Optimal Network Design for Inducing Effort

Pinar Yildirim, Yanhao Wei, Christophe Van den Bulte, Joy Lu*

Abstract

Many companies create and manage communities where consumers observe and exchange information about the effort expended by other consumers. Such communities are especially popular in the areas of fitness, education, dieting, and financial savings. We study how to optimally structure such consumer communities when the objective is to maximize the total or average amount of effort expended. Using mathematical modeling and assuming peer influence through conformity, we find that the optimal network design consists of a set of disconnected or very loosely connected sub-graphs, each of which is very densely connected within. Also, each sub-community in the optimal network consists of consumers selected such that their "standalone" propensity to exert effort negatively correlates with their propensity to conform and positively correlates with their propensity to influence others.

Keywords: customer communities, social networks, network design, network optimization.

*Yildirim: Assistant Professor of Marketing, The Wharton School, University of Pennsylvania, email: pyild@wharton.upenn.edu. Wei: Assistant Professor of Marketing, Marshall School of Business, University of Southern California. Van den Bulte: Gayfryd Steinberg Professor and Professor of Marketing, The Wharton School, University of Pennsylvania. Lu: Doctoral Candidate, The Wharton School, University of Pennsylvania. All correspondence about the manuscript can be directed to the first author. We thank Anthony Dukes and Upender Subramanian for comments on an earlier draft.

1 Introduction

Over the last two decades, marketers have become increasingly keen on leveraging peer influence among customers. Today, may companies go beyond viral marketing in already existing networks, and bring customers together in communities to help them better use their products and services, share information about their experiences, share or develop ideas and solutions, and support and motivate product and service use in general. In short, rather than merely leveraging pre-existing networks, marketers increasingly create networks.

Another notable change in marketers' mindset is the renewed appreciation by marketers that peer influence can be leveraged even among people who hardly know each other or even total strangers. In the absence of a clear benchmark about what is appropriate or desirable, people tend to conform to what others do, even if those are strangers. This insight goes back to famous experiments showing that whether people put into a group with strangers stated that a light dot was moving or not (Sherif, 1935) or that one line was longer than another (Asch, 1956) was influenced by what those strangers stated. More recently, field experiments have shown that people are more likely to re-use towels in a hotel, to make charitable donations, to consume less energy, to consume less water, to exercise, to pay overdue taxes, and to rate movies when they see or are told that others do so, even when they have no prior connection to these others (e.g., Frey and Meier, 2004; Goldstein et al., 2008; Martin and Randal, 2008; Shang and Croson, 2009; Chen et al., 2010; Allcott, 2011; Ferraro and Price, 2013; Ayres et al., 2013; Allcott and Rogers, 2014; Hallsworth et al., 2017). Along similar lines, a recent field study documented that physicians were more likely to keep prescribing a new drug if their colleagues in the same practice or hospital did so, even if those were not the peers they turned to for discussing treatment options or referring patients (Iyengar et al., 2015).

Numerous firms and platforms connect consumers who do not know each other into online communities and allow them to observe each others' activity to increase the level or consistency of the effort they exert to achieve a goal. BuddyUp helps college students to form study groups based on shared classes and majors. Quitnet helps smokers and ex-smokers make a daily pledge to quit (or stay quit), and share with others how well they are doing. Members of the ALS community on PatientsLikeMe share data allowing them to see what actions are taken by other patients and whether their progress is fast, slow, or about average. Opower, now part of Oracle, works with many utilities to share energy consumption information from similar households in one's vicinity to reduce energy consumption. Online platforms like Crowdrise and Razoo allow users to observe each others' charitable donations. Many other companies and organizations use face-to-face rather than online communities to help customers or members put in the effort necessary to achieve their goals. Weight Watchers and Alcoholics Anonymous are perhaps the most famous. Though the meetings are an opportunity to provide support by sharing tips and stories, they also provide a benchmark or descriptive social norm of how hard others are working towards the objective. In all these examples, the community manager is interested in leveraging norm-based peer influence to increase the overall or average effort exerted towards a goal.

In this paper, we study the optimal design of communities when consumers exert effort subject to normative peer influence. Specifically we address the following research questions: How should the community manager design the network of ties among the community members to maximize their total effort? Which consumer should share information with which other consumer? Is it more effective to have all community members exchange information with every other member, to organize members in distinct sub-communities, or to have star-like structures where some members are more central than others? Regardless of the pattern of the ties, should tie strength be varied to make about some consumer's effort or activity more prominent than that of others?

Our analysis focuses on communities with three characteristics. First, the community manager seeks to maximize the average (or total) level of effort that community members exert towards achieving their own goal, like losing weight or consuming less electricity. Second, the community manager has the opportunity to organize the network of ties among the members freely, e.g., by forming virtual or physical groups or by managing the person-to-person flow of information. Third, the decision to exert effort is subject to normative peer influence, such that people prefer to act in accordance to the descriptive norm set by the effort level of other community members they are connected to.

The first two characteristics are present in the company and platform examples we just provided and need no further discussion. The third characteristic, though quite familiar to behavioral researchers, may require some elaboration. Social conformity through adherence to the descriptive norm set by peers' behavior occurs quite frequently when people are not sure about what the adequate or appropriate behavior is or when people do not want to embarrass themselves by over- or under-achieving. People do so even when deviating from the norm is not subject to punitive action or when status concerns and other signaling considerations are ruled out (e.g., Asch, 1956; Cialdini and Goldstein, 2004; Cialdini et al., 1990; Festinger, 1954; Kelman, 1958; Zafar, 2011). As Bernheim and Exley (2015) note, conformity in such situations arises through a preference mechanism rather than a belief mechanism. Consequently, the consumer model in our analysis of optimal network design is set up differently from the model by Bernheim (1994) emphasizing belief rather than preference mechanisms of conformity.

Our analysis produces three main insights:

1. The optimal network structure of a community consists of a set of disconnected or very loosely connected sub-graphs or sub-communities, each of which is very densely connected within.

2. In the optimal network, each sub-community consists of members selected such that their "standalone" propensity to exert effort negatively correlates with their propensity to conform and positively correlates with their propensity to influence others.

3. The improvement from linking community members optimally increases with the size of the

community and with the heterogeneity in members' (i) propensity to exert effort in isolation, (ii) social influence, and (iii) susceptibility to social influence.

The key intuition for these three insights is that conformity to descriptive social norms pulls down the effort of these who are prone to exert high effort in the absence of peer influence, but pulls up the effort of those who are prone to exert low effort in the absence of peer influence. Hence, to maximize the upside and minimize the downside of conformity, the network designer will connect influential high-effort people with influenceable low-effort people. This matching can be best achieved within sub-communities between which no ties exist. Also, because every tie is conduit of influence, conformity is leveraged to the maximum in sub-communities that are internally fully connected. Finally, the greater the number of people in the network, and the greater the variation in the propensities to exert effort in isolation, to influence others, and to be influenced, the greater the degrees of freedom that the network has in assigning specific people to communities to achieve optimal matching.

The optimal pattern of ties is very easy to implement: It amounts to simply putting people into groups where everyone is exposed to information from everyone else. This can be done conveniently in online as well as offline networks. As to the propensities to exert effort, to be influential, and to be susceptible to influence, these can be assessed directly using questionnaires administered when people sign up to become members of the network (e.g., Flynn et al., 1996; Duckworth et al., 2007) or indirectly using easy-to observe correlates of these propensities (e.g., Aral and Walker, 2012).

The three insights are derived for the setting where peer influence operates through conformity. The optimal structure is very different when peer influence operates through competition. When individuals seek to out-compete their peers rather than to conform, the optimal network structure for inducing maximum total or average effort is the completely connected network where every member is connected to every other member.

The rest of the paper is organized as follows. Section 2 provides a summary of the relevant literature and this paper's contributions to it. Section 3 describes the benchmark network model. Section 4 characterizes the optimal community design and its characteristics. Section 5 studies contingencies that may effect our key results. Section 6 concludes with implications for practice and research.

2 Related Literature in Marketing and Network Science

Marketing practitioners and academics have long been keen on understanding how to best leverage peer influence in social networks. Research attention at first focused on documenting peer influence in field settings (e.g., Van den Bulte and Lilien, 2001; Aral et al., 2009). More recent research has focused on identifying the mechanisms at work (e.g., Chen et al., 2011; Iyengar et al., 2015), characterizing which customers, if any, exert a disproportional amount of influence on others (e.g., Aral and Walker, 2012; Goldenberg et al., 2009; Iyengar et al., 2011; Katona et al., 2011), and leveraging either homophily or peer influence through efficiently targeted marketing effort (e.g., Hill et al., 2006; Godes and Mayzlin, 2009; Goel and Goldstein, 2013; Bimpikis et al., 2016; Ascarza et al., 2017). This prior work takes the network structure, be it static or evolving, as a given for the marketer. Most recently, research has started to view the network as an outcome or even a choice variable. Some work investigates how marketing actions might alter the network to the marketers' benefit (Ansari et al., 2016) and how network structures result endogenously from members' choices affected by a marketer's or platform manager's policies (Iyer and Katona, 2015; Wei et al., 2016; Phan and Godes, 2017). Yet other work compares the effectiveness of various exogenously given network structures for a particular marketing objective (Murtha et al., 2014; Peres and Van den Bulte, 2014).

The present study raises the hitherto unaddressed question of optimal network design. More specifically, we investigate the following question: How can a network designer or central planner choose a network structure such that the overall effort expended by members is maximized? Unlike prior work involving comparative assessments of multiple network structures, we seek to directly characterize the optimal structure. Unlike prior work where the network emerges endogenously from members' choices in tie formation, we endogenize network structure as a decision made by a central planner or community manager. Unlike prior work where marketers aim to improve their marketing ROI by targeting specific nodes within a network or by increasing the virality of ties, we take the view of a marketer or community manager intervening by choosing the network structure; i.e., choosing the pattern of ties connecting the given set of nodes.

Network scientists and policymakers are becoming increasingly interested in using interventions that manipulate or create the network structure (Cerdeiro et al., 2017; Haag and Lagunoff, 2006; Kraut et al., 2012; Valente, 2012). These are a departure from interventions that leverage an existing structure by targeting particularly influential or influenceable nodes or by altering the virality of ties. Our study is part of this emerging research on effective networks design.

Our contribution to purposive social network design and forward engineering of networks (Alderson, 2008) assumes that members who join a particular community to achieve a specific goal like losing weight can be put into a designed network structure in spite of people's tendency to form homophilous ties in organic networks. More specifically, we make two assumptions about network membership and network rewiring.

First, purposive network design requires that people are willing to join and remain members of a community with an exogenously imposed network structure. Every year, millions of freshmen students across the globe willingly agree to be assigned—often in explicitly random fashion—to dormitory rooms, student cohorts, study groups and course project teams by the educational institutions they join in their pursuit of an education. Allcott and Kessler (2015) show that a majority of customers receiving Opower social comparison reports for free from their electric utility are willing to pay for such information without being able to select the people they receive benchmark information about. Five field experimental studies in the realm of health behavior and physical fitness by Centola (2010; 2011), Zhang et al. (2015a); Rovniak et al. (2016) and Zhang et al. (2016) similarly document people's willingness to relinquish their agency in network formation and allow themselves to be placed in a network structure designed by the manager of a community they join. In short, network design does not require some dictatorial Leviathan, and is willingly accepted by many people joining communities built and managed to help them reach a specific objective.

Second, we assume that after people are given a particular position in the network they joined to achieve a specific goal, they do not immediately and radically rewire their set of connections, which would make the initial design ineffective. This assumption is consistent with ample evidence that exogenous assignment to particular apartments, offices, or dorm rooms has long-run effects on both the subsequent pattern of organically developed social ties (e.g., Festinger et al., 1950; Van den Bulte and Moenaert, 1998; Hasan and Bagde, 2013) and the level of achievement many months after the initial assignment (e.g., Hasan and Bagde, 2013; Sacerdote, 2001). That network design affects interaction and outcomes even after network members form social ties organically is consistent with a core proposition of Festinger's (1954) theory of conformity through social comparison: People dislike being in a state of discordance with peers, and often resolve that state of cognitive dissonance by changing their attitude and behavior rather than severing the social ties with their discordant peers. It is this very tendency that underlies social conformity through descriptive norms, amply in both lab and field experiments. The same behavioral principle of discordance aversion explains the pervasive tendency for organic networks to exhibit tie transitivity, balance, and clustering (e.g., Cartwright and Harary, 1956; Davis, 1970).

3 Model Setup and Equilibrium

We consider a set of n individuals or consumers who decide how much effort to exert in order to reach an objective, and a network designer who wants to organize these individuals into a network in order to maximize the total, or equivalently, the average amount of effort exerted. Any two individuals may be connected to each other, and if i and j are connected, we denote the tie by $i \sim j$. A network, denoted generically by G, is simply the collection of the individuals and their ties. We start by characterizing the "out-of network" effort in the absence of any ties, and then move to establishing the equilibrium effort level in a network with ties.[1] We assume that the consumer's out-of-network utility from exerting effort at level y_i is:

$$u_i = \alpha_i y_i - \frac{1}{2} y_i^2 \tag{1}$$

[1] We use the terms effort and activity interchangeably.

where α_i is the idiosyncratic reward from working towards an objective, and there is a quadratic cost y_i^2 of effort. Consumers are heterogeneous with respect to how much they care about the objective. We assume $\alpha_i \in [0, 1]$. If the objective is weight loss, for instance, consumers vary in how much they trade off the health and physical appearance benefits of one hour of exercise against the time loss and physical discomfort involved. Results easily scale, mutatis mutandis, to larger bounded intervals for α_i.

A consumer exerting effort within a network or community of consumers gains utility not only from working towards the objective but also from conforming to the activity of others in the network. We represent the in-network utility as:

$$u_i = \alpha_i y_i - \frac{1}{2} y_i^2 + \frac{1}{2} \sum_{j \sim i} [1 - \rho_{ij}(y_i - y_j)^2] \tag{2}$$

In the expression, $j \sim i$ denotes any individual j that is connected to i and vice versa, and $\rho_{ij} > 0$ calibrates the social influence of j on i. The last term $\sum_{j \sim i}[1 - \rho_{ij}(y_i - y_j)^2]$ captures the utility from adhering to the descriptive norm, and is the component which represents peer influence. The influence parameter ρ_{ij} reflects both the ability of j to influence others, and the susceptibility of i to influence by others:

$$\rho_{ij} = \lambda_i \cdot \chi_j, \tag{3}$$

where $\lambda_i > 0$ calibrates how sensitive i is to observed peer behavior and $\chi_j > 0$ calibrates how much influence j has over others. We allow the influence between any two connected consumers to be asymmetric ($\rho_{ij} \neq \rho_{ji}$) and also allow the activity of i to be greatly influenced by j's activity ($\rho_{ij} \gg 0$) or at a marginal rate ($\rho_{ij} \to 0$).[2] These individual-level influence and susceptibility parameters will be central to the characterization of the optimal network design.

This framework captures an individual's desire to conform to the behaviors of others they are connected to. Specifically, conformity enters the utility function directly through the quadratic term, consistent with how Bernheim and Exley (2015, pg. 10) model descriptive normative influence as a preference mechanism. In 5.2, we investigate the optimal network design for an alternate formulation in which individuals want to out-compete their connected peers by exerting more effort than them, rather than adhere to the descriptive norm set by their peers.

When a customer i exerts effort alone, his utility-maximizing level of activity is α_i. Thus we also refer to α_i as the out-of-network effort level. However, when he is maximizing the utility given in Equation (2), the chosen effort level is a weighted average of α_i and the activity of his contacts:

$$y_i = \frac{\alpha_i + \sum_{j \sim i} \rho_{ij} y_j}{1 + \sum_{j \sim i} \rho_{ij}} \tag{4}$$

[2] Also ρ_{ij} need not equal ρ_{ik}, $k \neq j$, and i may be influenced by different individuals in his network at different rates. Note, ρ_{ij} and α_i are calibrated against the cost of effort, which is assumed to take a constant value of 1.

The optimal personal effort in (4) is not specific to the utility function formulation in (2). Alternate functional formulations, for instance $u_i = \alpha_i y_i - \frac{1}{2}y_i^2 - \frac{1}{2}\sum_{j\sim i}\rho_{ij}(y_i - y_j)^2$, also lead to the expression in (4). All of our subsequent analyses are based on (4) and hold with any utility specification that gives (4). Utility specification in (2) has the benefit that the individual is never better off outside the network, such that participation is guaranteed, if $\rho_{ij} \in [0, 1]$ for all i and j.

We consider a Nash equilibrium where all individuals choose their effort simultaneously. In what follows, y_i^* denotes the equilibrium effort of individual i and vector $\mathbf{y}^*$ collects the equilibrium effort levels of all the network or community members. Proposition 1 establishes the existence and uniqueness of an equilibrium, two important properties when designing a network.

Proposition 1. *For a fixed network G, there exists a unique equilibrium of individual activities, $\mathbf{y}^*$. Customer i's equilibrium activity satisfies:*

$$y_i^* = \alpha_i + \sum_{j\sim i}\rho_{ij}(y_j^* - y_i^*) \tag{5}$$

The equilibrium is given by

$$\mathbf{y}^* = \mathbf{L}^{-1}\boldsymbol{\alpha} \tag{6}$$

where $\mathbf{L}$ is a square "influence matrix" of size n, given by

$$\begin{aligned} L_{ij} &= \begin{cases} -\rho_{ij}, & \textit{if } i \sim j \\ 0, & \textit{otherwise} \end{cases} \\ L_{ii} &= 1 - \sum_{j\sim i} L_{ij}. \end{aligned}$$

Proposition 1 demonstrates (i) how an individual's in-network activity is linked to not only his own out-of-network activity, but also the activity of the other individuals, and (ii) how social influence is determined by both ρ_{ij} and the topology of the network (the set of other consumers he is connected to).

More specifically, Equation (5) states that at equilibrium, one's activity is boosted by each higher-activity neighbor, and decreased by each lower-activity neighbor. Equations (5) and (6) imply that an individual's equilibrium in-network effort can sometimes be lower than his out-of-network effort, that the same is true for the overall level of effort, and that any connection can influence anyone's equilibrium level of effort. Consequently, to increase the total collective effort exerted by customers or community members, managers may need to purposely design the community structure. For instance, in the area of physical exercise, club and platform managers may want to carefully consider how to organize their customer-members into physical groups or who to connect with whom through virtual peer connections.

4 Optimal Network or Community Design

Given the utility formulation in (2), the total effort $T \equiv \sum_{i \in G} y_i^*$ can be maximized by determining which consumers' information should be shared with which other consumer, and the network designer has to answer this question simultaneously for all customers. Suppose we consider only two individuals with $\alpha_1 > \alpha_2$. It is easy to see that placing a tie between them increases their joint activity iff consumer 1 influences consumer 2 more than consumer 2 influences consumer 1 (that is, iff $\rho_{21} > \rho_{12}$). When applying such a pairwise comparison rule for $n > 2$, any two individuals would be connected iff $(\rho_{ij} - \rho_{ji})(\alpha_j - \alpha_i) > 0$. However, the following example shows that this pairwise design rule could lead to worse outcomes than the empty network (i.e., where everyone is isolated), let alone a comprehensive design rule that considers not only the direct but also the indirect influence of consumers on each other.

Consider a 4-person network. The individuals are characterized by their individual parameters as follows:

$$\boldsymbol{\alpha} = \begin{pmatrix} 0.0 \\ 1.5 \\ 0.5 \\ 0.1 \end{pmatrix}, \boldsymbol{\lambda} = \begin{pmatrix} 2.0 \\ 1.5 \\ 0.5 \\ 0.1 \end{pmatrix}, \boldsymbol{\chi} = \begin{pmatrix} 1 \\ 1 \\ 1 \\ 1 \end{pmatrix}$$

The community design produced by the pairwise rule is given in the network below on the left. All consumers are placed in a star structure where consumer 1 is connected to every other consumer. The equilibrium activity of each consumer is written next to each node. Applying the pairwise design rule results in a network structure where the total effort exerted T is 1.86, an 11% decrease compared to the empty network where each individuals is disconnected from anyone else (in which case $T = \sum_i \alpha_i = 2.1$).

Now consider the network below on the right where the total activity is 2.32, a 10% increase compared with the empty network. In this network, the only tie is between consumers 1 and 3. Connecting these two results in a net increase in their effort from 0.5 to 0.71. The network joint on the right is actually the "optimal" design in the sense that there is no other network configuration that can lead to more than 10% increase in the total activity.

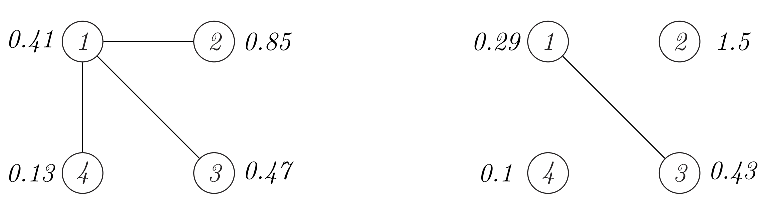

This example illustrates that designing a network based on direct pairwise influence alone can be not only suboptimal but even detrimental to population-level effort. Optimal network design requires going beyond the immediate returns to adding a link between two individuals and taking into account the spillovers to other connected nodes.

Designing a network without any simplifications implies searching for the best design, and this is not a trivial task. For n consumers to select from, there are $2^{\frac{n(n-1)}{2}}$ possible configurations.[3] Moreover, the problem is nonlinear, as evidenced by (6). To tackle this difficult problem, we temporarily relax the discreteness of tie creation. Specifically, we allow the firm to not only create a connection between two consumers but also to do so at varying levels of tie strength. For instance, a social network may change the frequency at which consumers hear from each other. Although this relaxation enlarges the set of possible configurations, it makes the mathematical analysis easier. Let $s_{ij} = s_{ji} \in [0, 1]$ be the strength of the link between i and j. When $s_{ij} = 0$, this implies that the two consumers do not hear from each other, or, equivalently, that there is no connection between these consumers. The equilibrium activity of a consumer then becomes a variant of that in Equation (4), weighted for tie strength:

$$y_i^* = \frac{\alpha_i + \sum_{j=1}^{n} s_{ij}\rho_{ij}y_j}{1 + \sum_{j=1}^{n} s_{ij}\rho_{ij}} \tag{7}$$

The Nash result in Equation (6) can be similarly expressed as $\mathbf{L}^{-1}\boldsymbol{\alpha}$ and using the following influence matrix $\mathbf{L}$:

$$\begin{aligned} L_{ij} &= -\rho_{ij}s_{ij}, \\ L_{ii} &= 1 - \sum_{j\sim i} L_{ij} \end{aligned}$$

We can now assess how changing the strength of the tie between two consumers affects the overall effort level in the network.

Lemma 1. *The marginal effect of the strength of the tie between individuals i and j on the activity of each network member, $\partial\mathbf{y}^*/\partial s_{ij}$, is given by:*

$$\frac{\partial \mathbf{y}^*}{\partial s_{ij}} = \mathbf{L}^{-1}\mathbf{v}_{ij}^*, \tag{8}$$

where the vector $\mathbf{v}_{ij}^*$ is defined as below. We use $\mathbf{e}_i$ to denote the i-th basis vector and $\mathbf{e}_j$ to denote the j-th basis vector (In other words, $\mathbf{e}_i$ is all zeros except for the i-th entry).

$$\mathbf{v}_{ij}^* \equiv \rho_{ij}(y_j^* - y_i^*)\mathbf{e}_i + \rho_{ji}(y_i^* - y_j^*)\mathbf{e}_j.$$

[3] This number is 64 for $n = 4$, larger than 32,000 for $n = 6$, and is considerably larger for any greater sized population (e.g., $n > 100$).

The change in the overall community activity in response to a change in a specific tie strength equals $\partial T/\partial s_{ij} = \sum_{k=1}^{n} \partial y_k^*/\partial s_{ij}$. Lemma 1 says that the marginal effect of increasing s_{ij} on the collective effort T is positive iff

$$\partial T/\partial s_{ij} = \mathbf{1}'\mathbf{L}^{-1}\mathbf{v}_{ij}^* > 0 \tag{9}$$

Linking two consumers exhibiting conformity raises the effort of the lower-activity consumer, and lowers that of the higher-activity consumer. Optimal network design consists of finding the best trade-off between these two opposing effects, operating not just directly on the two focal individuals but who are being connected also indirectly on others in the network who in turn affect the focal two. The entirety of these effects is taken into account and summarized in Equation (9).

Lemma 2. *Hold constant all the tie weights in a network except for s_{ij}. If $\partial T/\partial s_{ij} > 0$ at any particular value of s_{ij}, then $\partial T/\partial s_{ij} > 0$ at all values of $s_{ij} > 0$.*

So if the marginal effect of s_{ij} on T is positive at some point, it will remain positive as one increases (or decreases) s_{ij}. This remarkable result implies that creating medium-strength ties within is generally not optimal. If strengthening the connection between two consumers in a community increases the total activity (T) in the network, then the tie should be pushed to full strength $s_{ij} = 1$. Conversely, if weakening a tie between two individuals increases the activity T, the link should be removed completely. Put differently, the optimal tie strength is "bang-bang" and to maximize total activity, the network designer must add ties at the extensive rather than intensive margin: it is the presence or absence of a tie rather than its strength which helps to maximize total activity. With Lemma 2, we can lay out the main proposition about the value of inserting or removing a tie on a discrete network.

Proposition 2. ***(Value of a Tie to the Population)*** *Consider any given unweighted network for a set of individuals. Let $\mathbf{L}$ be as defined in Proposition 1 and $\mathbf{v}_{ij}^*$ be as defined in Lemma 1 under the current equilibrium. If consumers i and j are not yet connected, connecting them will increase total community activity iff*

$$\mathbf{1}'\mathbf{L}^{-1}\mathbf{v}_{ij}^* > 0. \tag{10}$$

If i and j are already connected, removing the tie between them increases the collective community activity iff the reverse holds:

$$\mathbf{1}'\mathbf{L}^{-1}\mathbf{v}_{ij}^* < 0. \tag{11}$$

Proposition 2 helps to compute the sign of the value of a single link in terms of its contribution to the collective activity level. It significantly simplifies the task of checking through all the pairs of individuals to determine connecting (or disconnecting) which pairs will improve T. In particular, this proposition avoids the need for computing any what-if equilibrium of adding or removing a tie. In the Online Appendix, we present and assess a computational algorithm for optimal network based on this proposition. The algorithm computes the effect of changing each tie, and changes ties in a sequential manner. In each step, the algorithm adds or removes ties so that the total activity T is increased, and it ends when further addition or removal of any single tie does not further increase total activity. This algorithm (which we call the Fast Multiple Link Altering (FMLA), performs remarkably well in terms of both speed and precision against several benchmarks. In the subsequent analysis, we use various numerical examples, obtained by applying this algorithm, to illustrate some of the theoretical results.

4.1 When is Network Design Most Beneficial?

We now present results on how the total effort exerted in an optimized network depends on (i) the members' propensities to influence, be influenced, and exert stand-alone effort, (ii) the heterogeneity across the members, and (iii) the size of the community. These results help network designers to understand under what circumstances they can benefit the most from purposively configuring the network structure.

Consider two platforms with a consumer base identical in size (n) and their isolated activity (α). If the influence potential (ρ) among the consumers of the first platform is greater than that of the second, the maximum level of activity that can be attained by the second platform can never exceed that of the first. The following proposition formalizes this, as a consequence of Lemma (2).

Proposition 3. ***(Strength of Conformity and Outcomes in an Optimized Network Structure)*** *Consider the consumers in two networks, A and B, characterized with* $A = (\boldsymbol{\alpha}, \boldsymbol{\rho}) = \{\alpha_i, \{\rho_{ij}\}_{j \neq i}\}_{i=1}^n$ *and* $B = (\boldsymbol{\alpha}, \boldsymbol{\rho}') = \{\alpha_i, \{\rho'_{ij}\}_{j \neq i}\}_{i=1}^n$*. Suppose* $\rho'_{ij}/\rho_{ij} = \rho'_{ji}/\rho_{ji} \geq 1$ *for all* (i, j) *pairs. Then, with each network being given its optimal configuration, the total effort level of the consumers in B is greater than or equal to that in A.*

This proposition means that stronger conformity provides greater opportunities for optimal design to boost the total level of effort exerted. The network designer can always find a structure that leverages greater conformity, and consequently can always find a design in which the positive motivating effect of conforming to others' effort weakly dominates the negative demotivating effect of conformity.

The following proposition quantifies the gain from placing consumers in a well-designed network structure as opposed to keeping them in isolation.

Proposition 4. ***(Population Heterogeneity and Gain from Communities)*** *Consider a con-*

sumer base $\{\alpha_i, \{\rho_{ij}\}_{j\neq i}\}_{i=1}^n$. Regardless of the network configuration,

$$\frac{1}{n}\left(T - \sum_{i=1}^{n}\alpha_i\right) \leq \frac{n-1}{2}\rho^+\alpha^+$$

where $\rho^+ \equiv \max\{|\rho_{ij} - \rho_{ji}|, i \neq j\}$ and $\alpha^+ \equiv \max\{|\alpha_i - \alpha_j|, i \neq j\}$ are maximum differences across all (i,j) pairs. In words, the upper bound on the improvement of an average consumer depends on the heterogeneity among consumers in their propensities to influence, be influenced, and exert effort stand-alone. The greater the heterogeneity, the higher the upper bound.

The proposition implies that, for networking to boost total effort, the influence and susceptibility to influence cannot be both equal for everyone, i.e., either χ_i or λ_i, or both, must vary across individuals. The reason is that ρ^+, and hence the gain from networking, is zero when each network member has the same influence χ and the same susceptibility to influence λ.

The proposition also implies that networks and sub-communities showing little variation in α_i provide lower potential for improvement through peer influence than heterogeneous networks and sub-communities. Hence, network structures formed organically where people tend to connect with others with a similar α (tie formation exhibiting homophily in α) provide little benefit from networking. This is an important insight, and emphasizes the importance of purposive network design or "central planning" for effective communities.

The proposition also has important implications for customer targeting. Most large firms prefer to segment consumers into externally heterogeneous and internally homogeneous groups, and then target each segment with possibly tailored offerings. Altogether, identifying and targeting internally homogeneous segments of customers is good practice. In the setting we study, in contrast, maintaining or even seeking heterogeneity among one's customers or members is beneficial and a firm or service provider with a homogeneous customer or membership base will see smaller gains from networking, as will its average customer or member.

Next, we study how the size of the customer or membership base affects the gains from network design, taking into account the heterogeneity of the base. To this end, we consider a network with n consumers whose characteristics $(\alpha_i, \lambda_i, \chi_i)$ are heterogeneous and drawn from a non-degenerate distribution Φ. For any set of n consumers, there is an optimal network that leads to an improvement in the community activity, $\sum_i(y_i^* - \alpha_i)$. We denote by $Q(\Phi, n)$ the expectation of this improvement over the distribution Φ, and by $q(\Phi, n) \equiv Q(\Phi, n)/n$ the expected per capita improvement.

Proposition 5. ***(Network Size and Gain from Network Optimization)*** *The expected per capita improvement in effort exerted from optimizing the network design satisfies the following in-*

equality:[4]

$$q(\Phi, n+1) > \frac{\sum_{k=1}^{n} q(\Phi, k)}{n} \quad , \; n = 1, 2, ...$$

The proposition provides a lower bound for the expected per capita improvement from optimally arranged communities. It states that the per capita improvement in a population of size $n+1$, i.e., $q(\Phi, n+1)$, exceeds the average of the per capita improvements of smaller networks or communities: $q(\Phi, n)$, $q(\Phi, n-1), ...,$ and $q(\Phi, 1)$. Therefore, working with more customers or members will generate greater improvements in the average level of effort exerted after network optimization.

Figure (1) illustrates the analytical results of Propositions 4 and 5 with a numerical example. It plots the expected average benefit $q(\Phi, n)$ as a function of n for two different heterogeneity distributions, where one has greater variations in α_i and λ_i than the other.[5] Clearly, the per capita improvement from placing consumers in optimized network structures increases with the number of members (Proposition 5) and the heterogeneity in the stand-alone effort α_i and the influenceability λ_i of these members (Proposition 4).

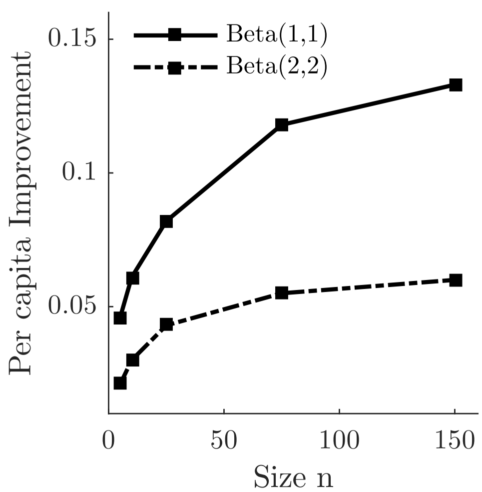

Size n	Per capita Improvement $q(\Phi, n)$	
	Beta (1,1)	Beta (2,2)
5	.046	.021
10	.061	.030
25	.082	.043
75	.118	.055
150	.133	.060

Notes: Each value for $q(\Phi, n)$ is evaluated by averaging across 100 communities. For each community, members are drawn from Φ and the optimal network is computed. Φ is specified as follows: for the first column, $(\alpha_i, \lambda_i) \sim \text{Beta}(1,1) \times \text{Beta}(1,1)$ and $\chi_i = 1$; for the second column, $(\alpha_i, \lambda_i) \sim \text{Beta}(2,2) \times \text{Beta}(2,2)$ and $\chi_i = 1$.

Figure 1: Expected per Capita Improvement $q(\Phi, n)$

The intuition for these analytic and computational results is straightforward. Larger and more heterogeneous sets of nodes provide more opportunities for the network designer to identify and

[4]In fact, a stronger result also holds: $q(\Phi, n+1) > \frac{\sum_{k=1}^{n} k \cdot q(\Phi,k)}{\sum_{k=1}^{n} k}$. The right hand side of this inequality is a weighted average and assigns more weight to the larger communities.

[5]In the Figure we assume that $(\alpha_i, \lambda_i) \sim \text{Beta}(2,2) \times \text{Beta}(2,2)$ and $\text{Beta}(1,1) \times \text{Beta}(1,1)$. Note that $\text{Beta}(2,2)$ is a symmetric bell shaped distribution over the unit interval centered around 0.5, appropriate for populations where there is a greater mass of individuals carrying close to median characteristics in the importance they give to their goals and the degree of susceptibility they are subject to. In contrast, $\text{Beta}(1,1)$ is the uniform distribution, suggesting a more homogeneous spread of individual types.

select a structure where, on balance, the positive motivating effect of conformity on low α_i nodes outweighs its negative demotivating effect on high α_i nodes.

One might further intuit that, to have the positive effect of conformity outweigh its negative effect, an optimal network tends to favor ties that involve customers that have (i) a high (low) stand-alone effort level combined with high (low) social influence, or (ii) a high (low) stand-alone effort level combined with low (high) social influenceability. The numerical example at the beginning of Section 4 illustrates the latter. More generally, one might intuit that, in an optimized network with a fixed number of customers n, the number of ties increases with the correlation between α and χ, and decreases with correlation between α and λ. The same would be true for the network density, defined as the actual number of ties divided by maximum number of possible ties $(n(n-1)/2)$. The numerical results reported in Figure 2 support that intuition: Optimal networks indeed become denser as the correlation between α_i and λ_i goes from 1 to -1.[6]

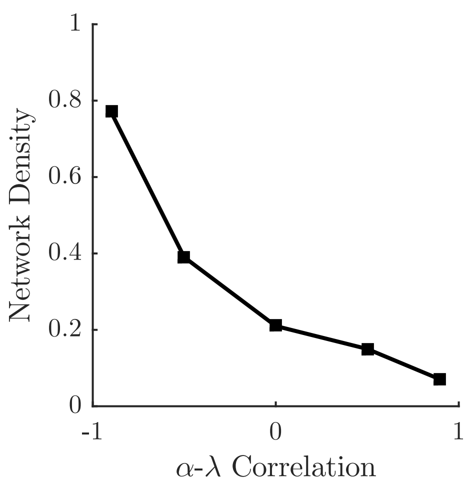

Correlation	Network Density
-0.9	0.77
-0.5	0.39
0	0.21
0.5	0.15
0.9	0.07

Notes: Each value of network density is evaluated by averaging across 100 communities. For each community, members are drawn from Φ and an optimal network is computed. Community size n is fixed at 25. The Φ is specified as follows: (α_i, λ_i) follows a bi-variate Beta(2,2) distribution with the correlation introduced using Plackett's method; $\chi_i = 1$ for all i.

Figure 2: Density of Optimally Structured networks for Different $\alpha - \lambda$ Correlations

4.2 What is the Optimal Network Design for Inducing Effort?

We now turn to the key question: How to structure the network or community of n individuals to maximize the total or average level of effort exerted? We first introduce a result for the most dense networks, i.e., complete networks where every two individuals are connected.

Lemma 3. *(i) For a complete network of individuals characterized by* $\{\alpha_i, \lambda_i\}_{i=1}^n$ *and* $\chi = 1$*, the total effort is given by:*

$$T = \frac{\sum_{i=1}^n \frac{\alpha_i}{1+n\lambda_i}}{\frac{1}{n}\sum_{i=1}^n \frac{1}{1+n\lambda_i}}. \tag{12}$$

[6]Correlation is introduced by a method used by Plackett (1965).

(ii) For a complete network of individuals characterized by $\{\alpha_i, \chi_i\}_{i=1}^n$ *and* $\lambda = 1$*, the total effort is given by:*

$$T = \frac{\sum_{i=1}^n \alpha_i \left(1 + n\chi_i\right)}{\frac{1}{n}\sum_{i=1}^n \left(1 + n\chi_i\right)}. \tag{13}$$

Equations (12) and (13) imply that the average effort level within a complete network, $T\,/\,n$, is a weighted average of the stand-alone effort levels (α_i), where individuals who are less influenceable (λ_i) or more influential (χ_i) are given a greater weight.

The last key building block to our main result is stated in Lemmas 4 and 5, characterizing the optimal network structure for the extreme case where the propensity to exert stand-alone effort (α_i) and the propensity to influence (χ_i) or be influenced (λ_i) are perfectly correlated.

Lemma 4. *Consider a set of individuals represented by* $\{(\alpha_i, \lambda_i), i = 1, ..., n\}$ *and* $\chi_i = 1$*. If* α_i *and* λ_i *are perfectly negatively correlated (positively correlated), then the complete (empty) network is locally optimal and the empty (complete) network is not optimal.*

Lemma 5. *Consider a set of individuals represented by* $\{(\alpha_i, \chi_i), i = 1, ..., n\}$ *and* $\lambda_i = 1$*. If* α_i *and* χ_i *are perfectly positively correlated (negatively correlated), then the complete (empty) network is locally optimal and the empty (complete) network is not optimal.*

Of course, perfect correlations are unlikely to occur within the customer or membership base of any company or organization. However, network designers may be able to organize the network members into multiple sub-communities or sub-graphs within which such correlations hold. Lemmas 4 and 4 state that these sub-communities should be very densely connected. Moreover, Proposition 5 suggests that members of such sub-communities should have high variance in stand-alone effort, influence, and influenceability. We combine and formalize these insights in Propositions 6 and 7. Note, these two propositions pertain to local rather than global optima. The reason is that we focus on the equilibrium network configuration where a manager does not have an incentive to remove or add any single tie. It is possible that there are multiple local optima, but numerical analyses presented below show that the proposed characteristics of optimal sub-communities also hold when searching for the globally optimal network design.

Proposition 6. ***(Membership to Sub-communities in the Optimal Network)*** *Consider* $n \to \infty$ *consumers who are distributed according to a density* $\Phi(\alpha, \lambda)$ *over the support* $[\underline{\lambda}, \overline{\lambda}] \times [\underline{\alpha}, \overline{\alpha}]$*. Under a mild regularity condition on* Φ*, it is a locally optimal network to organize consumers into sub-communities where any two consumers are linked iff they belong to the same sub-community. Each sub-community consists of consumers located along a line segment parallel to the negatively-sloped diagonal of the support as indicated by the solid line in Figure 3.*

Proposition 7. ***(Membership to Sub-communities in the Optimal Network)*** *Consider* $n \to \infty$ *consumers who are distributed according to a density* $\Phi(\alpha, \chi)$ *over the support* $[\underline{\chi}, \overline{\chi}] \times [\underline{\alpha}, \overline{\alpha}]$. *Under a mild regularity condition on* Φ, *it is a locally optimal network to organize consumers into sub-communities where any two consumers are linked iff they belong to the same sub-community. Each sub-community consists of consumers located along a line segment parallel to the positively-sloped diagonal of the support as indicated by the solid line in Figure 4.*

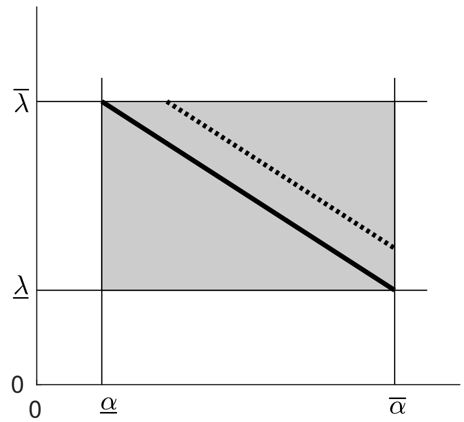

Figure 3: Individuals Located Along a Line on the α-λ Plane

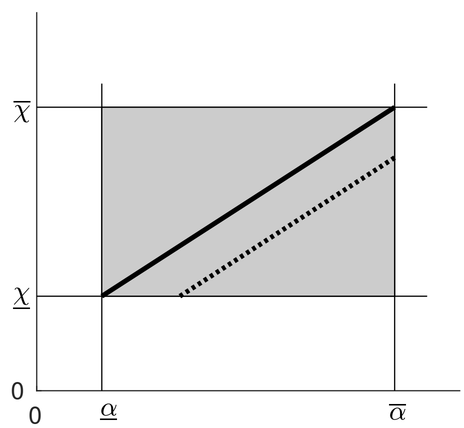

Figure 4: Individuals Located Along a Line on the α-χ Plane

Proposition 6 and 7 state that in the optimal design for a very large network, the network members will be sorted into a set of mutually disconnected sub-communities (i.e., graph components) where each sub-community is internally fully connected (i.e., the link density is 1 within a sub-community). Put differently, the complete or near-complete networks described in Lemma 4 still arise, but at the level of the smaller sub-communities. This implies that the network exhibits perfect transitivity (if $i \sim j$ and $j \sim k$, then $i \sim k$) and clustering.

As noted earlier, the key challenge in optimal network design under conformity is how to balance the two opposite effects of conformity. The effort of low-α types gets pulled up towards the norm but that of high-α types tends to get pulled down. Proposition 6 states that this trade-off can be handled by combining influential high-α types and non-influential low-α types into heterogeneous sub-communities (close to the main diagonal in Figure 3). Similarly, Proposition 7 states that this trade-off can be handled by combining non-susceptible high-α types and susceptible low-α types into heterogeneous sub-communities (close to the main diagonal in Figure 4). While Propositions 6 and 7 are proven using the asymptotic properties (under $n \to \infty$) to make the derivations tractable, their insights also hold for finite n, as we illustrate numerically.

To investigate and visually illustrate the typical optimal network structure in moderately sized networks, in Figures 5 through 9, we numerically investigate five scenarios in which we fix the network size to $n = 200$ and vary the relations between α_i, λ_i, and χ_i among the full set of 200 network members.

In the first scenario, we fix χ_i (ability to influence) to 1 and let α_i and λ_i vary independently of each other. In the second scenario, we fix λ_i (susceptibility to influence) to 1 and let α_i and χ_i vary independently of each other. In the third scenario, we vary all three parameters independently from each other. In the fourth scenario, we again vary all three parameters, but induce a negative correlation between χ_i and λ_i. Several theories imply such a negative correlation between influence and susceptibility to influence (e.g., Van den Bulte and Joshi, 2007), but empirical evidence typically shows little or only moderate negative correlation (e.g., Aral and Walker, 2012; Iyengar et al., 2011). Consequently, we deem the correlation structures of both scenarios 3 and 4 to be fairly realistic. Finally, in the fifth scenario, susceptibility to influence (λ_i) exhibits an inverse-U relation with respect to ability to influence (χ_i). This is consistent with the theory and evidence of middle-status conformity (e.g., Hu and Van den Bulte, 2014; Iyengar et al., 2015; Phillips and Zuckerman, 2001). In each scenario, we first identify the optimal network structure using the Fast Multiple Link Altering (FMLA) algorithm described in the Online Appendix, and identify the various sub-communities in the optimized network using the community detection algorithm proposed by Newman (2006). We report the optimized network graph using different colors to indicate different sub-communities.

As shown in Figures 5 through 9, the optimal network structure is remarkably similar across these five scenarios. As the graph on the left of each Figure shows, the optimal network always consists of a small number of sub-communities that are highly densely interconnected within but not, or only very weakly, connected across. As the plot on the top right of each Figure shows, the members of each sub-community fall on a diagonal band in the space defined by α_i on the horizontal axis and, depending on the parameters manipulated in the scenario, λ_i, χ_i, or $\lambda_i - \chi_i$ on the vertical axis.

The amount of effort exerted in these numerical illustrations exhibits a fourfold pattern consistent with the analytical results and their intuition. The first three patterns are shown in the middle plots on the right of Figures 5 through 9, whereas the first is shown in the bottom plots. The first pattern is that sub-communities constructed by combining people in the top right (as opposed to bottom left) corner in the top right plots in Figures 5, 7 and 8 exhibit higher effort, consistent with Figure 3. Similarly, sub-communities constructed by combining people in the bottom right (as opposed to top left) corner in the top right plots in Figures 6 and 9 exhibit higher effort, consistent with Figure 4.

The second pattern is that, in the optimal network, the effort expended increases with baseline effort α_i and the susceptibility to social influence λ_i, and decreases with social influence χ_i. The latter occurs because influential high-α types tend to be connected with susceptible low-α types, such that the influentials tend to be "pulled down" for the common good and the susceptibles tend to be "pulled up."

The third pattern is that not only the average but also the variance in effort exerted varies across sub-communities. That amount of variance is reflected in the range in y_i^* among members of each community in the middle plots of Figures 5-9. People located further away from the relevant diagonal in the parameter space tend to be organized into smaller sub-communities than people closer to the diagonal, and their sub-communities exhibit greater variance in effort than do sub-communities consisting of people located close to the diagonal. The reason is that people on the diagonal in Figures 3 and 4 (influential high-α and susceptible low-α types) tend to be put together into sub-communities counting both influential and influenceables, such that the central pull towards the sub-communities norm is strong and the variation in y_i^* is much smaller than α_i. People far away from the main diagonals (susceptible high-α and influential low-α types), in contrast, are organized in communities with people like themselves, resulting in little pull towards the norm such that the variation in y_i^* in these sub-communities tends to be almost as large as that in α_i.

The fourth pattern is that the network members who tend to exert more effort in the optimal network than as social isolates. These people for whom $y_i^* > \alpha_i$ are identified as solid nodes in the plot on the lower right of Figures 5-9 tend to be low-α types who are either highly influenceable or not very influential. In contrast, high-α types who are very influential or not very susceptible to influence tend to be "sacrificed for the greater good" and end up exerting less effort in the optimal network than as social isolates (identified as hollow nodes in the plot in the lower right of Figures 5-9). There are always more members with $y_i^* > \alpha_i$ than those with $y_i^* < \alpha_i$.

Overall, these numerical results for networks with $n = 200$ and various constellations of α_i, λ_i, and χ_i parameters exhibit optimal network structures very similar to those identified in Propositions 6 and 7 for $n \to \infty$. They do so even though the algorithm used in the numerical analyses relies only on Proposition 2. Hence, one can view these numerical results as a corroboration of Propositions 6 and 7 for finite networks.

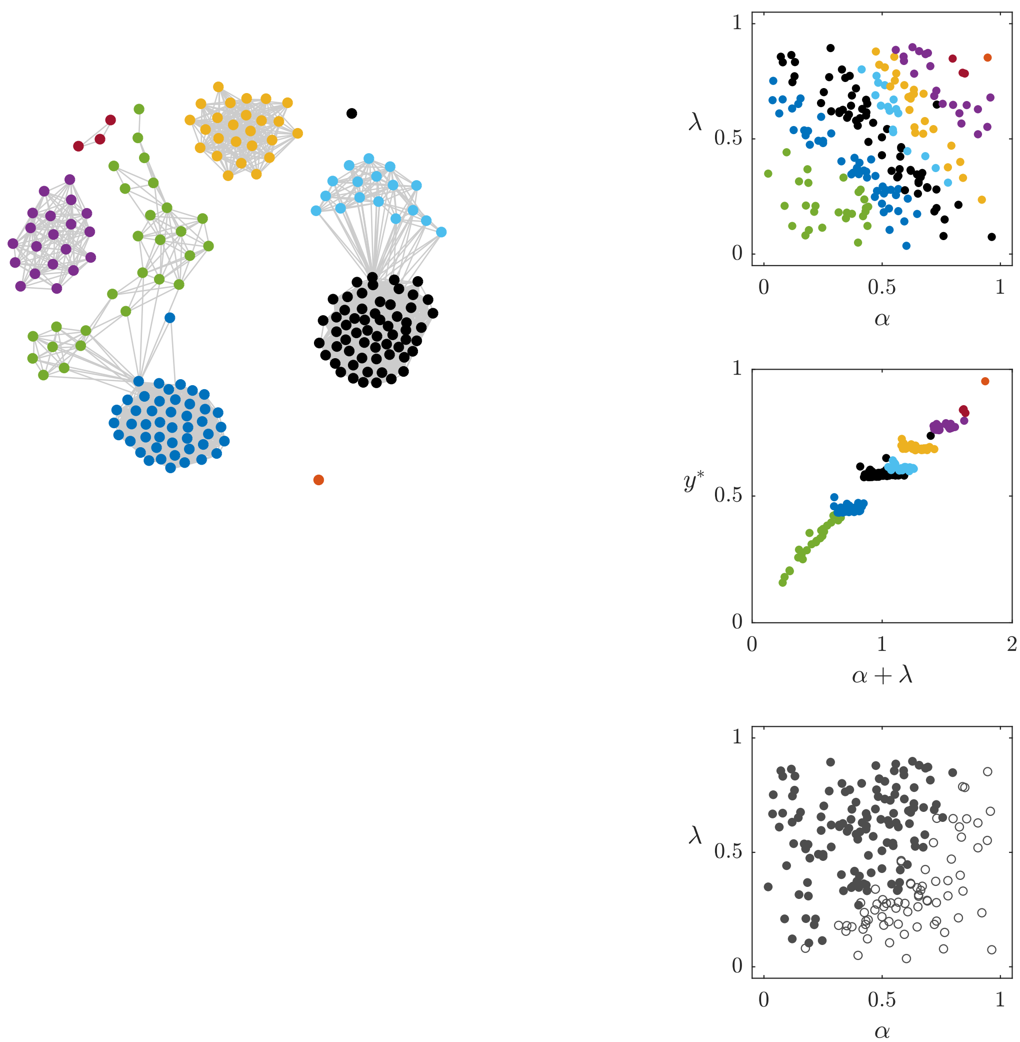

The plots characterize the optimal network for 200 individuals, where (α_i, λ_i) is drawn from Beta(2, 2)×Beta(2, 2); $\chi_i = 1$. The plot on the left side visualizes the optimal network. We use different colors to separate different subcommunities. The plots on the right show, from top to bottom: (i) the scatter plot of λ_i against α_i, (ii) the individual activity levels y_i^* in the network against $\alpha_i + \lambda_i$, and (iii) the scatter plot of λ_i against α_i with solid nodes indicating individuals with $y_i^* > \alpha_i$ and hollow nodes indicating otherwise. The percentage of individuals with $y_i^* > \alpha_i$ is 64%.

Figure 5: An Optimal Network Among 200 Individuals with Varying α_i and λ_i (Uncorrelated)

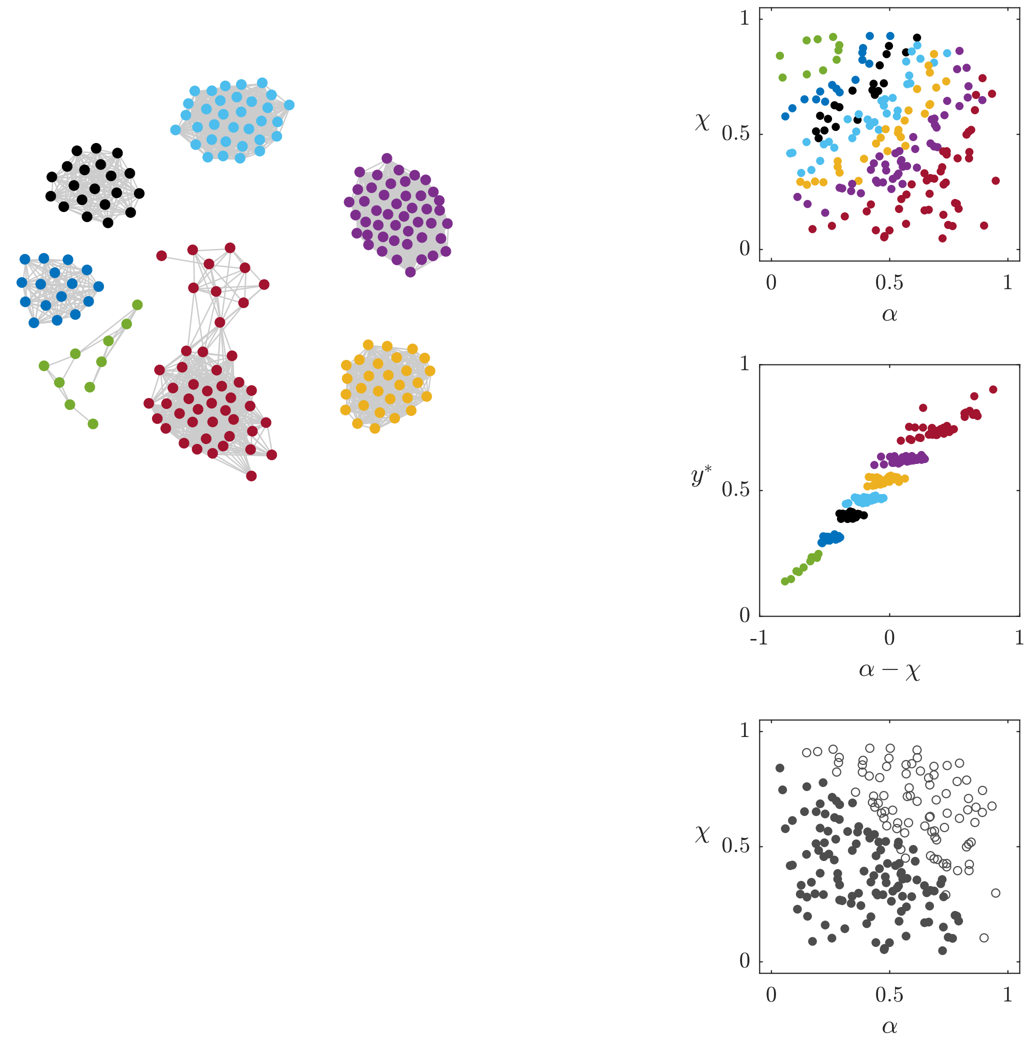

The plots characterize the optimal network for 200 individuals, where (α_i, χ_i) is drawn from Beta(2, 2)×Beta(2, 2); $\lambda_i = 1$. The plot on the left side visualizes the optimal network. We use different colors to separate different sub-communities. The plots on the right show, from top to bottom: (i) the scatter plot of χ_i against α_i, (ii) the individual activity levels y_i^* in the network against $\alpha_i - \chi_i$, and (iii) the scatter plot of χ_i against α_i with solid nodes indicating individuals with $y_i^* > \alpha_i$ and hollow nodes indicating otherwise. The percentage of individuals with $y_i^* > \alpha_i$ is 58%.

Figure 6: An Optimal Network Among 200 Individuals Varying in α_i and χ_i (Uncorrelated)

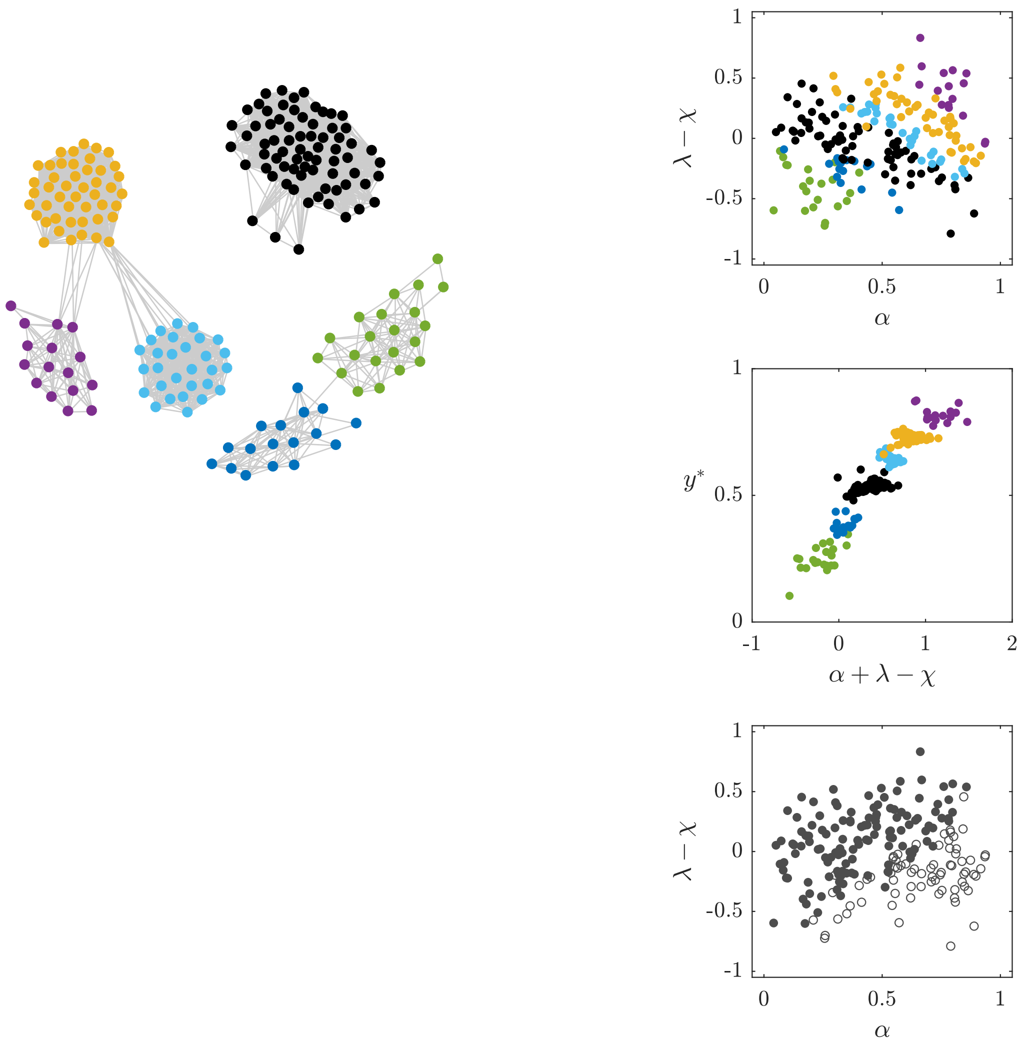

The plots characterize the optimal network for 200 individuals, where α_i, λ_i and χ_i are drawn independently from Beta(2, 2). The plot on the left side visualizes the optimal network. We use different colors to separate different sub-communities. The plots on the right show, from top to bottom: (i) the scatter plot of $\lambda_i - \chi_i$ against α_i, (ii) the individual activity levels y_i^* in the network against $\alpha_i + \lambda_i - \chi_i$, and (iii) the scatter plot of $\lambda_i - \chi_i$ against α_i with solid nodes indicating individuals with $y_i^* > \alpha_i$ and hollow nodes indicating otherwise. The percentage of individuals with $y_i^* > \alpha_i$ is 64.5%.

Figure 7: An Optimal Network among 200 Individuals Varying in α_i, χ_i and λ_i (Uncorrelated)

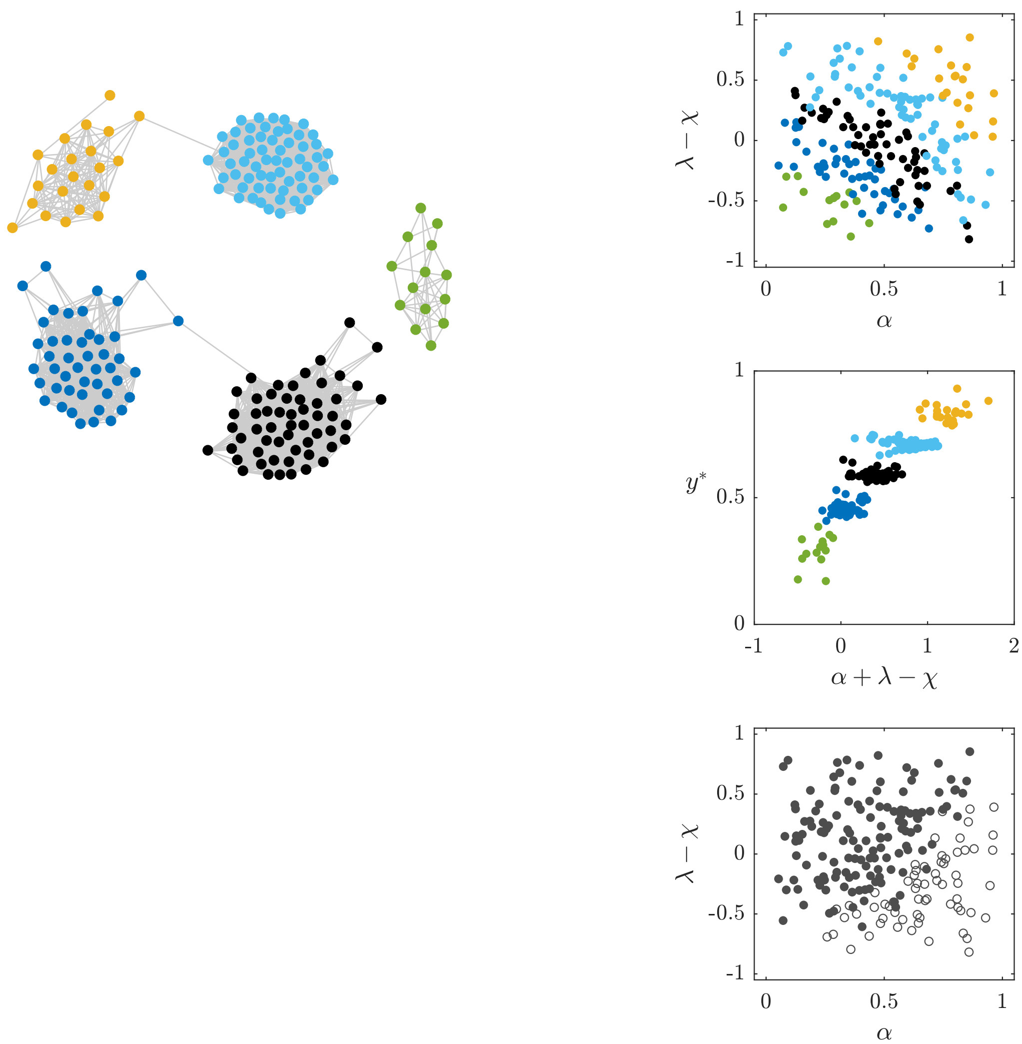

The plots characterize the optimal network for 200 individuals, where α_i, λ_i, and χ_i are all drawn from Beta(2, 2). A negative correlation of -0.6 is introduced between λ_i and χ_i using Plackett's method. The plot on the left side visualizes the optimal network. We use different colors to separate different sub-communities. The plots on the right show, from top to bottom: (i) the scatter plot of $\lambda_i - \chi_i$ against α_i, (ii) the individual activity levels y_i^* in the network against $\alpha_i + \lambda_i - \chi_i$, and (iii) the scatter plot of $\lambda_i - \chi_i$ against α_i with solid nodes indicating individuals with $y_i^* > \alpha_i$ and hollow nodes indicating otherwise. The percentage of individuals with $y_i^* > \alpha_i$ is 68%.

Figure 8: An Optimal Network among 200 Individuals Varying in α_i, χ_i and λ_i, with the Latter Two Negatively Correlated

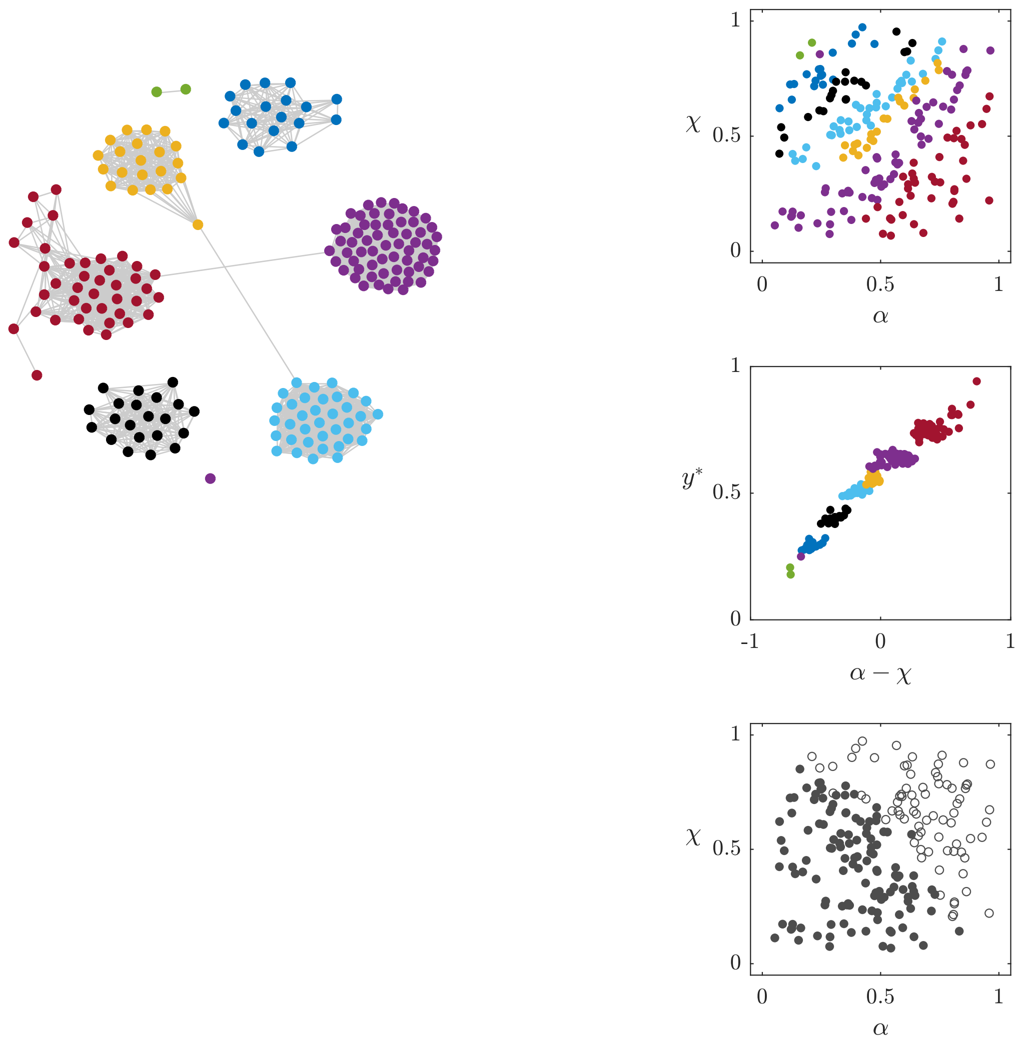

The plots characterize the optimal network for 200 individuals, where α_i and χ_i are drawn independently from Beta(2, 2) and $\lambda_i = 0.75 - 2(\chi_i - 0.5)^2$. The plot on the left side visualizes the optimal network. We use different colors to separate different sub-communities. The plots on the right show, from top to bottom: (i) the scatter plot of χ_i against α_i, (ii) the individual activity levels y_i^* in the network against $\alpha_i - \chi_i$, and (iii) the scatter plot of χ_i against α_i with solid nodes indicating individuals with $y_i^* > \alpha_i$ and hollow nodes indicating otherwise. The percentage of individuals with $y_i^* > \alpha_i$ is 62%.

Figure 9: An Optimal Network among 200 Individuals Varying in α_i, χ_i and λ_i, with the Latter Two Exhibiting Middle-status Conformity

5 Alternative Models

5.1 Activity Minimizing Networks

In some contexts, activity can be harmful (e.g., smoking, gambling, addictive gaming) and a network designer may want to minimize rather maximize it. It is easy to show that an activity-minimizing network for $(\boldsymbol{\alpha}, \boldsymbol{\rho})$ is, at the same time, an activity-maximizing network for $(-\boldsymbol{\alpha}, \boldsymbol{\rho})$. So it is straightforward to transform a minimization problem into a maximization problem, and the results presented in Section 4 apply as long as the negative value of the individual activity $(-\boldsymbol{\alpha})$ is used in the calculations and network design .

5.2 Social Competition

The model and analysis assume that peer influence operates through social conformity, specifically the desire to adhere to descriptive social norms. It is possible, however, that social influence operates through other mechanisms. In this section, we consider social competition or the desire to outperform others, which may be relevant at least for some people for intense physical exercise (Aral and Nicolaides, 2017) as opposed to managing one's weight and fitness through mindful eating and moderate exercise, reducing one's electricity consumption, reducing one's tobacco or drug use, making charitable donations, and studying. To account for peer influence through social competition rather than conformity, we modify the model setup given in equations (2) and (4) into:

$$u_i = \alpha_i y_i - \frac{1}{2} y_i^2 + \sum_{j \sim i} \left[\delta_{ij} (y_i - y_j) \right], \tag{14}$$

where $\delta_{ij} > 0$ measures how much individual i cares about doing better than individual j. For instance, he may want to exercise harder than his peers, lose more body weight than his buddies, or outperform other salespeople at work. He may care very little about being better than some of his peers and may care greatly about exceeding the level of effort or activity of some others. So, relationships are heterogeneous and $\delta_{ij} \neq \delta_{ik}$. The implied effort chosen by individual i is:

$$y_i = \alpha_i + \sum_{j \sim i} \delta_{ij}. \tag{15}$$

Clearly, when peer influence operates through social competition, effort is always higher in-network than out-of-network effort, social information should always be shared, and the completely connected network is the optimal structure. Social competition does not present the same design challenges and trade-offs as social conformity does.

5.3 Uni-directional ties

Our analysis assumed that influence flows both ways over social ties, even though the strength of influence need not be symmetric. What happens when the community manager has the ability to make the influence over ties flow in only one direction, e.g., by informing i about j's effort or actions, but not vice versa? The community manager will want to connect every higher-α type to every lower-α type. Since the downward pull of lower-α types can be avoided, there is no reason to have separate sub-communities, and the optimal network structure will be a single giant, fully connected graph component. The ability to make ties uni-directional renders the optimal network structure very easy to identify.

5.4 Rewiring of ties

Consistent with the empirical evidence reviewed in Section 2, we assumed that community members accept their assigned location in the network designed to help them achieve a specific goal. How critical is that assumption? Specifically, how different would the optimal network design be when, in a first stage, the network designer imposes a structure and people join and, in a second stage, community members experiencing tension from non-conforming peers are able to rewire their connections at a cost reflecting mere inertia and the hassle cost of identifying other potential partners?

One assumption is that in the second stage people seek to achieve conformity by connecting to those with a first-stage effort level very similar to theirs. Since in the current optimal design, within-community variance in y^* is small, and especially so in the large communities on the main diagonals, there will be only limited rewiring.

An alternative assumption is that people will seek to form connections with others having a parameter α similar to theirs. We expect more rewiring in this second scenario, but still expect that the optimal design will not be radically different from the one we have identified. In the current design rules, high-α types with low influence are already paired with high-α types, and low-α types with low susceptibility are already paired with low-α types. So, the communities consisting of people located towards the off-diagonal corners are already fairly homogeneous in α. The key question is how much the heterogeneous sub-communities on the main diagonals change once rewiring starts. That is likely to be a function of the cost of rewiring relative to the tension induced by the difference in α between two connected community members. Hence, rewiring does not drive the total benefit from network optimization to zero. Rather, the total gets diluted and the amount of dilution is a function of the relative cost of rewiring, the time lapsed until rewiring takes place, and the network designer's discount factor. Most importantly, we see no reason to expect the optimal design rules in this two-stage set-up to differ markedly from those we have identified.

6 Conclusion

In the absence of a clear benchmark about what is appropriate or desirable, people tend to conform to what others do, even strangers. Conformity to descriptive social norms has a sizable effect on the effort people exert in areas like energy conservation, health and charitable help, and is increasingly being leveraged by companies and platforms like Opower, Crowdrise and BuddyUp. Motivated by these important developments, and the broader emerging interest in network design, we address the question of how to design the network of ties among community members to maximize the total effort exerted. The key trade-off to be managed is that social conformity pushes up the effort of those who would otherwise exert relatively little effort while pushing down the effort of those who would otherwise exert relatively greater effort.

6.1 Recommendations

Our analysis results in three specific recommendations. First, the community should be organized as a set of sub-communities, each of which is fully connected within itself but disconnected from other sub-communities. Second, each sub-community should consists of members selected such that their susceptibility to influence (propensity to influence) and their "standalone" propensity to exert effort are negatively (positively) correlated. Third, there is no benefit in varying the strength of ties across dyads, e.g., by varying the visual salience across ties. Our analysis also shows that the benefits of following these design rules are greater in communities that count more members or exhibit greater variation in their members' social influence, susceptibility to social influence, and propensity to exert effort in isolation. The reason is that greater size and greater variance provides the network designer greater degrees of freedom when managing the key trade-off.

6.2 Implementation

One appeal of our recommendations is that the structure is easy to implement not only online but also in face-to-face networks. Organizing the network into sub-communities that are fully connected within but disconnected across amounts to creating distinct groups where every member is connected to every other member of their own group and to no one else. Not having to vary the visual salience, interaction frequency, or other facets of tie strength across dyads also makes the optimal network design easy to implement.

The key to assigning specific members to specific sub-communities requires information about their stand-alone propensity to exert effort, their social influence, and their susceptibility to influence. One practical way to obtain such information is through self-reports at the time people join the community. Stand-alone effort is likely correlated with both aptitude and grit, the latter being the passion and motivation for the particular goal the community is designed to foster (Duckworth et al. 2007). Influence and susceptibility to influence can be measured through simple questionnaires similar to those used to measure opinion leaders and seekers (Flynn et al. 1996), though using easy-to-observe correlates of influence and influenceability identified by analyzing prior member behavior may be a better procedure for managers of established communities (Aral and Walker 2012; Iyengar et al. 2011, 2015).

A final implementation issue deals with how sustainable the network design is. Left to their own devices, people tend to form networks that exhibit two key structural features: clustering and homophily. The tendency of organic networks to be clustered ("friends of friends are friends") does not pose a challenge, since the proposed design involves complete clustering. The tendency for people to prefer homophilous networks ("birds of a feather flock together") warrants more discussion. The concern is that community members may want to rewire their pattern of connections, linking only to others similar to them. To the extent that the bases of homophily correlate with α, such rewiring will dilute the benefits of network design. However, two important elements imply that our design rules are likely to be robust to homophily. First, in our analysis, people prefer being connected to others, even if they have a different α, as long as $\rho \leq 1$ (see Equation 2). Consequently, being connected to peers with a different α is no reason to exit the community. Second, the empirical evidence reviewed in Section 2 shows that (i) people joining an organization or interest-based community in pursuit of a goal are often willing to accept an imposed group or network structure and that (ii) such an imposed structure often has a lasting effect on subsequent network structure and individual behavior, even after people are allowed to rewire their network connections freely (e.g., Centola 2011; Festinger et al. 1950; Hasan and Bagde 2013; Sacerdote 2001; Zhang et al. 2015a). Consequently, though homophily-based rewiring will dilute the benefits of network design, it will not negate them.

6.3 Suggestions for Future Research

Network design is an emerging area of research (Cerdeiro et al. 2017; Valente 2012), and many questions remain open. We studied the problem of a network designer who seeks to maximize total effort, in a setting where community members conform to descriptive social norms, and the key trade-off is how to best balance the positive conformity effect on people who would exert less effort when disconnected from others against the negative conformity effect on those who would exert more effort when disconnected.

Network design is applicable to many other areas of marketing, presenting different combinations of design objective, nature of social interaction, and key trade-off to be managed. One such area is ideation platforms, where the manager brings customers or employees together to generate ideas. Connected people may improve on each other's ideas, and the trade-off is to connect people with experience and expertise that are sufficiently different but not too different (De Vaan et al. 2015; Toubia and Netzer 2016; Uzzi et al. 2013). Another area is learning about new products and technologies when the objective is knowledge spread and acquisition, customers learn through observation, and the trade-off is to have informative social learning that does not overpower private learning (Zhang et al. 2015b). Yet another area is managing the trade-off between motivation and demotivation in performance posting, e.g., when publicly ranking the performance of salespeople or customer reps (Cai 2015).

Network design need not be restricted to choosing a network structure and may, for instance, also include a policy of incentives given to high-α types. This was not important in our analysis where people prefer being connected to others, even if they have a different α, as long as $\rho \leq 1$, but it may be useful when $\rho > 1$. Giving incentives to high-α types may also be effective in settings where the social influence mechanism is not one of conformity but of status or knowledge spillover, such that people like to connect to high-α types but the latter may need additional incentives to accept invitations to connect from low-α types (Wei et al. 2016). Such incentives may include public tokens of status, special benefits like pre-release access to new products, membership in advisory councils, and various other incentives already implemented in many communities. As companies and other platforms increasingly rely on building and managing communities, network design is bound to offer increasing opportunities to improve business outcomes and consumer welfare.

References

Alderson, D. L. (2008). OR forum - catching the "network science" bug: Insight and opportunity for the operations researcher. *Operations Research*, 56(5):1047–1065.

Allcott, H. (2011). Social norms and energy conservation. *Journal of Public Economics*, 95(9):1082–1095.

Allcott, H. and Kessler, J. B. (2015). The welfare effects of nudges: A case study of energy use social comparisons. Technical Report w21671, National Bureau of Economic Research.

Allcott, H. and Rogers, T. (2014). The short-run and long-run effects of behavioral interventions: Experimental evidence from energy conservation. *American Economic Review*, 104(10):3003–3037.

Ansari, A., Stahl, F., Heitmann, M., and Bremer, L. (2016). Building a social network for success. Working Paper, Columbia University.

Aral, S., Muchnik, L., and Sundararajan, A. (2009). Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks. *Proceedings of the National Academy of Sciences*, 106(51):21544–21549.

Aral, S. and Nicolaides, C. (2017). Exercise contagion in a global social network. *Nature Communications*, 8.

Aral, S. and Walker, D. (2012). Identifying influential and susceptible members of social networks. *Science*, 337(6092):337–341.

Ascarza, E., Ebbes, P., Netzer, O., and Danielson, M. (2017). Beyond the target customer: Social effects of customer relationship management campaigns. *Journal of Marketing Research*, 54(3):347–363.

Asch, S. E. (1956). Studies of independence and conformity: I. A minority of one against a unanimous majority. *Psychological Monographs: General and Applied*, 70(9):1–70.

Ayres, I., Raseman, S., and Shih, A. (2013). Evidence from two large field experiments that peer comparison feedback can reduce residential energy usage. *Journal of Law, Economics, and Organization*, 29(5):992–1022.

Bernheim, B. D. (1994). A theory of conformity. *Journal of Political Economy*, 102(5):841–877.

Bernheim, B. D. and Exley, C. L. (2015). Understanding conformity: An experimental investigation. Technical report, Working paper 16-070, Harvard Business School.

Bimpikis, K., Ozdaglar, A., and Yildiz, E. (2016). Competitive targeted advertising over networks. *Operations Research*, 64(3):705–720.

Cai, J. (2015). Silent or salient? Perks and perils of performance posting. Working paper, Texas A&M University.

Cartwright, D. and Harary, F. (1956). Structural balance: A generalization of Heider's theory. *Psychological Review*, 63(5):277–293.

Centola, D. (2010). The spread of behavior in an online social network experiment. *Science*, 329(5996):1194–1197.

Centola, D. (2011). An experimental study of homophily in the adoption of health behavior. *Science*, 334(6060):1269–1272.

Cerdeiro, D. A., Dziubiński, M., and Goyal, S. (2017). Individual security, contagion, and network design. *Journal of Economic Theory*, 170:182–226.

Chen, Y., Harper, F. M., Konstan, J., and Xin Li, S. (2010). Social comparisons and contributions to online communities: A field experiment on MovieLens. *American Economic Review*, 100(4):1358–1398.

Chen, Y., Wang, Q., and Xie, J. (2011). Online social interactions: A natural experiment on word of mouth versus observational learning. *Journal of Marketing Research*, 48(2):238–254.

Cialdini, R. B. and Goldstein, N. J. (2004). Social influence: Compliance and conformity. *Annual Review of Psychology*, 55:591–621.

Cialdini, R. B., Reno, R. R., and Kallgren, C. A. (1990). A focus theory of normative conduct: recycling the concept of norms to reduce littering in public places. *Journal of Personality and Social Psychology*, 58(6):1015–1026.

Davis, J. A. (1970). Clustering and hierarchy in interpersonal relations: Testing two graph theoretical models on 742 sociomatrices. *American Sociological Review*, 35(5):843–851.

De Vaan, M., Stark, D., and Vedres, B. (2015). Game changer: The topology of creativity. *American Journal of Sociology*, 120(4):1144–1194.

Duckworth, A. L., Peterson, C., Matthews, M. D., and Kelly, D. R. (2007). Grit: perseverance and passion for long-term goals. *Journal of Personality and Social Psychology*, 92(6):1087.

Ferraro, P. J. and Price, M. K. (2013). Using nonpecuniary strategies to influence behavior: Evidence from a large-scale field experiment. *Review of Economics and Statistics*, 95(1):64–73.

Festinger, L. (1954). A theory of social comparison processes. *Human Relations*, 7(2):117–140.

Festinger, L., Schachter, S., and Back, K. (1950). *Social Pressures in Informal Groups: A Study of Human Factors in Housing*. Stanford University Press, Palo Alto, CA.

Flynn, L. R., Goldsmith, R. E., and Eastman, J. K. (1996). Opinion leaders and opinion seekers: Two new measurement scales. *Journal of the Academy of Marketing Science*, 24(2):137–147.

Frey, B. S. and Meier, S. (2004). Social comparisons and pro-social behavior: Testing "conditional cooperation" in a field experiment. *American Economic Review*, 94(5):1717–1722.

Godes, D. and Mayzlin, D. (2009). Firm-created word-of-mouth communication: Evidence from a field test. *Marketing Science*, 28(4):721–739.

Goel, S. and Goldstein, D. G. (2013). Predicting individual behavior with social networks. *Marketing Science*, 33(1):82–93.

Goldenberg, J., Han, S., Lehmann, D. R., and Hong, J. W. (2009). The role of hubs in the adoption process. *Journal of Marketing*, 73(2):1–13.

Goldstein, N. J., Cialdini, R. B., and Griskevicius, V. (2008). A room with a viewpoint: Using social norms to motivate environmental conservation in hotels. *Journal of Consumer Research*, 35(3):472–482.

Haag, M. and Lagunoff, R. (2006). Social norms, local interaction, and neighborhood planning. *International Economic Review*, 47(1):265–296.

Hallsworth, M., List, J. A., Metcalfe, R. D., and Vlaev, I. (2017). The behavioralist as tax collector: Using natural field experiments to enhance tax compliance. *Journal of Public Economics*, 148:14–31.

Hasan, S. and Bagde, S. (2013). The mechanics of social capital and academic performance in an indian college. *American Sociological Review*, 78(6):1009–1032.

Hill, S., Provost, F., and Volinsky, C. (2006). Network-based marketing: Identifying likely adopters via consumer networks. *Statistical Science*, 21(2):256–276.

Hu, Y. and Van den Bulte, C. (2014). Nonmonotonic status effects in new product adoption. *Marketing Science*, 33(4):509–533.

Iyengar, R., Van den Bulte, C., and Lee, J. Y. (2015). Social contagion in new product trial and repeat. *Marketing Science*, 34(3):408–429.

Iyengar, R., Van den Bulte, C., and Valente, T. W. (2011). Opinion leadership and social contagion in new product diffusion. *Marketing Science*, 30(2):195–212.

Iyer, G. and Katona, Z. (2015). Competing for attention in social communication markets. *Management Science*, 62(8):2304–2320.

Katona, Z., Zubcsek, P. P., and Sarvary, M. (2011). Network effects and personal influences: The diffusion of an online social network. *Journal of Marketing Research*, 48(3):425–443.

Kelman, H. C. (1958). Compliance, identification, and internalization three processes of attitude change. *Journal of Conflict Resolution*, 2(1):51–60.

Kraut, R. E., Resnick, P., Kiesler, S., Burke, M., Chen, Y., Kittur, N., Konstan, J., Ren, Y., and Riedl, J. (2012). *Building Successful Online Communities: Evidence-based Social Design*. MIT Press, Cambridge, MA.

Martin, R. and Randal, J. (2008). How is donation behaviour affected by the donations of others? *Journal of Economic Behavior & Organization*, 67(1):228–238.

Murtha, B. R., Bharadwaj, S. G., and Van den Bulte, C. (2014). Interlocking networks: How and when do connections between buying and selling teams affect customer solutions? 14-120. Marketing Science Institute Report, Cambridge, MA.

Newman, M. E. (2006). Modularity and community structure in networks. *Proceedings of the National Academy of Sciences*, 103(23):8577–8582.

Peres, R. and Van den Bulte, C. (2014). When to take or forgo new product exclusivity: Balancing protection from competition against word-of-mouth spillover. *Journal of Marketing*, 78(2):83–100.

Phan, T. Q. and Godes, D. (2017). The evolution of influence through endogenous link formation. *Marketing Science*. in press.

Phillips, D. J. and Zuckerman, E. W. (2001). Middle-status conformity: Theoretical restatement and empirical demonstration in two markets. *American Journal of Sociology*, 107(2):379–429.

Plackett, R. L. (1965). A class of bivariate distributions. *Journal of the American Statistical Association*, 60(310):516–522.

Rovniak, L. S., Kong, L., Hovell, M. F., Ding, D., Sallis, J. F., Ray, C. A., Kraschnewski, J. L., Matthews, S. A., Kiser, E., Chinchilli, V. M., et al. (2016). Engineering online and in-person social networks for physical activity: A randomized trial. *Annals of Behavioral Medicine*, 50(6):885–897.

Sacerdote, B. (2001). Peer effects with random assignment: Results for Dartmouth roommates. *Quarterly Journal of Economics*, 116(2):681–704.

Shang, J. and Croson, R. (2009). A field experiment in charitable contribution: The impact of social information on the voluntary provision of public goods. *Economic Journal*, 119(540):1422–1439.

Sherif, M. (1935). A study of some social factors in perception. *Archives of Psychology*, 187:23–46.

Toubia, O. and Netzer, O. (2016). Idea generation, creativity, and prototypicality. *Marketing Science*, 36(1):1–20.

Uzzi, B., Mukherjee, S., Stringer, M., and Jones, B. (2013). Atypical combinations and scientific impact. *Science*, 342(6157):468–472.

Valente, T. W. (2012). Network interventions. *Science*, 337(6090):49–53.

Van den Bulte, C. and Joshi, Y. V. (2007). New product diffusion with influentials and imitators. *Marketing Science*, 26(3):400–421.

Van den Bulte, C. and Lilien, G. L. (2001). Medical innovation revisited: Social contagion versus marketing effort. *American Journal of Sociology*, 106(5):1409–1435.

Van den Bulte, C. and Moenaert, R. K. (1998). The effects of R&D team co-location on communication patterns among R&D, marketing, and manufacturing. *Management Science*, 44(11-part-2):S1–S18.

Wei, Y., Yildirim, P., Van den Bulte, C., and Dellarocas, C. (2016). Credit scoring with social network data. *Marketing Science*, 35(2):234–258.

Zafar, B. (2011). An experimental investigation of why individuals conform. *European Economic Review*, 55(6):774–798.

Zhang, J., Brackbill, D., Yang, S., Becker, J., Herbert, N., and Centola, D. (2016). Support or competition? how online social networks increase physical activity: A randomized controlled trial. *Preventive Medicine Reports*, 4:453–458.

Zhang, J., Brackbill, D., Yang, S., and Centola, D. (2015a). Efficacy and causal mechanism of an online social media intervention to increase physical activity: results of a randomized controlled trial. *Preventive Medicine Reports*, 2:651–657.

Zhang, J., Liu, Y., and Chen, Y. (2015b). Social learning in networks of friends versus strangers. *Marketing Science*, 34(4):573–589.

Appendix: Proofs

Proof of Proposition 1: First we rearrange Equation (4) to write:

$$y_i^* \left(1 + \sum_{j \sim i} \rho_{ij}\right) - \sum_{j \sim i} \rho_{ij} y_j^* = \alpha_i$$

which in matrix form can be represented as

$$\mathbf{L}\mathbf{y}^* = \boldsymbol{\alpha}.$$

Since $\mathbf{L}$ is strictly diagonally dominant, it is non-singular and thus invertible. Consequently, $\mathbf{L}^{-1}\boldsymbol{\alpha}$ exists and is the unique equilibrium. In fact, using results on strictly diagonally dominant matrix, one can show that all the entries of $\mathbf{L}^{-1}$ are non-negative.

Proof of Lemma 1: Using the Inverse Function Theorem on matrix functions, we have

$$\frac{\partial \mathbf{L}^{-1}}{\partial s_{ij}} = -\mathbf{L}^{-1} \frac{\partial \mathbf{L}}{\partial s_{ij}} \mathbf{L}^{-1}$$

Hence

$$\frac{\partial \mathbf{y}^*}{\partial s_{ij}} = \frac{\partial \mathbf{L}^{-1}}{\partial s_{ij}}\boldsymbol{\alpha} = -\mathbf{L}^{-1}\frac{\partial \mathbf{L}}{\partial s_{ij}}\mathbf{y}^*,$$

Notice that s_{ij} represents the link strength between i and j, so it appears four times in $\mathbf{L}$ (in L_{ij}, L_{ji}, L_{ii}, and L_{jj}). Hence,

$$-\frac{\partial \mathbf{L}}{\partial s_{ij}} = \begin{pmatrix} & & & \\ & -\rho_{ij} & \rho_{ij} & \\ & & & \\ & \rho_{ji} & -\rho_{ji} & \\ & & & \end{pmatrix},$$

and

$$-\frac{\partial \mathbf{L}}{\partial s_{ij}}\mathbf{y}^* = (-\rho_{ij}y_i^* + \rho_{ij}y_j^*)\mathbf{e}_i + (\rho_{ji}y_i^* - \rho_{ji}y_j^*)\mathbf{e}_j,$$

which gives us what we need to show.

Proof of Lemma 2: It is clear that in the case where $y_i^* = y_j^*$, altering the link weight between i and j (s_{ij}) does not change the activity of these individuals. Therefore there will be no change in the total activity of the network. Next, without loss of generality, let's consider the case where $y_i^* < y_j^*$. Note that any change in the link weight s_{ij} cannot reverse the order between y_i^* and y_j^*. This is because if it could, we had to first reach the state where $y_i^* = y_j^*$, at which point any further change in s_{ij} will have no effects on the activity levels.

First, we show that $\frac{\partial(y_j^* - y_i^*)}{\partial s_{ij}} < 0$. In words, as the tie strength increases between i and j, the difference between their efforts should decrease. Recall that from Lemma 1,

$$\frac{\partial \mathbf{y}^*}{\partial s_{ij}} = \mathbf{L}^{-1}\mathbf{v}_{ij}^*.$$

This says that essentially, the value of $\frac{\partial \mathbf{y}^*}{\partial s_{ij}}$ equals to the equilibrium effort vector where the out-of-network effort levels are $\mathbf{v}_{ij}^*$ instead of $\boldsymbol{\alpha}$. Note that $\mathbf{v}_{ij}^*$ is all zero except for the ith entry, which is positive, and the jth entry, which is negative. It is not difficult to show that for such out-of-network effort levels as given by $\mathbf{v}_{ij}^*$, individual i will still exert more effort than j in equilibrium. To see this, without loss of generality, suppose that i exerts less or equal effort compared to j and i's effort is negative. For this to happen, i must have a contact whose effort is lower than i's. By the same token, this contact of i must have a further contact whose effort is even lower. Continue this argument and we will have an infinite sequence of distinct individuals, which is a contradiction because n is finite. As a result, $(\mathbf{L}^{-1}\mathbf{v}_{ij}^*)_i > (\mathbf{L}^{-1}\mathbf{v}_{ij}^*)_j$, which gives us $\frac{\partial(y_j^* - y_i^*)}{\partial s_{ij}} < 0$.

Next, we focus on the second derivative $\partial^2\mathbf{y}^*/\partial s_{ij}^2$. Again, using the equation in Lemma 1, we have

$$\begin{aligned}\frac{\partial^2 \mathbf{y}^*}{\partial s_{ij}^2} &= -\mathbf{L}^{-1}\frac{\partial \mathbf{L}}{\partial s_{ij}}\mathbf{L}^{-1}\mathbf{v}_{ij}^* + \mathbf{L}^{-1}\frac{\partial \mathbf{v}_{ij}^*}{\partial s_{ij}} \\ &= -\mathbf{L}^{-1}\frac{\partial \mathbf{L}}{\partial s_{ij}}\frac{\partial \mathbf{y}^*}{\partial s_{ij}} + \mathbf{L}^{-1}\frac{\partial \mathbf{v}_{ij}^*}{\partial s_{ij}}\end{aligned}$$

Notice that:

$$\begin{aligned}\frac{\partial \mathbf{v}_{ij}^*}{\partial s_{ij}} &= \frac{\partial (y_j^* - y_i^*)}{\partial s_{ij}} \cdot \left(\rho_{ij}\mathbf{e}_i^* - \rho_{ji}\mathbf{e}_j^*\right) \\ &= \frac{\frac{\partial y_j^*}{\partial s_{ij}} - \frac{\partial y_i^*}{\partial s_{ij}}}{y_j^* - y_i^*} \cdot \mathbf{v}_{ij}^* \\ &= -\frac{\partial \mathbf{L}}{\partial s_{ij}} \frac{\partial \mathbf{y}^*}{\partial s_{ij}}.\end{aligned}$$

The last line above uses the expression for $-\frac{\partial \mathbf{L}}{\partial s_{ij}}$ in the proof of Lemma 2. Hence,

$$\begin{aligned}\frac{\partial^2 \mathbf{y}^*}{\partial s_{ij}^2} &= 2\mathbf{L}^{-1}\frac{\partial \mathbf{v}_{ij}^*}{\partial s_{ij}} \\ &= 2\frac{\partial \mathbf{y}^*}{\partial s_{ij}} \times \frac{\frac{\partial y_j^*}{\partial s_{ij}} - \frac{\partial y_i^*}{\partial s_{ij}}}{y_j^* - y_i^*}.\end{aligned}$$

Because $T \equiv \mathbf{1}'\mathbf{y}^*$, we have

$$\frac{d^2T}{ds_{ij}^2} = \frac{2dT}{ds_{ij}} \times \frac{\frac{\partial y_j^*}{\partial s_{ij}} - \frac{\partial y_i^*}{\partial s_{ij}}}{y_j^* - y_i^*}.$$

This can be seen as a first-order differential equation for dT/ds_{ij}, and

$$\frac{\frac{\partial y_j^*}{\partial s_{ij}} - \frac{\partial y_i^*}{\partial s_{ij}}}{y_j^* - y_i^*} < 0,\ \forall s_{ij} \in [0, 1].$$

The fact that dT/ds_{ij} satisfies such a differential equation tells us that if dT/ds_{ij} is positive (or negative, respectively) at some particular value of s_{ij}, then dT/ds_{ij} stays positive (or negative, respectively) for all values of s_{ij}.

Proof of Proposition 2: The total effort $T \equiv \mathbf{1}'\mathbf{y}^*$. Suppose that we are dealing with a weighted network. Then Lemma 1 and Lemma 2 together say that for any given network, if $\mathbf{1}'\mathbf{L}^{-1}\mathbf{v}_{ij}^*$, evaluated at the given network and the associated equilibrium, is positive, then increasing s_{ij} all the way to 1 will be increasing T; on the other hand, if $\mathbf{1}'\mathbf{L}^{-1}\mathbf{v}_{ij}^*$ is negative, then decreasing s_{ij} all the way to 0 will be increasing T. It is straightforward to translate this result to an unweighted (or discrete) network, which is the result in the proposition.

Proof of Proposition 3: We only need to consider the case where $\boldsymbol{\rho}$ and $\boldsymbol{\rho}'$ are only different at a single pair (i, j). When there are two different pairs, the conclusion can be extended by considering a third intermediate population that is only one-pair different from either of the two populations.

Let the total activity of $(\boldsymbol{\alpha}, \boldsymbol{\rho})$ under its optimal network be T. Let the total activity of $(\boldsymbol{\alpha}, \boldsymbol{\rho}')$ under its optimal network be T'. Let the total activity of $(\boldsymbol{\alpha}, \boldsymbol{\rho}')$ under the optimal network for $(\boldsymbol{\alpha}, \boldsymbol{\rho})$ be T^o. By Lemma 2, in the optimal network for $(\boldsymbol{\alpha}, \boldsymbol{\rho})$, if there is a link between i and j, then increasing the strength of that link does not decrease total activity; if there is no link between i and j, increasing the strength of that link has no effect on total activity. Thus, $T^o \geq T$. It is by construction that $T' \geq T^o$. So $T' \geq T$ holds.

Proof of Proposition 4: For any network configuration, at the equilibrium, we have

$$y_i^* = \alpha_i + \sum_{j \sim i} \rho_{ij}(y_j^* - y_i^*),$$

as given in Equation (5). Adding this equation across all individuals, we obtain the total effort as

$$T = \sum_{i=1}^{n} \alpha_i + \sum_{i \sim j, i > j} (\rho_{ij} - \rho_{ji})(y_j^* - y_i^*).$$

The second summation on the right hand side sums over all the links in the network. The highest possible activity that one can carry in the network is $\max_i\{\alpha_i\}$ and the lowest is $\min_i\{\alpha_i\}$. In addition, the maximum number of links is $\frac{n(n-1)}{2}$. Therefore,

$$\sum_{i \sim j, i > j} (\rho_{ij} - \rho_{ji})(y_j^* - y_i^*) \leq \frac{n(n-1)}{2} \rho^+ \alpha^+.$$

This gives the expression in the proposition.

Proof of Proposition 5: Assume that for a given population, there is an optimal network for the first k individuals, and an optimal network for the remaining individuals. Combining these two networks together does not necessarily give the optimal network for the entire population. Hence it is easy to see that

$$Q(\Phi, n+1) > Q(\Phi, n+1-k) + Q(\Phi, k), \quad k = 1, 2, ..., n.$$

which implies

$$q(\Phi, n+1) > \frac{(n+1-k) \cdot q(\Phi, n+1-k) + k \cdot q(\Phi, k)}{n+1}, \quad k = 1, 2, ..., n.$$

For a series of ratios between positive numbers, the ratio between the sum of the nominators and the sum of the denominators must be smaller than the largest ratio in the series. So

$$\begin{aligned} q(\Phi, n+1) &> \frac{\sum_{k=1}^{n} [(n+1-k) \cdot q(\Phi, n+1-k) + k \cdot q(\Phi, k)]}{\sum_{k=1}^{n} (n+1)} \\ &= \frac{2 \sum_{k=1}^{n} k \cdot q(\Phi, k)}{n\,(n+1)} \\ &= \frac{\sum_{k=1}^{n} k \cdot q(\Phi, k)}{\sum_{k=1}^{n} k}. \end{aligned}$$

Using this inequality repeatedly in the following manner leads to the inequality in the proposition:

$$\begin{aligned}
q(\Phi, n+1) &> \frac{n \cdot q(\Phi, n) + (n-1) \cdot q(\Phi, n-1) + ... + q(\Phi, 1)}{n\,(n+1)\,/2} \\
&> \frac{\frac{n+1}{2} \cdot q(\Phi, n) + \frac{n+1}{n}\left[(n-1) \cdot q(\Phi, n-1) + ... + q(\Phi, 1)\right]}{n\,(n+1)\,/2} \\
&= \frac{q(\Phi, n)}{n} + \frac{(n-1) \cdot q(\Phi, n-1) + ... + q(\Phi, 1)}{n^2/2} \\
&> \frac{q(\Phi, n)}{n} + \frac{\frac{n}{2} \cdot q(\Phi, n-1) + \frac{n}{n-1}\left[(n-2) \cdot q(\Phi, n-2) + ... + q(\Phi, 1)\right]}{n^2/2} \\
&= \frac{q(\Phi, n)}{n} + \frac{q(\Phi, n-1)}{n} + \frac{(n-2) \cdot q(\Phi, n-2) + ... + q(\Phi, 1)}{n(n-1)/2} \\
&... \quad ... \\
&= \sum_{k=1}^{n} \frac{q(\Phi, k)}{n}.
\end{aligned}$$

Proof of Lemma 3: Part (i): For the equilibrium activity of individual i, we have

$$y_i^* = \frac{\alpha_i + \lambda_i\,(T - y_i^*)}{1 + \lambda_i(n-1)},$$

which implies

$$\begin{aligned}
(1 + n\lambda_1)\, y_1^* &= \alpha_1 + \lambda_1 T, \\
(1 + n\lambda_2)\, y_2^* &= \alpha_2 + \lambda_2 T, \\
&\vdots \\
(1 + n\lambda_n)\, y_n^* &= \alpha_n + \lambda_n T.
\end{aligned}$$

Dividing the ith equation by $1 + n\lambda_i$ and summing over all equations yield

$$T = \sum_{i=1}^{n} \frac{\alpha_i}{1 + n\lambda_i} + T \times \sum_{i=1}^{n} \frac{\lambda_i}{1 + n\lambda_i}.$$

Solving for T yields the expression.

Part (ii): For the equilibrium activity of individual k, we have

$$y_i^* = \frac{\alpha_i + \sum_{j \neq i} \chi_j y_j^*}{1 + \sum_{j \neq i} \chi_j},$$

which implies

$$y_i^* = \frac{\alpha_i + \sum_{j=1}^{n} \chi_j y_j^*}{1 + \sum_{j=1}^{n} \chi_j}.$$

Summing this equation across i, we have

$$T = \frac{\sum_{i=1}^{n} \alpha_i + n \sum_{i=1}^{n} \chi_i y_i^*}{1 + \sum_{i=1}^{n} \chi_i}.$$

Had we multiply each equation by χ_i before the summation, we would get

$$\sum_{i=1}^{n} \chi_i y_i^* = \frac{\sum_{i=1}^{n} \alpha_i \chi_i + \sum_{i=1}^{n} \chi_i \cdot \sum_{i=1}^{n} \chi_i y_i^*}{1 + \sum_{i=1}^{n} \chi_i},$$

which implies

$$\sum_{i=1}^{n} \chi_i y_i^* = \sum_{i=1}^{n} \alpha_i \chi_i.$$

Hence,

$$T = \frac{\sum_{i=1}^{n} \alpha_i + n \sum_{i=1}^{n} \alpha_i \chi_i}{1 + \sum_{i=1}^{n} \chi_i} = \frac{\sum_{i=1}^{n} \alpha_i \left(1 + n\chi_i\right)}{\frac{1}{n} \sum_{i=1}^{n} \left(1 + n\chi_i\right)}.$$

Proof of Lemma 4: It is straightforward that if α_i and λ_i follow a positive linear relation then the empty network is locally optimal.
For the rest of the proof we express the negative linear relationship between the two parameters by writing $\alpha_i = \theta - \gamma\lambda_i$ with $\theta > 0$ and $\gamma > 0$.

Consider any two nodes i and j in the complete network. Substituting T into the equilibrium effort given in Equation (4) yields:

$$y_i^* = \frac{\alpha_i + \lambda_i T}{1 + \lambda_i n} = \frac{\theta + \lambda_i (T - \gamma)}{1 + \lambda_i n}$$

Hence

$$y_i^* - y_j^* = \frac{(T - \gamma - n\theta)}{(1 + \lambda_i n)(1 + \lambda_j n)} \left(\lambda_i - \lambda_j\right)$$

From Lemma 3, we can set the upper bound for $T < n \cdot \max_i \{\alpha_i\} < n\theta$. Thus $T - \gamma - n\theta < 0$. As a result, in the equilibrium, the sign for $\lambda_i - \lambda_j$ is always the reverse of the sign for $y_i^* - y_j^*$.

Now, consider any two nodes i and j in the complete network. We have two cases:

- Case $\alpha_i = \alpha_j$: Then $\lambda_i = \lambda_j$ and $y_i^* = y_j^*$. From Proposition 2, we see that removing the link between i and j will lead to a zero change on T.

- Case $\alpha_i \neq \alpha_j$: Then $(\lambda_i - \lambda_j)(y_j^* - y_i^*) > 0$. Using the expression for T in Lemma 3, we have

$$\left(\lambda_i \frac{\partial T}{\partial \alpha_i} - \lambda_j \frac{\partial T}{\partial \alpha_j}\right) (y_j^* - y_i^*) < 0.$$

  Notice that $\frac{\partial T}{\partial \alpha_i} = \mathbf{1}'\mathbf{L}^{-1}\mathbf{e}_i$ and $\frac{\partial T}{\partial \alpha_j} = \mathbf{1}'\mathbf{L}^{-1}\mathbf{e}_j$, so the above becomes

$$\mathbf{1}'\mathbf{L}^{-1}(\lambda_i \mathbf{e}_i - \lambda_j \mathbf{e}_j)(y_j^* - y_i^*) < 0.$$

  With $\chi_i = \chi_j = 1$, this is equivalent to $\mathbf{1}'\mathbf{L}^{-1}\mathbf{v}_{ij}^* > 0$. Applying Proposition 2, we see that removing the link between i and j will decrease total activity.

Proof of Lemma 5: It is straightforward to see that if α_i and χ_i follow a negative linear relation then the empty network is locally optimal.
For the rest of the proof we express the positive linear relationship between the two parameters by writing $\alpha_i = \theta + \gamma\chi_i$ with $\theta > 0$ and $\gamma > 0$. Consider any two nodes i and j in the complete network. From the proof of Lemma 3 part (ii), we have

$$y_i^* = \frac{\alpha_i + \sum_{k=1}^n \alpha_k \chi_k}{1 + \sum_{k=1}^n \chi_k}.$$

Hence

$$y_i^* - y_j^* = \frac{\alpha_i - \alpha_j}{1 + \sum_{k=1}^n \chi_k} = \frac{\gamma(\chi_i - \chi_j)}{1 + \sum_{k=1}^n \chi_k}.$$

This equation indicates that the sign of $(y_i^* - y_j^*)$ is the same as $(\chi_i - \chi_j)$ in the complete network.
Now, consider any two nodes i and j. There are two cases:

- Case $\alpha_i = \alpha_j$: We have $\lambda_i = \lambda_j$and thus $y_i^* = y_j^*$. Removing the link between i and j does not alter the value of T.

- Case $\alpha_i \neq \alpha_j$: In this case $(\chi_j - \chi_i)(y_j^* - y_i^*) > 0$. Using the expressions in Lemma 3, we can get

$$\left(\chi_j \frac{\partial T}{\partial \alpha_i} - \chi_i \frac{\partial T}{\partial \alpha_j}\right)(y_j^* - y_i^*) > 0$$

  Notice that $\frac{\partial T}{\partial \alpha_i} = \mathbf{1}'\mathbf{L}^{-1}\mathbf{e}_i$ and $\frac{\partial T}{\partial \alpha_j} = \mathbf{1}'\mathbf{L}^{-1}\mathbf{e}_j$, so the above becomes

$$\mathbf{1}'\mathbf{L}^{-1}(\chi_j\mathbf{e}_i - \chi_i\mathbf{e}_j)(y_j^* - y_i^*) > 0.$$

  With $\lambda_i = \lambda_j = 1$, this is equivalent to $\mathbf{1}'\mathbf{L}^{-1}\mathbf{v}_{ij}^* > 0$. Applying Proposition 2, we see that removing the link between i and j will decrease total activity.

Proof of Proposition 6: Consider a sub-community whose members are located along a segment parallel to the diagonal line. Among these people, by Lemma 4 we know that the complete network is locally optimal, so removing any link between two members in this sub-community will not increase T.
Next, we consider adding a link between two individuals. Let m be the size of a sub-community. From the proof of Lemma 3, we know that each individual's activity in this sub-community is given by

$$\begin{aligned} y_i^* &= \frac{\alpha_i + \lambda_i T}{1 + \lambda_i m} \\ &\simeq \frac{\sum_{k=1}^m \frac{\alpha_i}{1+m\lambda_k}}{\sum_{k=1}^m \frac{1}{1+m\lambda_k}} \\ &\simeq \underbrace{\int_{\max\{\underline{\lambda}, r+\underline{\lambda}\}}^{\min\{\bar{\lambda}, r+\bar{\lambda}\}} \frac{\hat{\alpha}(\lambda, r)}{\lambda}\varphi(\lambda, r)d\lambda}_{\Gamma(r)} \cdot \underbrace{\left(\int_{\max\{\underline{\lambda}, r+\underline{\lambda}\}}^{\min\{\bar{\lambda}, r+\bar{\lambda}\}} \frac{1}{\lambda}\varphi(\lambda, r)d\lambda\right)^{-1}}_{\Lambda(r)^{-1}}. \end{aligned}$$

The above approximations become exact as $m \to +\infty$. On the other hand, using Lemma 3 on this sub-community, we have

$$\lambda_i \frac{\partial T}{\partial \alpha_i} = \frac{\frac{\lambda_i}{1+m\lambda_i}}{\frac{1}{m}\sum_{k=1}^{n}\frac{1}{1+m\lambda_k}}$$
$$\simeq \frac{1}{\frac{1}{m}\sum_{k=1}^{m}\frac{1}{\lambda_k}}$$
$$\simeq \Lambda(r)^{-1}.$$

Again, the approximations become exact as $m \to +\infty$. Take individual i in segment r_1 and j in segment r_2. By Proposition 2, the necessary and sufficient condition for the total activity to not increase when a link is added between i and j is

$$\mathbf{1}'\mathbf{L}^{-1}\left(\lambda_i \mathbf{e}_i - \lambda_j \mathbf{e}_j\right)\left(y_j^* - y_i^*\right) \le 0.$$

Since $\frac{\partial T}{\partial \alpha_i} = \mathbf{1}'\mathbf{L}^{-1}\mathbf{e}_i$, this condition can be re-written as:

$$\left(\lambda_i \frac{\partial T}{\partial \alpha_i} - \lambda_j \frac{\partial T}{\partial \alpha_j}\right)\left(y_j^* - y_i^*\right) \le 0,$$

which, if $n \to +\infty$, becomes:

$$\left[\Lambda(r_1)^{-1} - \Lambda(r_2)^{-1}\right] \cdot \left[\Gamma(r_1)\Lambda(r_1)^{-1} - \Gamma(r_2)\Lambda(r_2)^{-1}\right] \ge 0.$$

A sufficient condition for this inequality to hold is that both $\Lambda(r)$ and $\Lambda(r)/\Gamma(r)$ should be decreasing in r, and a regularity condition can be proposed to satisfy these criteria. We provide it next.

A regularity condition for Proposition 6: First, use the following function, parameterized by r, to denote the lines parallel to the diagonal line:

$$\hat{\alpha}(\lambda, r) \equiv \bar{\alpha} - \frac{\bar{\alpha} - \underline{\alpha}}{\bar{\lambda} - \underline{\lambda}}\left(\lambda - \underline{\lambda} - r\right),$$

where $r \in [\underline{\lambda} - \bar{\lambda}, \bar{\lambda} - \underline{\lambda}]$. Each r corresponds to one line; $\hat{\alpha}(\lambda, 0)$ is the diagonal itself. Define two functions of r:

$$\Gamma(r) \equiv \int_{\max\{\underline{\lambda}, r+\underline{\lambda}\}}^{\min\{\bar{\lambda}, r+\bar{\lambda}\}} \frac{\hat{\alpha}(\lambda, r)}{\lambda}\varphi(\lambda, r)d\lambda, \quad \Lambda(r) \equiv \int_{\max\{\underline{\lambda}, r+\underline{\lambda}\}}^{\min\{\bar{\lambda}, r+\bar{\lambda}\}} \frac{1}{\lambda}\varphi(\lambda, r)d\lambda$$

where φ is the density along the parallel line:

$$\varphi(\lambda, r) \equiv \frac{\Phi\left[\hat{\alpha}(\lambda, r), \lambda\right]}{\int_{\max\{\underline{\lambda}, r+\underline{\lambda}\}}^{\min\{\bar{\lambda}, r+\bar{\lambda}\}} \Phi\left[\hat{\alpha}(\lambda, r), \lambda\right] d\lambda}.$$

It is easy to check this condition numerically. In particular, one can verify that it holds virtually for any Φ that is the product of any two bell-shaped Beta distributions.

Proof of Proposition 7: The derivation of Proposition 7 follows the same steps in Proof of Proposition 6. A regularity condition for Proposition 7 can be provided in the following way.

First, use the following function, parameterized by r, to denote the lines parallel to the diagonal line:

$$\hat{\alpha}(\chi, r) \equiv \underline{\alpha} + \frac{\bar{\alpha} - \underline{\alpha}}{\bar{\chi} - \underline{\chi}} (\chi - \underline{\chi} - r),$$

where $r \in [\underline{\chi} - \bar{\chi}, \bar{\chi} - \underline{\chi}]$. Each r corresponds to one line; $\hat{\alpha}(\chi, 0)$ is the diagonal itself. Define two functions of r:

$$\Gamma(r) \equiv \int_{\max\{\underline{\chi}, r+\underline{\chi}\}}^{\min\{\bar{\chi}, r+\bar{\chi}\}} \chi \hat{\alpha}(\chi, r) \varphi(\chi, r) d\chi, \quad \Lambda(r) \equiv \int_{\max\{\underline{\chi}, r+\underline{\chi}\}}^{\min\{\bar{\chi}, r+\bar{\chi}\}} \chi \varphi(\chi, r) d\chi$$

where φ is the density along the parallel line:

$$\varphi(\chi, r) \equiv \frac{\Phi\left[\hat{\alpha}(\chi, r), \lambda\right]}{\int_{\max\{\underline{\chi}, r+\underline{\chi}\}}^{\min\{\bar{\chi}, r+\bar{\chi}\}} \Phi\left[\hat{\alpha}(\chi, r), \chi\right] d\chi}.$$

The proposition holds if both $\Lambda(r)$ and $\Lambda(r)/\Gamma(r)$ are increasing in r. It is easy to check this condition numerically, again, by verifying that it holds for any Φ that is the product of any two bell-shaped Beta distributions.

Online Appendix

The optimal network structure in the Example in Section 4 is found through enumeration (i.e., checking all possible networks). All other numerical analyses in the main text are based on applying the following Fast Multiple Link Altering (FMLA) algorithm with multiple starting guesses.

Algorithm FMLA (Fast Multiple Link Altering)

1. Start with a guess (e.g., the empty network).
2. Given the current network, use Proposition 2 to find all the pairs such that altering the link of each individually (i.e., remove the link if it exists, add otherwise) would increase total activity T. If no such pair exists, stop.
3. Alter the links of these pairs together: If T is increased, go back to Step 2; if T is not increased, alter the link of any single one of these pairs, then go back to Step 2.

This algorithm is an improved version of the more basic and "brute force" Link Altering (LA) algorithm:

Algorithm LA (Link Altering):

1. Start with a guess (e.g., the empty network).
2. Take a pair of individuals. Compute the equilibrium if their link is altered (i.e., remove the link if it exists, add otherwise). Going through pairs one by one until a pair is found for which link altering increases T. If no such pair can be found, stop.
3. Alter the link of the pair found in Step 2. Go back to Step 2.

Because there are a finite number of possible network configurations and the algorithms always increase T in each iteration, they are guaranteed to converge. Similar to most algorithms which maximize continuous functions, the solution guarantees a local optimal solution such that, no addition or removal of a single link will result in a better network.

Suppose the optimal network has m links. Because the LA algorithm changes only one link in one iteration, it has to run through at least m iterations to reach it. In each iteration, it also needs to compute at least one, likely many, what-if equilibria; the computation of each equilibrium requires the inversion of an influence matrix as defined in Proposition 1. The FMLA speeds up this process in two ways. First, in each iteration, it checks all the pairs simultaneously using Proposition 2, requiring one single matrix inversion. This matrix inversion is needed anyway to compute the current equilibrium. Second, FMLA changes multiple links simultaneously in each iteration. Because such simultaneous altering could result in a lower T, we reserve the option of reverting to single-link alteration that guarantees a higher T.

We compare the computational efficiency of these two algorithms under a bivariate Beta(2, 2) distribution Φ with various correlations and population sizes. The computation time of FMLA versus LA is shown in Figure B1. FMLA is always faster than LA, and the computational time advantage widens as n becomes bigger. As a general observation, the basic LA seems to have a time complexity of $O(n^4)$, while FMLA seems to be $O(n^3)$.

To also compare FMLA and LA on their ability to identify the network structure resulting in the highest level of effort, we gauge them against a benchmark. Since the true optimum is unknown,

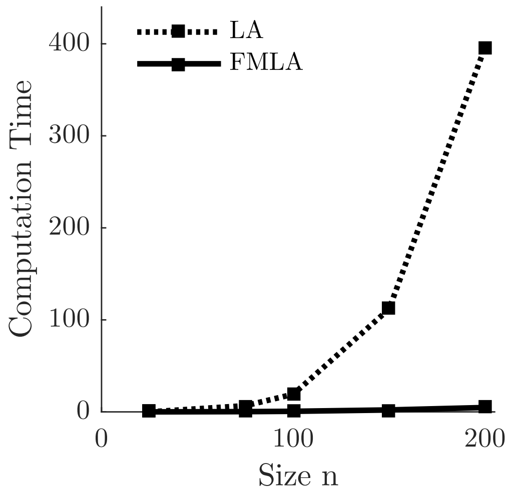

Size n	Run Time (sec)		
	LA	FMLA	Ratio
25	0.185	0.0150	12.3
75	6.71	0.273	24.5
100	19.2	0.634	30.3
150	111.9	2.11	53.0
200	396.4	4.75	83.5

Notes: Each run time is averaged across 500 communities. For each community, the members are drawn from $(\alpha_i, \lambda_i) \sim \text{Beta}(2, 2) \times \text{Beta}(2, 2)$ and χ_i is fixed at 1.

Figure B1: Computation Time Comparison

the ideal benchmark is a general method that is robust in reaching a good solution, even though it may do so very slowly. We choose simulated annealing (SA), a very general global optimization algorithm. SA is a stochastic algorithm and is not greedy; it allows temporary setback in the object to possibly go over local optima. Our implementation of SA follows Gastner and Newman (2006). In each iteration, it proposes to either: (i) add a link, (ii) remove a link, or (iii) move a link from one pair to another.

Figure B2 presents the test results on the quality of the solutions found by FMLA. Each dot represents a different population of individuals, $\{\alpha_i, \lambda_i\}_{i=1}^n$. For a more extensive test, we vary both the size n and the correlation between α_i and λ_i across populations. The top row compares FMLA against LA. It shows that despite the disparity in speed, their performances in improving effort are almost identical. The middle row compares FMLA against SA. The plots suggest that the networks found by FMLA are globally optimal or at least very close to being globally optimal. In the bottom row, we experiment with using multiple starting points for FMLA. The plots suggests that there are no discernible gains from using multiple starting points. Overall, these results support the use of FMLA to find optimal networks.

As a final check of the FMLA algorithm, we also compare it to an alternative that exploits a greater set of our analytic results about how to arrange network members into sub-communities by drawing parallel lines on the α-λ plane and assigning people between two lines into a community.

Diagonal Community Algorithm (DCA)

1. Find a rectangle on the α-λ plane that tightly contains all the individuals. Choose an integer k, an intended number of sub-communities.
2. Divide the rectangle into k strips by drawing $k - 1$ equidistant lines all parallel to the negatively-sloped diagonal of the rectangle.
3. Connect any two individuals within each strip, while leaving any two individuals unconnected if they are located in different strips.
4. Repeat Step 2 and 3 for many possible values of k (e.g., from 1 to 50). Choose the k that results in the highest collective effort.

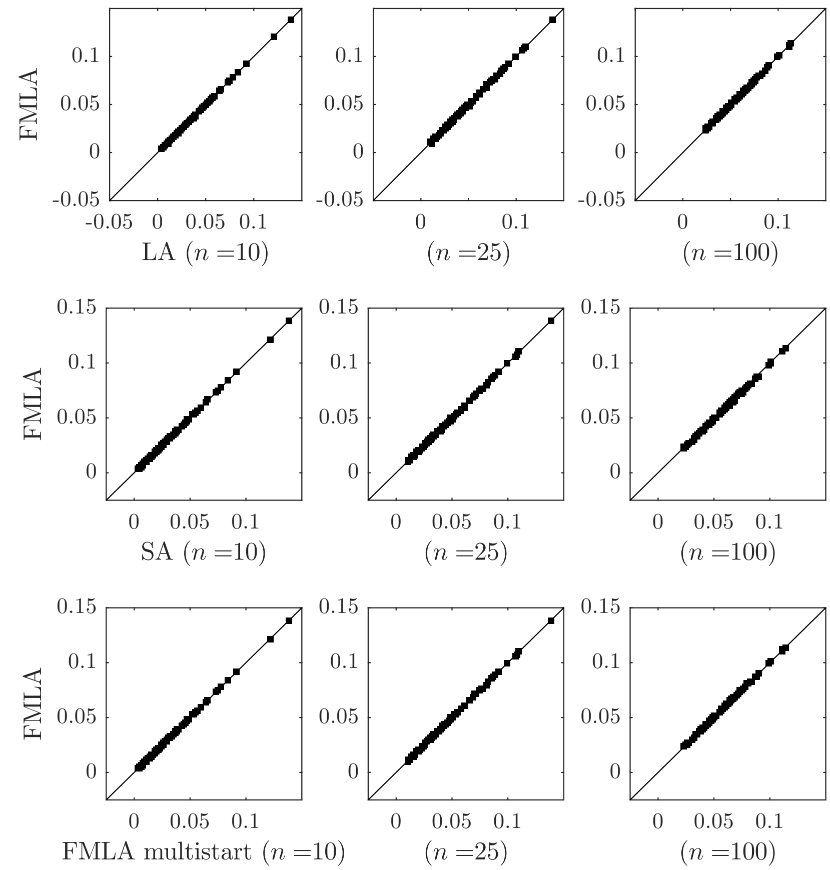

Notes: Each dot displays per capita improvements by two different algorithms for a population of n individuals. For each population, a correlation coefficient is first drawn uniformly from [-.5,.5] and each (α_i, λ_i) is drawn from a bi-variate Beta(2,2) with that correlation. The correlation is introduced using Plackett (1965). The top row compares FMLA against LA. The middle row compares FMLA with simulated annealing. The bottom row compares FMLA against FMLA with multiple starts. All the algorithms start with the empty network. For the FMLA with multi-starts, it additionally uses the following starting points: the complete network, the network given by the pairwise rule (see Sec.4), and 100 Erdős–Renyi random networks with various link formation probabilities. SA is implemented as in Gastner and Newman (2006), with a very slow cooling factor at 0.9999.

Figure B2: Performance of the FMLA Algorithm

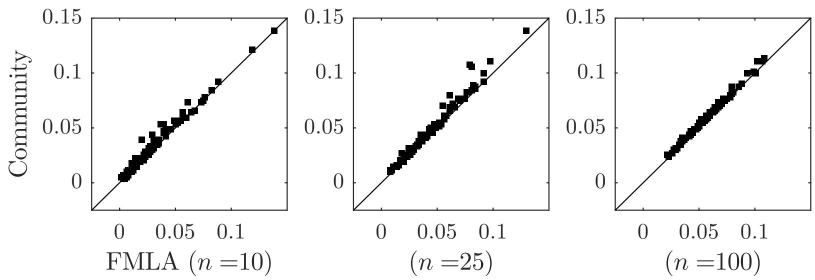

Notes: Each dot displays per capita improvement by DCA vs. FMLA for a population of n individuals. For each population, a correlation coefficient is first drawn uniformly from [-.5,.5] and each (α_i, λ_i) is drawn from a bi-variate Beta(2,2) with that correlation.

Figure B3: Performance of the Pairwise Rule and Community Arrangement

The upside of this Diagonal Community Algorithm (CDA) is that it is very fast and the complexity grows slowly with n. It does not attempt to find the optimal network but directly approaches a near-optimal solution. As reported in Figure B3, the performance of CDA is generally only slightly worse than that of FMLA. This is remarkable given the simplicity of CDA. In addition, notice that the performance gap diminishes as the size of population becomes large. This is consistent with the fact that our analytical results on the diagonal sub-communities (Proposition 6 and 7) are derived under $n \to +\infty$.

Reference for Online Appendix

Gastner, M. T. and Newman, M. (2006). Optimal design of spatial distribution networks. *Physical Review E*, 74(1):016117.