

# The Unchosen Adjacency

Flyxion  
Independent Researcher

September 2026

## Abstract

Social platforms can package user actions into bundles whose meaning the user does not control. An advertisement is placed beside, over, or immediately after content the user did not sponsor, and a block removes perception along with participation. This essay states the structure exactly: an action has an intended set of signed relational changes and an effect set, and the mismatch has two parts, surplus (what the action does beyond the intention) and deficit (what the intention required and the action does not do). Mismatch matters when observers attribute it to the user or link it with the user's content, or the user loses something they value, is made to bear a relation they reject, or fails to obtain a change they rate highly. An action is defective when a cleaner action is technically feasible and not offered. Applied to blocking, advertising adjacency, continued participation, and recommendation on one large platform, with another as comparator, the test returns provisional verdicts and excludes one case. A constructive section separates participation from perception, and a hypothesis about withdrawal is stated with the evidence that would test it.

## 1 Introduction

Two ordinary events on a large social platform look unrelated. In the first, a user publishes a post or Reel, and advertisements the user did not select can be placed among the user's posts, over the Reel, or immediately after it, where, on the controls reviewed, an ordinary user cannot prevent them. In the second, a user blocks another person to stop a stream of hostile comments. The block succeeds, and it also does something the user did not ask for: it removes the other person's ability to see anything the user does, including the conduct the user meant to keep visible. Intermediate controls exist, and Section 3 examines what they do and do not separate.

The first event adds a relation the user did not choose. The second removes a relation the user meant to keep. Both are cases of overreach: a single action, offered through a coarse interface, changes several relations at once, and the user is left with the combination. A third failure runs the other way, underrealization: an action may fail to make a change the user did intend, as when a restriction is imposed without the explanation or end date that correction requires. The essay states this structure exactly, examines it on one large platform with another as comparator, and argues that a good deal of what users experience as platform overreach follows from it.

The thesis is that platforms govern partly by deciding which relations can be changed separately. Their power lies not only in permitting or forbidding actions but in packaging actions into bundles whose meaning the user does not control. A user may want to speak without advertising, exclude someone from a conversation without excluding them from view, or take part in a network without lending it approval. Where the interface offers no action that does exactly this, the interface decides what the user's act means.

The vocabulary matters. *Compelled association* names the additive case, where an unwanted relation is attached. It also names a doctrine in constitutional law, which presupposes state action and does not apply to a private platform that users can leave. The essay does not rely on that doctrine. It uses *unchosen bundling* as the general term for overreach, with compelled association and compelled dissociation as its two directions, and it treats the availability of exit as part of the phenomenon: exit is the only disaggregation on offer, and its cost is what the absence of finer controls imposes. The title names the additive case, the advertisement attached to a post, while the general term is bundle.

Two features distinguish the account from a complaint about bad design. First, it distinguishes mismatches that matter from mismatches that do not. Every software action has side effects and limits, so the account asks further questions: whether observers read a side effect as the user’s own act, whether the user loses something they value or is made to bear a relation they reject, and whether the user fails to obtain a change they rate highly. Second, it is comparative. An interface is condemned only where a cleaner action is technically feasible, as shown by other systems, and is not offered.

**Related work.** The nearest conceptual neighbor is contextual integrity, which asks whether information flows accord with the norms of their context [16]. The present account asks a different question, whether distinct relational changes can be selected separately, and it covers relations that are not information flows, such as an added advertisement or a removed view. Context collapse names the merging of several audiences into a single context [17]. The collapse studied here is of the action space, in which several changes are tied to one act, and not of the audience. The graduated moderation sequence in Section 4 parallels Ostrom’s principle of graduated sanctions, under which sanctions are proportional to the severity of the violation [18], and Grimmelmann’s taxonomy of moderation techniques (exclusion, pricing, organizing, and norm-setting) [19] situates moderation as governance. The constitutional doctrine mentioned above belongs to the law of freedom of association, as in *Roberts v. United States Jaycees* [20], and it constrains the state.

The argument proceeds as follows. Section 2 states the definitions: signed relational changes, surplus and deficit, attribution and value, and the notion of a defective action. Section 3 applies them to blocking, advertising adjacency, continued participation, and recommendation, and shows that the framework returns different provisional verdicts for different cases and excludes at least one. Section 4 develops the constructive case, in which regulating participation is separated from regulating perception. Section 5 addresses objections. Section 6 states a hypothesis about withdrawal, that users who cannot disaggregate expression from unwanted association will reduce the expression, and identifies what evidence would test it. Section 7 lists candidate separable controls.

The essay makes no claim that any particular platform acts in bad faith or that the interface choices it discusses are accidental. The claim is structural: where meaning is bundled, someone other than the user is setting it, and the user’s remaining recourse is to act less or to leave.

## 2 Definitions

### 2.1 Relations, signed changes, and actions

**Definition 2.1** (Relations and signed changes). Let  $R_u$  be the set of platform-mediated relations involving a user  $u$  whose state may be changed by the platform or by an action exposed through its interface: who may view, comment, reply, tag, or message; what is recommended to whom; which commercial content appears beside the user’s material; whether a restriction carries a stated reason or an end; and how the user’s activity is attributed. A *signed change* is a pair  $(r, \varepsilon)$  with  $r \in R_u$

and  $\varepsilon \in \{+, -\}$ , meaning that  $r$  becomes present (+) or absent (−). Let  $\Delta_u = R_u \times \{+, -\}$ , and let  $\tau : R_u \rightarrow \mathcal{T}$  assign each relation a type.

**Definition 2.2** (Interface and effect set). Let  $\mathcal{A}$  be the set of actions the platform exposes to  $u$ , and  $\mathcal{A}^* \supseteq \mathcal{A}$  the set of actions that are technically feasible for the platform to expose. What  $u$  can select is represented by  $\mathcal{A}$ , and what the platform can vary by  $R_u$ . Each action  $a \in \mathcal{A}^*$  has an *effect set*  $E(a) \subseteq \Delta_u$ , the signed changes the platform actually makes.

## 2.2 Surplus and deficit

**Definition 2.3** (Surplus, deficit, and mismatch). For an action  $a \in \mathcal{A}^*$  and a set  $S \subseteq \Delta_u$  of changes the user intends, define

$$U(a, S) = E(a) \setminus S, \quad D(a, S) = S \setminus E(a), \quad M(a, S) = U(a, S) \cup D(a, S).$$

The *surplus*  $U$  is what the action does beyond the intention, and the *deficit*  $D$  is what the intention required and the action does not do. For  $a \in \mathcal{A}$  with intended set  $S_a$  write  $U_a$ ,  $D_a$ , and  $M_a$ . Write  $U^+$  and  $U^-$  for the elements of  $U$  with sign + and −.

No assumption is made that  $S_a \subseteq E(a)$ . An interface can fail an intention in two independent ways: by overreach ( $U_a \neq \emptyset$ ) and by underrealization ( $D_a \neq \emptyset$ ).

**Definition 2.4** (Forms of mismatch). An action  $a$  exhibits *unchosen bundling* if  $U_a \neq \emptyset$ . It exhibits *compelled association* if  $U_a^+ \neq \emptyset$ , and *compelled dissociation* if  $U_a^- \neq \emptyset$ . It exhibits a *shortfall* if  $D_a \neq \emptyset$ .

Publishing a post that is displayed with an advertisement adds the relation “ $u$  is adjacent to this advertisement” and is compelled association. Blocking a person to revoke commenting removes that person’s view of  $u$  and is compelled dissociation. A restriction that bars commenting but carries no stated reason or end, when the user intended both, is a shortfall.

*Remark 2.5.* No relation between surplus effects and intended changes is modeled beyond set difference. That a single placement may attach at once to authorship, publication, identity, and continued participation affects no verdict below, and so no incidence structure is defined.

## 2.3 Attribution, value, and importance

Mismatch alone is not condemnatory, since every software action has incidental effects and imperfect realizations. Five functions select the mismatch that matters: attribution and association by observers, the user’s value for a retained relation, the user’s aversion to an added relation, and the user’s importance for an intended change. Attribution is descriptive and depends on context. The parameters are normative and depend on the kind of claim at stake.

**Definition 2.6** (Attribution and association). For an observer  $o$ , a signed change  $b$ , and interface and interpretive conditions  $c$ , let

$$J_o(b, u \mid c) \in [0, 1]$$

be the degree to which  $o$  attributes  $b$  to the choice, approval, or normative stance of  $u$ , and let

$$A_o(b, u \mid c) \in [0, 1]$$

be the degree to which  $o$  links  $b$  with  $u$  or with  $u$ ’s content, without necessarily attributing choice or approval to  $u$ . Each action  $a$  has a *standard presentation context*  $c_a$ , the way its result is ordinarily displayed, including what shares its frame, what is adjacent to it in space, and what follows it in the viewing sequence.

*Remark 2.7.* Association is distinct from attribution. A viewer may know that a user did not choose an advertisement and still experience it as linked to the user's content, because it shares a frame with it or follows it immediately.  $J_o$  records the first judgment and  $A_o$  the second. Whether associative linkage without attributed agency should count as consequential is a normative question, addressed by the thresholds below.

**Definition 2.8** (Claim types and parameters). Let  $\kappa(b)$  classify the kind of attribution at issue for a signed change  $b$ , for example commercial endorsement, interpersonal rejection, institutional approval, or mere causal contribution. For each type  $\kappa$  fix

- a class  $\mathcal{O}_\kappa$  of *qualified observers*, whose readings are counted for claims of that type, with a probability measure  $\mu_\kappa$  on it;
- an attribution threshold  $\theta_\kappa \in (0, 1]$  and an association threshold  $\theta_\kappa^A \in (0, 1]$ ; and
- a proportion  $\rho_\kappa \in (0, 1]$ .

For each relation type  $t \in \mathcal{T}$  fix de minimis levels  $\lambda_t > 0$  for losses and shortfalls and  $\lambda_t^+ > 0$  for aversion.

**Definition 2.9** (Value, aversion, and importance). For a removed relation  $b = (r, -)$ , let  $V_u(r) \geq 0$  be the value  $u$  places on retaining  $r$ . For an added relation  $b = (r, +)$  and a context  $c$ , let  $W_u(r | c) \geq 0$  be the cost to  $u$  of acquiring or sustaining  $r$  under  $c$ . For an intended change  $s \in S$ , let  $I_u(s) \geq 0$  be the importance  $u$  places on  $s$  being made.

*Remark 2.10.* The parameters  $\theta_\kappa$ ,  $\theta_\kappa^A$ ,  $\rho_\kappa$ ,  $\lambda_t$ ,  $\lambda_t^+$  and the choice of qualified observers are openly normative. They are not psychological constants, and no single value is defended for all types. Reasonableness of a reading enters through the choice of  $\mathcal{O}_\kappa$  and the thresholds, while  $J_o$  records only what observers do attribute. The essay does not estimate the parameters. Declaring them makes clear why one eccentric reading or a trivial inconvenience does not condemn an interface.

**Definition 2.11** (Consequential mismatch). Let  $a \in \mathcal{A}^*$  and  $S \subseteq \Delta_u$ , and let  $c_a$  be the standard presentation context of  $a$ . An element of  $M(a, S)$  is *consequential* if it satisfies one of the following.

1. (*misattribution*) It is some  $b \in U(a, S)$  with  $\kappa = \kappa(b)$  and

$$\mu_\kappa(\{o \in \mathcal{O}_\kappa : J_o(b, u | c_a) \geq \theta_\kappa\}) \geq \rho_\kappa.$$

2. (*association*) It is some  $b \in U(a, S)$  with  $\kappa = \kappa(b)$  and

$$\mu_\kappa(\{o \in \mathcal{O}_\kappa : A_o(b, u | c_a) \geq \theta_\kappa^A\}) \geq \rho_\kappa.$$

3. (*loss*) It is some  $b = (r, -) \in U(a, S)$  with  $V_u(r) \geq \lambda_{\tau(r)}$ .

4. (*aversion*) It is some  $b = (r, +) \in U(a, S)$  with  $W_u(r | c_a) \geq \lambda_{\tau(r)}^+$ .

5. (*shortfall*) It is some  $s = (r, \varepsilon) \in D(a, S)$  with  $I_u(s) \geq \lambda_{\tau(r)}$ .

*Remark 2.12.* Aversion complements attribution and loss and does not replace them. An added relation can be consequential because observers read it as the user's act, or because the user objects to sustaining it, for instance by contributing content, attention, or legitimacy to a system the user rejects, even when no observer reads it as endorsement. The threshold  $\lambda_t^+$  prevents every disliked side effect from becoming consequential. Establishing consequence does not establish a defect, since feasibility and unavailability remain to be shown. The definitions are stated for a single user; platform-level evaluation would require a distribution over user types, which the essay leaves for later work.

## 2.4 Realizability and defective actions

**Definition 2.13** (Clean realization). An intended set  $S \subseteq \Delta_u$  is *cleanly realizable* at an interface  $\mathcal{B}$  if there is  $a' \in \mathcal{B}$  such that  $M(a', S)$  contains no consequential element. Exact realization  $E(a') = S$  is the special case with empty mismatch.

*Remark 2.14.* Exact realization is too strong a requirement, since a benign log entry would defeat it. Cleanliness is relative to consequence, and consequence is defined through  $J_o$ ,  $V_u$ , and  $I_u$  without reference to realizability, so the definition is not circular. An intended set for which no action in  $\mathcal{B}$  makes the required changes is simply not cleanly realizable there.

**Definition 2.15** (Withheld intention and defective action). An intended set  $S$  is *withheld* at  $\mathcal{A}$  if it is cleanly realizable at  $\mathcal{A}^*$  but not at  $\mathcal{A}$ . An action  $a \in \mathcal{A}$  is *defective* if  $S_a$  is withheld at  $\mathcal{A}$ . A defective action is a *defective bundle* if a consequential element of  $M_a$  lies in  $U_a$ , and a *defective shortfall* if one lies in  $D_a$ .

Because  $a$  itself is available and does not realize  $S_a$  cleanly, a defective action has at least one consequential element of  $M_a$ , so it is a defective bundle, a defective shortfall, or both. The withheld clause is the counterfactual criterion. It condemns an interface only when the platform could have offered an action that makes the intended changes without consequential mismatch and did not. Unavoidable side effects and unavoidable limits are excluded.

## 2.5 Exit

**Proposition 2.16** (Cost of the intended change). If  $S$  is withheld at  $\mathcal{A}$ , then every  $a' \in \mathcal{A}$  has a consequential element in  $M(a', S)$ .

*Proof.* Otherwise some  $a' \in \mathcal{A}$  would cleanly realize  $S$  at  $\mathcal{A}$ , contradicting that  $S$  is withheld.  $\square$

The user can therefore obtain  $S$  only by accepting some consequential mismatch, surplus or deficit, which need not be that of the original action. The only way to avoid consequential mismatch altogether is to forgo  $S$ , by refraining or leaving the platform. This is a condition on the interface, not a claim of legal compulsion. The availability of exit is not evidence that a clean realization exists.

# 3 Applying the Formalism

The definitions are worth having only if they discriminate. A framework that condemned every interface would be a grievance in notation. This section applies the defective-action test to four cases and reports where it returns a verdict, where it withholds one, and where the case falls outside the framework.

## 3.1 The verdict rule

An action  $a \in \mathcal{A}$  is judged defective only if three findings are made.

1. **Consequence.**  $M_a$  contains a consequential element: material attribution or association among a sufficient share of qualified observers, loss or aversion above the de minimis level, or a shortfall in an intended change that the user rates above that level.
2. **Feasibility.**  $S_a$  is cleanly realizable at  $\mathcal{A}^*$ . A comparator system that already offers such an action is the strongest evidence.

3. **Unavailability.**  $S_a$  is not cleanly realizable at  $\mathcal{A}$ . This is a claim about documented controls and should be dated.

A case that fails the first finding is not a mismatch problem. A case that fails the second is a limitation of the technology and not of the interface. A case that fails the third is already served. The third finding requires the first for every available action, so evidence of consequence is the shared bottleneck. A verdict is *undetermined* when a required finding has not been made.

### 3.2 Four cases

| Case                                                     | Intended $S_a$                                                            | Mismatch                                                                                                                                                              | Claim type $\kappa$                   | Provisional verdict                                                                                                                        |
|----------------------------------------------------------|---------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------|
| Block or profile controls to stop comments               | Bar one person's comments on public posts; state the norm; set a duration | Block removes their view of public activity ( $U^-$ ). Silent tools give no notice, reason, or duration ( $D$ )                                                       | Interpersonal rejection               | Provisionally defective on personal profiles: feasibility shown on other Meta surfaces; unavailability rests on the documentation reviewed |
| Ad or recommended content beside or after a post or Reel | Publish the post or Reel                                                  | Advertisement placed among the author's posts, overlaid on a Reel, or shown immediately after it, and platform-selected content that follows it in sequence ( $U^+$ ) | Commercial endorsement or association | Undetermined: feasibility supported, unavailability provisionally supported, consequence unmeasured                                        |
| Presence as support                                      | Publish the post                                                          | Use of the post for recommendation or growth beyond the chosen audience ( $U^+$ )                                                                                     | Institutional approval                | Consequential for some user types by aversion. Defect status undetermined: feasibility and unavailability unverified                       |
| Reels delivered to a viewer                              | None (no user action)                                                     | None                                                                                                                                                                  | Not applicable                        | Outside the framework                                                                                                                      |

Table 1: Verdicts under the defective-action test. All verdicts are provisional on the evidence noted in the text.

**Blocking and the controls around it.** The intended change is a restriction on one person's participation, and in the correction case it includes stating the norm and an end to the restriction. A block removes more than participation: the surplus is a loss of perception, and the user may value the perception that a block removes, since the point of removing a comment is often to keep the surrounding conduct visible. Comparators show that a clean action is technically feasible, including within Meta's own products. Discord's timeout lets a member read while barring them from sending messages, adding reactions, or joining voice or video, records a reason, and can be lifted early [1]. Facebook Groups let administrators temporarily suspend a member from posting, commenting, and

reacting [6]. Facebook Pages can ban a person from posting, liking, and commenting while the public Page remains visible to them [7]. Instagram’s Restrict hides a person’s comments from everyone but the commenter unless the account owner approves them [5]. These are moderator, administrator, or owner powers on different surfaces, so they establish feasibility and not a like-for-like feature of personal profiles.

For personal Facebook profiles the controls found are of a different kind. Meta’s Help Center describes the Restricted list as limiting what a person sees, namely public information and posts in which they are tagged, and not private posts [4]. That limits perception, the opposite of the separation sought. Users can also limit who may comment on public posts to Public, Friends, or those mentioned [21], which sets an audience and not a person. Several secondary guides describe a Facebook Restrict that hides a restricted person’s comments. That behavior is documented for Instagram’s Restrict, and the guides appear to conflate the two, so the essay does not rely on them. On the documentation reviewed, no per-person action bars commenting on a personal profile while preserving view, so the third finding is supported for that intention, provisionally, since Meta’s Help Center could not be searched exhaustively.

The shortfall issue applies to the tools that exist elsewhere. Instagram’s Restrict is silent: the restricted person is not told, and no reason is stated [5]. No duration is described either, and removal is manual, so it is an indefinite reversible state and not a time-bounded suspension with automatic restoration. A silent, open-ended separation conceals and does not correct, so notice, explanation, and a bounded duration are intended changes that it does not make, and they belong to  $D_a$ . The verdict for personal profiles is therefore a defective bundle, provisionally, with a possible defective shortfall in any silent replacement.

**Advertising adjacency.** Advertisements can be adjacent to a user’s content in three documented ways. Advertiser-facing sources describe a Facebook profile feed placement that shows advertisements in between the posts of a person’s profile [8], and Meta’s brand-safety tools let advertisers keep their advertisements off specific profiles [9]. On Facebook Reels, overlay banner ads are layered over the video, at the bottom of the Reel, and post-loop ads are video ads shown between Reels [13]. The overlay shares the frame of the user’s content, and the post-loop advertisement follows it immediately, so adjacency here is spatial, in-frame, and temporal. An overlay layered on the video is displayed on the screen, so a screenshot of the Reel captures it, and the adjacency travels with any copy shared outside the platform. A screenshot on file shows the case: a banner labeled “Ad” over the lower part of a Reel, above the creator’s name, audio credit, and follow control. The banner carries a control that lets the viewer dismiss it, and an ordinary author has no documented equivalent.

Sequence adjacency is not confined to advertisements. A Reel is followed in the viewing sequence by whatever the recommendation system selects next, and Meta’s Reels stream is described as algorithmically recommended [13]. The author has not chosen what follows, and where the following item comes from a source that neither the author nor the viewer follows, neither has chosen the connection. This adds a surplus of the same form, association with platform-selected content, and it can be consequential by aversion for authors who object to what tends to follow, whether or not any viewer attributes the following content to them. A user may decline to publish Reels for this reason. Reports of that kind are the stated reasons that the third prediction of the withdrawal section would collect, and they are not evidence for the hypothesis by themselves. Whether an author can publish a Reel that is not followed by platform-selected content was not found and is not shown here.

Whether such adjacency is consequential depends on how viewers read it, and two questions should be kept apart. The first is attribution: whether viewers take the author to have chosen

or endorsed the advertisement, which is a claim about  $J_o$ . The second is association: whether viewers experience the advertisement as linked to the content even while knowing the author did not choose it, which is a claim about  $A_o$ . The likely finding is that attribution of choice is low and association is not, and if so the second is where the consequence lies. Neither has been measured. An industry-cited survey reports that a majority of respondents associate a brand with the content around its advertisements on social media [12]. That concerns the association of a brand with content and not of an author with an advertisement, so it is suggestive of transfer through adjacency and does not measure the quantity at issue. The advertising industry’s own rationale for adjacency controls is that association is real and matters, and that rationale applies to authors as well as brands.

Feasibility is supported from the advertiser side. Meta documents inventory filters with Expanded, Moderate, and Limited settings, publisher and content block lists, and third-party content block lists for Feed and Reels placements [3]. Some creators have a control of their own: a 2022 report describes an opt-out of banner ads in Reels through Creator Studio for creators in Meta’s in-stream monetization program, with placement of sticker ads left to the creator and selection of the advertisement left to Facebook [14]. That report is dated and secondary, but it indicates that an adjacency control for content owners has been built for monetized creators. For ordinary users, Meta’s Help Center states that ads cannot be blocked entirely [10], and that ad preferences change which ads a person sees and not how many [11]. The same page lists a setting that adjusts who can see a person’s social interactions alongside ads, an existing control over one association between a user and advertising. No control that lets an ordinary profile owner shape the advertising placed beside their own posts or Reels was found. The search was not exhaustive, and other available actions, such as publishing to a restricted audience, have not been checked for whether they avoid adjacency. The asymmetry the essay can defend is narrow and stated as an absence of documented control: advertisers can decline particular profiles and some monetized creators can opt out of overlays, while no symmetrical control for ordinary posters was found. The verdict is undetermined: consequence is unmeasured, and unavailability rests on absence in the documentation reviewed.

**Presence as support.** Posting is itself a contribution of content and attention to the platform, and that contribution is partly constitutive of using it, so it is not obviously an effect bundled with publication. The narrower candidate surplus is the use of the post for recommendation, promotion, growth, or training beyond the audience the author selected. Two separations should be distinguished here. One keeps a post from being distributed as a recommendation to people who do not already follow the author. The other keeps interactions with the post from serving as signals in recommending other things. A platform may support one without the other, and the effect set should say which is meant. Observer attribution is unlikely to establish consequence here: the claim type is institutional approval, where the threshold is plausibly high and observers may not read participation as endorsement. Aversion can establish it for some users. A user who objects to sustaining a platform they reject bears an added relation whether or not anyone reads it as endorsement, and for users whose aversion exceeds the *de minimis* level the surplus is consequential. Consequence for some user types does not settle a defect. A clean action, such as publishing to a chosen audience without the material entering recommendation, may exist. Meta already sets eligibility for recommendation by category of content under its Recommendation Guidelines [15], so a content-level eligibility gate exists at platform level, but a user-set opt-out for one’s own posts was not found, and its availability is not shown here. The framework therefore declines to condemn the case, which shows that the test does not prove too much.

**Reels delivered to a viewer.** A recommendation stream delivers content to the viewer. No action of the viewer is paired with a mismatch, so there is nothing for  $S_a$ ,  $U_a$ , and  $D_a$  to describe. It is an imposition on perception and is excluded by the definitions. The framework thereby separates mismatch in a user’s own actions from unsolicited importation, a topic treated in the author’s essays on proxy capture in platform interfaces, “Safe Mode” and “Affect Before Architecture”. The exclusion applies to the viewer. Publishing a Reel is a user action, and the overlay and post-loop advertisements that accompany it, together with the platform-selected content that follows it, belong to the adjacency case above. Interaction with a delivered item, such as watching, can be modeled as an action with a training-signal surplus, but that analysis is not pursued here.

### 3.3 A condition for the ordering

The thresholds and proportions are normative, so it is worth stating exactly what the ordering of two verdicts depends on.

**Proposition 3.1** (Monotone ordering). Suppose that, over a common population of qualified observers, the attribution distribution for advertising adjacency first-order stochastically dominates that for presence as support, so that for every  $t$  the share of observers with  $J_o \geq t$  is at least as large for adjacency. If  $\theta_{\text{comm}} \leq \theta_{\text{inst}}$  and  $\rho_{\text{comm}} \leq \rho_{\text{inst}}$ , then whenever presence as support is consequential by misattribution, advertising adjacency is as well.

*Proof.* Let  $F_{\text{ad}}(t)$  and  $F_{\text{pr}}(t)$  be the shares with  $J_o \geq t$ . Then  $F_{\text{ad}}(\theta_{\text{comm}}) \geq F_{\text{ad}}(\theta_{\text{inst}}) \geq F_{\text{pr}}(\theta_{\text{inst}}) \geq \rho_{\text{inst}} \geq \rho_{\text{comm}}$ , using that  $F$  is non-increasing, dominance, the assumption that presence is consequential, and the ordering of the proportions.  $\square$

The proposition rests on two assumptions: that adjacency is at least as attributable as mere presence in a common observer population, and that institutional approval requires at least as much attribution and agreement as commercial endorsement. Under them, no choice of parameters reverses the ordering. It does not show that the two cases receive different verdicts, since both may be consequential or neither. Whether the assumptions hold, and how large the attribution gap is, are empirical and normative questions that the essay does not settle.

### 3.4 What remains to be established

The verdicts depend on findings not made in this essay: primary-source confirmation of the scope of Facebook’s controls on personal profiles, since Meta’s Help Center was only partly retrievable; viewer-side confirmation that advertisements are placed among ordinary users’ posts and after their Reels (an overlay on a Reel is shown by a screenshot on file); and a small screenshot study of how viewers read them, separating four readings (the advertisement was selected by the author, endorsed by the author, merely associated with the author, or selected entirely by the platform), which are different claims about  $J_o$  and  $A_o$ ; whether any available action, such as restricted-audience publishing, already avoids adjacency; whether a user-level opt-out of recommendation use exists or is feasible; whether an author can publish a Reel that is not followed by platform-selected content; and the choice of parameters  $\theta_\kappa$ ,  $\rho_\kappa$ ,  $\lambda_t$ , and  $\lambda_t^+$ . The table is offered as a structure for those findings and not as their result.

## 4 Separating Participation from Perception

The applications section found that a block removes perception along with participation, that the controls found on personal Facebook profiles do not separate the two, and that other Meta surfaces

do, silently in one case. This section describes what a transparent alternative looks like, in order to show that clean realization is a concrete design target and not an abstraction.

#### 4.1 Two kinds of relation

Partition the relations that concern another person  $v$  into two classes.

**Definition 4.1** (Participation and perception). Let  $R_u^{\text{part}}(v)$  be the relations by which  $v$  may act within  $u$ 's space: commenting, replying, tagging, messaging, and mentioning. Let  $R_u^{\text{perc}}(v)$  be the relations by which  $v$  may observe  $u$ 's space: viewing posts, viewing responses, and viewing the norms in force.

The partition is analytic. It says that the two classes are independently controllable in principle, and it identifies which relation a given action removes. A block removes both classes at once. The constructive proposal is an action that removes the participation the user intends to bar and preserves perception.

**Definition 4.2** (Asymmetric separation). Let  $P \subseteq R_u^{\text{part}}(v)$  be the participation relations the user intends to bar. An action  $a$  effects *asymmetric separation of  $v$  with respect to  $P$*  if

$$E(a) \supseteq \{(r, -) : r \in P\} \quad \text{and} \quad E(a) \cap \{(r, -) : r \in R_u^{\text{perc}}(v)\} = \emptyset.$$

Defining separation relative to  $P$  matters: the intention may be only to stop commenting, and an action that also removes messaging, tagging, and mentioning would bundle unwanted loss with the intended one.

Where a comparator exists, it supports feasibility. Discord's timeout, as documented by the service, leaves a member able to read while barring them from sending [1]. The comparator is a moderator power and not a user power, so it shows that the separation can be built and operated at scale and does not by itself show that a personal-profile equivalent has been deployed elsewhere.

*Remark 4.3.* Comparators on other Meta surfaces show that the separation is buildable within one company's products: temporary suspension of group members [6], bans on Pages that leave the Page visible [7], and Instagram's Restrict, which hides a person's comments from others unless approved [5]. On personal Facebook profiles the controls found differ. The Restricted list limits what a person sees, namely public information and tagged posts [4], and comment limits set an audience and not a person [21]. Secondary guides that describe a Facebook Restrict hiding comments appear to conflate it with Instagram's feature. Instagram's Restrict is silent, so the person is neither told which norm was applied nor asked to revise, and no reason or duration is described. It is an indefinite reversible state and not a time-bounded suspension, and so it is a concealment tool more than a correction tool.

#### 4.2 What the user values

The case for retained perception does not require that seeing good conduct reforms the person who sees it. That is an empirical claim and may be false in many cases. The framework asks only that the user place value on the retained relation, that is,  $V_u(r) \geq \lambda_{\tau(r)}$  for  $r \in R_u^{\text{perc}}(v)$ . A user may value it for several reasons that do not depend on reform: keeping a shared space open, avoiding the appearance of hiding from a person, or preserving the possibility of later reconciliation. The reform hypothesis is one reason among these and is not a premise.

### 4.3 A graduated moderation sequence

Moderation in ordinary social life is graduated. A person is asked to stop, then asked to leave the conversation, then asked to leave the gathering, and only in the last case excluded from the community. The interface can be described by the same structure. An illustrative sequence for conduct in comment spaces has two parts.

*Corrective stages*, which differ in how they treat an utterance and not in what access the person retains:

1. **Pre-posting prompt.** Before submission, a notice states the norm in force and asks the author whether to revise. No relation is removed.
2. **Return with explanation.** A posted comment is returned to its author with a stated reason and a request to remove or revise it. The comment is hidden pending response.
3. **Removal with notice.** The comment is removed and a public note states the norm applied, without naming the author.

*Access restrictions*, which differ in how much participation or perception they remove:

4. **Temporary thread exclusion.** The person may not comment on the user's posts for a stated period, while retaining view.
5. **Restricted interaction.** A longer or wider exclusion of participation, again preserving view.
6. **Block.** All relations are removed, reserved for stalking, threats, persistent harassment, or cases where visibility itself creates danger.

The corrective stages form a workflow rather than a chain, and the access restrictions form a ladder that can be ordered exactly.

**Definition 4.4** (Restriction set and ladder). For an access restriction  $a$  and target  $v$ , let

$$Q_v(a) = \{ r \in R_u^{\text{part}}(v) \cup R_u^{\text{perc}}(v) : (r, -) \in E(a) \}$$

be its *restriction set*, the participatory and perceptual relations it revokes. Order access restrictions by  $a \preceq a'$  if  $Q_v(a) \subseteq Q_v(a')$ . A *restriction ladder* is a chain  $a_0 \preceq a_1 \preceq \dots \preceq a_n$  in this order. Notices, reasons, and durations do not enter  $Q_v$ ; they are intended changes in their own right.

*Remark 4.5.* Complete effect sets are not ordered by inclusion. A prompt adds a warning, a removal deletes an utterance already made, a timeout disables future participation, and a block removes participation and perception, so their full signed effect sets are generally incomparable. Ordering by restriction sets isolates the dimension along which access severity increases. The corrective stages are not ordered by  $Q_v$  because they restrict an utterance and not a person's access. Enlarging  $Q_v$  to utterance-level relations and time-indexed relations would order them, at a cost in machinery that the argument does not need.

**Definition 4.6** (Local collapse). For an intended restriction  $S \subseteq R_u^{\text{part}}(v) \cup R_u^{\text{perc}}(v)$ , let  $\mathcal{A}_S = \{a \in \mathcal{A} : S \subseteq Q_v(a)\}$  be the available access restrictions that cover it. If  $\mathcal{A}_S = \emptyset$ , the restriction is unrealizable at  $\mathcal{A}$ . Otherwise, for  $a \in \mathcal{A}_S$  the *excess restriction* is

$$\Sigma(S, a) = Q_v(a) \setminus S.$$

The restriction  $S$  *collapses* at  $\mathcal{A}$  if  $\mathcal{A}_S \neq \emptyset$  and  $\Sigma(S, a) \neq \emptyset$  for every  $a \in \mathcal{A}_S$ . The intended restriction is then forced upward, and  $\Sigma$  measures how far.

*Remark 4.7.* Collapse is local to an intended restriction. An interface may offer several finer controls and still collapse for a particular intention, because no offered action covers that intention without excess. This avoids the false claim that an interface offers only the top rung. Collapse is the overreach failure of the access restrictions, and the omission of notice, reason, or duration is the underrealization failure. When  $S \subseteq R_u^{\text{part}}(v)$  collapses and every covering action has excess containing a perceptual relation  $r$  with  $V_u(r) \geq \lambda_{\tau(r)}$ , the excess is consequential loss, and if a covering action with no consequential excess is feasible at  $\mathcal{A}^*$ , the intention is withheld in the earlier sense.

#### 4.4 Norm-stating removals

The present arrangement also privatizes the lesson. When a user blocks someone, other participants see only that the disruption ended. They do not learn which norm was applied, and the person moderated cannot observe how the conversation proceeds. A removal notice of the form “a comment was removed because it prescribed another participant’s private conduct instead of addressing the topic” would make the norm public without exposing the individual. Moderation then also articulates a rule, which is a relation the community may want and cannot obtain from a silent block. This is an effect of a third kind, neither association nor dissociation: the production of public institutional memory. It can enter the intended set as an added relation, and an interface whose only covering action is a silent block cannot realize it.

#### 4.5 Limits

Three limits should be stated. The classification of a comment as addressing character rather than argument is contestable and error-prone, so pre-posting prompts risk false positives and paternalism; the proposal is a design possibility and not a claim about the accuracy of any classifier. Graduated systems have costs in complexity and in the labor of the person applying them, which returns to the simplicity objection above. And asymmetric separation does not prevent circumvention, since a restricted person may route around it through direct messages or another account, though neither does a block. The section does not claim that the sequence is optimal. It claims that the lower stages are realizable, that their absence is a design fact, and that their absence is what turns each mild intention into an exile.

### 5 Objections and Replies

Five objections are considered. Each is stated in its strongest form. Two of the replies concede something, and the concessions are carried into the framework and not merely acknowledged.

#### 5.1 “A private platform compels nothing; users can leave”

The objection is correct about doctrine. Constitutional compelled association requires state action, and a user who joins a private service on its terms is not compelled in that sense. The essay does not claim otherwise, which is why it uses *unchosen bundling* as its general term.

The objection does not reach the design claim. That a user may exit does not show that the interface has offered the separations the user needed. Under the exit proposition, when an intention is withheld, every available action carries consequential mismatch with it, and the only way to avoid all of it is to forgo the change. Exit is then the cost the interface imposes, and its availability

describes the problem without answering it. Nothing here depends on the exit being difficult. The claim holds even where leaving is easy.

## 5.2 “Advertising is the business model, and disaggregation costs revenue”

The objection is partly right, and the framework should say so. The defective-action test asks whether a cleaner action is technically feasible, not whether providing it is costless. A platform may have sound commercial reasons for withholding an action.

Two replies apply. First, feasibility is not merely asserted. Where a comparator exists, including advertiser-side controls over adjacency, the capability has been built and deployed. Second, a commercial justification is a reason for the design, and the test classifies the design without denying the reason. The framework therefore *classifies* bundles and does not by itself *mandate* their removal. Whether a defective action should be corrected is a further question that weighs the user-side cost of the bundle against the platform-side cost of the separation. The essay reports the first term and does not claim to have measured the second.

## 5.3 “Simple interfaces are a virtue”

A single block button is easy to understand. Each added control adds options, documentation, and error, and technical feasibility does not oblige a designer to expose every action.

The objection identifies a real cost that the definitions omit:  $\mathcal{A}^*$  counts what can be offered, not what can be offered without degrading usability. The reply is that simplicity and disaggregation are not opposed by nature. Bundling is a simplification that hides a choice, and the hidden choice is the one the user needs to make. Defaults, presets, and graduated escalation are familiar ways to reduce the burden of finer controls, so that the finer action is available without being demanded, although the essay does not survey them. The essay concedes, however, that the counterfactual criterion identifies where a better interface is possible and does not by itself show that it is better overall.

## 5.4 “Separating perception from participation enables stalking and brigading”

The strongest form is that a person barred from commenting but able to watch can still harass through other means, such as a second account, screenshots, or coordination elsewhere. Blocking exists precisely because visibility can itself be a harm.

The framework accommodates this and does not hide it. The loss criterion requires that the user place positive value on the retained relation. When a user wants the other person to lose sight of them, removing perception belongs to the intended set  $S_a$  and is not surplus. A block used against stalking or threats is therefore not a defective action, and the essay says that block remains the correct action in those cases. The circumvention objection also applies equally to blocks, which second accounts likewise evade, so it does not favor the coarser control. What the framework requires is that the two purposes be distinguishable, so that the user who wants only to bar participation is not forced into the remedy suited to the user who needs to disappear.

## 5.5 “The thresholds are arbitrary, so any verdict can be produced”

The parameters  $\theta_\kappa$ ,  $\rho_\kappa$ , and  $\lambda_t$  are normative and unestimated, and a framework that returns whatever verdict its author wants would be worthless.

Three replies apply. First, the parameters are declared and typed by claim, so a verdict can be traced to a stated choice. Second, the framework does not claim a verdict for every case. The applications section returns different results for different cases and excludes one, which is not the

behavior of an instrument that condemns everything. Third, the appropriate test of the result is whether the ordering of cases survives changes in the parameters. The applications section states a condition: if attribution for adjacency stochastically dominates attribution for presence, and the thresholds for institutional approval are at least as demanding as those for commercial endorsement, then the ordering cannot reverse. Whether those assumptions hold is an empirical and normative question that the essay does not settle. If they fail, the ordering may flip, and the essay should then report that.

## 5.6 What the replies change

Two concessions carry forward. The framework classifies mismatches and does not itself mandate their removal, which affects how the design implications are framed. And the counterfactual criterion identifies possible improvement without establishing that the improvement is net beneficial. The safety objection is handled inside the formalism, not by exception, through the requirement that the user value the relation that is lost.

## 6 Withdrawal

If some actions are defective, users who cannot disaggregate the meanings attached to them have three options: accept the surplus, act less, or leave. The last two are forms of withdrawal. This section states a hypothesis about withdrawal and is careful to separate what is observed from what is proposed.

### 6.1 Three levels

**Observed decline in contribution.** Ofcom’s *Adults’ Media Use and Attitudes* report, published 2 April 2026 [2], finds that 49 percent of UK adult social media users now post, share, or comment, down from 61 percent in 2024. Ofcom links the decline to video-focused features and to concern about the long-term consequences of past posts, and the share of adults worried that their posts could cause them problems later rose from 43 to 49 percent. This supports the claim that the dependent variable is real and measurable. It does not identify a cause, and the explanations Ofcom offers are competitors to the bundling mechanism proposed here and not instances of it. Ofcom’s measure is also the share of users who post at all, whereas the hypothesis below concerns how often a user performs a given action, so the data bear on the first margin and not the second.

**A proposed mechanism.** The hypothesis is that part of the decline is produced by bundling. Where an action carries consequential surplus that the user cannot separate, the user’s expected cost of performing it rises, and the rate of performing it falls.

**Hypothesis (withdrawal).** As the share of a user’s intended actions that are defective rises, and as attribution  $J_o$  for their consequential surplus rises, the rate at which the user performs those actions falls.

**The untested connection.** No evidence presented here links the mechanism to the decline. Competing explanations are plausible and should be listed: a shift from posting to passive consumption, algorithmic ranking that reduces the reach of ordinary posts, privacy and permanence concerns unrelated to bundling, fatigue, and migration to messaging. The hypothesis is one candidate among these.

## 6.2 What would test it

The hypothesis makes differential predictions that the competing explanations do not all share.

1. **Differential decline by action type.** Actions with more consequential mismatch, such as posting where an attached advertisement is likely to be read as the author's, should decline more than actions with less.
2. **Response to new controls.** When a platform introduces an action that separates a previously bundled relation, contribution of the affected action should rise relative to comparable unaffected actions. Such introductions are natural experiments.
3. **Stated reasons.** Users who reduce an action should report the unwanted association or lost relation as a reason more often than users who do not, controlling for age and platform. One kind of stated reason to record is an objection to the platform-selected content that follows a user's Reels.

Failure of these predictions would count against the hypothesis. A uniform decline unrelated to bundling would favor the competing explanations.

## 6.3 The commercial dimension

If the hypothesis holds, a platform that bundles advertising into user posts draws short-term inventory from the context it appropriates and gradually reduces the willingness to produce the content that makes the inventory valuable. This is a claim about cost to the platform and is conditional on the hypothesis. It is stated to indicate why the bundling may not serve the platform's own interest, not to assert that it does not.

## 7 Design Implications

The framework classifies bundles and does not by itself mandate removing them. What follows are candidate separations. Some are supported by comparators discussed in the earlier sections and others are not yet shown to be feasible. Their user-side benefit is described, and their platform-side cost is not measured here.

1. **Transparent participation exclusion that preserves view, on personal profiles.** An action that bars commenting, replying, tagging, and messaging by a chosen person while leaving public material visible, and that carries a stated reason and a duration. Facebook Groups already offer time-limited suspension [6], Pages offer bans that leave the Page visible [7], and Instagram offers a silent Restrict [5]. Personal profiles offer block, a perception-limiting Restricted list [4], and audience-wide comment limits. The separation should be explicit, so that it serves correction and not only concealment. Block is retained for stalking, threats, and cases where visibility itself creates danger.
2. **A graduated moderation sequence for conduct.** Pre-posting prompts, return with explanation, removal with notice, temporary exclusion, restricted interaction, and block, so that mild intentions do not collapse to the top rung of the access restrictions.
3. **Norm-stating removals.** Public notices that state the rule applied without naming the person, so that moderation leaves a reusable account.

4. **Attribution disclosure for attached content.** A visible indication that an advertisement was placed by the platform and not chosen by the author, which may reduce  $J_o$  for commercial endorsement without requiring the user to control the advertisement. It may not reduce  $A_o$ , since association can persist when viewers know the author did not choose the advertisement.
5. **Symmetric adjacency controls.** Where advertisers manage adjacency through inventory filters and blocklists, ordinary posters could be offered a corresponding control over what appears beside their posts and Reels, subject to the platform’s commercial constraints. A version already exists for some monetized creators [14].
6. **Audience-limited publishing without recommendation use.** An option to share with a chosen audience without the material entering the platform’s recommendation of content to others. Its feasibility is not shown here.
7. **An author-level option to end a Reel’s viewing sequence without platform-selected continuation.** An option that returns the viewer to the author’s content, or ends playback, instead of advancing to content selected by the recommendation system. Its feasibility and its effect on platform metrics are not shown here.

Each item separates a relation that is currently bundled. None is claimed to be costless, and the simplicity objection applies to all of them.

## 8 Conclusion

Interfaces govern in two ways. They permit and prohibit actions, and they decide which meanings can be enacted separately. The second form of power is easy to miss because it appears in no rule. It shows in what a single act carries with it: an advertisement displayed with a post, a block that removes more than it was meant to, participation that also lends support to a system the participant may not endorse, a restriction imposed without the explanation that would make it correction.

The account developed here gives that power a form. An action has an intended set and an effect set. The difference has two parts, surplus, what the action does beyond the intention, and deficit, what the intention required and the action does not do. Mismatch matters when observers attribute it to the user, when the user loses something they value, is made to bear a relation they reject, or fails to obtain a change they rate highly. An action is defective when a cleaner action is feasible and not offered. Applied to blocking, advertising adjacency, presence as support, and recommendation, the test returns different results. It finds the first case provisionally defective on personal profiles, since other Meta surfaces show that the separation is feasible, leaves the second undetermined pending evidence on attribution and availability, finds the third consequential for some users by aversion but undetermined as a defect, and excludes the fourth.

The remedy is not better judgment by users, who are already exercising judgment within a coarse action space. It is a larger action space, with finer separations between what a user does and what the platform makes of it. Where that space is small, one remaining option is to say less, and the loss of that speech is a cost that the design imposes and that may, if the withdrawal hypothesis holds, eventually fall on the platform.

## References

- [1] Discord. “Time Out FAQ.” Discord Support. <https://support.discord.com/hc/en-us/articles/4413305239191>. Accessed 19 September 2026.

- [2] Ofcom. *Adults' Media Use and Attitudes Report 2026*. Published 2 April 2026. Report: <https://www.ofcom.org.uk/siteassets/resources/documents/research-and-data/media-literacy-research/adults/adults-media-use-and-attitudes-2026/adults-media-use-and-attitudes-2026-report.pdf>. Summary: <https://www.ofcom.org.uk/media-use-and-attitudes/media-habits-adults/passive-social-media-use-ai-companionship-and-online-side-hustles-uk-adults-media-and-online-lives-revealed>. Accessed 19 September 2026.
- [3] "Meta enhances brand suitability controls with new third-party content block lists." *Mediaweek*. <https://mediaweek.com.au/news/meta-enhances-brand-suitability-controls-with-new-third-party-content-block-lists>. Secondary report of Meta documentation; Meta's own pages not retrieved. Accessed 19 September 2026.
- [4] Meta. "Add or remove someone from your Restricted list on Facebook." Facebook Help Centre. <https://www.facebook.com/help/206571136073851>. Accessed 19 September 2026 (search excerpt; page not fetched).
- [5] "Instagram to let you restrict bullies without notifying them." *Gulf News*, 2019. <https://gulfnews.com/business/instagram-to-let-you-restrict-bullies-without-notifying-them-1.1562655284065>. Secondary report of Instagram's announcement. Accessed 19 September 2026.
- [6] "Facebook rolls out new tools for Group admins to manage their communities and reduce misinformation." *TechCrunch*, 9 March 2022. <https://techcrunch.com/2022/03/09/facebook-rolls-out-new-tools-for-group-admins-to-manage-their-communities-and-reduce-misinformation>. Accessed 19 September 2026.
- [7] Swash Labs. "How to Ban Someone From Your Facebook Page." 15 January 2021. <https://swashlabs.com/blog/ban-someone-from-your-facebook-page>. Secondary description. Accessed 19 September 2026.
- [8] NestScale. "Facebook Ad Placements: All Types in 2025 & Examples." <https://nestscale.com/blog/facebook-ad-placement.html>. Secondary description. Accessed 19 September 2026.
- [9] "Meta strengthens brand safety with new ad controls and comment muting tools." *eMarketer*. <https://www.emarketer.com/content/meta-strengthens-brand-safety-with-new-ad-controls-comment-muting-tools>. Accessed 19 September 2026 (excerpt only).
- [10] Meta. "Can you block or hide ads showing on your Facebook account." Facebook Help Centre. <https://www.facebook.com/help/146952742043748>. Accessed 19 September 2026 (search excerpt; page not fetched).
- [11] Meta. "Your ad preferences and how you can adjust them on Facebook." Facebook Help Centre. <https://www.facebook.com/help/247395082112892>. Accessed 19 September 2026 (search excerpt; page not fetched).
- [12] "Ad Adjacency: Key Considerations and Tips for 2023." Peer39. <https://www.peer39.com/blog/ad-adjacency>. Industry blog reporting a 2021 Harris Poll survey. Accessed 19 September 2026.
- [13] "Meta Expands Reels Overlay Ads to More Brands." *Social Media Today*, 17 July 2024. <https://www.socialmediatoday.com/news/meta-expands-reels-overlay-ads-more-brands/721691/>. Secondary report quoting Meta. Accessed 19 September 2026.

- [14] Osmundson, B. “Monetize Facebook Reels With Overlay Ads.” *Search Engine Journal*, 24 February 2022. <https://www.searchenginejournal.com/monetize-facebook-reels/439449/>. Secondary and dated; describes an opt-out for eligible creators. Accessed 19 September 2026.
- [15] Meta. “Recommendation Guidelines.” About Meta, 31 August 2020. <https://about.fb.com/news/2020/08/recommendation-guidelines/>. Accessed 19 September 2026.
- [16] Nissenbaum, H. “Privacy as Contextual Integrity.” *Washington Law Review* 79(1): 119–158, 2004.
- [17] Marwick, A. E., and boyd, d. “I Tweet Honestly, I Tweet Passionately: Twitter Users, Context Collapse, and the Imagined Audience.” *New Media & Society* 13(1): 114–133, 2011.
- [18] Ostrom, E. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press, 1990.
- [19] Grimmelman, J. “The Virtues of Moderation.” *Yale Journal of Law and Technology*, 2015.
- [20] *Roberts v. United States Jaycees*, 468 U.S. 609 (1984).
- [21] Search Engine Journal. “Facebook Lets Users Limit Who Can Comment On Their Posts.” <https://searchenginejournal.com/facebook-lets-users-limit-who-can-comment-on-their-posts/401097>. Secondary report; Meta documentation not retrieved. Accessed 19 September 2026.