July, 2023 ERI-CES DP 11/10

Multinationals, R&D and productivity: evidence for UK manufacturing firms

> Dolores Añón Higón
Universitat de València and ERI-CES, Spain

> Miguel Manjón Antolín
Universitat Rovira I Virgili, Spain

> Juan A. Mañez
Universitat de València and ERI-CES, Spain

eri-ces www.erices.es

iscussion Papers
in Economic Behaviour

VNIVERSITAT DE VALÈNCIA

Buying Followers or Collaborators: Influencer Campaigns*

Penélope Hernández† and Neus Planells‡

November 2026

Abstract

Firms can inflate a client's audience by buying followers, or borrow one by paying other creators to collaborate with her. We study how such campaigns move a social network from one equilibrium to another when agents coordinate with the accounts they follow and weigh each account by the size of its audience. A campaign converts a dissenter exactly when the exposure it delivers exceeds her resistance, and conversions accumulate into a new equilibrium. A collaboration is worth at most two bought followers, and exactly two when it closes a triangle around one of the client's followers, so the premium for collaborating is the price of a triangle: at intermediate prices, the clustering of the client's audience, not prices, decides which instrument to use. In nested split graphs the required clustering is automatic. In core–periphery networks cascades are all or nothing and may run from users to influencers, and between communities money buys only the first step.

Keywords: Network games; Coordination; Network interventions; Triadic closure

JEL classification: C72, D85, M31, M37

1 Introduction

Influencer marketing sells attention. Firms pay creators to display their products, and what they buy is the attention of the creators' followers (Fainmesser and Galeotti, 2021). The size of an audience is public, displayed next to every post, and it shapes how an account is perceived: accounts with more followers are seen as more popular and more likable (De Veirman et al., 2017), and the engagement that a sponsored campaign generates in response to spending rises with the creator's audience (Leung et al., 2022). Because an audience is valuable and visible, it is also manufactured. A firm can inflate its client's audience with followers

*Hernández acknowledges financial support from grant PID2024-162635NB-I00 funded by MICIU/AEI/10.13039/501100011033/FEDER, UE, and from grant CIAICO/2023/031, funded by the Conselleria d'Innovació, Universitats, Ciència i Societat Digital, Generalitat Valenciana.

†ERI-CES, Department of Economics, University of Valencia, Spain – Penelope.Hernandez@uv.es

‡Department of Economics, University of Valencia, Spain – Neus.Planells@uv.es

attracted by promotional giveaways or bought from specialized vendors, or it can borrow the audiences of other creators by paying them to collaborate with its client. Regulators have taken notice of the first practice: in 2024 the US Federal Trade Commission prohibited the sale and purchase of fake indicators of social media influence (Federal Trade Commission, 2024). This paper asks when manufacturing an audience changes what people choose, and when a firm should buy followers for its client rather than collaborations with other creators.

We study a platform on which agents choose between two varieties of a product, a style or an opinion. Each agent has an innate preference for one variety and also wants to coordinate with the accounts she follows, as in network games with conflicting preferences (Hernández et al., 2013). In the status quo every agent plays her preferred variety, and this profile is a Nash equilibrium. A firm is hired by a client, whom we call the pivot, to spread the client's variety. It can neither change tastes nor choose actions; it can only change the network around its client, and platform practice offers it two ways of doing so. Under a *giveaway* it buys followers for the client, accounts that follow her but bring no audience of their own. Under a *collaboration* it links the client with a creator who shares her variety and has an audience. The client is given. Unlike in seeding and key-player problems (Kempe et al., 2003; Ballester et al., 2006), the firm does not choose whom to target; it chooses how to reshape the network around the account it has been hired to promote.

What makes an audience worth buying is the way agents weigh what they see. Each agent weighs a neighbour by the size of that neighbour's audience. The weights have a simple reading. An agent learns what her contacts do through the activity of their own audiences, the likes, shares and comments that bring their posts to her feed, so a contact with a larger audience leaves more traces of her behaviour and is over-represented in what the agent sees. This is the logic of the friendship paradox (Feld, 1991) and of the "majority illusion" (Lerman et al., 2016). That influence grows with connectedness is also a theme of models of social learning (DeMarzo et al., 2003; Golub and Jackson, 2010), and laboratory evidence suggests that people take the network position of others into account, if only partially, when weighing their opinions (Battiston and Stanca, 2015). The weighting is what separates our question from threshold models in which agents respond to the share of their neighbours who adopt (Granovetter, 1978; Morris, 2000). There, buying followers for one account changes nobody else's incentives, since her followers see exactly the same neighbours as before. Here, it raises the client's weight in the evaluation of every one of her followers. A bought follower is the limiting case: she has no audience, so she carries no weight in anybody's evaluation, and her only effect is to raise the client's.

Our first result is an accounting of what a campaign does to a dissenter, an agent who prefers the other variety. A campaign never touches the support she receives from like-minded neighbours. It only adds weight to the neighbours who play the client's variety, and this addition is her *exposure*. Exposure has two components. *Breadth* is the increase in the client's audience, which reaches every dissenter who follows the client and to which bought followers and recruited creators contribute alike. *Depth* is one unit for each recruited creator the dissenter follows, and only collaborations produce it. Inside the client's audience depth is triadic closure: a collaboration with a creator whom a follower of the client already follows closes a triangle around that follower. Against exposure stands the dissenter's *resistance*, the largest increase in opposing weight she can absorb, which is a statistic of the original network and does not depend on the campaign. A dissenter converts exactly when her exposure

exceeds her resistance, agents who share the client's variety never move, and conversions accumulate until the network settles in a new Nash equilibrium. Comparing two campaigns therefore reduces to comparing the exposures they deliver.

The second result uses this accounting to compare the instruments. A collaboration is a giveaway with depth attached. Inside the client's audience it is worth at most two bought followers, and exactly two to a follower around whom it closes a triangle; outside that audience only collaborations act. The premium that a collaboration costs over a follower, measured in followers, is therefore the price of a triangle, and it divides the firm's problem into three regimes. When triangles are free, the firm collaborates as much as its budget and the available creators allow. When a triangle costs at least a follower, collaborations survive only as bridges to dissenters whom no giveaway can reach. In between, prices no longer decide. Some optimal campaign recruits only groups of creators around which part of the client's audience is clustered, so the price ratio becomes a required degree of clustering. We construct two networks with the same prices, budget, audience, resistances and reach that call for opposite instruments. The reinforcement logic of complex contagion (Centola and Macy, 2007; Centola, 2010) thus reappears as a rule for designing the instrument: giveaways are a broadcast instrument, and collaborations a reinforcement instrument that pays only when it is clustered.

The third set of results follows a campaign beyond its first round in network structures observed on platforms. In nested split graphs, the architectures that network design tends to produce under complementarities, the clustering requirement is met automatically, and a client opens a new audience only by collaborating with a creator better connected than herself. In the ideal core–periphery network, the simplest of them, the cascade is all or nothing. A single inequality decides whether influencers or users tip first, and when the dissenting influencers have a large dissenting audience the cascade runs from the bottom up, as in the simulations of Watts and Dodds (2007). For a client who is a user, every collaboration closes a triangle around every target, and the choice of instrument is settled by prices alone: she should collaborate as much as possible when a collaboration costs less than two followers, and buy only followers otherwise. When each variety has its own community, as in the homophilic clusters that dominate interaction on large platforms (Cinelli et al., 2021), money buys only the first step. An influencer client converts her own followers in the other community as soon as her bought audience exceeds their resistance, but whether the conversion spreads further is decided by the network alone, and it requires her to be followed by more than half of that community's users. The same condition yields a paradox of strong cores: a community can be opened from outside only if the status quo allows heavy cross-following, and that requires its own influencers to be strong.

These results bear on the regulation of fake indicators. The FTC rule reaches followers generated by bots or hijacked accounts, but not real users who follow an account only to enter a giveaway. In our model the two are the same object, a tie that brings no audience of its own and raises the client's weight in the eyes of her followers. Our results also indicate where audience inflation matters. It acts directly only on the client's own dissenting audience, it is weakly dominated by collaborations whenever creators are available and a collaboration costs no more than a follower, and between communities it buys at most the first step. None of this relies on the undirected representation of links: the accounting holds verbatim when links are directed and weights are given by followers, and its qualitative content survives

any weighting that rises with audience size (Section 8). We study the effectiveness of these practices, not their welfare consequences, which would require a model of what agents lose when they abandon their preferred variety.

Related literature

The paper contributes first to the study of coordination in networks when agents have conflicting preferences. Binary coordination games on networks and the spread of an action through them have been studied extensively (Morris, 2000; López-Pintado, 2006; Young, 2009), as have network games with strategic complements or substitutes (Galeotti et al., 2010; Bramoullé et al., 2014), and Jackson et al. (2017) survey the economic consequences of network structure. Closer to our setting are models in which an intrinsic preference over actions may conflict with the coordination motive. Neary (2012) studies two groups that prefer different conventions, Hernández et al. (2013, 2017) characterize the equilibria of network games in which players have an idiosyncratic identity, and Orlova (2022) extends the characterization to a broad class of games with idiosyncratic action preferences. These papers characterize the equilibria of a given network, or the spread of an action from a given seed. We build on the framework of Hernández et al. (2013) with two departures. Neighbours are weighted by their audience, which is what makes a change in one agent's links matter to third parties, and the network is the object of an outside intervention, so that the question becomes which equilibrium the intervened network reaches from a fixed status quo.

The paper also belongs to the growing literature on interventions in networks. A planner can target agents, as in the key-player problem of Ballester et al. (2006) and in the targeting of private returns by Galeotti et al. (2020), extended to large networks by Parise and Ozdaglar (2023), or she can act on the links themselves. Efficient networks under local complementarities are nested split graphs (Belhaj et al., 2016), optimal weighted and directed networks generalize them (Li, 2023), and a planner who builds a network one link at a time optimally forms a nested split graph at every step (Sun et al., 2026). Structural interventions are equivalent to interventions on characteristics (Sun et al., 2023), joint interventions on characteristics and link weights have distinct welfare effects (Kor et al., 2025), and when agents choose their own links a planner may prefer to subsidize links rather than actions (Dasaratha and Shah, 2026). Most of this work studies continuous actions, where the effect of an intervention is summarized by centralities or eigenvalues. Nora and Winter (2024) study binary actions with complementarities, as we do, and show that complete core–periphery networks maximize the revenue of a principal who induces all agents to act. Our setting differs from this literature in three respects. Actions are binary and the outcome tips, so the question is which equilibrium the network ends in. The intervening agent is a firm acting for a given client rather than a planner who chooses her targets. And the two instruments differ in reciprocity, one raising the connectivity of a single agent and the other of two, so that the statistic that governs the choice between them is local, the clustering of the client's audience, and it is met automatically in the nested split graphs that network design tends to produce.

Finally, the paper relates to the economics and marketing of influence. The problem of whom to seed in order to maximize diffusion goes back to Kempe et al. (2003). Galeotti and Goyal (2009) show how optimal influence strategies depend on the content of social in-

teraction, Goyal et al. (2019) study firms that compete by seeding a network, and Banerjee et al. (2013) find that diffusion is higher when the first informed agents are more central. In marketing, Hinz et al. (2011) find that seeding well-connected people is the most successful strategy, while Watts and Dodds (2007) argue that large cascades are driven by a critical mass of easily influenced individuals rather than by influentials. Fainmesser and Galeotti (2021) model the market in which marketers pay influencers for endorsements, and Anand et al. (2018) the contract between an advertiser and an influencer who can inflate her follower count. A related empirical literature documents the manipulation of other public signals of quality, such as online reviews (Mayzlin et al., 2014; Luca and Zervas, 2016; He et al., 2022). These papers ask why signals are manipulated and how manipulation can be priced, detected or deterred. We ask instead what an inflated audience does to equilibrium behaviour, and we find that the answer depends on how the client's audience is clustered. This connects the question to the evidence that clustered ties spread behaviour that needs reinforcement (Centola and Macy, 2007; Centola, 2010; Katona et al., 2011), and to the finding that homophily slows down best-response dynamics (Golub and Jackson, 2012).

The rest of the paper is organized as follows. Section 2 presents the model and the best response, and Section 3 introduces the two instruments. Section 4 defines exposure and resistance and characterizes how a campaign moves the network to a new equilibrium, and Section 5 compares the instruments. Section 6 studies nested split graphs and the core–periphery network, and Section 7 the two-community network. Section 8 concludes and discusses the modelling choices. Proofs are in Appendix A, and Appendix B constructs the networks of Example 1.

2 Model

We consider a static game among n players, each represented by a node in a network. The network is an undirected graph $G = (N, L)$ in which each player i has a set of neighbours N_i, with $|N_i| \geq 2$ for all i.[1] Each player chooses an action $x_i \in \{0, 1\}$ and has a preference for one of the two actions. We call this preference agent i's *type* $\theta_i \in \{0, 1\}$ and we write $\bar{\theta} = (\theta_1, \ldots, \theta_n)$ for the profile of types. Player i's utility is

$$u_i(x_1, \ldots, x_n) = \left(\delta \sum_{j \in N_i} \omega_j^i \mathbf{1}_{\{x_i = x_j\}} + (1 - \delta) \sum_{j \in N_i} \omega_j^i \mathbf{1}_{\{x_i \neq x_j\}} \right) \left(1 + \lambda \mathbf{1}_{\{x_i = \theta_i\}}\right), \tag{1}$$

where $\delta \in (0, 1)$ represents the relative preference for coordination, $\lambda > 0$ measures the intensity with which the player values acting in accordance with her type, and ω_j^i is the weight that player i assigns to neighbour j. The weights reflect the importance that i assigns to each neighbour j once the reciprocal link is removed, so that they capture the net connectivity of j from i's perspective:

$$\omega_j^i = \begin{cases} \dfrac{|N_j| - 1}{\sum_{k \in N_i} (|N_k| - 1)} & \text{if } j \in N_i, \\ 0 & \text{if } j \notin N_i \text{ or } j = i. \end{cases} \tag{2}$$

[1] The condition $|N_i| \geq 2$ ensures that every agent's weight vector is well defined and non-degenerate.

The weights are normalized, $\sum_{j\in N_i}\omega_j^i = 1$. Writing $f_j = |N_j| - 1$ for the net connectivity of agent j and $F_i = \sum_{k\in N_i} f_k$ for its total over i's neighbourhood, we have

$$\omega_j^i \;=\; \frac{f_j}{F_i} \qquad \text{for } j \in N_i, \tag{3}$$

and $F_i > 0$ because every neighbour k of a player is herself a player, so that $|N_k| \geq 2$ and $f_k \geq 1$.

The weights have a simple reading in terms of what an agent observes. The total $F_i = \sum_{k\in N_i}(|N_k| - 1)$ is the number of paths of length two that start at i and do not return to her, $i \to k \to l$ with $l \neq i$, and f_j is the number of those paths that go through j. Hence ω_j^i is the share of i's two-step paths that pass through j. On a platform, a path $i \to j \to l$ is a trace of j's behaviour that reaches i through j's own audience: a like, a share or a comment by l on a post by j that appears in i's feed. If i draws one such trace at random, it concerns j with probability ω_j^i. Since the factor $1 + \lambda\mathbf{1}_{\{x_i=\theta_i\}}$ in (1) does not depend on j, the utility of i is her expected payoff from a coordination game with the author J of a randomly drawn trace,

$$u_i(x_1,\ldots,x_n) = \mathbb{E}\big[\delta\,\mathbf{1}_{\{x_i=x_J\}} + (1-\delta)\,\mathbf{1}_{\{x_i\neq x_J\}}\big]\big(1+\lambda\mathbf{1}_{\{x_i=\theta_i\}}\big), \qquad \Pr(J=j) = \omega_j^i. \tag{4}$$

This is the mechanism behind the friendship paradox of Feld (1991) and the "majority illusion" of Lerman et al. (2016): a contact reached through a link is over-represented in proportion to her audience. It also delivers the two features on which the paper rests. An account with no audience of her own generates no trace and therefore carries no weight, and every new follower of an account adds one trace of that account to the feed of each of her other followers. We return to the choice of weights, and to the direction of links, in Section 8.

Links in G are symmetric, but influence is not. By (3), the weight of j in i's evaluation is $\omega_j^i = f_j/F_i$, while the weight of i in j's evaluation is $\omega_i^j = f_i/F_j$. The weight matrix $W = (\omega_j^i)$ is therefore not symmetric, and it defines a weighted *directed* graph of influence on top of the undirected graph of links. Along each link,

$$\omega_j^i > \omega_i^j \quad\Longleftrightarrow\quad f_jF_j > f_iF_i,$$

so influence flows predominantly from the endpoint that is better connected and better surrounded to the other. This asymmetry is what licenses the language of followers used throughout the paper. We say that i *follows* j, and that i belongs to j's *audience*, when $j \in N_i$. The relation is formally symmetric, but it carries the meaning of following exactly when ω_j^i is large relative to ω_i^j, as it is between a user and an influencer. The limiting case is a degree-one agent. Her net connectivity is zero, so she attends to her single neighbour but carries zero weight in that neighbour's evaluation. Such a tie is two-sided in the graph and one-sided in influence. Players have degree at least two, so this case arises only for the nodes created by a promotional giveaway in Section 3, and it is precisely the tie that a giveaway creates.

We add an external agent, an influencer agency, that does not play the game but can alter it through the only channel that (3) leaves open. Because the weight j carries in i's evaluation is governed by f_j, and f_j is a property of the network rather than of the play, an agent's influence over her neighbours can be raised without touching preferences, payoffs or

actions: it suffices to raise her degree. The instruments of Section 3 are two ways of doing so, and they differ in whether the connectivity created accrues to one agent or to two.

Throughout, we restrict attention to the coordination region

$$\mathcal{R}_C \;=\; \left\{(\delta,\lambda)\in(0,1)\times(0,\infty):\; \delta \;>\; \frac{1+\lambda}{2+\lambda}\right\}. \tag{5}$$

Since $\frac{1+\lambda}{2+\lambda} > \frac{1}{2}$ for every $\lambda > 0$, membership in $\mathcal{R}_C$ implies $\delta > \frac{1}{2}$, so that actions are strategic complements throughout. Complementarity alone, however, is not enough. When $\frac{1}{2} < \delta \leq \frac{1+\lambda}{2+\lambda}$ actions are complements but playing one's type is a dominant strategy, so no change in the network can move anyone. The region $\mathcal{R}_C$ is exactly the set of parameters in which the coordination motive can override an agent's type, and Lemma 1 below shows that within it the threshold $r_{\delta,\lambda}$ is strictly positive. The restriction is therefore substantive rather than expositional. It is what gives a campaign both purchase and a predictable direction: adding connectivity to an agent who consumes a given variety can only draw her neighbours towards that variety, and it can draw them far enough to switch.

A pure-strategy profile $x^* = (x_1^*, \ldots, x_n^*)$ is a Nash equilibrium if $u_i(x_i^*, x_{-i}^*) \geq u_i(x_i, x_{-i}^*)$ for every player i and every $x_i \in \{0,1\}$, where x_{-i} denotes the action profile of all agents except i.

A campaign acts on the numerators of the weights (3) and not on the weights themselves, since raising the degree of one neighbour raises her numerator and every agent's denominator at once. We therefore state the best response in levels rather than in shares. For an action profile $\mathbf{x}$ and an agent i we write

$$F_i^+(\mathbf{x}) \;=\; \sum_{j\in N_i} f_j\,\mathbf{1}_{\{x_j=\theta_i\}}, \qquad F_i^-(\mathbf{x}) \;=\; \sum_{j\in N_i} f_j\,\mathbf{1}_{\{x_j\neq\theta_i\}}, \tag{6}$$

for the aggregate net connectivity of the neighbours who act as i prefers and of those who do not, so that $F_i^+ + F_i^- = F_i$. We call F_i^+ the *affinity mass* and F_i^- the *dissent mass* of agent i. Both are sums of degrees, and are therefore unaffected by the normalization. Define

$$r_{\delta,\lambda} \;=\; \frac{\delta(2+\lambda)-(1+\lambda)}{\delta(2+\lambda)-1}. \tag{7}$$

Lemma 1. *Let $(\delta,\lambda)\in\mathcal{R}_C$. Then $r_{\delta,\lambda}\in(0,1)$, and, with ties broken in favour of her type, the best response of agent i is*

$$BR_i(\mathbf{x}_{-i}) \;=\; \begin{cases} \theta_i, & \text{if } F_i^+(\mathbf{x}) \;\geq\; r_{\delta,\lambda}\,F_i^-(\mathbf{x}),\\ 1-\theta_i, & \text{if } F_i^+(\mathbf{x}) \;<\; r_{\delta,\lambda}\,F_i^-(\mathbf{x}). \end{cases} \tag{8}$$

Lemma 1 says that an agent consumes the variety she prefers precisely when the connectivity of the neighbours who consume it is at least a fixed multiple of the connectivity of those who do not. Three features of the condition matter for what follows. It compares masses rather than counts, so that a single well-connected dissenter can outweigh several peripheral allies. The multiple $r_{\delta,\lambda}$ is common to all agents and does not depend on degrees, so that what distinguishes one agent's position from another's is the composition of her neighbourhood and not its size. And the condition is homogeneous of degree zero in (F_i^+, F_i^-), so

only the ratio matters.[2] A campaign therefore has purchase on an agent in exactly one way, by shifting connectivity from one mass to the other, and the rest of the paper accounts for the two ways of doing so.

A Nash equilibrium is *homogeneous* if all agents choose the same action irrespective of their type, and *fully satisfactory* if $x^* = \bar{\theta}$. In $\mathcal{R}_C$ both homogeneous profiles are equilibria of every network and every distribution of types: an agent all of whose neighbours choose the same action has $F_i^- = 0$ if that action is her preferred one and $F_i^+ = 0$ otherwise, and in the second case the condition of Lemma 1 fails, so she follows the crowd. By the same lemma the fully satisfactory profile is an equilibrium if and only if

$$F_i^+ \;\geq\; r_{\delta,\lambda}\, F_i^- \qquad \text{for every } i \in N, \tag{9}$$

both masses being evaluated at $\bar{\theta}$. A network on which (9) holds and in which both types are present therefore carries at least three equilibria, and the campaigns we study operate in the gap between them.

Every campaign in this paper starts from the same *status quo*, which we maintain throughout without further mention: $(\delta, \lambda) \in \mathcal{R}_C$; the profile $\bar{\theta}$ is a fully satisfactory equilibrium of $\Gamma = (G, \bar{\theta})$, that is, (9) holds; and $\bar{\theta}$ is not homogeneous, so that both types are present. From now on F_i^+ and F_i^- without an argument denote these masses evaluated at $\bar{\theta}$ in the original network Γ.

3 The two mechanisms

A firm sells one of two product varieties on a platform whose users are simultaneously consumers and content creators. A user's type is her taste over the two varieties, and her action is the one she consumes and displays. The firm has a single client, whom we call the *pivot* and denote by v: a creator who genuinely prefers the firm's variety, so that θ_v is the action the firm wishes to spread. The pivot is the firm's customer and is given: the firm does not choose whom to promote, only how to change the network around the account it has been hired to promote. We normalize $\theta_v = 1$, the opposite case being symmetric, and we write

$$\Theta^+ = \{i \in N : \theta_i = \theta_v\}, \qquad \Theta^- = N \setminus \Theta^+, \qquad J^- = N_v \cap \Theta^-,$$

so that J^- collects the dissenting followers of the pivot, her *dissenting audience*.

The firm cannot change tastes and does not choose actions. It can only change the network, and platform practice offers it two ways of doing so. The first is a promotional giveaway: entry requires following the pivot's account, which produces a large number of ties from users who bring no audience of their own. The second is a collaboration between creators, a joint post or a cross-mention, which produces a tie between the pivot and a creator who does have an audience. Both raise the pivot's connectivity. They differ in whether anyone else's connectivity rises with it.

[2] Equivalently, i plays her type exactly when the share F_i^+/F_i of her attention that goes to like-minded neighbours is at least $\tau_{\delta,\lambda} = r_{\delta,\lambda}/(1 + r_{\delta,\lambda}) < \frac{1}{2}$. The share form is natural on a fixed network. The level form is natural here, because a campaign changes F_i^+ and F_i^- by known amounts, whereas its effect on the share runs through a denominator that the campaign also moves.

Formally, a *campaign* is a pair (α, S), where $\alpha \in \mathbb{Z}_{\geq 0}$ is a number of bought followers and S is a set of creators not yet linked to the pivot. We write $\sigma = |S|$ and $m = \alpha + \sigma$ for the *scale* of the campaign. The campaign adds α new nodes $Z_\alpha = \{z_1, \ldots, z_\alpha\}$ and links each of them, and each member of S, to the pivot. The resulting *intervened network* $\Gamma_{(\alpha,S)} = (\widetilde{N}, \widetilde{L}, \bar{\theta})$ has

$$\widetilde{N} = N \cup Z_\alpha, \qquad \widetilde{L} = L \;\cup\; \big\{\{v, z\} : z \in Z_\alpha\big\} \;\cup\; \big\{\{v, s\} : s \in S\big\},$$

and net connectivities $\widetilde{f}_v = f_v + m$, $\widetilde{f}_s = f_s + 1$ for $s \in S$, and $\widetilde{f}_i = f_i$ for every other $i \in N$. We refer to α as Mechanism I and to S as Mechanism II, and $\Gamma_{(0,\emptyset)} = \Gamma$.

Two features of the campaign are maintained throughout. The first concerns the bought followers: the nodes of Z_α are not players, each has degree one, and none revises. That a bought follower carries no weight is not an additional assumption. By (2), a node of degree one has net connectivity $f_z = 0$, so $\omega_z^v = 0$, and she belongs to no other agent's neighbourhood. She raises the pivot's audience f_v and nothing else, so her action is immaterial, and we set it to θ_v for definiteness. This is the limiting case of the asymmetry discussed in Section 2: the tie runs both ways in the graph, but influence runs only from the pivot to her. Degree one reflects that a user who follows an account to enter a draw brings no audience of her own. Non-participation reflects that she was recruited by a promotion rather than drawn into the strategic problem, and it is also what allows these nodes to have degree one without violating $|N_i| \geq 2$. A collaboration, by contrast, is bilateral both in the story and in the weights: both parties gain a link, and both gain weight in the evaluation of their own neighbours.

The second feature concerns the collaborations: the firm collaborates only with creators who share the pivot's taste, so that S is a subset of the set of *available creators*

$$\mathcal{A} \;=\; \Theta^+ \setminus \big(N_v \cup \{v\}\big).$$

This describes commercial practice rather than accommodating the argument. A firm recruits ambassadors who already prefer its variety, both because an endorsement from a creator who consumes something else is worthless and because the collaboration is itself a display of the product. It also has three analytical consequences: the pivot's own stability ceases to be a constraint, since every agent whose weight rises is already acting as the pivot does; the adjustment process moves in one direction only; and no agent is ever pushed away from θ_v, so the campaign has a sign.

4 Exposure and resistance

A campaign reaches an agent through exactly one channel. By Lemma 1 her behaviour turns on the ratio of her affinity mass to her dissent mass, so what a campaign can do to her is to add connectivity to one side of that ratio. This section answers the two questions that follow. What does a campaign deliver to an agent, and how much does she need? The first is *exposure* and the second is *resistance*, and the conversion condition is the comparison of the two.

Let $i \in \Theta^-$. The agents whose connectivity the campaign raises, the pivot and the recruited creators, all have type $\theta_v \neq \theta_i$, and therefore contribute to F_i^- whenever they are neighbours of i. Since i is neither the pivot nor a recruited creator, her neighbourhood is

unchanged. The pivot's net connectivity rises by m and that of each recruited creator by one, so the increase in her dissent mass is

$$E_i(\alpha, S) \;=\; \underbrace{m \cdot \mathbf{1}_{\{v \in N_i\}}}_{\text{breadth}} \;+\; \underbrace{\big|S \cap N_i\big|}_{\text{depth}}, \tag{10}$$

which we call the *exposure* of agent i. Two features of (10) organize everything that follows. First, α and σ enter the breadth term identically: a bought follower and a recruited creator raise the pivot's connectivity by the same amount, and are therefore perfect substitutes in the production of breadth. Second, only Mechanism II produces the depth term. The two instruments are not two goods but two technologies for producing a common good, one of which jointly produces a second good.

Exposure is zero exactly for the dissenters who follow neither the pivot nor any recruited creator; the bought followers never enter, since they are linked only to v. A campaign therefore reaches directly only two groups of dissenters: the pivot's dissenting audience J^-, each member of which receives at least m, and the dissenters who follow some recruited creator, each of whom receives one unit per recruited creator she follows. Under Mechanism I alone the second group is empty. Every other dissenter faces, after the campaign, the same masses as in Γ, and can change her behaviour only after some of her neighbours have changed theirs.

Inside the pivot's neighbourhood, depth has a structural interpretation. For an agent i let T_i denote the number of links among her neighbours, that is, the number of triangles of which i is a vertex, and let ΔT_i be its change under the campaign. Every link the campaign creates has v as an endpoint. For $j \in J^-$, whose degree the campaign leaves unchanged, a new link $\{v, s\}$ joins two of her neighbours exactly when $s \in N_j$, and then closes the single triangle $\{v, j, s\}$. A link to a bought follower closes none, since that follower has degree one. Hence

$$\Delta T_j \;=\; |S \cap N_j| \qquad \text{and} \qquad E_j(\alpha, S) \;=\; m \;+\; \Delta T_j \qquad \text{for every } j \in J^-, \tag{11}$$

while around a dissenter outside N_v no triangle is ever closed. Depth inside the pivot's audience is triadic closure.[3] The two instruments therefore buy different things. A giveaway buys the client degree and nothing else: a follower with no audience of her own cannot link two people who already follow the client, so no neighbourhood becomes denser. A collaboration buys degree and, whenever the recruited creator is followed by one of the client's own followers, closes a triangle around that follower, who now sees two sources she attends to visibly joined.

Exposure measures what a campaign delivers to a dissenter. Whether it is enough to change her behaviour depends on how much she can absorb. Let $i \in \Theta^-$. At the status quo she keeps θ_i because, by Lemma 1, $F_i^+ \geq r_{\delta,\lambda} F_i^-$. The campaign leaves her affinity mass unchanged and adds $E_i(\alpha, S)$ to her dissent mass, so the question is how large that addition can be before the inequality fails.

[3] Since $|N_j|$ is unchanged, (11) can be read in terms of the local clustering coefficient $c_j = T_j / \binom{|N_j|}{2}$, which rises by $|S \cap N_j| / \binom{|N_j|}{2}$. We work with triangle counts for the reason given before Lemma 1: a campaign moves the denominators of normalized measures as well as their numerators.

Definition 1 (Resistance). *For $i \in \Theta^-$, the* resistance *of i is*

$$\rho_i \;=\; \frac{F_i^+}{r_{\delta,\lambda}} \;-\; F_i^-. \tag{12}$$

The resistance is the largest increase in dissent mass that agent i can absorb while still preferring θ_i; it is nonnegative because $\bar{\theta}$ is an equilibrium. It is measured in the same units as exposure, net connectivity, so the two can be compared directly. It is a local statistic of the original network, known before any campaign is chosen: exposure belongs to the campaign, resistance to the agent. And it summarizes the role of preferences through the single number $r_{\delta,\lambda}$, which decreases in λ and increases in δ: a stronger attachment to one's own variety raises resistance, and a stronger motive to coordinate lowers it.

To follow a campaign over time, let $\mathbf{x}^0$ be the status quo $\bar{\theta}$, with the bought followers playing θ_v, and let $\mathbf{x}^{t+1}$ be the profile of best responses to $\mathbf{x}^t$ in $\Gamma_{(\alpha,S)}$. The *first-round converts* are the agents for whom their type is no longer a best response once the campaign is in place,

$$\mathcal{C}_1(\alpha, S) = \big\{i \in N :\; x_i^1 \neq \theta_i\big\},$$

and, when the sequence settles at a profile $\mathbf{x}^{(\alpha,S)}$, the *final converts* are $\mathcal{C}^*(\alpha, S) = \{i \in N : x_i^{(\alpha,S)} \neq \theta_i\}$.

Proposition 1. *For every campaign (α, S):*

(i) $\mathcal{C}_1(\alpha, S) = \big\{i \in \Theta^- :\; E_i(\alpha, S) > \rho_i\big\}$.

(ii) *Agents of type θ_v never abandon θ_v, and a dissenter who converts never returns. The sequence $(\mathbf{x}^t)$ is constant after at most $|\Theta^-|$ steps, its limit is a Nash equilibrium of $\Gamma_{(\alpha,S)}$, and $\mathcal{C}_1(\alpha, S) \subseteq \mathcal{C}^*(\alpha, S)$.*

(iii) *If $E_i(\alpha, S) \leq E_i(\alpha', S')$ for every $i \in \Theta^-$, then $\mathcal{C}_1(\alpha, S) \subseteq \mathcal{C}_1(\alpha', S')$ and $\mathcal{C}^*(\alpha, S) \subseteq \mathcal{C}^*(\alpha', S')$.*

Proposition 1 answers the question with which the paper began. A campaign moves the network from the equilibrium $\bar{\theta}$ of Γ to an equilibrium of the intervened network, and the move is governed, agent by agent, by a single comparison: a dissenter is won over at once exactly when what the campaign delivers to her exceeds what she can absorb. Under giveaways alone, for instance, exposure is α on J^- and zero elsewhere, so giveaways convert the pivot's dissenting followers in increasing order of resistance and nobody else in the first round. Because every switch is a switch to θ_v, the adjustment only accumulates conversions: it cannot unravel, and the first-round converts are a lower bound on the final change. Part (iii) is the tool for everything that follows. Resistances do not depend on the campaign, so comparing two campaigns requires only comparing the exposures they deliver: a campaign that gives every dissenter at least as much exposure wins over at least the same agents, both at once and in the new equilibrium.

5 Comparing the mechanisms

The firm has a budget and wants to change as much behaviour as possible. By Proposition 1(iii), the two mechanisms can be compared through what each adds to exposure, and three comparisons follow directly from (10). First, a collaboration is a giveaway with depth attached. When $\alpha \geq 1$, replacing one bought follower by a collaboration with a new creator $s \in \mathcal{A} \setminus S$ leaves the scale unchanged and adds one unit of depth to every dissenter who follows s:

$$E_i(\alpha - 1, S \cup \{s\}) \;=\; E_i(\alpha, S) + \mathbf{1}_{\{i \in N_s\}} \qquad \text{for every } i \in \Theta^-. \tag{13}$$

Second, inside the pivot's audience a collaboration is worth at most two followers. For $j \in J^-$,

$$E_j(\alpha, S) \;=\; \alpha + \sigma + |S \cap N_j| \;\leq\; \alpha + 2\sigma \;=\; E_j(\alpha + 2\sigma, \emptyset), \tag{14}$$

so a collaboration is worth exactly two followers to a member of the audience around whom it closes a triangle, and one follower to any other. Third, outside the pivot's audience only collaborations act, since there exposure is $|S \cap N_i|$.

Let a follower cost $c_I > 0$ and a collaboration $c_{II} > 0$, and let the firm choose a campaign with $S \subseteq \mathcal{A}$ and $c_I\alpha + c_{II}\sigma \leq B$ to maximize the number of converts, either $|\mathcal{C}_1(\alpha, S)|$ or $|\mathcal{C}^*(\alpha, S)|$. Writing

$$c_I\alpha + c_{II}\sigma \;=\; c_I\, m + (c_{II} - c_I)\, \sigma,$$

every link buys one unit of breadth at the price c_I, and every collaboration pays in addition a premium for its depth. We measure this premium in bought followers,

$$\kappa \;=\; \frac{c_{II} - c_I}{c_I}.$$

The firm's problem is posed for a given client, and its scope deserves a comment. The pivot is not chosen: she is the customer who hires the firm, and her position in the network is part of the data. This separates our question from the seeding and key-player problems of Kempe et al. (2003) and Ballester et al. (2006), in which a planner selects whom to target. Here the target is fixed, and the firm reshapes the network around her. Nor do we compute optimal campaigns network by network. Which creators to recruit, when the budget cannot cover all the useful ones, is a combinatorial selection problem of the coverage type, of interest in its own right for operations research. Our question is a different one. We ask which instrument an optimal campaign uses, and what kind of creator it recruits, as functions of prices and of the structure of the client's neighbourhood. The answer identifies transformations that can only improve a campaign, and hence properties that some optimal campaign must have.

Once the creators are chosen, the rest of the budget should buy followers, since by (10) exposure never falls with α. Writing $\alpha^B(\sigma) = \lfloor (B - c_{II}\sigma)/c_I \rfloor$ for every affordable number of collaborations, $c_{II}\sigma \leq B$, the resulting scale is

$$m^B(\sigma) \;=\; \sigma + \alpha^B(\sigma) \;=\; \left\lfloor \frac{B}{c_I} - \kappa\, \sigma \right\rfloor, \tag{15}$$

so each collaboration costs κ units of breadth on average, borne by every member of J^-, and in exchange adds one unit of depth to each of the creator's dissenting followers. Whether the

exchange pays depends on who receives the depth. We call an available creator s *inner* if all her dissenting followers belong to the pivot's audience, $N_s \cap \Theta^- \subseteq J^-$, and a *bridge* otherwise. Removing a bridge lowers the exposure of some dissenter whom no giveaway can reach, so bridges are bought for reach and their value turns on the resistances outside the audience. For an inner creator, by contrast, the comparison with giveaways is entirely internal to the pivot's audience. The following example shows that, for inner creators, the choice between the two instruments can be reversed by the network alone.

Example 1. Let $c_I = 2$, $c_{II} = 3$ and $B = 10$, so that $\kappa = 1/2$ and, by (15), $m^B(\sigma) = \lfloor 5 - \sigma/2 \rfloor$, which equals $5, 4, 4, 3$ for $\sigma = 0, 1, 2, 3$. The pivot's audience J^- has six members. Three of them, the *easy* ones, have resistance in $[4, 5)$, and the other three, the *hard* ones, have resistance in $[5, 6)$. Since exposures are integers, an easy member converts when her exposure is at least 5 and a hard member when it is at least 6. The firm can recruit three inner creators s_1, s_2, s_3, who together are followed by all six members of J^-. We compare two networks that differ only in how these follows are arranged.

- In the *clustered* network, every member of J^- follows all three creators.
- In the *scattered* network, each creator is followed by one easy and one hard member, and each member follows exactly one creator.

Appendix B constructs two such networks, with identical resistances, in which no conversion spreads beyond J^-, so that $\mathcal{C}^* = \mathcal{C}_1$ for every campaign. A member of J^- who follows q recruited creators has exposure $m^B(\sigma) + q$, and Table 1 collects the outcomes.

		Clustered network		Scattered network	
σ	$m^B(\sigma)$	exposure	conversions	exposure	conversions
0	5	5	3	5	3
1	4	5	3	5 or 4	1
2	4	6	6	5 or 4	2
3	3	6	6	4	0

Table 1: Exposure of the members of J^- and number of conversions in Example 1. In the scattered network the higher exposure accrues to the followers of the recruited creators, one easy and one hard member per creator.

In the clustered network the firm should collaborate: two or three creators convert the whole audience, while the pure giveaway converts only the easy members. In the scattered network it should buy followers: the pure giveaway converts the three easy members, every campaign with collaborations converts fewer, and recruiting all three creators converts nobody. The two networks share the prices, the budget, the pivot's audience, the resistances of its members, and the set of members that the creators reach. They differ only in how the follows of the audience are distributed across the creators,[4] and yet the optimal

[4]They therefore also differ in the total number of triangles that the creators can close, eighteen in the clustered network and six in the scattered one. The criterion of Proposition 2 is stated follower by follower: what matters is how many triangles fall on the same member of J^-.

instrument is opposite, so the reversal can be produced neither by prices nor by the number of people a collaboration reaches. In the clustered network depth accumulates on the same members until it takes the hard ones past their threshold. In the scattered network each collaboration gives one unit of depth to two members and takes a fraction of a unit of breadth from all six, and the depth never accumulates. The clustered network also shows that collaborations must be evaluated together: on its own the first collaboration converts nobody beyond the pure giveaway, and only the second, landing on the same members, does. A firm that evaluated collaborations one at a time, starting from the giveaway, would stop before recruiting the first creator.

To state the rule in general, write

$$\Delta T_j(U) \;=\; |U \cap N_j|$$

for the number of triangles at $j \in J^-$ closed by the collaborations in a set U of recruited creators, as in (11).

Definition 2. *A set R of inner creators is* κ-concentrated *if every nonempty $U \subseteq R$ has a member $j \in J^-$ with $\Delta T_j(U) > \kappa\,|U|$.*

In words, for every group of recruited inner creators, some follower of the pivot follows more than a share κ of the group, or equivalently gains more than κ triangles per collaboration in the group.

Proposition 2. *For either objective, $|\mathcal{C}_1|$ or $|\mathcal{C}^*|$, the firm's problem has a solution $(\alpha^B(\sigma), S)$ with the following properties.*

(i) *If $\kappa \le 0$, then $\sigma = \min\{|\mathcal{A}|, \lfloor B/c_{II} \rfloor\}$. In particular, the firm buys followers only if it recruits every available creator.*

(ii) *If $0 < \kappa < 1$, the inner creators in S form a κ-concentrated set. Hence, if no member of J^- follows more than d available inner creators, then either S contains no inner creator or it contains fewer than $d/\kappa = d\,c_I/(c_{II} - c_I)$ of them; in particular, if $d = 0$, every creator in S is a bridge.*

(iii) *If $\kappa \ge 1$, every creator in S is a bridge.*

The proof of parts (ii) and (iii) rests on a pruning argument that conveys the logic of the result; the details are in Appendix A. Start from an optimal campaign and spend its residual budget on followers, so that its scale is $m^B(\sigma)$. Take a group U of u recruited inner creators and replace it by the followers its budget buys. The scale rises to $m^B(\sigma - u)$, so every member j of the pivot's audience gains the breadth $\beta(U) = m^B(\sigma - u) - m^B(\sigma)$ and loses the $\Delta T_j(U)$ triangles that U closed around her. Dissenters outside the audience lose nothing, because inner creators have no followers there. If no member of J^- loses more triangles than the breadth she gains, the replacement gives every dissenter weakly more exposure, and by Proposition 1(iii) the new campaign is still optimal. Each replacement lowers σ, so repeating it while possible must stop, and at that point every group U of recruited inner creators has a member j of J^- with $\Delta T_j(U) > \beta(U)$. By (15), $\beta(U) \ge \lfloor \kappa u \rfloor$, and since $\Delta T_j(U)$ is an integer this gives $\Delta T_j(U) > \kappa u$: the surviving set is κ-concentrated. When $\kappa \ge 1$ no single

creator passes the test, since she closes at most one triangle around each follower, and only bridges survive. In Example 1 the argument accounts for the whole table. In the clustered network it drops only the third collaboration, which releases one unit of breadth against one triangle per member. In the scattered network it also starts by dropping one collaboration, on the same terms. With two creators left, however, no single collaboration can be dropped, since $m^B(1) = m^B(2)$ and the breadth released would be zero; the pair can, since it releases one unit of breadth and costs each member at most one triangle. This is why Definition 2 quantifies over every group of creators and not only over single ones.

Proposition 2 is a statement about the price of a triangle. Each inner collaboration costs κ units of breadth, borne by the pivot's whole audience, and buys one triangle for each of its followers in that audience. When triangles are free the firm takes every one it can, and buys followers only with what is left. When a triangle costs at least a follower it takes none, and collaborations survive only as bridges to dissenters that giveaways cannot reach; if there are no such dissenters, the pure giveaway is optimal. In between, prices no longer decide. A group of collaborations is worth its price only if some follower of the pivot gains more than κ triangles per collaboration in the group, so the price ratio becomes a required degree of clustering of the client's audience around the creators. The requirement moves continuously across the regimes: as the premium vanishes the bound d/κ grows without limit, and as it approaches the price of a follower some follower must follow every creator in the group. This is the logic of complex contagion (Centola and Macy, 2007; Centola, 2010), in which behaviour that needs reinforcement from several sources spreads through clustered ties, applied here to the design of the instrument rather than to the diffusion process: giveaways are a broadcast instrument, and collaborations a reinforcement instrument that pays only when it is clustered. The managerial rule is to concentrate depth and diversify only for reach. The network decides whether collaborations can pay; the resistances decide whether they do, and which concentrated set of creators to recruit is the selection problem we leave aside.

In one important class of networks, the nested split graphs that network design tends to produce, the concentration requirement is met automatically. We show this in Section 6.

6 Influencers: nested split graphs and the core–periphery network

The previous sections hold for any network. This section and the next fix the structure of the network, for two purposes. The first is to see which properties of the instruments are inherited from the architecture of the network rather than from prices. The second is to follow a campaign beyond its first round, which requires computing every mass in closed form. We proceed from the general to the particular. We first consider nested split graphs, the class of networks that design tends to produce when actions are complements, and show that in them the concentration required by Proposition 2 holds automatically. We then turn to the simplest member of the class, the ideal core–periphery network, in which the cascade can be followed to the end and the choice between the instruments is settled by prices alone.

6.1 Nested split graphs

A *nested split graph* is a network whose neighbourhoods are nested: for any two agents i and j, either $N_i \setminus \{j\} \subseteq N_j$ or $N_j \setminus \{i\} \subseteq N_i$. Nested split graphs are the efficient architectures of network games with local complementarities (Belhaj et al., 2016), their generalization to weighted and directed links plays the same role in Li (2023), and a planner who adds one link per period optimally forms a nested split graph at every period (Sun et al., 2026). In this sense they are the networks that design tends to produce when actions are complements.

Lemma 2. *Let Γ be a nested split graph.*

(i) *For every campaign (α, S) and any two dissenters i, i' with $|N_i| \leq |N_{i'}|$, $E_i(\alpha, S) \leq E_{i'}(\alpha, S)$.*

(ii) *An available creator with no more neighbours than the pivot is inner. Every bridge has more neighbours than the pivot and is followed by every member of J^-.*

(iii) *Every available creator followed by some member of J^- is followed by a member of J^- with the largest number of neighbours. Hence every set of inner creators, each followed by some member of J^-, is κ-concentrated for every $\kappa < 1$.*

In a nested split graph depth is ranked in the same way as connectivity. Whatever creators a campaign recruits, it exposes better-connected dissenters at least as much as less connected ones, and every inner collaboration that reaches the client's audience closes a triangle around its best-connected member. In the terms of Section 5, the network always allows collaborations to pay, and whether they do is left entirely to the resistances: to how far down the client's audience, ranked by connectivity, depth has to reach. Part (ii) adds that a bridge is always better connected than the client and reaches her whole audience as well as new dissenters, so that a client opens a new audience only by collaborating upwards.

6.2 The core–periphery network

To follow a campaign beyond the first round we need a nested split graph that is simple enough to compute every mass in closed form, and close enough to the platforms where influencers operate to make the results meaningful. It is the ideal core–periphery network of Borgatti and Everett (2000): a small group of agents who are all linked to each other and to everybody else, and a large group of agents who are linked only to the first. Social media discussion networks display this structure clearly, and in them the agents at the core concentrate public attention (Yang and Saffer, 2021). In our model the structure has a further consequence. Since an agent's weight in her neighbours' evaluations is her net connectivity, the core dominates what the periphery sees, which is the majority illusion of Lerman et al. (2016).

Formally, Γ is an *ideal core–periphery network* if $N = C \cup P$ with $C \cap P = \emptyset$, $|C| \geq 2$, and ij is a link if and only if $i \in C$ or $j \in C$. We call C the *core*, its members *influencers*, and P the *periphery*, its members *users*. We write $k^{\pm} = |C \cap \Theta^{\pm}|$ and $n^{\pm} = |P \cap \Theta^{\pm}|$, with $k = k^+ + k^-$ and $n = n^+ + n^-$, and we assume $k^- \geq 1$ and $n^- \geq 1$, so that there are dissenters in both groups. It is the simplest nested split graph: for a user p and an influencer

c, $N_p \setminus \{c\} \subseteq N_c$, and two influencers have the same neighbours apart from each other. An influencer is linked to the other $k-1$ influencers and to the n users, and a user is linked to the k influencers. Hence all influencers share the same net connectivity, and so do all users:

$$f_C \;=\; k+n-2, \qquad f_P \;=\; k-1.$$

In a user's evaluation each influencer has weight $1/k$, while in an influencer's evaluation each user has weight $1/(k+2n-2)$: influence runs from the core to the periphery. A dissenting user sees only influencers, so her affinity and dissent masses are $k^- f_C$ and $k^+ f_C$. A dissenting influencer sees the other influencers and all users, so her masses are $(k^- - 1)f_C + n^- f_P$ and $k^+ f_C + n^+ f_P$. All dissenting users therefore share the resistance ρ_P, and all dissenting influencers the resistance ρ_C:

$$\rho_P = f_C\Big(\frac{k^-}{r_{\delta,\lambda}} - k^+\Big), \qquad \rho_C = \frac{(k^- - 1)f_C + n^- f_P}{r_{\delta,\lambda}} - \big(k^+ f_C + n^+ f_P\big). \tag{16}$$

The status quo requires both to be nonnegative, together with the analogous inequalities for the agents of Θ^+. A user's resistance depends only on the composition of the core, so users are entirely in the hands of the influencers, while a dissenting influencer is protected by her own audience. Subtracting,

$$\rho_C - \rho_P \;=\; \frac{f_P\,(n^- - r_{\delta,\lambda} n^+) - f_C}{r_{\delta,\lambda}}, \tag{17}$$

so influencers are harder to convert than users exactly when the dissenting part of the periphery, net of the pull of the aligned part, outweighs one influencer.

The position of the pivot fixes what each mechanism can do. If the pivot is an influencer, she is already linked to every other agent, so no creator is available and only Mechanism I can be used; every dissenter follows her and receives exposure α. If the pivot is a user, her neighbours are the influencers, so $J^- = C \cap \Theta^-$, and the available creators are the other aligned users. Each of them is followed by every influencer, so every collaboration closes a triangle around every dissenting influencer, while no dissenting user is linked to the pivot or to a recruited user. The exposure of a dissenting influencer is then $\alpha + 2\sigma$, and that of a dissenting user is zero. In either position a collaboration cannot bring the client a follower she does not already have: an influencer client is linked to everybody, and the creators a user client can recruit are followed only by the influencers, all of whom already follow her. The comparison of the instruments below is therefore exact even if collaborations transfer part of the creator's audience to the client (Section 8).

Proposition 3. *Let Γ be an ideal core–periphery network.*

(i) *If $v \in C$, then $\mathcal{C}^*(\alpha, \emptyset) = \Theta^-$ if $\alpha > \min\{\rho_P, \rho_C\}$, and $\mathcal{C}^*(\alpha, \emptyset) = \emptyset$ otherwise. In the first round, users convert first if $\rho_P < \rho_C$ and influencers convert first if $\rho_C < \rho_P$.*

(ii) *If $v \in P$, then $\mathcal{C}^*(\alpha, S) = \Theta^-$ if $\alpha + 2\sigma > \rho_C$, and $\mathcal{C}^*(\alpha, S) = \emptyset$ otherwise.*

The cascade in the core–periphery network is all or nothing, and its tipping point is the lower of the two resistances. When the influencers switch, their followers lose all support, since

they see nothing else; when the users switch first, the dissenting influencers lose the audience that protected them, and the loss always exceeds what kept them from switching. Which group tips first is decided by the single inequality (17). When the dissenting influencers have a large dissenting audience they are the more resistant group, and the cascade runs from the bottom up, the pattern that Watts and Dodds (2007) find in simulations, where large cascades are driven by a critical mass of easily influenced individuals rather than by influentials; when the periphery is mostly aligned it runs from the top down. The position of the client decides the role of each mechanism. An influencer client can only buy followers, but she reaches every dissenter and needs to overcome only the weaker group. A user client reaches only the influencers and must overcome ρ_C, whatever the users' resistance. For her, however, every collaboration closes a triangle around every target, so any set of creators is κ-concentrated for every $\kappa < 1$: this is the case of maximal concentration in Proposition 2, a strong form of Lemma 2(iii) in which every target, and not only the best-connected one, follows every recruited creator. A collaboration is then worth exactly two followers to every target, and the choice between the instruments is settled by prices alone.

Corollary 1. *Let Γ be an ideal core–periphery network with $v \in P$, and let $\bar{\sigma} = \min\{n^+ - 1, \lfloor B/c_{II} \rfloor\}$. For either objective, $|\mathcal{C}_1|$ or $|\mathcal{C}^*|$, the firm's problem has a solution $(\alpha^B(\bar{\sigma}), S)$ with $|S| = \bar{\sigma}$ if $c_{II} < 2c_I$, and the solution $(\lfloor B/c_I \rfloor, \emptyset)$ if $c_{II} > 2c_I$.*

Corollary 1 completes Proposition 2 for a user client. Every creator she can recruit passes the concentration test of part (ii), and the premium for depth is compared with the single triangle that each collaboration closes around every target. When a triangle costs less than a follower, $\kappa < 1$, the firm recruits as many creators as its budget and the network allow and buys followers only with what is left; when it costs more, it recruits none. The intermediate regime of Proposition 2, $0 < \kappa < 1$, in which clustering decides, here collapses into the first: nothing is left to the network, because in the ideal core–periphery network every member of the client's audience follows every creator she can recruit.

7 Two communities

In the network of Section 6 both types share the same influencers. On social media, however, users who disagree tend to gather in separate clusters, each with its own influencers, and the links between clusters are few (Cinelli et al., 2021). Homophily of this kind is known to slow down best-response dynamics (Golub and Jackson, 2012). We now ask whether a campaign launched in one cluster can change behaviour in the other. Each type forms its own core–periphery community, and the only links between them run from dissenting users to the influencers of the pivot's community: users who follow, outside their own cluster, some of the accounts that dominate the other.

Formally, Γ consists of *two communities* if $\Theta^{\pm} = C^{\pm} \cup P^{\pm}$ with $C^{\pm} \cap P^{\pm} = \emptyset$, each $\Theta^{\pm}$ is an ideal core–periphery network with core $C^{\pm}$ and periphery $P^{\pm}$, and the only other links join P^- and C^+: every member of P^- is linked to $\ell \geq 1$ members of C^+, and every member of C^+ to b members of P^-. We write $k^{\pm} = |C^{\pm}|$ and $n^{\pm} = |P^{\pm}|$, so that $k^+ b = n^- \ell$. Unlike the network of Section 6, this one is not a nested split graph: since each core has at least two members, an influencer of C^- and one of C^+ each have a neighbour the other lacks,

another influencer of her own core. Agents in the same group share the same neighbourhood structure, and hence the same net connectivity:

$$f_C^+ = k^+ + n^+ + b - 2, \qquad f_P^+ = k^+ - 1, \qquad f_C^- = k^- + n^- - 2, \qquad f_P^- = k^- + \ell - 1.$$

A dissenting user sees the k^- influencers of her community and ℓ influencers of the other, so her affinity and dissent masses are $k^- f_C^-$ and ℓf_C^+. A dissenting influencer sees only her own community, so her affinity mass is $(k^- - 1) f_C^- + n^- f_P^-$ and her dissent mass is zero. To keep them apart from the resistances of Section 6, we write $\hat{\rho}_P$ and $\hat{\rho}_C$ for the resistances of dissenting users and dissenting influencers in this network:

$$\hat{\rho}_P = \frac{k^- f_C^-}{r_{\delta,\lambda}} - \ell f_C^+, \qquad \hat{\rho}_C = \frac{(k^- - 1) f_C^- + n^- f_P^-}{r_{\delta,\lambda}}. \tag{18}$$

The status quo requires $\hat{\rho}_P \geq 0$, which limits how many influencers of the other community a dissenting user can follow, $r_{\delta,\lambda}\, \ell\, f_C^+ \leq k^- f_C^-$, together with the analogous condition for the members of C^+, who have b dissenting followers. Dissenting influencers, by contrast, are fully protected at the status quo, since nothing in their neighbourhood plays θ_v.

The pivot can now be of two kinds. A pivot in C^+ is linked to every agent of her community, so no creator is available and only Mechanism I can be used. Her dissenting followers are the b users of P^- who follow her, and they receive exposure α; every other dissenter receives none. A pivot in P^+ is linked only to C^+, and the available creators, the other members of P^+, have no dissenting followers, so every dissenter has exposure zero whatever the campaign. As in Section 6, a collaboration cannot bring the client new followers: a client in C^+ has no creator to recruit, and the creators available to a client in P^+ are followed only by members of C^+, all of whom already follow her.

Proposition 4. *Let Γ consist of two communities.*

(i) *If $v \in C^+$, then*

$$\mathcal{C}^*(\alpha, \emptyset) = \begin{cases} \emptyset & \text{if } \alpha \leq \hat{\rho}_P, \\ J^- & \text{if } \alpha > \hat{\rho}_P \text{ and } (1 + r_{\delta,\lambda})\, b\, f_P^- \leq r_{\delta,\lambda}\, \hat{\rho}_C, \\ \Theta^- & \text{if } \alpha > \hat{\rho}_P \text{ and } (1 + r_{\delta,\lambda})\, b\, f_P^- > r_{\delta,\lambda}\, \hat{\rho}_C. \end{cases}$$

(ii) *If $v \in P^+$, then $\mathcal{C}^*(\alpha, S) = \emptyset$ for every campaign.*

(iii) *If $\mathcal{C}^*(\alpha, \emptyset) = \Theta^-$ for some α, then $b > n^-/(1 + r_{\delta,\lambda})$ and*

$$r_{\delta,\lambda}\, k^+ f_C^+ \;<\; (1 + r_{\delta,\lambda})\, k^- f_C^-. \tag{19}$$

Money buys only the first step. A giveaway campaign by an influencer converts her own dissenting followers as soon as $\alpha > \hat{\rho}_P$, and any further spending has no effect; whether the conversion spreads to the rest of the other community is decided by the network alone, by the size b of the pivot's audience there against the resistance $\hat{\rho}_C$ of the influencers who hold it together, and a user of the pivot's community cannot reach the other community at all. Unlike in Section 6, partial outcomes appear: the campaign may win the pivot's own

dissenting audience while the rest of the other community, protected by its influencers, stays where it was. The condition for the whole community to fall is demanding, since by part (iii) the pivot must already be followed by more than a share $1/(1+r_{\delta,\lambda}) > 1/2$ of its users, and this yields a paradox of strong cores. By (19), a community can be opened from outside only if its own influencers are strong enough, in total connectivity, relative to those of the pivot's community. When they are weak, the status quo cannot tolerate much cross-following, the clusters are nearly closed, and a closed cluster is out of reach; when they are strong, the status quo tolerates heavy cross-following, and that same cross-following is what allows a campaign to pull the users away. Mechanism II plays no role here, because a pivot in an ideal core is already linked to every agent of her community; collaborations become available as soon as the influencers of the pivot's community are not all linked to each other, a case we leave for future work.

8 Concluding remarks

We have studied how a firm hired by a given client can move a network from one equilibrium to another by reshaping the network around her. Because agents weigh their neighbours by the size of their audiences, buying followers changes the incentives of the client's followers, and collaborating with other creators changes them through two channels, breadth and depth. Comparing campaigns then reduces to comparing the exposures they deliver against resistances that belong to the agents. Two lessons stand out. The choice between instruments is a choice between breadth and depth, settled by prices at the extremes and by the clustering of the client's audience in between. And how far a campaign travels beyond its first round depends less on the budget than on the network: on where the client sits and on the strength of the influencers who hold each community together.

These lessons rest on two modelling choices, that influence is weighted by audience and that links are undirected, and it is natural to ask how much of them survives when either choice is relaxed. The first is the essential one, although its precise form is not. All the analysis needs is that the weight of a neighbour rises with her audience. Suppose that $\omega_j^i = g(f_j)/\sum_{k\in N_i} g(f_k)$ for an increasing function g with $g(0)=0$, the model being the case $g(f)=f$. The proof of Lemma 1 goes through with the masses $F_i^{\pm}$ computed from $g(f_j)$ instead of f_j, resistance is defined as before, and a campaign adds

$$E_i^g(\alpha, S) \;=\; \big[g(f_v+m)-g(f_v)\big]\mathbf{1}_{\{v\in N_i\}} \;+\; \sum_{s\in S\cap N_i}\big[g(f_s+1)-g(f_s)\big]$$

to the dissent mass of a dissenter. Bought followers and collaborations remain perfect substitutes in breadth, only collaborations produce depth, and Proposition 1 holds with E_i^g in place of E_i. What belongs to linear weights alone is the exchange rate between the two instruments. With linear weights a unit of depth and a unit of breadth are worth the same, which is why a collaboration is worth at most two bought followers. With concave weights, a unit of depth from a creator with a small audience is worth more than a marginal unit of breadth for a client with a large one, so a collaboration with a small creator can be worth more than two bought followers, and the thresholds of Proposition 2 would shift in favour of such collaborations. At the other extreme, if weights did not depend on audiences at

all, as in threshold models in which agents respond to the share of their neighbours who adopt, a bought follower would change nobody's incentives and giveaways would have no effect. The mechanism we study therefore rests on the sensitivity of influence to audience size documented by De Veirman et al. (2017) and Leung et al. (2022), and on nothing more specific than that.

The second choice, undirected links, is largely one of convenience. Following on most platforms is directed, but the two instruments differ precisely in reciprocity: a giveaway produces a tie that is one-sided in influence, since the bought follower has no audience of her own, while a collaboration produces a tie that is two-sided, since both creators gain weight in their own audiences. With audience-based weights a single primitive, the undirected link, delivers both kinds of tie, and the difference between the instruments follows from the weights rather than from separate assumptions on the direction of each tie. Nothing in the accounting of Sections 4 and 5 depends on this choice. If N_i is the set of accounts that i follows and f_j the number of followers of j other than i, which is the same for every follower of j, a giveaway adds α accounts that follow the pivot, and a collaboration with an aligned creator s who neither follows nor is followed by the pivot makes the two follow each other. The pivot's audience grows by m and that of each recruited creator by one, so exposure is again given by (10). The proof of Proposition 1 uses only two facts, that a campaign raises the weight of agents of type θ_v alone and that it leaves unchanged the set of accounts each dissenter follows, and the proof of Proposition 2 uses only (10); both propositions therefore hold verbatim, and depth still counts closed triangles. A bought follower even carries no weight for a simpler reason than in the undirected model, since nobody follows her. The undirected structure matters only in Lemma 2, whose nestedness condition is stated for undirected neighbourhoods, and in Sections 6 and 7, where a link is read as mutual attention. There an influencer also attends to her own audience, and this is what allows the users who switch first to tip the influencers in Proposition 3. If influencers attended only to one another, users could not protect their influencers, and that bottom-up channel would disappear.

A related simplification concerns what a collaboration does. On platforms, a collaboration may also bring some of the creator's followers to the client, a channel our accounting ignores. When the new followers share the client's taste, the channel can only help the campaign. Each of them raises the client's audience by one, which adds breadth for every member of J^-, and her own audience by one, which adds depth for each of her dissenting followers, while no dissenter's neighbourhood changes. Every dissenter's exposure then weakly rises, and the argument of Proposition 1(iii) shows that the conversions we compute are a lower bound on those of the campaign. When some of the new followers are dissenters, the channel has no definite sign: they join the client's audience and become exposed to her, but the audience they gain also raises their weight in the evaluation of other dissenters, so the effect of a collaboration depends on how the creator's audience is connected to the rest of the network. In the networks of Sections 6 and 7 the question does not arise, since every follower of a creator the client can recruit already follows her.

Beyond these robustness questions, several extensions remain open. We have identified properties of optimal campaigns without computing them, and selecting the best set of creators in a given network is a combinatorial problem in which, because conversions follow thresholds, the standard guarantees for greedy selection (Kempe et al., 2003) need not apply. In the two-community network collaborations play no role because the client's core is

complete; when it is not, recruited creators can act as bridges between communities, and breadth and depth interact with the community structure. We have also kept collaborations aligned, and recruiting creators of the other type would remove the sign of the campaign. A rival firm promoting the other variety would turn the campaign into a contest, in the spirit of Goyal et al. (2019). Finally, we have studied the effectiveness of audience inflation, not its welfare consequences. Agents who abandon their preferred variety lose utility, platforms detect and remove fake accounts, and followers may unfollow. Evaluating rules such as that of Federal Trade Commission (2024) requires a model that includes these responses, and the accounting developed here is a natural place to start.

A Proofs

Proof of Lemma 1. Fix i and a profile $\mathbf{x}$. By (3) and (6), the accumulated weight of the neighbours whose action coincides with θ_i is F_i^+/F_i, and that of the remaining neighbours is F_i^-/F_i. Hence

$$u_i(\theta_i, \mathbf{x}_{-i}) = \Big(\delta \tfrac{F_i^+}{F_i} + (1-\delta)\tfrac{F_i^-}{F_i}\Big)(1+\lambda),$$
$$u_i(1-\theta_i, \mathbf{x}_{-i}) = \delta \tfrac{F_i^-}{F_i} + (1-\delta)\tfrac{F_i^+}{F_i}.$$

Agent i weakly prefers θ_i if and only if the first expression is at least the second. Multiplying both by $F_i > 0$ and collecting terms,

$$F_i^+\big[\delta(2+\lambda) - 1\big] \;\geq\; F_i^-\big[\delta(2+\lambda) - (1+\lambda)\big].$$

On $\mathcal{R}_C$ we have $\delta(2+\lambda) > 1+\lambda > 1$, so both brackets are strictly positive, and dividing by the first gives $F_i^+ \geq r_{\delta,\lambda} F_i^-$. The numerator of $r_{\delta,\lambda}$ is positive and, since $\lambda > 0$, strictly smaller than its denominator, so $r_{\delta,\lambda} \in (0,1)$. Ties are broken in favour of θ_i. □

Proof of Proposition 1. Fix a campaign (α, S). Let $\mathcal{X}$ be the set of profiles in which every agent of Θ^+ plays θ_v, and for a profile $\mathbf{x}$ write $F_i^{\pm}(\mathbf{x})$ for the masses computed with the connectivities of the original network Γ. Throughout, $\mathbf{x} \in \mathcal{X}$, and $\mathcal{C}(\mathbf{x}) = \{i \in \Theta^- : x_i = \theta_v\}$ denotes its set of converted dissenters.

Step 1: masses in the intervened network. Only the pivot and the members of S change connectivity, and at $\mathbf{x} \in \mathcal{X}$ they all play θ_v. For $i \in \Theta^-$, every neighbour playing θ_i is a dissenter with unchanged connectivity, so her affinity mass in $\Gamma_{(\alpha,S)}$ is $F_i^+(\mathbf{x})$ and her dissent mass is $F_i^-(\mathbf{x}) + E_i(\alpha, S)$. By Lemma 1,

$$\theta_v \text{ is } i\text{'s best response to } \mathbf{x} \iff F_i^+(\mathbf{x}) \;<\; r_{\delta,\lambda}\big(F_i^-(\mathbf{x}) + E_i(\alpha, S)\big). \tag{20}$$

For $i \in \Theta^+$, the neighbours playing the other variety are unconverted dissenters with unchanged connectivity, so her dissent mass is at most F_i^-, and every other neighbour plays θ_v with connectivity at least its original one, so her affinity mass is at least F_i^+. By (9) and Lemma 1, θ_v is her best response.

Step 2: monotonicity. If $\mathcal{C}(\mathbf{x}) \subseteq \mathcal{C}(\mathbf{y})$, then $F_i^+(\mathbf{y}) \leq F_i^+(\mathbf{x})$ and $F_i^-(\mathbf{y}) \geq F_i^-(\mathbf{x})$ for every $i \in \Theta^-$, so the condition in (20) is preserved from $\mathbf{x}$ to $\mathbf{y}$.

Part (i). $\mathbf{x}^0 \in \mathcal{X}$ and $F_i^{\pm}(\mathbf{x}^0) = F_i^{\pm}$. By Step 1, no agent of Θ^+ converts, and $i \in \Theta^-$ converts if and only if $F_i^+ < r_{\delta,\lambda}(F_i^- + E_i(\alpha, S))$, that is, dividing by $r_{\delta,\lambda} > 0$, if and only if $E_i(\alpha, S) > \rho_i$.

Part (ii). Write $\mathcal{C}_t = \mathcal{C}(\mathbf{x}^t)$. By Step 1, $\mathbf{x}^t \in \mathcal{X}$ for every t. We show $\mathcal{C}_t \subseteq \mathcal{C}_{t+1}$ by induction. It holds for $t = 0$, since $\mathcal{C}_0 = \emptyset$. If $\mathcal{C}_{t-1} \subseteq \mathcal{C}_t$ and $i \in \mathcal{C}_t$, then (20) holds at $\mathbf{x}^{t-1}$, hence at $\mathbf{x}^t$ by Step 2, so $i \in \mathcal{C}_{t+1}$. The sets $\mathcal{C}_t$ are thus nested in Θ^-. If $\mathcal{C}_{t+1} = \mathcal{C}_t$, then $\mathbf{x}^{t+1} = \mathbf{x}^t$, the sequence is constant from then on, and its limit is a profile of mutual best responses, a Nash equilibrium of $\Gamma_{(\alpha,S)}$. Since $\mathcal{C}_t$ grows strictly until it stops, this happens after at most $|\Theta^-|$ steps. Finally, the first-round converts are $\mathcal{C}_1(\alpha, S) = \mathcal{C}_1$ and the final converts are $\mathcal{C}^*(\alpha, S) = \mathcal{C}_\infty$, so $\mathcal{C}_1(\alpha, S) \subseteq \mathcal{C}^*(\alpha, S)$.

Part (iii). Let $\mathcal{C}_t$ and $\mathcal{C}'_t$ be the converted sets along the processes of (α, S) and (α', S'), and suppose $\mathcal{C}_t \subseteq \mathcal{C}'_t$, which holds at $t = 0$. If $i \in \mathcal{C}_{t+1}$, then by (20), Step 2 and $E_i(\alpha, S) \leq E_i(\alpha', S')$,

$$F_i^+(\mathbf{x}'^t) \leq F_i^+(\mathbf{x}^t) < r_{\delta,\lambda}\big(F_i^-(\mathbf{x}^t) + E_i(\alpha, S)\big) \leq r_{\delta,\lambda}\big(F_i^-(\mathbf{x}'^t) + E_i(\alpha', S')\big),$$

so $i \in \mathcal{C}'_{t+1}$. Hence $\mathcal{C}_t \subseteq \mathcal{C}'_t$ for every t, which gives both inclusions. □

Proof of Proposition 2. By Proposition 1(iii), any feasible change that gives every dissenter weakly more exposure turns a solution into a solution, for either objective. Since exposure is nondecreasing in α, any solution (α, S) may be replaced by $(\alpha^B(\sigma), S)$, which is feasible and has the same set of creators.

(i) Let $\kappa \leq 0$, so $c_{II} \leq c_I$, and let $\bar{\sigma} = \min\{|\mathcal{A}|, \lfloor B/c_{II} \rfloor\}$. Take a solution $(\alpha^B(\sigma), S)$. While $\sigma < \bar{\sigma}$, add a creator of $\mathcal{A} \setminus S$, which exists since $\sigma < |\mathcal{A}|$. If $\alpha \geq 1$, pay for her by removing one follower: the cost changes by $c_{II} - c_I \leq 0$, and exposure weakly rises by (13). If $\alpha = 0$, she is affordable, since $c_{II}(\sigma + 1) \leq c_{II}\lfloor B/c_{II} \rfloor \leq B$, and she raises m. Then spend the residual budget on followers. The procedure ends at a solution with $\sigma = \bar{\sigma}$. If $\bar{\sigma} = \lfloor B/c_{II} \rfloor < |\mathcal{A}|$, the residual budget is below $c_{II} \leq c_I$, so $\alpha^B(\bar{\sigma}) = 0$; hence followers are bought only if $S = \mathcal{A}$.

(ii)–(iii) Let $\kappa > 0$, and take a solution $(\alpha^B(\sigma), S)$. Let U be a nonempty set of u inner creators in S, and let $\beta(U) = m^B(\sigma - u) - m^B(\sigma)$. The campaign $(\alpha^B(\sigma - u), S \setminus U)$ is feasible, and its scale exceeds that of $(\alpha^B(\sigma), S)$ by $\beta(U)$. By (10), the exposure of $j \in J^-$ changes by $\beta(U) - \Delta T_j(U)$, and the exposure of a dissenter outside N_v does not change, since no such dissenter follows an inner creator. Hence, if $\Delta T_j(U) \leq \beta(U)$ for every $j \in J^-$, the new campaign is again a solution. Moreover, with $x = B/c_I - \kappa\sigma$, (15) gives $\beta(U) = \lfloor x + \kappa u \rfloor - \lfloor x \rfloor \geq \lfloor \kappa u \rfloor$.

Apply this replacement while some nonempty set U of inner creators in the current campaign satisfies $\Delta T_j(U) \leq \beta(U)$ for every $j \in J^-$. Each replacement lowers σ, so the procedure stops at a solution in which every nonempty set U of recruited inner creators has some $j \in J^-$ with $\Delta T_j(U) > \beta(U) \geq \lfloor \kappa|U| \rfloor$. Since $\Delta T_j(U)$ is an integer, $\Delta T_j(U) \geq \lfloor \kappa|U| \rfloor + 1 > \kappa|U|$, so the set R of recruited inner creators is κ-concentrated. If, moreover, no member of J^- follows more than d available inner creators, then $\Delta T_j(R) \leq d$ for every j. If R is nonempty, concentration applied to $U = R$ requires $d \geq \Delta T_j(R) > \kappa|R|$ for some j, that is, $|R| < d/\kappa$; in particular R is empty when $d = 0$. This proves (ii). If $\kappa \geq 1$, a single inner creator s gives $\Delta T_j(\{s\}) \leq 1 \leq \kappa$ for every j, so no nonempty set of inner creators is κ-concentrated, $R = \emptyset$, and every recruited creator is a bridge. This proves (iii). □

Proof of Lemma 2. We first show that $|N_i| \leq |N_j|$ implies $N_i \setminus \{j\} \subseteq N_j$. If this inclusion fails, then $N_j \setminus \{i\} \subseteq N_i$, and since $j \notin N_j$, in fact $N_j \setminus \{i\} \subseteq N_i \setminus \{j\}$. Because $i \in N_j$ if and only if $j \in N_i$, the two sets differ in size by $|N_j| - |N_i| \geq 0$, so the inclusion is an equality and $N_i \setminus \{j\} = N_j \setminus \{i\} \subseteq N_j$, a contradiction.

(i) Let $|N_i| \leq |N_{i'}|$, so that $N_i \setminus \{i'\} \subseteq N_{i'}$. The pivot and the members of S have type θ_v and are therefore different from i'. Hence $v \in N_i$ implies $v \in N_{i'}$, and $S \cap N_i \subseteq S \cap N_{i'}$. By (10), $E_i(\alpha, S) \leq E_{i'}(\alpha, S)$.

(ii) Let $s \in \mathcal{A}$, so that $s \neq v$ and $s \notin N_v$. If $|N_s| \leq |N_v|$, then $N_s \subseteq N_v$, so every dissenting follower of s belongs to J^- and s is inner. If s is a bridge, then $|N_s| > |N_v|$, so $N_v \subseteq N_s$ and every member of J^- follows s.

(iii) Let j^* be a member of J^- with the largest number of neighbours, and let $s \in \mathcal{A}$ be followed by $j \in J^-$. Since $|N_j| \leq |N_{j^*}|$, we have $N_j \setminus \{j^*\} \subseteq N_{j^*}$, and $s \neq j^*$ because $s \in \Theta^+$. Hence $s \in N_{j^*}$. If R is a set of inner creators, each followed by some member of J^-, then every nonempty $U \subseteq R$ satisfies $\Delta T_{j^*}(U) = |U| > \kappa|U|$ whenever $\kappa < 1$, so R is κ-concentrated. □

Proof of Proposition 3. By Proposition 1(ii), agents of Θ^+ never move and converts never return. Two facts are used repeatedly. (a) Once every dissenting influencer plays θ_v, a dissenting user has affinity mass zero and positive dissent mass, so she converts by Lemma 1. (b) If $\mathcal{C}_1 = \emptyset$, the process stops at $\bar{\theta}$ and $\mathcal{C}^* = \emptyset$.

(i) Every dissenter has exposure α, so by Proposition 1(i) users convert in the first round if and only if $\alpha > \rho_P$, and influencers if and only if $\alpha > \rho_C$. If $\alpha > \rho_C$, influencers convert in the first round and users follow by (a). If $\rho_P < \alpha \leq \rho_C$, only users convert in the first round. In the second round a dissenting influencer has affinity mass $(k^- - 1)f_C$ and dissent mass $k^+ f_C + n f_P + \alpha$, so she converts if and only if $\alpha > \rho_C - (1 + r_{\delta,\lambda})n^- f_P / r_{\delta,\lambda}$. By (17), $\rho_C - \rho_P < n^- f_P / r_{\delta,\lambda}$, so this threshold lies below $\rho_P < \alpha$, and influencers convert. If $\alpha \leq \min\{\rho_P, \rho_C\}$, apply (b).

(ii) Dissenting users have exposure zero and do not convert in the first round, while dissenting influencers convert if and only if $\alpha + 2\sigma > \rho_C$. In that case users follow by (a); otherwise apply (b). □

Proof of Corollary 1. By Proposition 3(ii) and its proof, both $\mathcal{C}_1(\alpha, S)$ and $\mathcal{C}^*(\alpha, S)$ are empty if $\alpha + 2\sigma \leq \rho_C$ and equal to $C \cap \Theta^-$ and Θ^-, respectively, otherwise. A feasible campaign that maximizes $\alpha + 2\sigma$ therefore solves the problem for either objective. The available creators are the $n^+ - 1$ aligned users other than v, and with $\alpha = \alpha^B(\sigma)$ the campaign delivers $V(\sigma) = 2\sigma + \lfloor (B - c_{II}\sigma)/c_I \rfloor$ for $0 \leq \sigma \leq \bar{\sigma}$. For $\sigma < \bar{\sigma}$, write $y = B - c_{II}\sigma$, so that

$$V(\sigma + 1) - V(\sigma) \;=\; 2 - \left(\left\lfloor \frac{y}{c_I} \right\rfloor - \left\lfloor \frac{y - c_{II}}{c_I} \right\rfloor \right).$$

For every real a and $b > 0$, $\lfloor b \rfloor \leq \lfloor a \rfloor - \lfloor a - b \rfloor \leq \lceil b \rceil$. With $b = c_{II}/c_I$, the term in brackets is at most $\lceil c_{II}/c_I \rceil \leq 2$ when $c_{II} < 2c_I$, and at least $\lfloor c_{II}/c_I \rfloor \geq 2$ when $c_{II} > 2c_I$. Hence V is nondecreasing in the first case, so that $\sigma = \bar{\sigma}$ maximizes it, and nonincreasing in the second, so that $\sigma = 0$ does, with $\alpha^B(0) = \lfloor B/c_I \rfloor$. □

Proof of Proposition 4. By Proposition 1(ii), agents of Θ^+ never move, converts never return, and the process stops at $\bar{\theta}$ when $\mathcal{C}_1 = \emptyset$. This settles (ii), since every dissenter then has exposure zero, and the case $\alpha \leq \hat{\rho}_P$ of (i).

Let $v \in C^+$ and $\alpha > \hat{\rho}_P$. In the first round exactly the b members of J^- convert. Dissenting users outside J^- are linked only to influencers, none of whom has moved, so their masses do not change. Each dissenting influencer is linked to every member of P^-, so her affinity mass falls to $(k^- - 1)f_C^- + (n^- - b)f_P^-$ and her dissent mass rises to bf_P^-. By Lemma 1 she converts if and only if $(k^- - 1)f_C^- + (n^- - b)f_P^- < r_{\delta,\lambda} b f_P^-$, that is, if and only if $(1 + r_{\delta,\lambda})bf_P^- > r_{\delta,\lambda}\hat{\rho}_C$. If she does not, the process stops and $\mathcal{C}^* = J^-$. If she does, every remaining dissenting user has affinity mass zero and positive dissent mass in the next round, so she converts, and $\mathcal{C}^* = \Theta^-$. This proves (i).

For (iii), $\mathcal{C}^* = \Theta^-$ requires the last case of (i). Since $r_{\delta,\lambda}\hat{\rho}_C = (k^- - 1)f_C^- + n^- f_P^- \geq n^- f_P^-$, it implies $(1 + r_{\delta,\lambda})b > n^-$. Using $k^+ b = n^- \ell$, this gives $\ell > k^+/(1 + r_{\delta,\lambda})$, and combining it with the status-quo bound $r_{\delta,\lambda}\ell f_C^+ \leq k^- f_C^-$ yields (19). □

B The networks of Example 1

Let $\delta = 7/9$ and $\lambda = 1/4$, which lie in $\mathcal{R}_C$ and give $r_{\delta,\lambda} = 2/3$, so that $\rho_i = \frac{3}{2}F_i^+ - F_i^-$. Both networks share the following structure. There is a clique K^- of 21 dissenters and a clique K^+ of 60 agents of type θ_v. The pivot v is linked to the six members of J^-, three easy members e_1, e_2, e_3 and three hard members h_1, h_2, h_3, and to three members of K^+, so that $f_v = 8$. Each member of J^- is linked to one member of K^-, a different one for each. Each hard member h_k is, in addition, linked to a dissenter w_k of degree two, whose other neighbour is a further member of K^-, again a different one for each k. The creators s_1, s_2, s_3 have type θ_v, are not linked to v, and are linked to their followers in J^- and to members of K^+ not linked to v, with different creators linked to different members of K^+.

- In the clustered network, each s_k is linked to all six members of J^- and to one member of K^+, so that $f_{s_k} = 6$.
- In the scattered network, s_k is linked to e_k, to h_k and to 17 members of K^+, so that $f_{s_k} = 18$.

Resistances. The neighbour of $j \in J^-$ in K^- has degree 21, so $f = 20$, and w_k has $f = 1$. Hence $F_{e_k}^+ = 20$ and $F_{h_k}^+ = 21$ in both networks. The dissent mass of every member of J^- is $f_v + 3 \cdot 6 = 26$ in the clustered network and $f_v + 18 = 26$ in the scattered one. Therefore $\rho_{e_k} = 30 - 26 = 4$ and $\rho_{h_k} = 31.5 - 26 = 5.5$ in both networks.

Status quo. Members of $K^\pm$ and the agents w_k have no neighbour of the other type, so their dissent mass is zero. The members of J^- have nonnegative resistance. The pivot has affinity mass at least $3 \cdot 59$ and dissent mass $\sum_{j \in J^-} f_j$, which is 27 in the clustered network and 15 in the scattered one. Each creator has affinity mass at least 59 and dissent mass at most 27. Since $r_{\delta,\lambda} < 1$, (9) holds for every agent, and every agent has degree at least two.

No spread beyond J^-. Members of J^- are not linked to one another. By Proposition 1(ii), agents of type θ_v never move. The remaining dissenters outside J^- are the members of K^- and the agents w_k. After any conversions in J^-, a member of K^- keeps an affinity mass of

at least $19 \cdot 19$ against a dissent mass of at most 5, and w_k keeps affinity mass 20 against dissent mass at most 5. Neither converts, so $\mathcal{C}^* = \mathcal{C}_1$ for every campaign.

Other creators. The members of K^+ not linked to v are also available, but they have no dissenting followers. Recruiting one adds no depth and costs breadth, so by the replacement step in the proof of Proposition 2 it is weakly dominated, and we may ignore them. With $\alpha = \alpha^B(\sigma)$ and $S \subseteq \{s_1, s_2, s_3\}$, the exposures and conversions are those of Table 1.

References

Anand, A., Dutta, S., Mukherjee, P., 2018. Influencer marketing with fake followers. SSRN Working Paper 3306088. doi:10.2139/ssrn.3306088.

Ballester, C., Calvó-Armengol, A., Zenou, Y., 2006. Who's who in networks. Wanted: The key player. Econometrica 74, 1403–1417. doi:10.1111/j.1468-0262.2006.00709.x.

Banerjee, A., Chandrasekhar, A.G., Duflo, E., Jackson, M.O., 2013. The diffusion of microfinance. Science 341, 1236498. doi:10.1126/science.1236498.

Battiston, P., Stanca, L., 2015. Boundedly rational opinion dynamics in social networks: Does indegree matter? Journal of Economic Behavior & Organization 119, 400–421. doi:10.1016/j.jebo.2015.08.013.

Belhaj, M., Bervoets, S., Deroïan, F., 2016. Efficient networks in games with local complementarities. Theoretical Economics 11, 357–380. doi:10.3982/TE1742.

Borgatti, S.P., Everett, M.G., 2000. Models of core/periphery structures. Social Networks 21, 375–395. doi:10.1016/S0378-8733(99)00019-2.

Bramoullé, Y., Kranton, R., D'Amours, M., 2014. Strategic interaction and networks. American Economic Review 104, 898–930. doi:10.1257/aer.104.3.898.

Centola, D., 2010. The spread of behavior in an online social network experiment. Science 329, 1194–1197. doi:10.1126/science.1185231.

Centola, D., Macy, M., 2007. Complex contagions and the weakness of long ties. American Journal of Sociology 113, 702–734. doi:10.1086/521848.

Cinelli, M., De Francisci Morales, G., Galeazzi, A., Quattrociocchi, W., Starnini, M., 2021. The echo chamber effect on social media. Proceedings of the National Academy of Sciences 118, e2023301118. doi:10.1073/pnas.2023301118.

Dasaratha, K., Shah, A., 2026. Network interventions: Targeting agents or targeting links? arXiv preprint 2602.12897. doi:10.48550/arXiv.2602.12897.

De Veirman, M., Cauberghe, V., Hudders, L., 2017. Marketing through Instagram influencers: the impact of number of followers and product divergence on brand attitude. International Journal of Advertising 36, 798–828. doi:10.1080/02650487.2017.1348035.

DeMarzo, P.M., Vayanos, D., Zwiebel, J., 2003. Persuasion bias, social influence, and unidimensional opinions. Quarterly Journal of Economics 118, 909–968. doi:10.1162/00335530360698469.

Fainmesser, I.P., Galeotti, A., 2021. The market for online influence. American Economic Journal: Microeconomics 13, 332–372. doi:10.1257/mic.20200050.

Federal Trade Commission, 2024. Trade regulation rule on the use of consumer reviews and testimonials. Final rule, *Federal Register* 89(163), 68034–68079, 22 August 2024; codified at 16 CFR Part 465.

Feld, S.L., 1991. Why your friends have more friends than you do. American Journal of Sociology 96, 1464–1477. doi:10.1086/229693.

Galeotti, A., Golub, B., Goyal, S., 2020. Targeting interventions in networks. Econometrica 88, 2445–2471. doi:10.3982/ECTA16173.

Galeotti, A., Goyal, S., 2009. Influencing the influencers: a theory of strategic diffusion. RAND Journal of Economics 40, 509–532. doi:10.1111/j.1756-2171.2009.00075.x.

Galeotti, A., Goyal, S., Jackson, M.O., Vega-Redondo, F., Yariv, L., 2010. Network games. Review of Economic Studies 77, 218–244. doi:10.1111/j.1467-937X.2009.00570.x.

Golub, B., Jackson, M.O., 2010. Naïve learning in social networks and the wisdom of crowds. American Economic Journal: Microeconomics 2, 112–149. doi:10.1257/mic.2.1.112.

Golub, B., Jackson, M.O., 2012. How homophily affects the speed of learning and best-response dynamics. Quarterly Journal of Economics 127, 1287–1338. doi:10.1093/qje/qjs021.

Goyal, S., Heidari, H., Kearns, M., 2019. Competitive contagion in networks. Games and Economic Behavior 113, 58–79. doi:10.1016/j.geb.2014.09.002.

Granovetter, M., 1978. Threshold models of collective behavior. American Journal of Sociology 83, 1420–1443. doi:10.1086/226707.

He, S., Hollenbeck, B., Proserpio, D., 2022. The market for fake reviews. Marketing Science 41, 896–921. doi:10.1287/mksc.2022.1353.

Hernández, P., Martínez-Cánovas, G., Muñoz-Herrera, M., Sánchez, Á., 2017. Equilibrium characterization of networks under conflicting preferences. Economics Letters 155, 154–156. doi:10.1016/j.econlet.2016.12.004.

Hernández, P., Muñoz-Herrera, M., Sánchez, Á., 2013. Heterogeneous network games: Conflicting preferences. Games and Economic Behavior 79, 56–66. doi:10.1016/j.geb.2013.01.004.

Hinz, O., Skiera, B., Barrot, C., Becker, J.U., 2011. Seeding strategies for viral marketing: An empirical comparison. Journal of Marketing 75, 55–71. doi:10.1509/jm.10.0088.

Jackson, M.O., Rogers, B.W., Zenou, Y., 2017. The economic consequences of social-network structure. Journal of Economic Literature 55, 49–95. doi:10.1257/jel.20150694.

Katona, Z., Zubcsek, P.P., Sarvary, M., 2011. Network effects and personal influences: The diffusion of an online social network. Journal of Marketing Research 48, 425–443. doi:10.1509/jmkr.48.3.425.

Kempe, D., Kleinberg, J., Tardos, É., 2003. Maximizing the spread of influence through a social network, in: Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ACM. pp. 137–146. doi:10.1145/956750.956769.

Kor, R., Liu, Y., Zenou, Y., Zhou, J., 2025. Welfare and distributional effects of joint intervention in networks. Journal of Economic Theory 230, 106109. doi:10.1016/j.jet.2025.106109.

Lerman, K., Yan, X., Wu, X.Z., 2016. The "majority illusion" in social networks. PLoS ONE 11, e0147617. doi:10.1371/journal.pone.0147617.

Leung, F.F., Gu, F.F., Li, Y., Zhang, J.Z., Palmatier, R.W., 2022. Influencer marketing effectiveness. Journal of Marketing 86, 93–115. doi:10.1177/00222429221102889.

Li, X., 2023. Designing weighted and directed networks under complementarities. Games and Economic Behavior 140, 556–574. doi:10.1016/j.geb.2023.04.010.

López-Pintado, D., 2006. Contagion and coordination in random networks. International Journal of Game Theory 34, 371–381. doi:10.1007/s00182-006-0026-5.

Luca, M., Zervas, G., 2016. Fake it till you make it: Reputation, competition, and Yelp review fraud. Management Science 62, 3412–3427. doi:10.1287/mnsc.2015.2304.

Mayzlin, D., Dover, Y., Chevalier, J.A., 2014. Promotional reviews: An empirical investigation of online review manipulation. American Economic Review 104, 2421–2455. doi:10.1257/aer.104.8.2421.

Morris, S., 2000. Contagion. Review of Economic Studies 67, 57–78. doi:10.1111/1467-937X.00121.

Neary, P.R., 2012. Competing conventions. Games and Economic Behavior 76, 301–328. doi:10.1016/j.geb.2012.06.003.

Nora, V., Winter, E., 2024. Exploiting social influence in networks. Theoretical Economics 19, 1–27. doi:10.3982/TE5068.

Orlova, O., 2022. Idiosyncratic preferences in games on networks. Games and Economic Behavior 131, 29–50. doi:10.1016/j.geb.2021.10.006.

Parise, F., Ozdaglar, A., 2023. Graphon games: A statistical framework for network games and interventions. Econometrica 91, 191–225. doi:10.3982/ECTA17564.

Sun, Y., Zhao, W., Zhou, J., 2023. Structural interventions in networks. International Economic Review 64, 1533–1563. doi:10.1111/iere.12639.

Sun, Y., Zhao, W., Zhou, J., 2026. Sequential network design. Journal of Economic Theory 232, 106145. doi:10.1016/j.jet.2026.106145.

Watts, D.J., Dodds, P.S., 2007. Influentials, networks, and public opinion formation. Journal of Consumer Research 34, 441–458. doi:10.1086/518527.

Yang, A., Saffer, A.J., 2021. Standing out in a networked communication context: Toward a network contingency model of public attention. New Media & Society 23, 2902–2925. doi:10.1177/1461444820939445.

Young, H.P., 2009. Innovation diffusion in heterogeneous populations: Contagion, social influence, and social learning. American Economic Review 99, 1899–1924. doi:10.1257/aer.99.5.1899.