

Content Exploration Beyond the Feed: Creator Supply and the Shared Corpus

Yuanyuan Shen
Snap Inc.
New York, NY, USA
yshen2@snapchat.com

Wenjie Li*
Snap Inc.
Palo Alto, CA, USA
wenjie.li@snapchat.com

Yiren Yan
Snap Inc.
Palo Alto, CA, USA
yyan5@snapchat.com

Chunhui Zhu
Snap Inc.
Palo Alto, CA, USA
chunhui.zhu@snapchat.com

Abstract

Industrial recommenders give new content initial views through budgeted exploration, then use early performance to decide further delivery. On many short-video platforms, exploration is the primary route by which new videos reach viewers. Viewer-side tests measure consumption, while the published budget objectives we review omit creator response. We analyze four experiments on a major short-video platform. An eight-month creator ablation finds that production exploration raises videos posted per creator by 8.55% and creators posting at least once by 7.10% relative to a minimal floor. A budget-matched reallocation raises creator participation with no detectable short-run viewer-side change. A year-long viewer ablation separately finds 1.74% more video views but 2.13% less view time. A delivered view creates immediate feed value, can trigger organic take-up, and can induce creator supply. Take-up and supply replenish a shared corpus, creating two measurement limits. Viewer-side A/B tests cancel the corpus effect when both arms consume the same corpus. Giving each arm its own corpus avoids cancellation, but turnover still controls the horizon. If the corpus turns over at rate ω per posting cycle, a t -cycle experiment expresses at most ωt of the eventual corpus effect. More users reduce noise without making the corpus turn faster. Before the corpus path visibly bends, data cannot distinguish a modest fast effect from an arbitrarily large slow one, so a valid confidence interval may lack a finite upper endpoint. As predicted, the three-week co-diverted experiment cannot determine whether the eventual corpus effect is positive or negative. Within the window, it identifies the direct feed effect, and an exploratory cohort analysis detects organic lift after exploration ends. Together, the experiments establish a positive creator response, measure the gross corpus flow visible within three weeks, and show the design and duration needed to identify total value.

CCS Concepts

• **Information systems** → **Recommender systems**; *Collaborative filtering*; • **Computing methodologies** → Multi-agent systems.

*Work done while at Snap Inc.

Keywords

content cold-start, content distribution, creator economy, two-sided platforms, bandits, short-video

1 Introduction

Short-video platforms receive enormous volumes of new uploads, most of which initially reach few viewers. Ranking models favor content with an observed response, so new videos lack the history needed to compete. *Exploration* gives them initial views. On many short-video platforms, exploration is the primary distribution channel for most new videos.

Platforms commonly give each new video an *exposure budget* and use early performance to decide whether to release more views [10, 16, 19, 28, 33]. We study systems that allocate exposure in stages before each video’s exploration period ends. A video that clears an early bar receives more exposure. A video that misses it leaves the exploration pool (§2).

This design creates an evaluation problem. Viewer-side A/B tests can show mixed or negative consumption effects even though exploration supplies most new-video distribution. Related systems report the same short-window pattern [11, 30]. Longer one-sided tests share the same corpus across arms, so its contribution cancels at every horizon.

On the major short-video platform we study, an eight-month creator ablation shows that production exploration raises videos posted per creator by 8.55% and creators posting at least once by 7.10% relative to a minimal floor. A separate budget-matched reallocation raises creator participation with no detectable short-run viewer-side change. A year-long viewer ablation finds 1.74% more video views and 2.13% less view time relative to disablement. The viewer path shows no sustained buildup over twelve months. These results establish a positive creator response and isolate the immediate viewer trade-off.

Prior isolated-submarket experiments establish two corpus links: exploration expands the discoverable corpus, and random thinning shows that reducing the corpus lowers satisfied-user counts [30]. We ask how a deployed budgeted mechanism can measure the resulting value. The creator ablation measures the posting response that replenishes the corpus, and the viewer ablation measures the immediate feed trade-off. Co-diversion jointly randomizes matched creator and viewer submarkets, so each arm develops its own corpus. Our three-week probe lands at the estimation floor predicted by the

theory. At that horizon, a cohort analysis measures gross corpus flow, while the aggregate asymptote remains unsigned. Figure 1 maps the three channels to the designs that identify them.

Published cold-start systems score and budget videos with viewer and corpus quantities such as predicted discoverability and click-through. None of the systems we review includes the creator’s later posting response in its *budget* objective. Production feed ranking outside this class has modeled creator incentives for years [32] (§2).

Cold-start systems must decide within the delivery window whether a video receives more exposure or leaves the pool. At that point, the system observes only fast viewer signals. The delayed posting response and corpus turnover therefore require offline estimates. Historical creator experiments and corpus-age data supply those estimates, and budget objectives can use them as priors. A horizon-matched creator experiment identifies the supply response. Co-diversion retains the shared-corpus channel (§5.3 and §5.5).

We model the ecosystem that exploration feeds and test its channels with four experiments on a major short-video platform (§3.1). The linear model supports the value accounting. The timing results require only corpus turnover and survive without linearity (§4.2). The two long-running ablations compare the mechanism with disablement or a minimal floor, so their estimands are invariant to allocator details. The budget-matched experiment separately changes allocation at a fixed nominal budget. We contribute a valuation and measurement framework for these rules. Algorithm design lies outside our scope. We make three contributions:

- (1) **Signed creator response.** An eight-month ablation of production exploration raises videos posted per creator by 8.55% and creators posting at least once by 7.10% above a minimal floor. The published cold-start budget objectives we review omit this response. A budget-matched reallocation raises creator participation with no detectable short-run viewer cost (§5).
- (2) **Value and identification.** We separate a delivered view’s value into the immediate view, organic take-up, and induced creator supply, then map each channel to its identifying randomization. A posting loop with gain $\mathcal{R}$ amplifies posting by $1/(1 - \mathcal{R})$. One-sided designs identify the direct and supply channels but exclude the shared corpus by construction. These are distinct estimands, and the designs determine which net signs are identified (§3.2, §4.1, §5).
- (3) **Corpus horizon.** Only isolated cells retain the corpus contribution in their contrast, and they must wait for the corpus itself to turn over. Under a monotone stock transition with turnover ω , a t -cycle experiment expresses at most ωt of the asymptote. Additional users reduce noise but cannot advance that transition. The three-week probe lands at the predicted estimation floor. An exploratory cohort analysis measures gross corpus flow at the feasible horizon, while the aggregate asymptote and net value after displacement remain unresolved. This result separates what three weeks can identify from what requires a longer isolated trajectory (Theorem 1 and Proposition 3; §5.5).

The long-running ablations estimate a viewer-feed effect after substitution and an aggregate posting response. These effects use different units, so we report them separately. A total-value estimate requires a common scale. Unlike one-off traffic grants, the ablations measure a standing mechanism on both market sides over months. The shared-corpus value requires co-diversion at the prescribed

horizon [30]. Below a curvature threshold, that design can lower-bound an expressed increment while the asymptote still lacks a finite upper bound (§5.5).

2 Related Work

Content Cold-Start in Recommender Systems. Cold-start is among collaborative filtering’s oldest problems [20, 22, 27], met on video platforms with a dedicated exploration phase before full recommendation [12]. We treat that phase as a first-class distribution channel carrying creator as well as viewer value.

Long-Term and Slow-Expressing Effects. Short experiments are known to mismeasure long-run impact, with slow expression treated as a *user-side* phenomenon to be estimated around. Existing methods include cohort and stepped-wedge designs for user learning [17], surrogate indices [2], and temporal designs for carryover [4]. Here the slowness is structural, a two-sided supply loop with a corpus stock, and the question is when the long-run effect is observable at all. Together, Theorem 1 and Corollary 1, the $N^{-1/3}$ detection law of Proposition 3, and the $N^{-1/5}$ estimation law of Theorem 2 bound what a finite-length experiment can see of the equilibrium effect. The gain–relaxation relation is classical in Markov mixing and critical slowing down. The phenomenon of unbounded confidence sets behind Theorem 2 is classical in weakly identified models [14, 15]. We specialize these results to the exponential-approach family, derive an operational curvature constant and horizon laws, and tie them to exploration’s equilibrium value.

Industrial Cold-Start Allocation. Published cold-start systems share two features. New items receive a per-item exposure allocation, and continued exposure is gated on early performance. Wang et al. [33] allocate per-item traffic under a global budget with a floor and ceiling by inverting a learned discoverability probability. Wang et al. [34] screen candidates using a Beta posterior that balances satisfaction against item quality. Jeon et al. [19] reserve quota below an impression threshold, and Chen et al. [10] pair an exposure floor with flow steered toward quality content. The full mechanism analyzed here adds a bounded delivery window with escalation and withdrawal, as in the tiered system of Shen et al. [28]. The platform we study runs this mechanism class, although published accounts differ in how much of its structure they report.

Objectives and Their Timing. The systems differ in what determines the budget. Wang et al. [33] maximize the number of items becoming “discoverable post exploration.” The tiered system above budgets on predicted click-through and gates escalation on realized CTR. None of the systems reviewed above includes a term for the creator’s subsequent supply behavior *in its budget objective*, a statement about the published record that the structural argument of §1 predicts. That system motivates its design partly by producer incentives, yet no producer variable enters its budget rule.

The gap is specific to this class. Tu et al. [32] reshape a production feed’s feedback distribution to nurture creators, raising creation at no measured consumer cost. Wang et al. [35] derive allocation policies that price creators’ future growth. The settings differ in cadence. Feed ranking can adapt against long-run outcomes, whereas a cold-start budget must be set, escalated, and withdrawn within a bounded window whose exit gate returns to viewer-side evidence.




Figure 1: Value channels and identifying designs. Exploration has an immediate feed effect, can change creator supply, and feeds a shared corpus through both take-up and new posts. The shaded channels express slowly. Each dashed design reaches a different channel. A one-sided design shares the corpus across arms, so that channel cancels. A separate budget-matched co-diverted variant perturbs allocation at fixed nominal budget and measures the creator response (§5.4).

Removing that gate reportedly degrades viewer-side metrics [28], yet its creator-side cost is unmeasured.

Two prior results sit closest to ours. With two separate interventions, Su et al. [30] establish that exploration grows the discoverable corpus through joint user and corpus co-diversion. They then show that random corpus thinning costs satisfied users, with the benefit accruing over months. Shen et al. [29] characterize the cost imposed by co-diversion itself. We identify the creator-supply response to a standing budgeted mechanism, separate structural cancellation from horizon attenuation, and derive the design and information limits that govern direct measurement of the corpus-mediated consequence.

Adjacent Lines. Prior systems in this family used bandit exploration [11, 21]. It valued a view for information about an arm. Budgeted allocation instead prices the view as a distribution quantity (Proposition 2). Ad delivery under budgets [3, 24] contributes the mechanics but prices impressions in a common currency and omits the creator’s behavioral response. Field experiments on two-sided platforms establish that exposure causally shapes creator supply. Randomized traffic raises production by 5.87% [18]. Additional impressions produce higher-quality and more diverse output [36], as do peer awards [8]. Theory connects demand allocation to participation [5, 9, 25, 26]. Closest to our reading, Yao et al. [37] formalize the trade-off in theory and simulation. Our ablations price a standing production mechanism on both market sides over months. Prior experiments study one-off traffic grants over weeks. Fairness work on provider-side exposure [1, 7] studies minimum-exposure guarantees as commitments, while this class grants floors for operational reasons.

3 The Mechanism and Its Ecosystem

3.1 Budgeted Exploration as Practiced

Creators submit content items into an *exploration pool* $\mathcal{P}_t$. Each item s remains eligible in the pool for at most $T_{\max}$, its time-to-live (TTL). The pool competes for an aggregate exploration capacity of V_t views. A *view budget* $B(s) \in \mathbb{Z}_{\geq 0}$ is the target number of

exploration views to deliver to s within its TTL, subject to capacity and continued eligibility.

DEFINITION 1 (BUDGETED EXPLORATION PROBLEM). *Given pool $\mathcal{P}_t$, capacity V_t , and a value function $U(s, B)$ for delivering B views to content item s , choose budgets to*

$$\begin{aligned} \max_{\{B(s)\}_{s \in \mathcal{P}_t}} \quad & \sum_{s \in \mathcal{P}_t} U(s, B(s)) \\ \text{s.t.} \quad & \sum_{s \in \mathcal{P}_t} B(s) \leq V_t, \quad B_{\min} \leq B(s) \leq B_{\max} \quad \forall s, \end{aligned} \quad (1)$$

where $B_{\min}$ is a minimum budget guarantee (cold-start views) for every content item and $B_{\max}$ a per-item ceiling.

Industrial practice resolves this generically (§2). Every eligible item receives the guarantee $B_{\min}$ before any viewer response exists. Views above it, *escalation*, are conditioned on early response inside the window. Delivery of the remainder stops, *withdrawal*, when the budget is exhausted, expires, or repeatedly misses viewer-response criteria. Both decisions must rely on the fastest-expressing quantities available, and Theorem 1 will show that decisions made on those quantities are exactly the ones blind to the creator-supply and corpus-stock channels of Proposition 2. The platform we measure runs a standard member of this class. How budgets are set is platform-specific. The theory conditions on the resulting budgets, while the ablations compare the mechanism with its absence or a minimal floor. Both analyses apply without specifying the platform’s particular budget-setting rule.

Budget assignment is decoupled from delivery. The contract requires only that each content item receive its target views from relevant viewers within its TTL, and a dedicated funnel, ranking boosts, quotas, or guaranteed-delivery pacing honor it equally, so the valuation and evaluation transfer to any of them.

3.2 An Ecosystem Model

Whether budgeted exploration pays for itself is put to viewer-side A/B tests, whose consumption results can be mixed or negative

(§5.2). The viewer-side contrast identifies only the immediate feed trade-off. Feedback value compounds over periods and grows with persistence. No short window can observe a slow effect, whether it is an A/B test's measurement window or the mechanism's own gate.

States. We work in discrete time. One *period* is a submission-feedback cycle that includes submission, exploration views within TTL, and the creator's next posting decisions. Its calendar length is a platform property external to the model. Two state variables describe the ecosystem at period t . The first is Λ_t , the new submission rate or *supply*. The second is K_t , the *organic corpus* of content items remaining in organic distribution after their exploration window.

Model ingredients. Six ingredients parameterize the model. The baseline submission rate is Λ_0 . The *creator response* $g(\cdot)$ gives expected next-period submissions as a function of views received. We assume it is increasing and *concave*, consistent with the saturating responses reported in field experiments [18, 36]. Expected total views are $\bar{V}(B) = B + q(B) V_{\text{org}}$. This expression includes B delivered exploration views plus, with *organic take-up* probability $q(B)$, a further V_{org} organic views. We assume q is nondecreasing and all view-value weights are nonnegative. The corpus turns over at per-period rate $\omega \in (0, 1)$. Viewer value is η per corpus unit and period, so aggregate consumption (total view time, say) is locally affine in the stock, $E_t = \eta K_t + \text{const}$, the constant carrying consumption outside the corpus. Writing η_o for gross value per organic view, our normalization sets $\eta/\omega = \eta_o V_{\text{org}}$: both sides are the lifetime organic value of one taken-up item.

Dynamics. These ingredients give a pair of linear recursions.

$$\Lambda_{t+1} = \Lambda_0 + \mathcal{R} \Lambda_t, \quad K_{t+1} = (1 - \omega) K_t + q(B) \Lambda_t. \quad (2)$$

The first says next period's supply is the baseline plus a fed-back share of this period's supply. The second says the corpus retains a fraction $1 - \omega$ and gains the taken-up share of new submissions.

Where the recursions settle, and what one more delivered view is worth once the loop has run its course, follow directly.

PROPOSITION 1 (EXPLORATION MULTIPLIER). *For $\mathcal{R} < 1$, the dynamics (2) converge to*

$$\Lambda^* = \frac{\Lambda_0}{1 - \mathcal{R}}, \quad K^* = \frac{q(B)}{\omega} \Lambda^*,$$

and one additional delivered view yields $\delta/(1 - \mathcal{R})$ cumulative additional submissions over the ecosystem's response. Every induced submission attracts views that induce further submissions, producing the multiplier.

PROOF. The affine recursion for Λ_t contracts to its unique fixed point $\Lambda_0/(1 - \mathcal{R})$. Substituting into the K recursion gives K^* . A marginal view induces δ submissions, which induce $\delta\mathcal{R}$, then $\delta\mathcal{R}^2$, summing to $\delta/(1 - \mathcal{R})$. $\square$

Gain and sensitivity. The creator response enters through two distinct quantities. The *loop gain* is $\mathcal{R} \triangleq g(\bar{V}(B))$, the expected number of *further* submissions one submission induces and the eigenvalue of the supply recursion. The *sensitivity* is $\delta \triangleq g'(\bar{V}(B))$. It is the marginal posting response to one additional delivered view, and the quantity that values a view. The recursions are mean-field approximations. Both δ and q are population averages, and

the subscripted δ_c , q'_s of Proposition 2 are their individual-level counterparts. A healthy platform is *subcritical*, with $\mathcal{R} < 1$. Each submission then induces less than one additional submission, so the supply loop remains stable. The budget B is the lever through which the platform sets $\mathcal{R}$.

Supply loop and corpus stock. Stacked as $x_t = (\Lambda_t, K_t)$, the recursions read $x_{t+1} = Mx_t + u$ with eigenvalues $\mathcal{R}$ (the *supply loop*) and $1 - \omega$ (the *corpus stock*). The two play different roles. The supply loop governs how much a view is worth. The corpus stock governs how long an experiment must run before that value is visible. No result below requires the corpus to be slower than the supply loop: because $1 - \omega$ is mechanical turnover and involves g nowhere, the horizon bound holds for any creator response and any ordering of the two eigenvalues (Theorem 1).

REMARK 1 (FIXED CAPACITY). *Fixing the per-item budget is compatible with finite capacity. Under (1), a larger pool lowers views per item through V/Λ , adding negative feedback. Ignoring that congestion can overstate the supply multiplier. The timing statements remain unchanged because they depend on ω alone. The local elasticity derivation is in Appendix A.1. Concavity is needed only near current budgets, where the field evidence above is most direct.*

The value of a view and its multiplier. The dynamics converge to $\Lambda^* = \Lambda_0/(1 - \mathcal{R})$ and $K^* = q(B) \Lambda^*/\omega$. One marginal delivered view induces δ submissions, whose descendants form the geometric sum $\delta \sum_{k \geq 0} \mathcal{R}^k = \delta/(1 - \mathcal{R})$ (Proposition 1). Valuing a view by its direct response alone therefore understates its worth by the factor $1/(1 - \mathcal{R})$, which grows without bound as the platform approaches criticality, $\mathcal{R} \rightarrow 1$.

REMARK 2 (SCOPE OF THE LINEAR MODEL). (2) is the *tangent description of a subcritical ecosystem at steady state, including V_{org}* . The multiplier applies locally at that steady state. Theorem 1's bound carries no such restriction. The η 's are gross values. An on/off contrast measures the gap net of displaced impressions, the decision-relevant quantity. Additional nonnegative feedbacks would raise the spectral radius of M . Opposite-sign feedbacks such as crowd-out or quality dilution can lower it and lie outside this model (Appendix A.1).

4 Value and Measurability

4.1 What a Delivered View Is Worth

The multiplier of §3.2 values the supply loop as a whole. An allocator needs something finer, the value of one more view to one particular content item. Under the same model, that value separates into three channels. They are the posting that the view induces, the view itself, and organic take-up.

PROPOSITION 2 (MARGINAL VALUE OF A VIEW). *Under (2), the cumulative consumption value of one additional view delivered to content item s from creator c with budget $B(s)$ is*

$$\text{MV}(s) = \underbrace{\delta_c v^*}_{\text{direct creator channel}} + \underbrace{\eta_e}_{\text{the view itself}} + \underbrace{q'_s(B) V_{\text{org}} (\eta_o + \delta_c v^*)}_{\text{take-up channel}},$$

where η_o, η_e are gross per-view consumption values of organic and exploration delivery, $v^* = (\eta_o \bar{q} V_{\text{org}} + \eta_e \bar{B})/(1 - \mathcal{R})$ is the total consumption value of one submission, and $\bar{q}, \bar{B}$ are pool averages.

Proof sketch. A marginal view raises total views by $1 + q'_s V_{\text{org}}$ and induces $\delta_c(1 + q'_s V_{\text{org}})$ submissions, each worth v^* , which already sums the descendant chain (Appendix A.1).

The take-up term multiplies a small probability change by a large reward. One extra view barely changes the probability of organic take-up, but conditional on take-up, the item receives V_{org} additional organic views in expectation. Two readings of Proposition 2 matter later:

(1) Budget scale and ranking. At the resolution identified here, a pool-level creator sensitivity re-prices the budget's shadow value but supplies no new within-pool ranking signal. Creator-term re-ranking requires variation in δ_c across content items, which none of our designs identifies. The proposition shows where that heterogeneity would enter. The view-itself term has no item subscript and drops from the allocator's comparison. The varying terms are proportional to q'_s and δ_c . Organic reach V_{org} weights take-up and enters v^* , which values induced supply. A view is worth more near the take-up margin, where q'_s is largest, and when conditional organic reach is larger. The latter can represent demand per unit supply, coverage, diversity, or freshness. When organic delivery already reaches nearly every item, q'_s approaches zero and an extra view buys little beyond itself.

(2) Sign. A randomized contrast measures *net* of the displaced impression, so a negative measured direct effect is consistent with $\eta_e > 0$ (Remark 2). Subtracting the same opportunity cost from every candidate leaves their ordering unchanged, but varying that cost across candidates can change the ranking.

4.2 What a Measurement Can See

Proposition 2's take-up and creator channels are deferred, and this subsection explains why they are hard to observe. The creator channel arrives a posting cycle late and compounds over the supply loop. The take-up channel pays off only through organic distribution that the corpus has yet to deliver. We analyze the extreme member of the intervention family, switching exploration off entirely, because it perturbs the corpus stock with the largest possible asymptotic gap. Any milder test, such as a budget change or parameter tweak, moves the same stock at the same turnover toward a smaller gap. The switch-off case is a best-case benchmark for what a fixed horizon can express. Consumption quantities such as E_t enter below only as observables of the corpus stock K_t , which is not directly observed.

Multiplier and horizon. For a feedback loop with eigenvalue $\lambda \in (0, 1)$, define its equilibrium multiplier $A(\lambda) \triangleq 1/(1 - \lambda)$, the steady-state response to one unit of input added each period. Its *relaxation time* is $\tau(\lambda) \triangleq 1/\ln(1/\lambda)$, so a switch decays as λ^t . Both grow without bound as λ approaches 1: stronger feedback also means slower forgetting. The timing results concern the transient after a switch, so production need not remain at equilibrium. The dynamics (2) contain two such modes. The supply mode has eigenvalue $\mathcal{R}$, the number of posts induced per post, and carries the comparatively fast posting multiplier. The corpus mode has eigenvalue $1 - \omega$, its per-period retention, and sets the experiment's clock. Beyond the immediate slot value, viewer-side effects are corpus-mediated. The model locates them in time. A short window reveals only a fraction of a slow corpus effect even without noise.

THEOREM 1 (HORIZON BLINDNESS). *Disable exploration at $t = 0$ from equilibrium and write E^* for pre-switch equilibrium consumption. With consumption affine in the corpus as above, the gap decomposes as $E^* - E_t = \Delta_{\text{direct}} + \Delta_{\text{corpus}} F(t)$, the direct term present from the first period. If the corpus retains a fraction $1 - \omega$ per period and its post-switch inflow stays at or above its eventual level, then*

$$F(t) \leq 1 - (1 - \omega)^t \leq \omega t \quad \text{for every integer } t \geq 1.$$

Proof sketch. Among inflow paths satisfying this condition, an immediate jump to the eventual post-switch inflow maximizes the corpus gap at every date. Iterating the stock recursion then gives the normalized gap $1 - (1 - \omega)^t$. Bernoulli's inequality gives the second bound. The bound is the operational result. At horizon t , at most a ωt fraction of the eventual corpus gap has materialized, even without noise. The linear model locates the delay. If take-up changes immediately, the first period expresses a fraction ω . If corpus inflow changes only after posting responds, onset is quadratic. When the supply response settles before the corpus has appreciably turned over, the gap thereafter follows an approximately linear ramp whose extrapolated delay lies between zero and one post-switch supply-loop multiplier $A(\mathcal{R}_{\text{off}}) = 1/(1 - \mathcal{R}_{\text{off}})$, where $\mathcal{R}_{\text{off}}$ is the post-switch loop gain. The delayed endpoint attains that upper limit. The closed-form path, including mixtures of the two cases, is in Appendix A.1. This timing reconciles mixed or negative short viewer-side readings with a valuable mechanism: the corpus-mediated demand value arrives on the slow eigenvalue, after the window closes.

Equilibrium gain and relaxation time coincide to within one period. A coverage target multiplies that time only by its logarithm.

COROLLARY 1 (VALUE AND MEASURABILITY DUALITY). *Consider a stable nonnegative ecosystem $x_{t+1} = Mx_t + u$ and a feedback loop with real eigenvalue $\lambda \in (0, 1)$.*

- (i) *the equilibrium multiplier $A(\lambda)$ and relaxation time $\tau(\lambda)$ coincide to within one period, so $\tau(\lambda) \leq A(\lambda) \leq \tau(\lambda) + 1$. Observing a $(1 - \epsilon)$ fraction of a switch requires the horizon $t_\epsilon = \ln(1/\epsilon) \tau(\lambda)$.*
- (ii) *consequently an experiment of length t reads the partial sum $\sum_{k < t} \lambda^k \leq t$, whatever the true multiplier $A(\lambda)$.*

Within the linear model, no parameter regime of such a loop combines large equilibrium value with a short measurement horizon.

Part (ii) is the operational form (full proof in Appendix A.1). What a test can read is bounded by its own duration, so a gap settling at Δ_∞ reads early as $\Delta_\infty \omega t$. The multiplier is achieved by the slowness. The corpus stock instantiates the corollary at eigenvalue $1 - \omega$. Its gain is $1/\omega$, and its relaxation time lies within one period of that value. The more evergreen the corpus, the more stock a unit of inflow buys and the longer any test must run to see a fixed fraction of it.

Adding users does not change that clock. It only sharpens the reading of wherever the clock stands. Appendix D works these results at illustrative magnitudes.

PROPOSITION 3 (SAMPLE SIZE CANNOT BUY HORIZON). *Consider an on/off experiment with N users per arm on a corpus-stock metric with asymptotic gap Δ_∞ , read at horizon t in the early regime $\omega t \ll 1$.*

- (i) *The expressed fraction $\Delta(t)/\Delta_\infty \approx \omega t$ is independent of N . Larger samples estimate the same attenuated number more precisely.*
- (ii) *Even mere detection of the attenuated gap is horizon-bound. With*

per-user noise scale σ and detection level z , a snapshot analysis first detects at $t_{\text{det}} = \frac{z\sigma\sqrt{2/N}}{\Delta_\infty\omega} \propto N^{-1/2}$, and a cumulative analysis with independent per-period noise at $t_{\text{det}} = \left(\frac{2\sqrt{2}z\sigma}{\Delta_\infty\omega\sqrt{N}}\right)^{2/3} \propto N^{-1/3}$.

PROOF SKETCH. (i) follows from Theorem 1, whose statement contains no N . For (ii), equate the gap $\Delta_\infty\omega t$ (snapshot) or its running mean (cumulative, standard error $\sigma\sqrt{2/(Nt)}$) to z times its standard error and solve for t . $\square$

Detecting a nonzero finite-horizon effect is easier than estimating its asymptote.

THEOREM 2 (ESTIMATION FLOOR). *Let y_t be the observed between-arm contrast in a corpus-stock metric at horizon t . For $t = 1, \dots, T$, suppose*

$$y_t = c + \Delta_\infty(1 - e^{-\omega t}) + \varepsilon_t, \quad \varepsilon_t \stackrel{\text{iid}}{\sim} N(0, s^2),$$

with level c and rate ω unknown and $s^2 = 2\sigma_{\text{eff}}^2/N$ for N users per arm, where σ_{eff} is the effective per-user transient-noise scale. Define the curvature signal-to-noise ratio

$$C = \frac{\Delta_\infty\omega^2\sqrt{S_4}}{2s}, \quad S_4 = \sum_{t \leq T} t^4.$$

Whenever $C \leq 1$, every confidence procedure for Δ_∞ with uniform coverage $1 - \alpha$ over positive amplitudes and rates returns an infinite upper endpoint with probability at least $\frac{1}{2} - \alpha$. Raising N moves the crossing $C = 1$ only as $T \propto N^{-1/5}$, improving to $N^{-1/3}$ once ω is estimated from auxiliary data.

The proof, a two-point total-variation indistinguishability argument over a matched-product family, appears in Appendix A.2, with the panel reduction at Lemma 1 and the formal statement at Theorem 3. The construction holds the early slope $\Delta_\infty\omega$ fixed while sending $\omega \rightarrow 0$ and $\Delta_\infty \rightarrow \infty$. Below the threshold these paths are statistically indistinguishable. Until the trajectory visibly bends, monotonicity can still supply an increment lower bound. The slow component is the corpus stock. The regimes in which exploration is most valuable are those in which feasible-horizon inference may admit only an increment lower bound, and §5.5 uses a horizon-matched endpoint accordingly.

Figure 3 and Figure 7 visualize the estimation floor and delay family.

The theorem treats the contrast series as independent draws with a single noise scale. Panel data earns that form: persistent per-user differences drop out once the level is profiled, serial correlation is absorbed into σ_{eff} , and only a unit root would change the exponents.

LEMMA 1 (PANEL REDUCTION). *Let per-user outcomes carry persistent user effects of arbitrary variance σ_α^2 and stationary transient noise, so that $\text{Cov}(\eta_s, \eta_t) = \frac{2}{N}[\sigma_\alpha^2 + \gamma(s-t)]$. Then: (i) because $\mathbf{1}$ is an eigenvector of the persistent component, profiling any nuisance set containing the intercept Δ_{direct} renders every information quantity below independent of σ_α^2 , and persistent effects can only enlarge the impossibility region of Theorem 3; (ii) for AR(1) transient noise with parameter ρ and marginal variance σ^2 , all statements hold with $\sigma_{\text{eff}}^2 = \sigma^2(1+\rho)/(1-\rho)$, up to a factor $1 + O(\rho/((1-\rho)T))$, by the exact tridiagonal form of the AR(1) inverse; (iii) stationarity is necessary: under random-walk user drift the T^5 information accumulation below degrades to T^3 , because a unit root mimics the ramp.*

Stationarity is diagnosable in logs: the variance of user-mean-centered outcomes must not grow with the window.

The proof is in Appendix A.2. Part (iii) is why §5 reports a stationarity check rather than assuming it.

Drift. A changing environment moves the steady state without changing timing or identification. Theorem 1 uses only the corpus stock's retention and holds pathwise under time-varying turnover, the gain-relaxation relation survives because both are monotone in the same λ , and the exponents in Proposition 3 follow from that algebra alone (Appendix A.1). Unmodelled drift only widens the region with no finite upper bound. The multiplier and Proposition 2's level are local statements. The probe in §5.5 imports drift from the concurrent ablation. Each design contrasts arms at a common date, so shared changes in viewer taste or the corpus drop out.

4.3 What Each Design Identifies

Each randomization holds a different part of the ecosystem in common between its arms. What is common cancels.

PROPOSITION 4 (WHAT EACH RANDOMIZATION DESIGN IDENTIFIES). *In the model (2), suppose an experiment holds a fraction p of one side of the market at the perturbed condition (the floored arm). Treat immediate within-feed substitution as part of the direct channel, and suppose any remaining change in organic consumption is mediated by the shared corpus state. The following statements hold.*

- (i) **Viewer-side.** *Disabling exploration delivery in enrolled users' feeds identifies the direct channel exactly at every horizon and arm size. Corpus-mediated consumption is common to both arms and cancels.*
- (ii) **Creator-side.** *Holding enrolled creators' content at the floor identifies the equilibrium supply response $\Lambda_0[(1-\mathcal{R})^{-1} - (1-\mathcal{R}_f)^{-1}]$ at the perturbed aggregate, up to a capacity-release term $\epsilon(p)$. Here $\mathcal{R}_f$ is the floored arm's loop gain. The term $\epsilon(p)$ is nonnegative because flooring can only raise delivery in the non-floored arm relative to full rollout. It scales with the floored market share and is absent from the co-diverted design, where an isolated submarket exchanges no capacity with the platform. Whatever consumption the induced corpus generates is common to both arms and therefore cancels from the contrast.*
- (iii) **Co-diverted.** *Isolating matched creator and viewer submarkets retains the corpus-mediated component in the contrast, attenuated by the general $F(t)$ of Theorem 1. This differs from the one-sided cancellation of (i)–(ii). Flooring the submarket's own creators lowers its own supply, so the cell sits toward the theorem's delayed end.*

In particular, the corpus channel enters no one-sided contrast at any horizon or sample size. It enters only under co-diversion and expresses a fixed fraction at horizons of order $1/\omega$.

For (i)–(ii), subtracting arm means cancels the common K_t . Under co-diversion each cell evolves its own K_t , so subtraction retains their stock gap. The capacity-release sign in (ii) follows because flooring enrolled traffic frees nonnegative capacity for the remaining pool. Appendix A.1 gives the full algebra.

These designs identify population aggregates. Per-item q'_s and δ_c remain unidentified. An external turnover estimate improves corpus estimation from $N^{-1/5}$ to $N^{-1/3}$.

Table 1: Identification map for the four randomized designs. The final column gives the sample-size law. For the fast channels it sharpens the estimate at a fixed horizon. For the corpus rows it sets the required horizon.

Target	Design	Required horizon	Rate in N
Direct effect	Viewer ablation	Immediate	$N^{-1/2}$
Supply response	Creator ablation	Supply loop $\tau(\mathcal{R})$	$N^{-1/2}$
Allocation response	Budget reallocation	Supply loop $\tau(\mathcal{R})$	$N^{-1/2}$
Corpus ramp	Co-diversion	Early ramp	$N^{-1/3}$
Corpus asymptote	Co-diversion	Curvature threshold	$N^{-1/5}$

5 Testing the Identification Map

The four experiments identify complementary channels. Each tests the channel and time path assigned to it by Table 1. Two long-running one-sided ablations identify the direct link and the supply response in sign, although shared capacity can overstate the latter under full rollout. Two co-diverted experiments test budget reallocation (§5.4) and the delayed corpus response (§5.5), which one-sided designs exclude by construction. In both ablations, randomization and analysis units coincide, so effects are differences of window-aggregated means at the randomized-unit level, with unit-level standard errors and no clustering. Enrollment is fixed at randomization. Where used, variance reduction regresses outcomes on their pre-experiment values [13]. The creator first stage is mechanical. Exploration views per video rise by the budget-to-floor gap (levels withheld). Arm shares agree within hundredths of a percentage point, with no sample-ratio mismatch. We report effect sizes with 95% intervals. Tabulated rows are a subset of each analysis’s full outcome family. The probe reports the endpoints specified in its experiment design. Each co-diverted experiment realizes one pair of matched submarkets, so its intervals condition on that pair.

5.1 Qualitative Predictions

The theory implies three signatures independent of effect size. They are a direct step, a supply response that can build after onset, and a corpus ramp as isolated stocks diverge. Cumulative and non-overlapping period estimates reveal whether a gap ramps or stays level. One-sided ablations reach the first two responses. Only co-diversion can express the corpus ramp. Figure 2 shows the two observed time paths and the turnover sensitivity below.

5.2 Viewer Ablation of the Direct Channel

The viewer ablation ran for over a year, enrolling eight percent of users randomized between production exploration and full disablement. Creators are common to both arms, so each arm draws from the same corpus supplied by the surrounding creator population. The contrast isolates the *direct channel*, which is the effect of exploration delivery in enrolled viewers’ feeds after within-feed substitution. It excludes corpus effects. Its length separates persistent response from launch novelty and spans a full seasonal cycle.

Table 2 (Panel A) reports a mixed direct effect. Video views rise 1.74% while view time falls 2.13%. Views under 1.5 s rise 7.5%, shifting the view mix toward brief views. Favorites move with view time (−2.3%). Exploration increases viewing starts while reducing viewing duration at this operating point. This is the direct-feed

Table 2: Randomized effects at the tested operating points. Relative arm differences have 95% CIs. Positive values favor production in Panels A–B and reallocation in Panel C. The panels isolate a mixed direct trade-off, a signed supply response, and a budget-matched participation gain.

Metric	Rel. change (%)	95% CI
Panel A: Viewer ablation (§5.2)		
Video views	+1.74	[+1.57, +1.90]
View time	−2.13	[−2.28, −1.98]
Views under 1.5 s	+7.49	[+7.28, +7.70]
Video favorites	−2.32	[−2.72, −1.92]
Panel B: Creator ablation (§5.3)		
Videos posted per creator	+8.55	[+7.14, +9.96]
Creators posting at least once	+7.10	[+6.60, +7.59]
Panel C: Budget-matched reallocation (§5.4)		
Creators posting at least once	+0.77	[+0.35, +1.19]
Videos posted per creator	+0.92	[+0.05, +1.79]
View time	+0.01	[−0.14, +0.16]
Video views	−0.06	[−0.21, +0.09]
Video favorites	+0.15	[−0.44, +0.74]

component of the allocation problem. Proposition 2 also requires the creator response.

The cumulative and non-overlapping period series in Appendix C confirm the reading. The gap is close to full size within weeks. Thereafter the period contrast stays level across most of the window and then attenuates, tracking a decline in exploration’s share of the feed. Cumulative video views move from +2.02% early in the window to +2.32% at the peak and +1.74% at the end. Immediacy without a sustained ramp is the direct channel’s signature under the reading rule of §5.1. The shared corpus contributes zero to this contrast.

The concurrent creator ablation randomizes a different population, yet its consumption changes sign at the same calendar boundary (+1.11% on video views just before it), as shown in Appendix C. Because the experiments began months apart, the same calendar-date break indicates a common operating-point shift. Since both designs exclude the corpus term (Proposition 4(i)–(ii)), this concurrent movement also supplies the drift estimate imported by §5.5.

5.3 Creator Ablation of Content Supply

The creator ablation ran for eight months, enrolling eight percent of the platform’s creators, with about four percent in each arm. This share also sizes the interference term, with $p/(1-p) \approx 0.04$. Creators in the *floored* arm receive a minimal exploration budget per video, and their videos remain eligible for organic retrieval. Creators in the *production* arm receive production exploration budgets (absolute levels withheld for confidentiality). This two-dose comparison falls short of a strict on/off but measures the supply response directly. Only the budget above the floor differs between arms.

Across eight months, exploration raises videos posted per creator by 8.55% (Table 2, Panel B). It also raises the number of creators posting at least once by 7.10%. The two outcomes identify aggregate




Figure 2: Observed time paths and turnover sensitivity. (a) The year-long viewer contrast appears within weeks and stays level until a common operating-point shift, matching the direct-step prediction. Intervals are approximate. (b) The eight-month posting response, indexed to its first measured period, is descriptive and follows onset with a plateau. (c) The corpus-asymptote conclusion depends on the assumed turnover. The 36-day case rules out offsetting either direct view-time cost (the ablation’s -2.13% , drawn dashed, and the probe’s own -1.02% step), whereas the 360-day case admits either sign. Panels (a)–(b) use finer analysis periods than the appendix’s series. Their shapes agree (Appendix C).

and extensive-margin supply responses. Unlike randomized traffic grants lasting a few weeks, this contrast compares the standing production mechanism with a minimal floor over eight months.

The effect on videos posted per creator is $+5.9\% \pm 3.4$ (95%) in the first measured period, below the $+8.55\%$ whole-window average. Together with the later plateau in Figure 2b, this pattern is consistent with an onset period. The increase in creators posting at least once locates part of the response on the extensive margin. Because the two creator arms compete in a shared marketplace, Proposition 4(ii) implies capacity-release inflation at order $p/(1-p)$, so the contrast upper-bounds the full-rollout effect at this operating point (Appendix A.1). The co-diverted probe later provides a descriptive comparison without that shared-capacity term. Recovering the guarantee’s total value would require a zero-exploration creator arm.

5.4 Budget-Matched Reallocation of Exposure

The ablations identify the effects of exploration presence. Proposition 2 also characterizes fixed-budget allocation. A two-week co-diverted experiment holds nominal budget fixed while spreading exposure across more videos in matched cells. Its design fixes the arm contrast, horizon, and endpoints. By concavity, the marginal posting response is larger at lower exposure. Promotion thresholds implement the shift. A penultimate-day spend diagnostic puts realized per-creator budgets within about 3% across arms, with the reallocated arm slightly lower, so any residual dose imbalance biases its posting response downward. The experiment identifies the randomized reallocation contrast and leaves pure breadth elasticity open. The horizon targets fast supply. Creators posting at least once increase by 0.77%, and videos posted per creator increase by 0.92%. Every viewer-side interval includes zero (Table 2, Panel C). This is the predicted signature of higher creator participation without a detectable short-run viewer cost (Appendix C.2).

5.5 Co-Diverted Probe of the Corpus Response

The probe co-diverts viewers and creators into isolated production and floored submarkets [6, 23]. Consumption uses viewer-level means. Posting uses creator-level means. With one partition per arm, intervals condition on these submarkets. Matched isolation cancels the common consumption cost while letting cells approach different steady states [29].

We test video views and view time for a step followed by a delayed ramp. Video favorites per delivered view provide a quality endpoint sensitive to composition. The experiment runs for three weeks from switch-on. The corpus estimator omits the first seven days because videos already in their exploration windows are only partially exposed to the intervention.

The model predicts an immediate step plus a delayed corpus ramp at rate ω , while the concurrent one-sided ablation stays flat (Theorem 1). Because early data identify only the slope $\Delta_{\text{corpus}} \omega$, we report asymptote intervals conditional on turnover and a shape-robust increment. The 36-day timescale is the measured kernel’s single-rate equivalent. The 360-day timescale is the eligibility horizon. They are fixed-turnover sensitivity cases. Kernel calibration and estimator details are in Appendix B.2. Platform policy caps matched-submarket isolation at three weeks because its ecosystem cost grows with duration [29]. At this horizon, Theorem 2 predicts an identified direct step but a potentially unsigned corpus asymptote (Figure 3).

Table 3 collects the readout. The direct step repeats the viewer ablation’s signs and is smaller, consistent with the probe’s later operating point. Video favorites per delivered view confound reception with composition. The estimates use an equal-length pre-period difference because video views and favorites were imbalanced before the switch.

Table 3: Co-diverted three-week readout at the calibrated estimation floor. Entries are relative changes with 95% CIs. Panel B gives intervals for the set-identified corpus asymptote under two turnover timescales. Panel C gives feasible-horizon gross flow; its implied whole-cell contribution is a derived range. Posting is quoted as estimate $\pm$ margin. Panels C–D retain the inferential qualifications stated in text.

Quantity	Estimate
Panel A: Direct step	
Video views	+0.43 [+0.30, +0.57]
View time	−1.02 [−1.14, −0.89]
Video favorites / delivered view	−0.87 [−1.21, −0.53]
Panel B: Corpus asymptote by turnover timescale	
Video views, 36-day timescale	[−0.28, +0.70]
View time, 36-day timescale	[−0.93, +0.30]
Video views, 360-day timescale	[−2.77, +7.04]
View time, 360-day timescale	[−9.30, +2.95]
Panel C: Feasible-horizon corpus flow	
Organic view time in cohort slice	+37.9 [+37.1, +38.7]
Implied whole-cell contribution	+0.1+0.2 pp
Panel D: Videos posted per creator at a matched 21-day horizon	
In-cell design	+0.63 $\pm$ 0.44
One-sided design	+2.79 $\pm$ 0.89

The corpus readout lands at the predicted estimation floor. Without fixing a turnover timescale, neither the ramp fit nor the shape-robust increment yields a finite two-sided upper bound on the corpus asymptote (Appendix B.2). Under the 36-day turnover timescale, the auxiliary analysis rules out a corpus benefit large enough to offset the direct view-time cost. Under the 360-day timescale, either sign remains possible (Table 3, Panel B). The aggregate corpus asymptote remains unsigned at three weeks.

Feasible-horizon corpus flow. Before the aggregate trajectory bends, the theory targets the effect expressed at horizon t . An extrapolated asymptote is premature. A cohort decomposition measures feasible-horizon gross flow among videos posted during the experiment and later served organically after their exploration windows. By construction, exploration can no longer deliver these videos (Appendix C.3). In matched cells, organic view time is 37.9% higher in a slice representing 0.3–0.5% of cell consumption, equivalent to 0.1–0.2 percentage points of the whole-cell contrast (Table 3). Videos posted before the experiment show no detectable difference after the first seven days. The pattern first appeared in a 1% discovery sample. Because the full-population calculation reuses those users, we retain its exploratory status. It measures gross corpus flow by the feasible horizon. The net asymptote after displacement remains unresolved (Appendix C.3).

In-cell supply. With a comparable exploration first stage and the same horizon and instrument, in-cell posting is roughly a quarter of the one-sided response (Table 3, Panel D). The ratio is descriptive because the populations, seasons, reach, and exposure composition differ. It leaves $\epsilon(p)$ and the rollout asymptote unidentified.

5.6 What the Four Experiments Establish

The experiments form an identification sequence across distinct estimands. The measured timescales match the model’s assignment of roles: the posting response plateaus within a few measured periods (Figure 2b), while the two candidate corpus timescales range from comparable to roughly an order of magnitude longer than the posting response (§5.5). Table 4 records the findings and nearest limits.

6 Discussion

The creator-side response is consistent with creators on the margin of continued participation, where modest distribution separates occasional from regular posting [5, 25]. Organic delivery accrues mostly to videos that have demonstrated strong early performance, which is why exploration is the primary channel for most new videos (§1) and the platform’s direct lever over the long tail. Proposition 2 shows where the creator response enters value accounting. Our designs identify only its gross aggregate. Per-item sensitivity and the common net-value scale required to prescribe an objective remain unidentified.

Implications for Practice. Each result yields a design implication.

- (1) **A mixed viewer A/B identifies only the direct trade-off.** A viewer-side contrast excludes creator response by construction. On the platform we study, production exploration raises videos posted per creator by 8.55% and creators posting at least once by 7.10%.
- (2) **A neutral result leaves shared-corpus value unresolved.** The shared-corpus term cancels from a one-sided viewer contrast. A short co-diverted contrast identifies only the finite-horizon increment (Theorem 1). The asymptote remains unresolved. Neither result alone identifies total value.
- (3) **Estimate the slow factors offline and import them.** Traffic cannot rescue the gate’s short window (Proposition 3). Budget policies follow a planning cadence that can incorporate offline estimates from past randomized cohorts or quasi-experimental variation. Historical data on yield by video age can estimate corpus turnover.

Table 4: What the four experiments establish. Signs apply at the realized horizon. Co-diverted rows condition on the realized submarkets.

Experiment	Identified finding	Leaves open
Viewer ablation	Direct feed trade-off: video views +, time −	Shared-corpus effect
Creator ablation	Supply response: posting +	Full-rollout magnitude and corpus value
Budget reallocation	Reallocation contrast: participation +; viewer CIs span 0	Breadth elasticity and long-run corpus value
Co-diverted	Direct step mixed; exploratory gross cohort flow +	Net corpus asymptote at three weeks

- (4) **Match the design to the channel.** Viewer-side randomization measures the direct channel, and creator-side randomization measures the aggregate supply response at the tested operating point. Only co-diversion retains the corpus channel required to estimate its asymptote (Table 1). The exploratory cohort analysis measures gross corpus flow at a feasible horizon using videos posted during the experiment and served organically after their exploration windows (Table 3; §5.5).
- (5) **Plan in calendar time.** A tenfold increase in users reduces the required horizon by at most a factor of 3.2. Below the curvature threshold (Theorem 2), use the finite-horizon increment and any model-restricted lower bound as the reported endpoints. Plan horizons from turnover, effect-size, and noise inputs.
- (6) **Treat allocation as a separate experimental margin.** At a fixed nominal budget, reallocation raises creator participation with no detectable viewer-side change (§5.4).

Limitations. The ablations estimate the budget’s effects at the tested operating point. The reallocation experiment tests one direction of allocation there, leaving the allocation rule unidentified (§5.4). Open questions include posting-response persistence, generalizability beyond one platform, and whether a learned allocation rule improves on the production rule [19, 31]. Nothing measured here depends on that rule (§3.1).

7 Conclusion

Content exploration reaches beyond the feed by changing creator participation and replenishing the shared corpus. Standard one-sided viewer tests identify only the direct feed channel. Our decomposition separates a delivered view into immediate value, organic take-up, and induced creator supply. The latter two pay off through the shared corpus. Horizon blindness bounds what any experiment or the mechanism’s own gate can observe of that deferred value. Relative to a minimal floor, production exploration raises videos posted per creator by 8.55% and creators posting at least once by 7.10%. At a matched nominal budget, reallocation raises creator participation with no detectable short-run viewer-side change. The viewer ablation raises video views while lowering view time, a mixed direct trade-off. These outcomes use different units, and total viewer value remains unidentified because one-sided designs cancel the corpus term. The co-diverted probe reaches that term, but its sign remains unresolved under both reported turnover scenarios. Evaluating exploration as an ecosystem intervention requires matching the randomization to the target channel and planning the horizon from corpus turnover.

References

- [1] Himan Abdollahpour, Robin Burke, and Bamshad Mobasher. 2019. Managing Popularity Bias in Recommender Systems with Personalized Re-Ranking. In *Proceedings of the Thirty-Second International Florida Artificial Intelligence Research Society Conference (FLAIRS)*. AAAI Press, 413–418.
- [2] Susan Athey, Raj Chetty, Guido W. Imbens, and Hyunseung Kang. 2026. The Surrogate Index: Combining Short-Term Proxies to Estimate Long-Term Treatment Effects More Rapidly and Precisely. *Review of Economic Studies* 93, 4 (2026), 2284–2312. doi:10.1093/restud/rdof087
- [3] Santiago R. Balseiro and Yonatan Gur. 2019. Learning in Repeated Auctions with Budgets: Regret Minimization and Equilibrium. *Management Science* 65, 9 (2019), 3952–3968. doi:10.1287/mnsc.2018.3174
- [4] Guillaume W. Basse, Yi Ding, and Panos Toulis. 2023. Minimax Designs for Causal Effects in Temporal Experiments with Treatment Habituation. *Biometrika* 110, 1 (2023), 155–168. doi:10.1093/biomet/asac024
- [5] Hemant K. Bhargava. 2022. The Creator Economy: Managing Ecosystem Supply, Revenue Sharing, and Platform Design. *Management Science* 68, 7 (2022), 5233–5251. doi:10.1287/mnsc.2021.4126
- [6] Jennifer Brennan, Yahu Cong, Yiwei Yu, Lina Lin, Yajun Peng, Changping Meng, Ningren Han, Jean Pouget-Abadie, and David Holtz. 2025. Reducing Symbiosis Bias Through Better A/B Tests of Recommendation Algorithms. In *Proceedings of the ACM on Web Conference 2025 (WWW ’25)*. ACM, New York, NY, USA, 3702–3715. doi:10.1145/3696410.3714738
- [7] Robin Burke. 2017. Multisided Fairness for Recommendation. *CoRR* abs/1707.00093 (2017). arXiv:1707.00093
- [8] Gordon Burtch, Qinglai He, Yili Hong, and Dokyun Lee. 2022. How Do Peer Awards Motivate Creative Content? Experimental Evidence from Reddit. *Management Science* 68, 5 (2022), 3488–3506. doi:10.1287/mnsc.2021.4040
- [9] Bernard Caillaud and Bruno Jullien. 2003. Chicken & Egg: Competition among Intermediation Service Providers. *The RAND Journal of Economics* 34, 2 (2003), 309–328. doi:10.2307/1593720
- [10] Gaode Chen, Ruina Sun, Yinjie Jiang, Tianxiang Li, Yinlong Dai, Qifan Shi, Xiaojie Qin, Jialong Fu, Piaoyi Chen, Ronggeng Huang, Na Li, Qi Zhang, Jian Liang, Han Li, and Kun Gai. 2025. A Cold-start Recommendation System at Kuaishou Designed from the Short-video Perspective. In *Companion Proceedings of the ACM on Web Conference 2025 (Sydney, NSW, Australia) (WWW Companion ’25)*. ACM, New York, NY, USA, 124–132. doi:10.1145/3701716.3715205
- [11] Minmin Chen, Yuyan Wang, Can Xu, Ya Le, Mohit Sharma, Lee Richardson, Su-Lin Wu, and Ed H. Chi. 2021. Values of User Exploration in Recommender Systems. In *Fifteenth ACM Conference on Recommender Systems (RecSys)*. ACM, 85–95. doi:10.1145/3460231.3474236
- [12] Paul Covington, Jay Adams, and Emre Sargin. 2016. Deep Neural Networks for YouTube Recommendations. In *Proceedings of the 10th ACM Conference on Recommender Systems (RecSys)*. ACM, 191–198. doi:10.1145/2959100.2959190
- [13] Alex Deng, Ya Xu, Ron Kohavi, and Toby Walker. 2013. Improving the Sensitivity of Online Controlled Experiments by Utilizing Pre-Experiment Data. In *Proceedings of the 6th ACM International Conference on Web Search and Data Mining (WSDM ’13)*. 123–132. doi:10.1145/2433396.2433413
- [14] Jean-Marie Dufour. 1997. Some impossibility theorems in econometrics with applications to structural and dynamic models. *Econometrica* 65, 6 (1997), 1365–1387.
- [15] Leon Jay Gleser and Jiunn T Hwang. 1987. The nonexistence of $100(1 - \alpha)\%$ confidence sets of finite expected diameter in errors-in-variables and related models. *The Annals of Statistics* 15, 4 (1987), 1351–1362.
- [16] Olafur Gudmundsson, Bo Zhao, Huayi Liao, Anna Kiyantseva, Sai Xiao, Heath Vinicombe, Mostafa Keikha, Luke DeLuccia, Zihao Chen, Junpeng Hou, WeiJie Jiang, Bhawna Juneja, Andreanne Lemay, Wei-Ting Lin, Keyvan Moghadam, Jiaxing Qu, Zhiqing Rao, and Zhihua Zhang. 2026. PinEqualizer: Full Funnel Content Exploration and Debiasing System at Pinterest. In *Proceedings of the 32nd ACM SIGKDD Conference on Knowledge Discovery and Data Mining, V.2 (KDD ’26)*. ACM, 7279–7288. doi:10.1145/3770855.3818461
- [17] Henning Hohnhold, Deirdre O’Brien, and Diane Tang. 2015. Focusing on the Long-term: It’s Good for Users and Business. In *Proceedings of the 21st ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*. 1849–1858. doi:10.1145/2783258.2788583
- [18] Qinlu Hu, Ao Huang, Ni (Nina) Huang, and Renyu Zhang. 2026. Attention as Investment: Traffic Allocation for Long Tail Creators. doi:10.2139/ssrn.4888995 SSRN Working Paper 4888995.
- [19] Hong Jun Jeon, Songbin Liu, Yuanlong Li, Jie Lyu, Hunter Song, Ji Liu, Peng Wu, and Zheqing Zhu. 2024. Epinet for Content Cold Start. arXiv:2412.04484 [cs.IR]
- [20] Hoyeop Lee, Jinbae Im, Seongwon Jang, Hyunsouk Cho, and Sehee Chung. 2019. MeLU: Meta-Learned User Preference Estimator for Cold-Start Recommendation. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. ACM, 1073–1082. doi:10.1145/3292500.3330859
- [21] Lihong Li, Wei Chu, John Langford, and Robert E. Schapire. 2010. A Contextual-Bandit Approach to Personalized News Article Recommendation. In *Proceedings of the 19th International Conference on World Wide Web (WWW)*. ACM, 661–670. doi:10.1145/1772690.1772758
- [22] Pasquale Lops, Marco de Gemmis, and Giovanni Semeraro. 2011. Content-based Recommender Systems: State of the Art and Trends. In *Recommender Systems Handbook*. Springer, 73–105. doi:10.1007/978-0-387-85820-3_3
- [23] Lorenzo Masoero, Suhas Vijaykumar, Thomas S. Richardson, James McQueen, Ido Rosen, Brian Burdick, Pat Bajari, and Guido Imbens. 2026. Multiple Randomization Designs: Estimation and Inference with Interference. *Journal of the Royal Statistical Society Series B* 88, 3 (2026), 958–977. doi:10.1093/rjssb/qqaf073
- [24] Aranyak Mehta, Amin Saberi, Umesh V. Vazirani, and Vijay V. Vazirani. 2007. AdWords and Generalized Online Matching. *Journal of the ACM* 54, 5, Article 22 (2007). doi:10.1145/1284320.1284321
- [25] Kun Qian and Sanjay Jain. 2024. Digital Content Creation: An Analysis of the Impact of Recommendation Systems. *Management Science* 70, 12 (2024), 8668–8684. doi:10.1287/mnsc.2022.03655
- [26] Jean-Charles Rochet and Jean Tirole. 2003. Platform Competition in Two-Sided Markets. *Journal of the European Economic Association* 1, 4 (2003), 990–1029.

- doi:10.1162/154247603322493212
- [27] Andrew I. Schein, Alexandrin Popescul, Lyle H. Ungar, and David M. Pennock. 2002. Methods and Metrics for Cold-Start Recommendations. In *Proceedings of the 25th annual international ACM SIGIR conference on Research and development in information retrieval*. ACM, 253–260. doi:10.1145/564376.564421
 - [28] Qijie Shen, Yuanchen Bei, Zihong Huang, Jialin Zhu, Keqin Xu, Boya Du, Jiawei Tang, Yuning Jiang, Feiran Huang, Xiao Huang, and Hao Chen. 2025. AliBoost: Ecological Boosting Framework in Alibaba Platform. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, V.2 (KDD '25). ACM, 4827–4838. arXiv:2506.00954 doi:10.1145/3711896.3737188
 - [29] Yuanyuan Shen, Yiren Yan, Wenjie Li, and Chunhui Zhu. 2026. The Price of Isolation: Estimating the Ecosystem Cost of Symmetric Two-Sided A/B Testing. *arXiv preprint arXiv:2608.04432* (2026). arXiv:2608.04432 [cs.IR]
 - [30] Yi Su, Xiangyu Wang, Elaine Ya Le, Liang Liu, Yuening Li, Haokai Lu, Benjamin Lipshitz, Sriraj Badam, Lukasz Heldt, Shuchao Bi, Ed H. Chi, Cristos Goodrow, Su-Lin Wu, Lexi Baugher, and Minmin Chen. 2024. Long-Term Value of Exploration: Measurements, Findings and Algorithms. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining (WSDM '24)*. ACM, 636–644. doi:10.1145/3616855.3635833
 - [31] Adith Swaminathan and Thorsten Joachims. 2015. The Self-Normalized Estimator for Counterfactual Learning. In *Advances in Neural Information Processing Systems 28 (NIPS)*. Curran Associates, Inc., 3231–3239.
 - [32] Ye Tu, Chun Lo, Yiping Yuan, and Shaunak Chatterjee. 2019. Feedback Shaping: A Modeling Approach to Nurture Content Creation. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD '19)*. 2241–2250. doi:10.1145/3292500.3330764
 - [33] Dong Wang, Junyi Jiao, Arnab Bhadury, Yaping Zhang, and Mingyan Gao. 2024. Item Level Exploration Traffic Allocation in Large-Scale Recommendation Systems. arXiv:2505.09033 Presented at the VideoRecSys+LargeRecSys workshop at RecSys '24.
 - [34] Dong Wang, Junyi Jiao, Arnab Bhadury, Yaping Zhang, Mingyan Gao, and Onkar Dalal. 2025. Item-Centric Exploration for Cold Start Problem. In *Proceedings of the Nineteenth ACM Conference on Recommender Systems (RecSys) Industry Track*. ACM, 987–990. doi:10.1145/3705328.3748113
 - [35] Zhengli Wang, Lin Feng, and Zhixi Wan. 2025. Dynamic Traffic Allocation for Revenue Maximization on Creator Economy Platforms. doi:10.2139/ssrn.6419698 Working paper, SSRN 6419698.
 - [36] Tianli Xia. 2025. Supporting Content Creators on Two-Sided Markets: Experimental Evidence from a Short-Form Video Platform. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5118719 SSRN working paper 5118719.
 - [37] Fan Yao, Yiming Liao, Jingzhou Liu, Shaoliang Nie, Qifan Wang, Haifeng Xu, and Hongning Wang. 2024. Unveiling User Satisfaction and Creator Productivity Trade-Offs in Recommendation Platforms. In *Advances in Neural Information Processing Systems 37 (NeurIPS 2024)*, Vol. 37. Curran Associates, Inc., 86958–86984. doi:10.52202/079017-2759

Appendix roadmap. The appendices carry the material the conference submission holds as separate supplementary matter, bound in here so that this version is self-contained. They follow the order of the sections they support. Section A derives the ecosystem dynamics of §3 and proves the information limits stated in §4. Section B states the results omitted from the main text for length and specifies the estimator used in the co-diverted analysis. Section C reports the additional empirical analyses behind §5, together with their inferential status. Section D gives worked numerical illustrations of the theory.

A Supporting Propositions and Proofs

This appendix collects the formal statements summarized in Section 3.2 and Section 4.2. Notation is as in the main text.

A.1 Ecosystem dynamics

The two results of §3–§4 that describe where the ecosystem settles, Proposition 1 and Proposition 2, are proved here in the order they are stated.

PROOF OF PROPOSITION 2. A marginal view raises the content item’s total views by $1 + q'_s(B)V_{\text{org}}$. It does so once directly and again through the take-up probability, since a taken-up item receives a further V_{org} organic views. The marginal exploration view is itself

consumed and is worth η_e . The induced organic views are worth η_o each. Because the model credits all of an item’s views with inducing posting, the same increment induces $\delta_c(1 + q'_s(B)V_{\text{org}})$ submissions from that creator. Each submission is worth v^* , which satisfies $v^* = \eta_e \bar{B} + \eta_o \bar{q}V_{\text{org}} + \mathcal{R}v^*$, so v^* already contains the whole descendant chain and no further multiplier is applied. Collecting the posting term $\delta_c v^*(1 + q'_s V_{\text{org}})$ with the consumption term $\eta_e + q'_s V_{\text{org}} \eta_o$ gives the statement. $\square$

Theorem 1 on horizon blindness and Corollary 1 on the value and measurability duality are stated in Section 4.2. Their proofs are below. The detection-scaling result (Proposition 3) is stated with proof there too.

Scope of the linear model: the discussion behind Remark 2. Three effects are deliberately outside the recursions (2).

First, V_{org} is held at its steady-state level, in the budget and in the corpus, and the two dependences thereby suppressed are both signed. If extra exploration also improves how a content item performs organically *conditional* on take-up, the marginal view earns a further $q(B)V'_{\text{org}}(B) \geq 0$, so Proposition 2 understates it and the take-up channel here is the extensive margin only. If attention is finite, V_{org} instead falls as the corpus grows, which lowers $\bar{V}(B)$ and hence $\mathcal{R}$. That is the negative feedback capacity rationing already produces (Remark 1), with the same consequence, an overstated multiplier and untouched timing statements. The stock-timing bound holds structurally. Since V_{org} enters nowhere in the corpus recursion, whose inflow $q(B)\Lambda_t$ is counted in content items, Theorem 1 uses only the stock’s retention and inflow staying at or above its eventual level, and is unaffected by any time path of V_{org} . Translating that stock bound to consumption retains the main text’s local affine observation model. Second, the posting response g is concave and saturating, a modeling assumption consistent with the field literature. We linearize this response only at its current steady state. Together with a positive intercept and the assumed subcritical slope, concavity gives a unique stable equilibrium. Its local dynamics are governed by the loop gain $\mathcal{R} = g(\bar{V}(B))$, and concavity signs the linearization error. Comparative statics in B overstate the response to budget increases and understate it to decreases, while the recursion in Λ itself is exact at fixed $\bar{V}$. Third, attention is finite. Exploration views displace organic impressions, so a global budget change induces crowd-out that a linear model ignores. An on/off contrast between randomized arms measures the equilibrium gap net of this reallocation. The model also omits the viewer-side loop, in which a richer corpus draws more viewing and hence more deliverable capacity (the corpus-to-viewing link, distinct from the exploration-to-viewing link measured in Section 5.2). Including it adds a nonnegative coupling to M and can only raise its spectral radius, hence the multiplier, so its omission is conservative in the direction of the paper’s claims. The same monotonicity covers network effects generally: further nonnegative feedbacks (views accumulating followers, thicker corpora matching viewers better) add stocks or couplings to the nonnegative M and can only raise its spectral radius, so each slows the slowest component and enlarges the multiplier. The model’s role is to explain the timing and magnitude structure of an on/off contrast. Forecasting levels from micro-parameters alone is outside its scope.

Fixed capacity is self-damping: the argument behind Remark 1. Under fixed aggregate capacity V the pool of Λ content items shares it, so $\bar{V} = V/\Lambda$ and the induced-submission map is $\varphi(\Lambda) = \Lambda g(V/\Lambda)$, the *perspective transform* of g . Write $u = V/\Lambda$ for per-item views. Differentiating,

$$\varphi'(\Lambda) = g(u) - u g'(u).$$

Three consequences. (i) *Monotone and concave.* Concavity of g with $g(0) \geq 0$ gives $g(0) \leq g(u) - u g'(u)$, so $\varphi' \geq g(0) \geq 0$; and φ inherits concavity from g under the perspective transform. (ii) *Unique globally stable equilibrium.* $\Lambda \mapsto \Lambda_0 + \varphi(\Lambda)$ is concave and increasing with $\varphi(0) = 0$, so $\varphi(\Lambda)/\Lambda$ is decreasing and the fixed point is unique; it is stable whenever $\varphi'(\Lambda^*) < 1$, which subcriticality already gives since $\varphi'(\Lambda^*) < \mathcal{R} < 1$ by (iii). (iii) *Strictly smaller gain.* At the same per-item view level $\mathcal{R} = g(u)$, so

$$\varphi'(\Lambda) = g(u) - u g'(u) = \mathcal{R} \left(1 - \frac{u g'(u)}{g(u)} \right),$$

whose bracketed subtrahend is the *local elasticity* of the response at u ; concavity with $g(0) \geq 0$ puts it in $(0, 1)$, so $\varphi'(\Lambda) < \mathcal{R}$ strictly. The elasticity is constant only for a power law: for a saturating response such as the logarithmic one it falls as the view level rises, so the damping is strongest where per-item views are thinnest. Fixed capacity therefore lowers the loop gain and the multiplier $1/(1 - \mathcal{R})$. It leaves ω unchanged, so Theorem 1 and Corollary 1 apply unchanged to the corpus stock, and the timing conclusions hold a fortiori. This is the one negative feedback the paper accounts for explicitly (Remark 2).

PROOF OF COROLLARY 1. The geometric sum gives $\sum_{k \geq 0} \lambda^k = 1/(1 - \lambda) = A(\lambda)$, while setting $\lambda^t = \varepsilon$ gives the expression horizon $t = \ln(1/\varepsilon)/\ln(1/\lambda) = \ln(1/\varepsilon) \tau(\lambda)$. On $(0, 1)$, $1 - \lambda \leq \ln(1/\lambda) \leq (1 - \lambda)/\lambda$; taking reciprocals, the left inequality gives $\tau(\lambda) \leq A(\lambda)$ and the right gives $A(\lambda) - 1 = \lambda/(1 - \lambda) \leq \tau(\lambda)$. For part (ii), the partial sum is $(1 - \lambda^t)/(1 - \lambda)$, which is bounded by t for every $\lambda \in (0, 1)$ and tends to t as $\lambda \rightarrow 1$. $\square$

Full solution behind Theorem 1. Write Λ_{off}^* and K_{off}^* for the post-switch supply and corpus equilibria, $\Delta s = \Lambda^* - \Lambda_{\text{off}}^*$ for the supply gap, and $q_{\text{off}} = q(0)$ for post-switch take-up. The submission recursion solves to $\Lambda_t = \Lambda_{\text{off}}^* + \Delta s (\mathcal{R}_{\text{off}})^t$. Substituting into the corpus recursion and solving the driven linear recursion gives, for $\mathcal{R}_{\text{off}} \neq 1 - \omega$ (the coincident case is the limit of what follows, in which the $(1 - \omega)^t$ term acquires a factor t),

$$K_t = K_{\text{off}}^* + (K^* - K_{\text{off}}^* + \kappa)(1 - \omega)^t - \kappa (\mathcal{R}_{\text{off}})^t, \quad \kappa \triangleq \frac{q_{\text{off}} \Delta s}{1 - \omega - \mathcal{R}_{\text{off}}},$$

which satisfies $K_0 = K^*$. The expressed fraction of the corpus gap is therefore exactly

$$\frac{K^* - K_t}{K^* - K_{\text{off}}^*} = 1 - (1 - \omega)^t - \frac{\kappa}{K^* - K_{\text{off}}^*} \left[(1 - \omega)^t - (\mathcal{R}_{\text{off}})^t \right].$$

The corpus gap satisfies $K^* - K_{\text{off}}^* = (q(B) \Lambda^* - q_{\text{off}} \Lambda_{\text{off}}^*)/\omega \geq q_{\text{off}} \Delta s/\omega$ because $q(B) \geq q(0)$. When $\mathcal{R}_{\text{off}} < 1 - \omega$, the correction coefficient $c \triangleq \kappa/(K^* - K_{\text{off}}^*)$ obeys $0 \leq c \leq \omega/(1 - \omega - \mathcal{R}_{\text{off}})$, which is $O(\omega/(1 - \mathcal{R}_{\text{off}}))$ in the regime $1 - \mathcal{R}_{\text{off}} \gg \omega$. That regime is where the closed form is most readable. At this platform's measured parameters, it is comfortable under the 360-day corpus turnover timescale and marginal under the 36-day timescale, where $1 - \omega$ approaches $\mathcal{R}_{\text{off}}$ and κ and the bound on c both grow. The theorem

itself is unaffected, since the paragraph below derives it without (3). The corpus-mediated fraction is, exactly and in discrete time,

$$F(t) = \left[1 - (1 - \omega)^t \right] - c \left[(1 - \omega)^t - (\mathcal{R}_{\text{off}})^t \right]. \quad (3)$$

All three statements of Theorem 1 follow from (3) directly, with no continuous-time step. *Bound:* for $\mathcal{R}_{\text{off}} \leq 1 - \omega$ both c and the second bracket are nonnegative; in the opposite ordering both change sign together, so the subtracted product is nonnegative either way and $F(t) \leq 1 - (1 - \omega)^t \leq \omega t$ for every t and every admissible c .

The bound without the linear model. That derivation runs through (3) and so through the linearized supply loop, but the bound needs neither. Let the corpus be any stock with per-period retention $1 - \omega$ and inflow I_t , so $K_{t+1} = (1 - \omega)K_t + I_t$, and let I_∞ be the post-switch inflow with $K_\infty = I_\infty/\omega$. Writing $D_t = K_t - K_\infty$ and subtracting the fixed-point identity from the recursion,

$$D_{t+1} = (1 - \omega)D_t + (I_t - I_\infty).$$

If the inflow stays at or above its eventual level, $I_t \geq I_\infty$, then $D_{t+1} \geq (1 - \omega)D_t$, so $D_t \geq (1 - \omega)^t D_0$, and for a switch-off ($D_0 > 0$) the expressed fraction obeys

$$F(t) = 1 - D_t/D_0 \leq 1 - (1 - \omega)^t \leq \omega t.$$

Conditional on $I_t \geq I_\infty$, no property of g is used. The bound therefore covers nonlinear creator responses and supply loops. Only the corpus's constant fractional turnover and the stated inflow condition enter. Both are mild here. An increasing concave self-map with positive intercept and slope below one at its fixed point has a unique fixed point, and iterates converge to it monotonically (Remark 2), which gives $I_t \geq I_\infty$; and the bound is tight, with equality exactly when the inflow steps to I_∞ at once. Turnover may also vary: with rates ω_s the same argument gives $F(t) \leq 1 - \prod_{s < t} (1 - \omega_s) \leq \sum_{s < t} \omega_s$, which is what a drifting platform satisfies. The argument also allows stochastic inflow and turnover. It is an inequality on D_t realization by realization, so the bound holds *pathwise* and the realized gap never exceeds it. A finite-sample *estimate* of the gap can, by sampling error alone, which is why detection is treated separately in Section 4.2. With constant turnover, exceeding the bound requires an inflow path that falls below its eventual level. The monotone concave loop of §3.2 cannot produce that path. The shape does require constant coefficients. The two-term $F(t)$ of (3), the delay range, and Corollary 1 all need the constant coefficients, and Proposition 2 is a gradient and so is local by construction. *The two ends:* at $t = 1$, (3) gives exactly $F(1) = \omega - c(1 - \omega - \mathcal{R}_{\text{off}})$, so c is precisely what is subtracted from the corpus's own first-period response. At $c = 0$ (take-up falls independently of any change in submissions, e.g. the exploration-window signal itself stops sustaining organic take-up at the *same* supply level) $F(1) = \omega$: the pure corpus response, undelayed. At the upper bound $c = \omega/(1 - \omega - \mathcal{R}_{\text{off}})$, attained when the switch moves the supply loop alone ($q(B) = q(0)$, take-up rate unchanged, only submission volume falls) $F(1) = 0$ exactly: this is the case of maximal delay, and the one the theorem's sharper reading names. *Onset at that upper bound:* the first period expresses nothing at all, which is the mechanical reading of the delay: the corpus cannot begin to thin until the submissions feeding it have themselves fallen, and the onset is therefore *quadratic*.

Expanding (3) to first order in ω at $c = c_{\max}$ gives

$$F(t) = \omega \left[t - \frac{1 - (\mathcal{R}_{\text{off}})^t}{1 - \mathcal{R}_{\text{off}}} \right] + O((\omega t)^2),$$

which vanishes at $t = 1$, as $F(1) = 0$ requires, and grows as $\omega(1 - \mathcal{R}_{\text{off}}) t(t-1)/2$ while $(1 - \mathcal{R}_{\text{off}})t \ll 1$. We state the discrete form. The continuous $\omega t^2/(2\tau(\mathcal{R}_{\text{off}}))$ requires both $t \gg 1$ and $\mathcal{R}_{\text{off}} \rightarrow 1$. At $\omega = 0.005$, $\mathcal{R}_{\text{off}} = 0.9$ it overstates $F(2)$ by a factor of two and $F(5)$ by half, even inside the regime $t \ll \tau(\mathcal{R}_{\text{off}})$ where the continuous form is often invoked. *Ramp*: once $(\mathcal{R}_{\text{off}})^t$ is negligible and $\omega t \ll 1$, substituting $(1 - \omega)^t = 1 - \omega t + O((\omega t)^2)$ gives $F(t) = (1 + c)\omega t - c + O((\omega t)^2)$, a line crossing zero at $t_0 = c/[(1 + c)\omega]$ exactly, for every admissible c . The common c/ω expression omits the $(1 + c)$ factor in the slope. At the upper bound $c = \omega/(1 - \omega - \mathcal{R}_{\text{off}})$, $1 + c$ simplifies to $(1 - \mathcal{R}_{\text{off}})/(1 - \omega - \mathcal{R}_{\text{off}})$ and t_0 to $1/(1 - \mathcal{R}_{\text{off}}) = A(\mathcal{R}_{\text{off}})$ exactly, with no approximation beyond the linearization already made. The inferred delay is one supply-loop multiplier whatever ω is relative to $1 - \mathcal{R}_{\text{off}}$. The intercept is an extrapolation. At $\omega = 0.005$, $\mathcal{R}_{\text{off}} = 0.9$ the exact $F(A(\mathcal{R}_{\text{off}})) = 0.017$, almost exactly the fast term $c(\mathcal{R}_{\text{off}})^{A(\mathcal{R}_{\text{off}})} = 0.018$ that the ramp discards. At $c = c_{\max}$ the only zeros of F are $t = 0$ and $t = 1$. Only the ramp's *slope* away from that crossing, $\omega(1 - \mathcal{R}_{\text{off}})/(1 - \omega - \mathcal{R}_{\text{off}})$, needs $\omega \ll 1 - \mathcal{R}_{\text{off}}$ to reduce to the bare ω of the theorem's statement; under that regime $F(t) = \omega(t - A(\mathcal{R}_{\text{off}}))(1 + o(1))$, which Corollary 1 places within one period of $\tau(\mathcal{R}_{\text{off}})$. Replacing the discrete factors in (3) by $e^{-\tilde{\omega}t}$ and $e^{-\alpha t}$, with rates $\tilde{\omega} = \ln(1/(1 - \omega))$ and $\alpha = \ln(1/\mathcal{R}_{\text{off}})$, recovers the familiar continuous-time form $1 - [\alpha e^{-\tilde{\omega}t} - \omega e^{-\alpha t}]/(\alpha - \omega)$; since $1 - \lambda \leq \ln(1/\lambda) \leq (1 - \lambda)/\lambda$, the two agree to first order in ω and $1 - \mathcal{R}_{\text{off}}$, and we keep the discrete form because in it the delay appears directly as the theorem's $A(\mathcal{R}_{\text{off}})$, where the continuous form returns $\tau(\mathcal{R}_{\text{off}})$. The consumption gap adds the immediate direct-channel step and inherits this time path for its corpus-mediated component, since consumption is linear in the corpus stock.

PROOF OF PROPOSITION 4. (i) A user's consumption in arm a decomposes as $E_t^a = \eta_e X_t^a + \eta K_t$, using §3.2's affine consumption map, where the exploration delivery X_t^a differs by arm and organic consumption depends on the corpus state K_t , a single stock driven by aggregate delivery and submissions. Enrollment perturbs those aggregates by $O(p)$, but K_t enters both arms identically, so the perturbation subtracts out and the contrast equals $\eta_e(X_t^{\text{ctrl}} - X_t^{\text{abl}}) = \Delta_{\text{direct}}$ for every t and N . (ii) Each creator population follows its own supply recursion, $\Lambda_{t+1}^a = \Lambda_0 + g(\tilde{V}^a) \Lambda_t^a$, whose eigenvalue is that arm's loop gain $\mathcal{R}^a = g(\tilde{V}^a)$ and where $\tilde{V}^a$ is views per item under that arm's budgets; aggregate viewer attention moves the two $\tilde{V}^a$ together by $O(p)$. The per-creator posting contrast at equilibrium is therefore $\Lambda_0[(1 - \mathcal{R})^{-1} - (1 - \mathcal{R}_f)^{-1}] + O(p)$. The consumption generated by the induced corpus accrues to viewers of both arms through the shared K_t and cancels, as in (i). The $O(p)$ term has a determinate sign: the above-floor budget freed by flooring returns to the shared capacity pool and redistributes over the remaining $1 - p$ share, so per-item delivery in the production arm rises relative to full rollout by a factor of order $p/(1 - p)$ times the reallocated budget share. Since g is increasing, this inflates the measured contrast, and $\epsilon(p) = O(p/(1 - p))$ with $\epsilon(p) \geq 0$. (iii) Within an isolated submarket the full two-component system runs

closed, so the corpus stock expresses in the cell contrast on its own timescale (Theorem 1). The final claim follows because the contrasts in (i) and (ii) contain no corpus term at any t or N . $\square$

A.2 Information Limits for Valuing the Corpus Multiplier

This section proves the information limits stated in Section 4.2 (Theorem 3 and Corollary 2) and records the noise model behind them. In the co-diverted application of §5.5, the level is the direct channel and the asymptote the corpus term, so we write the gap series as

$$y_t = \Delta_{\text{direct}} + \Delta_{\text{corpus}}(1 - e^{-\omega t}) + \eta_t, \quad t = 1, \dots, T,$$

with N users per arm. The estimand is Δ_{corpus} ; Δ_{direct} and the turnover ω are unknown nuisance parameters.

PROOF OF LEMMA 1. Write the transient covariance of the arm-mean series as $\Sigma_0 = \frac{2}{N}\Gamma$ with $\Gamma_{st} = \gamma(s - t)$, and the full covariance as $\Sigma_a = \Sigma_0 + aJ$, where $J = \mathbf{1}\mathbf{1}^\top$ and $a = \frac{2}{N}\sigma_\alpha^2$. (i) With the intercept free, every likelihood ratio and information quantity below depends on the data law only through the profiled precision of its covariance, $P_\Sigma = \Sigma^{-1} - \Sigma^{-1}\mathbf{1}(\mathbf{1}^\top \Sigma^{-1}\mathbf{1})^{-1}\mathbf{1}^\top \Sigma^{-1}$ for generic Σ . Let $s = \mathbf{1}^\top \Sigma_0^{-1}\mathbf{1} > 0$. Sherman-Morrison gives $\Sigma_a^{-1} = \Sigma_0^{-1} - \frac{a}{1+as} \Sigma_0^{-1} J \Sigma_0^{-1}$, hence $\Sigma_a^{-1}\mathbf{1} = \Sigma_0^{-1}\mathbf{1}/(1 + as)$ and $\mathbf{1}^\top \Sigma_a^{-1}\mathbf{1} = s/(1 + as)$. Substituting,

$$P_{\Sigma_a} = \Sigma_0^{-1} - \left[\frac{a}{1+as} + \frac{1}{s(1+as)} \right] \Sigma_0^{-1} J \Sigma_0^{-1} = \Sigma_0^{-1} - \frac{1}{s} \Sigma_0^{-1} J \Sigma_0^{-1} = P_{\Sigma_0},$$

identically in a : the profiled model is the same for every σ_α^2 . For unprofiled quantities, $\Sigma_a \succeq \Sigma_0$ in the Loewner order, so $\Sigma_a^{-1} \preceq \Sigma_0^{-1}$ and information can only shrink, enlarging the impossibility region. (ii) For $\text{AR}(1)$, $\Gamma = \sigma^2[\rho^{|s-t|}]$ admits the exact decomposition

$$m^\top \Gamma^{-1} m = \frac{(1 - \rho)^2 \sum_{t=1}^T m_t^2 + \rho \sum_{t=2}^T (m_t - m_{t-1})^2}{\sigma^2(1 - \rho^2)} + \frac{\rho(1 - \rho)(m_1^2 + m_T^2)}{\sigma^2(1 - \rho^2)}.$$

verified by collecting coefficients against the tridiagonal inverse. The first term is $\sigma_{\text{eff}}^{-2} \sum_t m_t^2$ with $\sigma_{\text{eff}}^2 = \sigma^2(1 + \rho)/(1 - \rho)$, and both corrections vanish at $\rho = 0$. For the polynomial regressors underlying S_k , where $|m_t - m_{t-1}| \leq C|m_t|/t$ and $m_T^2 = O(\sum_t m_t^2/T)$, the increment term is a relative $O(\rho/((1 - \rho)^2 T^2))$ and the boundary term a relative $O(\rho/((1 - \rho)T))$. Every statement therefore holds with σ^2 replaced by σ_{eff}^2 up to a relative $O(\rho/((1 - \rho)T))$. (iii) Under random-walk drift, $\Gamma_{st} = \sigma_w^2 \min(s, t)$, and first differencing whitens the noise while reducing the curve to its increments. The ramp's information then accumulates from increments of the quadratic onset, $\sum_{t \leq T} t^2 \asymp T^3$, in place of $S_4 \asymp T^5$; equivalently, level-plus-random-walk sample paths span slowly varying trajectories that mimic the ramp, leaving identification to the increments alone. The diagnostic in the statement follows since under stationarity the variance of user-mean-centered outcomes is bounded in T , while under a unit root it grows linearly. $\square$

In what follows η_t is accordingly taken i.i.d. $N(0, s^2)$ with $s^2 = 2\sigma_{\text{eff}}^2/N$, and $S_k \triangleq \sum_{t \leq T} t^k$.

PROOF OF THEOREM 3. Let the true parameters be $\Delta_{\text{direct}}, \Delta_0$ (the statement's true Δ_∞) and ω_0 , and write $1 - e^{-x} = x\phi(x)$ with $\phi(x) = (1 - e^{-x})/x$, $\phi(0) = 1$. Then $0 \leq -\phi'(x) \leq \frac{1}{2}$ for all $x \geq 0$: $-x^2\phi'(x) = 1 - (1+x)e^{-x}$, whose derivative is $xe^{-x} \in [0, x]$, so $0 \leq 1 - (1+x)e^{-x} \leq x^2/2$. For $M \in (\Delta_0, \infty]$ consider the *matched-product* alternative $(\Delta_{\text{direct}}, M, \omega_M)$ with $\omega_M = \omega_0\Delta_0/M$; the case $M = \infty$ is the linear curve $\Delta_{\text{direct}} + \Delta_0\omega_0 t$. With $\beta = \Delta_0\omega_0$, the mean curves differ by

$$|f_M(t) - f_0(t)| = \beta t |\phi(\omega_M t) - \phi(\omega_0 t)| \\ \leq \frac{1}{2} \beta t^2 (\omega_0 - \omega_M) \leq \frac{1}{2} \Delta_0 \omega_0^2 t^2,$$

uniformly in M . Hence $D_M^2 = \sum_t (f_M - f_0)^2 \leq \frac{1}{4} \Delta_0^2 \omega_0^4 S_4 = C^2 s^2 \leq s^2$, so for the product Gaussian measures $\text{KL}(P_0 \| P_M) = D_M^2 / (2s^2) \leq \frac{1}{2}$ and $\text{TV} \leq \sqrt{\text{KL}/2} \leq \frac{1}{2}$, simultaneously for every M . Coverage at θ_M transfers: $P_0(M \in \text{CI}) \geq P_M(M \in \text{CI}) - \text{TV} \geq \frac{1}{2} - \alpha$. The events $\{U \geq M\} \supseteq \{M \in \text{CI}\}$ decrease to $\{U = +\infty\}$ as $M \uparrow \infty$, and monotone convergence gives the claim. $\square$

The certified constant is conservative in two removable ways. Replacing Pinsker's inequality with the exact Gaussian total variation, $\text{TV}(P_0, P_M) = 2\Phi(D_M/(2s)) - 1$ with Φ the standard normal c.d.f., certifies the same conclusion up to $C \leq 2\Phi^{-1}(3/4) \approx 1.35$; and allowing the alternative to re-profile the level and slope shrinks D_M further, since the leading difference $\propto t^2$ is far from orthogonal to $\text{span}\{1, t\}$. Both enlargements extend the impossibility region, so the theorem's direction is unaffected; the calibrated Monte Carlo of Figure 3 accordingly places the practical estimability transition roughly an order of magnitude above the certified threshold.

PROOF OF COROLLARY 2. (i), lower bound: choose M in the proof of Theorem 3 so that $D_M = s$. The pair is then separated by

$$|M - \Delta_0| \gtrsim \frac{s}{\omega^2 \sqrt{S_4}}, \\ |M - \Delta_0| \asymp \frac{s}{\omega^2 T^{5/2}}.$$

while $\text{TV} \leq \frac{1}{2}$, and the standard two-point reduction converts indistinguishability into the risk bound. For (i)'s attainment claim: expand the mean curve's derivatives to second order in ωt and profile the intercept and ω . The Gaussian Fisher information then leaves

$$I^*(\Delta_{\text{corpus}}) = \frac{\omega^4 T^5}{720s^2} (1 + O(\omega T)).$$

The Gaussian regression has an analytic mean and nonsingular information at every interior point with $\omega > 0$, so standard parametric theory gives asymptotic attainment by profile maximum likelihood. Inverting I^* for a target relative error gives the $N^{-1/5}$ law. (ii) Delta method on $\Delta_{\text{corpus}} = \beta/\omega$, with $\hat{\beta}$ the intercept-profiled slope: the slope's information is $\omega^2 T^3 / (12s^2)$ to leading order, so the error ratio between (ii) and the known- ω problem is $\sqrt{720/12}/(\omega T) \approx 7.7/(\omega T)$, which the external estimate removes. $\square$

REMARK 3 (ESTIMATOR CHOICE). A numerical check at $\omega T \leq 0.3$ confirms that profile likelihood attains the bound at the probe's scale. The quadratic plug-in estimator $\hat{\Delta} = \hat{u}^2 / (2\hat{v})$ from a polynomial fit is truncation-biased at fixed ωT . Wherever the single-rate form is fit, profile likelihood is used for this reason. The probe's primary analysis

avoids the single-rate form entirely, regressing on the measured age-yield kernel and the posting-path convolution (§5.5).

B Omitted Results and the Co-Diverted Estimator

This appendix states the results omitted from the main text for length and details the estimator used in the co-diverted analysis. Proofs are in Section A. It is self-contained given the notation of §3.1.

B.1 Finite-horizon information limits

The following two results are stated without proof in §4.2.

Detecting a nonzero finite-horizon effect is easier than estimating its asymptote. Let y_t be the observed between-arm contrast in a corpus-stock metric at horizon t . Suppose it has an unknown level plus a single slow component,

$$y_t = c + \Delta_\infty (1 - e^{-\omega t}) + \varepsilon_t, \quad t = 1, \dots, T,$$

with ε_t i.i.d. $N(0, s^2)$, $s^2 = 2\sigma_{\text{eff}}^2/N$ for N users per arm, where σ_{eff} is the effective per-user transient-noise scale (Lemma 1, Appendix A.2, which also shows that persistent user heterogeneity costs nothing once the level is profiled), and $S_4 = \sum_{t \leq T} t^4$. Both the level c and the rate ω are unknown nuisance parameters.

THEOREM 3 (NO FINITE UPPER CONFIDENCE BOUND BELOW THE CURVATURE THRESHOLD). Define the curvature signal-to-noise ratio $C = \Delta_\infty \omega^2 \sqrt{S_4} / (2s)$. If $C \leq 1$, then any confidence procedure for Δ_∞ with uniform coverage $1 - \alpha$ over positive amplitudes and rates satisfies, under the true parameters,

$$P(\text{upper endpoint} = +\infty) \geq \frac{1}{2} - \alpha.$$

COROLLARY 2 (HORIZON-SAMPLE LAW FOR THE ASYMPTOTE). (i) Any estimator of Δ_∞ errs by $\gtrsim s/(\omega^2 T^{5/2})$ on some parameter within total variation $\frac{1}{2}$ of the truth. To leading order in ωT this rate is attained by profile likelihood, so a fixed relative-error target requires $T \propto N^{-1/5}$: a tenfold increase in traffic shortens the required horizon by a factor of 1.6. (ii) With ω estimated from auxiliary data to relative precision ρ_ω , $\Delta_\infty = \beta/\omega$ inherits the slope's law: relative error $\approx \sqrt{\text{relerr}(\hat{\beta})^2 + \rho_\omega^2}$, the horizon law improves to $T \propto N^{-1/3}$, and an error inflation of order $1/(\omega T)$ is removed.

Proofs are in Appendix A.2. The theorem uses an elementary two-point argument along the family of curves with matched product $\Delta_\infty \omega$, which includes the straight line (the asymptote- ∞ limit) within total variation $\frac{1}{2}$ of the truth whenever $C \leq 1$. For the ecosystem, the slow component is the corpus stock: the regimes in which exploration is most valuable (Corollary 1) are those in which its asymptotic value can, at feasible horizons, admit only an increment lower bound. §5.5 runs exactly this estimation problem with a horizon-matched endpoint.

B.2 The estimator used in the co-diverted analysis

The results above analyze the design-stage problem, where the corpus timescale is unknown. The estimator used in §5.5 instead pins the response shape from historical platform data. For horizon




Figure 3: Finite-upper-endpoint probability for a 95% profile procedure in semi-synthetic Monte Carlo at the measured noise scale. Blue is the exact single-rate family of Theorem 3. Its $C \leq 1$ region obeys the theorem’s $\frac{1}{2} + \alpha$ cap. Orange uses the measured age–yield kernel under free timescale dilation. It shares the horizontal coordinate through the kernel’s 36-day single-rate equivalent. At the probe’s 21-day window ($C \approx 2$), its finite-endpoint probability is 0.21. Markers give each curve’s median crossing.

t in the co-diverted contrast, the model is

$$y_t = c + \delta_t + \Delta_{\text{direct}} \mathbf{1}\{t \geq 0\} + \Delta_{\text{corpus}} (k_{\text{emp}} * \pi)(t) + \varepsilon_t,$$

where k_{emp} is the measured age–yield kernel, π the measured in-cell posting path normalized to unit long-run level, δ_t the drift imported from the concurrent one-sided ablation (§5.5), and ε_t independent with the per-estimate sampling variances. Writing F_{emp} for the kernel’s normalized cumulative response, the eleven monthly cross-sections give

$$F_{\text{emp}}(7) = 0.471, \quad F_{\text{emp}}(21) = 0.583, \quad F_{\text{emp}}(90) = 0.755.$$

These values display the front-loaded, non-exponential shape behind the 36-day single-rate equivalent used in the main text.

Interpreting k_{emp} as the intervention response uses four assumptions. (A1) A change in exploration delivery scales cohort inflow without changing the normalized age–yield shape. (A2) Ranker competition between cohorts leaves that shape unchanged at the probe’s perturbation size. (A3) The historical cross-sections transport to the cells. Both run the production ranker, and the dispersion across the eleven cross-sections enters the variance below. (A4) The observed in-cell posting path is the relevant intervention input. Once the ramp is estimable, (A1)–(A2) can be assessed by comparing the cell’s own age profile with k_{emp} .

With k_{emp} and π pinned and δ_t imported, the model is linear in the intercept c and the two coefficients Δ_{direct} and Δ_{corpus} . The estimator is weighted least squares under the sampling variances. The step and ramp regressors are linearly independent once three or more post-onset estimates exist, which gives algebraic identification. Reporting also requires the stability and interval checks described

below. Uncertainty combines three components, summed as variances under independence across sources. (i) Sampling error uses the per-estimate variances, with estimates over non-overlapping windows treated as independent. Shared-user serial correlation is largely removed by the pre-period differencing, and as a check the variances are inflated by the estimated lag-one autocorrelation of the residual series. (ii) Kernel uncertainty treats the eleven cross-sectional kernels as draws: the fit is repeated per kernel and the across-kernel variance of $\hat{\Delta}_{\text{corpus}}$ is added. (iii) Drift enters as a Gaussian prior centered on the imported series with its reported margins, propagated through the weighted fit. Coverage of the combined interval is checked on synthetic data with known truth under the specified procedure. In 300 replicates of each nondegenerate scenario, combined-gap coverage is 93.7% (Monte Carlo s.e. 1.4 points), while the lower-bound construction is valid in 96–97%. If the posting path is exactly zero, the supply regressor is dropped and coverage falls to 89.0%. The reported analysis uses the nonzero-path specification, whose kernel–supply regressor correlation is 0.994. This nominally passes the pre-specified collinearity cutoff of 0.995, fixed before estimation. Moreover, the primary interval includes zero and large offsetting step and ramp coefficients expose intercept–kernel near-collinearity. The pre-specified joint rule rejects a two-sided ceiling estimate. The main text reports the well-conditioned robustness intervals instead. The estimation procedure distinguishes two targets. Conditional on (A1)–(A4), the WLS coefficient is a two-sided estimate of Δ_{corpus} . The shape-robust companion uses an increment. An envelope regressor would leave the coefficient unbounded once the direct-step nuisance is profiled. For horizons $t_1 < t_2$ after the exploration window, the drift-adjusted increment has mean $\Delta_{\text{corpus}} [F(t_2) - F(t_1)]$: the direct step and the level cancel from the difference, and under the stock model F is nondecreasing with $F \leq 1$, so $0 \leq F(t_2) - F(t_1) \leq 1$ for any admissible response shape. The increment has the same sign whenever the stock changes and never exceeds $|\Delta_{\text{corpus}}|$. The nonnegative-stock model imposes $\Delta_{\text{corpus}} \geq 0$ on the gross-flow target (Proposition 2). This direction was fixed before the ramp was estimable, and we use it only for the model-restricted lower bound. The policy-relevant net asymptote in the main text is reported without this restriction and remains unsigned. Under the directional restriction, the robust bound is the lower endpoint of the one-sided 95% interval on the drift-adjusted increment. If the estimate contradicts that direction, we report the two-sided 95% interval and claim a nonzero corpus term only if it excludes zero. The bound costs sharpness, paying $1 - [F(t_2) - F(t_1)]$ of the asymptote, but assumes nothing about the shape beyond monotonicity.

C Additional Empirical Analyses

C.1 The Observed Viewer Time Path

Section 5.2 rests part of its identification argument on the shape of the viewer ablation over time. The main text summarizes that path in Figure 2. Figure 4 shows seven cumulative estimates in faint lines and their six non-overlapping period differences in solid lines. The final cumulative estimates reproduce the video-view and view-time entries in Table 2, Panel A.

Three features matter, and one of them is easy to misread from the cumulative series alone.

No sign flip or ramp. Each series holds its sign across all seven estimates, and neither exhibits a sustained upward ramp. A growing corpus contribution would build monotonically on the turnover timescale (Theorem 1). Nothing here does. This is the feature Section 5.2 uses to read the gap as the direct, per-user effect.

Cumulative smoothing hides attenuation. Video views rise from +2.02 to a peak of +2.32 around month seven, then fall to +1.74. Read this way the excursion looks small, 0.58 percentage points against effects of 1.7 to 2.3. A cumulative reading necessarily looks small because each estimate contains every earlier one. Consecutive points are constrained to agree whatever the underlying effect is doing, making a cumulative series the wrong instrument for a claim about shape. Figure 4 differences the same seven estimates into six non-overlapping periods, which share no data.

$$\Delta_{(t_1, t_2]} = \frac{c_2 t_2 - c_1 t_1}{t_2 - t_1},$$

exact when the cumulative denominator grows in proportion to elapsed time, i.e., per-period exposure is stable. With balanced arms of fixed membership and equal-length windows the deviation is second order. Without covariances between nested cumulative estimates, the plotted bands scale the cumulative margins under independent per-period noise, which overstates increment variance when period-level variances are stable. We label the bands approximate. On that reading video views hold between +2.27 and +2.45 for the first three periods and then fall to +0.13, and view time deepens to -2.87 before attenuating to -1.14. The movement is several times larger than the cumulative series suggests.

Both shapes lie outside the corpus contribution, which a one-sided design cannot express at any horizon (Proposition 4(i)), so the reading of §5.2 is unaffected. The path is a stable regime followed by attenuation, which coincides with a decline in exploration’s share of feed consumption over the same span. An ablation’s effect size scales with the share of delivery it removes, and that share is an operating choice. These figures report an operating-point price.

Precision without convergence. The confidence bands narrow monotonically as the cumulative sample grows, from ± 0.23 to ± 0.16 percentage points on video views. The point estimate moves by more than that narrowing over the same window, which is Proposition 3(i) in miniature. More data measures the system’s state at the time of measurement more precisely without pinning down a quantity that is itself drifting.

Figure 5 repeats the period reading for the concurrent creator ablation’s *consumption* series, the second, independently randomized witness of the operating-point shift cited in §5.2. Period effects on video views hold near +1% up to the same boundary (+1.11% in the final period before it) and turn negative after.

C.2 Budget-Matched Reallocation

A separate randomized co-diverted experiment tests reallocation (§5.4). As in §5.5, viewers and creators were assigned to matched isolated cells. Intervals condition on these realized cells. The design holds the nominal total budget fixed and spreads exploration across a broader set of videos. Promotion thresholds implement the shift, and a spend comparison on the penultimate day finds realized budget per creator matched within approximately 3%, the




Figure 4: The same seven cumulative estimates differenced into six non-overlapping periods (solid, with approximate bands), against the cumulative reading (faint). The dashed vertical marker is the operating-point shift both experiments turn at (§5.2). Positive = higher under production. Consecutive periods share no data, so their shape is the effect’s shape. The cumulative series is smooth by construction. Period bands are approximate upper bounds derived from the cumulative margins under independent per-period noise, scaling the cumulative margin by $\sqrt{t_2/(t_2 - t_1)}$. No closed-form margin for nested increments is available from the per-horizon estimates.




Figure 5: The creator ablation’s consumption series read per period (solid, approximate bands) against its cumulative reading (faint), as in Figure 4. The dashed marker is the same operating-point boundary both experiments turn at. Period effects on video views are positive up to it and negative after, the coincident break cited in §5.2.

reallocated arm slightly lower, so any residual dose imbalance biases its posting response downward. Snapshot video counts differ because reallocation changes pool residence. This one-day diagnostic bounds the observed residual. Window-wide equivalence remains unverified, so the estimate is the randomized arm contrast, leaving pure breadth elasticity open. The experiment runs for two weeks, a horizon matched to the fast supply margin it targets. The production-versus-reallocation contrast and the two-week horizon belong to the original experiment design, which was run independently of this study. A second planned reallocation arm is omitted because a configuration audit showed that it did not implement the intended budget-matched contrast. The intervention configuration determined this exclusion before the outcomes were considered. Over the full window, creators posting at least once

increase by $0.77\% \pm 0.42$, and videos posted per creator increase by $0.92\% \pm 0.87$. Viewer-side consumption shows no detectable change (view time $+0.01\% \pm 0.15$, video views $-0.06\% \pm 0.15$, and video favorites $+0.15\% \pm 0.59$). The two-week window captures only the fast margin. By Theorem 1, the corpus consequence of the induced supply is attenuated and its asymptote unidentified in this window. The result identifies the direction of the short-horizon trade-off. Long-run value and the allocation rule remain open.

C.3 Feasible-Horizon Corpus Flow in the Co-Diverted Cells

Splitting each co-diverted cell's consumption by video posting cohort and video age at viewing isolates organic consumption of videos posted during the experiment after their exploration windows. By construction, exploration can no longer deliver these videos. Among videos posted by creators assigned to the same cell as the viewers, production raises organic view time by 37.9% [37.1, 38.7] relative to the floored cell and raises video views by 41.7%. Videos posted before the switch show no effect once the exploration-window period is excluded. All subsequent estimates for pre-experiment videos are within $\pm 3\%$. The pooled reading that includes the exploration-window period falls outside that band. We report both versions. Figure 6 plots the daily series against the placebo band. Video favorites per viewer on in-cell videos rise 29.8% [24.3, 35.3]. Per view, marginal corpus videos receive $\approx 0.91\times$ the incumbent viewing duration and have $\approx 0.71\times$ the incumbent favorite rate. Both are marginal ratios of the outcome change to the change in video views, relative to the incumbent per-view rate. The affected slice consists of in-cell videos posted during the experiment and viewed after their exploration windows. It carries 0.3–0.5% of a cell's consumption at this horizon. Together, all in-cell videos account for 8–14% of a cell's consumption. The $+37.9\%$ effect on it moves the aggregate contrast by roughly 0.1–0.2 percentage points, a magnitude consistent with the whole-cell null.

Inferential status. The decomposition was discovered in an exploratory analysis of a 1% user sample. Applying the same endpoints, estimators, and windows to the full population reproduces the numerical pattern without creating an independent confirmatory test. The pooled pre-experiment cohort also misses the placebo band when the exploration-window period is included. Excluding that period, during which the intervention can still affect delivery, restores the placebo pattern. We report both versions and retain exploratory status. A shape-based step-plus-ramp decomposition of the net term is also underidentified because its collinearity criterion fails at the available horizon, so we report a bound. Accordingly, the main text uses this decomposition to measure feasible-horizon gross corpus flow. The sign of aggregate corpus value remains unresolved.

D Worked Illustrations

The paper states its results as bounds and scaling laws. This appendix works four of them through numerically at one illustrative parameter setting and gives each an operational reading. Every parameter is chosen for legibility and differs from production values. The parameters stated below determine every figure in this section.




Figure 6: In-cell organic video views and view time for videos posted during the experiment and viewed after their exploration windows, production versus floored cell, daily with 95% intervals (user-level sample). The pooled full-population view-time contrast is $+37.9\%$ [37.1, 38.7]. The gray band is the placebo interval from videos posted before the experiment.




Figure 7: Theorem 1's delay family, from (3) (§A) at $\omega = 0.005$, $\mathcal{R}_{\text{off}} = 0.9$. Four values of the interpolating coefficient, $c/c_{\text{max}} \in \{0, \frac{1}{3}, \frac{2}{3}, 1\}$, light to dark. (a) The onset differs. At $c = 0$ the corpus responds at rate ω from the first period. At $c = c_{\text{max}}$ the first period expresses nothing, and the ramp extrapolates back to zero at exactly $A(\mathcal{R}_{\text{off}})$. (b) The same four curves over the corpus timescale, where the distinction has washed out.

The running example. One period is a submission-feedback cycle. We take corpus turnover $\omega = 0.005$ per period, so the corpus half-life is $\ln 2 / \ln(1/(1-\omega)) \approx 138$ periods and $\tau_\omega = 1/\omega = 200$. We use a post-switch loop gain $\mathcal{R}_{\text{off}} = 0.9$, so one supply-loop multiplier is $A(\mathcal{R}_{\text{off}}) = 1/(1-\mathcal{R}_{\text{off}}) = 10$ periods. The two timescales are separated by a factor of 20, the regime §3.2 assumes.

D.1 Theorem 1: the cost of delay

Theorem 1 leaves one coefficient free, $c \in [0, \omega/(1-\omega-\mathcal{R}_{\text{off}})]$. How much turns on it? At the running parameters $c_{\text{max}} = 0.0526$, and the answer splits by horizon.

Figure 7(a) shows the first 40 periods. The two ends have genuinely different shapes. At $c = 0$, $F(1) = \omega$ exactly and the response is linear from the start. At $c = c_{\text{max}}$, $F(1) = 0$ exactly, the onset is quadratic, and the ramp extrapolates back to zero at $t_0 = 10.000$ periods, which is $A(\mathcal{R}_{\text{off}})$ to four decimals, as the algebra behind (3) in §A requires. That intercept is what a fitted ramp reports. The




Figure 8: Corollary 1: the multiplier a feedback loop delivers and its relaxation time coincide to within one period. Both are plotted against a single axis. A target expression horizon adds the factor set by that target. In particular, $t_{90\%} = \ln(10) \tau(\lambda)$.

exact $F(10) = 0.017$, essentially the fast term the ramp discards. A three-week window would report materially different numbers depending on where a platform sits.

Figure 7(b) shows the same four curves out to 800 periods. By τ_ω they agree to within 1.9% of the asymptote, and by $t = 800$ to within 0.1%. The delay is a transient on the fast timescale. The approach to the asymptote is governed by ω alone.

Planning consequence. Two consequences point in opposite directions. Planning requires no estimate of c . The horizon is set by ω , estimable from turnover data without running the experiment. Reading a short experiment does require it, since the attenuation factor depends on c . A short window's estimate is attenuated by an amount the model leaves unspecified, which is why §5.5 fits the delayed onset from data.

D.2 Corollary 1: reading value off the horizon axis

Corollary 1 says $A(\lambda)$ and $\tau(\lambda)$ coincide to within one period. Figure 8 makes the coincidence visible: the gap $A - \tau$ is 0.56 periods at $\lambda = 0.5$ and 0.50 at $\lambda = 0.99$, never exceeding one, with the ratio approaching 1.

A loop with $\lambda = 0.9$ delivers a 10× multiplier and takes 9.5 periods to express 63% of a switch, 21.9 to express 90%. A loop with $\lambda = 0.99$ delivers 100× and takes 99.5 and 229 periods respectively.

Feasibility check. Figure 8 doubles as a feasibility check on a research plan. Read the multiplier a proposal hopes to demonstrate off the vertical axis, find the λ that delivers it, and read off $t_{90\%}$. That is the shortest honest study. A plan to establish a 50× ecosystem multiplier inside one turn of the corpus is internally inconsistent, because the multiplier is achieved by the slowness. This is also the argument for the prescription of §1: since the multiplier and the delay cannot be separated, a mechanism that must decide quickly has to *import* the slow terms. The remaining question is whether the offline estimate is good enough.




Figure 9: Proposition 3 and Corollary 2: required horizon against traffic per arm, both relative, for the three exponents the paper derives. Markers give the traffic multiple that halves the horizon.




Figure 10: Proposition 2's three channels as shares of the marginal value of a view, against loop gain $\mathcal{R}$. Illustrative units ($\eta_e = \eta_o = 1$, $V_{\text{org}} = 10$, $\bar{q} = 0.1$, $\bar{B} = 100$, $q' = 0.01$, $\delta = 0.01$). The shares depend on these. The monotonicity is structural. Stack order is bottom to top as listed.

D.3 Proposition 2 and Corollary 2: pricing a traffic increase

The three exponents in the paper ($N^{-1/2}$ for snapshot detection, $N^{-1/3}$ for cumulative detection, $N^{-1/5}$ for estimating the asymptote) are easier to act on inverted. Figure 9 plots the horizon a given traffic multiple buys. Tenfold traffic shortens the required horizon by 3.2×, 2.2×, or 1.6× respectively; halving the horizon costs 4×, 8×, or 32× the traffic.

Pricing a traffic ramp. This prices the standard proposal to fix a slow measurement by ramping the experiment. Going from 1% to 10% of traffic is tenfold, and buys at most 3.2×, but only if the estimand is mere detection of an attenuated gap under snapshot analysis. If the estimand is the asymptotic multiplier, the same ramp buys 1.6×, and reaching half the horizon would require 32× the traffic. The exponent is the binding constraint. The remaining lever is calendar time or an external estimate of ω , which by Corollary 2(ii) moves the law from $N^{-1/5}$ to $N^{-1/3}$ and is the cheapest available improvement.

D.4 Proposition 1: which channel carries the value

Proposition 2 separates the marginal value of a view into a direct creator channel, the view itself, and a take-up channel. Figure 10 plots the three as shares against the loop gain. At $\mathcal{R} = 0$ the creator channel is 46% of the total, rising to 70% at $\mathcal{R} = 0.7$ and 87% at $\mathcal{R} = 0.95$. Counting the δ -dependent part of the take-up channel as well, the share of marginal value omitted from published budget objectives in this class runs from 50% to 95% across the same range.

The monotonicity is structural across choices of illustrative units. Both the take-up and creator channels carry the submission value

$v^* \propto 1/(1 - \mathcal{R})$. The view-itself term lacks that multiplier, so as $\mathcal{R} \rightarrow 1$ the view-itself term’s share vanishes and the creator channel approaches $1/(1 + q'V_{\text{org}})$, which is 91% here. Only that limit and the crossover location depend on the chosen $q'V_{\text{org}}$.

Design consequence. Read against Corollary 1, this is the paper’s design argument: the omitted channel dominates at high $\mathcal{R}$, which is exactly where it is slowest to measure. An allocator scoring only viewer-side signals discards the majority of marginal value where that majority is largest, and the mechanism’s own feedback cannot reveal the loss (Theorem 1).