Platform Control and the Political Economy of Audience Access

Flyxion
Independent Researcher

November 2026

Abstract

A page or creator that publishes on a large social platform builds a relationship with followers, and the platform then determines how far each post travels while selling placement in the same feed. This essay argues that two forms of control follow from that arrangement: control over whether a producer can use a relationship already established with followers, and control over whether a user's refusal persists across repeated exposures. Both concern how far a recorded preference governs future delivery: who can register it, which decision it enters, and how long the system gives it effect. A follow can fail to secure delivery through ranking, and the Canadian restriction on news shows that it can fail earlier, when a source's eligibility is withdrawn. The claims are structural and conditional, and they concern incentives and not the intent of any firm or the quality of any product. A simple auction model supplies a benchmark in which, with identical valuations, competition for fixed placement transfers the gross value of the placements to the platform, a comparative static in which added bidders weakly raise the price of being seen and weakly lower each incumbent's surplus, and a compatibility result showing that platform revenue and negative buyer value can coexist when willing buyers keep arriving. Further results state that, when paid selection exhibits a smaller exposure-probability gap between quality classes than organic selection and the route-specific probabilities are held fixed, increasing its weight narrows the pooled gap, which is a different quantity from what a reader infers from one exposure, how a bid-ranked allocation can diverge from welfare when recipients bear a cost that no bid reflects even if every advertiser bids its true value, why entity-level hiding cannot be durable against replacement, and when an attribution-based control lowers a user's burden. Two observations summarize what the formal analysis adds: a losing bidder can affect revenue because its bid can set the price a winner pays, and a recipient can bear the cost of a placement without appearing in the auction. The essay closes with remedies grouped by the form of control they address, and it separates what is proved, what is documented, what is argued, and what remains open.

1 Introduction: Built Audiences and Rented Access

A page, a creator or a small business publishes on a social platform, answers questions, and gives people a reason to follow. A follower makes a deliberate choice, and the choice is recorded as a relationship between the follower and the producer. The platform then decides how far each post travels, and it sells placement in the same feed in which the post would appear. The arrangement is familiar enough that its structure is easy to overlook. A relationship that the

producer built, at the cost of effort and time, is carried by a delivery system that the producer does not control, and the party that controls the delivery system also sells placement in the same feed. The platform determines the organic reach of a producer's posts and separately sells placements that may reach followers as well as other users, so the producer can use the relationship only to the extent that the platform's delivery rule allows.

The essay organizes this arrangement around a single distinction between access that a follower requests and avoidance that a reader requests. A follow expresses a request to receive a producer's posts, and a hide or a block expresses a request not to receive a particular exposure again. In both cases the platform determines how far the request carries. A follow can establish eligibility for distribution without securing delivery, and a hide can suppress one advertiser without resolving what happens when a similar exposure arrives under another identity. The same control over distribution therefore makes audience access a product for the producer and makes avoidance a recurring task for the reader. These are two faces of one arrangement, and the essay treats them together because the remedies that address one often fail to address the other.

Put as a question about any recorded preference, the arrangement asks who can register the preference, which decision the preference enters, and how long the system gives it effect. A follow registers a preference whose effect depends on whether the source remains eligible, on ranking, and on which feed the reader uses. A hide registers a preference whose effect depends on which future exposures the platform treats as coming from the same source. The auction analysis explains how a platform can earn revenue from allocating delivery, and the attribution analysis states what it would take for a refusal to persist. The essay uses the auction results as a benchmark and as the source of two observations that return in the conclusion, and it gives the central policy argument to the avoidance and attribution results. The framing does not depend on whether delivery to every follower was ever general practice, a question that the sources reviewed do not settle.

The claims are structural and conditional. The essay makes no judgment about whether any particular product, page or business is good, and it attributes no motive to any firm. Where a statement depends on an institutional fact, such as what a platform announced, how its auction is described, or what its controls cover, the statement is limited to what published sources establish, and the remainder is flagged as unverified. Where a statement depends on a model, the conditions are stated and the result is proved under them. The essay closes with a section that separates the proved from the documented, the argued and the open, so that a reader can see which parts rest on which kind of support.

2 Ranked Delivery and Paid Placement

The follow relationship can be described through three arrangements, and the sources available support some of them much better than others, so the arrangements are ordered here for exposition and not as a chronology. In the first arrangement, a follower's choice to follow a page would, on a simple reading, result in that page's posts appearing in the follower's feed. In the second arrangement, delivery is ranked, so that each post competes with other eligible

posts for a place in the feed and reaches only a fraction of followers. In the third arrangement, ranked delivery coexists with paid placement as an additional route to reach. Whether the first arrangement was ever general practice is not established by the sources reviewed here, and the essay treats delivery to all followers by default as an institutional benchmark, a reference case against which later arrangements can be compared, and not as a documented past state of the platform.

The record does establish that the platform was publicly addressing concerns about declining organic reach for pages by 2014. A post on the platform's business site that year acknowledged that many people were concerned about declines in organic reach, defined organic reach as the number of people a page can reach for free by posting, and attributed the decline in part to the growth in the amount of content competing for places in each feed [3]. The same post denied that the decline was motivated by revenue and presented paid media as a way for businesses to reach audiences more predictably than organic content allows. The post is evidence for two claims and for no more. It shows that the platform treated organic reach as variable and subject to ranking, and it shows that the platform described paid distribution as a counterpart to organic distribution. It does not show that payment had become the practical route to reach for any defined population, and a claim of that kind would need a stated population, period and threshold, together with a source for any figure.

A further announcement in January 2018 supports a more specific claim about ranked delivery. The head of News Feed described an update to ranking that would prioritize posts expected to prompt conversation and interaction among people, and the announcement stated that pages might see decreases in reach, video watch time and referral traffic, with the impact varying from page to page [4]. The announcement added that pages making posts that people generally do not react to or comment on could see the largest decreases in distribution, and that people who wanted to see more posts from pages they follow could choose to see them first through a feed preference. Two features of this text matter for the argument. The first is that a follow did not, by itself, secure placement, since the announcement describes delivery as depending on a prediction about how people would interact with the post. The second is that a control existed through which a follower could raise the priority of a followed page, which shows that the platform has at times offered a mechanism that lets a follower's request carry further, and which gives the later discussion of guaranteed delivery a precedent within the platform's own design.

A 2022 announcement adds a qualification that the history should preserve. The platform introduced a Feeds tab for seeing the most recent posts from friends, pages and groups, with a Favorites list for filtering them, and it named the default tab Home, describing Home as a discovery surface driven by recommendation [5]. The announcement states that ads are included in Feeds, and press coverage repeated the point [5, 6]. The history therefore needs to distinguish access through the default ranked surface from access through a separate feed. A follower's request can carry further in the second setting than in the first, and the producer's exposure to the platform's allocation rule differs between them. The sources reviewed here leave open how much of the audience's attention each surface captures, and the essay makes no claim about the proportion.

The section ends by fixing the terms used in what follows. An audience is the set of

people who have chosen to follow a producer, and a follower relationship is the recorded choice itself. Reach is the number of distinct people to whom a post is delivered, and placement is a position in a feed that a bidder can obtain through an allocation rule. An exposure route is the mechanism through which a particular item reached a particular person, either organic selection through the ranking of eligible items or paid allocation through an auction. Payment is the amount a producer remits for a delivered outcome, which the model below stipulates to be per impression or per click until the relevant billing terms are verified. Gross value captured is the part of an advertiser's value for placements that a platform collects through the clearing price, and the essay uses this language in place of rent because calling the entire payment economic rent would require an account of the platform's costs and of the return needed to supply the service.

3 News Distribution and the Cost of Reporting

The Canadian restriction on news shows the eligibility side of the follow relationship more plainly than any ranking change. In August 2023 the platform announced that it had begun ending news availability for people in Canada in order to comply with the federal Online News Act, with the change implemented over the following weeks, and the announcement stated that people in Canada would no longer be able to view or share news content on Facebook and Instagram [1]. The platform identified news outlets by the language of the legislation, and the announcement noted that its tests would affect outlets both inside and outside Canada. Its stated rationale was that the legislation misrepresents the value that outlets receive from the platform, since outlets share content voluntarily to expand their audiences, and that people do not come to the platform for news [1]. In November 2026 coverage still described the restriction as in effect, and reported the platform's Canadian policy lead as saying that resolution through Canada-U.S. trade talks after the American midterm elections was *very likely* [2]. For the present argument the case matters for a structural reason and not because of the dispute over the statute. A reader who followed a news outlet recorded a choice, and the platform changed which sources were eligible to reach that reader. For an affected outlet the change occurs before any ranking decision is made, so a follow does not secure delivery even in the weak sense of securing eligibility. The essay takes no position on whether the legislation is well designed, and the platform's stated reading of the exchange is an empirical claim that the sources reviewed here do not settle.

A platform with data on what millions of people read, share and discuss can use those observations to improve its predictions about which stories a reader might want to see. Those predictions concern the circulation of information that is already available to the platform. The essay makes no claim about what the platform believes regarding its ability to supply that information from another source, and the structural point holds whichever way that belief runs. Consider a municipal council decision. Someone has to attend the meeting or obtain the record, read the budget, ask officials to explain it, and check an account against the documents. If nobody does that work, a ranking system can still surface commentary about the decision, but it has a thinner record from which to select. The distinction matters especially for local

news, where the relevant facts may be known to a few people and may enter the public record only because someone took the time to ask.

A reported story can circulate at low marginal cost once it exists. The next story still requires someone to obtain the records, interview the people involved and verify what happened, and reporters spend years acquiring the skills and local knowledge that allow them to recognize a significant fact in the first place. A newsroom pays for that capacity before it knows which inquiry will yield a story or how widely the story will travel. Reporters need salaries that cover housing and food during the weeks when the work produces little that can yet be published. Better ranking can help readers find the finished account, while the cost of producing it remains with the newsroom. When a platform decides whether those organizations can reach readers, it also shapes the conditions under which they can recover some of that investment. In the terms of this essay the newsroom is a producer whose audience was built at its own cost and whose reach depends on a delivery rule that it does not control, and the Canadian case is the limiting instance in which eligibility itself was withdrawn.

Two recommendations follow, and they rest on different grounds. The block demonstrates control over eligibility, and that supports the recommendation that the platform restore access to news in Canada, so that a reader's choice to follow an outlet can again take effect. The case for payment rests on a separate point, which is sustaining the reporting that gives such access its value. The platform's referral argument deserves some force here, since a referral can help an outlet gain a subscriber. The cost of maintaining a reporting beat, however, comes due whether or not that week's work generates traffic, and the platform's ability to identify a useful story depends on someone having had the time and resources to report it. The agreements proposed should therefore compensate original reporting, with terms that protect editorial independence and allow local and independent outlets as well as national publishers to participate, and they should preserve readers' ability to follow and share the sources they choose. Understood this way, the recommendation concerns the financing of journalism and is distinct from a charge for each link shared. It is also the one recommendation in the essay that takes a position on a specific policy outcome and not on a structural feature of the arrangement, and it rests on the reporting-cost argument and not on the model. The opposing considerations are real, since a payment arrangement could favour large publishers over small ones or place conditions on editorial choices if it were poorly designed, and the argument does not require settling the value of referrals, because whatever their value, the decision to withdraw eligibility shows how much of the producer's position is held by the party that controls distribution.

4 A Model of Placement Under Fixed Capacity

The analytical core of the essay is a simple auction model that isolates competition for attention. The model is a deliberate abstraction. The platform's published description of its auction states that the winner is determined by a total value that combines the advertiser's bid, an estimated rate at which the targeted person will take the advertiser's chosen action, and a measure of ad quality, and that the highest bidder therefore does not always win [7, 8]. The Instagram help page, which is linked, describes the same rule, and it carries the claim where

the Facebook business page lacks a verified address. The model below omits the second and third components and treats placements as allocated by bids alone. This omission simplifies the proofs and removes features that matter in practice, including the fact that an advertiser with a higher predicted action rate can win with a lower bid. The section returns to these omissions after the results.

Suppose there are K identical placements and a set of bidders N with $|N| \geq K + 1$. Each bidder wants at most one placement and bids $b_i \geq 0$. The K highest bids win, and every winner pays a uniform price p equal to the highest losing bid, which is the $(K + 1)$-st highest bid. With unit demand, bidding one's true value is a weakly dominant strategy in this format, which follows the standard argument for second-price auctions [9], and position-auction research supplies a precedent for treating allocation rules and bidder values explicitly [10, 11]. The first result is a benchmark.

Proposition 1 (Gross value captured under identical valuations)**.** *Suppose every bidder values a placement at the same amount v and bids truthfully. Then the clearing price is v, each winner obtains zero surplus, and platform revenue is Kv.*

Proof. All bids equal v, so the $(K + 1)$-st highest bid is v and each of the K winners pays v. A winner's surplus is its value minus its payment, which is zero, and revenue is the sum of K payments of v. Which of the identical bids wins does not affect revenue or surplus. □

This is a clean case of the platform collecting the gross value of the placements to advertisers. A low cost of making something to advertise is one reason to expect many bidders, but the auction rule and the common valuation do the work in the proof, and the model does not connect production cost to the number of bids above the clearing price. The proposition speaks to the gross value of placement and not to economic rent in the full sense. The essay connects it to the literature on rent seeking and contests, in which competitive entry into a prize can dissipate the prize [12, 13, 20], and it uses the language of gross value captured because rent in that literature is measured net of opportunity cost, and the proposition does not model the platform's own costs.

The second result is a comparative static with heterogeneous bids, and it gives the supply claim its proper scope. Fix a tie-breaking rule that preserves the relative priority of existing bidders, and assume that the incumbent bidders $1, \ldots, n$, with $n \geq K + 1$, bid truthfully, with bids $b_{(1)} \geq b_{(2)} \geq \cdots \geq b_{(n)}$ and clearing price $p = b_{(K+1)}$. An entrant now submits a bid b_0.

Proposition 2 (Entry weakly raises price and weakly lowers incumbent surplus)**.** *After the entrant's bid is added, the clearing price is*

$$p' = \min\{\, b_{(K)},\ \max\{b_0,\ p\}\,\},$$

so $p' \geq p$. Platform revenue Kp' weakly rises, and the surplus of each incumbent bidder weakly falls. The clearing price strictly rises if and only if $b_{(K)} > p$ and $b_0 > p$.

Proof. Adding an element to a finite list cannot lower its $(K + 1)$-st largest element, which gives $p' \geq p$. For the formula, consider three cases. If $b_0 \leq p$, the new bid ranks at or below

the old $(K+1)$-st bid and leaves the $(K+1)$-st highest bid at p, and the formula returns $\min\{b_{(K)}, p\} = p$. If $p < b_0 < b_{(K)}$, the new bid becomes the $(K+1)$-st highest bid, and the formula returns b_0. If $b_0 \geq b_{(K)}$, the new bid enters the top K, the old K-th bid becomes the $(K+1)$-st, and the formula returns $b_{(K)}$. In each case the formula matches. The price is the value of the $(K+1)$-st highest bid and so does not depend on how ties are broken. The tie-breaking rule decides only which of two tied bidders wins, and so bears on whether an incumbent is displaced. Because $p' \geq p$, revenue weakly rises when K placements are sold at the higher price. An incumbent who still wins pays $p' \geq p$ and so has weakly lower surplus. An incumbent who is displaced has surplus zero, and a truthful incumbent who won at price p had surplus $b_{(i)} - p \geq 0$, so displacement also weakly lowers surplus. Strict increase requires $p' > p$, which holds exactly when both $b_{(K)} > p$ and $b_0 > p$, since $p' = p$ whenever either fails. □

A simple illustration has one placement and incumbent bids of 10 and 5, so $K = 1$ and $p = 5$. An added bid of 7 gives $p' = \min\{10, 7\} = 7$ and raises the price from 5 to 7 without changing the winner. An added bid of 4 leaves the price at 5. This result is separate from the benchmark, where the price has already reached v and cannot rise further, and it supplies the claim that more competition for fixed placement can raise the price of being seen while the demand for the advertised products stays fixed. The result is stated at the level of the incumbent, since averages over entrants change with the population being averaged, and a statement about average returns would mix a change in each bidder's return with a change in which bidders are counted.

The mechanism that the screening and attribution sections both use can be stated once.

Lemma 1 (A losing bid can set the winner's price)**.** *Suppose $n \geq K+2$ truthful bids, and let the $(K+1)$-st highest bid $b_{(K+1)} = p$ be strictly above the $(K+2)$-nd, $b_{(K+2)} < p$. Removing the $(K+1)$-st bid, with all other bids fixed, lowers the clearing price from p to $b_{(K+2)}$ and lowers revenue by $K\big(p - b_{(K+2)}\big)$, although the removed bidder wins no placement.*

Proof. The K highest bids are unchanged by the removal, so the same bidders win. The $(K+1)$-st highest bid among the remaining bids is the former $(K+2)$-nd, which is $b_{(K+2)}$, and each of the K winners now pays that amount. □

With one placement and bids of 10, 5 and 3, removing the losing bid of 5 lowers the price from 5 to 3. The lemma is the sense in which an entrant that buys nothing in a given auction can still contribute to platform revenue, and it is a statement about the affected auctions with the remaining bids held fixed.

The third result concerns payment per attempt, and it is where the essay switches from an auction with an endogenous clearing price to a posted-price example with an assumed supply of willing buyers. A seller charging a posted price p for each of K placements rations the placements among the buyers who are willing to pay that price. An entrant is willing to buy an impression or a click at price p because its believed expected contribution margin $\hat{v}$ exceeds p, while its actual expected contribution margin v falls below p. The entrant's belief may be too optimistic because it rests on a single product, on advice that promotion is the next step, or on a failure to account for the gap between an impression and a purchase.

Proposition 3 (Revenue persists under replenishment)**.** *Suppose a seller posts a price $p > 0$ for each of K placements in each period and rations the placements among willing buyers. In each period $\lambda \geq 1$ new entrants arrive, each with $\hat{v} > p > v$, so that each is willing to buy one placement per period at price p. Suppose that, under an assumed turnover rule, each entrant stays in the market for $T \geq 1$ periods and is then replaced. If $\lambda T \geq K$, then once the market has reached steady state all K placements sell in every period at price p, platform revenue per period is $Kp > 0$, and every purchase has negative expected net value $v - p < 0$ for its buyer. Entrants who do not obtain a placement make no purchase.*

Proof. Once the market is in steady state, the number of active entrants is $\lambda T \geq K$, and each is willing to buy at price p, so the number of willing buyers is at least the number of placements. The seller therefore sells all K placements each period at price p and collects Kp, which is positive because $p > 0$. Each purchase yields expected net value $v - p < 0$ by assumption, and the entrants who go without a placement incur no loss in that period. □

The proposition is a compatibility result and not a prediction. Given the assumption that at least K willing buyers are always present, the sales claim is close to definitional, and the content is that positive platform revenue and negative expected value for every purchase are compatible when a continuing inflow of willing buyers keeps demand at or above capacity, and it says nothing about the empirical size of the effect or about how often the condition holds. The positive revenue uses $p > 0$, which $\hat{v} > p > v$ does not imply when the expected margin is negative, and the sales guarantee applies in steady state, since a market that starts empty has fewer than K active entrants in its first periods. The turnover rule is an assumption. A negative expected return does not by itself guarantee that every entrant observes a loss and learns from it within T periods, so T and λ stand for whatever mix of exit and arrival keeps the number of willing buyers at or above K. A further feature of the setting should be stated and does not enter the proof. Exposure is delivered to a recipient, and the recipient did not choose to receive it. When an entrant's offering has no chance of interesting the recipient, the exposure has negative value to the recipient, and a share of what the entrant pays buys an imposition on a stranger. The recipient's cost is an externality separate from the entrant's loss, and a later section states it as a result about what a bid-ranked allocation ignores.

The entrant in the model need not be a firm. It can be an individual who has made one piece of art, a book, a song or a short video, and who concludes that the next step is to advertise it. For such an entrant the believed margin could rest on one product and on no evidence of demand, so the gap between the believed and actual value could be large, and the replenishment condition would be met if each exit were followed by another person finishing a first piece. This is an illustration of the supply that the model needs, and it is not a description of observed advertisers. The account of why such entrants keep arriving belongs to a literature this essay does not survey, and the model needs only the supply and not the explanation.

The section closes by naming what the abstraction omits. The actual allocation rule includes the estimated action rate and the ad quality, so an advertiser with a higher predicted action rate can win with a lower bid, and the results above hold in the sense that they describe the bidding component of a richer rule. The model also treats each entrant as buying at most one placement, treats capacity as fixed, and sets aside any strategic behavior beyond truthful

bidding. Whether the platform's costs leave a positive profit after serving advertisers is outside the model, and so is the question of whether the platform ever caps its own supply. A further limit concerns the relation between the two routes. The propositions describe competition for a fixed number of paid placements. They do not show that selling paid placement reduces organic delivery, that advertisers are buying access to their own followers, or that cheaper content production raises the number of bids above the clearing price, and a claim of any of these kinds would need a model that connects paid and organic delivery. The institutional argument rests on the coexistence of ranked delivery and paid access, which the cited announcements document.

5 Incentives for Viability Screening

The third result applies to onboarding. An advertiser's contribution to platform revenue can include both its own payments and its effect on the prices other advertisers pay, and neither depends on whether the advertiser has a business plan or a business at all. This is a statement about the structure of the revenue relation and carries no claim about any firm's practice. The relevant comparison for a screening decision is between the incremental revenue from retaining an entrant, after accounting for what a replacement advertiser would spend, and the cost of serving the entrant together with any longer-run effects on users, on other advertisers, or on compliance. Screening itself costs money, since it requires review, and the cost is incurred whether or not the review changes an outcome.

The model gives a precise form to the first part of the comparison, and Proposition 2 applied in reverse needs care at the clearing price. Removing a bid strictly below the clearing price leaves the price unchanged. Removing a bid at the clearing price may lower it, because the removed bid may be the one setting the price, even if it lost, as in Lemma 1. Removing a bid above the clearing price weakly lowers the price as well. An advertiser can therefore contribute to platform revenue by setting the price that another advertiser pays, even when it wins no placement in a given auction, and this is a reason a screen can reduce revenue from entrants who buy nothing in that auction. Because several bidders can jointly determine the price, the effects of removing individual bidders need not add, and the condition must be stated for the whole set removed. Let R be the set of applicants that a screen would remove, let $\Delta_{\text{rev}}(R)$ be the revenue lost from removing the whole set after replacement and reallocation, let $L(R)$ be the platform cost avoided by excluding R, including any attrition of users or advertisers and any compliance cost, and let c be the cost of reviewing each of the N applicants who are screened. Proposition 2 gives the weak direction of the effect in the affected auctions, with the remaining bids held fixed and enough bidders left to clear, and $\Delta_{\text{rev}}(R)$ is the full counterfactual change once replacement and reallocation are allowed. The notable implication is that a losing bidder can matter to the platform because its bid sets what the winner pays, which gives the screening argument a mechanism beyond the spending of the applicants who are screened. Screening raises the platform's payoff only if $L(R) > \Delta_{\text{rev}}(R) + Nc$. The cost term counts applicants who pass as well as those who fail, since the review is incurred for both. The inequality is conditional, and it names the quantities on which the platform's incentive

depends. It supports a conditional explanation for weak viability screening and leaves the empirical size of the incentive open.

A further question is what the platform's checks are designed to predict. A platform can estimate the likelihood that a targeted person will take an advertiser's chosen action and can assess the quality of an ad, and both estimates can be made while leaving the advertiser's profitability unresolved. These are separate tasks. The first concerns whether a person will click, buy or sign up, and the second concerns whether the advertiser earns more from the resulting action than the advertiser paid to obtain it, which depends on costs and margins that the platform does not observe. The evidence reviewed so far supports the first claim, since the published description of the auction includes an estimated action rate and an ad quality term [7, 8], and it does not support any statement about whether the platform assesses the second. A reproduced help page describes boosting a page post with a budget, but the frequency and recipients of promotion prompts, the scope of any viability review, and the rules for geographic expansion are not established by the sources reviewed, so the essay asserts nothing about them.

The range of entrants that a single user encounters shows what this means in practice, and it is described here by category and without naming individuals or small businesses. Local-service firms in other provinces reach a user whose location the platform holds, and judging by their locations they may be unable to serve that user. Regional chains and national brands appear beside individual creators, coaches, travel agents and product-promotion pages, and pages that appear to publish generated content appear beside them. Nothing visible to the recipient distinguishes an entrant with a viable business from one with a single product and a hope, which is the screening question in concrete form. This description records what appeared in the feed. It cannot establish who chose the audience, whether the advertiser selected a national or regional audience or the platform widened it, or whether the business could serve the viewer, and each of those questions depends on documentation of the advertiser-side targeting and audience-expansion options that remains to be checked.

6 Organic Selection and What an Exposure Means

The informational argument concerns what a recipient can infer from an exposure. Suppose two posts appear in a feed, one selected through audience response and the other through a paid allocation rule. The two appear alike on the screen and carry different evidence. Where spread depends on people choosing to react to and share an item, appearing in a feed is correlated, imperfectly, with something about how people valued the item. Where an item appears because a producer purchased the placement, appearing is correlated with the producer's willingness and ability to pay. Paid selection is not uninformative, since a producer's willingness to pay can itself reflect something about the product, and a visible label lets a reader condition on the route of delivery. The question is how much the pooled exposure tells a reader about a specified quality measure when the reader sees a mixture of the two routes. The section states the claim conditionally and defines the quantity it measures, a difference in exposure probability, and it separates that quantity from the likelihood ratio that a reader uses in updating. The essay

does not claim that paid placement destroys the signal. Good content can fail to spread, weak content can spread, and organic spread is an imperfect signal even when no purchase is involved. This is the kind of selection by differential cost that the signaling literature describes, in which an action is informative about type to the extent that its cost or availability differs by type [14].

Fix two classes of item, $Q = 1$ and $Q = 0$, within a comparable set of eligible items, where Q is the specified quality measure. Each eligible item is assigned, before exposure is realized, an unobserved route $R \in \{o, p\}$, organic or paid, and the route is drawn with the same probability $\alpha = \Pr(R = p)$ in both classes. Let $E = 1$ denote that a given recipient is exposed to the item. Write $\pi_r(q) = \Pr(E = 1 \mid Q = q, R = r)$ for the exposure probability of a class-q item under route r, and let

$$s_o = \pi_o(1) - \pi_o(0), \qquad s_p = \pi_p(1) - \pi_p(0)$$

be the differences in exposure probability between the two classes under organic and paid delivery. The quantity s_r measures how much more likely exposure is for a high-quality item than for a low-quality item when delivery runs through route r. The route-specific exposure probabilities $\pi_o(q)$ and $\pi_p(q)$ are assumed not to depend on α.

Proposition 4 (Dilution of discrimination by the paid route)**.** *If the route is unobserved and independent of Q, the pooled difference in exposure probability between the two classes is*

$$s(\alpha) = (1 - \alpha)s_o + \alpha s_p.$$

When $0 \le s_p < s_o$, $s(\alpha)$ is strictly decreasing in α, with $s(0) = s_o$ and $s(1) = s_p$.

Proof. By the law of total probability and independence of the route from Q, $\Pr(E = 1 \mid Q = q) = (1 - \alpha)\pi_o(q) + \alpha\pi_p(q)$. Subtracting the expression for $q = 0$ from the expression for $q = 1$ gives $s(\alpha)$. The derivative is $s_p - s_o < 0$ when $s_p < s_o$, and the endpoint values follow by substitution. □

The proposition gives an exact result and a reason to distinguish two quantities that are easily confused. Here $\alpha = \Pr(R = p)$ is the probability that an eligible item is assigned to the paid route before exposure is realized, which is not the probability of the joint event that the item is on the paid route and is exposed. The paid share among delivered impressions is a different quantity, because it depends on how the delivered impressions are weighted across items, and the proposition makes no claim about it. A label identifies the route for a delivered item, and the dilution in Proposition 4 describes an observer who sees only the mixture. Estimating s_o or s_p themselves is a different task, because it requires exposure rates among eligible items, including items that never appeared, and neither a reader nor an outside researcher observes those. This is a reason for treating disclosure of the route as a remedy that identifies the selection mechanism for each delivered item and leaves the allocation rule alone, though what a reader can infer from the label depends on what the reader knows about selection through each route, and the platform's own data on eligible items matter for any measurement of discrimination.

The difference in exposure probability is not the quantity a reader uses in updating. After seeing an item, the relevant quantity is the likelihood ratio $\Pr(E = 1 \mid Q = 1)/\Pr(E = 1 \mid Q = 0)$, which depends on the baseline exposure rate as well as on the difference. Suppose organic exposure probabilities are 0.6 and 0.3 for the two classes, so that $s_o = 0.3$ and the likelihood ratio is 2, and suppose paid exposure probabilities are 0.02 and 0.001, so that $s_p = 0.019$ and the likelihood ratio is 20. The paid route has the smaller difference and the larger ratio. With an unobserved route and $\alpha = 0.5$, the pooled difference is $s(0.5) = 0.1595$, about half of s_o, while the pooled ratio is $0.31/0.1505 \approx 2.06$, close to the organic ratio. Proposition 4 is therefore a statement about how the difference in exposure probability moves with the paid weight, and it does not state how much a reader learns from an exposure. A reader who sees the label can condition on the route, and what the reader then infers depends on the route-specific probabilities π_o and π_p, which do not depend on α, and on what the reader knows about selection through that route.

The corresponding claim about engagement counts is that a count without a denominator or a breakdown by traffic source conflates the amount of exposure with the response to it. An engagement count is the product of the number of people exposed and the rate at which exposed people respond. If the count is of distinct people who engaged, dividing by distinct people reached gives the fraction of those reached who engaged, and if it is of actions and people can act more than once, the same denominator gives actions per person reached, so the two should not be mixed. When some of the exposure was purchased, a high count can reflect a large purchase and a low count can reflect a small one, and the count does not separate this from the rate of response among those exposed. Reach by route supplies the denominators, and engagement by route supplies the numerators, so a producer needs both in order to compare response rates. A producer or an advertiser who uses the count to judge what works then has a measurement that moves with spending. This is a statement about what the number measures and does not claim that any particular reported figure is inflated.

7 The Recipient's Side of the Exchange

The auction model has three parties, and only two of them appear in its prices. The entrant pays the platform a clearing price, and the platform collects it. The recipient, who receives the exposure that the price purchased, neither pays nor is paid, and yet the recipient's attention is the resource being sold. Suppose each exposure costs a recipient an amount $\kappa > 0$ in attention, and suppose that the recipient gains an amount $w > 0$ if the offer turns out to be of interest, which happens with probability θ_i for the offer of entrant i. The recipient's expected value of an exposure to entrant i is then $\rho_i = \theta_i w - \kappa$, which is negative whenever $\theta_i < \kappa/w$. This is a statement about the structure of the exchange and makes no claim about how often the inequality holds. Let v_i be entrant i's true expected margin from the exposure, let $c \geq 0$ be the platform's cost of serving a placement, and write $b_i = \hat{v}_i$ for the entrant's bid, equal to its believed margin. The valuation error is $o_i = b_i - v_i$, which is positive for the entrants of Proposition 3. It differs from the entrant's expected net loss on a purchase, which is $p - v_i$ when $p > v_i$, so for the replenishment entrants, where $\hat{v}_i > p > v_i$, the error exceeds the loss

by exactly $\hat{v}_i - p$, while outside those assumptions the two can differ in either direction. The payment from entrant to platform is a transfer and drops out of the total, so the welfare of a set S of K winners is

$$W(S) = \sum_{i \in S} (v_i + \rho_i - c).$$

Proposition 5 (What a bid-ranked allocation ignores)**.** *For any set S of winners,*

$$W(S) = \sum_{i \in S} b_i - \sum_{i \in S} o_i + \sum_{i \in S} \rho_i - Kc.$$

The allocation that awards placements to the K highest bids maximizes the first term only, and so takes no account of the valuation errors or of the recipient's value. There exist configurations in which the bid-ranked allocation yields lower welfare than another feasible allocation even when every entrant bids its true value, so that every $o_i = 0$, and there exist configurations in which it yields negative welfare while another feasible allocation yields positive welfare.

Proof. The identity follows by substituting $v_i = b_i - o_i$ into the definition of $W(S)$. The bid-ranked allocation selects S to maximize $\sum_{i \in S} b_i$, which is the first term, and nothing in the selection rule depends on o_i or ρ_i. For the existence claims, take $K = 1$ and $c = 0$ with two entrants in each case. In the first case both entrants bid their true values. Entrant A has $b_A = v_A = 5$ and recipient value $\rho_A = -6$, and entrant B has $b_B = v_B = 4$ and $\rho_B = 1$. The bid-ranked allocation awards the placement to A, giving $W = 5 - 6 = -1$, while awarding it to B gives $W = 4 + 1 = 5$. In the second case the recipient values are equal across entrants, $\rho_A = \rho_B = 0$, and the entrants misjudge their margins. Entrant A has bid $b_A = 10$ and true margin $v_A = 2$, and entrant B has bid $b_B = 6$ and true margin $v_B = 6$. The bid-ranked allocation awards the placement to A, giving $W = 2$, while awarding it to B gives $W = 6$. □

The numbers in the proof are illustrative and carry no claim about frequency or size. The first case shows the independent force of the recipient externality, since no entrant misjudges its own margin and the allocation still diverges from the welfare ranking, because ρ_i enters welfare and does not enter the bid. The second case isolates the valuation error, since recipient values are equal and only the gap between bid and margin reverses the ranking. The sign of the divergence does not depend on the particular magnitudes. With one placement, any bids, margins and recipient values for which the higher bid belongs to the entrant with the lower value of $v_i + \rho_i$ produce it, and the numbers in the proof fix only its size. The platform's published rule adds an estimated action rate and an ad quality term to the bid, and those terms can partly proxy for the recipient's value to the extent that they track what recipients want [7, 8]. The proxy is imperfect for a reason that can be stated without any assumption about the platform's methods. A prediction that a person will take an action is not a measure of whether the person is better off for having seen the offer, because an exposure that does not lead to the action still cost the recipient attention, and an action induced by a misleading presentation can leave the person worse off. Even a well-calibrated action probability would estimate only something like θ_i, and since $\rho_i = \theta_i w - \kappa$, the recipient's benefit w from a relevant exposure and the attention cost κ would still have to be estimated. Nothing in the

argument makes these quantities unmeasurable in principle, only unmeasured in the sources reviewed, and platform data on attention or downstream outcomes might bear on them. The essay therefore does not assume that the extra terms close the gap, and it leaves the size of the gap to measurement.

A hypothetical case makes the second source concrete. Consider an individual who has made one piece of art and decides to advertise it, and suppose that the true margin is small, that the believed margin is large relative to it, so that o_i is large, and that the offer reaches recipients to whom it has little reason to appeal, so that ρ_i is negative for most of them. These are properties assumed for the case. The feed observations described earlier establish only what appeared to the recipient, and the essay presents no evidence on the margins, beliefs or recipient values of any actual advertiser. A further consequence concerns the recipient's means of response. A recipient cannot bid against an unwanted placement, and the only signal the recipient can send is a hide or a block. Each such action costs the recipient effort, and it prices the externality in the recipient's own labour and not in money. This is the connection between the externality in this section and the burden of avoidance in the next, since the recipient is both the party whose cost is left out of the allocation and the party who must act to reduce it.

8 Avoidance Costs and Entry Costs

The avoidance result concerns the scope at which a user's refusal operates and the scope at which an unwanted exposure is reproduced. The documented controls are limited in what they cover. A user can hide a single ad, and a user can hide all ads from a given advertiser, and the platform's help pages have described the second control as a list of advertisers from which the user can choose to stop seeing ads [18]. A separate control lets a user ask to see fewer ads on a topic, and the help page in the archived version reviewed here lists a short set of topics, namely alcohol, parenting, pets, and social issues, elections or politics, and describes a topic choice as influencing which ads a user sees [18]. An earlier version of the help page had stated that there is no way to hide all ads [19]. Whether the current scope of these controls matches the archived descriptions, and whether an advertiser exclusion persists indefinitely or lapses, are matters the essay leaves to verification against the current documentation.

The model states the result under explicit conditions. Suppose a user's hide action excludes exactly one entity, and suppose an operator can supply an unbounded sequence of distinct entities, each of which reaches the user until hidden, with exclusions that do not carry over from one entity to the next.

Proposition 6 (No finite sequence of entity-level hides is durable)**.** *Under entity-specific hiding, excluding n separately treated entities takes n user actions. If the operator can continually replace excluded entities and exclusions do not carry over, then after any finite number m of user actions at least one entity remains unexcluded and can still reach the user, so no finite sequence of actions guarantees durable avoidance. If instead a control operates at the level of the operator, attribution is exact, and the exclusion remains in force and is applied whenever one of the operator's entities becomes eligible for delivery, a single action excludes all of the operator's present and future entities.*

Proof. Each action excludes one entity, so after m actions at most m entities are excluded. The supply of entities is unbounded, so at least one entity remains unexcluded. Under an operator-level control with exact attribution, persistence and application at eligibility, the single exclusion applies to every entity bearing the operator's identity, including entities created after the action. □

The proposition contains the policy question in a single sentence: does the control operate at the same level as the level at which the unwanted exposure is reproduced? It is a statement about scope and says nothing about whether any operator behaves in the way the proposition allows. The cost comparison sharpens the point. For the user, the cost of each hide is a decision and an action, and the cost recurs for each entity, so the total scales with the number of entities the operator supplies. For the operator, the cost of creating one more entity can be small relative to the user's cost of excluding it, and where it is, the burden lies with the party who did not choose to be exposed. This is where the labour of pruning belongs, since the user supplies classification work repeatedly that the system could in principle preserve across exposures. A rule that groups entities too aggressively has its own failure mode, because it can hide unrelated publishers, and the remedies below therefore depend on reliable attribution.

A single user's blocked list gives the argument a concrete case, and it is presented here as a self-reported record of first encounters under active filtering. The list, kept over several years, runs to thousands of separate entries, and the user reports blocking on the order of a hundred new entities per visit. The list is a blocking list that groups people and pages under one heading, which differs from the advertiser-level hide in the ad preferences, and the essay does not assume that the two controls behave alike. The list establishes a reported number of distinct entities encountered and blocked during the period. A roughly flat rate of newly blocked entities is compatible with ongoing replenishment, a large pre-existing pool, or changes in ranking and session length, and a declining rate is compatible with depletion among other explanations, so the rate is suggestive and does not identify replenishment. Inferring a rate of new page creation would require creation dates, and inferring replacement by a common operator would require evidence linking the entities. One publisher's brand appears as five separate entries, four city-specific pages and one site page, which illustrates how often the user had to exercise a control while leaving ownership open. A business in the user's own province also required its own block, which shows that distance from the user is not what produces the need for avoidance. The list does not distinguish entities drawn from a pre-existing pool from newly created ones, and it does not estimate the size of the advertiser pool.

9 Attribution and the Scope of Operator-Level Controls

The question can be put in words before any notation is introduced. When a user hides a page, what should happen when another page run by the same operator becomes eligible for delivery? Answering requires deciding what the original refusal applies to and what evidence licenses carrying it forward. If the refusal applies to the entity, nothing carries over, and the user repeats the action for each replacement. If it applies to the operator, something must establish that the new page belongs to the operator, and every such judgment can err in two ways: a

replacement can be missed, which costs the user another exposure and another hide, and an unrelated publisher can be swept in, which costs the user something the user wanted. Both consequences belong in the calculation, and a platform that reports catching most replacements has not yet shown that the user is better off, since enough wrongly suppressed publishers can outweigh the exposures avoided. The model below makes that comparison exact so that it can be checked against measured errors. A contest model of entry could supply the burden functions once they are given, but it would not say that the work of repeated blocking exists, that it recurs because the control is scoped to individual entities, or that broadening the control introduces a cost of its own, and those three observations are what turn the question into the threshold stated in Proposition 7.

Proposition 6 assumed that an operator-level control attributes entities to their operator exactly and that the exclusion persists and is applied when the entities become eligible, and attribution is not exact in practice. A control that excludes an operator must classify each new entity as belonging to the operator or not, and any classification rule makes two kinds of error. Let $r \in [0,1]$ be its recall, the probability that an entity run by the blocked operator that would otherwise reach the user is flagged, and let $f \in [0,1]$ be its false-grouping rate, the probability that a wanted, unrelated publisher that would otherwise reach the user is flagged. These are the populations counted by the arrival rates below, so the error rates that would test the threshold are those measured among entities that would actually reach the user. Both depend on the signals the rule uses, which could include names, shared administrators, shared creative assets or payment relationships. A platform holds account-level information of this kind that a user does not, and the user can see only what the interface displays, such as the page transparency fields and the name history [16]. The section states the condition under which a rule with given errors lowers the user's burden, and it makes no claim about the errors of any rule that exists.

Consider a user over a period in which the blocked operator creates new entities that would reach the user at an expected rate λ_o. Each such entity that is not excluded in advance costs the user an exposure cost κ_o and a hide cost c_h, since the user must act to remove it after the first exposure. Unrelated publishers that the user would want to see arrive at expected rate λ_u, and each one that is wrongly flagged costs the user a value g. Under entity-level hiding no new entity is excluded in advance and no unrelated publisher is wrongly flagged, so the user's expected burden per period is $B_e = \lambda_o(\kappa_o + c_h)$. Under an operator-level control with recall r and false-grouping rate f, a fraction r of the operator's new entities is excluded before it reaches the user, and a fraction f of wanted, unrelated publishers is wrongly suppressed, so the burden per period is $B_o = (1-r)\lambda_o(\kappa_o + c_h) + f\lambda_u g$, setting aside the one-time cost of the initial action.

Proposition 7 (When an attribution-based control lowers the burden)**.** *Suppose* $\lambda_o(\kappa_o + c_h) > 0$. *The operator-level control lowers the user's expected burden if and only if*

$$r\,\lambda_o(\kappa_o + c_h) > f\,\lambda_u\,g.$$

If $f = 0$, any $r > 0$ strictly lowers the burden. If $f > 0$, the condition is equivalent to

$$\frac{r}{f} > \frac{\lambda_u\, g}{\lambda_o(\kappa_o + c_h)}.$$

Proof. The difference in burden is $B_e - B_o = r\lambda_o(\kappa_o + c_h) - f\lambda_u g$, which is positive exactly when the stated inequality holds. Dividing both sides by $f\lambda_o(\kappa_o + c_h) > 0$ gives the second form when $f > 0$. □

The premise $\lambda_o(\kappa_o + c_h) > 0$ excludes the degenerate case in which the operator's entities impose no burden on the user, where a control cannot lower the burden and can only raise it through false grouping. The proposition turns the remedy's failure mode into a threshold. The ratio of recall to false-grouping must exceed the ratio of two weighted rates, the rate $\lambda_u g$ at which wanted publishers are at risk of being lost, weighted by what each loss costs, and the rate $\lambda_o(\kappa_o + c_h)$ at which the operator's entities impose a burden, weighted by what each one costs. Multiplying the first by f gives the expected loss from false grouping. A rule that is reliable in this sense needs more than a high recall, and a rule with a high recall and a high false-grouping rate can raise the burden when the user has many publishers at risk of wrongful suppression. The point also bears on matching by name alone. Two unrelated businesses can share a common word in their names, so a rule that groups on name similarity has a false-grouping rate that depends on how common the names are, while an operator that reuses a brand name across several entities would be caught at high recall. Whether ownership is common is a separate fact that names cannot establish, which is why the essay treats the repeated-name entries on the blocked list as a pattern of naming and leaves ownership open.

The platform's incentive to supply attribution deserves a statement of its own, and it needs the same counterfactual as the screening condition. Applying Proposition 2 in reverse, removing an operator's bids from a given user's auctions weakly lowers revenue in the affected auctions, holding the remaining bids fixed and assuming that enough bidders remain to clear. The mechanism is the one in Lemma 1: a bid that loses can still set what the winner pays, so the revenue effect can arise even from entities that would have won no placement. This is the revenue effect with the remaining bids held fixed, and it is not the full counterfactual, which would include replacement by other bidders and reallocation of the freed placements. A platform that weighs this revenue effect against any longer-run benefit has a conditional reason to offer weaker controls than the user would prefer. The statement is conditional on the revenue effect dominating, and the size of the effect is open. It also suggests why the measurement of the errors r and f matters for policy, since the condition in Proposition 7 can be checked for any proposed rule once the errors and rates are estimated, and for that user, under the specified rates and costs, a rule below the threshold raises expected burden while a rule at the threshold leaves it unchanged.

10 Built Relationships and Portability

A follower relationship held on a platform can be compared with one held in a portable form, such as a mailing list or a subscription feed. In the portable case, the producer holds a means

of reaching each follower and chooses when and how to use it, and delivery depends on the producer and the follower and on the channel that carries the message. In the platform case, the producer holds a record that the platform maintains, and delivery depends on the platform's allocation rule applied to each post. The difference lies in who controls delivery. A producer with a portable relationship can, in principle, carry the audience to a different channel if the first channel changes its terms, which is the option that exit provides [15]. A producer with a platform-held relationship faces the platform's ranking changes without that option.

The section asks what a producer would need in order to carry an audience elsewhere. The producer would need a stable identifier for each follower that the follower has agreed to share, a channel that accepts that identifier, and a way to signal to followers that the new channel exists. The sources reviewed here do not establish whether any platform exports a page's follower list in a usable form, and the essay asks only what the arrangement would require and what the current design allows within the platform. One mechanism within the platform deserves mention. The 2018 announcement described a preference through which a follower could choose to see a followed page's posts first [4], and the 2022 Feeds tab and its Favorites list give a follower a second way to raise a producer's priority [5]. These controls allow a follower's request to carry further, though only when the follower finds and uses them, and they leave the producer's reach dependent on follower action taken on the platform's terms. Exit is used here descriptively, since the model contains no exit option, and the section marks the point where the access argument rests on institutional design and not on a proved result. The section prepares the policy discussion without prescribing a remedy.

11 Remedies and Their Costs

Four measures are examined, grouped by the form of control they address, with attention to the tradeoffs and failure modes of each. The first two address control over an established relationship. Guaranteed delivery of posts to followers who opted in would restore the follow relationship as the thing the follower chose, and it would constrain feed design, since a feed that must deliver every followed post cannot also rank them freely. It raises a concern about volume, because a follower who follows many producers could face a feed dominated by the most frequent posters, and any guarantee would need a rule for frequency or a means for the follower to tune it. The existing See First preference and the Feeds tab are partial versions of this idea that operate at the follower's initiative. Portability of the follower graph would restore exit, and it depends on interoperability that a platform has little reason to build when the follower graph is a source of its own advantage. A portability measure would need to protect followers who did not agree to share their identifiers, which makes consent for the transfer the central design problem.

The other two measures address control over refusal and inference. Disclosure of the route of each delivered item, and of reach broken down by route, would bear on the inference argument. A route label identifies the selection mechanism for the delivered item, and how much a recipient can infer from it depends on what the recipient knows about selection through that route, as the likelihood-ratio example after Proposition 4 shows. Reach by route supplies

the denominators for route-specific response rates, and a producer or advertiser would also need engagement counts by route to compute them. The cost is that disclosure of this kind requires the platform to publish figures that it may regard as commercially sensitive, and the figures can be misread if their definitions vary. Category-level and operator-level avoidance controls address the asymmetry in Proposition 6. A control whose scope matches the source of repeated exposure could preserve a user's refusal across identities that the same operator runs, so that one action excludes what entity-level hiding would require the user to repeat. The remedy depends on reliable attribution, since a control that cannot identify a common operator either fails to catch the replacements or groups unrelated publishers by mistake, and its chief failure mode is suppressing unrelated publishers. Proposition 7 states the threshold at which the tradeoff favors the control, in terms of recall and false-grouping. A user-facing category control faces a similar tradeoff between breadth and precision.

The model supports the scope-matched controls most directly. The avoidance result states the condition under which entity-level hiding cannot be durable, and the inference result states how the pooled difference in exposure probability falls with the paid weight, while leaving open how much a reader can infer from a route label. The access argument supports guaranteed delivery and portability, though those measures depend on design constraints and interoperability that the model does not address. On the producer side, the corresponding distinction is between a relationship that remains usable through a chosen delivery channel and one whose reach is repeatedly allocated by the platform.

12 Evidence Plan

The essay distinguishes what is established from what is unverified, and this section sets out what would settle the remaining questions. The primary documents to collect are the platform's business onboarding and help pages, its auction description, its ad review and verification policies, its feed ranking announcements, and its statements on organic reach. Secondary sources are peer-reviewed work on position auctions, on rent dissipation in contests, and on declines in organic reach. For each institutional claim in the essay, the relevant document must be identified by title and date, and any claim without one is dropped or marked as unsupported.

One measurement is feasible in principle. The platform's Ad Library tools page, updated in August 2023, describes ads that deliver an impression in the European Union or associated territories as displayed while active and archived for one year after their last impression [17], which supports a regional study of repeated names within the available window. Matching names would establish a pattern of naming only, and common ownership would need separate evidence. A 2019 announcement added page transparency information to the library and to each page, including the date the page was created, previous page merges and name changes [16], and those fields could support a comparison between page creation and the earliest ad observed in the available archive. The first-ever ad date would require coverage extending back to the page's first campaign, and an archive window of one year means that the earliest observed ad is a bound on the first ad and not its date. A study of this interval must therefore

verify that both dates are observable under the relevant coverage, and it should treat any observed interval as censored by the window. The study must also establish what the library covers in the relevant regions before promising any measurement.

Anecdotes, including repeated ads from identically named pages and the single-user blocked list described in the section on avoidance, may motivate the argument but are labeled as anecdotes and are not used as evidence. The user's location is known to the platform, so the targeting question matters. Whether the advertiser selected a national or regional audience, or the platform widened the audience, determines which claim the distant-business examples support, and the advertiser-side targeting and audience-expansion documentation must be checked before the essay asserts either. Whether a block applies to a page, a person or an advertiser, since the blocking interface groups them under one heading, and whether the platform prompts individuals to promote content and on what terms, remain on the list of items to verify.

13 Objections and Replies

Five objections deserve a reply. The first is that advertising funds a free service. This is true, and it does not answer the question of who captures the gross value of the placements, since the benchmark in Proposition 1 describes how a competitive allocation moves value to the party that sells placement without addressing whether the service is worth funding. The second is that some small businesses benefit from paid reach. The model allows this. Its incumbent-level result says that surplus weakly falls as bidders are added, which does not say that every return is negative and does not rely on an average over a changing population, and the compatibility result in Proposition 3 requires a continuing supply of enough overoptimistic buyers to fill capacity at the posted price. The entry comparison also holds product demand fixed, and the model leaves open how product demand responds to changes in advertising competition.

The third objection is that declines in reach may reflect competition for feed space and not strategy. The 2014 post itself offered growth in content as a reason for the decline [3]. This affects the history and leaves the model untouched, since the structure produces the same incentives under either explanation: whatever the cause, a producer's established relationship delivers less than the producer might expect, and paid placement is available as a counterpart. The fourth is that some people do want the content that appears, including content that is generated by software. The argument does not deny this, since the entry comparison holds product demand fixed, and a person who wants a given item is not an exception to that comparison. The fifth is that ads can be relevant to users. This is compatible with the avoidance result, because the problem identified there lies in the granularity of the control and not in whether any given ad is relevant, and a user who finds an ad relevant is not the user for whom the control is needed.

14 Scope of the Claims

The essay makes four kinds of claim, and they rest on different support. What is proved is limited to the model results under their stated assumptions: the benchmark in Proposition 1, the entry result in Proposition 2, the replenishment result in Proposition 3, the pooled-difference result in Proposition 4, the result on what a bid-ranked allocation ignores in Proposition 5, the scope result in Proposition 6, and the attribution threshold in Proposition 7. Each is a statement about a stripped-down setting and does not follow from the institutional facts alone. What is documented is whatever the cited announcements and help pages establish: that ranking changes can reduce reach, that paid distribution is offered alongside organic distribution, that the auction combines a bid with an estimated action rate and an ad quality term, that advertiser-level hiding and a limited topic control exist, that the ad library and page transparency expose certain fields, and that the platform ended news availability for people in Canada beginning in August 2023 and was reported to be maintaining the restriction in November 2026. What is argued is the policy reading of control granularity and route of delivery, the reading of the blocked-list case as a record of first encounters, and the reading of the Canadian news restriction as a decision about eligibility together with the position that news organizations should be paid for the reporting the platform distributes. The Canadian news argument rests on institutional analysis and not on the model, and it is argued and not proved. What remains open is the empirical size of each effect, the extent to which the replenishment condition holds, whether a common operator stands behind the repeated entities, and the effect of any remedy.

15 Conclusion

The formal analysis adds two observations to the institutional history. An advertiser can affect the price without winning a placement, since a losing bid at the clearing price can set the price another advertiser pays, and an entrant that buys nothing in a given auction can still contribute to platform revenue. This is why screening can reduce revenue even when the screened entrant wins nothing, and why the screening condition has to be stated for the set removed. A recipient can bear the cost of a placement without appearing in the auction, since the recipient does not bid and no price reflects the cost. This is why a user must supply repeated labour to register a cost that the allocation does not directly price, and why the scope of the control matters. Taken together, the platform's revenue can depend on participants who never receive a placement, and the cost of the placements it sells can fall on a participant who never appears in the auction.

These observations lead back to the follow and the hide. The essay asks who can register a preference, which decision that preference enters, and how long the system gives it effect. A follow registers a preference whose effect depends on whether the source remains eligible, on ranking, and on which feed the reader uses, and guaranteed delivery and portability address how far it carries, while the Canadian restriction shows that eligibility itself can be withdrawn after readers have chosen to follow. A hide registers a preference whose effect depends on which future exposures the platform treats as coming from the same source, and avoidance

controls scoped to the source of repeated exposure, subject to the attribution threshold in Proposition 7, address how long it lasts. A route label identifies how an item reached the recipient. For producers, reach by route supplies the denominators for response rates, and engagement by route supplies the numerators. The claim does not depend on whether delivery to every follower was once general practice, and it does not require that selling paid placement reduces organic delivery, which the model does not show. The same control over distribution makes audience access a product for the producer and avoidance a recurring task for the reader, and the remedies that follow from the analysis are those that let a request, once made, carry as far as the person who made it intended.

References

[1] Meta Platforms. 2023. "Changes to News Availability on Our Platforms in Canada." Meta Newsroom, updated August 1. https://about.fb.com/news/2023/06/changes-to-news-availability-on-our-platforms-in-canada/.

[2] Shankar, Bradly. 2026. "News 'Very Likely' Coming Back to Meta Platforms in Canada." MobileSyrup, as published by Daily Hive, November 5. https://dailyhive.com/canada/meta-news-coming-back-canada.

[3] Boland, Brian. 2014. "Organic Reach on Facebook: Your Questions Answered." Facebook for Business, June 5. Archived at https://webarchive.library.unt.edu/dentonfracking/20141107005901/https://www.facebook.com/business/news/Organic-Reach-on-Facebook.

[4] Mosseri, Adam. 2018. "Bringing People Closer Together." Facebook Newsroom, January 11. https://about.fb.com/news/2018/01/news-feed-fyi-bringing-people-closer-together/.

[5] Meta Platforms. 2022. "Introducing Home and Feeds on Facebook." Meta Newsroom, July 21. https://about.fb.com/news/2022/07/home-and-feeds-on-facebook/.

[6] Fox Business. 2022. "Facebook Rolling out New Feeds Tab." July 21. https://www.foxbusiness.com/technology/facebook-rolling-new-feeds-tab.

[7] Meta Platforms. n.d. "About Ad Auctions." Meta Business Help Center. Undated help page, accessed November 2026; no stable address was verified.

[8] Meta Platforms. n.d. "How Instagram Ads Use Machine Learning." Instagram Help Center. Undated help page, accessed November 2026. https://help.instagram.com/119516847899397.

[9] Vickrey, William. 1961. "Counterspeculation, Auctions, and Competitive Sealed Tenders." *Journal of Finance* 16(1): 8–37.

[10] Varian, Hal R. 2007. "Position Auctions." *International Journal of Industrial Organization* 25(6): 1163–1178.

[11] Edelman, Benjamin, Michael Ostrovsky, and Michael Schwarz. 2007. "Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords." *American Economic Review* 97(1): 242–259.

[12] Krueger, Anne O. 1974. "The Political Economy of the Rent-Seeking Society." *American Economic Review* 64(3): 291–303.

[13] Posner, Richard A. 1975. "The Social Costs of Monopoly and Regulation." *Journal of Political Economy* 83(4): 807–827.

[14] Spence, Michael. 1973. "Job Market Signaling." *Quarterly Journal of Economics* 87(3): 355–374.

[15] Hirschman, Albert O. 1970. *Exit, Voice, and Loyalty: Responses to Decline in Firms, Organizations, and States*. Cambridge, MA: Harvard University Press.

[16] Shukla, Satwik. 2019. "A Better Way to Learn About Ads on Facebook." Meta Newsroom, March. `https://about.fb.com/news/2019/03/a-better-way-to-learn-about-ads/`.

[17] Meta Platforms. 2023. "Meta Ad Library Tools." Meta Transparency Center, page updated August 24. `https://transparency.meta.com/researchtools/ad-library-tools`.

[18] Facebook. 2021. "What Are My Ad Preferences and How Can I Adjust Them on Facebook?" Facebook Help Center. Archived version, October 23. `https://web.archive.org/web/20211023130603/https://www.facebook.com/help/247395082112892`.

[19] Facebook. 2015. "About Facebook Ads." Facebook Help Center. Archived version, June 3. `https://webarchive.library.unt.edu/web/20150603012618/https://www.facebook.com/help/769828729705201`.

[20] Tullock, Gordon. 1980. "Efficient Rent Seeking." In *Toward a Theory of the Rent-Seeking Society*, edited by James M. Buchanan, Robert D. Tollison, and Gordon Tullock, 97–112. College Station: Texas A&M University Press.