{u} = \pi u_a + (1 - \pi) u_b.$$

Let W_z^P be provider value when p accepts and the campaign launches with $\{z, p, s\}$. Let B_z be provider value when p declines, integrating the probability that r completes $\{z, s, r\}$, campaign expiry if it does not, and all ordinary post-decline behavior by p. I impose a further exclusion: conditional on z and the complete realized branch, including acceptance or decline, the reserve outcome, success or expiry, and subsequent activity, the earlier display has no direct effect on gross provider value. Thus W_z^P and B_z are common across C and R. Define

$$g_z = W_z^P - B_z > 0.$$

These values integrate policy-common $Q(\cdot \mid S)$. The premise $g_z > 0$ is a realized branch-value estimand, not an implication of pledges, affinity, or centrality. Together with policy-common

O_z, the full-path branch-value exclusion, and equal campaign and display costs, this makes Proposition 2 a recruitment/cohort-formation comparison. A simple network interpretation uses $u_z = \chi + \omega_{pz}$ for a common term χ, with $\omega_{pa} > \omega_{pb}$, while W_z^P and B_z value the realized activity and active ties of p and the reserve using the provider's own weights. Thus the same destination network can make anchor a more relevant to p and make a prepared launch by p more valuable than the reserve branch without equating private affinity with provider value.

4 Results

4.1 Assurance Lift

The formation probability can be derived in a simple, uncalibrated benchmark with exactly K invitees, all of whom are required for success. This benchmark is separate from the five-account witness. Each account receives a fixed success surplus $b > 0$ relative to ordinary continuation. Acceptance requires an immediate preparation bundle: creating or linking the destination account, completing account-control authentication, configuring its profile and lawful import, and staging the campaign-specific launch. Its cost is sunk on acceptance; its prevalence and magnitude are pilot estimands. With independent costs drawn from continuous cdf F and a policy-invariant announced order, let $\alpha_{K+1} = 1$. Backward induction gives

$$\alpha_t = F(b\alpha_{t+1})\,\alpha_{t+1}, \qquad \alpha = \alpha_1.$$

Position t accepts exactly when its cost is no greater than $b\alpha_{t+1}$. In a separate stylized specialization among selected signers, iid transitory launch costs independent of preparation imply voluntary launch probability $\varphi = F_\varepsilon(v)$; iid unit contributions then give the binomial tail for at least m launches. These are signer, not population, probabilities. The online appendix gives heterogeneous derivations, the finite-game extension beyond K invitees, and correlated execution through Q.

Proposition 1 (Behavioral assurance lift). Under the additional binary-value approximation above, suppose ordinary migration is classified low, so $J_N = W_L$, assurance can be classified high only after campaign success and a strict crossing, and the account-equivalent launch contribution has no atom at $y^u - y^-$. For either $d \in \{C, R\}$,

$$J_d - J_N = \alpha_d q_d (W_H - W_L) - A_d - D_d.$$

The assurance lift is strictly positive if $\alpha_d q_d (W_H - W_L) > A_d + D_d$, zero if equality holds, and strictly negative if the inequality reverses.

Within this approximation, α_d and q_d separate recruitment from execution, and correlated execution enters through q_d. If expired, subcritical, or same-basin paths differ in horizon value, the comparison instead uses $\mathbb{E}_d[\mathcal{W}_d] - A_d - D_d - J_N$. A larger threshold can raise conditional crossing while reducing formation; the model neither optimizes K nor makes pledges directly valuable. The formula is an accounting identity under the stated approximation, not evidence that the basin model is realistic.

More generally within the binary approximation, $J_d - J_N = (b_d - b_N)(W_H - W_L) - A_d - D_d$, where b_d and b_N are assurance and ordinary high-classification probabilities; Proposition 1 sets $b_N = 0$ and $b_d = \alpha_d q_d$. Gross lift is zero if ordinary organizing reproduces the record, notice, and launch distribution, and also if ordinary activity is already high and assurance leaves it there. With increased crossing, net lift is zero at cost equality and negative above it. Proposition 1 is a platform translation of established assurance logic, not a claim of novelty or universal profitability.

The order of the two decisions matters. For the first decision, compare each assurance policy's expected full-path value with matched N; under the binary-value approximation, Proposition 1's strict positive condition provides the decision criterion. A favorable C/R recruitment difference cannot rescue assurance if its better version still has $\Delta_A \leq 0$. Only after positive assurance lift does disclosure determine recruited rosters; Proposition 2 isolates that channel under common execution, outside behavior, branch values, and costs.

4.2 Count Versus Roster

Under these exclusions, the fixed preparation cost κ is paid on acceptance, and p accepts under C exactly when $\kappa \leq \bar{u}$ and under R with displayed anchor z exactly when $\kappa \leq u_z$. If p arrives before the ready anchor, both interfaces are empty, its posterior is again π, and it accepts under either policy exactly when $\kappa \leq \bar{u}$. These cutoffs compare accepting with a genuine ordinary-migration outside option; declining does not remove p from the platform. If C succeeds, the roster is revealed before execution, so p learns z after the preparation cost is sunk but before deciding whether to launch. The maintained $L_b > 0$ is an uncalibrated voluntary-launch condition and pilot estimand; when it holds, no sanction is needed.

Because $u_a > u_b > 0$ and $\pi \in (0,1)$, the witness satisfies $u_b < \bar{u} < u_a$. It therefore contains two nonempty parameter regions in which the provider's preferred live message reverses.

Proposition 2 (Random-order disclosure reversal in recruitment). In the five-account witness, under the policy-common execution, outside-option, and full-path provider-branch restrictions, with equal campaign costs and $g_a, g_b > 0$,

$$u_b < \kappa < \bar{u} \quad \Longrightarrow \quad J_C - J_R = \beta(1-\pi)g_b > 0,$$

whereas

$$\bar{u} < \kappa < u_a \quad \Longrightarrow \quad J_R - J_C = \beta\pi g_a > 0.$$

Away from indifference points, $J_C = J_R$ when $\kappa < u_b$ or $\kappa > u_a$. Rankings at equality follow the stated tie rule and are not needed for either strict result. Under the stated common execution, outside-option, branch-value, and cost restrictions, this ranks recruitment and cohort formation, not direct display channels, total product value, or welfare.

In the first region, R gives p better information: after b, p rejects preparation, whereas C's pooled count induces preparation it would reject if the identity were known. Policy C therefore triggers with $\{b, p, s\}$, while R follows the reserve branch valued at B_b. The identity-sensitive history has probability β, which gives the first difference. In the second region, the pooled count is too weak for p, while observing a is sufficient. Policy R then triggers with

$\{a, p, s\}$ where C follows the reserve branch, giving the second difference. Histories in which p arrives before the ready anchor are identical under the two policies. Live disclosure changes recruitment, campaign completion, and cohort composition after averaging over the common random order.

A transparent numerical witness sets $\beta = \pi = \rho = \frac{1}{2}$, $L_a = 7$, $L_b = 5$, $O_a = O_b = 3$, and $g_a = g_b = 1$. Then

$$u_b = 2, \qquad \bar{u} = 3, \qquad u_a = 4.$$

At $\kappa = \frac{5}{2}$, C gains $\frac{1}{4}$ by inducing preparation after b that R's better information would deter. At $\kappa = \frac{7}{2}$, R gains $\frac{1}{4}$. Even at the larger cost, $L_b - \kappa = \frac{3}{2} > 0$, so a prepared p voluntarily launches after either anchor is revealed. These exact, uncalibrated values form an existence witness and integrate all six early-block orders; they are not forecasts or effect sizes.

As Figure 1 illustrates, equal counts can create unequal launch value inside the appendix's screened positive-spillover specialization. Scored nodes and ties are desired and policy-compliant, feasible to host or reconstruct, and nonnegative in marginal provider contribution; legality, authorization, or reconstructability alone does not establish benefit. Exact uncalibrated witnesses show that a three-mover hub or triangle can cross where three dispersed movers fail, and that a bridge can cross where a peripheral account fails. Adding only such nonnegative ties, holding weights and the threshold fixed, weakly expands the family of sufficient launch sets. Active-node and active-tie measures and network-position effects are rollout estimands, and neither private affinity nor centrality establishes $g_z > 0$.

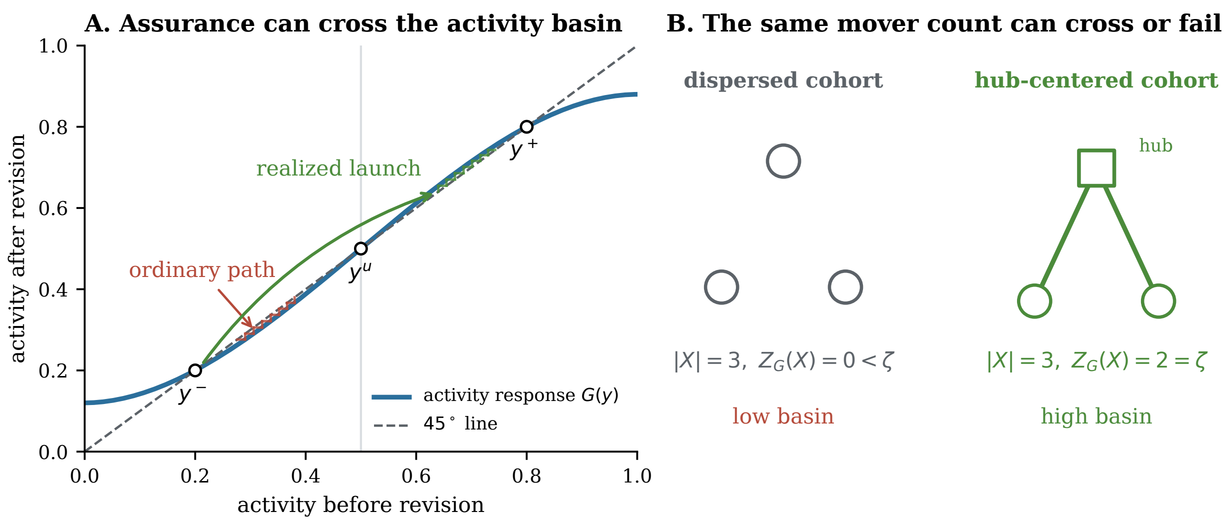

Figure 1: Assurance lift and cohort composition. Panel A illustrates the common activity response: ordinary revision from below y^u remains in the low basin, while a sufficiently large realized launch places activity above y^u and on the high path. Panel B holds the mover count fixed inside the nonnegative-score specialization: a hub-centered reconstructed cohort reaches the tie threshold while an equal-sized dispersed cohort does not. The figure is a schematic, uncalibrated existence witness.

The count-favoring region has a sharp interpretation. Roster policy R gives p better information. In state b, p would reject preparation if R disclosed the identity because the incremental value of the prepared common launch is too small to cover κ; C's pooling instead makes ex ante preparation acceptable. Eventual roster disclosure does not reverse that decision because the preparation cost is already sunk; $L_b > 0$ still makes actual launch voluntary. Account p could instead have declined and retained the ordinary path represented by O_b, and that path is included in B_b. Thus the count advantage is strategic opacity, not the suppression of ordinary migration. Positive provider value does not imply that withholding the live identity improves participant surplus or social welfare.

The opposite region is the more familiar value of composition information. The pooled count averages the high and low anchors and cannot justify preparation, while the live roster reveals the high-value state. More generally, the roster's better information can lower recruitment when weak composition deters preparation that C's pooling would induce; it can help when a valued creator, dense local group, or bridge account stimulates later recruitment and expected activity-and-tie value. The proposition does not claim that any one topology dominates in observed Bluesky communities.

Finally, count sufficiency requires the complete continuation environment to be invariant to identity permutations: at matched count-and-clock histories, relabeling cannot change future private types, readiness, opportunities, voluntary execution, administrative transitions, action payoffs, or full-path provider value. With equal C/R costs and the same permutation-invariant tie rule or equilibrium selection, finite-node backward induction makes the roster an action-irrelevant refinement and gives $J_C = J_R$. Cardinality-only provider value is not sufficient because names could still change behavior or transitions. This strong boundary keeps count versus roster secondary to whether assurance itself creates net value.

5 Managerial Implications and Scope

For a screened bounded segment, the provider should first compare assurance with matched N using the general finite-horizon difference $J_d - J_N = \mathbb{E}_d[\mathcal{W}_d] - A_d - D_d - \mathbb{E}_N[\mathcal{W}_N]$, validated activity-and-tie value and incremental-cost measures, and uncertainty. Only the binary-value approximation reduces this to $\alpha_d q_d(W_H - W_L) - A_d - D_d$ when N is low. Targeting requires evidence of the focal segment, matched ordinary opportunities, consequential preparation, comprehension, voluntary execution, residual pooling, $g_z > 0$, and desired, policy-compliant, feasible topology with nonnegative contributions.

Only after positive assurance lift compare C with R. With common execution, outside behavior, full-path provider branch values, and equal costs, Proposition 2 ranks recruitment and cohort formation: the roster gives better information; the count region is provider-optimal strategic opacity because pooling can induce preparation a disclosed identity would deter. This is no privacy or welfare benefit. Total-product choice separately measures naming effects on trust, privacy, consent, safety, moderation, execution, outside behavior, and costs.

A controlled rollout could randomize N, C, and R at the bounded-community or invitation-wave level. Because linked accounts can transmit treatment, the design would use separated network clusters or prespecify exposure and spillover estimands for both the focal account and its linked network. Audits of implementation fidelity would verify which functions the operator controls, the matched baseline stack, eligibility, cross-wave overlap, outreach, the opportunity-order distribution, account-control authentication, count integrity, reconstruction, removals, and direct name-induced activity. The protocol would specify in advance how to measure comprehension and usability, notice delivery and exposure, bot and fraud handling, attrition, missing data, and outcomes at multiple horizons. It would also distinguish preparation, acceptance, displayed state, and threshold process measures from validated activity and active-tie endpoints. Power calculations would address rare basin crossings and the precision of both policy contrasts. Uncertainty, heterogeneity, and transport across community and network types would be reported before deployment decisions. Interpreting the C/R contrast as a recruitment channel would additionally require measuring or suppressing direct effects of naming on trust, privacy, consent, moderation, execution, ordinary outside behavior, and costs. Random assignment under these controls would identify the policy contrasts for the tested design. It would not identify α_d, q_d, $W(S)$, or the channel decomposition without further measurement and assumptions. No cited evidence establishes segment prevalence, willingness to undertake consequential preparation, comprehension of the interface, or how often a live count leaves invitees uncertain who has already pledged; migration and starter-pack studies validate neither policy contrast.

Ordinary self-organization that reproduces the record and notice erases lift; contamination can obscure both contrasts. The model excludes waiting, repeat visits, endogenous reach or timing, and future-signer conditioning. Imperfect or correlated execution, leakage, incomplete reconstruction, and policy-specific removals for consent, fraud, authentication, law, harassment, or safety require their effects on outcomes, costs, and value to be measured rather than absorbed into pledge counts.

J is expected narrow provider value from destination activity and active ties net of administration and delay. Publicity and data objectives, risk and ambiguity, hard safety constraints, participant and incumbent payoffs, X displacement, and social welfare fall outside that objective. Their omission does not make them valueless, and hard constraints can override the ranking.

6 Conclusion

For the operator of a particular destination-side migration service or client, the first question in a screened bounded wave is whether assurance improves expected activity-and-tie value over matched ordinary migration. The binary-value approximation links this lift to formation, voluntary execution, crossing, cost, and delay; general evaluation uses expected full-path value. Only after positive lift does count versus roster arise. Holding execution and outside behavior fixed across policies conditional on the same roster, with equal costs, names reveal composition while a count can induce preparation the roster's better information would deter. This uncalibrated count advantage is strategic opacity, not a total-product or welfare claim. For Bluesky, the managed-wave product remains untested; direct display channels and true open enrollment require separate evidence and models.

Funding and Competing Interests

Full funding and competing-interest disclosures will be supplied separately to the editor and are withheld from this review manuscript to preserve double anonymity.

References

Bagnoli M, Lipman BL (1989) Provision of public goods: Fully implementing the core through private contributions. Rev. Econ. Stud. 56(4):583–601.

Balduf L, Sokoto S, Baronchelli A, Castro I, Król M, Tyson G, Pavlou G, Scheuermann B, Ascigil O (2025) Bootstrapping social networks: Lessons from Bluesky starter packs. Proc. Int. AAAI Conf. Web Soc. Media 19:178–192.

Ballester C, Calvó-Armengol A, Zenou Y (2006) Who's who in networks. Wanted: The key player. Econometrica 74(5):1403–1417.

Braghieri L, Bursztyn L, Fasnacht J (2026) Threshold disclosure in collective decisions. Working Paper, NBER Working Paper 34827. https://doi.org/10.3386/w34827.

Burtch G, Ghose A, Wattal S (2016) Secret admirers: An empirical examination of information hiding and contribution dynamics in online crowdfunding. Inf. Syst. Res. 27(3):478–496.

Chwe MSY (2001) Rational ritual: Culture, coordination, and common knowledge (Princeton University Press, Princeton, NJ).

Dybvig PH, Spatt CS (1983) Adoption externalities as public goods. J. Public Econ. 20(2):231–247.

Farrell J, Saloner G (1985) Standardization, compatibility, and innovation. RAND J. Econ. 16(1):70–83.

Farrell J, Saloner G (1986) Installed base and compatibility: Innovation, product preannouncements, and predation. Am. Econ. Rev. 76(5):940–955.

Goldenberg J, Han S, Lehmann DR, Hong JW (2009) The Role of Hubs in the Adoption Process. J. Mark. 73(2):1–13.

He J, Zia HB, Castro I, Raman A, Sastry N, Tyson G (2023) Flocking to Mastodon: Tracking the great Twitter migration. Proceedings of the 2023 ACM on Internet Measurement Conference (IMC '23) (Association for Computing Machinery), 111–123.

Hinz O, Skiera B, Barrot C, Becker JU (2011) Seeding strategies for viral marketing: An empirical comparison. J. Mark. 75(6):55–71.

Hu M, Shi M, Wu J (2013) Simultaneous vs. sequential group-buying mechanisms. Manag. Sci. 59(12):2805–2822.

Jing X, Xie J (2011) Group Buying: A New Mechanism for Selling Through Social Interactions. Manag. Sci. 57(8):1354–1372.

Katz ML, Shapiro C (1985) Network externalities, competition, and compatibility. Am. Econ. Rev. 75(3):424–440.

Landsman V, Stremersch S (2011) Multihoming in two-sided markets: An empirical inquiry in the video game console industry. J. Mark. 75(6):39–54.

Le Nagard-Assayag E, Manceau D (2001) Modeling the impact of product preannouncements in the context of indirect network externalities. Int. J. Res. Mark. 18(3):203–219.

Lin M, Xiao S, Yu J (2024) Showing the forest, not the trees: Backer list transparency and crowdfunding dynamics. SSRN working paper, 4935295, SSRN. https://doi.org/10.2139/ssrn.4935295.

Markus ML (1987) Toward a "Critical Mass" Theory of Interactive Media. Commun. Res. 14(5):491–511.

Palfrey TR, Rosenthal H (1984) Participation and the provision of discrete public goods: A strategic analysis. J. Public Econ. 24(2):171–193.

Quelle D, Denker F, Garg P, Bovet A (2025) Simple contagion drives population-scale platform migration. arXiv preprint, arXiv:2505.24801, arXiv. https://doi.org/10.48550/arXiv.2505.24801.

Tabarrok A (1998) The private provision of public goods via dominant assurance contracts. Public Choice 96(3–4):345–362.

Van den Bulte C, Joshi YV (2007) New Product Diffusion with Influentials and Imitators. Mark. Sci. 26(3):400–421.

Wu J, Shi M, Hu M (2015) Threshold Effects in Online Group Buying. Manag. Sci. 61(9):2025–2040.

Xie J, Shugan SM (2001) Electronic Tickets, Smart Cards, and Online Prepayments: When and How to Advance Sell. Mark. Sci. 20(3):219–243.

Zhong C, Salehi M, Shah S, Cobzarenco M, Sastry N, Cha M (2014) Social bootstrapping. Proc. 23rd Int. Conf. World Wide Web:305–314.

Online Appendix:

I'll Move If You Move: Assurance Campaigns for Platform Migration

This appendix specifies the finite sequential campaign, derives formation and voluntary destination-side execution in a transparent specialization, proves the assurance-lift result, and gives the complete staged-random-order count-versus-roster solution for the five-account witness. The comparison throughout is the destination provider's expected value from realized activity and ties at a fixed horizon. Pledges and threshold events have no direct value in that objective, and no social-welfare ordering is asserted.

A. A finite sequential campaign

A.1 Matched technologies and policies

The provider is the operator of one particular destination-side migration service or client. It can govern its service's eligibility and outreach, interface, account-control checks, pledge record, roster publication, notices, and operator-supplied migration tools. It cannot compel users or govern protocol components it does not operate, third-party clients or services, or users outside the managed wave.

Before policy assignment, the operator precommits to one treatment-invariant eligibility rule for a bounded community. The rule fixes a finite set E of eligible origin accounts and is auditable for exclusions, community overlap, linked-account matching, duplicate accounts or people, and network coverage. Let $K \in \{1, \dots, |E|\}$ be the campaign threshold. The three policies hold fixed destination-account creation, matching, lawful data import, starter packs, public activity, ordinary announcements, onboarding, and any partial reconstruction the operator can supply. Reconstruction is lawful, opt-in where required, and identical across policies. Thus the comparison neither weakens ordinary migration nor attributes third-party technology to the operator.

Ordinary migration (N). Accounts may create and announce destination accounts, use the common matching and onboarding tools, organize around public dates, and move in the ordinary way. Policy N has no authenticated threshold pledge record and no threshold-contingent common notice.

Count assurance (C). Policy C adds one common nonpersonalized pledge, threshold K, deadline D, a live authenticated count, release of the complete successful roster when the threshold is reached, and a common success notice. Acceptance requires the announced preparation bundle: creating or linking the destination account, completing account-control authentication, configuring the profile and lawful import, and staging the campaign-specific launch. The pledge records agreement to make the destination the default venue for the named context. Enrollment obtains informed consent to publish the authenticated account identity on the successful roster. The pledge is not exclusive and has no term contingent on a future named signer.

Roster assurance (R). Policy R is identical to C except that the identities in the authenticated pledge record are visible while pledges accrue. Enrollment under R therefore obtains informed consent to immediate live identity display after acceptance as well as successful-roster publication. The complete successful roster and common notice are the same under C and R; the treatment difference is the consented pre-success display of names.

Authentication establishes current control of one eligible linked account and prevents that account from occupying more than one record position. It does not establish a unique person, eliminate multiple accounts controlled by one person or matching error, or establish sincerity, readiness, or future performance; duplicate and matching-error rules therefore remain audit requirements. At first hit, C and R publish the same roster under successful-roster consent already obtained. Actual destination-side action remains voluntary under every policy: there is no bond, legal enforcement, penalty, or technological lock.

A.2 Opportunity times, histories, and messages

The formal environment is one bounded, managed, expiring, one-shot invitation wave. Nature draws a state θ and each account's private information from a commonly known joint distribution. Account i has one scheduled opportunity at $\tau_i \in [0, D)$ under an announced, policy-invariant invitation rule; opportunity-time ties have probability zero. An opportunity after first-hit closure is preempted. No account can defer, return after declining, revise an accepted pledge strategically, or condition a personal pledge on a signer who may arrive later.

The equations formalize the short-window append-only strategic benchmark; operational consent withdrawals and administrative decrements are stated below and in G.3, not added as strategic actions. Let $P_t \subseteq E$ be the right-continuous, post-event durable set of pledges accepted by time t, with $P_{0^-} = \emptyset$, let $c_t = |P_t|$, and let ℓ_t record whether the campaign is open, successful, or expired. At an opportunity $t = \tau_i$ reached while the campaign remains open, account i observes its private preparation cost κ_i, its other private information, and the left-limit public history. Later invitees retain the earlier public states, so while the campaign is open their messages are the complete public paths together with the current clock:

$$m_C(h_{t^-}) = \Big(t, \big((s, c_s, \ell_s)\big)_{0 \le s < t}\Big), \qquad m_R(h_{t^-}) = \Big(t, \big((s, c_s, P_s, \ell_s)\big)_{0 \le s < t}\Big). \tag{OA.1}$$

A decline is not publicly labeled. If i accepts, it incurs its preparation cost and $P_t = P_{t^-} \cup \{i\}$; under R, the identity enters the live display under consent obtained at acceptance. If it declines, $P_t = P_{t^-}$. The clock, invitation rule, prior, authentication standard, threshold, deadline, and all policy-independent public events are common knowledge. So is the operator's commitment not to fabricate counts, suppress acceptances, alter authenticated identities, selectively reveal names, or change eligibility or invitation timing after observing the state. An audit log records eligibility, authentication, acceptances, withdrawals and removals, displayed states, first hit, publication, and notice attempts; strategic operator messaging is outside the game.

The success time is

$$T = \inf\{t < D : c_t = K\}, \tag{OA.2}$$

where $\inf \emptyset = +\infty$. Because opportunities do not tie, a finite T produces $S = P_T$ with exactly K accounts. The theorem takes this exact first-hit rule, with no tied opportunities, as a primitive. An implementation with nearly simultaneous events must announce its clock or batching tolerance and deterministic within-batch order before launch; the theorem does not replace S with an oversubscribed batch. At T, the campaign closes, publishes S under consent already obtained, and sends the common notice rather than waiting until D. Notice attempts, delivery, exposure, and timing must be logged; missed or staggered notices enter Q rather than being presumed received. If $T = +\infty$, the campaign expires at D. Destination action remains voluntary whether the campaign succeeds or expires.

Operationally, a signer may withdraw consent required by the assigned display policy before T. The operator then deletes the pledge, removes an R identity from prospective display, excludes the account from the eventual successful roster, and lowers the count; it cannot erase copies of an identity already shown under R. Announced removals for fraud, lost authentication, law, harassment, or safety have the same record effect. These rules do not contradict the formal state: they are operational exceptions suppressed in the append-only benchmark and are not options in the monotone five-account witness.

The timing defines a finite-horizon extensive-form Bayesian game with a mechanically committed operator. A behavioral pledge strategy maps an account's private information and observed message into acceptance or decline. The solution concept is perfect Bayesian equilibrium: actions are sequentially optimal at every information set, beliefs follow Bayes' rule at every reached information set, and off-path beliefs are completed wherever needed. The disclosure comparison therefore changes information at a decision node, not the underlying eligible set, invitation technology, or state distribution.

A.3 Voluntary execution and provider value

For accounts i and j, let $\omega_{ij} \geq 0$ be the directed value that i places on j's destination activity after feasible tie reconstruction. A common thickness term may also enter participant utility. The provider instead places nonnegative weights on hosted activity and active destination ties. Its weights need not equal the participants' private affinities: an account can be persuasive to one prospective signer without being the provider's most valuable account, and vice versa.

For a successful roster S, let $Q(X \mid S)$ be the joint probability that exactly the subset $X \subseteq S$ voluntarily undertakes the destination-side launch action after the common notice, after integrating over the signer-type posterior induced by campaign participation. The law $Q(\cdot \mid S)$ may contain idiosyncratic failure, common shocks, clustered nonperformance, and arbitrary correlation across participants. As a maintained exclusion restriction, it is the same under C and R conditional on the same successful roster. This rules out an additional direct effect of the earlier display history or trigger time on post-notice execution; Section G states the boundary.

For each $d \in \{N, C, R\}$, let $\mathcal{W}_d$ be the random gross provider value at the fixed evaluation horizon from realized destination activity and ties on the full path. For a successful assurance roster S and executing subset X, let $V(X,S)$ be the corresponding conditional expected value under the common post-notice technology. Then, for $S \neq \emptyset$ and either $d \in \{C, R\}$,

$$W(S) = \mathbb{E}_d[\mathcal{W}_d \mid \widehat{S}_d = S] = \sum_{X \subseteq S} Q(X \mid S) V(X,S), \tag{OA.3}$$

which is common to C and R under the maintained exclusions. Any destination activity by accounts outside S that is common across policies can be included in V. A participant's continuation value is formed analogously by integrating its payoff over voluntary execution; paying a preparation cost does not remove its later choice.

Let $\widehat{S}_d$ equal the successful roster under assurance policy $d \in \{C, R\}$ and equal $\emptyset$ after expiry. The empty symbol states only that the campaign expired without a provider-coordinated success notice; $\mathcal{W}_d$ remains defined and may vary across full expired paths. With incremental administrative cost $A_d \geq 0$ and incremental campaign-induced delay $D_d \geq 0$, both measured relative to N in provider-value units, the policy values are

$$J_d = \mathbb{E}_d[\mathcal{W}_d] - A_d - D_d, \quad d \in \{C, R\}, \qquad J_N = \mathbb{E}_N[\mathcal{W}_N]. \tag{OA.4}$$

Effects already included in $\mathcal{W}_d$ are not counted again in A_d or D_d. Pledges and success affect the objective only through the realized path and incremental costs; direct publicity, symbolic, or data-acquisition objectives are excluded rather than declared valueless. All policies use the same horizon and common technology. The criterion ranks expected narrow provider value in commensurable units, not risk, hard safety constraints, participant payoffs, or social welfare; empirical use requires validated endpoints and prespecified horizon checks.

B. Assurance relative to ordinary migration

B.1 A common activity path and its basin boundary

The following uncalibrated binary-basin benchmark uses a common activity response rather than different technologies for N and assurance. Let y_t be destination activity in account-equivalent units at public revision date t, and let

$$y_{t+1} = G(y_t). \tag{OA.5}$$

Assume G is continuous and strictly increasing on the relevant interval and has exactly three fixed points $y^- < y^u < y^+$. Assume further that

$$G(y) - y \begin{cases} > 0, & y < y^-, \\ < 0, & y^- < y < y^u, \\ > 0, & y^u < y < y^+, \\ < 0, & y > y^+. \end{cases} \tag{OA.6}$$

Because G is increasing and fixes y^u, every state below y^u maps to another state below y^u; the analogous statement holds above y^u. If $y_0 \in (y^-, y^u)$, equation (OA.6) makes the sequence strictly decreasing, while monotonicity of G bounds it below by y^-. Its limit exists and, by continuity, must be the only fixed point in the closed interval $[y^-, y_0]$, namely y^-. The cases $y_0 < y^-$ and $y_0 > y^u$ follow symmetrically. Thus an ordinary path beginning below y^u converges asymptotically to y^-, while a realized campaign launch that changes the post-notice state from y^- to a value strictly above y^u converges asymptotically to y^+. Both paths then use the same map G. Convergence identifies eventual basin membership; it does not imply equal activity or value for paths in the same basin at the fixed horizon H. Those transient differences remain in $\mathcal{W}_d$.

Let M_d denote the account-equivalent contribution of the subset that actually launches under policy d. The high-basin event after campaign success is $y^- + M_d > y^u$. This need not coincide with reaching the pledge threshold: K counts recorded agreements, whereas M_d is generated by voluntary destination activity. Network position can be represented by heterogeneous contribution weights inside M_d, so that the same number of launchers need not have the same crossing effect. The crossing event is strict. Proposition 1 assumes no probability mass at $M_d = y^u - y^-$; if a discrete illustration has such an atom, equality is assigned to the noncrossing low class and its probability is included there.

B.2 Assurance-lift decomposition

Proposition 1 (behavioral assurance lift). Under the additional binary-value approximation, set $\mathcal{W}_d = W_L$ for every noncrossing outcome, including expiry and successful but subcritical execution, and $\mathcal{W}_d = W_H > W_L$ for every crossing outcome. These equal values are imposed classifications, not exact finite-H activity levels implied by convergence. Suppose ordinary migration is classified low, under assurance policy d the high class can be reached only after campaign success, and the launch contribution has no probability mass exactly at $y^u - y^-$. Let

$$\alpha_d = \Pr_d(\widehat{S}_d \neq \emptyset), \qquad q_d = \Pr_d(\text{high basin} \mid \widehat{S}_d \neq \emptyset). \tag{OA.7}$$

Set $q_d = 0$ when $\alpha_d = 0$. Then

$$J_d - J_N = \alpha_d q_d (W_H - W_L) - A_d - D_d. \tag{OA.8}$$

Consequently, assurance strictly raises, leaves unchanged, or strictly lowers provider value according as $\alpha_d q_d (W_H - W_L)$ is greater than, equal to, or less than $A_d + D_d$.

Proof. Under the stated binary-value approximation, $\mathcal{W}_d = W_H$ with probability $\Pr_d(\text{high class}) = \alpha_d q_d$ and $\mathcal{W}_d = W_L$ otherwise. The latter assignment includes expiry and subcritical execution; it is not inferred from the absence of a success notice or from asymptotic convergence. Hence expected gross value is

$$\alpha_d q_d W_H + (1 - \alpha_d q_d) W_L = W_L + \alpha_d q_d (W_H - W_L). \tag{OA.9}$$

Ordinary migration has $\mathcal{W}_N = W_L$ in this specialization, so $J_N = W_L$. Subtracting incremental administration and delay gives (OA.8); comparing its gross term with its incremental costs gives the three regions. $\square$

Outside the binary-value approximation, different expired, subcritical, or same-basin transient values are compared through $\mathbb{E}_d[\mathcal{W}_d] - A_d - D_d - J_N$; asymptotic convergence alone does not reduce that expectation to two values. Within the binary approximation, the fair-counterfactual boundary does not require normalizing ordinary migration's crossing probability to zero. If b_d and b_N denote the high-classification probabilities under assurance and ordinary migration, respectively, the general binary-basin identity is

$$J_d - J_N = (b_d - b_N)(W_H - W_L) - A_d - D_d, \tag{OA.10}$$

where A_d and D_d are incremental to N and exclude effects already counted in $\mathcal{W}_d$. Equation (OA.8) is the case $b_N = 0$ and $b_d = \alpha_d q_d$.

Three boundaries follow directly. First, if ordinary organizing creates the same contingent record, common notice, and launch distribution as assurance, then $b_d = b_N$ and lift is zero before incremental cost; positive cost makes net lift negative. Second, if ordinary activity already lies in the high basin and assurance leaves it there, $b_d = b_N = 1$, so there is no crossing gain. Third, when assurance increases crossing, net lift is zero if the gross gain equals administration and incremental delay and is negative if those costs exceed it. This is an expected narrow provider-value accounting result under a disclosed approximation, not evidence that the basin specification is realistic and not a ranking of risk, safety, or welfare.

B.3 Formation and execution from individual choices

The probabilities in Proposition 1 need not be free parameters. Consider a transparent, uncalibrated all-or-none recruitment specialization with exactly K eligible accounts and threshold K, so every account is required. This independent-cost illustration is not the formation model for the five-account witness. A successful prepared campaign gives account i a fixed expected optimized continuation surplus $b_i > 0$ relative to its ordinary path, before preparation cost; in this benchmark, that surplus does not vary with cosigner identities. Acceptance requires an announced destination-account bundle: creating or linking the account, completing account-control authentication, configuring the profile and lawful import, and staging the campaign-specific launch. Its incremental time and setup cost κ_i is incurred immediately and is not refunded if the campaign expires; any reusable account value remains in ordinary continuation. The bundle's prevalence and cost are pilot estimands. Let κ_i have continuous cdf F_i and be independent across accounts. Normalize the incremental campaign value of expiry or declining to zero.

For a policy-invariant announced invitation order $\sigma = (\sigma_1, \dots, \sigma_K)$, define $\alpha_{K+1}(\sigma) = 1$ and work backward:

$$\alpha_t(\sigma) = F_{\sigma_t}\Big(b_{\sigma_t}\alpha_{t+1}(\sigma)\Big)\,\alpha_{t+1}(\sigma), \qquad t = K, \dots, 1. \tag{OA.11}$$

Conditional on all earlier accounts accepting, $\alpha_t(\sigma)$ is the probability that the campaign eventually succeeds. Account σ_t receives b_{σ_t} only if all later accounts accept, an event with probability $\alpha_{t+1}(\sigma)$, so it accepts exactly when

$$\kappa_{\sigma_t} \leq b_{\sigma_t} \alpha_{t+1}(\sigma). \tag{OA.12}$$

Independence then gives (OA.11). Correlated preparation costs require the relevant joint conditional calculation, while cost atoms require an explicit action at equality. Under the stated benchmark, campaign formation is

$$\alpha(\sigma) = \alpha_1(\sigma). \tag{OA.13}$$

If the provider randomizes over announced orders with policy-invariant distribution ν, then $\alpha = \sum_\sigma \nu(\sigma)\alpha_1(\sigma)$. In the homogeneous case, order is irrelevant and the same scalar recursion applies without announcing individual positions. With more than K eligible accounts, only the one-shot finite game is covered: backward induction tracks the accepted count, remaining invitation nodes, and clock. No scalar product recursion is claimed, and this extension excludes repeated opportunities and endogenous arrivals.

The fixed successful-campaign surplus can itself come from a separate, stylized, unenforced launch option. After the common notice, selected signer i observes a transitory launch cost ε_i and receives gross destination value v_i if it launches, with nonlaunch payoff normalized to zero. This benchmark assumes ε_i is independent of preparation cost κ_i; equivalently, the displayed formula may be read conditionally when option value is invariant across selected preparation-cost types. The signer voluntarily launches exactly when $\varepsilon_i \leq v_i$, so

$$\varphi_i = F_{\varepsilon_i}(v_i), \qquad b_i = \mathbb{E}[(v_i - \varepsilon_i)_+], \tag{OA.14}$$

where F_{ε_i} is the launch-cost cdf. Thus φ_i is a probability among selected successful signers, not a population-user rate. If preparation and launch costs are related, selection into signing changes the relevant execution distribution, which must instead be measured among signers and represented in Q. The same value in the successful state supports the pledge cutoff and later voluntary action without a bond or a penalty for reneging and without assuming that signers execute. The migration evidence cited in the main text supports the distinction between registration and later action only indirectly, through multihoming and heterogeneous destination activity; it does not estimate this launch process or pledge follow-through. In the homogeneous independent case, substituting the binomial crossing probability below and $\alpha = \alpha_1$ into (OA.8) gives the assurance gain from the benchmark primitives. Common or clustered launch shocks instead enter through the induced q_d under Q.

This specialization also clarifies the ordinary benchmark. The focal residual accounts prefer the incumbent to the destination's low state, and the common map keeps ordinary activity below y^u. Assurance adds the prepared contingent record and first-hit notice that can create a launch batch; it does not change the post-launch map. If ordinary organizers can reproduce the same record, notice, preparation, and launch distribution, equation (OA.10) assigns assurance no gross advantage.

B.4 Independent-execution benchmarks

The joint law Q is the maintained model. Independence is useful only as a transparent specialization. Suppose a successful campaign has exactly K selected signers, each launches independently with common conditional probability φ_d, and each launcher contributes one account-equivalent unit. Define

$$m = \min\{z \in \mathbb{Z}_+ : y^- + z > y^u\}. \tag{OA.15}$$

Then

$$q_d = \begin{cases} \sum_{j=m}^{K} \binom{K}{j} \varphi_d^j (1-\varphi_d)^{K-j}, & m \le K, \\ 0, & m > K. \end{cases} \tag{OA.16}$$

For these independent bounded Bernoulli launches, if $K\varphi_d > m$, Hoeffding's inequality gives the conservative bound

$$q_d \ge 1 - \exp\left[-\frac{2(K\varphi_d - m)^2}{K}\right]. \tag{OA.17}$$

For a fixed successful roster S, let X_i be mutually independent Bernoulli launches with fixed nonnegative bounded contributions z_i. If $\delta > 0$ and $\sum_{i\in S} z_i \mathbb{E}[X_i] \ge y^u - y^- + \delta$, the corresponding bound is

$$\Pr\left(\sum_{i\in S} z_i X_i > y^u - y^- \,\middle|\, S\right) \ge 1 - \exp\left[-\frac{2\delta^2}{\sum_{i\in S} z_i^2}\right]. \tag{OA.18}$$

These expressions illustrate the distinction between the pledge threshold and realized crossing. The binomial formula and both Hoeffding bounds are independent-launch benchmarks; none applies to the common, clustered, or other correlated execution allowed by Q.

B.5 Finite-network crossing

The scalar boundary has a finite-network counterpart as an analyst-side positive-spillover specialization. Let A_N be the ordinary active set, let $\Phi_\theta^H(A)$ be the active set at the common horizon when the common post-launch activity process starts from A, and fix a target active set $\bar{A} \supsetneq A_N$. For the realized launch subset X_d, define

$$E_d(\bar{A}) = \{\bar{A} \subseteq \Phi_\theta^H(A_N \cup X_d)\}. \tag{OA.19}$$

When the horizon map is monotone in its initial active set, the minimal target-reaching launch sets are

$$\mathcal{L}_\theta(A_N, \bar{A}) = \left\{L \subseteq E : L \cap A_N = \emptyset,\ \bar{A} \subseteq \Phi_\theta^H(A_N \cup L) \text{ and no } L' \subsetneq L \text{ has this property}\right\}. \tag{OA.20}$$

Because E is finite and the horizon map is monotone, the target is reached exactly when X_d contains at least one member of this family. The family is an analyst's description, not a claim that the provider observes it or every private affinity. Signed spillovers, congestion, nonlinear activation, safety losses, or moderation effects can break monotonicity; those environments require branch-specific expected values W and, for Proposition 2, a reverified g_z rather than this containment characterization.

The one-shot role of screened topology can be made exact. Here a scored node or tie is desired and policy-compliant, feasible to host or reconstruct, and screened for nonnegative marginal provider contribution. Lawful authorization or reconstructability makes it eligible for screening, not beneficial. Let $G^R = (E, L^R)$ be the resulting scored destination graph, let A_N be the existing active set, and define the one-shot contribution of mover set X by

$$Z_{G^R}(X) = \sum_{i \in X} z_i + \gamma \sum_{\{i,j\} \in L^R} \mathbf{1}\{i, j \in A_N \cup X\}, \qquad z_i, \gamma \geq 0. \tag{OA.21}$$

For a threshold $\zeta > 0$, take the finite-network specialization in which the target basin is reached exactly when $Z_{G^R}(X) \geq \zeta$, and call such an X sufficient. This one-shot linear-threshold score is exact inside the specialization, not a claim that general platform dynamics reduce to one index.

Four exact, uncalibrated constructions cover the homogeneous, hub, triangle, dispersed, bridge, and peripheral cases. First, in the homogeneous benchmark $\gamma = 0$ and $z_i = 1$, so only mover count matters and names are physically irrelevant. Second, set $A_N = \emptyset$, let h be a hub linked to leaves 1 and 2 with no other scored edges incident to the compared accounts, set $z_i = 0$, $\gamma = 1$, and $\zeta = 2$. The three-mover set $\{h, 1, 2\}$ has score two and is sufficient, while a three-mover set of accounts with no edges among or incident to them has score zero. Third, under the same empty-baseline convention, a reconstructed triangle of three local contacts has score three at $\zeta = 3$, whereas three accounts with no scored incident edges again fail. Fourth, let A_N contain active accounts x and y in different local groups and let the only scored edges in this comparison be those stated. A bridge b with reconstructed edges $\{b, x\}$ and $\{b, y\}$ is sufficient by itself at $\zeta = 2$; a peripheral account with only one reconstructed edge is not. These are proof witnesses, not calibrated Bluesky network effects.

Finally, let $L_0^R \subseteq L_1^R$ represent two screened feasible reconstruction levels with the same nodes, ordinary active set, node weights, edge weight, and threshold. Write $\mathcal{S}(L^R) = \{X \subseteq E : X \cap A_N = \emptyset,\ Z_{(E,L^R)}(X) \geq \zeta\}$ for the full sufficient-set family. Nonnegative weights imply

$$Z_{(E,L_0^R)}(X) \leq Z_{(E,L_1^R)}(X) \quad \text{for every } X. \tag{OA.22}$$

Hence $\mathcal{S}(L_0^R) \subseteq \mathcal{S}(L_1^R)$, and the minimum sufficient cardinality, defined as $+\infty$ when the family is empty, weakly falls. In the bridge witness, one scored bridge edge leaves $\{b\}$ insufficient, while two make it sufficient. This conclusion applies only to added edges that pass the nonnegative-contribution screen; legality or technical feasibility alone gives no ordering. Operational active-node and active-tie definitions, bot and fraud handling, missing protected ties, and score validation are rollout requirements; the score and network-position effects are estimands. Private affinity ω_{ij} remains distinct from the score and may be only imperfectly proxied.

C. Exact five-account existence witness

C.1 State and policy-invariant random order

Set $K = 3$ and $E = \{a, b, p, s, r\}$. This is an uncalibrated existence construction with a binary anchor state. Accounts a and b are possible anchors whose authenticated identities p is assumed to recognize under R; recognition is a pilot estimand, and noisy recognition requires a partial-information extension. Account p is pivotal, s is a reliable signer, and r is a reserve signer. Exactly one anchor is ready in the campaign window. Write $Z \in \{a, b\}$ for that identity and let

$$\Pr(Z = a) = \pi \in (0, 1), \qquad \Pr(Z = b) = 1 - \pi. \tag{OA.23}$$

The two anchors and p receive their unique opportunities in an initial block according to a timing law that is independent of Z, symmetric under exchanging the two anchor labels, and common across policies. Invitation identities and declines are not public under C. The unready anchor strictly declines and the ready anchor strictly accepts. Let

$$\beta = \Pr(\tau_Z < \tau_p) \in (0, 1) \tag{OA.24}$$

be the probability that the ready anchor precedes p. Under a uniform random permutation of a, b, and p, all six orders are integrated and $\beta = 1/2$. If p precedes Z, it observes an empty pledge record and the same pooled anchor prior under either policy. If Z precedes p, the history contains one accepted pledge. Anchor-label symmetry makes the complete count-and-clock likelihood of either pooled event the same in both anchor states; independence from Z alone would not be sufficient.

After the early block, account s receives its opportunity and strictly accepts. If p accepted, s supplies the third pledge, the first-hit rule closes the campaign, and the successful roster is $\{Z, p, s\}$. If p declined, s is only the second signer, so the campaign remains open. Account r then receives the next opportunity and is ready with probability $\rho \in (0, 1)$, independently of Z and the early order in the baseline witness. A ready r strictly accepts and completes $\{Z, s, r\}$; a nonready r strictly declines, and the campaign expires with two recorded pledges. Whether r is reached is determined only by ordinary first-hit closure, not by an adaptive invitation or wait-list rule.

The strict actions outside p can be supported without enforcement. For example, give each ready anchor and s zero preparation cost and a strictly positive continuation surplus from campaign success, give a ready r zero preparation cost and a strictly positive surplus from the success it completes, and give an unready account a strictly negative payoff from preparing. Since $\rho > 0$, the ready anchor and s strictly prefer acceptance; a ready r strictly prefers to complete the cohort. These policy-common auxiliary payoffs support strict equilibrium actions only; their magnitudes are not forecasts and do not enter the provider ranking.

C.2 Sequential information and the pivotal pledge

Conditional on the ready anchor preceding p, the observed count-and-clock path has the same likelihood in both anchor states. Under C, p observes one authenticated pledge but no identity, so Bayes' rule gives

$$\Pr(Z = a \mid m_C) = \pi. \tag{OA.25}$$

Under R, the live roster is $\{Z\}$ and the posterior is degenerate on the displayed identity. Thus the count pools the two anchor states at the pivotal history, while the roster separates them. On the complementary early-block histories, p precedes Z and observes an empty record under both policies. Anchor-label symmetry again leaves posterior π, so the decision is pooled and identical across policies.

Measure p's payoffs relative to a common continuation baseline that may include continued X use. If p accepts and the campaign succeeds with $\{z, p, s\}$, normalize refusal of destination launch to zero and let $L_z > 0$ be the incremental destination-launch value from acting voluntarily. If p declines the pledge, let $O_z \geq 0$ be its optimized ordinary continuation payoff. This outside option integrates over r's readiness, campaign success or expiry, later roster publication on success, and any ordinary post-decline migration by p before the horizon. It is common across C and R conditional on z, parallel to policy-common $Q(\cdot \mid S)$ conditional on a successful roster S; these exclusions isolate live disclosure's effect on recruitment rather than direct effects on execution or ordinary post-decline behavior. Define

$$u_z = L_z - O_z, \qquad u_a > u_b > 0. \tag{OA.26}$$

Acceptance incurs the fixed cost $\kappa \geq 0$ immediately. If p accepts after Z, s completes the campaign; if p accepts before Z, the ready anchor and then s complete it. Thus acceptance guarantees campaign success in either early order, while declining preserves O_z and leaves success to the reserve branch. Define

$$\bar{u} = \pi u_a + (1 - \pi) u_b. \tag{OA.27}$$

With weak acceptance at indifference, sequential rationality gives

$$\begin{aligned} p \text{ accepts under } C &\iff \kappa \leq \bar{u}, \\ p \text{ accepts under } R \text{ after anchor } z &\iff \kappa \leq u_z, \\ p \text{ accepts before the ready anchor under either policy} &\iff \kappa \leq \bar{u}. \end{aligned} \tag{OA.28}$$

Under the policy-common continuation exclusions, the payoff difference between acceptance and decline at a pooled history is $\bar{u} - \kappa$, whereas under R after z it is $u_z - \kappa$. Together with the strict auxiliary actions in C.1, these rules and the stated Bayesian beliefs form a perfect Bayesian equilibrium under each policy; only p's action at an exact cutoff uses the weak-acceptance tie rule. The exact Bayesian cutoffs are equilibrium implications, not behavioral forecasts. The maintained $L_b > 0$ is an uncalibrated condition and pilot estimand; after success, sunk preparation then leaves destination action voluntarily optimal in either state. The model therefore preserves both of p's real options: it can decline and use ordinary migration, or it can prepare now and still decline to launch later if launch utility were negative.

A minimal destination-network microfoundation uses the effective affinity weights from A.3. Let

$$L_z = \ell + \omega_{pz} + \omega_{ps}, \qquad O_z = (1-\rho)o_L + \rho\left[o_H + \delta(\omega_{pz} + \omega_{ps} + \omega_{pr})\right], \quad 0 \le \delta < 1. \tag{OA.29}$$

The parameter δ captures the fraction of relevant tie value that p can recover through a later ordinary response after declining. If $\omega_{pa} > \omega_{pb}$, then

$$u_a - u_b = (1 - \rho\delta)(\omega_{pa} - \omega_{pb}) > 0. \tag{OA.30}$$

Thus the high anchor can be a prominent account valued by many users or one member of p's own local group; the ordering comes from participant-specific destination ties rather than raw cohort size.

D. Sequential disclosure reversal

For $z \in \{a, b\}$, let W_z^P be expected provider value when p accepts and the successful roster is $\{z, p, s\}$. Let B_z be expected provider value when p declines. The latter integrates the probability ρ of successful reserve roster $\{z, s, r\}$, expiry when r is not ready, and every ordinary post-decline action by p. Conditional on any successful roster S, both policies use the same $Q(\cdot \mid S)$, horizon activity process, and provider weights; O_z is likewise common across C and R conditional on z. In addition, conditional on the complete realized branch, including z, the pivotal accept/decline decision, the reserve outcome, success or expiry, and subsequent activity, the earlier display has no direct effect on gross provider value. This full-path exclusion makes W_z^P and B_z policy-common, including on the expiry branch. Define

$$g_z = W_z^P - B_z > 0. \tag{OA.31}$$

For example, if expiry returns to W_L and the successful reserve branch is worth W_z^R, then $B_z = (1-\rho)W_L + \rho W_z^R$. The value W_z^R may include ordinary activity by p even though p is not in the campaign roster. Assume that campaign administration, delay, naming, and interface costs are the same under C and R, so they cancel in the disclosure comparison.

One concrete positive-spillover provider specification is

$$V_H = \sum_i \lambda_i x_{iH} + \sum_{i \ne j} \eta_{ij} \ell_{ij,H}, \qquad \lambda_i, \eta_{ij} \ge 0, \tag{OA.32}$$

where x_{iH} and $\ell_{ij,H}$ denote proposed horizon-activity and active-tie measures that require rollout validation. Taking expectations under each branch gives W_z^P and B_z. The witness assumes $g_z > 0$, a premise to be tested by estimating the realized values of those branches; it is not implied by $u_z > 0$, pledges, participant affinity, or centrality. With signed or nonlinear value, congestion, safety losses, or moderation costs, branch-specific W_z^P and B_z must be evaluated and $g_z > 0$ reverified.

The participant-affinity microfoundation and B.5 bridge case supply an exact uncalibrated existence construction, not evidence for the branch gains. Put two incumbent active groups in A_N, give p one scored edge to each group, and give reserve r only one. At $\zeta = 2$, a launching p scores two while a launching r scores one, so branch values can be specified with $g_a, g_b > 0$. Independently setting $\omega_{pa} > \omega_{pb}$ gives $u_a > u_b$ without equating private and provider value. The construction supports both disclosure regions without assigning pledge totals direct value; whether p is such a bridge and whether the gains are positive are rollout estimands.

Proposition 2 (random-order disclosure reversal in recruitment). In the five-account witness, impose common execution and outside options across policies, require provider value to be common conditional on each complete realized branch, and assume equal campaign and display costs across policies. Then $u_b < \bar{u} < u_a$, and after integrating the policy-invariant early-block order, two nonempty open regions have opposite strict rankings:

$$u_b < \kappa < \bar{u} \quad \Longrightarrow \quad J_C - J_R = \beta(1-\pi)g_b > 0, \tag{OA.33}$$
$$\bar{u} < \kappa < u_a \quad \Longrightarrow \quad J_R - J_C = \beta\pi g_a > 0. \tag{OA.34}$$

Although R gives p better information, C can preserve recruitment in the low-anchor state by pooling identities and inducing preparation that p would reject if b were known; R can instead recruit p in the high-anchor state when the pooled count is not persuasive enough. These are recruitment/cohort-formation rankings under the stated common participant, execution, provider-branch, and cost restrictions, not total-product or welfare rankings.

Proof. Because $0 < \pi < 1$ and $u_a > u_b$, their strict convex combination satisfies $u_b < \bar{u} < u_a$. Let J^0 be provider value when p declines in both states:

$$J^0 = \pi B_a + (1-\pi)B_b - A - D^{\text{delay}}, \tag{OA.35}$$

where A and D^{delay} are the common campaign terms. The pledge rules in (OA.28) imply

$$J_C = J^0 + \mathbf{1}\{\kappa \le \bar{u}\}\left[\pi g_a + (1-\pi)g_b\right], \tag{OA.36}$$
$$J_R = J^0 + (1-\beta)\mathbf{1}\{\kappa \le \bar{u}\}\left[\pi g_a + (1-\pi)g_b\right] \tag{OA.37}$$
$$+ \beta\left[\pi g_a \mathbf{1}\{\kappa \le u_a\} + (1-\pi)g_b\mathbf{1}\{\kappa \le u_b\}\right]. \tag{OA.38}$$

The corresponding campaign-success probabilities are

$$\alpha_C = \rho + (1-\rho)\mathbf{1}\{\kappa \le \bar{u}\}, \tag{OA.39}$$
$$\alpha_R = \rho + (1-\rho)\left\{(1-\beta)\mathbf{1}\{\kappa \le \bar{u}\} + \beta\left[\pi\mathbf{1}\{\kappa \le u_a\} + (1-\pi)\mathbf{1}\{\kappa \le u_b\}\right]\right\}. \tag{OA.40}$$

Thus disclosure changes both the probability of first-hit success and the identity composition of successful rosters on the random-order histories in which Z precedes p; g_z values the complete branch difference. In the first open region, R's better information makes p reject after b, whereas C induces that preparation at every pooled history; p accepts before Z under either policy, yielding the gain in (OA.33). In the second, p declines at every pooled history but accepts under R after a, yielding (OA.34). When p precedes Z, its information and action are identical under the two policies, so those provider increments cancel. Multiplying the conditional gaps by the common reach probability β integrates every early-block order. All other account strategies and the reserve readiness probability are common across policies. □

Under the stated weak-acceptance convention, equations (OA.36)–(OA.38) give the complete piecewise ranking for every $\kappa \geq 0$:

$$J_C - J_R = \begin{cases} 0, & 0 \leq \kappa \leq u_b, \\ \beta(1-\pi)g_b, & u_b < \kappa \leq \bar{u}, \\ -\beta\pi g_a, & \bar{u} < \kappa \leq u_a, \\ 0, & \kappa > u_a. \end{cases} \tag{OA.41}$$

The values at $\kappa = \bar{u}$ and $\kappa = u_a$ use acceptance at indifference; at $\kappa = u_b$, both policies recruit p in the low-anchor state. Reversing the tie rule can change these isolated boundary points but cannot change either strict open region. Proposition 2 isolates recruitment under equal costs and policy-common O_z, Q, and full-path provider branch values. A total product comparison must add display-specific effects on trust, privacy, consent, moderation, exposure, execution, ordinary outside behavior, and costs.

The count-favoring region uses strategic opacity. Roster policy R gives p better information: in the low-anchor state, $u_b < \kappa$ makes p reject preparation if b is known, whereas C's pooled value $\bar{u} > \kappa$ induces preparation before the identity is learned. The outside option has not been removed: it is built into both $u_b = L_b - O_b$ and provider branch B_b. The result establishes a difference in provider value through the recruitment channel; it does not rank total product value or prove that opacity benefits p or maximizes aggregate surplus.

E. Count sufficiency

The reversal requires identities to matter somewhere in the complete continuation environment. For the strong count-sufficiency boundary, suppose that, within each relevant observable account class and at every pair of histories with the same complete count-and-clock path, identity permutations leave unchanged (i) the joint continuation law of future private types, readiness, remaining opportunity times, voluntary execution, and administrative transitions; (ii) each remaining invitee's payoff from every available action; (iii) the public transition law; and (iv) gross full-path provider value. Accepted-set effects in (ii)–(iv) are cardinality-only. Suppose also that C and R have equal administration and delay costs and use the same permutation-invariant tie rule or equilibrium selection. Then the named roster is an action-irrelevant refinement of the count and

$$J_C = J_R. \tag{OA.42}$$

Proof. Work backward through the finite remaining opportunity nodes. At a terminal node, (i), (ii), and (iv) make participant continuation payoffs and gross provider value identical across named refinements of the same count-and-clock history. Suppose continuation values agree after the current node. Assumptions (i) and (ii) give the current invitee the same distribution of payoff-relevant states and the same payoff from every action across named refinements; the common selection rule therefore chooses the same action. Assumption (iii) gives the same distribution of next count-and-clock nodes, where the induction hypothesis applies. Thus C and R induce the same terminal full-path distribution up to identity permutation and the same expected gross provider value. Equal costs complete (OA.42). □

These conditions are stronger than cardinality-only provider value, which is not sufficient by itself. Names can still change payoffs, actions, opportunities, execution, or transitions and thereby change success and full-path value even when same-sized terminal cohorts are valued alike. Conversely, equal recruitment probabilities do not imply equivalence if names affect provider value conditional on roster size. The equality is a strong counterfactual boundary, not a description of identity-rich communities.

F. Voluntary execution and exact numbers

The pivotal account makes two distinct decisions. At the pledge opportunity, it compares the incremental value $u_z = L_z - O_z$ with the preparation cost. Payoffs are measured relative to a common continuation baseline that may include continued X use, and L_z is incremental destination-launch value. After success, κ is sunk and p freely chooses whether to undertake destination-side action. The maintained, uncalibrated condition $L_b > 0$ ensures launch is optimal even in the low-anchor successful cohort and is a pilot estimand. A sufficient condition for positive total payoff from every successful prepared launch, after counting the sunk cost, is

$$\kappa < \min\{L_a, L_b\} \quad \Longrightarrow \quad L_z - \kappa > 0 \quad \text{for both } z \in \{a, b\}. \tag{OA.43}$$

This condition does not say that count disclosure is participant-optimal. In the count-favoring low-anchor state, R gives p better information and $u_b < \kappa$ means it would reject preparation if b were known; C pools the identity and induces that preparation. Even so, $L_b - \kappa$ is positive and launch remains voluntary after the sunk preparation cost.

For the transparent numerical witness, set

$$\pi = \rho = \tfrac{1}{2}, \qquad \beta = \tfrac{1}{2}, \qquad L_a = 7, \qquad L_b = 5, \qquad O_a = O_b = 3, \qquad g_a = g_b = 1. \tag{OA.44}$$

Then

$$u_b = 2, \qquad \bar{u} = 3, \qquad u_a = 4. \tag{OA.45}$$

At $\kappa = \frac{5}{2}$, conditional on the ready anchor preceding p, R's better information makes p reject preparation after b, whereas C pools the identity and induces preparation in both states. Averaging the uniform early-block order gives

$$J_C - J_R = \beta(1 - \pi) g_b = \tfrac{1}{4}. \tag{OA.46}$$

At $\kappa = \frac{7}{2}$, C does not recruit p and R recruits it only after a, so

$$J_R - J_C = \beta\pi g_a = \tfrac{1}{4}. \tag{OA.47}$$

Both displayed costs satisfy $\kappa < \min\{L_a, L_b\} = L_b = 5$. Even at $\kappa = \frac{7}{2}$, the successful-launch payoff is positive in both states because its lower value is $L_b - \kappa = \frac{3}{2} > 0$. These exact values from the uncalibrated existence witness illustrate the two open regions; they are neither estimates nor behavioral forecasts.

G. Robustness and limits

G.1 Imperfect and correlated execution

No assumption of independent execution is needed. Idiosyncratic, common, and clustered failures are already integrated into $Q(X \mid S)$ and hence into $W(S)$, g_a, g_b, and the basin-crossing probability q_d. Proposition 1 remains valid for any such joint law because it uses only the induced high-basin probability. Proposition 2 remains valid when $Q(\cdot \mid S)$ and full-path branch values are common across C and R and the induced increments satisfy $g_a, g_b > 0$. If disclosure directly changes post-success execution or branch value after conditioning on the realized branch, that effect is a separate treatment channel and equations (OA.36)–(OA.38) no longer isolate recruitment.

G.2 Nonuniform state, invitation, and completion probabilities

Equal state or invitation probabilities are unnecessary. Let $\ell_z > 0$ be the policy-invariant probability that p reaches the focal one-pledge history in state z; in the uniform early block, $\ell_a = \ell_b = \beta = 1/2$. Together with $\pi \in (0,1)$, strict positivity keeps both states possible at the pooled history. Reserve readiness may depend on z only through the policy-common conditional branch values O_z and B_z, and hence u_z and g_z; display-dependent reserve readiness is excluded. Conditional on reaching a pooled count history, Bayes' rule gives

$$\hat{\pi} = \frac{\pi \ell_a}{\pi \ell_a + (1-\pi)\ell_b}, \qquad \bar{u} = \hat{\pi} u_a + (1-\hat{\pi}) u_b. \tag{OA.48}$$

For one such pooled history, the disclosure difference contributed by the pivotal decision is

$$\begin{aligned} J_C - J_R = \pi \ell_a g_a \left[\mathbf{1}\{\kappa \le \bar{u}\} - \mathbf{1}\{\kappa \le u_a\}\right] \\ + (1-\pi)\ell_b g_b \left[\mathbf{1}\{\kappa \le \bar{u}\} - \mathbf{1}\{\kappa \le u_b\}\right]. \end{aligned} \tag{OA.49}$$

Thus the two strict regions survive whenever the labels satisfy $u_a > u_b$, both states have positive conditional probability, and the relevant provider increment is positive. With several pooled clock histories, the calculation applies history by history and is integrated using their policy-invariant reach probabilities. This nonuniform-order result excludes degenerate pooled posteriors and policy-dependent reach: if disclosure changes who is invited or when they arrive, the value difference combines information with a different recruitment process.

G.3 Leakage and administrative removals

Suppose, conditional on the ready anchor preceding p, that its identity leaks perfectly after its pledge but before p's decision with probability $\lambda \in [0,1]$, independently of the state and other events, and that leakage has no direct effect beyond reproducing the roster message at that history. On leaked histories, behavior is the same as under R; on unexposed histories, it is the same as under C. Pre-anchor signals are excluded because revealing Z before any pledge would be more informative than R's empty interface. Holding costs fixed,

$$J_{C,\lambda} - J_R = (1-\lambda)(J_C - J_R). \tag{OA.50}$$

Perfect leakage therefore attenuates either informational ranking and eliminates it at $\lambda = 1$. This clean benchmark applies only to exogenous, state-independent, post-anchor leakage with no direct effect. Equation (OA.50) does not extend to state-dependent or pre-anchor leakage or to leakage that changes trust, reach, timing, execution, outside behavior, or costs. Residual pooling and leakage are pilot estimands; operational relevance requires positive probability on histories that remain meaningfully pooled.

The operational product permits withdrawal before first hit of consent required by the assigned display policy. The operator then deletes the account from P_t, removes an R identity from prospective display, makes the pledge ineligible for that wave's successful roster, and decrements the count. This is an operational consent right, not a strategic action or reentry option in the append-only benchmark or exact witness. Announced pre-success removals for fraud, lost authentication, law, harassment, or safety have the same record effect. After first hit, roster publication and notice use consent already obtained, while destination action remains voluntary; under R, deletion cannot erase identity copies already made public.

A removal kernel that is common across policies and does not reveal identity can enter the public transition law, with post-removal beliefs, success probabilities, and conditional branch values replacing the corresponding baseline quantities; the strict rankings persist only if their inequalities remain strict. Common kernels and the exact witness's exclusion of removals during the early anchor/pivotal block are testable exclusions. Removals that reveal the hidden identity under C, or direct utility and behavior effects of the policy-specific consent terms, harassment, safety, execution, timing, or reserve effects, lie outside the baseline identities and are additional channels in a total product comparison.

G.4 Managed waves and true open enrollment

The analysis covers one bounded, managed, expiring, one-shot invitation wave in which each account has one pledge opportunity and no repeat opportunity. It provides no formal option to wait, revisit a declined decision, withdraw strategically, or condition a pledge on a future named signer. The pre-trigger consent withdrawal in G.3 is an operational deletion, not a modeled timing strategy. Ordinary post-decline migration remains available through O_z and B_z.

An exogenous, policy-invariant order with no repeat opportunities fits the environment even when its probabilities are nonuniform. Proposition 2 integrates every permutation of the early anchor/pivotal block, and equation (OA.49) integrates pooled histories under a common reach distribution. The witness remains a staged invitation design because s and r are contacted after the early block. True open enrollment, strategic waiting, repeat visits, endogenous sharing or reach, and timing-equilibrium selection are deferred. The paper supplies neither a dynamic open-enrollment model nor a robustness claim for that domain. Proposition 2 therefore ranks matched random-order managed waves, not open processes in which disclosure changes who arrives or when.