

Whom Should a Platform Amplify?

Accuracy Mandate, User Choice, and Engagement

Zaruhi Hakobyan

University of Luxembourg

Network Science and Economics Conference
Luxembourg, October 24, 2026

Who receives attention?



Notre-Dame fire, 2019

Donald Trump

“Perhaps flying water tankers could be used.”

French Civil Security

Water-bombing aircraft could risk the building's collapse.

Visibility does not necessarily reflect expertise.

Platforms mediate source reach

WHO launches a chatbot on Facebook Messenger to combat COVID-19 misinformation

15 April 2020 | Reading time: Less than a minute (191 words)

WHO information on Messenger

WHO, April 2020

🕒 This article is more than **3 years old**

Elon Musk reportedly forced Twitter algorithm to boost his tweets after Super Bowl flop

A tweet from Joe Biden got triple the impressions than Musk's game time post, sending engineers scrambling to boost their boss

Reported Twitter amplification

The Guardian, February 2023


Research questions

If a platform chooses—or is required by regulators—to prioritize **truth-tracking accuracy**, **whom should it amplify?**

Can accuracy-oriented recommendations be implemented when users can reject them? What changes when platforms instead maximize **engagement**?

Amplifying a source both informs and influences

The platform adds source j to user i 's feed ($M_{ij} : 0 \rightarrow 1$).



- **Information channel:** An additional signal about the state θ .
- **Influence channel:** A change in coordination weights, allowing directional biases to propagate through the network.

More information need not imply greater accuracy.

The most informative source is not necessarily the optimal source to amplify.

Key takeaways

1. **Information.** In the symmetric benchmark without directional biases, broadcasting a private signal **eliminates the Morris–Shin non-monotonicity**. With scarce broadcast capacity, the optimal policy selects the **most precise sources** under the monotonicity condition.
2. **Whom to amplify.** The most informative source need not be optimal: information gains must be weighed against **network-wide bias risk**, which depends on how sources' influence overlaps.
3. **Segregation.** Integration improves accuracy in the two-community benchmark, yet users may prefer segregation. The disagreement threshold **shrinks as communities grow**.
4. **Adoption.** Even with complete information, users reject unilateral cross-camp coordination changes that improve accuracy. **Private incentives can block accuracy-improving design**.
5. **Engagement.** In the symmetric full-information benchmark, maximizing expected user utility favors **echo chambers**, whereas accuracy favors **cross-camp coordination**. The switching threshold is intervention-specific.

Model

Network. There are N users connected through a directed feed network M :

$$M_{ij} = 1 \iff \text{user } i \text{ observes source } j\text{'s signal.}$$

Information. Users are uncertain about an underlying state $\theta \in \mathbb{R}$.

► Public signal: $y = \theta + \eta$, $\eta \sim \mathcal{N}(0, \alpha^{-1})$.

► Private signal: $x_i = \theta + \varepsilon_i$, $\varepsilon_i \sim \mathcal{N}(0, \beta_i^{-1})$.

User i observes y , x_i , and the signals $\{x_j : j \in \mathcal{N}_i(M)\}$ appearing in her feed.

Bayesian updating

Under a diffuse prior,

$$\mathbb{E}_i[\theta] = \frac{\alpha y + \beta_i x_i + \sum_{j \in \mathcal{N}_i(M)} \beta_j x_j}{\alpha + \beta_i + \sum_{j \in \mathcal{N}_i(M)} \beta_j}.$$

Platform's instrument: who observes whom, not signal quality or realizations.

Model: Preferences and Coordination

Actions. Each user i chooses $a_i \in \mathbb{R}$, caring about both her preferred assessment and coordination with others.

$$u_i(a, \theta) = -(1 - r_i) \underbrace{(a_i - \theta - b_i)^2}_{\text{own preferred action}} - r_i \underbrace{\sum_{j \neq i} \phi_{ij}(M)(a_j - a_i)^2}_{\text{coordination}}$$

- ▶ **Directional preferences:** $b_i \in \mathbb{R}$. User i 's ideal action is $\theta + b_i$. In the two-camp benchmark: $b_i \in \{-b, +b\}$, $b > 0$. *Bias affects preferences, not Bayesian beliefs.*
- ▶ $r_i \in (0, 1)$: strength of the coordination motive (beauty-contest preferences).
- ▶ $\phi_{ij}(M)$: weight user i places on coordinating with j . A share q_i is assigned to feed sources, and $1 - q_i$ to other users.

Truth-tracking mandate

$$W(M) = -\frac{1}{N} \sum_{i=1}^N \mathbb{E}[(a_i - \theta)^2]$$

Expected user utility serves as a proxy for engagement.

Engagement objective

$$\Pi(M) = \frac{1}{N} \sum_{i=1}^N \mathbb{E}[u_i(a, \theta)]$$

Equilibrium: information and directional bias separate

Bayesian best response:

$$a_i = (1 - r_i)(\mathbb{E}_i[\theta] + b_i) + r_i \sum_{j \neq i} \phi_{ij}(M) \mathbb{E}_i[a_j]$$

Theorem 1: unique linear Bayes–Nash equilibrium

$$a_i = \underbrace{\sum_k c_{ik} x_k + c_{iy} y}_{\text{signal component}} + \underbrace{d_i}_{\text{directional component}}$$
$$d = \underbrace{(I - R\Phi(M))^{-1}(I - R)}_{L(M)} b$$



Signal coefficients do not depend on preferences b .

Directional components do not depend on signal precisions.

$$W(M) = - \underbrace{\bar{V}(M)}_{\text{variance}} - \underbrace{\frac{1}{N} \|d(M)\|^2}_{\text{bias risk}}$$

Broadcasting and information aggregation

Symmetric benchmark without directional biases ($b_i = 0$).




Theorem: broadcasting eliminates the non-monotonicity

$$W_B(\alpha) = W_{MS}(\alpha + \beta), \quad W'_B(\alpha) > 0.$$

When does amplification reduce accuracy?

Suppose source c is the only user with directional bias $b_c \neq 0$.



$$W(M) = -\bar{V}(M) - b_c^2 B_c(M)$$

$$\Delta W = \underbrace{\Delta V}_{\text{information}} - \underbrace{b_c^2 \Delta B_c}_{\text{bias risk}}$$

Proposition: bias threshold

If $\Delta V > 0$ and $\Delta B_c > 0$, amplification weakly improves accuracy iff

$$|b_c| \leq \sqrt{\frac{\Delta V}{\Delta B_c}}.$$

A sufficiently biased source can be harmful despite providing valuable information.

Optimal amplification in general networks

Theorem: information–bias decomposition

For any intervention $M_0 \rightarrow M_1$,

$$\Delta W = \underbrace{\Delta V}_{\text{information gain}} - \underbrace{b^\top \Delta Q b}_{\text{change in bias risk}}.$$

where $Q(M) = \frac{1}{N} L(M)^\top L(M)$ and $L(M) = (I - R\Phi(M))^{-1}(I - R)$.

$$b^\top Q(M) b = \sum_c b_c^2 B_c(M) + 2 \sum_{c < h} b_c b_h \underbrace{C_{ch}(M)}_{\text{influence overlap}}.$$

- ▶ **Source selection:** for one amplification slot, choose the candidate maximizing $S_c = \Delta V_c - b^\top \Delta Q_c b$ relative to a common baseline.
- ▶ Same-sign biases **reinforce**; opposite-sign biases can **offset** one another.

Neither precision nor individual bias alone determines optimal amplification.

Limits of network design

Theorem: irreducible directional distortion

For any feasible coordination network and common $r \in (0, 1)$,

$$\|d\|_{\infty} \geq \frac{1-r}{1+r} \max_i |b_i|.$$


Full cross-camp coordination



Attains the bound in the symmetric two-camp case

- ▶ Reallocating coordination weights can **reduce** directional distortion.
- ▶ But for $r < 1$ and $b \neq 0$, it cannot **eliminate** the distortion of all users.

Recommendations and user choice



$$M_{t+1} = \begin{cases} M' & \text{if user } i \text{ accepts,} \\ M_t & \text{otherwise.} \end{cases}$$


Theorem: an implementation constraint remains

In the symmetric two-camp echo chamber, even with *complete information*, any unilateral cross-camp coordination reallocation improves accuracy ($\Delta W > 0$) but lowers the adopting user's utility ($\Delta U_i < 0$).

Each update solves a new static equilibrium; users do not plan over future rounds.

full-information result

User choice shapes network evolution



97%

Accepted with no bias

32%


Accepted with two camps

74%

Rejected links crossing camps

Accuracy-oriented recommendations maintain higher truth-tracking accuracy than friends-of-friends and random rewiring, but **user choice filters out many cross-camp links**.

Accuracy versus engagement



Accuracy-oriented platform

Higher truth-tracking accuracy W .

Engagement-oriented platform

Higher expected user utility Π .

Engagement-oriented recommendations generate **more polarization and larger cross-camp action gaps** in the simulations.

Conclusion

1. **Source selection is a network problem.** Amplification improves accuracy when **information gains** exceed changes in **network-wide bias risk**.
2. **Accuracy-oriented design faces two limits.** Network design cannot eliminate all directional distortion; users may **reject accuracy-improving links**, even with complete information.
3. **Platform objectives shape network outcomes.** Engagement-oriented recommendations can generate **more segregation and lower accuracy**.

The network that best tracks the truth need not be the one users choose—or the platform prefers.

Network influence and robust amplification

Source c 's directional influence profile:

$$\ell_c = (1 - r_c) \sum_{k=0}^{\infty} (R\Phi)^k e_c, \quad B_c = \frac{1}{N} \|\ell_c\|^2, \quad C_{ch} = \frac{1}{N} \ell_c^\top \ell_h.$$

Robustness when $\|b\| \leq B$:

$$\Delta V \geq B^2 \max\{\lambda_{\max}(\Delta Q), 0\}$$

is necessary and sufficient for weak improvement for every admissible bias profile. **Interpretation:** a candidate source cannot be scored from individual precision and bias alone; influence profiles can interact across the network.

[◀ return to source selection](#) [conclusion](#)

The segregation threshold

Let $N = 2n$, $V_I = (\alpha + N\beta)^{-1}$, $V_S = (\alpha + n\beta)^{-1}$, $\kappa = (1 - r)/(1 + r/(N - 1))$.

$$W^I = -V_I - \kappa^2 b^2 > W^S = -V_S - b^2.$$

$$U^S = -(1 - r)V_S, \quad U^I = -(1 - r)[V_I + (1 - \kappa)^2 b^2] - \frac{4rn}{N - 1} \kappa^2 b^2.$$

$$\tilde{b}^2 = \frac{(1 - r)(V_S - V_I)}{(1 - r)(1 - \kappa)^2 + \frac{4rn}{N - 1} \kappa^2} = \frac{V_S - V_I}{r(2 - r)}(1 + o(1)).$$

With $\alpha = 30$, $\beta = 20$, $r = 0.7$, $n = 10$, we obtain $\tilde{b} \approx 0.047$.

[◀ return to network evolution](#)

[conclusion](#)

Coordination design with complete information

With complete information, let $s \in [0, 1]$ be the cross-camp coordination-weight share, and $b_i \in \{-b, +b\}$.

$$\kappa_s = \frac{1-r}{1-r+2rs}, \quad d_i = \kappa_s b_i.$$

$$W(s) = -V_I - \kappa_s^2 b^2,$$

$$\Pi(s) = -(1-r)[V_I + (1-\kappa_s)^2 b^2] - 4rs\kappa_s^2 b^2.$$

Within this symmetric coordination-weight class, $\Pi(s)$ is uniquely maximized at $s = 0$, while $W(s)$ is strictly increasing and maximized at $s = 1$.

[◀ return to recommendations](#)

[conclusion](#)

Mixed platform objectives

Mixed platform objective:

$$O_\lambda = (1 - \lambda)W + \lambda\Pi, \quad \lambda \in [0, 1].$$

For an intervention with $\Delta W < 0 < \Delta\Pi$, the platform switches at

$$\lambda^* = \frac{-\Delta W}{\Delta\Pi - \Delta W}.$$


This cutoff is **intervention-specific**, rather than a universal regulatory threshold. **Recommendation simulations**
(from submitted manuscript):

Accepted accuracy-improving recommendations, no directional bias	97%
Accepted with two camps ($b_i = \pm 0.6$)	32%
Rejected links that are cross-cutting	74%

These are simulated quantities, not the universal conclusions of theorems.

[◀ return to engagement](#) [conclusion](#)

Recommendation dynamics: additional outcomes



$N = 20$ (top), $N = 50$ (bottom): accuracy, engagement, polarization, action gap.