An Equilibrium Model of Experimentation on Networks*

Simon Board† and Moritz Meyer-ter-Vehn‡

November 1, 2021

Preliminary and Incomplete - Please do not circulate

Abstract

We introduce a model of strategic experimentation on social networks where forward-looking agents learn from their own and neighbors' successes. In equilibrium, a private discovery phase is followed by a social diffusion phase; the anticipation of future social information crowds out agents' own experimentation. We first study tree-like networks, characterize learning dynamics via ODEs, and draw tight comparisons between directed, undirected and clustered networks. We then turn to general large random networks, with density ranging from sparse trees to dense cliques. We show that information aggregates if network density is low, which motivates private discovery. In contrast, welfare attains a second-best benchmark if network density is intermediate, which allows for a quicker diffusion of discoveries.

1 Introduction

The discovery and diffusion of innovations are important drivers of long-term economic growth. This is illustrated by the seminal papers of Griliches (1957) and Coleman, Katz, and Menzel (1957) that document the spread of new technologies by farmers and doctors. From the perspective of societal welfare, discovery and diffusion are complements: Mokyr (1992) argues that both are required for sustained economic progress. From an individual strategic perspective, they are substitutes: Grossman and Stiglitz (1980) famously point out that if society aggregates information efficiently, then individual agents have no incentive to privately invest into producing information in the first place. Economic theory has made large strides in understanding information acquisition and aggregation in centralized settings such as financial markets, auctions, and collective experimentation. However, there is far less understanding about these incentives in decentralized settings, where information slowly diffuses through society. This project seeks to reconcile these perspectives in a parsimonious equilibrium model of experimentation on networks.

*We have received useful comments from Andy Atkeson, Alessandro Bonatti, Johannes Horner, and Oleg Itskhoki. We also thank seminar audiences at UCLA. Keywords: networks, experimentation, social learning. JEL codes: D83, D85.

†UCLA, http://www.econ.ucla.edu/sboard/

‡UCLA, http://www.econ.ucla.edu/mtv/

The classic paper on this topic, Bala and Goyal (1998), restricts attention to myopic, non-Bayesian agents, stating that if agents were more sophisticated "their incentives for strategic behavior (such as free riding on the information generated by other agents) would also interact with the imperfect monitoring of the rest of society in very complex ways". We propose a tractable model of experimentation on networks with forward-looking, Bayesian agents. While each agent faces a rich strategy space, her social learning curve is described by a simple function of time, and her problem reduces to choosing a single number: the total amount of individual experimentation, as captured by a cutoff time. This insight allows for a clean characterization of initial experimentation and the subsequent contagion in terms of ordinary differential equations, opening the gate to a myriad of questions about experimentation on networks.

In this project, we investigate how canonical network structures affect agents' speed of learning, their welfare, and whether society ultimately learns the underlying state of nature. In the clique, free-riding incentives are so strong that social information crowds out private experimentation one-for-one. In core-periphery networks, we show that information aggregation requires a sufficiently small core. In large random regular networks, we provide a tight comparison between directed, undirected and clustered networks. A consistent theme is that agents' distance in the network induces the information frictions that mitigate free-riding and sustain private experimentation.

In the model, a group of agents (Iris, John, Kata...) that are connected by an exogenous network (e.g. clique, core periphery, tree) can experiment with a new technology whose quality/state is high or low. Experimentation generates successes at random times iff the state is high. Agents learn from their own and neighbors' successes. This simple model captures a number of applications: Consider landowners learning about the presence of oil from nearby frackers, consumers learning about the effectiveness of a new diet from friends, or researchers learning about a new tool from colleagues.

In Section 2.1 we characterize Iris's best-response to arbitrary strategies of other agents. Observing a success perfectly reveals high quality and essentially ends the game for her. Before this time, Iris's experimentation decision is based on the expected effort of her neighbors conditional on her not having observed a success, generating her *social learning curve.* We show that Iris's potentially very complicated dynamic experimentation problem is solved by a simple *cutoff strategy:* in the absence of success, Iris experiments until some cutoff time τ and then stops. An increase in social information crowds out Iris's private experimentation, lowering τ: More past social information makes Iris pessimistic, while more future social information raises her opportunity cost of experimentation.

In Section 3 we introduce three canonical classes of networks, for which we can completely characterize learning dynamics. First, as a benchmark, we study cliques, where all agents can observe one another. The unique equilibrium features complete crowding out: the agents collectively experiment as much as a single agent would by herself. Adding agents speeds up experimentation and spreads its cost, but does not raise the amount of information produced.

Next are core-periphery networks, where core agents are connected to everyone, while peripheral agents are only connected to core agents. Core agents have more social information than peripherals, and so experiment less but have higher equilibrium value. Indeed, with enough peripherals, the core agents do not experiment at all, and only work after a success, acting as information brokers.

Third, we consider regular trees including directed trees (e.g. Twitter), undirected trees (e.g. LinkedIn), and trees of triangles that capture clustering (e.g. Facebook). Trees are interesting because they approximate large random networks, and they are tractable because neighbors' behavior is independent. We characterize social learning in the contagion phase by simple ordinary differential equations. For example, in a directed line, social information arrives at a constant rate, whereas in a directed tree with degree $d \geq 2$, the arrival rate rises over time. Using these equations, we show that the value of an agent in a directed tree with degree d is sandwiched between the value in undirected trees with degree d and $d+1$. Thus, agents prefer directed to undirected links, but even more strongly prefer one more neighbor. Similarly, we show that the value of an agent in a triangle tree with degree d is sandwiched between the value in undirected trees with degree $d-1$ and d.

In Section 6 we study how information aggregation depends on the network structure. We first ask which networks exhibit *eventual information aggregation* (EIA), meaning that social learning eventually reveals the true state. This fails in the clique, but it obtains in tree networks. Intriguingly, core-periphery networks exhibit EIA iff the periphery grows large while the core stays below a critical threshold. Intuitively, frictions are necessary in order to incentivize agents to generate their own information, which then diffuses through society.

Next we ask which networks exhibit *immediate information aggregation* (IIA), meaning social learning immediately reveals the true state. Clearly, IIA requires an exploding number of neighbors. Additionally, since IIA fully undermines agents' incentives to learn the state themselves, it requires the prior belief to support experimentation even for myopic agents. In trees and core-periphery networks, these conditions are also sufficient.

These results are new, and we are still investigating which other classes of networks admit tractable and insightful characterizations. Over the course of the next two years, we plan to develop our results into a more comprehensive theory of networked experimentation. We also plan to complement our equilibrium analysis with the mechanism design perspective of a utilitarian planner and investigate the implications for network formation. Section 7 discusses these two directions.

1.1 Literature

At the core of the paper is a "perfect good news" model of strategic experimentation with unobserved actions, observed outcomes, and private payoffs. Keller, Rady, and Cripps (2005) study a model with observed actions and private payoffs, while Bonatti and Hörner (2011) consider unobserved actions and public payoffs. Our assumption of unobserved actions and private payoffs is a natural way to model a network of oil frackers whose externalities are purely informational. The above papers characterize a unique symmetric equilibrium, where agents gradually phase out their experimentation as the public belief approaches the exit threshold. In our model, cutoff strategies are optimal and easy to characterize. This allows us to go beyond the clique and solve for equilibria in rich classes of networks.

Experimentation on networks was pioneered by Bala and Goyal (1998) who study myopic, non-Bayesian agents and provide conditions under which (i) agents reach a consensus and (ii) the agents learn the true state. Subsequent work has generalized these two limit results. Rosenberg, Solan, and Vieille (2009) considers a very general model with forward-looking Bayesian agents that encompasses strategic experimentation on networks, and shows that all agents eventually play the same action. Camargo (2014) considers a continuum of forward-looking Bayesian agents with “random sampling”, and shows that information aggregates if each action is myopically optimal for a positive measure of agents' heterogeneous priors. By focusing on good news learning, we can characterize learning dynamics at each point in time, rather than restricting attention to long-run behavior. This is important since empirical researchers must identify economic models from finite data, and because governments and firms care about *when* innovations take off, not just *if* they take off.[1]

The complexity of Bayesian updating has led some authors to consider reduced-form models of information acquisition and aggregation. For example, Galeotti and Goyal (2010) consider a local public goods game where each agent chooses a contribution level, and benefits from her neighbors' contributions. Since our agents' optimally choose a deterministic stopping time, we recover the tractability of the reduced-form models of experimentation in a model of Bayesian learning.

In seeking to characterize learning dynamics on networks, the paper is related to Board and Meyer-ter-Vehn (2021). In that paper, myopic agents sequentially choose to acquire information at a single point in time. Here, forward-looking agents simultaneously choose to acquire information at every point in time. Since agent are forward-looking, they anticipate the arrival of future social information which crowds out their private experimentation. And since they can choose repeatedly, the model gives rise to the clean distinction between an *experimentation phase* and a *contagion phase* that speaks to the discovery and diffusion of innovations. We also plan to investigate additional questions about network formation and the planner's problem that did not arise in Board and Meyer-ter-Vehn (2021).

The project also complements a recent and growing empirical literature that studies how people learn about innovations from their neighbors. Conley and Udry (2010), Banerjee et al. (2013), BenYishay and Mobarak (2019) and Beaman et al. (2021) study the spread of new production techniques and financial innovations in developing countries. Fetter et al. (2018) and Hodgson (2018) study the diffusion of fracking and oil exploration decisions. And Moretti (2011) and Finkelstein, Gentzkow, and Williams (2021) explore the adoption of new products. Such empirical analysis lacks a simple framework with forward-looking Bayesian agents that can be estimated and used for counterfactuals. This project proposes such a framework.

[1]A parallel literature considers dynamic learning games where private information is initially endowed to agents, instead of being learned over time. Gale and Kariv (2003) show that consensus must emerge when agents are Bayesian and myopic. Mossel, Sly, and Tamuz (2015) extend this result to forward-looking agents, and also show that agents eventually learn the true state if the network is not too connected (e.g. the network is undirected with bounded degree). Another classic literature considers agents who move in sequence, learning from (a subset of) prior agents. Acemoglu et al. (2011) show that society learns the true state if signals are unbounded and agents (indirectly) observe an unbounded number of agents. Mossel et al. (2020) unify many of the results in these literatures by looking at steady-state asymptotic behavior.

2 Model

Network. There are I agents $i \in \mathcal{I}$, connected by an exogenous, directed random network $\mathcal{G}$, with realizations $G \subseteq \mathcal{I} \times \mathcal{I}$ that represents who observes whom. If i observes j, we write $i \to j$ or $(i,j) \in G$, and call j a *neighbor* of i. The set of i's neighbors is $N_i(G)$.

Game. The agents seek to learn about the effectiveness of a new technology as captured by the state $\theta \in \{L, H\} = \{0, 1\}$. Time is continuous, $t \in [0, \infty)$. At time $t = 0$, agents share a common prior $\Pr(\theta = H) = p_0 \in (0, 1)$. At each time t, agent i chooses effort $A_{i,t} \in [0, 1]$ at flow cost c. This effort results in successes that arrive at rate $A_{i,t}\mathbb{I}_{\{\theta=H\}}$. Agent i's histories $h_{i,t}$ consist or her own and her neighbors' successes before t.

Payoffs. Agents receive payoff $x > c$ from their own successes. Payoffs are discounted at rate $r > 0$, so the expected discounted value equals

$$V_i = \max_{\{A_{i,t}\}_{t\geq 0}} \mathbb{E}\left[\int_0^\infty e^{-rt} A_{i,t}\left(x\mathbb{I}_{\{\theta=H\}} - c\right) dt\right] \tag{1}$$

where the expectation is taken over quality θ and information $h_{i,t}$. Our solution concept is Nash equilibrium.

- **Remark.** We impose more structure on our networks in Section 3. Much of our paper focuses on deterministic networks; we then identify the degenerate random network $\mathcal{G}$ with its realization G. When studying non-degenerate random networks $\mathcal{G}$, we assume that agents observe nothing about the realization G, not even the identity or number of their neighbos $N_i(G)$.

Notational issues

- Drop the $_0$ in p_0?
- Degree d looks odd when multipliying with t so dt

2.1 Best-Responses

In this section, we characterize the best response of a generic agent, Iris, given arbitrary strategies of other agents. After Iris observes a single success, maximal effort $A_{i,t} = 1$ is a dominant strategy, with continuation value $y := (x - c)/r$. We can thus restrict attention to times before that. So motivated, write T_i for Iris's first success time, S_i for her neighbors' first success, $\{a^{\varnothing}_{i,t}\}_{t\geq 0}$ for her *experimentation*, i.e. her effort before $\min\{T_i, S_i\}$, and

$$b_{i,t} := E\left[\sum_{j\in N_i(\mathcal{G})} A_{j,t} | t < T_i, S_i\right] \tag{2}$$

for Iris's *social learning curve*, where the expectation is taken over the random network $\mathcal{G}$ Since her experimentation is unobservable to others and own success effectively ends the game for her, Iris takes $\{b_{i,t}\}$ as given. We thus study the best response $\{a^{\emptyset}_{i,t}\}$ to $\{b_{i,t}\}$, and drop the i subscript for the rest of the section.[2,3]

When Iris has not observed a success by time-t, she uses Bayes' rule to update her belief to

$$p_t = P^{\varnothing}\left(\int_0^t (a^{\varnothing}_s + b_s)ds\right) \qquad \text{where} \qquad P^{\varnothing}(x) := \frac{p_0 e^{-x}}{p_0 e^{-x} + (1-p_0)}$$

Truncating her objective function (1) at the first observed success, we get

$$V = \max_{\{a^{\varnothing}_t\}_{t\geq 0}} \int_0^{\infty} \left[\left(a^{\varnothing}_t(x+y) + b_s y\right) p_t - a^{\varnothing}_t c\right] e^{-\int_0^t \left(r + (a^{\varnothing}_s + b_s)p_s\right)ds} dt. \tag{3}$$

Intuitively, Iris receives $x+y$ when she succeeds, y when a neighbor succeeds, and incurs effort cost of c when she works. These payoffs are discounted at the interest rate plus the success rate, $(a^{\varnothing}_s + b_s)p_s$. Equation (1) implies that Iris experiments, $a^{\emptyset}_t = 1$, when p_t exceeds the myopic threshold belief $\bar{p} := c/x$. Conversely, equation (3) implies that she does not experiment, $a^{\emptyset}_t = 0$, when p_t falls below the single-agent threshold belief $\underline{p} := c/(x+y)$. Her optimal experimentation for intermediate beliefs $p_t \in [\underline{p}, \bar{p}]$ then depends on her social learning.

We first claim that Iris uses a *cutoff strategy*. That is, she experiments maximally until some cutoff time τ and then stops, $a^{\emptyset}_t = \mathbb{I}_{\{t\leq\tau\}}$.[4] To see why, suppose she shirks at time t but works at time $t+\delta$, and consider the effect of front-loading effort ϵ from $t+\delta$ to t. This has two consequences. First, if the effort pays off, i now gets to enjoy the success earlier, raising her value by $r\delta(p_t(x+y)-c)\epsilon > 0$. Second, if one of her neighbors succeed over $[t, t+\delta]$, she ends up working at both t and $t+\delta$, raising her value by $p_t b_t \delta\epsilon(x-c) > 0$. Thus, Iris always prefers to front-load experimentation, giving rise to a cutoff time τ with cutoff belief $p_\tau \in [\underline{p}, \bar{p}]$.

Next, we characterize the optimal cutoff τ. Define Iris's (terminal) *experimentation incentives*

$$\psi_t := p_t\left(x + \frac{r}{r+\beta_t}y\right) - c, \tag{4}$$

where discounted average social learning β_t is defined implicitly by

$$\frac{r}{r+\beta_t} = r\int_t^{\infty} e^{-\int_t^s (r+b_u)du} ds. \tag{5}$$

To understand equations (4) and (5), suppose that successes from Iris's neighbors arrive at constant rate b. If she raises the cutoff from t to $t+\delta$, she gains the expected payoff from a success $p_t(x+y)\delta$,

[2]The analysis in this Section 2.1 immediately generalizes beyond our deterministic, known networks G to more general, random and/or time-varying networks.

[3]Our model is equivalent to a model where each agent can only succeed once for a payoff of $x+y$; while agents do not get to observe their neighbors' repeated successes in this model variant, this does not matter since the first success reveals $\theta = H$ perfectly .

[4]Of course, "stopping" is provisional in the sense that Iris starts to work again when she observes one of her neighbors succeed at some $t > \tau$.

but forgoes the expected benefit of future social learning $p_t(\frac{b}{r+b}y)\delta$, and incurs marginal effort cost $c\delta$. The experimentation incentives are the sum of these three effects. When the social learning curve $\{b_s\}_{s\geq t}$ is time-varying, we substitute its discounted average β_t for the constant b.

We summarize these results as follows:

Theorem 1. *Given social information $\{b_t\}$, the agent's optimal experimentation is given by the cutoff strategy $a_t^{\varnothing} = \mathbb{I}_{\{t\leq\tau\}}$ where the cutoff time τ solves $\psi_\tau = 0$.*

Theorem 1 reduces the potentially very complicated dynamic experimentation problem of a forward-looking, Bayesian agent to choosing one number τ, which is characterized rather explicitly by (4).[5] This opens the gate to a myriad of questions concerning experimentation on networks.

Before taking a stance on the network G, we can already provide useful comparative statics on Iris's value and optimal cutoff as a function of her social learning. Experimentation incentives (4) depend on past information $B_\tau := \int_0^\tau b_s ds$ via $p_\tau = P^{\varnothing}(\tau + B_\tau)$ as well as future information via β_τ, both of which crowd out her own private experimentation. Thus,

Lemma 1. *Higher social learning $\{\int_0^t b_s\}_{t\geq 0}$ increases value V and decreases the stopping time τ.*

- The optimal cutoff τ is thus maximized in the absence of social learning, $b_s \equiv 0$, i.e. the single-agent problem, where its level $\bar{\tau}$ satisfies $P^{\emptyset}(\bar{\tau}) = \underline{p}$.[6]

Lemma 1 suggests that Iris's value V and experimentation τ are negatively associated as we vary $\{\int_0^t b_s\}_{t\geq 0}$. More precisely, truncating (3) at τ we can write value V — somewhat unusually — as a function or her endogenous cutoff and her initial social information, $\mathcal{V}(\tau, B_\tau)$,[7] and show:

Lemma 2. *$\mathcal{V}(\tau, B_\tau)$ is decreasing in τ and B_τ. Thus, an agent who chooses not to work, $\tau = 0$, has higher value than an agent who chooses to work, $\tau' > 0$.*

2.2 Information Aggregation and Welfare

- The guiding questions of our paper are information aggregation and the value of social learning
- Taking the perspective of the worst-off agent, we measure information in network $\mathcal{G}$ and equilibrium $\{\tau_i\}_{i\in\mathcal{I}}$ by $\inf_i \int_0^\infty b_{i,t}dt$ and welfare by $\inf_i V_i$. Let V^{FB} be the supremum of $\inf_i V_i$, taken over all networks $\mathcal{G}$ and strategy profiles, and V^{SB} be the supremum when restricting attention to equilibria $\{\tau_i\}$.

[5] In contrast to Theorem 1, the seminal papers on strategic experimentation in the clique network, Keller, Rady, and Cripps (2005) and Bonatti and Hörner (2011), both find that agents phase out effort gradually in equilibrium, working too late as well as too little. The reason for this difference is that free-riding incentives are greater in their models. Keller, Rady, and Cripps (2005) considers public actions and Markovian strategies, so when Iris experiments more without succeeding, other agents get more pessimistic and experiment less. Bonatti and Hörner (2011) assume there can only be one success with public payoffs, ending the game; front-loading is then a waste of effort if others are about to succeed.

[6] FWIW: $\bar{\tau} = \log \frac{p_0(1-\underline{p})}{(1-p_0)\underline{p}}$.

[7] For $\tau = 0$, the expression $\mathcal{V}(\tau, B) = \mathcal{V}(0, 0)$ presumes that Iris is indifferent about experimentation at $\tau = 0$.

- In finite networks with I agents, social learning is bounded above, $\int_0^\infty b_{i,t}dt \leq \sum_{j\neq i}\tau_j \leq (I-1)\bar{\tau} < \infty$, which also bounds welfare away from V^{FB} by Lemma 1. We thus consider either sequences of finite networks, or alternatively infinite networks outright.

- We say that the sequence of I-agent networks $\mathcal{G}^{(I)}$ with associated equilibrium cutoffs, learning curves $\{b_{i,t}^{(I)}\}$, and equilibrium values $\{V_i^{(I)}\}$ *aggregates information* if

$$\lim_{I\to\infty}\inf_i \int_0^\infty b_{i,t}^{(I)}dt = \infty$$

  and that it *achieves first-best* if $\lim_{I\to\infty}\inf_i V_i^{(I)}$ equals V^{FB}, and second-best if it equals V^{SB}.

- More simply, an infinite network $\mathcal{G}$ aggregates information if $\inf_i \int_0^\infty b_{i,t}dt = \infty$ and achieves second-best if $\inf_i V_i = V^*$.

- In strongly connected networks, every agent eventually learns about the success of every other agent, and so $\int_0^\infty b_{i,t}dt = \sum_{j\neq i}\tau_j$, simplifying the test for information aggregation.

Lemma 3. *First-best is given by* $V^{FB} = p_0 y$*, and second-best by* $V^{SB} = \min\{p_0 y, p_0(x+y) - c\}$*.*

Proof for first-best:

- $p_0 y$ is the value of an agent who learns the state θ immediately, and thus clearly an upper bound for equilibrium value V_i

- This upper bound is attained by I-cliques where individual experimentation vanishes, yet collective experimentation diverges $\tau_I \to 0, I\tau_I \to \infty$, say $\tau_I = \sqrt{I}$. But in equilibrium $\tau_I = \bar{\tau}/I$, so collective experimentation is constant at $\bar{\tau}$, and welfare converges to $\mathcal{V}(0,\bar{\tau}) < V^{FB}$.

Proof for second-best

- The additional upper bound $V^{SB} \leq p_0(x+y) - c$, which is less than $p_0 y$ for $p_0 < \bar{p}$ derives from the fact that some agent j has to experiment, $\tau_j > 0$, for others to socially learn; by Lemma 2, j's equilibrium value is bounded above by $\mathcal{V}(\tau_j, B_j) < \mathcal{V}(0,0) = p_0(x+y) - c$.

- Theorems 3 and 4 give instances of networks that attain this upper bound.

3 Regular Networks

- In the last Section we studied how Iris's best response τ_i depends on her social learning $\{b_{i,t}\}_{t\geq 0}$. To close the model in equilibrium we must study how individual cutoffs $\{\tau_j\}$ aggregate into social learning curves $\{b_{i,t}\}_{t\geq 0}$. To do so, we impose more structure on the networks $\mathcal{G}$.

- As a warm up, Section 3.1 considers the clique. This illustrates how social learning crowds out private experimentation, equilibrium values are bounded away from second-best, and serves as a comparison to classic experimentation papers (e.g. Keller, Rady, and Cripps (2005), Bonatti and Hörner (2011)).
- In Section 3.2 we consider (infinite) regular tree networks and compare directed networks, undirected networks, and clustered networks. Information aggregates in all of these networks, and we study individual experimentation and welfare across these networks.
- In Section we introduce regular random networks, that encompass the clique and (in the limit) trees, as in Sadler (2020) and Board and Meyer-ter-Vehn (2021). We find that both information aggregation and welfare hinge on agents' degree as function of the size of the network, $d(I)$. Generally speaking, second-best welfare requires denser networks than information aggregation, but when the network gets too dense, i.e. $d(I)/I$ does not vanish, crowding out is sufficiently strong to also prevent second-best welfare. Information aggregatoin is additionally aided by experimentation being myopically optimal at the prior belief, $p_0 > \bar{p}$.

3.1 Clique

Recall the single-agent experimentation level, defined by $P^{\varnothing}(\bar{\tau}) = \underline{p}$.

Lemma 4. *There is a unique equilibrium with cutoffs $\tau_j = \bar{\tau}/I$ for all j.*

Crowding out is one-for-one: the I agents collectively experiment as much as a single agent would by herself. Intuitively, the agent who experiments the longest quits when the public belief reaches the single-agent threshold $\underline{p}$; since all agents are indifferent at $\underline{p}$ and agents prefer to front-load experimentation, they all experiment until that point, splitting the total experimentation evenly.[8]

- In the limit, $I \to \infty$, individual experimentation vanishes and all learning is social. Equilibrium values converge to $V^C = \mathcal{V}(0, \bar{\tau}) = \frac{p_0 - \underline{p}}{1 - \underline{p}} y < V^{SB}$.[9] Diffusion chokes off discovery too fast.

3.2 Trees

We next turn to (infinite) regular trees, where each agent has the same degree, d, as illustrated in Figure 2. In a *directed tree* $\vec{\mathcal{T}}^{(d)}$, there is at most one directed path between any two agents; this resembles users following one another on Twitter. In an *undirected tree* $\bar{\mathcal{T}}^{(d)}$, there is at most one undirected path between any two agents; this resembles the connections between acquaintances on LinkedIn. And in a *triangle tree* $\hat{\mathcal{T}}^{(d)}$, agents are connected in triangles; this resembles clusters of

[8] When thinking of information acquisition as a public goods problems, one might instead have expected the allocation of contributions across agents to be indeterminate (e.g. Galeotti and Goyal (2010)). In our setting, impatience resolves this indeterminacy.

[9] Careful: The arguments in $\mathcal{V}(0, \bar{\tau})$ mean no own experimentation $\tau = 0$ and social information equal to the single-agent problem, $B = \bar{\tau}$.

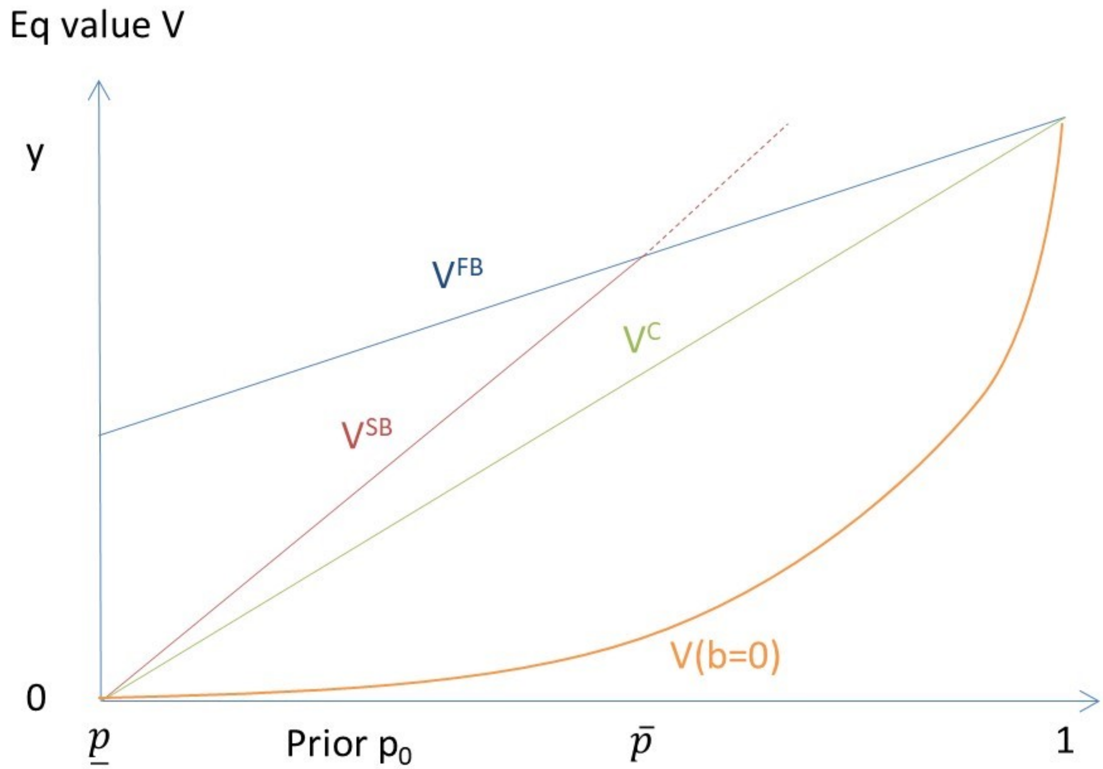

Figure 1: **Equilibrium values.** Second-best coincides with first-best for $p_0 \geq \bar{p}$, but is lower for $p_0 < \bar{p}$. Exploding cliques do not converge to second best for $p_0 \in (\underline{p}, 1)$, but do exceed the value function of a solo agent.

friends on Facebook. Trees are appealing to study since they approximate large random networks, and because the independence across neighbors (or triangles, in $\hat{\mathcal{T}}^{(d)}$) makes them highly tractable. With this in mind we ask: How does social learning depend on the structure of the tree?

Example 1 (Directed Line). Suppose agents are connected via an infinite directed line,

$$\ldots \rightarrow i \rightarrow j \rightarrow k \rightarrow \ldots$$

In a symmetric equilibrium all agents experiment for time τ. Suppose Kata succeeds in this experimentation phase, while Iris and John do not. After time τ, we enter the contagion phase where Kata and John continue to work, since they have both seen Kata's success, while Iris shirks. Eventually John also succeeds, and then all three work forever after.

To solve for the equilibrium cutoff τ, we calculate i's social information during the contagion phase, $\{b_{i,t}\}_{t \geq \tau}$. Conditioning on $\theta = H$,

$$b_{i,t} = E\left[A_{j,t} | t < T_i, S_i\right] = E\left[A_{j,t} | t < T_j\right] = \Pr\left(T_k < t | t < T_j\right) = 1 - e^{-\tau}. \tag{6}$$

The second equality uses that John is Iris's only neighbor, so $S_i = T_j$, and that his effort $A_{j,t}$ is independent of Iris's (lack of) success, $t < T_i$. The third equality relies on the observation that in

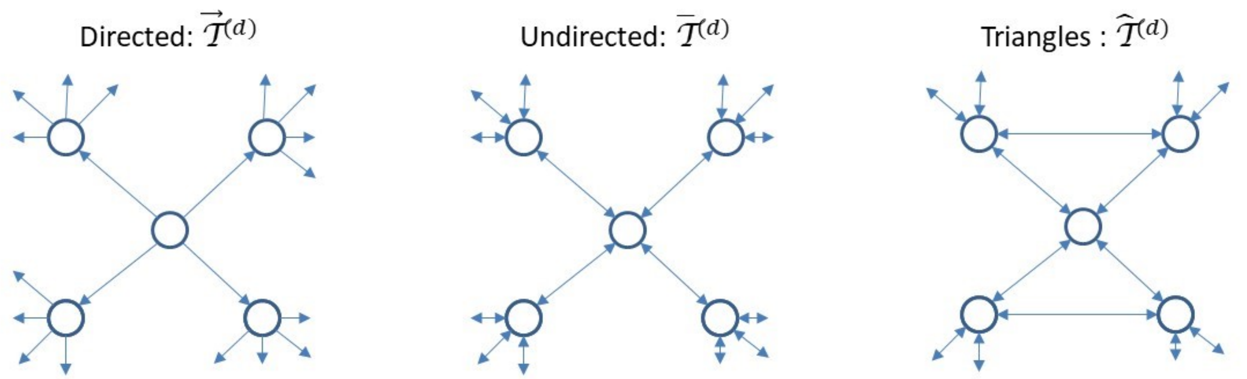

Figure 2: **Regular Trees** with degree $d = 4$.

the contagion phase, $t \geq \tau$, John works iff Kata has succeeded. The last equality uses Bayes' rule,

$$\Pr(t < T_k | t < T_j) = \frac{\Pr(t < T_j | t < T_k)\Pr(t < T_k)}{\Pr(t < T_j)} = \Pr(t < T_j | t < T_k) = e^{-\tau}.$$

Thus, social information arrives at constant rate $\beta_{i,t} \equiv b_{i,t} \equiv 1 - e^{-\tau}$. Using equation (4), the equilibrium stopping time τ solves

$$\psi_\tau = P^{\varnothing}(2\tau)\left(x + \frac{r}{r + (1 - e^{-\tau})}y\right) - c = 0. \tag{7}$$

Experimentation τ rises in the payoff x, and falls in the interest rate r and the effort cost c. △

Example 2 (Undirected Line). Now, consider the infinite undirected line

$$\ldots \leftrightarrow i \leftrightarrow j \leftrightarrow k \leftrightarrow \ldots$$

As in Example 1, conditioning on $\theta = H$,

$$b_{i,t} = 2E\left[A_{j,t} | t < T_i, S_i\right] = 2\Pr(T_k < t | t < T_j) = 2(1 - e^{-\tau}). \tag{8}$$

Using (4), the equilibrium stopping time τ is implicitly given by

$$\psi_\tau = P^{\varnothing}(3\tau)\left(x + \frac{r}{r + 2(1 - e^{-\tau})}y\right) - c = 0. \tag{9}$$

Comparing (7) and (9), agents experiment less in the undirected line, which has more sources of social information and hence greater crowding out. △

We now study trees more generally. By Lemma 1 the game has strategic substitutes, so:

Lemma 5. *Any regular tree, $\vec{\mathcal{T}}^{(d)}$, $\bar{\mathcal{T}}^{(d)}$ or $\hat{\mathcal{T}}^{(d)}$, admits a unique symmetric equilibrium with cutoff* $\tau > 0$.

- Clearly then, information aggregates in any regular tree, $\sum_i \tau_i = I\tau = \infty$.

We wish to compare stopping times and values across trees. First, we derive the arrival rate of social information in the contagion phase $t \geq \tau$ as in (6) and (8).

- In a directed tree, $\vec{\mathcal{T}}^{(d)}$, Iris's expectation of her neighbor John's effort $a_t = E[A_{j,t}|t < T_i, S_i]$ follows the ODE
$$\dot{a} = (d-1)a(1-a) \tag{10}$$
with initial condition $a_\tau = (1-e^{-d\tau})$ given by the probability that one of John's d neighbors succeeded in the experimentation phase. Subsequently, i's expectation rises because of j's expected inflow of information, but falls because of j's observed lack of success. The net effect is captured by the factor $(d-1)$ in (10): The more neighbors John has, the faster he observes a success, and the faster rises Iris's social information.

- In an undirected tree, $\bar{\mathcal{T}}^{(d)}$, Iris's expectation of John's effort follows
$$\dot{a} = (d-2)a(1-a) \tag{11}$$
with initial condition $a_\tau = (1-e^{-(d-1)\tau})$. Intuitively, i knows the backward link $j \to i$ does not produce information for j, because she conditions on not having seen a success herself. This lowers both the initial condition and the rate of increase by one degree.

- In a triangle tree, $\hat{\mathcal{T}}^{(d)}$, Iris's expectation of John's effort follows
$$\dot{a} = (d-3)a(1-a) \tag{12}$$
with initial condition $a_\tau = (1-e^{-(d-2)\tau})$. Intuitively, i knows the triangle links $j \to i, k$ do not provide information for j, because she conditions on not having seen a success by i or k herself. This lowers both the initial condition and the rate of increase by an additional degree.

To see how this difference in social learning feeds back into the equilibrium cutoff τ, we rewrite experimentation incentives (4) as a function of the degree d and neighbors' expected effort in the contagion phase, $\{a_t\}_{t\geq\tau}$
$$P^{\varnothing}((d+1)\tau)\left(x + ry\int_{s=\tau}^{\infty}\exp\left(-\int_{t=\tau}^{s}(r+da_t)dt\right)ds\right) - c = 0.$$

Substituting the solutions of the ODEs (10), (11) and (12) for $\{a_t\}_{t\geq\tau}$ allows us to compare stopping times across networks. Lemma 2 then delivers the comparative statics across values.

Theorem 2. *Equilibrium cutoff times for regular trees are ranked as follows:*
$$\hat{\tau}^{(d+2)} < \bar{\tau}^{(d+1)} < \vec{\tau}^{(d)} < \bar{\tau}^{(d)} < \hat{\tau}^{(d)}.$$

Equilibrium values are ranked in the opposite way:
$$\hat{V}^{(d+2)} > \bar{V}^{(d+1)} > \vec{V}^{(d)} > \bar{V}^{(d)} > \hat{V}^{(d)}.$$

This result provides a tight relationship between the value of different network structures and the value of extra neighbors. Intuitively, for fixed τ, the directed network $\vec{\mathcal{T}}^{(d)}$ has the same number of neighbors as the undirected network $\bar{\mathcal{T}}^{(d)}$, but more social information per neighbor since the neighbor's backward link is wasted. This extra social information provides value and crowds out experimentation. Conversely, the undirected network $\bar{\mathcal{T}}^{(d+1)}$ has the same social information per neighbor as the directed network $\vec{\mathcal{T}}^{(d)}$ but more neighbors. Again, this extra social information provides value and crowds out the agent's effort.

3.3 Regular Random Networks

- Focus on sequences of finite, undirected networks
- Constructing $\mathcal{R}(I,d)$
  - I agents each draw $d = d^{(I)} \geq 1$ link stubs
  - Randomly connect these, avoiding self-links and multi-links; drop unconnectable stubs
  - Assume for simplicity that agents don't observe anything about network realization[10]
- Open issues
  - Show that these networks look as expected: As $I \to \infty$, any agent i almost surely has d links, is part of a giant component, ...[11]
  - Notation: Better d_I? It's usually lighter than $d^{(I)}$, but becomes awkward when adding other subscripts, like time and agents, $b_{I,j,t}$
- Special cases
  - If $d^{(I)} = I - 1$, we recover the clique
  - If $d^{(I)} \equiv d$, we approach the tree $\bar{\mathcal{T}}^{(d)}$ as $I \to \infty$.
- $d^{(I)}/I$ is a measure of density of the network; cliques are maximally dense, trees are very sparse
- By symmetry, any such network $\mathcal{R}(I, d^{(I)})$ admits a unique symmetric equilibrium with cutoff $\tau^{(I)}$ and value $V^{(I)}$
- Information aggregation then obtains if $B := \lim I\tau^{(I)} = \infty$; second-best welfare convergence requires $V := \lim V_I = V^{SB}$.[12]

Conjecture 1. *Assume symmetric cutoffs τ_I (not necessarily equilibrium), asympotic information B, and degrees $\lim(d_I/\log I) = 1/\sigma \in [0,\infty]$. If $B < \infty$, the random time S_I at which a given*

[10]Introducing private signals ξ_i, say, telling agents when one of their stubs was truncated, would complicate the analysis beyound the single-dimensional cutoff τ for all agents. And omitting such private information seems practically innocuous in our regular networks, where most (asymptotically proportion one?) agents have degree d.

It would also raise the question whether we require EIA or IIA for every realization of ξ_i or in expectation. The latter would raise the problem that one could cheat to get first best for all agents i by randomly designating an exploding number (but vanishing proportion) of guinea pig agents who observe no one, but are observed by everyone.

[11]Alternatively, one might consider Erdos-Renyi networks

* Advantages: Look more realistic, and there's a larger theory on their asymptotic properties
* Disadvantages: Assuming that agents don't know their neighbors would be hokey.
* Interesting thought: It seems that information aggregation is impossible.
  · If $d^{(I)}/\log I \to 0$, the network is not connected, and the agents outside the giant component don't learn.
  · And otherwise the density of the network chokes off experimentation.

[12]We assume throughout that all limits exist, possibly equal to ∞, which is wlog subject to taking a subsequence.

agent observes the first success converges to a two-point distribution: $S=\infty$ with probability e^{-B} and $S = \sigma$ with probability $1 - e^{-B}$. If $B = \infty$, the support of S is bounded above by σ.

Argument

- Let $B < \infty$. With probability e^{-B}, nobody succeeds during the experimentation phase and thus nobody ever observes a success
- With the residual probability, there are finitely many seeds. The number of agents who have observed a success first grows exponentially at rate d_I to size $J_{\tau+t} \approx J_\tau e^{d_I t}$, until the contagion process starts running out of new nodes. At that time all remaining agents are infected asymptotically immediately because $d_I \to \infty$. The exponential growth rate means that that time converges to $\sigma = \lim(\log I/d_I)$.
- If $B = \infty$, there are infinitely many seeds with probability one, and the contagion process may reach the entire population before σ. For instance, if individual experimentation is bounded below, $\tau_I \equiv \tau > 0$, the number of seeds explodes, $\tau I \to \infty$; if additionally degrees explode $d_I \to \infty$, all agents observe a success immediately, $S \equiv 0$, irrespective of σ.[13]
- FWIW, with $B = \infty$ and $\tau_I \to 0$, S converges to a point distribution at $\hat{\sigma} := -\lim(\log \tau_I/d_I) < \sigma$.[14]

Lemma 6. *Individual experimentation vanishes, $\tau^{(I)} \to 0$, iff degrees explode, $d^{(I)} \to \infty$.*

Proof

- "If" follows by $\tau^{(I)} \leq \bar{\tau}/(d^{(I)} + 1) \to 0$
- "Only if" follows because social learning with bounded $d^{(I)}$ is bounded above by social learning in the undirected d-tree; and the boundary condition of (11) $a_\tau = 1 - \exp(-(d-1)\tau^{(I)})$ vanishes as $\tau^{(I)} \to 0$.

Implications

- Bounded degrees are sufficient for information aggregation, $I\tau^{(I)} \to \infty$.
- Exploding degrees are necessary for welfare maximization
  - With optimistic prior $p_0 > \bar{p}$, welfare maximization requires immediate information aggregation and hence an exploding degree $d^{(I)}$
  - With pessimistic prior $p_0 \leq \bar{p}$ the welfare bound equals $V^* = \mathcal{V}(0, 0)$ and $\mathcal{V}(\tau, B)$ is boundedly lower for boundedly positive τ.

Our upcoming main result shows that information aggregates when the networks are sufficiently sparse, while welfare converges to second-best when their density is intermediate.

[13] This is not an equilibirum since social learning during in the experimentation phase explodes, $d_I\tau_I \to \infty$.

[14] Indeed, $I = J_{\hat{\sigma}} = J_0 e^{d_I \hat{\sigma}} = I\tau_I e^{d_I \hat{\sigma}}$, and $\sigma - \hat{\sigma} = \lim((\log I\tau_I)/d_I) > 0$.

- For $p_0 < \bar{p}$, define $\sigma^* \in (0, \infty)$ such that socially learning the state with certainty at time σ^* (so a post-experimentation continuation value of $e^{-r\sigma^*}y$) renders an agent indifferent at $t = 0$

$$p_0 \left(x + (1 - e^{-r\sigma^*})y \right) = c. \tag{13}$$

Theorem 3. *The regular networks $\{\mathcal{R}(I, d^{(I)})\}$ converge to second-best iff $d^{(I)} \to \infty$ and $d^{(I)}/I \to 0$. Information aggregation additionally depends on the prior: If $p_0 \geq \bar{p}$ it occurs iff $d^{(I)}/I \to 0$; if $p_0 < \bar{p}$, it occurs if more strongly $\lim(d^{(I)}/\log I) \leq 1/\sigma^*$.*

Welfare:

- By Lemma 6, $d^{(I)} \to \infty$ is necessary for second-best convergence, so we assume this lower bound throughout.
- The proof (but not the result) for the upper bound $d^{(I)}/I \to 0$ depends on whether the prior p_0 is above or below the myopic threshold
- Optimism: $p_0 > \bar{p}$
  - Total learning in the experimentation phase is bounded, $(d^{(I)} + 1)\tau \in [\underline{\tau}, \bar{\tau}]$, where $P^{\emptyset}(\underline{\tau}) = \bar{p}$.[15]
  - If $\lim d^{(I)}/I = \epsilon > 0$, then total asymptotic information is bounded above by $I\tau = d^{(I)}\tau^{(I)}/\epsilon < \bar{\tau}/\epsilon$. So information does not aggregate, and welfare does not converge to second-best $p_0 y$.
  - Conversely, if $\lim d^{(I)}/I = 0$, then total asymptotic information $I\tau > I\underline{\tau}/(d^{(I)} + 1)$ diverges, and one can show more strongly that agents learn an exploding amount of information before any $t > 0$.[16,17,18]

[15] FWIW $\underline{\tau} = \log \frac{p_0(1-\bar{p})}{(1-p_0)\bar{p}}$

[16] See Section 7.2.2 in 210831_Repository.

[17] Conjecture 1 is not so useful here. With $B = \infty$, it implies immediate learning only if additionally $\lim(d^{(I)}/\log I) = 1/\sigma = \infty$.

[18] One could try to argue more strongly that (for any prior): Asymptotically, the ratio of social learning before the cutoff to total social learning equals the ratio of degree to total nodes: For every $t > 0$,

$$\lim_{I \to \infty} \frac{\tau^{(I)} d^{(I)}}{\int_0^t b_s^{(I)} ds} = \lim_{I \to \infty} \frac{d^{(I)}}{I}$$

* This is not to say that agents necessarily learn non-neighbors' information right after τ. (By Conjecture 1, immediate learning more strongly requires $\hat{\sigma} = -\lim((\log \tau_I)/d_I) = 0$). In an infinite tree they don't. But that ratio still goes to 0.
* Then, for $p_0 \leq \bar{p}$:
  - If the LHS is zero, $d\tau$ must vanish; for otherwise, agents would learn θ immediately after τ, eroding experimentation incentives, $\psi_\tau < 0$.
  - If this limit is finite, $d\tau$ must not vanish; for otherwise, social learning after τ also vanishes, but then ψ_τ approaches $p_0(x + y) - c > 0$.
* This is more elegant in that it avoids the repetition of the argument across the cases $p_0 > \bar{p}$ and $p_0 \leq \bar{p}$.
* But it's somewhat heavy algebraically, and also not so obvious how to prove it rigorously

- Pessimism: $p_0 \leq \bar{p}$
  - Since $V^* = \mathcal{V}(0,0)$, Lemma 2 implies that welfare $\mathcal{V}(\tau, d\tau)$ approximates second best iff social learning during the experimentation phase $d^{(I)}\tau^{(I)}$ vanishes.
  - We argue by contradiction: Assume $d^{(I)}/I \to \epsilon > 0$ and second-best convergence. Then, pre-cutoff learning vanishes, $d^{(I)}\tau^{(I)} \to 0$, so total learning vanishes $I\tau^{(I)} < d^{(I)}\tau^{(I)}/\epsilon \to 0$, and agent value converges to its minimum $\mathcal{V}(\bar{\tau}, 0)$, contradicting the second-best assumption.
  - For $d^{(I)}/I \to 0$ assume counterfactually that $\lim d^{(I)}\tau^{(I)} > 0$. The same argument as in the case $p_0 > \bar{p}$ implies that agents learn an exploding amount of information before any $t > 0$. For $p_0 < \bar{p}$, this chokes off information generation immediately, so $\tau_I = 0$ for large I, contradicting our assumption $\lim d^{(I)}\tau^{(I)} > 0$. For $p_0 = \bar{p}$, asymptotic immediate learning pushes the cutoff belief to the prior $p_{\tau_I} \to \bar{p} = p_0$, and so $p_{\tau^{(I)}} = P^{\emptyset}((d^{(I)}+1)\tau^{(I)})$ implies the contradiction $d^{(I)}\tau^{(I)} \to 0$.

Information aggregation

- Unlike second-best welfare, information aggregation does not require a lower bound on the degree $d \geq 1$, and obtains for all bounded $d^{(I)}$ by Lemma 6. So we henceforth assume $d_I \to \infty$.
- For $p_0 \geq \bar{p}$, the upper bound $d^{(I)}/I \to 0$ for information aggregation follows by the same arguments as for second-best welfare
  - Second-best welfare $V^{(I)} \to p_0 y$ implies that all agents perfectly learn the state at any time $t > 0$, so a fortiori in the long-run. The upper bound $d^{(I)}/I \to 0$ is thus also sufficient for information aggregation.
  - In the proof of second-best welfare, we argued that this bound is also necessary.
- Turning to $p_0 < \bar{p}$,
  - Warmup: The upper bound $d^{(I)}/I \to 0$ is still necessary for information aggregation (by the same argument) but no longer sufficient. If the network is so dense that the contagion process covers the entire network immediately, $d^{(I)}/\log I \to 1/\sigma = \infty$, information cannot aggregate. If it did, $B = \infty$, social learning would fully crowd out learning incentives, so $\psi_0 = p_0 x - c < 0$, nipping information generation in the bud.
  - Necessary: More generally, if $\lim(\log I/d_I) = \sigma < \sigma^*$ and, by ways of contradiction $B = \infty$, Conjecture 1 implies that experimentation incentives at $t = 0$ turn negative

    $$\lim \psi_{I,0} = p_0 \left(x + (1 - E[e^{-rS}])y\right) - c < p_0 \left(x + (1 - e^{-r\sigma^*})y\right) - c = 0$$

    contradicting $B = \infty$.

– Sufficient: Conversely, for $\lim(\log I/d_I) = \sigma \geq \sigma^*$ assume by ways of contradiction, $B < \infty$. Given $d_I \to \infty$ and $p_0 \leq \bar{p}$, finite asymptotic information, $B < \infty$, implies that social learning during the experimentation phase vanishes, $d_I \tau_I \to 0$, and hence $p_{\tau_I} = P^{\emptyset}((d_I + 1)\tau_I) \to p_0$. Conjecture 1 then implies that experimentation incentives at the equilibrium cutoff are strictly positive:

$$\lim \psi_{I,\tau_I} = p_0 \left(x + y(1 - (1 - e^{-B})e^{-r\sigma})\right) - c > p_0 \left(x + (1 - e^{-r\sigma^*})y\right) - c = 0.$$

This contradiction implies that asymptotic information must be infinite, $B = \infty$.[19]

3.3.1 Beyond the benchmark result, Theorem 3:

We now argue more strongly that aggregate information falls with network density, while welfare is hump-shaped. We start with B.

- Theorem 3 and its proof establish that $B = \infty$ if $\lim(d_I/\log I) \leq 1/\sigma^*$.
- For $1/\sigma^* < \lim(d_I/\log I)$ and $d_I/I \to 0$, $B < \infty$ solves

$$p_0 \left(x + y(1 - (1 - e^{-B})e^{-r\sigma})\right) = c; \tag{14}$$

  this falls in $\lim(d_I/\log I) \to 1/\sigma$ as $1/\sigma$ rises to ∞, and is then constant, solving $p_0 \left(x + e^{-B}y\right) = c$ for all $\{d_I\}$ with $\lim(d_I/\log I) \to \infty$ and $d_I/I \to 0$.
- We next argue that when $d_I/I \to \alpha \in (0, 1]$, $B < \infty$ solves

$$P^{\emptyset}(\alpha B)\left(x + e^{-(1-\alpha)B}y\right) = c. \tag{15}$$

  – By Conjecture 1, asymptotically each agent learns all information B immediately.

  – Within this vanishing time interval of learning, information from the αI direct neighbors is learnt before τ_I and enters as argument of $P^{\emptyset}$ in (15), while the residual $(1 - \alpha)I$ is learnt immediately after and enters as an opportunity cost in (15).

  – Since (15) is more sensitive to pre-cutoff learning that to post-cutoff learning,[20] a rise in α decreases the LHS of (15), which must be compensated by a fall in total information B.

  – For $\alpha = 1$, we recover the infinite clique, with total crowding out $B = \bar{\tau}$.

Turning to welfare V:

- Theorem 3 shows it is maximized for intermediate degrees, that satisfy both $d_I \to \infty$ and $d_I/I \to 0$.

[19] FWIW, as long as $\lim(d^{(I)}/\log I) = 1/\sigma > 0$, individual cutoffs must satisfy $\lim(\log \tau_I/d_I) = -\sigma^*$

[20] Lemma 4 in 210831_repository

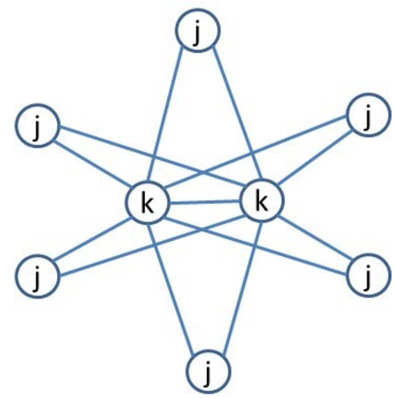

Figure 3: **Core-Periphery Network** with $K = 2$ core agents, and $J = 6$ peripherals.

- For finite d, Theorem 2 strongly suggests that V increases in d.[21]
- For $d_I/I \to \alpha > 0$, the fact that $B = B(\alpha)$ decreases implies that $V = V(\alpha)$ decreases, too.

4 Core-Periphery

- In Section 4 we study core-periphery networks, that are popular in the financial markets literature (e.g. Babus and Kondor (2018)) and arise endogenously in Galeotti and Goyal (2010).

The core-periphery network $\mathcal{CP}(J, K)$ consists of K core agents who are connected to everyone, and J peripheral agents who are only connected to core agents. See Figure 3 for an illustration. Writing τ_k and τ_j for the stopping time of core and peripheral agents in a symmetric equilibrium, we get

Lemma 7. *In the core-periphery network $\mathcal{CP}(J, K)$, core agents work less, $\tau_k < \tau_j$, and have higher values, $V_k > V_j$.*

Core agents work less than peripherals because of their greater social information. After τ_k, the core agents shirk, waiting for one of the peripheral agents to succeed. After such a success, the core agents start working, while the peripheral agents experiment until time τ_j. Finally, once a core agent succeeds, everyone learns the state is high and works. The core agents thus serve as information brokers.

- We now consider sequences of core-periphery networks with I agents, $K_I \geq 1$ of whom are in the core, $\mathcal{CP}(I - K_I, K_I)$.
- Mirroring Theorem 3, the next result shows that information aggregates for small core sizes, while welfare (of the worst-off peripheral agents) converges to second-best for intermediate core sizes.

[21] A formal proof would require to show that the equilibrium welfare is continuous as $\mathcal{R}(I, d) \to \bar{\mathcal{T}}^{(d)}$.

- For $p_0 < \bar{p}$, let $K^* \in (0, \infty)$ solve $p_0 \left(x + \frac{r}{r+K^*} y \right) = c$.

Theorem 4. *If $p_0 \geq \bar{p}$, information aggregates in core-periphery networks $\{\mathcal{CP}(I - K_I, K_I)\}$ iff $K_I/I \to 0$; convergence to second-best additionally requires $K_I \to \infty$. If $p_0 < \bar{p}$, information aggregates if more strongly* $\lim K_I \leq K^*$*; convergence to second-best occurs iff* $\lim K_I \geq K^*$ *and* $K_I/I \to 0$.

Proof

- Information aggregation for $p_0 < \bar{p}$ is Theorem 2 in 210832_Repository
- Welfare is Theorem 3(b); for $p_0 \geq \bar{p}$ this implies the information aggregation results

4.0.1 Beyond the benchmark result, Theorem 4:

We now argue more strongly that aggregate information falls with network density, while welfare is hump-shaped. We start with welfare for $\lim K_I = K < K^*$.

- Core agents learn immediately and so $b_{j,t} \equiv K$
- Thus, welfare increases in K

Next, consider information $B = \lim J\tau_j < \infty$ for pessimistic prior $p_0 < \bar{p}$ and intermediate core size, $K^* < \lim K_I$ and $K_I/I \to 0$

- $\lim K_I = K < \infty$: Core agents' experimentation starts (asymptotically) at $a_0 = 1 - e^{-B}$ and then decays at rate $\dot{a} = -Ka(1-a)$, so $p_0(x + re^{-\int_0^\infty r + Ka_t dt}) = c$. Since K is the learning speed, B falls in K.
- $\lim K_I = \infty$: Peripherals learn all information right after τ, so $p_0(x + e^{-B}y) = c$, as for regular random networks.[22]

Finally, consider welfare and information aggregation in networks with large cores, $K_I/I \to \alpha \in (0, 1]$

- Indifference condition for core agent

$$P^{\emptyset}(I\tau_k) \left(x + e^{-J(\tau_j - \tau_k)} y \right) = c.$$

- Indifference condition for peripheral agent

$$P^{\emptyset} \left(K(\tau_k + \int_{\tau_k}^{\tau_j} a_t dt) + \tau_j \right) \left(x + e^{-(J-1)\tau_j + K \int_{\tau_k}^{\tau_j} a_t dt} y \right) = c$$

  where $a_{\tau_k} = 1 - e^{J\tau_k}$ and $\dot{a} = (J - 1 - K)a(1-a)$

[22]FWIW: For finite $\lim K_I = K$, core agents' time-0 experimentation incentives are strictly lower than peripherals', so $\tau_k = 0$ and equilibrium is unique. But for $\lim K_I = \infty$, this is not obvious.

- Thus $K \int_{\tau_k}^{\tau_j} a_t dt + \tau_j = J\tau_k$
- Thus τ_k does not vanish
- $B = K\tau_k + J\tau_j$ is less than the solution B of $p_0(x + e^{-B}y) = c$ and one guesses it falls in α, reaching V^C for $\alpha = 1$, but it's not clear how to prove this.

5 Leftover: Regular Networks

5.0.1 Obsolete

- not all agents can learn the state immediately, since this would nip experimentation in the bud. Indeed, the upper bound $d^{(I)}/I \to 0$, no longer suffices for information aggregation: If $I = d^2$, say, asymptotically all social learning happens occurs right after τ, so the limit of total information solves $p_0(x + e^{-\lim I\tau_I}y) = c$ and is hence finite.
- Conversely, we know that information aggregates if d_I is bounded, by Lemma 6.[23]
- With exploding degrees $d^{(I)} \to \infty$, the proof of the welfare results allows us to restrict attention to $d^{(I)}/I \to 0$, and $d^{(I)}\tau^{(I)} \to 0$.[24]
- The limit of neighbors' first success times $S = \lim S_I$, must satisfy the equilibrium condition

$$p_0 \left(x + \left(1 - E[\exp(-rS)]\right) y\right) = c \tag{16}$$

THEOREM: Value $V = \lim V^{(I)}$ is

- (a) increasing in $\lim d^{(I)}$ and below second-best when $\lim d^{(I)} < \infty$
- (b) constant, equal to second-best V^* when $\lim d^{(I)} = \infty$ and $\lim d^{(I)}/I = 0$
- (c) decreasing in $\lim d^{(I)}/I$ and below second-best when $\lim d^{(I)}/I \in (0, 1]$, reaching V^C when $\lim d^{(I)}/I = 1$.
- If $p_0 < \bar{p}$, total information $B = \lim I\tau^{(I)}$ is
- (a) infinte if $\lim d^{(I)}/\log I < 1/t^{**}$ where t^{**} solves

$$p_0 \left(x + (1 - e^{-rt^{**}})y\right) = c$$

- (b) decreasing in $\lim d^{(I)}/\log I = 1/t^* \geq 1/t^{**}$ defined by $p_0 \left(x + (1 - (1 - e^{-B})e^{-rt^*})y\right) = c$

[23] We also know that information aggregates in infinite d-trees and their limit as as $d \to \infty$, which corresponds to the double-limit $\lim_{d\to\infty} \lim_{I\to\infty} \mathcal{R}(I, d)$.

[24] If $d^{(I)}/I \to \epsilon > 0$, total information is bounded by $I\tau^{(I)} \leq \bar{\tau}/\epsilon < \infty$.

If $d^{(I)}\tau^{(I)} \to \epsilon > 0$, then $d^{(I)}/I \to 0$ implies that the state is revealed immediately after τ, undermining the incentives to experiment, since $p_\tau < p_0 < \bar{p}$.

- (c) constant, defined by to $p_0\left(x+e^{-B}y\right)=c$ if $\lim d^{(I)}/\log I=\infty$ and $\lim d^{(I)}/I=0$.

- (d) decreasing in $\lim d^{(I)}/I=\alpha\in(0,1]$, defined by $P^{\emptyset}(\alpha B)\left(x+e^{-(1-\alpha)B}y\right)=c$, reaching $\bar{\tau}$ when $\lim d^{(I)}/I=1$.

- If $p_0\geq\bar{p}$, we have $B=\infty$ also in cases (b) and (c), that is whenever $\lim d^{(I)}/I=0$.

Previously

- This reduces the problem to a pure "contagion on networks" problem that may have an off-the-shelf solution in the literature: At what time S_I does a contagion process with seeding probability τ_I, so $s\sim f(s|\tau_I I)$ seeds, and neighbor-infection rate 1 reach a given agent?

- The growth rate of the expected number of "infected" agents J_t (i.e. agents i with $T_i\leq t$) is bounded above by $d-1$ and the expected number of seeds at time τ equals $I\tau$, so $J_{\tau+t}\leq\sum_s e^{(d-1)t}sf(s|\tau_I I)$.

- Thus, the chance of observing a neighbor succeed before time t is bounded above via

$$\Pr(S\leq\tau+t)\leq dJ_{\tau+t}/I\leq\sum_s de^{-t}e^{dt-\log I}sf(s|\tau_I I)\tag{17}$$

- To see that $I\tau$ must explode if $d/\log I\to 0$, assume to the contrary that $I\tau$ is bounded. Then the RHS vanishes for any $t>0$ as $I\to\infty$. Thus $E[\exp(-rS)]$ vanishes, contradicting (16), recalling $p_0>\underline{p}=c/(x+y)$.

- Conversely, if $d/\log I\to\infty$ is bounded

5.0.2 Obsolete (and probably false)

- We conjecture that information aggregation requires more strongly that $d(I)^{d(I)}/I\to 0$, and that information aggregation must obtain if $d(I)^{f(d(I))}/I$ is bounded away from 0 for some f with $f(d)/d\to\infty$.

- The key argument is that the time it takes for a success to transmit through a tree with span d and depth $r(d)$ has expectation $r(d)/d$, and converges in probability to $\lim r(d)/d$.

- If $d(I)^{d(I)}/I\to 0$, most nodes are

- If for some $M<\infty$ and all large I we have $I\leq Md(I)^{d(I)}$,
  - a non-vanishing fraction of nodes is contained in a tree of span $d^{(I)}$ and depth $d^{(I)}$ around any given agent i.
  - The expected time for a success to percolate through this tree to i is bounded above and below:

        * The transmission time at each step is exponential with parameter and expectation $1/d^{(I)}$
        * And there are $d^{(I)}$ such steps
        * Indeed, as $d^{(I)} \to \infty$, the total expected time should converge to 1
    - Since information aggregation requires an exploding number of successes before τ, this finite transmission time undermines information generation

- Conversely if $I \geq \epsilon d(I)^{f(d(I))}$
    - an exploding proportion of nodes are further than $f(d^{(I)})$ away from i
    - the expected transmission time from these nodes $f(d^{(I)})/d^{(I)}$ diverges

5.1 Next Steps

To bring the tree networks in this section closer to networks observed in practice and allow for empirical and policy analysis, one might want to introduce heterogeneity pertaining to agents' degree d, costs c, or benefits x. Preliminary attempts in this direction indicate that the equilibrium analysis readily accommodates such heterogeneity, but comparative statics are complicated by the higher dimensionality of the policy space $\tau = \tau(d, c, x)$. We also plan to study equilibria in time-varying networks. Specifically, we conjecture that re-sampling links should improve social learning by avoiding the redundancy of observing the same neighbor repeatedly.

6 OLD: Information Aggregation

We now study how the network structure affects society's ability to aggregate information. We propose two notions of information aggregation. Consider a sequence of exploding sets of agents, $\mathcal{I}^{(n)} \subseteq \mathcal{I}^{(n+1)}$, $\mathcal{I} := \bigcup_{n\geq 1} \mathcal{I}^{(n)}$ and $|\mathcal{I}| = \infty$, with associated networks and equilibrium stopping times, $(\mathcal{I}^{(n)}, G^{(n)}, \{\tau_i^{(n)}\})$. The sequence exhibits *eventual information aggregation (EIA)* if every agent eventually receives perfect social information,

$$\forall i : \lim_{n\to\infty} \int_0^\infty b_{i,s}^{(n)} ds = \infty.$$

This notion is familiar from discrete-time games with repeated actions (e.g. Bala and Goyal (1998), Golub and Jackson (2010)).

The sequence exhibits *immediate information aggregation (IIA)* if every agent immediately receives perfect social information,

$$\forall i, t > 0 : \lim_{n\to\infty} \int_0^t b_{i,s}^{(n)} ds = \infty.$$

This notion is familiar from static trading games (e.g. Grossman and Stiglitz (1980)) and continuous-time games of social learning (e.g. Board and Meyer-ter-Vehn (2021)).

6.1 Eventual Information Aggregation

Our networks in Section 3 offer a sharp contrast:

- The n-agent cliques fail EIA as $n \to \infty$. Crowding out is one-for-one, so the total amount of social information is $\int_0^\infty b_{i,s}^{(n)} ds = (n-1)\bar{\tau}/n$ is bounded above (and converges to) $\bar{\tau}$.

- Any infinite regular tree exhibits EIA.[25] For example in the directed line, $b_{i,t}$ equals 1 for $t \le \tau$ and $1 - e^{-\tau}$ for $t \ge \tau$, so $\int_0^\infty b_{i,s}^{(n)} ds = \infty$.

In the core-periphery network, we can delineate the boundaries of EIA as a function of the core size:

Theorem 5. *Core-periphery networks $\mathcal{CP}(J^{(n)}, K^{(n)})$ satisfy EIA iff the periphery explodes while the core is bounded, i.e. $J^{(n)} \to \infty$ and there exists K with*

$$p_0 \left(x + \frac{r}{r+K} y \right) \ge c \tag{18}$$

such that $K^{(n)} \le K$ for almost all n.

Intuitively, core agents jointly generate less information than a solo agent, $K^{(n)}\tau_k \le \bar{\tau}$, so information aggregation relies on peripherals. Since their social information is bounded above by $\beta \le K$, condition (18) guarantees that initial experimentation incentives (4)

$$\psi_{j,0}^{(n)} = p_0 \left(x + \frac{r}{r + \beta_{j,0}^{(n)}} y \right) - c$$

are positive. Aggregate experimentation by peripherals then diverges, $J^{(n)}\tau_j^{(n)} \to \infty$, implying EIA.

Underlying these arguments for our three classes of networks is the idea that sparse networks induce the information frictions which mitigate crowding out and sustain information aggregation. For low K, it takes a long time for information to pass from one peripheral agent to another (as in the infinite line), so each one generates some original information, even if $J = \infty$. But for high K, they crowd out each other's work, as in the clique, and information aggregation fails.

6.2 Immediate Information Aggregation

Immediate information aggregation requires that each agent has an exploding amount of information immediately at hand. Such information crowds out all experimentation incentives, and chokes off information generation for priors below the myopic threshold $p_0 < \bar{p} = c/x$. Thus,

Lemma 8. *Exploding degrees, $|N_i| \to \infty$ for all i, and an optimistic prior, $p_0 \ge \bar{p}$, are jointly necessary for IIA.*

In trees and core-periphery networks, these conditions are also sufficient:

[25] Instead of jumping to the limit by directly studying the infinite tree, we could instead alternatively consider an infinite sequence of finite tree networks, starting with the root agent, adding her d neighbors, then adding their d^2 neighbors, and so on. The equilibria of these finite trees converge to the symmetric equilibrium of the infinite tree.

Theorem 6. *Suppose $p_0 \geq \bar{p}$. Then IIA holds in*

(a) Regular trees with exploding degree, $\lim_{d\to\infty} \vec{\mathcal{T}}^{(d)}$, $\bar{\mathcal{T}}^{(d)}$ *or* $\hat{\mathcal{T}}^{(d)}$.

(b) Core-periphery network with a large core and larger degree, $\lim_{K\to\infty} \lim_{J\to\infty} \mathcal{CP}(J,K)$.

6.3 Next Steps

Moving forward, we hope to obtain tight conditions for eventual and immediate information aggregation (in the spirit of Theorems 5 and 6) in more comprehensive classes of networks. Small steps in this direction suggest that the failure of EIA in the clique extends to "finite sets of interconnected, exploding cliques" including "replica networks"; the success of EIA should extend from the infinite line to networks with "infinite diameter and bounded degree".[26] More ambitiously, we might extend our analysis to random networks. For example, consider the regular configuration model with degree d and $I(d)$ total agents; these networks encompass the tree ($I = \infty$, d finite) and the clique ($I = d$). We conjecture that IIA obtains as $d \to \infty$ iff agents' direct neighbors comprise of a small part of society $d/I(d) \to 0$.

7 Extensions

In this section, we sketch two avenues for future work. First, we ask which networks arise endogenously when agents can buy links. To this end, we conjecture that cliques arise when links are cheap, whereas stars can arise when they cost a bit more. Our preliminary analysis suggests that an increase in link costs can make agents better off.

Second, we introduce a utilitarian planner who chooses the experimentation cutoffs on behalf of the agents. Since the planner internalizes the benefits of private experimentation on other agents, she generally wants to raise cutoffs. We conjecture that the clique network maximizes this divergence between private an social incentives, while in core-periphery networks the planner may want to raise cutoffs for core agents but lower them for peripherals.

7.1 Network Formation

Consider the baseline model and suppose that, before the game, agents simultaneously choose to link to one another. Galeotti and Goyal (2010) studies such network formation, modeling the subsequent information discovery and dispersion in reduced form as a local public goods game. They show that equilibrium networks exhibit the "Law of the Few" in that: (i) a small subset of agents acquire information, and (ii) a majority of individuals get their information from this group of informed agents. We wish to revisit this question using a dynamic, Bayesian model of experimentation. Preliminary analysis suggests that part (ii) of the Law of the Few obtains, while part (i) does not.

[26] There are other interesting examples of networks. For example, unlike Bala and Goyal (1998), we obtain EIA in their "royal family" example, and can construct networks with infinitely many "locally independent" agents where EIA fails.

There are two large hurdles in solving such a model. First, in the network formation stage, agent face a combinatorial portfolio choice problem when choosing with whom to link. Second, in the experimentation stage, we need to characterize equilibria on every network. We skirt these problems by guessing a natural equilibrium network where we can ignore correlation across agents and rank them by their individual informativeness; we then verify there is no deviation.[27]

We conjecture that the following networks emerge as equilibria

- When κ is "high", agents form a star, with the central agent acting as "information broker" and doing no work.
- When κ is "middling", agents form a core-periphery with $K > 1$ core agents doing no work.[28]
- When κ is "low", agents form a clique, with everyone doing a little work.

As in Galeotti and Goyal (2010), agents link to the most informed core agents. But in contrast, core agents do not generate original information; rather they aggregate information generated by peripherals.

This result implies that linking costs κ may be socially beneficial by preventing over-clustering. Indeed, assuming (18) fails for $K = 1$, social learning in the star network crowds out experimentation by each peripheral agent, $\tau_j \to 0$ as $J \to \infty$, and so equilibrium values converge to $\mathcal{V}(0, 0)$, cf Lemma 2. But in the clique network, which must arise for $\kappa = 0$, equilibrium values converge to $\mathcal{V}(0, \bar{\tau})$, which is less than $\mathcal{V}(0, 0)$ by Lemma 2. Intuitively, agents link excessively because they ignore the equilibrium impact of their links on others' experimentation.

The above arguments showcase the tractability of our baseline model, but are only a starting point for an analysis of network formation, that would seek to characterize all equilibrium networks and propose a way to select from this set. One might also try to relax some of the more pragmatic modeling assumptions, e.g. that links are undirected yet costs are unilateral. We are still at the early stages of thinking about this, so it is a little unpredictable where this will go.

7.2 Planner's problem

Consider a utilitarian planner who takes the network G as given and chooses effort $A_{i,t}$ on behalf of the agents. As in traditional experimentation papers (e.g. Keller, Rady, and Cripps (2005)), the planner internalizes the informational spillovers of private experimentation and wants to correct for agents' free-riding. How does this discrepancy between private and social perspective depend on the network? To get a sense of the planner's problem consider a simple example:

[27] Formally, we follow Galeotti and Goyal (2010) by assuming that agents do not observe the realized network, so do not react to deviations in linking strategies by others; we also adopt their assumptions that links are undirected but that the link cost κ only has to be paid by the originating agent.

[28] More precisely, the network we have in mind has a large group of J "peripheral" agents that all link to K "core" agents, but unlike in Section 4, the core agents do not link to each other. One might alternatively call this structure an (asymmetric) bipartite graph.

Example 1 (Directed Line). Suppose the planner chooses a symmetric cutoff $\tau > 0$ (which one would need to show is optimal). At τ, the marginal social benefit of i's experimentation equals

$$\psi_\tau^{FB} = p_\tau \left(x + y + e^{-\tau} Z - \frac{b}{r+b} Z \right) - c$$

where $b = 1 - e^{-\tau}$ is the arrival of social information, and $Z = \frac{r+1}{r+(1-e^{-\tau})} y$ is the benefit to $\{i-1, i-2, \ldots\}$ when i succeeds and $i-1$ has not succeeded yet. In comparison, the private benefit is given by (7). $\triangle$

We conjecture that private and social incentives are less aligned in more connected networks. Intuitively, in a highly connected network the planner benefits from the fast communication of information, raising her choice of experimentation. However, in equilibrium, social information crowds out private information, leading to less information in the long run (e.g. see the clique in Section 3.1). One challenge with this conjecture is to formalize the notion of the discrepancy between private and social outcomes.

Another question is how experimentation is distributed in the planner's problem as compared to the equilibrium. For example, consider the two-person network $i \to j$. The planner would raise τ_j above $\bar{\tau}$ because of its positive externality on i, but then *lower* τ_i to its new equilibrium value. Indeed, it is not obvious whether aggregate experimentation $\tau_i + \tau_j$ must rise, especially if there are multiple copies of i. Turning to richer networks, one wonders about the planner's solution for the star network. We conjecture that core agents do most of the work in the planner's problem, whereas peripheral agents do most of the work in equilibrium.

Studying the planner's problem raises broader questions about how the government should reward innovators. Should the government "push" innovation by lowering the cost of innovation, c. Or should it "pull" innovation by raising the reward for success x? Following Galeotti, Golub, and Goyal (2020), one could analyze the effect of small changes in payoffs on incentives and equilibrium. This analysis could inform current policy debates about how to reward innovators with intellectual property and monopoly rents, that have greatly contributed to the current levels of inequality.

References

Acemoglu, D., M. A. Dahleh, I. Lobel, and A. Ozdaglar (2011): "Bayesian Learning in Social Networks," *Review of Economic Studies*, 78(4), 1201–1236.

Babus, A., and P. Kondor (2018): "Trading and information diffusion in over-the-counter markets," *Econometrica*, 86(5), 1727–1769.

Bala, V., and S. Goyal (1998): "Learning from Neighbours," *Review of Economic Studies*, 65(3), 595–621.

Banerjee, A., A. G. Chandrasekhar, E. Duflo, and M. O. Jackson (2013): “The Diffusion of Microfinance,” *Science*, 341(6144), 1236498.

Beaman, L., A. BenYishay, J. Magruder, and A. M. Mobarak (2021): “Can Network Theory-based Targeting increase Technology Adoption?,” *American Economic Review*, 111(6), 1918–43.

BenYishay, A., and A. M. Mobarak (2019): “Social Learning and Incentives for Experimentation and Communication,” *Review of Economic Studies*, 86(3), 976–1009.

Board, S., and M. Meyer-ter-Vehn (2021): “Learning Dynamics in Social Networks,” *Econometrica*, forthcoming.

Bonatti, A., and J. Hörner (2011): “Collaborating,” *American Economic Review*, 101, 632–663.

Camargo, B. (2014): “Learning in Society,” *Games and Economic Behavior*, 87, 381–396.

Coleman, J., E. Katz, and H. Menzel (1957): “The Diffusion of an Innovation among Physicians,” *Sociometry*, 20(4), 253–270.

Conley, T. G., and C. R. Udry (2010): “Learning about a New Technology: Pineapple in Ghana,” *American Economic Review*, 100(1), 35–69.

Fetter, T. R., A. L. Steck, C. Timmins, and D. Wrenn (2018): “Learning by Viewing? Social Learning, Regulatory Disclosure, and Firm Productivity in Shale Gas,” Discussion paper, National Bureau of Economic Research, No. 25401.

Finkelstein, A., M. Gentzkow, and H. Williams (2021): “What Drives Prescription Opioid Abuse? Evidence from Migration,” Working paper, Stanford University.

Gale, D., and S. Kariv (2003): “Bayesian Learning in Social Networks,” *Games and Economic Behavior*, 45(2), 329–346.

Galeotti, A., B. Golub, and S. Goyal (2020): “Targeting Interventions in Networks,” *Econometrica*, 88(6), 2445–71.

Galeotti, A., and S. Goyal (2010): “The Law of the Few,” *American Economic Review*, 100(4), 1468–92.

Golub, B., and M. O. Jackson (2010): “Naive Learning in Social Networks and the Wisdom of Crowds,” *American Economic Journal: Microeconomics*, 2(1), 112–49.

Griliches, Z. (1957): “Hybrid corn: An Exploration in the Economics of Technological Change,” *Econometrica*, pp. 501–522.

Grossman, S. J., and J. E. Stiglitz (1980): “On the Impossibility of Informationally Efficient Markets,” *American Economic Review*, 70(3), 393–408.

Hodgson, C. (2018): "Information Externalities, Free Riding, and Optimal Exploration in the UK Oil Industry," Working paper, Yale University.

Keller, G., S. Rady, and M. Cripps (2005): "Strategic Experimentation with Exponential Bandits," *Econometrica*, 73(1), 39–68.

Mokyr, J. (1992): *The Lever of Riches: Technological Creativity and Economic Progress.* Oxford University Press.

Moretti, E. (2011): "Social Learning and Peer Effects in Consumption: Evidence from Movie Sales," *Review of Economic Studies*, 78(1), 356–393.

Mossel, E., M. Mueller-Frank, A. Sly, and O. Tamuz (2020): "Social Learning Equilibria," *Econometrica*, 88(3), 1235–1267.

Mossel, E., A. Sly, and O. Tamuz (2015): "Strategic Learning and the Topology of Social Networks," *Econometrica*, 83(5), 1755–1794.

Rosenberg, D., E. Solan, and N. Vieille (2009): "Informational Externalities and Emergence of Consensus," *Games and Economic Behavior*, 66(2), 979–994.

Sadler, E. (2020): "Diffusion Games," *American Economic Review*, 110(1), 225–270.