Experimentation in Networks*

Simon Board[†] and Moritz Meyer-ter-Vehn[‡]

November 8, 2022

Abstract

We introduce a model of strategic experimentation on social networks where forward-looking agents learn from their own and neighbors' successes. In equilibrium, private discovery is followed by social diffusion. Social learning crowds out own experimentation, so total information decreases with network density; we determine density thresholds below which agents asymptotically learn the state. In contrast, agent welfare is single-peaked in network density, and achieves a second-best benchmark level at intermediate levels that achieve a balance between discovery and diffusion. We also show how learning and welfare differ across directed, undirected and clustered networks.

1 Introduction

The discovery and diffusion of innovations are key 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 prices aggregate information efficiently, then individual agents have no incentive to privately generate such information. Economic theory has made large strides in understanding information

*We have received useful comments from Andy Atkeson, Alessandro Bonatti, Roberto Corrao, Krishna Dasaratha, Glenn Ellison, Drew Fudenberg, Ben Golub, Yannai Gonczarowski, Johannes Hörner, Matt Jackson, Ben Jones, Emir Kamenica, Andreas Kleiner, Oleg Itskhoki, Phil Reny, and Basil Williams. We also thank seminar audiences at ASU, Cambridge-INET, CETC 2022, Chicago, CNSE 2022, Columbia, Haas, Harvard/MIT, Monash, HKUST, Kellogg, LA Theory Day, SAET 2022, University of Arizona, UT Austin, UCLA, VSET (Bristol/Carlos III/Essex), Wisconsin, Yale. Keywords: networks, experimentation, social learning. JEL codes: D83, D85.

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

[‡]UCLA, http://www.econ.ucla.edu/mtv/

acquisition and aggregation in centralized settings such as financial markets, auctions, and collective experimentation. These incentives are less well understood in decentralized settings, where information slowly diffuses through society. This paper seeks to reconcile these forces in a parsimonious equilibrium model of experimentation on networks.

The classic paper on this topic, Bala and Goyal (1998), restricts attention to myopic, non-Bayesian agents, shortcutting strategic considerations and allowing them to solve the model as a sequence of static decision problems. In contrast, our agents are forward-looking, and fully Bayesian, so both past and future social learning crowds out private experimentation. The key simplifying assumption is that agents learn via perfect good news events: This reduces each agent's problem to choosing a deterministic cutoff time, with social learning described by ordinary differential equations, opening the gate to a myriad of questions about experimentation on networks.

We use this new approach to study how asymptotic information and welfare depend on network density, as measured by either the size of the core in core-periphery networks or by the degree in regular random networks. For either measure we show that agents' asymptotic information decreases monotonically in network density and they learn the truth when the network is sufficiently sparse. In contrast, welfare is single-peaked in network density and attains a second-best welfare benchmark when density is intermediate; such networks both encourage generation of information and quickly diffuse the discoveries. Finally, we provide a tight comparison between directed, undirected and clustered networks. Collectively, these results paint a clear picture about learning dynamics, information aggregation, and welfare in networks when agents are forward-looking.

In the model, a group of I agents (Iris, John, Kata...) are connected by an exogenous network (e.g. clique, tree, core-periphery). They can each experiment with a new technology whose state is high or low; experimentation generates successes at random times iff the state is high. Agents learn from own and neighbors' successes but do not observe neighbors' actions. This simple model captures a number of applications: Consider farmers learning about the success of a new crop from neighbors, doctors learning about a new drug from colleagues, or landowners learning about the presence of oil from nearby frackers.

In Section 3.1, we first characterize Iris's best-response to arbitrary strategies of other agents. Observing a success perfectly reveals the high state and essentially ends the game for her. Before this time, Iris's experimentation decision is based on her *social learning curve*, i.e. the expected effort of her neighbors. We show that Iris's dynamic experimentation problem is solved by a simple *cutoff strategy:* In the absence of success, Iris stops experimenting at some cutoff time τ_i. An increase in social information crowds out Iris's private experimentation, lowering her cutoff time: Unsuccessful past social learning makes Iris pessimistic, while future social information lowers the information value of own experimentation.

Section 3.3 illustrates in examples how to aggregate individuals' cutoff times into social learning curves for their neighbors. In the clique network, where all agents 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 generated. In the line network, an initial experimentation phase is followed by a contagion phase. Agents asymptotically learn the state, but welfare is limited by the slow rate of information diffusion.

In Section 4, we study the effect of network density on asymptotic information and welfare. Specifically, we consider two canonical types of networks (regular random networks and core-periphery) as $I \to \infty$. To study aggregate information, define the *asymptotic information* to be the total information created by society; there is *asymptotic learning* if asymptotic information is unbounded, meaning that the agents learn the state. To study welfare, we propose a second-best benchmark that provides an upper bound on equilibrium utility (of the worst-off agent) across all networks. Neither the clique nor the line attain this benchmark: The former generates too little information, while the latter diffuses information too slowly. Nevertheless, we show below that the benchmark is attained by some networks.

We first study large regular random networks with degree n^I. This model encompasses sparse trees, where $n^I \equiv n$, and dense cliques, where $n^I/I \to 1$. Theorem 1 shows that asymptotic information decreases in network density, and asymptotic learning obtains if density is below a threshold; specifically the agents fully learn if the *time-diameter* (the typical time for information to travel between two agents) exceeds σ^*, which is the time such that perfectly learning the state at $t = \sigma^*$ renders agents indifferent about experimentation at $t = 0$. Welfare is single-peaked in network density and attains the second-best benchmark if $n^I \to \infty$ and $n^I/I \to 0$. Intuitively, asymptotic learning requires sparsity to sustain private experimentation incentives; high welfare additionally requires density to promptly diffuse news across society.

To study the role of network position on experimentation incentives we next turn to core-periphery networks, where K^I *core* agents are connected to everyone while $I - K^I$ *peripheral* agents are only connected to core agents. Such networks are used to describe financial markets, with dealers or banks as core agents (e.g. Li and Schürhoff (2019)). In equilibrium, core agents have more social information than peripherals, so experiment less and have higher utility. While core agents experiment little themselves (if at all), they serve an important role as information brokers connecting the peripherals. As $I \to \infty$, asymptotic learning and welfare exhibit similar properties to large random networks, with core size substituting for the degree. Theorem 2 shows that asymptotic information decreases in network density, and asymptotic learning obtains if K^I remains below a threshold κ^*. Welfare is single-peaked in network density and attains the second-best benchmark if K^I exceeds κ^* and $K^I/I \to 0$. The threshold κ^* renders peripherals indifferent about experimentation at $t = 0$ when κ^* core agents work forever.

Our analysis of large random networks and core-periphery networks points to a fundamental tradeoff between social learning and welfare. These goals are often thought to be aligned: Hayek (1945) famously emphasizes the importance of information aggregation for allocative efficiency. In our model, agents must be incentivized to acquire information, so the fast diffusion required for second-best welfare can lower total information. Indeed, for core-periphery networks the two goals are mutually exclusive.

Our two families of networks differ in their network structure and thus exhibit different social learning dynamics. In large random networks, the typical pair of agents has distance $\log I/\log n^I$; as the degree grows, social learning occurs in a single burst at a fixed time σ. In contrast, in core-periphery networks, all peripheral agents are two links apart; social learning occurs as the core agents slowly transmit the information created by the periphery. Cumulative social learning curves are thus convex for large random networks but concave for core-periphery networks, as illustrated in Figures 4 and 6.

In Section 5, we study different types of links in the context of regular tree networks, including directed trees (e.g. Twitter), undirected trees (e.g. LinkedIn), and trees of triangles that capture clustering (e.g. Facebook). As discussed above, trees approximate large random networks. Trees are also highly tractable because neighbors' behavior is independent; this allows us to 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 $n \geq 2$, the arrival rate rises over time. Theorem 3 provides tight bounds on the utility of agents across different networks. The utility of an agent in a undirected tree with degree n is sandwiched between her value in directed trees with degree $n-1$ and n. Thus, agents prefer directed to undirected links, but even more strongly prefer to be in a tree with one more neighbor. Similarly, the utility of an agent in a triangle tree with degree n is sandwiched between her value in undirected trees with degree $n-1$ and n. Link types are thus of second-order importance for dense networks, where the distinction between, say, $n = 74$ friends and $n' = 78$ friends becomes blurry. Collectively, these results demonstrate the qualitative and quantitative importance of network structure.

1.1 Literature

At the core of the paper is a "perfect good news" model of strategic experimentation with private actions and payoffs. In the context of a clique, Keller, Rady, and Cripps (2005) study a good-news model with observed actions and private payoffs, Bonatti and Hörner (2011) consider good-news model with unobserved actions and public payoffs, and Bonatti and Hörner (2017) consider a bad-news model with unobserved actions and private payoffs. In all of these papers, agents use mixed strategies. Specifically, in the first two papers, agents gradually phase out their experimentation as the public belief approaches the exit threshold. In our model, agents use simple cutoff strategies that allow us to go beyond the clique and solve for equilibria in rich classes of networks. We also think that the assumptions of unobserved actions and private payoffs is a natural way to model a network of farmers, doctors or oil frackers whose externalities are purely informational.

Observational learning on networks was pioneered by Bala and Goyal (1998) who study myopic, non-Bayesian agents and provide conditions on the network under which (i) agents reach a consensus and (ii) the agents learn the state.[1] Subsequent work has generalized these two limit results in models with forward-looking, Bayesian agents who incorporate the future value of information when choosing to experiment. Rosenberg, Solan, and Vieille (2009) consider a very general model that encompasses strategic experimentation on networks, and shows that all agents eventually play the same action. Camargo (2014) considers a continuum-agent model 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 because agents care about when innovations diffuse and not just if they diffuse; indeed, this consideration underlies the contrast between sparse networks that aggregate information and the denser networks that maximize welfare.[2]

Most closely related to our model, Salish (2015) embeds a discrete-time version of Keller, Rady, and Cripps's (2005) strategic experimentation model in a network. Neighbors observe each others' actions, which thus signal successes of second neighbors; Salish side-steps such

[1] Sadler (2020b) characterizes outcomes more completely in Bala-Goyal's model with Brownian learning.

[2] 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 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 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.

signaling by introducing an additional learning channel, whereby successes are automatically transmitted across the network, one link per period. The paper shows that experimentation tends to phase out over time, and that ring and star networks aggregate more information than the clique. In contrast, our best-responses are determined by simple cutoffs, allowing us to characterize equilibrium, aggregate information, and welfare.

The complexity of Bayesian updating has led some authors to consider reduced-form models of information acquisition and aggregation. For example, Bramoullé and Kranton (2007) and 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. The different models give rise to different economic forces: The forward-looking agents in this paper anticipate the arrival of future social information which crowds out their private experimentation, and the repeated choices gives rise to the clean distinction between an experimentation phase and a contagion phase. This paper also focuses on a different question: How does aggregate information and welfare change with network density? The results in Section 5 correspond most closely to our prior paper, where we studied different types of links in the configuration model.[3]

The project also complements a 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 (2021) 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 paper proposes such a framework.

[3]More specifically, in Board and Meyer-ter-Vehn (2021) we consider a more general configuration model with multiple types. We show that social learning was greater in directed networks than comparable undirected networks, and greater in tree networks than than comparable triangle networks. Section 5 only considers regular networks but provides a tighter characterization of the value of link type.

2 Model

Network. Agents $i = \{1, ..., I\}$ are connected by a network $g \subseteq I^2$ that represents who observes whom. If i (Iris) observes j (John), we write $i \to j$ or $(i,j) \in g$, and call j a *neighbor* of i. The set of Iris's neighbors is $N_i(g)$. The network may be directed or undirected. It may be deterministic or random; denote the random network by G with realization 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 $\Pr(\theta = H) = p_0$. At each time t, agent i privately chooses effort $A_{i,t} \in [0, 1]$ at flow cost c. This effort results in successes at random arrival times $(T_i^1, T_i^2, ...)$ with arrival rate $A_{i,t}\mathbb{I}_{\{\theta=H\}}$. Agent i observes her own and her neighbors' past successes, but not others' actions. If the network is random, she knows G but nothing about the realization g, not even her own degree.[4]

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

$$V_i = \max_{\{A_{i,t}\}_{t \geq 0}} E\left[x \sum_{\iota=1}^{\infty} e^{-rT_i^\iota} - c\int_0^\infty e^{-rt} A_{i,t} dt\right] \tag{1}$$

where the expectation is taken over quality θ, network G, and arrival times $\{T_i^\iota\}$. We solve for weak perfect Bayesian equilibria, where agents who have observed a success infer that $\theta = H$.

Remarks. Agents observe neighbors' successes but not their actions. This makes the model tractable by focusing on a single mechanism of information transmission. We think this is reasonable in many applications (e.g. a farmer doesn't know if her neighbor experimented with a new crop, but observes the results of a successful harvest). Agents also cannot communicate with each other directly. An agent has little incentive to reveal her failures since this will tend to make her neighbors more pessimistic and lower their experimentation.

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

[4]The main role of this assumption is to ensure our agents are symmetric in our large regular random networks. It has no impact on the analysis of the best-response (Section 3.1), our deterministic examples (Section 3.3) or the core-periphery (Section 4.3).

3 Preliminary Analysis

3.1 Best-Responses: Cutoff Strategies

In this section, we characterize the best response of a generic agent Iris, given arbitrary strategies of other agents.

As a benchmark, consider Iris's single-agent experimentation problem, or equivalently her problem when she has no neighbors. After her first success, she chooses maximal effort $A_{i,t} = 1$ and obtains continuation value $y := (x - c)/r$. Before that, her posterior belief evolves according to

$$p_t = P^{\emptyset}(t) := \frac{p_0 e^{-t}}{p_0 e^{-t} + (1 - p_0)}.$$

Iris thus experiments until time $\bar{\tau}$ when her belief hits the *single-agent threshold belief* $p_{\bar{\tau}} = \underline{p} := c/(x + y)$. It is also useful to define the *myopic threshold belief* $\bar{p} := c/x$, where Iris would stop if she ignored the future benefit of success, y.

Now, consider the general problem where Iris learns from her neighbors, $N_i(G)$. Write $T_i = T_i^1$ for Iris's first success time, and $S_i := \min_{j \in N_i(G)} T_j$ for her neighbors' first success. After Iris observes a success at $\min\{T_i, S_i\}$, she chooses maximal effort and receives continuation value y. We can thus restrict attention to earlier times, and write $\{a_{i,t}^{\emptyset}\}_{t \geq 0}$ for her *experimentation*, i.e. her effort before $\min\{T_i, S_i\}$. Also write

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

for Iris's rate of social learning,[5] where the expectation is taken over the random network G and success times $\{T_j\}$, conditional on $\theta = H$. We also define Iris's cumulative social learning $B_{i,t} := \int_0^t b_{i,s} ds$, and abuse terminology by referring to both $\{b_{i,t}\}$ and $\{B_{i,t}\}$ as Iris's *social learning curve*. Since Iris's experimentation is unobservable to others and her own success effectively ends the game for her, Iris takes $\{b_{i,t}\}$ as given. We thus study the best response $\{a_{i,t}^{\emptyset}\}$ to $\{b_{i,t}\}$, and drop the i subscript for the rest of the section.[6]

When $\theta = H$, the random time $\min\{T, S\}$ has hazard rate $a_t^{\emptyset} + b_t$, and so the chance of not observing a success before t equals $\exp(-\int_0^t (a_s^{\emptyset} + b_s) ds)$. Bayes's rule then implies the posterior belief

$$p_t = P^{\emptyset} \left(\int_0^t (a_s^{\emptyset} + b_s) ds \right).$$

[5] Iris additionally knows her past actions $\{A_{i,s}\}_{s<t}$, but since these are deterministic given $t < T_i, S_i$ there is no need to include them in the conditional expectation.

[6] The analysis in this section immediately applies to other models of social learning that give rise to a social learning curve $\{b_{i,t}\}$, say time-varying networks and/or random networks with private information.

Truncating her objective function (1) at the first observed success reduces it to a deterministic control problem

$$V = \max_{\{a_t^\emptyset\}_{t\geq 0}} \int_0^\infty \left(\left(a_t^\emptyset(x+y) + b_t y\right) p_t - a_t^\emptyset c\right) e^{-\int_0^t \left(r + (a_s^\emptyset + b_s)p_s\right) ds} dt. \tag{3}$$

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, $(a_s^\emptyset + b_s)p_s$.

Clearly, Iris experiments above the myopic threshold, $p_t \geq \bar{p}$. Conversely, equation (3) implies that Iris stops to experiment below the single-agent threshold, $p_t \leq \underline{p}$. For intermediate beliefs $p_t \in [\underline{p}, \bar{p}]$, her choice depends on her social learning. To avoid trivialities, we restrict $p_0 \in (\underline{p}, 1)$. We say Iris's prior is *optimistic if* $p_0 > \bar{p}$ and *pessimistic* if $p_0 < \bar{p}$;[7] the upshot of this distinction is that an optimistic agent always engages in some experimentation, no matter her social learning curve.

We first claim that Iris uses a *cutoff strategy* in that she experiments maximally until some cutoff time τ and then stops, $a_t^\emptyset = \mathbb{I}_{\{t\leq\tau\}}$.[8] Intuitively, it makes no sense to stop experimenting at some τ', but then resume it after neighbors' lack of success over $[\tau', \tau'']$. For a more rigorous argument, suppose Iris 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 $r\delta(p_t(x+y)-c)\epsilon$, which is positive in the relevant range of posteriors $p_t > \underline{p}$. Second, if one of her neighbors succeeds 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 $p_\tau \in [\underline{p}, \bar{p}]$.

To characterize the optimal cutoff τ, define Iris's *experimentation incentives* at time-t,

$$\psi_t := p_t \left(x + ry \int_t^\infty e^{-\int_t^s (r+b_u) du} ds \right) - c. \tag{4}$$

To understand (4), suppose that successes from Iris's neighbors arrive at constant rate b, so (4) simplifies to $p_t(x + \frac{r}{r+b}y) - c$. If she raises the cutoff from t to $t+\delta$, she gains the expected payoff from a success $p_t(x+y)\delta$, 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.

We summarize this discussion as follows:

[7] We are opportunistic about calling the boundary case $p_0 = \bar{p}$ optimistic or pessimistic.

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

Proposition 1. *Given social information $\{b_t\}$, the agent's optimal experimentation is given by the cutoff strategy $a_t^\emptyset = \mathbb{I}_{\{t \leq \tau\}}$, where the cutoff time $\tau \in (0, \bar{\tau}]$ uniquely solves $\psi_\tau = 0$ if $\psi_0 > 0$, and $\tau = 0$ if $\psi_0 \leq 0$.*

Proof. See Appendix A.1. □

Proposition 1 reduces the potentially complicated dynamic experimentation problem of a forward-looking, Bayesian agent to choosing one number, τ, which is characterized explicitly by setting (4) to zero. This tractability allows us to characterize equilibria for rich classes of networks. In contrast to Proposition 1, the seminal papers on strategic experimentation in the clique network, Keller, Rady, and Cripps (2005) and Bonatti and Hörner (2011), both find agents gradually phase out effort in equilibrium. This difference arises because free-riding incentives are greater in their models: In Keller, Rady, and Cripps (2005), Iris's neighbors observe her experimentation and makes them pessimistic (when it fails); in Bonatti and Hörner (2011), the payoff is public, so Iris does not want to exert effort if others are about to succeed.

3.2 Best-Responses: Comparative Statics

This section derives two useful comparative statics on Iris's value and her optimal cutoff as a function of social learning. The first result shows that more social information raises Iris's value and leads her to stop earlier. Thus, this is a game of strategic substitutes. The optimal cutoff τ is maximized in the absence of social learning, $B_t \equiv 0$, where it coincides with the single-agent solution, $\tau = \bar{\tau}$.

Lemma 1. *Higher social learning $\{B_t\}_{t \geq 0}$ increases value V and decreases the cutoff τ.*

Proof. Clearly, a rise in $\{B_t\}$ constitutes Blackwell-superior information, which raises V. Experimentation incentives (4) fall both in pre-cutoff learning B_τ which lowers the cutoff belief $p_\tau = P^\emptyset(\tau + B_\tau)$ and in future learning $\{b_t\}_{t \geq \tau}$. To show that ψ_τ falls in the integral, $\{B_t\}$, we need to compare the impact of "early" and "late" increases in b_t.

Specifically, differentiating time-τ experimentation incentives (4) with respect to time-t social learning, we get

$$-\frac{\partial \psi_\tau}{\partial b_t} = \begin{cases} p_\tau \left(ry \int_\tau^\infty e^{-\int_\tau^s (r+b_u) du} ds + x - c \right) & \text{for } t < \tau \\ p_\tau ry \int_t^\infty e^{-\int_\tau^s (r+b_u) du} ds & \text{for } t > \tau, \end{cases} \tag{5}$$

where the case $t < \tau$ uses $\frac{\partial p_\tau}{\partial b_t} = -p_\tau(1-p_\tau)$ and $(1-p_\tau)(x + ry \int_\tau^\infty e^{-\int_\tau^s (r+b_u) du} ds) = x + ry \int_\tau^\infty e^{-\int_\tau^s (r+b_u) du} ds - (\psi_\tau + c)$. Clearly, (5) is positive and falls in t, weakly for $t < \tau$ and discontinuously at $t = \tau$.

Thus, incentives ψ_τ fall as a function of cumulative learning $\{B_t\}$. Since ψ_t strictly single-crosses from above (cf. proof of Proposition 1), the solution τ of $\psi_\tau = 0$ falls in $\{B_t\}$. □

Equation (5) tells us that pre-cutoff learning B_τ crowds-out the agent's experimentation more than post-cutoff learning $\{b_t\}_{t\geq\tau}$. After the cutoff, it crowds out the option value of own experimentation $ry\int_\tau^\infty e^{-\int_\tau^s (r+b_u)du} ds$, as seen in the second line of (5) for $t=\tau$. Before the cutoff, the additional term $x-c$ in the first line of (5) represents the reduced opportunity of achieving a first success at τ, conditional on $\theta = H$.[9]

Iris's value depends on her social information $\{B_t\}$. The following lemma shows that a sufficient statistic for $\{B_t\}$ is her pre-cutoff learning B_τ and her optimal cutoff time τ: Before τ, Iris exerts effort anyway, so does not care about the timing of social learning $\{B_t\}_{t\leq\tau}$. Post-cutoff learning $\{B_t\}_{t\geq\tau}$ matters only via the continuation value $V_\tau = p_\tau y\int_\tau^\infty b_s e^{-\int_\tau^s (r+b_u)du} ds = p_\tau(x+y) - c$,[10] which is a function of (τ, B_τ) since $p_\tau = P^\emptyset(\tau + B_\tau)$.

Lemma 2. *Assume the agent is willing to experiment initially, $\psi_0 \geq 0$. There exists a continuous function $\mathcal{V} : \mathbb{R}_+^2 \to \mathbb{R}_+$ that falls in both arguments, such that for any social learning curve $\{B_t\}$ with associated optimal cutoff τ, the agent's value equals $\mathcal{V}(\tau, B_\tau)$.*

Proof. See Appendix A.2. □

By Lemma 2, the agent's value $\mathcal{V}(\tau, B_\tau)$ *falls* in B_τ. This sounds counterintuitive, but it arises because we fix the optimal stopping time τ, as characterized by $\psi_\tau = 0$. An increase in pre-cutoff learning B_τ must be compensated by a fall in post-cutoff learning $\{b_t\}_{t\geq\tau}$ in order to keep τ constant. Moreover, since pre-cutoff learning has a discontinuously larger effect on ψ_τ than post-cutoff learning by (5), we must reduce the latter by a larger amount to compensate. In contrast to (5), the effect of social learning on value, $\partial V/\partial b_t$, is continuous in t, so the combination of a small raise of $b_{\tau-\epsilon}$ and a large drop of $b_{\tau+\epsilon}$ decreases value.

Lemma 2 is the key tool to compare equilibrium welfare across networks since τ and B_τ are easily characterized in equilibrium. For example, suppose we take two networks G and G' such that an agent's best-response is to shirk in the first, $\tau = 0$, and work in the second, $\tau' > 0$. Since $B_\tau = 0$ and $B_{\tau'} \geq 0$, the agent has utility $\mathcal{V}(0,0)$ under G, which exceeds her utility $\mathcal{V}(\tau', B_{\tau'})$ under G'.

Lemma 2 assumes that the agent is willing to experiment initially, so her stopping time is characterized by $\psi_\tau = 0$. Otherwise, if the agent is receiving too much social information,

[9] For optimistic agents, $p_0 > \bar{p}$, this asymmetry is stark: Full crowding out of incentives, is achieved by a finite amount $B_\tau = \bar{\tau}$ of pre-cutoff learning (inducing $p_\tau < \underline{p}$ and so $\psi_\tau < 0$), but by no amount of post-cutoff learning (since $\psi_0 > p_0 x - c > 0$ for any $\{B_t\}$).

[10] The first equality leverages the fact that all learning after τ is social $\{b_t\}_{t\geq\tau}$, and the second leverages the indifference condition $\psi_\tau = 0$.

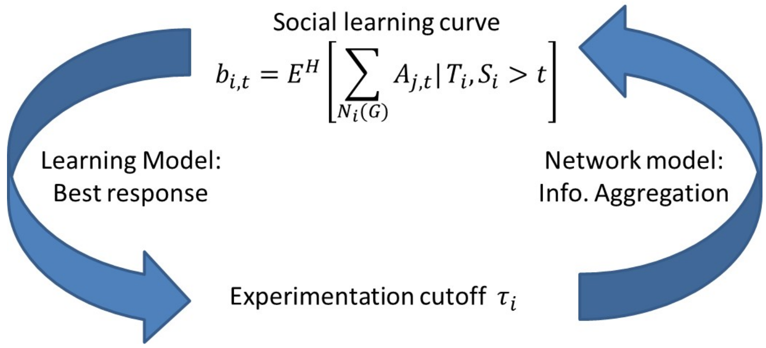

Figure 1: **Equilibrium Analysis**.

her value strictly exceeds $\mathcal{V}(0,0)$.[11] In the random networks in Section 4.2, all agents exert some effort and so $\psi_\tau = 0$. In the core-periphery networks in Section 4.3, we focus on the welfare of the worst-off peripheral agents, who also exert some effort.

3.3 Equilibrium: Examples

In the prior sections, we studied Iris's best response τ_i as a function of social learning $\{b_{i,t}\}$. To close the model in equilibrium we must study how individual cutoffs $\{\tau_j\}$ aggregate into social learning curves $\{b_{i,t}\}$, cf. Figure 1. Here we illustrate this aggregation in three canonical example networks, foreshadowing the more general analysis in Section 4.

Example 1 (Clique). Assume that all agents observe each other. We claim there is a unique equilibrium in which all agents equally divide the single-agent experimentation between them. That is, each agent i uses cutoff $\tau_i = \bar{\tau}/I$, recalling the single-agent experimentation cutoff $\bar{\tau}$ that solves $P^{\emptyset}(\bar{\tau}) = \underline{p}$. The resulting social learning curve is shown in Figure 2(a). Aggregate information is constant in I. Welfare rises with I as the agents share the total cost.

We prove our claim in two steps. First, the agents collectively experiment as much as a single agent would by herself $\sum_i \tau_i = \bar{\tau}$. This is because the agent who experiments the longest expects no social information after her cutoff, $b_{i,s} = 0$ for $s > \tau_i$. Hence she faces the first-order condition of the single-agent problem, $P^{\emptyset}(\sum_j \tau_j)(x+y) - c = 0 = \underline{p}(x+y) - c$. Second, the agents must split total experimentation evenly, $\tau_j = \bar{\tau}/I$. This follows because all agents are indifferent at $\underline{p}$ and, since agents prefer to front-load experimentation, they all experiment until that point. This is formally shown in Proposition 3 for all strongly symmetric networks.[12] △

[11] Formally, $V = V_0 = p_0 y \int_0^\infty b_s e^{-\int_0^s (r+b_u)du} ds = p_0(x+y) - c - \psi_0 = \mathcal{V}(0,0) - \psi_0 > \mathcal{V}(0,0)$.

[12] The uniqueness of equilibrium is notable. In public good problems with linear costs there is a continuum of equilibria; Bramoullé and Kranton (2007) select via a stability criterion while Galeotti and Goyal (2010)

Example 2 (Infinite Directed Line). Consider the following network:

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

In the unique symmetric equilibrium, an initial *experimentation phase* of length τ is followed by a *contagion phase*.[13] For example, suppose Kata succeeds in the experimentation phase, while Iris and John do not. After τ, Kata's success means that Kata and John continue to work while Iris shirks. Eventually John also succeeds, and all three work thereafter.

To solve for the equilibrium cutoff τ, we calculate Iris's social information during the contagion phase,

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

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 after τ, John works iff Kata has succeeded. The last equality uses Bayes' rule,

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

Thus, social information arrives at constant rate $b_{i,t} \equiv 1 - e^{-\tau}$, as illustrated in Figure 2(b). While the unconditional probability that Kata has succeeded, and hence John works, rises over time, this positive effect is exactly cancelled by conditioning on the bad news event that John has not succeeded yet, $t < T_j$.

Using (4), the equilibrium stopping time τ solves

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

Example 3 (Star). In the star, one core agent, Kata, k, has an undirected link to L peripheral agents, Lili, ℓ, who have no other links. By symmetry, Proposition 3 implies that peripherals share a common cutoff, τ_ℓ. Moreover, Kata learns faster than the peripherals, and so experiments less herself, $\tau_k < \tau_\ell$ (see Lemma 5 in Section 4.3). For large L, Kata

select via a network formation game. We resolve this determinacy through impatience. In experimentation papers there are also asymmetric equilibria (e.g. Keller, Rady, and Cripps (2005), Bonatti and Hörner (2011)). As discussed above, free-riding incentives are weaker in our paper, leading putative asymmetric equilibria to unravel.

[13] In Section 5 we approximate this infinite network with a sequence of finite random networks, that generate circles of random, exploding length. Corollary 1 shows that the unique equilibrium of each finite random networks is symmetric, and Proposition 4 shows that these equilibria converge to the symmetric equilibrium described here.

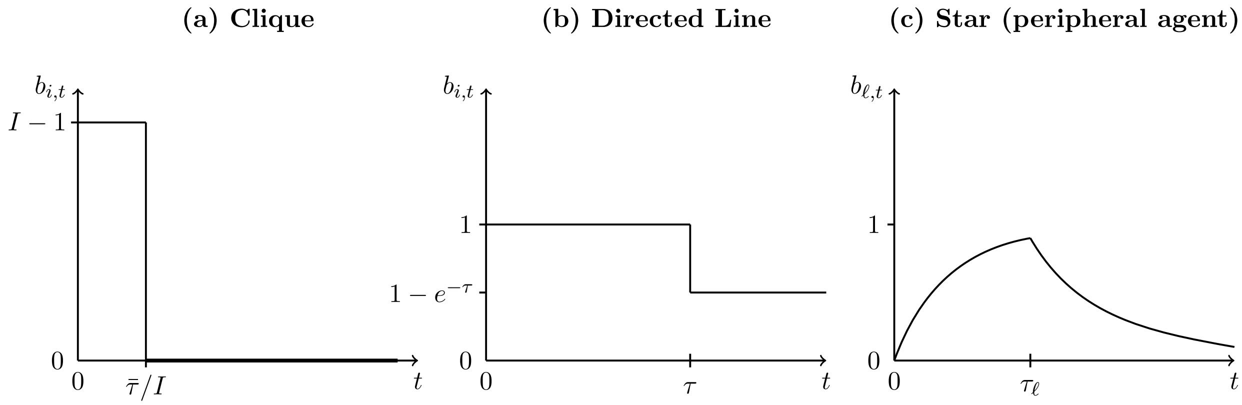

Figure 2: **Social Learning Curves.** This picture shows the rate of social learning b_{it} as defined in equation (2) for Examples 1-3, as described in the text.

does not experiment herself, $\tau_k = 0$. Information is thus primarily generated by peripherals, but flows via Kata who serves as an information broker.

If a peripheral succeeds before τ_ℓ, Kata starts to work; her eventual success then triggers all other peripherals to work. The resulting social learning curve for peripheral agent Lili $b_{\ell,t}$ undergoes two phases, illustrated in Figure 2(right) and derived in (13) for $K = 1$. Up to time τ_ℓ, it increases because of the effort of other peripherals, but tempered by Kata's lack of success. After τ_ℓ, no more additional information is created, and $b_{\ell,t}$ falls as the information filters through Kata, and Lili becomes pessimistic about Kata having seen a success. These dynamics are analogous to water flowing into a reservoir while the peripherals experiment, and slowly draining out through a bottleneck as Kata conveys the information. △

3.4 Equilibrium: Existence and Uniqueness

We round off our preliminary analysis by establishing equilibrium existence and presenting a limited uniqueness result.

Proposition 2. *Equilibrium exists.*

Proof. We argue that Nash's best-response mapping in cutoff vectors $\{\tau_j^*\} : [0, \bar{\tau}]^I \to [0, \bar{\tau}]^I$ is continuous, which implies equilibrium existence by Brouwer's fix point theorem. First note that i's social learning curve $\{b_{i,t}\}_{t\geq 0}$, as in (2), is pointwise continuous in $\{\tau_j\}_{j\neq i}$ for all $t \neq \tau_j$. Then, Lebesgue's dominated convergence theorem implies that incentives $\psi_{i,t}$ in (4) are also continuous in $\{\tau_j\}_{j\neq i}$ for all t. Finally, since $\psi_{i,t}$ strictly single-crosses in t (see the proof of Proposition 1), its root $\tau_i^*(\{\tau_j\}_{j\neq i})$ is also continuous. □

Uniqueness is more difficult. We cannot prove equilibrium uniqueness in general, but can show uniqueness for "strongly symmetric" networks. For a deterministic network g and agents $i \neq j$, define $g_{i\leftrightarrow j}$ to be the same network when switching i and j.[14] For a random network G, define $G_{i\leftrightarrow j}$ by $\Pr^{G_{i\leftrightarrow j}}(g) = \Pr^G(g_{i\leftrightarrow j})$ for all g.

Proposition 3. *If i, j are symmetric in G, i.e. $G_{i\leftrightarrow j} = G$, then in any equilibrium $\tau_i = \tau_j$.*

Proof. See Appendix A.3. □

For an intuition, consider a deterministic, undirected network where i and j are not connected. By contradiction assume $\tau_i > \tau_j$. Since i's additional learning over $[\tau_j, \tau_i]$ is more immediate to i than to j, who only benefits indirectly via some other agent k, we can argue that $\min\{T_i, S_i\}$ is smaller than $\min\{T_j, S_j\}$. This greater chance of learning the state depresses i's experimentation incentives below j's, leading to the contradiction that $\tau_j > \tau_i$.

Say that network G is *strongly symmetric* if $G_{i\leftrightarrow j} = G$ for any pair of agents i, j.

Corollary 1. *Strongly symmetric networks have a unique equilibrium, characterized by a cutoff $\tau \in (0, \bar{\tau})$, such that $\tau_i = \tau$ for all i.*

Proof. Proposition 3 implies that all agents must share the same cutoff τ. Uniqueness of the cutoff follows from strategic substitutes: Indeed, consider two cutoffs $\tau' > \tau$ where the latter constitutes a symmetric equilibrium: $\psi_{i,\tau} = 0$ when $\tau_j = \tau$ for all $j \neq i$. When others $j \neq i$ use the higher cutoff $\tau_j = \tau'$, Iris learns more $B'_{i,t} \geq B_{i,t}$ for all t, and thus chooses to stop prior to $\tau < \tau'$ by Lemma 1. Hence τ' does not constitute a symmetric equilibrium. □

Caveat: Our notion of symmetry is so strong that only two deterministic networks satisfy it: the clique, and the empty network. Indeed, the infinite directed line (Example 2) or a finite directed circle violate it since only agent i observes $i+1$, so $g \neq g_{i\leftrightarrow j}$ for any $j \neq i$. With this said, many natural classes of random networks, such as the configuration networks studied in Section 4.2, do satisfy strong symmetry and Corollary 1 applies. Moreover, Proposition 3 is useful beyond strongly symmetric networks; for instance, equilibria in core-periphery networks in Section 4.3 are characterized by one cutoff τ_k for all core-agents, and another cutoff τ_ℓ for all peripherals.

[14] Formally, given g, we can define $g_{i\leftrightarrow j}$ by three types of links. First, links involving i and j: $(i, j) \in g_{i\leftrightarrow j}$ iff $(j, i) \in g$, and analogously, $(j, i) \in g_{i\leftrightarrow j}$ iff $(i, j) \in g$. Second, links involving one third party: $(i, k) \in g_{i\leftrightarrow j}$ iff $(j, k) \in g$ and three analogous conditions, replacing i and j and switching the direction of these links. Third, links involving two third parties: $(k, \ell) \in g_{i\leftrightarrow j}$ iff $(k, \ell) \in g$.

4 Density of Links

We now turn to the main question of the paper: How learning and welfare depend on the network density. Section 4.1 defines some terminology and introduces a second-best benchmark for welfare. Section 4.2 studies large random networks. Section 4.3 studies core-periphery networks. The results for asymptotic learning and welfare in these two sections run parallel to one another, but the learning dynamics differ.

4.1 Bounds on Learning and Welfare

Information Aggregation. Iris eventually learns the successes of all agents with a path from i to j. Lemma 1 implies $\tau_j \leq \bar{\tau}$, and so Iris's (eventual) social information, as $t \to \infty$, is bounded by $B_{i,\infty} \leq (I-1)\bar{\tau}$. To study large networks, we consider sequences of networks $\{G^I\}_{I \in \mathbb{N}}$ with associated social information $B^I := \min_j B^I_{j,\infty}$ corresponding to the least-informed agent. If G^I admits multiple equilibria, we consider the infimum values of B^I. In the limit, as $I \to \infty$, define *asymptotic information* as $B = \liminf B^I$. There is *asymptotic learning* if $B = \infty$, so all agents eventually learn the state.

Welfare. Iris's value in trivially bounded above by the value of learning the state perfectly immediately, $V_i < p_0 y$. Another, less obvious, upper bound on agents' value comes from the fact that for i to socially learn, some other agent $j \neq i$ must generate that social information. By Lemma 2, this implies that $\min_j V_j < \mathcal{V}(0,0) = p_0(x+y) - c$,[15] motivating the following Rawlsian welfare upper bound

$$V^* := \min\{p_0 y, p_0(x+y) - c\},$$

illustrated in Figure 3 as a function of p_0. Given a sequence of networks $\{G^I\}_{I \in \mathbb{N}}$, let $V^I := \min_j V^I_j$ be the welfare of the worst-off agent. If G^I admits multiple equilibria, we consider the infimum values of V^I. In the limit, as $I \to \infty$, define *asymptotic welfare* as $V = \liminf V^I$. Our main results, Theorem 1 and 2 show that sequences of random networks and core-periphery networks with the appropriate, intermediate network density can attain the welfare upper bound, $V = V^*$.[16]

[15] This upper bound relies on agents using equilibrium strategies. In a sequence of clique networks, and symmetric cutoffs τ^I that vanish individually $\lim \tau^I = 0$ but explode in aggregate $\lim I\tau^I = 0$, agents' payoffs approach $p_0 y$, which exceeds $p_0(x+y) - c$ for pessimistic priors $p_0 < \bar{p}$. However, in equilibrium, an agent in such a network would have a strict incentive to shirk for large I (see Example 1).

[16] In both cases, the proof of the Theorem characterizes unique limit points for $\{B^I\}$ and $\{V^I\}$, so the $\liminf$ equals the ordinary limit. In the large random networks equilibria are unique, so taking the infimum over equilibrium values is moot. In finite core-periphery networks, we do not know whether equilibrium is unique, but the unique characterization of the limit points does not rely on taking the infimum over equilibria

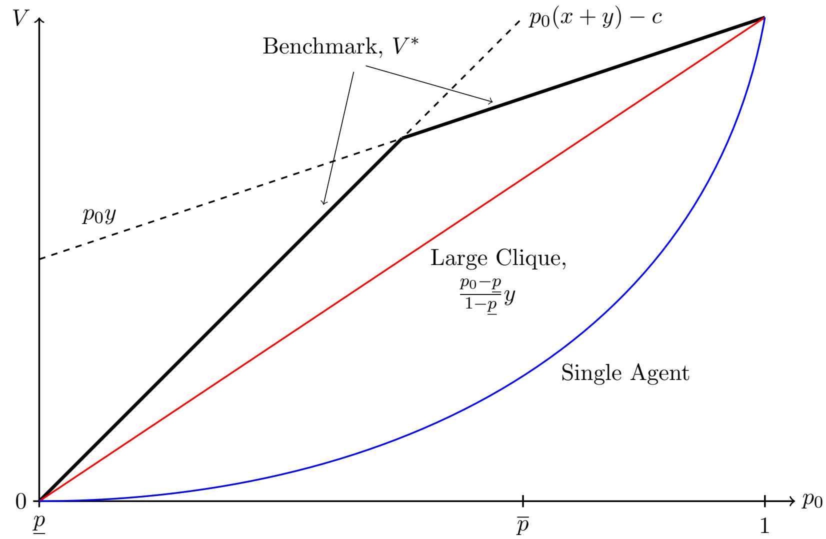

Figure 3: **Value Bounds.** The figure shows the bound on the value V^* as a function of the agent's prior, p_0. This upper bound is piece-wise linear, with a downward kink at $\bar{p}$, the myopic cutoff. The figure also illustrates agents' values in a large clique and when by themselves.

While asymptotic learning and the welfare benchmark are both driven by social learning $\{B_t\}$, they are distinct concepts since asymptotic learning focuses on the long run, while welfare incorporates discounting. For optimistic priors $p_0 \geq \bar{p}$, the welfare bound requires agents learn the state immediately, so clearly they also learn asymptotically. For pessimistic priors $p_0 < \bar{p}$ we will see that, for core-periphery networks, asymptotic learning and the welfare upper bound are mutually exclusive. To illustrate the benchmarks, we return to the three examples in Section 3.3.

Example 1, Continued (Clique). In a finite network, the agents divide up the single-agent experimentation, $\tau_i = \bar{\tau}/I$. As $I \to \infty$, individual experimentation vanishes and all learning is social with asymptotic information $B = \bar{\tau}$. Agents' receive all their social information before stopping, $B_\tau = \bar{\tau}$, and their asymptotic value equals $\mathcal{V}(0, \bar{\tau})$. More concretely, their beliefs instantly jump to 1 (if there is a success) or drop to $\underline{p}$ (if there is no success). The payoff to the former is y, so the equilibrium values converge to $\mathcal{V}(0, \bar{\tau}) = (p_0 - \underline{p})y/(1-\underline{p}) < V^*$, as illustrated in Figure 3. Thus, the speed of diffusion in the clique chokes off discovery and means that agents neither asymptotically learn nor obtain the welfare benchmark. $\triangle$

Example 2, Continued (Infinite Directed Line). In this infinite network, each agent experiments for time $\tau > 0$, where τ solves (7). This network aggregates information: Each

and rather applies for any equilibrium selection.

agent learns from infinite agents who each perform a strictly positive amount of experimentation. More explicitly, an agent's social learning curve $B_t = \tau + (1 - e^{-\tau})(t - \tau)$ is unbounded. However, agents learn too slowly and they do not attain the welfare benchmark. Specifically, each agent experiments for τ and learns an additional τ from her neighbor before stopping, so $B_\tau = 2\tau$ and welfare equals $\mathcal{V}(\tau, 2\tau) < V^*$. $\triangle$

Example 3, Continued (Star). When there is a large number of peripherals, the core agent Kata shirks, $\tau_k = 0$. The peripherals thus do all the experimentation and have the lowest information and welfare, so we focus on them. We show in Section 4.3 that agents asymptotically learn iff $p_0 \geq p^s$. This threshold is defined so that a peripheral agent will just work at $t = 0$ if he thinks Kata will instantly learn the state and choose $b_{k,t} \equiv 1$ thereafter,

$$\psi_{\ell,0} = p^s \left(x + \frac{r}{r+1} y \right) - c = 0.$$

Note that $p^s \leq \overline{p}$, so agents asymptotically learn if they have an optimistic prior.

The welfare result is exactly the opposite: Agents attain the welfare benchmark iff $p_0 \leq p^s$. For a high prior, $p_0 > p^s$, the peripheral agents experiment in the limit, $\tau_\ell > 0$, meaning Kata instantly learns. Thus a peripheral agent learns $B_{\tau_\ell} = \tau_\ell$ before stopping and has value $\mathcal{V}(\tau, \tau) < V^*$.[17] For a low prior, $p_0 \leq p^s$, the peripherals stop experimenting in the limit, $\tau_\ell = 0$, and since $B_{\tau_\ell} \leq \tau_\ell$, their value converges to its upper bound $V^* = \mathcal{V}(0, 0)$.[18] Thus, asymptotic learning and the welfare benchmark are not only distinct concepts, but in fact mutually exclusive (for generic priors with $p \neq p^s$). $\triangle$

4.2 Random Networks

We first study large random networks. This is a tractable canonical class of networks that captures realistic contagion dynamics. For simplicity we focus on regular networks, where agents all have the same number of neighbors, and comment after our main result which insights generalize to non-degenerate degree distributions. This class is rich enough to encompass the clique and trees, as in Sadler (2020a) and Board and Meyer-ter-Vehn (2021).

We construct a *regular random network* as follows. Each of the I agents has $\hat{n}^I \geq 2$ link stubs. We randomly draw pairs of stubs and connect them into undirected links. We then prune self-links (from i to i), multi-links (from i to j), and if $\hat{n}^I I$ is odd the single leftover stub. We assume that agents observe nothing about the network realization, not even their

[17]This presumes $p_0 \leq \bar{p}$ so that $V^* = \mathcal{V}(0, 0)$. For $p_0 > \bar{p}$, the welfare bound $V^* = p_0 y$ requires immediate perfect social information, which is clearly impossible with a single neighbor.

[18]It may appear paradoxical that the welfare upper bound V^* is achieved for low, but not high prior beliefs p_0. This is because V^* itself rises as function of p_0, and is hence a more demanding benchmark for high p_0.

own degree; omitting such information seems innocuous since agents' asymptotically know their degree, cf. Lemma 3.

By construction, the random network is strongly symmetric, and so Corollary 1 implies there is a unique equilibrium. Denote the symmetric cutoff by τ^I and agents' value by V^I. For tractability we consider sequences of such networks with degrees $\{\hat{n}^I\}$, and assume existence of the limits $\nu := \lim \hat{n}^I$, $\lambda := \lim \hat{n}^I / \log I$ and $\hat{\rho} := \lim \hat{n}^I / I$, possibly equal to ∞.

Let N^I be the number of realized links of a random agent. Some stubs fail to form links, so N^I is random with expectation $n^I := E[N^I] < \hat{n}^I$. We now argue that we can ignore this complication as $I \to \infty$.

Lemma 3. *As the network grows large, $I \to \infty$,*

(a) Realized degree: $N^I/n^I \xrightarrow{D} 1$.

(b) Expected degree: $n^I/I \to 1 - e^{-\hat{\rho}}$. If $\hat{\rho} = 0$, then $n^I/\hat{n}^I \to 1$.

(c) Information at the cutoff time: $\lim B^I_{\tau^I} = \lim n^I \tau^I$.

Proof. See Appendix B.1. □

Lemma 3(a) means agents essentially know their realized degree N^I. As a result, agents do not update N^I during the experimentation phase $t < \tau^I$, as stated in part (c). Part (b) means we can ignore the distinction between stubs and links when $\nu < \infty$ or $\lambda < \infty$ (which imply $\hat{\rho} = 0$), and so $\nu = \lim n^I$, $\lambda = \lim n^I / \log I$.

Our main result, Theorem 1, shows that the relevant measures of limit network density are ν, λ and $\rho := \lim n^I/I = 1 - e^{-\hat{\rho}}$. For sparse networks, where agents have a bounded number of links, we have $\nu \in [2, \infty)$. Proposition 4 in Section 5 shows that such networks approach a tree as $I \to \infty$. For intermediate networks, where information spreads across the network in finite time, as in Milgram's six degrees of separation, we have $\lambda \in (0, \infty)$. Indeed, Lemma 4 will show that the inverse $1/\lambda$ measures the network's *time-diameter,* i.e. the time for social information to travel between two random agents in the network.[19] For dense networks, where agents have finite proportion of links, we have $\rho \in (0, \infty)$ and agents are at most two links apart. For $\rho = 1$ we approximate the clique. The set of network (limit) densities is the disjoint union $\mathbb{N} \cup \{\infty | 0 \cdot \log I\} \cup (0, \infty) \cup \{\infty \cdot \log I | 0 \cdot I\} \cup (0, 1]$, which we endow with its natural order.[20]

[19] This time-diameter $1/\lambda = \lim(\log I / n^I)$ is smaller than the typical diameter estimate for large random networks $\lim(\log I / \log n^I)$. The smaller diameter reflects a faster contagion process: contagion in our model does not travel one link in every discrete time period; rather each link transmits continuously with rate one. Much like compound interest, this allows nodes infected at $t' \in [t, t+1]$ to begin transmitting immediately, instead of having to wait until $t + 1$.

[20] This order treats many sequences of networks as equally dense. For instance, $n^I = \log \log I$ or $n^I = (\log I)^{1/2}$ both correspond to $\nu = \infty, \lambda = 0$. Theorem 1 shows that asymptotic information B and welfare V of a sequence of networks $\{n^I\}$ only depends on its limit density.

We now define the threshold density for asymptotic learning. For pessimistic priors, $p_0 < \bar{p}$, let $\sigma^* \in [0, \infty)$ be such that perfectly learning the state at time σ^* renders an agent indifferent about experimenting at $t = 0$

$$\psi_0 = p_0 \left(x + (1 - e^{-r\sigma^*}) y \right) - c = 0. \tag{8}$$

Here, $e^{-r\sigma^*} y$ is the agent's post-experimentation continuation value. For optimistic priors, $p_0 \geq \bar{p}$, set $\sigma^* = 0$.

Theorem 1. *In large random networks $\{n^I\}$:*

(*a*) *Asymptotic information B is a decreasing function of network density: It attains asymptotic learning $B = \infty$ iff $\lambda \leq 1/\sigma^*$, and strictly falls for $\lambda \geq 1/\sigma^*$.*

(*b*) *Welfare V is a single-peaked function of network density: It first strictly rises in ν, attains the benchmark V^* iff $\nu = \infty$ and $\rho = 0$, and then strictly falls in ρ.*

Proof. See Appendix B.2 □

Asymptotic learning requires sparse networks. Intuitively, denser networks accelerate learning, crowd out experimentation, and undermine learning in the long run. Welfare attains the benchmark when network density is intermediate. Intuitively, welfare discounts the future and so relies on both information generation and its quick dissemination. Figure 4 illustrates Theorem 1 for $p_0 < \bar{p}$.[21] The top and middle panels sketch asymptotic information B and welfare V as functions of network density. The bottom panel illustrates cumulative social learning curves $\{B_t\}$ for three qualitatively different regions of network density.[22]

In Figure 4(i) agents have finite links, $\nu < \infty$, so the network resembles a tree with independent information across Iris's neighbors. Social learning in the contagion phase $\{B_t\}_{t \geq \tau}$ is convex with increasing rate b_t described by a first-order ODE (see Section 5). This convexity reflects the fact that an agent has ν first-degree neighbors, $\nu(\nu - 1)$ second-degree neighbors, $\nu(\nu - 1)^2$ third-degree neighbors, and so on; so contagion accelerates over time. Each agent experiments for a bounded time $\tau > 0$, which ensures asymptotic learning,[23] while welfare $\mathcal{V}(\tau, \nu\tau)$ falls short of the benchmark, V^*. In the limit, as $\nu \to \infty$, individual experimentation vanishes, $\tau \to 0$, and the convex cumulative learning curve $\{B_t\}$ converges to a step function, constant equal to 0 below σ^*, and ∞ above.

To understand this, we need a brief detour. When $\nu = \infty$, the success times of Iris's neighbors may no longer be independent. Instead, the key analytical tool is an accounting

[21]For $p_0 \geq \bar{p}$, experimentation incentives are higher, and asymptotic learning obtains as long as $\rho = 0$.

[22]The more intuitive rates of social learning $\{b_t\}$ in Figure 2 fail to exist for $\nu = \infty$.

[23]Recall that an agent eventually learns the successes of all agents in her component. Given $\nu \geq 2$, the component of a typical agent has size proportional to I almost surely.

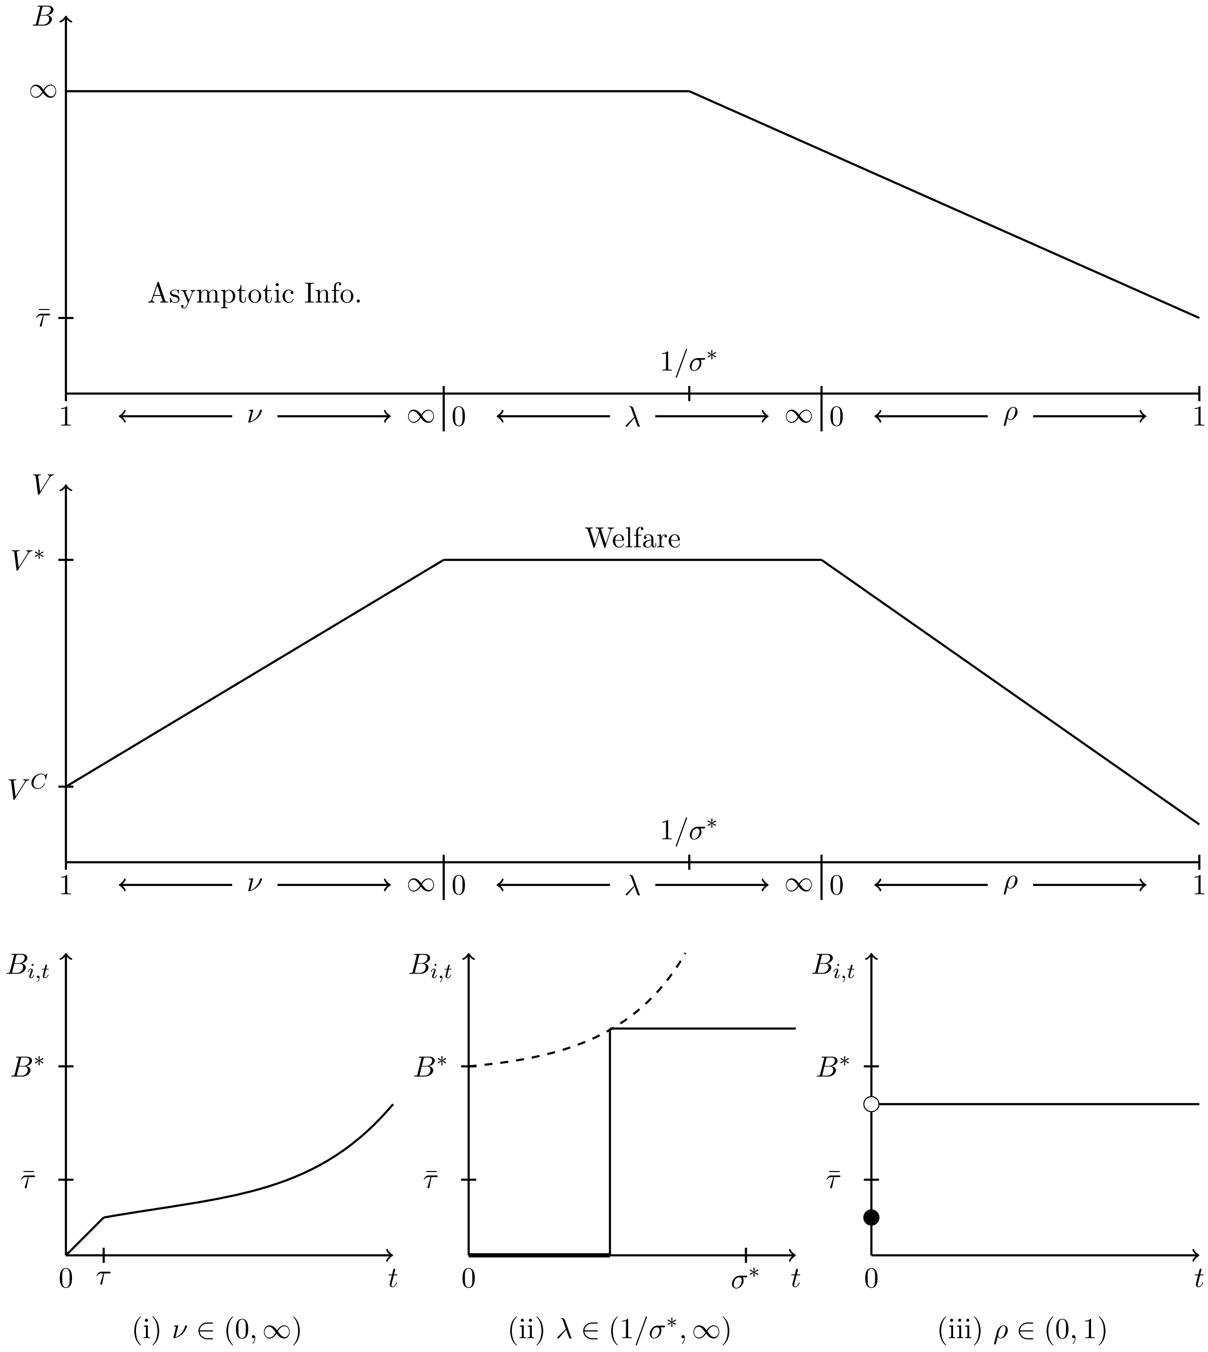

Figure 4: **Large Random Networks for Pessimistic Priors, $p_0 \leq \bar{p}$.** The **top** panel shows asymptotic information B as a function of network density, as described in Theorem 2(a). The **middle** panel shows welfare V as a function of network density, as described in Theorem 2(b). The **bottom** panel shows the cumulative learning curves of a typical agent in three canonical cases, as discussed in the text. In particular, σ^* is defined by (8) and B^* is defined by (11).

identity, that states that asymptotically all agents observe the first success at the same, deterministic time. To state the result, consider any sequence of cutoffs $\{\tau^I\}$ (not necessarily equilibrium) such that the limit $\sigma := \lim \frac{-\log \tau^I}{n^I} \in [0, \infty]$ exists. Let S^I be the random time at which a given agent i gets exposed, let $\bar{B} := \lim I\tau^I$ be total information, and S be the binary random time with $\Pr(S = \sigma) = 1 - e^{-\bar{B}}$ and $\Pr(S = \infty) = e^{-\bar{B}}$.[24]

Lemma 4. *Assume $\nu = \infty$. Asymptotically, i gets exposed at time σ or never: $S^I \xrightarrow{D} S$.*

Proof. See Appendix B.3. □

As a Corollary, Lemma 4 implies that agents learn asymptotically all generated information, $B = \bar{B}$. When B is finite, so is the expected number of "seeds" $1 - e^{-B}$, and the exposure time coincides with the network's time-diameter

$$\sigma = \lim \frac{\log I\tau^I - \log \tau^I}{n^I} = \lim \frac{\log I}{n^I} = \frac{1}{\lambda}. \tag{9}$$

To understand Lemma 4, suppose Iris's neighbors are a negligible share of the population, $\rho = 0$. At τ^I, there are approximately $I\tau^I$ exposed agents and the probability at least one agent learns the state is $1 - e^{-I\tau^I}$. The contagion then grows geometrically at rate n^I, so there are approximately $\tau^I I e^{n^I t}$ exposed agents at time t and, heuristically, everyone is exposed when $\tau^I I e^{n^I t} = I$, or $t = \frac{-\log \tau^I}{n^I} =: \sigma$.[25] This argument slightly overstates exposures because of double-counting. But this problem scales with the share of exposed agents, and we only need the argument as long as this share is negligible: once a fixed share of the population is exposed, all agents are exposed immediately since $n^I \to \infty$. The proof uses Chernoff bounds to make these arguments rigorous.

Returning to our equilibrium characterization, consider networks with $\lambda \in (0, 1/\sigma^*)$, so a time-diameter $1/\lambda \geq \sigma^*$, as illustrated in Figure 4(ii). By Lemma 4, Iris learns at time σ. To ensure equilibrium indifference (8), the rate at which τ^I goes to 0 adjusts to keep $\sigma = \sigma^*$. These networks are sparse enough to accommodate asymptotic learning. At the same time, they are dense enough to fully crowd out agents' pre-cutoff social learning, $\lim n^I \tau^I = 0$, and thus attain the welfare benchmark V^*.

For denser networks $\lambda \in (1/\sigma^*, \infty)$,[26] so a time-diameter $1/\lambda \leq \sigma^*$, learning is too fast to sustain perfect learning. Agents are exposed at $\sigma = 1/\lambda$ and equilibrium information

[24] Asymptotically the distinction between $\min\{T^I, S^I\}$ and S^I vanishes. If $\tau^I \to 0$ (as will be the case in equilibrium), i's experimentation is certain to fail; if $\lim \tau^I > 0$, social learning is asymptotically immediate and perfect, $\sigma = 0$ and $\bar{B} = \infty$, so $S = 0$.

[25] Recalling footnote 19, here we see the difference between typical discrete-time contagion models where exposed agents grow like $e^{(\log n^I)t}$ and our continuous-time model with the faster rate $e^{n^I t}$.

[26] For optimistic priors, $p_0 \geq \bar{p}$, this region is empty.

$B < \infty$ falls to maintain the indifference condition

$$p_0 \left(x + \left(1 - e^{-r\sigma}(1 - e^{-B})\right) y\right) = c. \tag{10}$$

Thus asymptotic learning fails but welfare attains its benchmark, V^*. As λ grows, the corner of the step function slides along the dashed line in Figure 4(ii). When $\lambda = \infty$ and $\rho = 0$, agents learn the state with probability $1 - e^{-B^*}$ immediately; the total information B^* is determined by agents' indifference condition,

$$p_0(x + e^{-B^*} y) = c. \tag{11}$$

The third type of equilibrium occurs when $\rho \in (0, 1)$, as illustrated by Figure 4(iii). Such networks are analogous to the clique, so given total information B, agents learn ρB before stopping and $(1 - \rho)B$ immediately after stopping. Total information B is then given by agents' indifference condition,

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

As $\rho \to 1$, we approach the clique and $B \to \bar{\tau}$.

Theorem 1 is stated for regular, undirected networks. The analysis immediately extends to regular directed networks, or regular triangular network (as described in Section 5). For networks with nondegenerate degree distributions where agents know their own degree, experimentation falls with an agent's degree. Equilibrium is thus characterized by a multi-dimensional fixed point argument which no longer guarantees uniqueness and undermines our sharp comparative statics. Non-regular networks also introduce a novel possibility for asymptotic learning to fail: An agent may be isolated, or more generally the size of her limit component may be finite. For instance, this arises with positive probability in Erdos-Renyi networks where links realize with probability q^I and $q^I I$ is bounded. While our analysis does not capture such networks, we nevertheless anticipate the spirit of the result to extend because of the same underlying economic forces. Indeed, the next section derives a result analogous to Theorem 1 for core-periphery networks, that are arguably more different from our regular random networks than the variants just discussed.

4.3 Core-Periphery Networks

In this section we study core-periphery networks. Theorem 2 shows that asymptotic information falls with network density while welfare is single-peaked, echoing Theorem 1 for random networks. This analysis serves three purposes. First, core-periphery networks are of intrinsic

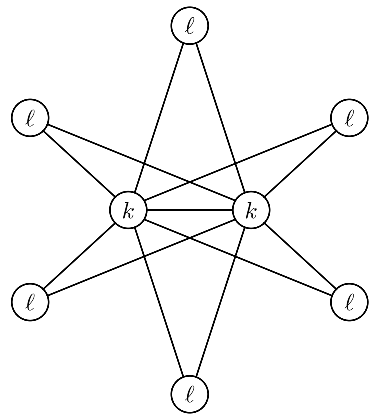

Figure 5: **Core-Periphery Network** with $K = 2$ core agents, and $L = 6$ peripherals.

interest: They are used to describe financial markets (e.g. Li and Schürhoff (2019)) and can arise endogenously in network formation models (Galeotti and Goyal (2010)). Second, core-periphery networks allow us to examine the role of network position for information generation. Third, core-periphery networks have a different neighborhood structure, with relatively few first neighbors in the core constituting a bottleneck for transmitting the information generated by the more numerous peripherals. As a result, social learning curves are then concave rather than convex in the contagion phase.

A *core-periphery network* is an undirected, deterministic network that consists of K core agents and $L = I - K$ peripheral agents. The core agents k are connected to everyone. The peripheral agents ℓ are only connected to core agents. See Figure 5 for an illustration. When $K = 1$, we have the star from Example 3.

Lemma 5. *Any equilibrium in a core-periphery network is characterized by two cutoffs, τ_k for all agents in the core, and τ_ℓ for all peripherals. Core agents work less, $\tau_k < \tau_\ell$, and have higher values, $V_k > V_\ell$.*

Proof. By symmetry and Proposition 3, equilibrium is characterized by cutoffs (τ_k, τ_ℓ). Core agents k observe all successes immediately, and so have greater total information than peripherals who observe some successes with delay, $B_{k,t} + \min\{t, \tau_k\} > B_{\ell,t} + \min\{t, \tau_\ell\}$ for all $t > 0$. Lemma 7 in Appendix A.3 implies $\tau_k < \tau_\ell$.[27] Since peripherals experiment more, core agents have greater social learning, $B_{k,t} > B_{\ell,t}$ for all $t > 0$, and so $V_k > V_\ell$ by Lemma 1. □

[27] There is a subtlety here. Lemma 1 tells us that more social learning leads to less experimentation, but this is insufficient to conclude that core agents experiment less. For example, consider the star network and assume peripherals do not experiment; the core agent then has no social information but the same amount of total information as peripherals. Lemma 7 adapts the arguments from Lemma 1 to show that greater total learning (including self-learning) implies less experimentation.

We now characterize the equilibrium cutoffs. Core agents k observe all successes immediately, so their social learning follows $b_{k,t} \equiv (K-1)\mathbb{I}_{\{t\leq\tau_k\}} + L\mathbb{I}_{\{t\leq\tau_\ell\}}$. Experimentation incentives (4) are given by

$$\psi_{k,\tau_k} = P^{\emptyset}(I\tau_k)\left(x + y\left(1 - \left(1 - e^{-(r+L)(\tau_\ell-\tau_k)}\right)\frac{L}{r+L}\right)\right) - c \tag{12}$$

where the opportunity cost is the continuation value from having L peripherals experiment over $[\tau_k, \tau_\ell]$. In equilibrium, $\psi_{k,\tau_k} \leq 0$ with equality if $\tau_k > 0$.

Peripheral agents ℓ only observe the successes of core agents, so their social learning $b_{\ell,t}$ equals K before τ_k and then drops to Ka_t where $a_t := \Pr^H(T_{\ell'} < t \text{ for at least one } \ell' \neq \ell | t < T_\ell, t < T_k \text{ for all } k)$ is the conditional probability that some other peripheral agent has succeeded by t and hence core agents are working. This follows

$$\frac{\dot{a}}{1-a} = (L-1)\mathbb{I}_{\{t\leq\tau_\ell\}} - Ka = \begin{cases} L-1-Ka & t \in (\tau_k, \tau_\ell) \\ -Ka & t > \tau_\ell \end{cases} \tag{13}$$

with boundary condition $a_{\tau_k} = 1 - e^{-(L-1)\tau_k}$.[28] Before τ_ℓ, social learning a_t rises because of experimentation by the other $L-1$ peripherals, tempered by the lack of success by the K core agents. After τ_ℓ, only the latter effect remains, so social learning $b_{\ell,t} = Ka_t$ slows down. Using equation (4), peripherals' cutoff $\tau_\ell > 0$ then solves

$$\psi_{\ell,\tau_\ell} = P^{\emptyset}\left(K\left(\tau_k + \int_{\tau_k}^{\tau_\ell} a_t dt\right) + \tau_\ell\right)\left(x + ry\int_{\tau_\ell}^{\infty} e^{-\int_{\tau_\ell}^{t}(r+Ka_s)ds}dt\right) - c = 0.$$

For general, finite core-periphery networks we do not know if the cutoffs (τ_k, τ_ℓ) are unique.

In order to cleanly characterize how social information and welfare depend on the network density, we consider sequences of core-periphery networks which we index by $I \in \mathbb{N}$. Each network is determined by its core size K^I; the number of peripherals is then $L^I = I - K^I$. We assume the following two limits exist. Define $\kappa := \lim K^I \in \mathbb{N} \cup \{\infty\}$ as the limit of the absolute core size, and $\rho := \lim K^I/I \in [0,1]$ as the limit of the relative core size, as a proportion of the population. The set of network densities is the disjoint union $\mathbb{N} \cup \{\infty | 0 \cdot I\} \cup (0,1]$ which we endow with its natural order, concatenating the standard orders on

[28] To see (13), we apply Bayes' rule

$$1-a_t = \frac{\Pr^H(\forall k, \ell' : t < T_k, T_{\ell'})}{\Pr^H(t < T_\ell, \forall k : t < T_k)} = \begin{cases} \frac{\exp(-(K+L)t)}{\exp(-(K+1)t)} = \exp(-(L-1)t) & t < \tau_k \\ \frac{\exp(-K\tau_k - Lt))}{\exp\left(-K\left(\tau_k + \int_{\tau_k}^t a_s ds\right) - t\right)} = \exp\left(-(L-1)t + K\int_{\tau_k}^t a_s ds\right) & t \in (\tau_k, \tau_\ell) \\ \frac{\exp(-K\tau_k - L\tau_\ell))}{\exp\left(-K\left(\tau_k + \int_{\tau_k}^t a_s ds\right) - \tau_\ell\right)} = \exp\left(-(L-1)\tau_\ell + K\int_{\tau_k}^t a_s ds\right) & t > \tau_\ell \end{cases}$$

and then differentiate wrt t.

$n \in \mathbb{N} \cup \{\infty\}$ and $\rho \in [0, 1]$ at $\infty|0 \cdot I$.

We now define a threshold on core size that is critical for both asymptotic learning and welfare. For pessimistic priors $p_0 < \bar{p}$, define $\kappa^* \in (0, \infty)$ such that learning from κ^* core agents who experiment forever, $b_{\ell,t} \equiv \kappa^*$, renders a peripheral agent indifferent about experimenting at $t = 0$,

$$\psi_{\ell,0} = p_0 \left(x + \frac{r}{r + \kappa^*} y \right) - c = 0. \tag{14}$$

For optimistic priors, $p_0 \geq \bar{p}$, set $\kappa^* = \infty$.

Theorem 2. *In core-periphery networks $\{n^I\}$ and any equilibrium selection $\{\tau_k^I, \tau_\ell^I\}$:*

(a) Asymptotic information B is a decreasing function of network density: It attains asymptotic learning $B = \infty$ iff $\kappa \leq \kappa^*$, and strictly falls for $\kappa \geq \kappa^*$.

(b) Welfare V is a single-peaked function of network density: First, for $\kappa \leq \kappa^*$, it strictly rises, it attains the benchmark V^* iff $\kappa \in [\kappa^*, \infty]$, and then strictly falls in ρ.

Proof. See Appendix B.4. □

Asymptotic learning is achieved when the core is sufficiently small, and welfare attains the benchmark when core size is intermediate. Figure 6 illustrates Theorem 2. The top and middle panels sketch asymptotic information B and welfare V as functions of core size. The bottom panel illustrate three canonical social learning curves $\{B_{\ell,t}\}$. While asymptotic learning and the second-best welfare may *a priori* seem to be related goals, Theorem 2 shows that for pessimistic priors they are in fact mutually exclusive. This tension can be understood with help of the value function $\mathcal{V}(\tau, B)$ from Lemma 2. Value $\mathcal{V}$ falls in τ, so the welfare benchmark $V^* = \mathcal{V}(0, 0)$ requires peripherals' experimentation to vanish, $\tau_\ell^I \to 0$, undermining perfect learning (except in the knife-edge case $\kappa = \kappa^*$).[29]

As with random networks, there are three regions of network density with qualitatively different social learning dynamics. First, consider a small core $\kappa < \kappa^*$, as illustrated in Figure 6(i). The exploding number of peripherals experiment for a bounded time interval, $\tau_\ell > 0$, and collectively create an exploding amount of information in an instant. This crowds out experimentation by core agents. Peripherals choose to experiment since the flow of social information is clogged by the small core size. Formally, $B_{\ell,t} = Kt$ so equation (14) implies $\psi_{\ell,0} > 0$ given than $\kappa < \kappa^*$. Eventually the network aggregates information; but since each

[29] The "pessimistic prior" assumption is important. For optimistic priors, $p_0 \geq \bar{p}$, it is easier to motivate agents to experiment and our welfare benchmark requires asymptotic learning also holds. Moreover, both these goals are obtained if $\kappa = \infty$ and $\rho = 0$. This is a single point in our density order, but there any many sequences that satisfy both conditions (e.g. $K^I = \log I$, $K^I = (n^I)^{1/2}$).

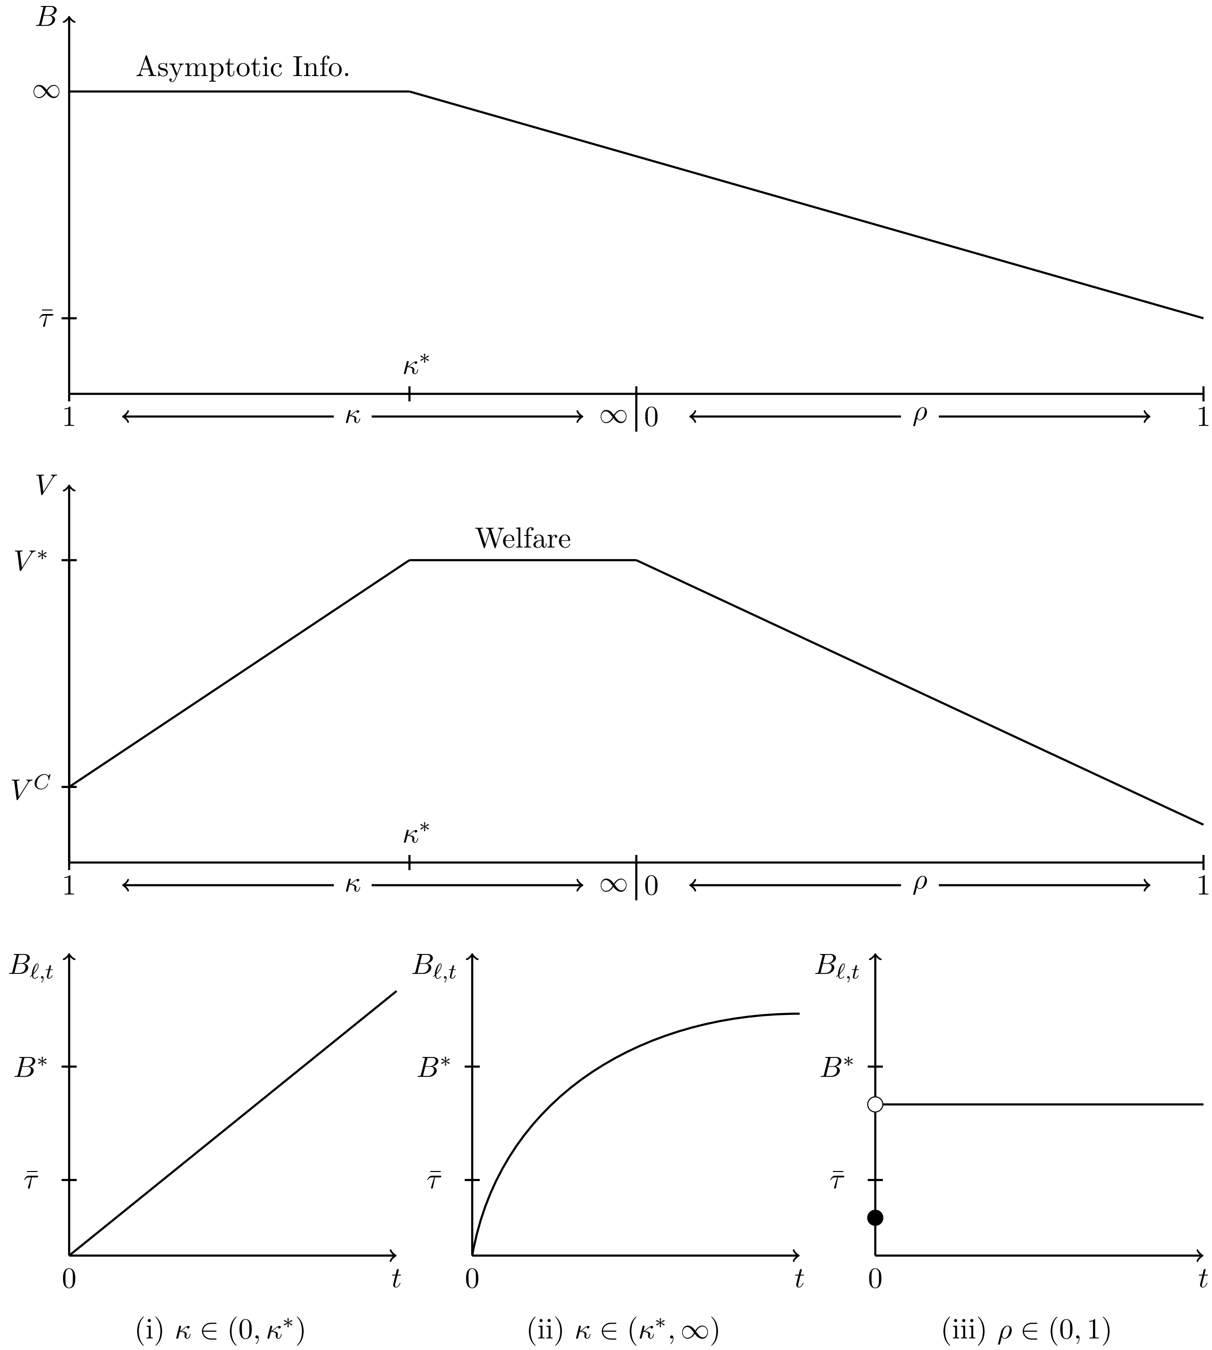

Figure 6: **Core-Periphery for Pessimistic Priors, $p_0 \leq \bar{p}$.** The **top** panel shows asymptotic information B as a function of network density, as described in Theorem 2(a). The **middle** panel shows welfare V as a function of network density, as described in Theorem 2(b). The **bottom** panel shows the learning curves of a peripheral agent in three canonical cases, as discussed in the text. In particular, B^* is defined by (11) and κ^* is defined by (14).

peripheral generates a non-vanishing amount of information, their utility falls short of the benchmark $\mathcal{V}(\tau_\ell, \kappa\tau_\ell) < V^*$.

Second, consider an intermediate core $\kappa \in (\kappa^*, \infty)$, as illustrated in Figure 6(ii).[30] With this core size, perfect information from peripherals would crowd out peripherals' experimentation incentives. In equilibrium, peripheral agents lower their cutoffs, limiting their total information $B = \lim L^I \tau_\ell^I < \infty$. The level of B is determined by peripheral agents' indifference condition at $t = 0$,

$$\psi_{\ell,0} = p_0 \left(x + ry \int_0^\infty e^{-\int_0^t (r + b_{\ell,s}) ds} dt \right) - c = 0$$

where ℓ's social learning curve satisfies

$$1 - e^{-B_{\ell,t}} = (1 - e^{-B})(1 - e^{-\kappa t}).$$

Intuitively, ℓ learns the state if some peripheral learned it and a core agent succeeds. As in the star, $b_{\ell,t}$ falls over time as agents grow pessimistic about the chance that one of them succeeded. Asymptotic learning fails, but agents do obtain the welfare benchmark, $\mathcal{V}(0,0)$, as pre-cutoff learning $(\kappa+1)\tau_\ell^I$ vanishes. For large κ, the core transmits information increasingly fast, reinforcing the crowding out and reducing asymptotic information. When $\kappa = \infty$ but $\rho = 0$, $B = B^*$ solves (11), so ℓ's social learning curve jumps to B^* and remains constant thereafter.

Third, consider a large core $\rho \in (0, 1]$, as illustrated by Figure 6(iii). Now core agents generate a non-vanishing share of total information. Social learning is asymptotically immediate, $B_{\ell,t} = B_{k,t} = B$ for all $t > 0$, and core agents' indifference condition becomes

$$P^{\emptyset}(B_{k,\tau}) \left(x + ye^{-(B - B_{k,\tau})} \right) - c = 0$$

with pre-cutoff learning $B_{k,\tau} = \lim I \tau_k^I$. This equation together with the analogous, but more involved expression for peripherals' pre-cutoff learning $\lim B^I_{\ell,\tau_\ell^I}$, pin down asymptotic information B, which is shown to fall in ρ with $B = \bar{\tau}$ for $\rho = 1$, as in the clique.

These results are reassuringly parallel to the ones for random networks in Section 4.2. In both cases, asymptotic information decreases in density, while welfare is single-peaked. However, significant differences arise from the higher ratio of second neighbors to first neighbors. First, the contrast between asymptotic information and welfare is starker: With $p_0 < \bar{p}$, asymptotic learning and second-best welfare are mutually exclusive under core-periphery networks, yet overlap under random networks. Second, social learning slows down over time in core-periphery networks with a finite core, as the information trickles through the core; in

[30] As in footnote 26, for optimistic priors, $p_0 \geq \bar{p}$, this region is empty.

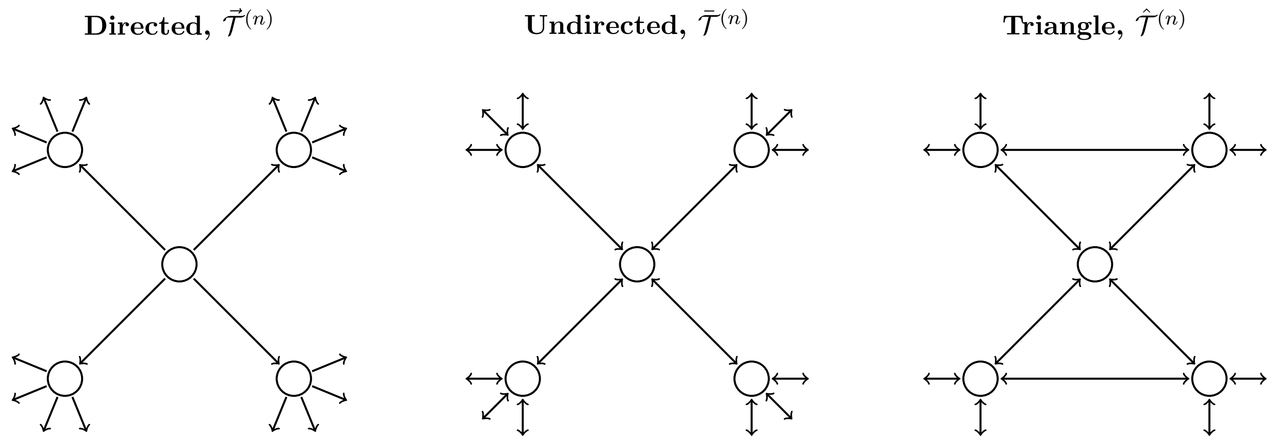

Figure 7: **Regular Trees** with degree $n = 4$.

contrast, social learning speeds up over time in random networks, as the number of indirect neighbors grows exponentially with path length.

5 Types of Links

In Section 4 we examined the effect of the number of links on experimentation and welfare; we now consider the effect of the type of links. We study this question in the context of regular trees, as illustrated in Figure 7. In a directed tree $\vec{\mathcal{T}}^{(n)}$, there is at most one directed path between any two agents; this resembles users following each other on Twitter. In a undirected tree $\bar{\mathcal{T}}^{(n)}$, there is at most one undirected path between any two agents; this resembles the connections between acquaintances on LinkedIn. And in a triangle tree $\hat{\mathcal{T}}^{(n)}$, agents are connected in triangles; this resembles clusters of friends on Facebook. Trees approximate large random networks; they are tractable because of the independence across neighbors, e.g. in a triangle tree, Iris's neighbors j and k have independent information given Iris has not observed a success. Proposition 4 shows that large random networks with finite degree converge to trees. Theorem 3 then shows that utility in a directed tree with n neighbors is greater than utility in an undirected tree with n neighbors, but less than an undirected tree with $n + 1$ neighbors. A similar comparison applies between undirected trees and triangle trees. For large n, these tight imply that the type of the link is of second-order importance for agents' behavior and utility.[31]

[31] In Board and Meyer-ter-Vehn (2021) we compared directed, undirected and tree networks in an adoption model, where each agent acts once. We showed that random networks converge to trees and compared welfare across networks, analogous to Theorem 3. In the experimentation model here, each agent is forward-looking and chooses an action at each time, so the proofs are novel.

5.1 Trees

We first consider *directed trees*, $\vec{\mathcal{T}}^{(n)}$, where every agent has n neighbors. Assuming symmetric cutoffs τ, let $a_t = E^H[A_{j,t}|t < S_i]$ be Iris's expectation of her neighbor John's effort in the contagion phase, $t \geq \tau$, in the absence of a success. As in equation (6) in Example 2,

$$1-a_t = \Pr{}^H\left(t < T_k, \forall k \in N_j | t < T_j\right) = \frac{\Pr^H\left(t < T_j, T_k, \forall k \in N_j\right)}{\Pr^H\left(t < T_j\right)} = \frac{\exp\left(-(n+1)\tau - n\int_\tau^t a_s ds\right)}{\exp\left(-\tau - \int_\tau^t a_s ds\right)}. \tag{15}$$

For the denominator, the hazard rate of John's first success time T_j equals 1 in the experimentation phase $t \leq \tau$ and a_t in the contagion phase $t > \tau$. For the numerator, the hazard rate of the success time $\min\{T_j, T_k, T_{k'}, \ldots\}$ equals $n+1$ in the experimentation phase $t \leq \tau$; in the contagion phase $t > \tau$, the lack of success by $\{T_k, T_{k'}, \ldots\}$ implies $A_{j,t} = 0$, so the hazard rate drops to a_t for each $k \in N_j$. Differentiating, Iris's belief follows the ODE

$$\dot{a}_t = (n-1)a(1-a) \tag{16}$$

with initial condition $a_\tau = 1 - e^{-n\tau}$ given by the probability that one of John's n neighbors succeeded in the experimentation phase.[32] In the case of the directed line in Example 2, $n = 1$, Iris's belief $a_t = 1 - e^{-n\tau}$ is constant over time. If $n \geq 2$, Iris's belief rises over time because the good news from John's expected inflow of information outweighs the bad news from his observed lack of success. The net effect is captured by the factor $(n-1)$ in (16): The more neighbors John has, the faster he observes a success, and the faster rises Iris's rate of social learning.

In order to study undirected and triangle trees we must address "backward" links, where agent i reasons about j who simultaneously learns from i's successes (or lack thereof).

Example 4 (Undirected Line). Consider the infinite undirected line

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

As in the directed line (Example 2), let a_t be i's expectation of j's effort at time $t > \tau$ conditional on not seeing a success; this coincides with i's expectation that k has succeeded:

$$a_t := E^H\left[A_{j,t}|t < T_i, S_i\right] = 1 - \Pr{}^H\left(t < T_k | t < T_i, T_j\right). \tag{17}$$

Calculating this conditional expectation is more subtle than in the directed line, where we

[32]Inverting, (16) admits the closed-form solution $a_t = 1/(1+\exp(-(d-2)(t+c)))$ with constant c.

could simply drop the "$t < T_i$" term. Now that j also observes i but has not seen i succeed, it is useful to introduce the expectation E^{-i} over others' success times $\{T_j\}_{j\neq i}$, given a symmetric cutoff τ, no successes of i, and $\theta = H$. Then, $a_t = E^{-i}[A_{j,t}|t < T_j]$ and

$$\Pr{}^H(t < T_k|t < T_i, T_j) = \Pr{}^{-i}(t < T_k|t < T_j) = \frac{\Pr^{-i}(t < T_j, T_k)}{\Pr^{-i}(t < T_j)} = \frac{\exp\left(-2\tau - \int_\tau^t a_s ds\right)}{\exp\left(-\tau - \int_\tau^t a_s ds\right)} = e^{-\tau}. \tag{18}$$

For the denominator, the hazard rate of John's first success time T_j equals 1 in the experimentation phase $t \leq \tau$ and a_t in the contagion phase $t > \tau$, under expectation E^{-i}. For the numerator, the hazard rate of the success time $\min\{T_j, T_k\}$ equals 2 in the experimentation phase $t \leq \tau$; in the contagion phase $t > \tau$, the lack of success by i, j, k implies $A_{j,t} = 0$, so the hazard rate drops to $E^H[A_{k,t}|t < T_j, S_j] = a_t$. Substituting back into (17) yields $a_t = 1 - e^{-\tau}$, just like in the directed line. Thus, the additional backward link does not affect the information i learns from j. Intuitively, while j has one more link, he is no more informed once we condition on i not having succeeded. That is, j's link to i does not help i since she cannot learn from herself.

In contrast to the directed line, i now has two neighbors, so using (4), the equilibrium stopping time τ solves

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

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

We now generalize this example and consider an *undirected tree* $\bar{\mathcal{T}}^{(n)}$ in which everyone has n neighbors. Adapting (15) and (18), Iris's belief of John's effort follows the ODE

$$\dot{a} = (n-2)a(1-a) \tag{19}$$

with initial condition $a_\tau = 1 - e^{-(n-1)\tau}$. In the case of the undirected line, $n = 2$, Iris's belief $a_t = 1 - e^{-n\tau}$ is constant over time, as in Example 4. When $n \geq 3$, Iris's belief increases over time. Compared to the directed tree, undirected links lower both the initial condition and the rate of increase of John's expected effort by one degree. Intuitively, i knows the backward link $j \to i$ does not generate information for j, because she conditions on not having seen a success herself.

Finally, consider a triangle tree, $\hat{\mathcal{T}}^{(n)}$. Adapting (15), Iris's belief of John's effort follows

$$\dot{a} = (n-3)a(1-a) \tag{20}$$

with initial condition $a_\tau = 1 - e^{-(n-2)\tau}$. When compared to the undirected tree, this lowers both the initial condition and the rate of increase by an additional neighbor. 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.

To see how this difference in social learning feeds back into the equilibrium cutoff τ, write experimentation incentives (4) as a function of social learning

$$\psi_\tau(\{na_t\}) = P^\emptyset((n+1)\tau)\left(x + ry\int_{s=\tau}^{\infty} e^{-\int_{t=\tau}^{s}(r+na_t)dt}ds\right) - c = 0. \tag{21}$$

Substituting the solutions of the ODEs (16), (19) and (20) for $\{a_t\}_{t\geq\tau}$ yields unique equilibrium cutoffs $\vec{\tau}, \bar{\tau}, \hat{\tau}$ with associated social learning curves $\{\vec{a}_t\}, \{\bar{a}_t\}, \{\hat{a}_t\}$.

Theorem 3. *Equilibrium cutoff times for regular trees are ranked as follows:*

$$\hat{\tau}^{(n+2)} < \bar{\tau}^{(n+1)} < \vec{\tau}^{(n)} < \bar{\tau}^{(n)} < \hat{\tau}^{(n)}.$$

Equilibrium values are ranked in the opposite way:

$$\hat{V}^{(n+2)} > \bar{V}^{(n+1)} > \vec{V}^{(n)} > \bar{V}^{(n)} > \hat{V}^{(n)}.$$

Proof. See Appendix C.1. □

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 $\vec{\mathcal{T}}^{(n)}$ has the same number of neighbors as the undirected network $\bar{\mathcal{T}}^{(n)}$, 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 $\bar{\mathcal{T}}^{(n+1)}$ has the same social information per neighbor as the directed network $\vec{\mathcal{T}}^{(n)}$ but more neighbors. Again, this extra social information provides value and crowds out the agent's effort.

This result is important for two reasons. First, it provides comparative statics across canonical networks in terms of experimentation, social learning, and welfare. In contrast, the rest of the literature typically focuses on long run considerations (e.g. whether the agents reach consensus or aggregate information). Second, it allows us to quantitatively assess the importance of network structure. This matters for network design (e.g. reducing clustering is useful, but adding connections is even better). It can also be useful for empirical work since we need not worry about exactly specifying the network (at least, within a class) if agents have many neighbors.

We end the section by showing that trees are indeed the limit of large random networks. The construction of the sequences of regular, random networks follows Section 4.2. For directed networks, suppose there are I agents, each of whom has n stubs that we randomly connect to other agents. Define the equilibrium cutoff $\vec{\tau}^I$, which is unique by Proposition 3, and social learning $\{\vec{b}_t^I\}_t$. For undirected networks, where we connect pairs of stubs into undirected links, define the equilibrium cutoff $\bar{\tau}^I$ and social learning $\{\bar{b}_t^I\}_t$. For triangle networks each agent has $n^I/2$ stub pairs; we randomly connect triples of stub pairs into triangles, and define the equilibrium cutoff $\hat{\tau}^I$ and social learning $\{\hat{b}_t^I\}_t$.

To see how these random networks approximate trees as $I \to \infty$, fix a "root agent" i and "uncover the network" from this root. That is, first randomly connect i's link stubs, then the stubs of i's n neighbors, and so on. For fixed n and $I \to \infty$, i's "local network" is almost surely a tree. The next result formalizes the idea that i's social learning only depends on her local network, and thus the network equilibria converge to the heuristic tree equilibria. The comparisons in Theorem 3 thus apply equally to large random networks.

Proposition 4. *The equilibria of the large random networks with degree $n^I \equiv n$ converge to the equilibria of the respective infinite regular n-trees:*

(a) *Directed networks: $\vec{\tau}^I \to \vec{\tau}$ and $\vec{b}_t^I \to n\vec{a}_t$ for all $t \neq \vec{\tau}$.*

(b) *Undirected networks: $\bar{\tau}^I \to \bar{\tau}$ and $\bar{b}_t^I \to n\bar{a}_t$ for all $t \neq \bar{\tau}$.*

(c) *Triangle networks: $\hat{\tau}^I \to \hat{\tau}$ and $\hat{b}_t^I \to n\hat{a}_t$ for all $t \neq \hat{\tau}$.*

Proof. See Appendix C.2. □

6 Discussion

This paper studies a simple model of experimentation in networks. We characterize individual experimentation, social learning curves, asymptotic information, and welfare in large random networks and core-periphery networks. While asymptotic information falls in network density, welfare is single-peaked. We also compare directed, undirected and clustered links in trees. Compared to the literature, we go beyond long-term outcomes by describing learning dynamics and welfare, and perform comparative statics across networks.

Our main result considers two canonical classes classes of networks, which admit natural density measures and give rise to unique limit behavior. But the economic forces underlying our analysis transcend these two classes. For example, the general conflict between learning and welfare for pessimistic priors $p_0 < \bar{p}$, is apparent from the welfare benchmark,

$V^* = \mathcal{V}(0,0)$, which requires individual experimentation to vanish, undermining asymptotic learning. The details of how these forces play out does depend on the structure of the network. For example, in core-periphery networks, asymptotic learning and second-best welfare are generically incompatible, but in random networks the larger diameter allows them to coexist for a range of intermediate network density.

This paper focuses on the role of networks in facilitating social learning. One can also use the model to study the impact of communication more directly, by assuming that agents observe each other imperfectly. The simplest such model has two agents, Iris and John, linked (undirectedly) with probability γ. Thus Iris and John do not know if they are experimenting by themselves, or if the other is working on the same problem. Over time, failure to observe a success makes Iris pessimistic about her chance of being linked to John, so social learning $b_{i,t} = E^H[j \in N_i(G) | t < T_i, S_i] = \frac{\gamma e^{-2t}}{\gamma e^{-2t} + (1-\gamma)e^{-t}}$ falls over $[0, \tau]$ before dropping to 0 forever. An increase rise in γ raises social learning $\{B_t\}$ for fixed τ, so lowers the equilibrium cutoff τ, and hence raises welfare $\mathcal{V}(\tau, \bar{\tau} - \tau)$.

A related idea is a random matching model, where a continuum of agents observe the current successes of n IID randomly sampled other agents at each t. When n is finite, the expected effort of i's random neighbor $\tilde{a}_t = E^H[A_t]$ follows $\dot{a} = na(1-a)$ with initial condition $a_\tau = 1 - e^{-(n+1)\tau}$, where the unique equilibrium cutoff $\tilde{\tau}^{(n)}$ solves (21) and induces value $\tilde{V}^{(n)}$. One can adapt the proof of Theorem 3 to show that random matching induces higher values than the directed tree network, but the additional value is less than one link: $\vec{V}^{(n+1)} > \tilde{V}^{(n)} > \vec{V}^{(n)}$ and $\vec{\tau}^{(n+1)} < \tilde{\tau}^{(n)} < \vec{\tau}^{(n)}$. Intuitively, random matching improves social information by avoiding auto-correlation of successes. But for large n, fixed networks are well approximated by random matching.

A third possibility is that $I \to \infty$ agents observe others' successes with a fixed delay $\sigma > 0$. Fix $p_0 < \bar{p}$ and define σ^* as in equation (8). When $\sigma > \sigma^*$, initial experimentation incentives are positive, so $\tau := \lim \tau^I > 0$ solves $P^\emptyset(\tau)(x + (1 - e^{-r(\sigma-\tau)})y) = c$. Agents learn perfectly at σ, but welfare is below second-best $\mathcal{V}(\tau, 0) < V^*$. Conversely, when $\sigma < \sigma^*$, perfect learning at σ would eliminate experimentation incentives for finite I, so total information B is finite and solves (10). Since $\tau^I \approx B/I \to 0$, welfare is second-best $\mathcal{V}(0,0) = V^*$. While this analysis is broadly parallel to the large random networks, the latter derives the delay σ endogenously and consequently features the wide range of network densities that achieve both benchmarks.

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.

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.

——— (2017): "Learning to Disagree in a Game of Experimentation," *Journal of Economic Theory*, 169, 234–269.

Bramoullé, Y., and R. Kranton (2007): "Public Goods in Networks," *Journal of Economic Theory*, 135(1), 478–494.

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., AND S. GOYAL (2010): "The Law of the Few," *American Economic Review*, 100(4), 1468–92.

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.

HAYEK, F. (1945): "The Use of Knowledge in Society," *American Economic Review*, 35(4), 519–530.

HODGSON, C. (2021): "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.

LI, D., AND N. SCHÜRHOFF (2019): "Dealer networks," *Journal of Finance*, 74(1), 91–144.

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

——— (2020b): "Innovation Adoption and Collective Experimentation," *Games and Economic Behavior*, 120, 121–131.

SALISH, M. (2015): "Learning Faster or More Precisely? Strategic Experimentation in Networks," Working paper, SSRN.

A Appendix: Proofs from Section 3

A.1 Proof of Proposition 1 (Cutoff strategies)

To formalize the discussion surrounding the statement of Proposition 1, we write Iris's payoff from an arbitrary experimentation strategy as $\Pi = \Pi(\{a_s\}, \{b_s\})$. We will show that front-loading incentives are positive, equal to

$$-\frac{d}{dt}\frac{\partial \Pi}{\partial a_t} = \big(r\left(p_t(x+y)-c\right) + p_t b_t(x-c)\big)e^{-\int_0^t r+p_u(a_u+b_u)du} \tag{22}$$

As discussed before the proposition statement, $r\left(p_t(x+y)-c\right)$ is the time-value of front-loading own experimentation from $t+\delta$ to t, while $p_t b_t(x-c)$ captures the value of additional experimentation that arises when a neighbor succeeds in $[t, t+\delta]$; finally, the discount-term $e^{-\int_0^t r+p_u(a_u+b_u)du}$ reflects that (22) evaluates the effect of front-loading time-t effort from the time-0 perspective. Equation (22) implies that agents maximally front-load their effort, and so cutoff strategies are optimal.

To establish the second derivative (22), we first derive convenient expression for Π and its various first derivatives. Truncating (3) at time-t, we get

$$\Pi = \int_0^t \big(p_s(a_s(x+y)+b_s y) - a_s c\big)e^{-\int_0^s r+p_u(a_u+b_u)du}ds + \Pi_t e^{-\int_0^t r+p_u(a_u+b_u)du}. \tag{23}$$

We next establish two convenient expressions for the continuation payoff

$$\Pi_t = \int_t^\infty \big(p_s(a_s(x+y)+b_s y) - a_s c\big)e^{-\int_t^s r+p_u(a_u+b_u)du}ds \tag{24}$$

$$= \int_t^\infty e^{-r(s-t)}\left(p_t e^{-\int_t^s (a_u+b_u)du}\big(a_s(x+y-c)+b_s y\big) - (1-p_t)a_s c\right)ds \tag{25}$$

Equation (24) follows the same logic as (3): the discounted chance of no success on $[t,s]$ is $e^{-\int_t^s r+p_u(a_u+b_u)du}$; at time-$s$ we condition on θ and realize expected flow benefits $a_s(x+y)+b_s y$ if $\theta = H$ and flow costs $a_s c$ for either θ. Equation (25) conditions on the true state already at time-t. For $\theta = H$ no success arrives by s with probability $e^{-\int_t^s (a_u+b_u)du}$, and net flow benefits are $a_s(x+y-c)+b_s y$; for $\theta = L$, no success arrives ever, and Iris incurs flow costs $a_s c$. Write $\alpha_t := \int_t^\infty e^{-r(s-t)}\left(a_s c\right)ds$ with time-derivative $\dot{\alpha}_t = r\alpha_t - a_t c$. By (25), $\Pi_t + \alpha_t$ is a linear function of the posterior belief p_t, and so

$$\frac{\partial \Pi_t}{\partial p_t} = \frac{1}{p_t}\left(\Pi_t + \alpha_t\right) \tag{26}$$

To compute $\partial \Pi/\partial a_t$, define the derivative of the posterior belief $p_t = P^\emptyset\left(\int_0^t (a_s+b_s)ds\right)$

with respect to "experimentation just before t",

$$\frac{\partial p_t(\{a_s\}_{s\geq 0})}{\partial a_t} = \lim_{\epsilon\to 0}\frac{1}{\epsilon}\left(p_t(\{a_s^{t,\epsilon}\}_{s\geq 0}) - p_t(\{a_s\}_{s\geq 0})\right) = -p_t(1-p_t)$$

where $a_s^{t,\epsilon} := a_s + \mathbb{I}_{\{s\in[t-\epsilon,t]\}}$. Similarly differentiating payoff (23) wrt a_t and using (26),

$$\begin{aligned}\frac{\partial \Pi}{\partial a_t} &= \left(p_t(x+y) - c + \frac{\partial \Pi_t}{\partial p_t}\frac{\partial p_t}{\partial a_t} - p_t\Pi_t\right) e^{-\int_0^t r + p_u(a_u+b_u)du} \\ &= \big(p_t(x+y) - c - (1-p_t)\alpha_t - \Pi_t\big)e^{-\int_0^t r+p_u(a_u+b_u)du} \qquad (27)\end{aligned}$$

Turning to the time-derivatives, we first note $\dot{p}_t = -(a_t+b_t)p_t(1-p_t)$, differentiate (24)

$$\dot{\Pi}_t = -\big(p_t(a_t(x+y) + b_t y - a_t c\big) + \big(r + p_t(a_t+b_t)\big)\Pi_t.$$

and then (27)

$$\begin{aligned} e^{\int_0^t r+p_u(a_u+b_u)du}\frac{d}{dt}\frac{\partial \Pi}{\partial a_t} =& - p_t(1-p_t)(a_t+b_t)(x+y+\alpha_t) - (1-p_t)(r\alpha_t - a_t c) \\ &+ \big(p_t(a_t(x+y)+b_t y) - a_t c\big) - \big(r + p_t(a_t+b_t)\big)\Pi_t \\ &- (r + p_t(a_t+b_t))\big(p_t(x+y) - c - (1-p_t)\alpha_t - \Pi_t\big) \\ =& - r\left(p_t(x+y) - c\right) - p_t b_t(x-c)\end{aligned}$$

which is (22).

Having established that cutoff strategies are optimal, we now show that the optimal cutoff is the unique solution of $\psi_\tau = 0$. For cutoff strategies $a_s = \mathbb{I}_{\{s\leq t\}}$ we have $\alpha_t = \int_t^\infty e^{-r(s-t)}(a_s c)\,ds = 0$ and (25) simplifies to

$$\Pi_t = p_t y\int_t^\infty e^{-r(s-t)} b_s e^{-\int_t^s b_u du} ds = p_t y\left(1 - r\int_t^\infty e^{-\int_t^s (r+b_u)du} ds\right),$$

where the last equality uses integration by parts. Then (27) simplifies to

$$e^{\int_0^t r+p_u(a_u+b_u)du}\frac{\partial \Pi(\mathbb{I}_{\{s\leq t\}})}{\partial a_t} = p_t(x+y) - c - \Pi_t = p_t\left(x + ry\int_t^\infty e^{-\int_t^s (r+b_u)du} ds\right) - c = \psi_t. \qquad (28)$$

Differentiating the LHS wrt t, we see that ψ_t strictly single-crosses from above since

$$\dot{\psi}_t = (r + p_t(a_t+b_t))\psi_t + e^{\int_0^t r+p_u(a_u+b_u)du}\frac{d}{dt}\frac{\partial \Pi(\mathbb{I}_{\{s\leq t\}})}{\partial a_t}$$

is negative whenever $\psi_t = 0$ by (22). □

For future reference we summarize some properties of

$$\psi_\tau(\{B_t\}) = P^\emptyset(\tau + B_\tau)\left(x + ry\int_t^\infty e^{-r(s-\tau)-(B_s-B_\tau)}ds\right) - c. \tag{29}$$

First, note that while (4) and (5) express ψ instead as a function of the social learning rate $\{b_t\}$, the definition and most properties of ψ extend to any increasing (not necessarily continuous or positive) cumulative social learning curve B_t.

Lemma 6. *Properties of $\psi_\tau(\{B_t\})$.*

(a) (29) falls in $\{B_t\}$, and thus also in $\{b_t\}$ with partial derivative given in (5).

(b) (29) strictly single-crosses from above in τ, and is equi-Lipschitz continuous in τ for all uniformly bounded $\{b_t\}$.

(c) The root τ of $\psi_\tau = 0$ falls in $\{B_t\}$, and strictly falls in $\{b_t\}$.

A.2 Proof of Lemma 2

We can write the agent's value function as

$$\begin{aligned} V &= \left(p_0\int_0^\tau e^{-\int_0^t(r+b_s+1)ds}(x+(b_t+1)y-c)dt\right) - \left((1-p_0)\int_0^\tau e^{-rt}cdt\right) + e^{-r\tau}\left(p_0e^{-B_\tau-\tau} + (1-p_0)\right)V_\tau \\ &= p_0y(1-e^{-B_\tau-(r+1)\tau}) - (1-p_0)c\frac{1-e^{-r\tau}}{r} + e^{-r\tau}\left(p_0e^{-B_\tau-\tau}(x+y-c) - (1-p_0)c\right) \\ &= \frac{p_0x-c}{r} + e^{-r\tau}\left(p_0e^{-B_\tau-\tau}(x-c) - (1-p_0)c\frac{r-1}{r}\right) =: \mathcal{V}(\tau, B_\tau) \end{aligned} \tag{30}$$

The first line conditions on θ at time-0 and truncates flow payoffs (3) at $t = \tau$. The second line evaluates the first integral using $x - c = ry$, and the last term using $p_0e^{-B_\tau-\tau} + (1-p_0) = p_0e^{-B_\tau-\tau}/p_\tau$ by Bayes' rule, and $V_\tau = p_\tau y\int_\tau^\infty b_t \exp^{-\int_t^\infty(r+b_s)ds} dt = p_\tau(x+y) - c$ (using $\psi_\tau = 0$). The last line uses $y = (x-c)/r$ and reorders terms.

The monotonicity in B_τ is immediate from (30). To see the monotonicity in τ, note that the first term in (30) is the payoff from experimenting forever. Thus, the second term is the opportunity value of stopping earlier, which must be positive. Then

Testing

$$\begin{aligned} \partial_\tau\mathcal{V} &= -re^{-r\tau}\left(p_0e^{-B-\tau}(x-c) - (1-p_0)c\frac{r-1}{r}\right) - e^{-r\tau}p_0e^{-B-\tau}(x-c) \\ &< -e^{-r\tau}p_0e^{-B-\tau}(x-c) = \partial_B\mathcal{V} < 0. \end{aligned} \tag{31}$$

A.3 Proof of Proposition 3 (Equal Cutoffs of Equals)

We first establish two Lemmas that are of some independent interest and clarify the proof logic. For social learning $\{B_t\}$ and the associated optimal cutoff τ, define total learning $B_t + \min\{t, \tau\}$. So defined, $\Pr^H(\min\{S, T\} \le t) = 1 - \exp(-(B_t + \min\{t, \tau\}))$.

Lemma 7. *Higher total learning, $B_t + \min\{\tau, t\} \ge \hat{B}_t + \min\{\hat{\tau}, t\}$ for all t, is associated with lower cutoffs, $\tau \le \hat{\tau}$.*

This is closely related to Lemma 1, that lower social learning $\{B_t\} \le \{\hat{B}_t\}$ implies higher cutoffs $\tau \ge \hat{\tau}$. Lemma 7 shows additionally that the higher cutoff cannot lead to higher total learning. Intuitively, all learning (both social and own) crowds out incentives.

Lemma 8. *Fix a network G, cutoffs $\{\tau_k\}_{k \ne i,j}$ and $\tau_* < \tau^*$, and write k's first success time as $\{T_k\}$ if $\tau_i = \tau^*, \tau_j = \tau_*$, and $\{T'_k\}$ if $\tau_i = \tau_*, \tau_j = \tau^*$.* $\min\{T_i, S_i\} \overset{D}{\preceq} \min\{T'_i, S'_i\}$.[33]

Lemma 8 is intuitive: Additional experimentation during $[\tau_*, \tau^*]$ is more immediate and useful to i when done by i herself instead of j.

Proof of Proposition 3. By contradiction, assume $\tau_i > \tau_j$. Symmetry, $G_{i \leftrightarrow j} = G$, implies $\min\{T_j, S_j\} \overset{D}{=} \min\{T'_i, S'_i\}$. Lemma 8 then implies $\min\{T_i, S_i\} \overset{D}{\preceq} \min\{T_j, S_j\}$. Noting the connection between total learning and the time of the first observed success, $\Pr^H(\min\{S, T\} \le t) = 1 - \exp(-(B_t + \min\{\tau, t\}))$, this implies $\{B_{i,t} + \min\{\tau_i, t\}\} \ge \{B_{j,t} + \min\{\tau_j, t\}\}$ and so, by Lemma 7, $\tau_i \le \tau_j$. □

Proof of Lemma 7. Lemmas 1 and 6 study incentives ψ_τ as a function of social learning $\{B_t\}$; we now study ψ_τ as a function of total learning $\{B_t + \min\{t, \tau\}\}$.

By contradiction assume that $B_t + \min\{\tau, t\} \ge \hat{B}_t + \min\{\hat{\tau}, t\}$ for all t, yet $\tau > \hat{\tau}$. Define $\tilde{B}_t := \hat{B}_t - (\tau - \hat{\tau})$; clearly $\tilde{B}_t \le B_t$, and so

$$\psi_\tau(\{\tilde{B}_t\}) \ge \psi_\tau(\{B_t\}) = 0.$$

Since $\tilde{B}_\tau + \tau = \hat{B}_\tau + \hat{\tau}$ and $\tilde{b}_u = \hat{b}_u$ for $u \ge \tau$, time-τ experimentation incentives for the social learning curve $\{\hat{B}_t\}$ are also positive

$$e^{\int_0^\tau r + p_u(\hat{a}_u + \hat{b}_u)du} \frac{\partial \hat{\Pi}(\mathbb{I}_{\{t \le \hat{\tau}\}})}{\partial a_\tau} = P^\emptyset(\hat{B}_\tau + \hat{\tau}) \left(x + ry \int_\tau^\infty e^{-\int_\tau^s (r + \hat{b}_u)du} ds \right) - c = \psi_\tau(\{\tilde{B}_t\}) \ge 0$$

[33]As always, $S_i = \min_{j \in N_i(G)}\{T_j\}$ and $S'_i = \min_{j \in N_i(G)}\{T'_j\}$.

where the first equality follows as in (28), using $\hat{a}_u = 0$ at $u \geq \tau$ since $\tau > \hat{\tau}$. Front-loading, (22), then implies

$$\frac{\partial \hat{\Pi}(\mathbb{I}_{\{t \leq \hat{\tau}\}})}{\partial a_{\hat{\tau}}} > \frac{\partial \hat{\Pi}(\mathbb{I}_{\{t \leq \hat{\tau}\}})}{\partial a_{\tau}} \geq 0$$

contradicting the optimality of cutoff $\hat{\tau}$. □

Proof of Lemma 8. As a baseline, write $\bar{T}_k$ for k's first success time in network G when i and j both use cutoff τ_*. For each realization of $(\bar{T}_i, \bar{S}_i)$, we dynamically realize $\{T_k, T'_k\}_k$ as follows. In a first step, raising τ_i (or τ_j) from τ_* to τ^* begets new success opportunities on $[\tau_*, \min\{\tau^*, \bar{T}_i, \bar{S}_i\}]$ (successes after $\min\{\bar{T}_i, \bar{S}_i\}$ have already been realized). Thus, we draw an exponential random variable $Z \sim Exp(1)$, and set

$$T_i, T'_j = \begin{cases} \tau_* + Z & \text{if } \tau_* + Z \leq \min\{\tau^*, \bar{T}_i, \bar{S}_i\}, \\ \bar{T} & \text{otherwise.} \end{cases}$$

In subsequent steps, we trace the effects of additional successes in the first step through the network. Since this cascade starts at $\tau_* + Z$ (if at all) and successes are not instantaneous, we have $T_k \in (\tau_* + Z, \bar{T}_k]$ for all $k \neq i$ and $T'_k \in (\tau_* + Z, \bar{T}_k]$ for all $k \neq j$.

So defined, if $\tau_* + Z > \min\{\tau^*, \bar{T}_i, \bar{S}_i\}$, no additional successes realize, so $T_k = T'_k = \bar{T}_k$ for all k; a fortiori $\min\{T_i, S_i\} = \min\{T'_i, S'_i\} = \min\{\bar{T}_i, \bar{S}_i\}$. If $\tau_* + Z \leq \min\{\tau^*, \bar{T}_i, \bar{S}_i\}$ we have $\min\{T_i, S_i\} = \tau_* + Z \leq \min\{T'_i, S'_i\}$ with equality iff j is a neighbor of i. All told, $\min\{T_i, S_i\} \overset{D}{\preceq} \min\{T'_i, S'_i\}$ with equality iff j is a neighbor of i. □

B Appendix: Proofs from Section 4

B.1 Proof of Lemma 3

Part (a): We will show separately that for every $\epsilon > 0$

$$\Pr\left[N^I \geq (1+\epsilon)I(1 - e^{-\hat{n}^I/I})\right] \to 0, \tag{32}$$

$$\Pr\left[N^I \leq (1-\epsilon)I(1 - e^{-\hat{n}^I/I})\right] \to 0. \tag{33}$$

This implies that the number of links converges to $1 - e^{-\hat{n}^I/I}$ in distribution, $N^I/(I(1 - e^{-\hat{n}^I/I})) \overset{D}{\to} 1$, and a fortiori in expectation, $n^I/(I(1 - e^{-\hat{n}^I/I})) \to 1$.

Start with the upper bound, (32). We can restrict attention to $\hat{\rho} = \lim \hat{n}^I/I < \infty$; for $\hat{\rho} = \infty$, we have $(1+\epsilon)I(1 - e^{-\hat{n}^I/I}) > I$ for any $\epsilon > 0$ and large enough I, so trivially

$\Pr[N^I \geq (1+\epsilon)I(1-e^{-\hat{n}^I/I})] = 0$.

Realize Iris's $\hat{n}^I$ stubs k one after another, and keep track of the number of stubs $K^I(m)$ used to reach degree m; if i has less than m neighbors set $K^I(m) := \hat{n}^I + 1$. When connecting Iris's k^{th} stub to her m^{th} neighbor, $I-m$ potential new neighbors with $\hat{n}^I(I-m)$ stubs compete with $\hat{n}^I m - (2k-1)$ remaining stubs of Iris and her $m-1$ neighbors, sandwiching the success rate between $\frac{I-m}{I}$ and $\frac{I-m}{I-2}$. Writing X_ℓ^I for independent (shifted) geometric random variables with success rate $\frac{I-\ell}{I}$ we can thus upper-bound $K^I(m) \overset{D}{\preceq} \sum_{\ell=1}^m X_\ell^I$.

The chance of m or more neighbors is then upper-bounded by

$$\begin{aligned}\Pr\left[N^I \geq m\right] &= \Pr\left[K^I(m) \leq \hat{n}^I\right] \leq \Pr\left[\sum_{\ell=1}^m X_\ell^I \leq \hat{n}^I\right] \leq \inf_{\xi\geq 0} \exp\left(\xi \hat{n}^I + \sum_{\ell=1}^m \log E[e^{-\xi X_\ell^I}]\right) \\ &= \inf_{\xi \geq 0} \exp\left(\xi(\hat{n}^I - m) - \sum_{\ell=1}^m \log \frac{1-e^{-\xi}\ell/I}{1-\ell/I}\right) \qquad (34)\end{aligned}$$

where the second inequality is a Chernoff-bound, and the final equality evaluates the moment generating function of the shifted geometric distribution, $E[e^{-\xi X_\ell^I}] = \frac{e^{-\xi}(1-\ell/I)}{1-e^{-\xi}\ell/I}$.

Since $\log \frac{I-e^{-\xi}\ell}{I-\ell}$ rises in ℓ, the last term in (34) is lower-bounded by

$$\begin{aligned}\sum_{\ell=1}^m \log \frac{1-e^{-\xi}\ell/I}{1-\ell/I} &\geq \int_0^m \left(\int_{1-\ell/I}^{1-e^{-\xi}\ell/I} \frac{1}{x}dx\right) d\ell = \int_{1-m/I}^1 \left(\int_{I(1-x)}^{\min\{e^\xi I(1-x), m\}} \frac{1}{x} d\ell\right) dx \\ &= \int_{1-m/I}^{1-e^{-\xi}m/I} \frac{m - I(1-x)}{x} dx + \int_{1-e^{-\xi}m/I}^1 \frac{I(1-x)(e^\xi - 1)}{x} dx \\ &= I\left[(1-m/I)\log(1-m/I) - e^\xi\left(1-e^{-\xi}m/I\right)\log(1-e^{-\xi}m/I)\right].\end{aligned}$$

For any $\epsilon > 0$, we now set $m = m^I := \left\lceil (1+\epsilon)I(1-e^{-\hat{n}^I/I})\right\rceil$, substitute back into the term in parentheses in (34), and divide by I

$$\xi\frac{\hat{n}^I - m^I}{I} - (1-m^I/I)\log(1-m^I/I) + e^\xi\left(1-e^{-\xi}m^I/I\right)\log(1-e^{-\xi}m^I/I) =: \Gamma^I(\xi,\epsilon)$$

with limit $\Gamma(\xi,\epsilon) = \lim \Gamma^I(\xi,\epsilon)$. So defined, (34) becomes

$$\Pr\left[N^I \geq (1+\epsilon)I(1-e^{-\hat{n}^I/I})\right] \leq \inf_{\xi\geq 0} \exp\left(I\Gamma^I(\xi,\epsilon)\right) \qquad (35)$$

The derivative $\Gamma_\xi(0,\epsilon) = \hat{\rho} + \log(1-(1+\epsilon)(1-e^{-\hat{\rho}}))$ vanishes for $\epsilon = 0$ and falls in ϵ. Thus, for any $\epsilon > 0$ we have $\Gamma_\xi(0,\epsilon) < 0$. Also, $\Gamma(0,\epsilon) = 0$, and so $\Gamma(\xi,\epsilon) < 0$ for small ξ, and $\Gamma^I(\xi,\epsilon) < 0$ uniformly for large I. Thus, (35) vanishes for $I \to \infty$, implying the upper bound (32).

The lower bound (33) follows analogously.

Part (b): Recall that $\hat{\rho} = \lim \hat{n}^I/I$ and, from part (a), $I(1 - \exp^{-\hat{n}^I/I})/n^I \to 1$, allowing us to asymptotically ignore the difference between n^I and

$$I(1 - \exp^{-\hat{n}^I/I}) = I\left(\frac{\hat{n}^I}{I} - \frac{1}{2}\left(\frac{\hat{n}^I}{I}\right)^2 + \frac{1}{6}\left(\frac{\hat{n}^I}{I}\right)^3 - ...\right).$$

Thus, $\lim n^I/I = 1 - e^{-\hat{\rho}}$. And $n^I/\hat{n}^I - 1$ is of order $\hat{n}^I/I$, which vanishes for $\hat{\rho} = 0$.

Part (c): Since $A_t = 1$ for $t < \tau$, we have $B^I_{\tau^I} = \int_0^{\tau^I} b^I_t dt = \int_0^{\tau^I} E^H[N^I|t < T^I, S^I]dt$. For I finite, $E^H[N^I|t < T^I, S^I] < n^I$ (and so $B^I_{\tau^I} < n^I\tau^I$) because lack of success, $t < T^I, S^I$, indicates fewer neighbors N^I. To bound the effect of such updating, we note that conditional on $|N^I - n^I| \le \epsilon n^I$, and so $N^I \le (1+\epsilon)n^I$, we have $\Pr^H(t < T^I, S^I) \ge e^{-((1+\epsilon)n^I+1)t}$. Thus

$$\frac{\Pr^H(|N^I - n^I| \le \epsilon n^I|t < T^I, S^I)}{\Pr^H(|N^I - n^I| \ge \epsilon n^I|t < T^I, S^I)} \ge \frac{\Pr^H(|N^I - n^I| \le \epsilon n^I)}{\Pr^H(|N^I - n^I| \ge \epsilon n^I)} e^{-((1+\epsilon)n^I+1)t}. \tag{36}$$

We show below that $n^I\tau^I$ is bounded. This bounds $e^{-((1+\epsilon)n^I+1)t}$ away from 0 for all $t \le \tau^I$. Thus, as the prior likelihood-ratio of $|N^I - n^I| \le \epsilon n^I$ on the RHS of (36) diverges as $I \to \infty$ (by part (a)), so does the posterior likelihood-ratio on the LHS of (36), implying $E^H[N^I|t < T^I, S^I]/n^I \to 1$ and so $B^I_{\tau^I}/(n^I\tau^I) \to 1$, finishing the proof of part (c).

To show that $n^I\tau^I$ is bounded, assume it was not. Then we could choose $\hat{\tau}^I < \tau^I$ such that $n^I\hat{\tau}^I$ is bounded, but with limit $\lim n^I\hat{\tau}^I > \bar{\tau}$. Applying the above argument to $n^I\hat{\tau}^I$ instead of $n^I\tau^I$, we get $\lim B^I_{\hat{\tau}^I} = \lim n^I\hat{\tau}^I > \bar{\tau}$, leading to the contradiction that $p_{\tau^I} < p_{\hat{\tau}^I} = P^\emptyset(B^I_{\hat{\tau}^I} + \hat{\tau}^I) < \underline{p}$ for large I.

B.2 Proof of Theorem 1

For finite degrees $\nu < \infty$, the proof of Theorem 1 relies on results in Section 5. Proposition 4 shows that the component size of a typical node explodes with I; Theorem 3 shows that equilibrium cutoffs $\tau = \lim \tau^I > 0$ (and fall in ν) so $B = \infty$, and that V rises in ν.

For $\nu = \infty$, equilibrium cutoffs must vanish, $\tau^I \to 0$; otherwise the posterior-belief at the cutoff $P^\emptyset((n^I + 1)\tau^I) \to 0$, choking off experimentation incentives. The accounting identity, Lemma 4 characterizes the limit of social learning curves as step functions $B_t = B\mathbb{I}_{\{t\ge\sigma\}}$; if $\sigma = 0$, we can distinguish two bursts, with pre-cutoff information $B_\tau := \lim n^I\tau^I$, and post-cutoff information $B - B_\tau$. The equilibrium indifference condition becomes

$$\psi_\tau = \lim \psi^I_{\tau^I} = p(B_\tau)\left(x + ry\int_0^\infty e^{-rt+B_t}dt\right) - c$$
$$= p(B_\tau)\Big(x + \big(1 - e^{-r\sigma}(1 - e^{-(B-B_\tau)})y\big)\Big) - c = 0. \quad (37)$$

To solve for B_τ, B, σ we complement (37) with the simple observation that

$$\frac{B_\tau}{B} = \frac{\lim n^I\tau^I}{\lim I\tau^I} = \lim\frac{n^I}{I} = \rho, \quad (38)$$

and two conditions linking the learning time $\sigma = \lim\frac{-\log\tau^I}{n^I}$ to pre-cutoff learning B_τ and total learning B: First, bounded total learning $B < \infty$, implies the learning time equals the network's time-diameter, $\sigma = 1/\lambda$, as seen in (9). Second, non-vanishing pre-cutoff learning implies immediate learning

$$\text{If } B_\tau = \lim n^I\tau^I > 0, \text{ then } \sigma = \lim\frac{\log(n^I\tau^I) - \log\tau^I}{n^I} = 0. \quad (39)$$

Case 1: $\rho = 0$ and $p_0 > \bar{p}$. The optimistic prior together with the equilibrium condition (37) require non-vanishing pre-cutoff learning, $B_\tau > 0$, and so by (39) immediate learning, $\sigma = 0$. The sparsity of the network together with (38) implies perfect learning $B = \infty$. Perfect immediate learning, $B = \infty, \sigma = 0$, in turn implies the welfare benchmark $V = p_0y = V^*$.

Case 2: $\rho = 0$ and $p_0 \leq \bar{p}$. We first observe $B_\tau = 0$. Otherwise, if $B_\tau > 0$, the proof for Case 1 implies $B = \infty$ and $\sigma = 0$, and so experimentation incentives $\psi_\tau < p_0x - c \leq 0$, contradicting equilibrium. This implies the welfare benchmark, $\lim\mathcal{V}(\tau^I, n^I\tau^I) = \mathcal{V}(0,0) = V^*$.

Turning to asymptotic information, we now show that B attains the benchmark ∞ iff $\lambda \leq 1/\sigma^*$, and strictly decrease above. For sparse networks $\lambda \leq 1/\sigma^*$,[34] if by contradiction learning was imperfect $B < \infty$, social learning happens too late, at $\sigma = 1/\lambda \geq \sigma^*$ by (9), so experimentation incentives are strictly positive

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

contradicting equilibrium.

Conversely, for dense networks $\lambda > 1/\sigma^*$, the social learning time $\sigma = \lim\frac{\log I - \log I\tau^I}{n^I} \leq \lim\frac{\log I}{n^I} = 1/\lambda$ is before σ^*, and equilibrium indifference

$$p_0\left(x + (1 - e^{-r\sigma}(1 - e^{-B}))y\right) - c = \psi_\tau = 0 = p_0\left(x + (1 - e^{-r\sigma^*})y\right) - c$$

[34]For $p_0 = \bar{p}$, we have $\sigma^* = 0$, so this condition is always satisfied.

requires $B < \infty$. Moreover, $B = B(\sigma)$ rises in σ, and so falls in $\lambda = 1/\sigma$.

Case 3: $\rho > 0$. Then $B_\tau = \hat{\rho}B > 0$; (39) then implies $\sigma = 0$, and so (37) becomes

$$p(\hat{\rho}B)\left(x + e^{-(1-\hat{\rho})B}y\right) = c.$$

In the notation of (51) in the proof of Theorem 2, this means $\Psi(\hat{\rho}B, (1-\hat{\rho})B) = 0$, so by (52), information $B = B(\hat{\rho})$ falls in $\hat{\rho}$. Then also welfare $V = p_0(1 - e^{-B})$ falls in $\hat{\rho}$.

B.3 Proof of Lemma 4

With probability $e^{-I\tau^I}$, no agent succeeds by τ^I, and so $S^I = \infty$; from here on we condition on the complementary event that at least one agent succeeds during experimentation, triggering a contagion process. For now, we also restrict attention to $\lim \hat{n}^I/I = 0$, so that $\hat{n}^I/n^I \to 1$ by Lemma 3(a). This allows us to work with $\hat{n}^I$ for finite I, but switch to n^I in the limit where $\sigma := \lim \frac{-\log \tau^I}{n^I}$. We discuss the case $\lim \hat{n}^I/I > 0$ later.

The overarching proof strategy is to separate the "geographical"/network aspects of the contagion process from its timing. Specifically, we realize the randomness of the network G^I as agents succeed. To emphasize the analogy to epidemiological SI contagion processes, we refer to agents who have succeeded as *infected.* When k agents are infected, let E_k^I be the random number of *exposed agents*, i.e. that have observed a success but have yet to succeed themselves. Clearly $E_k^I \leq \hat{n}^I k$; a (relative) *exposure gap*, $\Gamma_k^I := \frac{\hat{n}^I k - E_k^I}{\hat{n}^I k} > 0$, opens up after an exposed j agent succeeds (because the exposing agent i already succeeded and cannot be re-exposed), or a stub of a succeeding agent connects to an already exposed agent. For $\epsilon > 0$, write $\mathcal{E}^I(\epsilon) := \{\Gamma_k^I < 3\epsilon \text{ for all } k \leq \epsilon I/\hat{n}^I\}$ for the event that the gap process remains bounded in early stages of the contagion

Lemma 9. *For any* $\epsilon > 0$, $\lim_{I\to\infty} \Pr\left(\mathcal{E}^I(\epsilon)\right) = 1$.

We postpone the proof of Lemma 9; the idea is that with $E_k^I \leq \epsilon I$ exposed agents, ϵ small, and $\hat{n}^I$ large, most stubs expose new agents.

For small ϵ, Lemma 9 means that after the approximately $\tau^I I$ initial infections in the experimentation phase, the contagion process resembles a collection of tree networks emanating from these "seeds" at exponential rate $\hat{n}^I$. We now argue that as $\hat{n}^I \to \infty$, this contagion process reaches a negligible fraction of all agents at any $\underline{t} < \sigma = \lim \frac{-\log \tau^I}{\hat{n}^I}$, but approximately all agents at any $\bar{t} > \sigma$.

Specifically, write T_k^I for the k^{th} infection time, and K^I for the (random) number of infected agents at the cutoff time τ^I. Also define inter-arrival times in the contagion phase

$\Delta_k^I := T_{k+1}^I - T_k^I$ for $k > K^I$ and $\Delta_k^I := T_{k+1}^I - \tau^I$ for $k = K^I$. The proof idea is to apply Chernoff bounds to $T_k^I - \tau^I = \sum_{\ell=K^I}^{k-1} \Delta_\ell^I$. Towards this goal, note that conditional on the realization of the "geographical exposure process" $\{E_k^I\}_{k \in [K^I, \epsilon I/\hat{n}^I]}$, inter-arrival times Δ_k^I are independent with arrival rate E_k^I. Conditional on $\mathcal{E}^I(\epsilon)$ we have $E_k^I \in [(1-3\epsilon)\hat{n}^I k, \hat{n}^I k]$, and so

$$\mathbb{E}[e^{-\xi \Delta_k^I} | \mathcal{E}^I(\epsilon)] \leq \frac{\hat{n}^I k}{\hat{n}^I k + \xi} \qquad \text{for all } \xi \geq 0, \tag{40}$$

$$\mathbb{E}[e^{\xi \Delta_k^I} | \mathcal{E}^I(\epsilon)] \leq \frac{(1-3\epsilon)\hat{n}^I k}{(1-3\epsilon)\hat{n}^I k - \xi} \qquad \text{for all } \xi \in [0, (1-3\epsilon)\hat{n}^I k). \tag{41}$$

We now derive upper and lower bounds for the k^{th} success time T_k^I in the contagion phase $k \in [K^I, \epsilon I/\hat{n}^I]$; in the limit $I \to \infty$ these bounds are then shown to imply vanishing chances of getting exposed before σ and after σ, respectively. The upper bound is as follows

$$\begin{aligned}
\Pr(T_k^I \leq \tau^I + \delta | \mathcal{E}^I(\epsilon), K^I) &= \Pr(\sum_{\ell=K^I}^{k-1} \Delta_\ell^I \leq \delta | \mathcal{E}^I(\epsilon)) \leq \inf_{\xi \geq 0} e^{\xi\delta} \prod_{\ell=K^I}^{k-1} \mathbb{E}[e^{-\xi \Delta_\ell^I} | \mathcal{E}^I(\epsilon)] \\
&\leq \inf_{\xi \geq 0} \exp\left(\xi\delta - \sum_{\ell=K^I}^{k-1} \Big(\log(\hat{n}^I \ell + \xi) - \log(\hat{n}^I \ell)\Big)\right) \\
&\leq \inf_{\xi \in [0, \hat{n}^I]} \exp\left(\xi\delta - \sum_{\ell=K^I}^{k-1} \frac{\xi}{\hat{n}^I}\Big(\log(\hat{n}^I(\ell+1)) - \log(\hat{n}^I \ell)\Big)\right) \\
&= \inf_{\xi \in [0, \hat{n}^I]} \exp\left(\xi\left(\delta - \frac{\log k - \log K^I}{\hat{n}^I}\right)\right)
\end{aligned} \tag{42}$$

The first equality drops the τ^I to focus on time since the cutoff, the first inequality is a Chernoff-bound, the second uses (40), the third uses the concavity of the logarithm, and the final equality collapses the telescopic sum.

Next, we argue that for fixed $\epsilon > 0$ and the integer floor $k = \lfloor \epsilon I/\hat{n}^I \rfloor$, the fraction on the RHS of (42) (which approximates the time for the contagion process to reach k agents) converges to $\sigma = \lim \frac{-\log \tau^I}{\hat{n}^I}$:

$$\frac{\log \lfloor \epsilon I/\hat{n}^I \rfloor - \log K^I}{\hat{n}^I} \xrightarrow{D} \sigma \tag{43}$$

For $\bar{B} = \lim I\tau^I < \infty$, this follows because K^I is almost surely bounded above, so as $\hat{n}^I \to \infty$, all terms other than $\frac{\log I}{\hat{n}^I}$ vanish, and $\lim \frac{\log I}{\hat{n}^I} = \lim \frac{\log I - \log B}{\hat{n}^I} = \lim \frac{-\log \tau^I}{\hat{n}^I} = \sigma$. For $B = \infty$, it follows because, by the law of large numbers, $\frac{K^I}{I\tau^I} \xrightarrow{D} 1$; equivalently, $\log K^I - \log I - \log \tau^I \xrightarrow{D} 0$ so the LHS of (43) becomes $\frac{-\log \tau^I}{\hat{n}^I}$, whose limit is σ.

Exposing any positive fraction $\epsilon > 0$ of nodes requires infecting at least $\epsilon I/\hat{n}^I$ agents, and the chance of this at any time $\underline{t} < \sigma$ vanishes

$$\lim_{I\to\infty} \Pr(T^I_{\lfloor \epsilon I/\hat{n}^I \rfloor} \leq \tau^I + \underline{t}) = \lim_{I\to\infty} \Pr(T^I_{\lfloor \epsilon I/\hat{n}^I \rfloor} \leq \tau^I + \underline{t} | \mathcal{E}^I(\epsilon)) \leq \inf_{\xi \geq 0} \exp\left(\xi\left(\underline{t} - \sigma\right)\right) = 0$$

where the equality uses Lemma 9, and the inequality (42) and (43). A fortiori, $\Pr(T^I_{\lfloor \epsilon I/\hat{n}^I \rfloor} \leq \underline{t}) \to 0$.

Finally, for any population share $\epsilon > 0$, the probability that a given agent i has been exposed by time $\underline{t}$ is bounded above by the sum of that share ϵ and the probability that more than share ϵ has been exposed by time $\underline{t}$, $\Pr(S^I \leq \underline{t}) \leq \epsilon + \Pr(T^I_{\lfloor \epsilon I/\hat{n}^I \rfloor} \leq \underline{t})$. Since this inequality holds for any $\epsilon > 0$, we have

$$\lim_{I\to\infty} \Pr(S^I \leq \underline{t}) \leq \lim_{\epsilon\to 0} \lim_{I\to\infty} \left(\epsilon + \Pr(T^I_{\lfloor \epsilon I/\hat{n}^I \rfloor} \leq \underline{t})\right) = 0. \tag{44}$$

Turning to the lower bound for T^I_k, using the same steps as for (42), but with (41) substituting for (40) for the second inequality

$$\begin{aligned}
\Pr(T^I_k \geq \tau^I + \delta | \mathcal{E}^I(\epsilon), K^I) &\leq \inf_{\xi\geq 0} e^{-\xi\delta} \prod_{\ell=K^I}^{k-1} \mathbb{E}[e^{\xi \Delta^I_\ell} | \mathcal{E}^I(\epsilon)] \\
&\leq \inf_{\xi\geq 0} \exp\left(-\xi\delta + \sum_{\ell=K^I}^{k-1} \left(\log((1-3\epsilon)\hat{n}^I \ell) - \log((1-3\epsilon)\hat{n}^I \ell - \xi)\right)\right) \\
&\leq \inf_{\xi\in[0,(1-3\epsilon)\hat{n}^I]} \exp\left(-\xi\delta + \sum_{\ell=K^I}^{k-1} \frac{\xi}{(1-3\epsilon)\hat{n}^I} \left(\log((1-3\epsilon)\hat{n}^I \ell)) - \log((1-3\epsilon)\hat{n}^I(\ell-1))\right)\right) \\
&= \inf_{\xi\in[0,(1-3\epsilon)\hat{n}^I]} \exp\left(-\xi\left(\delta - \frac{\log(k-1) - \log(K^I-1)}{(1-3\epsilon)\hat{n}^I}\right)\right)
\end{aligned}$$

As for the upper bound, for $k = \lfloor \epsilon I/\hat{n}^I \rfloor$ the fraction on the RHS converges, $\frac{\log(\epsilon I/\hat{n}^I - 1) - \log(K^I - 1)}{(1-3\epsilon)\hat{n}^I} \overset{D}{\to} \sigma/(1-3\epsilon)$, so for any $\bar{\delta} > \sigma/(1-3\epsilon)$ in the limit

$$\lim_{I\to\infty} \Pr(T^I_{\lfloor \epsilon I/\hat{n}^I \rfloor} \geq \tau^I + \bar{\delta} | \mathcal{E}^I(\epsilon)) \leq \inf_{\xi\geq 0} \exp\left(-\xi\left(\bar{\delta} - \frac{\sigma}{1-3\epsilon}\right)\right) = 0.$$

Conditional on $\mathcal{E}^I(\epsilon)$, $\lfloor \epsilon I/\hat{n}^I \rfloor$ infections guarantee $\epsilon(1-3\epsilon)I$ exposures by $\tau^I + \bar{\delta}$. For $\epsilon' > 0$ small, approximately $\epsilon'\epsilon(1-3\epsilon)I$ of these get infected by $\tau^I + \bar{\delta} + \epsilon'$, generating approximately $\hat{n}^I \epsilon' \epsilon(1-3\epsilon)I$ new exposure possibilities; that is, an exploding number $\hat{n}^I \epsilon' \epsilon(1-3\epsilon) \to \infty$ for every agent. Now, for any $\bar{t} > \sigma$, we choose $\epsilon, \epsilon' > 0$ small enough, and I large enough that $\tau^I + \bar{\delta} + \epsilon' < \bar{t}$ for $\bar{\delta} := \sigma/(1-3\epsilon) + \epsilon' > \sigma$. As $I \to \infty$, all remaining nodes get exposed before $\tau^I + \bar{\delta} + \epsilon'$ and thus before $\bar{t}$ with probability

$$\lim_{I\to\infty} \Pr(S^I \leq \bar{t}) = 1. \tag{45}$$

Jointly, (44) and (45) for any $\underline{t} < \sigma < \bar{t}$ establish Lemma 4.

The case $\lim \hat{n}^I/I > 0$. So far we assumed $\lim \hat{n}^I/I > 0$ so that $\lim n^I/\hat{n}^I = 1$. Otherwise, we have $\rho = \lim n^I/I = 1 - \exp(-\lim \hat{n}^I/I) > 0$, implying $B_\tau = \rho B > 0$ and so the desired learning time equals $\sigma = 0$ by (39). To see that learning is indeed immediate, note that the first infection exposes fraction $\rho > 0$ of nodes. The paragraph preceding (45) then implies that indeed everyone else gets exposed immediately thereafter. Similarly, if $\lim \tau^I > 0$, a non-vanishing proportion of agents gets infected, and the entire population gets exposed at any $t > 0$, consistent with $\sigma = 0$.

Proof of Lemma 9. We will construct $p(\epsilon) < 1$ such that for large I and any $k \leq \epsilon I/\hat{n}^I$ the chance of a large exposure gap is bounded above via

$$\Pr(\Gamma_k^I > 3\epsilon) < p(\epsilon)^{\hat{n}^I k}. \tag{46}$$

Since $\mathcal{E}^I(\epsilon)$ is the complement of the union of these events over $k \geq 1$, Boole's inequality implies $1 - \Pr(\mathcal{E}^I(\epsilon)) \leq \sum_{k=1}^{\infty} p(\epsilon)^{\hat{n}^I k} = p(\epsilon)^{\hat{n}^I}/(1 - p(\epsilon)^{\hat{n}^I}) \to 0$, which implies (46).

We construct $p(\epsilon)$ and show (46) with the help of Chernoff bounds. The increment $E_k^I - E_{k-1}^I$ counts the newly exposed agents at the k^{th} infection, when agent j. If j was exposed himself, he exposes $\hat{n}^I - 1$ others and is himself deducted from E_k^I; if j was not exposed, he exposes $\hat{n}^I$ others. Each agent exposed by j was already exposed with probability at most $k\hat{n}^I/I$. Thus, writing X_ν for iid binary random variables with $\Pr(X_\nu = 1) = k\hat{n}^I/I$, and $X_\nu = 0$ else, we can upper bound the absolute exposure gap

$$\hat{n}^I k \Gamma_k^I = \hat{n}^I k - E_k^I = \sum_{\ell=1}^{k} \left(\hat{n}^I - (E_\ell^I - E_{\ell-1}^I)\right) \overset{D}{\preceq} 2k + \sum_{\nu=1}^{k\hat{n}^I} X_\nu \tag{47}$$

Now define $p(\epsilon) := \inf_{\xi\geq 0}\left(\frac{\mathbb{E}[e^{X_\nu \xi}]}{e^{2\epsilon\xi}}\right)$. We have $p(\epsilon) < 1$ since $\mathbb{E}[X_\nu] = k\hat{n}^I/I < \epsilon$, and so $\frac{\mathbb{E}[e^{X_\nu \xi}]}{e^{2\epsilon\xi}} \approx \frac{1+\mathbb{E}[X_\nu]\xi}{1+2\epsilon\xi} \leq \frac{1+\epsilon\xi}{1+2\epsilon\xi} < 1$ for small $\xi > 0$. For I large, such that $\epsilon\hat{n}^I > 2$, we then get the following Chernoff-upper bound for the RHS of (47)

$$\Pr\left(2k + \sum_{\nu=1}^{k\hat{n}^I} X_\nu > 3\epsilon\hat{n}^I k\right) \leq \Pr\left(\sum_{\nu=1}^{k\hat{n}^I} X_\nu > 2\epsilon\hat{n}^I k\right) \leq \inf_{\xi\geq 0}\left(\frac{\mathbb{E}[e^{X_\nu\xi}]}{e^{2\epsilon\xi}}\right)^{\hat{n}^I k} = p(\epsilon)^{\hat{n}^I k}$$

which together with (47) implies (46), and hence Lemma 9. □

B.4 Proof of Theorem 2

The challenge with this proof is the complexity of characterizing two outcome variables, asymptotic information and welfare, for a myriad of cases. Specifically we must consider six different network densities $\kappa \lesseqgtr \kappa^*$, $\rho = 0, \in (0,1)$, or $= 1$, and pessimistic priors $p_0 < \bar{p}$ as well as optimistic ones. While some arguments apply to all of these cases, each case also has its idiosyncrasies.

We structure the exposition in order of increasing network density, characterizing asymptotic information and welfare in parallel and emphasizing the case of pessimistic priors $p_0 < \bar{p}$. But to avoid repetitions, we sometimes break this linear narrative by bracketing out arguments that apply more broadly.

As in the paper body, we superscript variables in finite networks with the network size I, e.g. τ_ℓ^I, and drop the superscript in the limit, e.g. $\tau_\ell := \lim_{I \to \infty} \tau_\ell^I$. A priori the limit is well-defined only for some subsequence, but the analysis characterizes all limits under consideration uniquely.

Asymptotic information equals $B = \lim B^I = \lim(K^I \tau_k^I + L^I \tau_\ell^I)$ since the network is connected and each agent's own experimentation $\tau_{k,\ell}^I$ (which in principle is excluded from the social information B) is negligible as $I \to \infty$. It will be useful to decompose B into core agents' pre-cutoff learning $\Upsilon_k^I := I\tau_k^I$ and post cutoff learning $\Upsilon_\ell^I := L^I(\tau_\ell^I - \tau_k^I)$.

We can already note two bounds on $\Upsilon_k, \Upsilon_\ell$: Total information $B = \Upsilon_k + \Upsilon_\ell$ is strictly positive: By contradiction, $B = 0$ means agents face the single-agent problem, choose $\tau_k = \tau_\ell = \bar{\tau} > 0$ and so $B = \infty$. Any agent's pre-cutoff learning B_τ is no larger than $\bar{\tau}$, recalling from (4) that $P^\emptyset(B_\tau)(x+y) - c \geq \psi_\tau = 0$. For core agents, this means $\Upsilon_k \leq \bar{\tau}$. Thus, there is asymptotic learning iff $\Upsilon_\ell = \infty$; a sufficient (but not necessary) condition is $\tau_\ell > 0$.

B.4.1 Case 1: Bounded core size $\kappa < \infty$

Preliminaries. We first establish a necessary and sufficient condition for maximal social learning by peripherals

$$B_{\ell,t} \equiv \kappa t \qquad \text{iff} \qquad \Upsilon_\ell = \infty. \tag{48}$$

If $\Upsilon_\ell = \infty$, core agents immediately observe a peripheral succeed, and then work forever after. If $\Upsilon_\ell < \infty$, the probability of a success $1 - e^{-(\Upsilon_k + \Upsilon_\ell)}$ is less than one, bounding above $b_{\ell,t} \leq \kappa(1 - e^{-(\Upsilon_k + \Upsilon_\ell)}) < \kappa$ for $t > \tau_k$.

By Lemma 1, the social learning upper-bound (48) implies an incentive lower-bound

$$\psi_{\ell,0} \geq \underline{\psi}_{\ell,\kappa,0} := p_0 \left(x + \frac{r}{r+\kappa} y \right) - c \tag{49}$$

with equality iff $\Upsilon_\ell = \infty$.

We distinguish three cases, $\kappa \lesseqgtr \kappa^*$; for optimistic priors $p_0 \geq \bar{p}$, we have $\kappa^* = \infty$, and so only case 1a) $\kappa < \infty$ is relevant.

Case 1a: $\kappa < \kappa^*$**.** Since $\underline{\psi}_{\ell,\kappa,0}$ falls in κ, we have $\underline{\psi}_{\ell,\kappa,0} > \underline{\psi}_{\ell,\kappa^*,0} = 0$, so $\psi_{\ell,0} > 0$, and continuity of $\psi_{\ell,0}$ implies $\tau_\ell > 0$, and asymptotic learning $\Upsilon_\ell = \infty$. By Lemma 2, welfare is bounded away from second-best $\mathcal{V}(\tau_\ell, \kappa\tau_\ell) < \mathcal{V}(0,0) = V^*$. Quantitatively, $\Upsilon_\ell = \infty$ and (48) imply $B_{\ell,t} = \kappa t$, so welfare increases in κ by Lemma 1.

For $p_0 \geq \bar{p}$, only one argument needs adapting: the welfare benchmark now equals $V^* = p_0 y$ which requires immediate and perfect social learning, $B_t = \infty$ for $t > 0$. Clearly, $B_{\ell,t} = \kappa t$ falls short of this benchmark.

Case 1b: $\kappa = \kappa^*$**.** Now $\underline{\psi}_{\ell,\kappa,0} = 0$. We show asymptotic learning, $\Upsilon_\ell = \infty$, by contradiction: By (49) $\Upsilon_\ell < \infty$ would imply $\psi_{\ell,0} > 0$ and so $\tau_\ell > 0$, leading to the contradiction that $\Upsilon_\ell = \infty$. In turn, $\Upsilon_\ell = \infty$ implies by (48) and (49) that $\psi_{\ell,0} = \underline{\psi}_{\ell,\kappa,0} = 0$ and so $\tau_\ell = 0$ and $\kappa\tau_\ell = 0$, attaining the welfare upper bound $\mathcal{V}(0,0) = V^*$.

Case 1c: $\kappa \in (\kappa^*, \infty)$**.** Now $\underline{\psi}_{\ell,\kappa,0} < 0$. Asymptotic learning fails because $\Upsilon_\ell = \infty$ would imply by (48) and (49) that $\psi_{\ell,0} = \underline{\psi}_{\ell,\kappa,0} < 0$ and so $\tau_\ell^I = 0$ and $\Upsilon_\ell = 0$. In turn, $\Upsilon_\ell < \infty$ implies $\tau_\ell = 0$ and $\psi_{\ell,0} = 0$. To quantify information, we first claim that $\Upsilon_k = \lim I\tau_k^I = 0$: Indeed, core agents receive all social information immediately, $B_{k,t} = \Upsilon_k + \Upsilon_\ell$ for all $t > 0$, while peripherals' learning is bounded by $B_{\ell,t} \leq \kappa t$. This bounds incentives of core agents above $\psi_{k,0} < \psi_{\ell,0} = 0$, and so $\tau_k^I = 0$ for large I.[35]

Social information thus equals Υ_ℓ. We now show this falls in κ: Peripherals observe a success by time t iff at least one peripheral succeeds during experimentation, and then a core agent succeeds during $(0,t]$; thus $1 - e^{-B_{\ell,t}} = (1 - e^{-\Upsilon_\ell})(1 - e^{-\kappa t})$.[36] Since the RHS rises with both κ and Υ_ℓ and experimentation incentives $\psi_{\ell,0}$ fall in $\{B_{\ell,t}\}$, the equilibrium condition $\psi_{\ell,0} = 0$ implies that a rise in information transmission κ must be compensated by a fall in aggregate information Υ_ℓ. For future reference, we note that as $\kappa \to \infty$, the learning curve $B_{\ell,t}$ converges to Υ_ℓ for each $t > 0$, and so peripherals' indifference condition converges to

[35] We also get $\tau_k^I = 0$ for large I and $\Upsilon_k = 0$ in cases 1a,b with $p_0 < \bar{p}$, where $\psi_{k,0} < 0$ is ensured by $B_{k,t} = \infty$ for all $t > 0$.

[36] Solving for $B_{\ell,t}$ and differentiating yields $b_{\ell,t} = \kappa \frac{e^{-\kappa t}(1-e^{-\Upsilon_\ell})}{e^{-\kappa t}(1-e^{-\Upsilon_\ell}) + e^{-\Upsilon_\ell}}$, generalizing (48).

$p_0(x + e^{-\Upsilon_\ell} y) = c$, pinning down aggregate information Υ_ℓ.

Finally, since $\tau_\ell = \kappa\tau_\ell = 0$, welfare attains the upper bound $\mathcal{V}(0,0) = V^*$.

B.4.2 Case 2: Exploding core $\kappa = \infty$

Preliminaries. For simplicity we first cover the case $\rho < 1$, and separate the analysis for $\rho = 1$. We prepare the ground with two preliminary lemmas.

Lemma 10. *Assume $\kappa = \infty$, $\rho < 1$, and any prior $p_0 > \underline{p}$.*

(*a*) *Individual learning vanishes: $\tau_k^I, \tau_\ell^I \to 0$.*

(*b*) *Social learning is immediate: For all $t > 0$, $B_{k,t}^I, B_{\ell,t}^I \to \Upsilon_k + \Upsilon_\ell$.*

Proof. Part (a) follows by the upper bound on pre-cutoff learning $B_\tau \le \bar{\tau}$. For core agents, $B_{k,t}^I = (I-1)\tau_k^I \le \bar{\tau}$. For peripherals,

$$B_{\ell,\tau_\ell^I}^I := K^I \tau_k^I + \int_{\tau_k^I}^{\tau_\ell^I} K^I a_t^I dt \tag{50}$$

where core agents' expected effort a_t^I from (13) drifts towards $\min\{(L^I - 1)/K^I, 1\}$ and is hence bounded away from 0 by our assumption that $\rho < 1$. The upper bound, $B_{\ell,\tau_\ell}^I < \bar{\tau}$ thus requires the domain to vanish, $\tau_\ell^I \to 0$, as the integrand explodes, $K^I \to \infty$.

Turning to part (b), the conditional probability that some agent i has observed a neighbor succeed is bounded via

$$\left(1 - \exp(-(I\tau_k^I + (L^I - 1)\tau_\ell^I))\right)\left(1 - \exp(-K^I t/2)\right) < 1 - \exp(-B_t^I) < 1 - \exp(-(\Upsilon_k^I + \Upsilon_\ell^I))$$

The upper bound is the probability that any agent ever succeeds. The lower bound is the probability that some agent $j \neq i$ succeeds before $t/2$, times the probability that a core agent succeeds in $(t/2, t)$. Both bounds converge to $1 - \exp(-(\Upsilon_k + \Upsilon_\ell))$ as $I \to \infty$. □

Lemma 10(b) implies that welfare of both core agents and peripherals equals

$$V_k = V_\ell = \left(1 - \exp\left(-(\Upsilon_k + \Upsilon_\ell)\right)\right) p_0 y$$

and thus rises in social information $\Upsilon_k + \Upsilon_\ell$. All of our results for social information (monotonicity and attainment of benchmarks) thus apply equally to welfare.

Lemma 10b implies that social learning of both core agents and peripherals occurs in two bursts: one before the cutoff and one immediately after, and both approaching $t = 0$. For such learning with burst sizes B^- and B^+, the indifference condition (4) becomes

$$\Psi(B^-, B^+) := P^\emptyset(B^-)(x + e^{-B^+} y) - c = 0. \tag{51}$$

Recalling the effects of social learning on experimentation incentives (5) and $ry = x - c$, the solution of (51) has slope

$$-\frac{dB^+}{dB^-} = \frac{\partial_{B^-}\Psi}{\partial_{B^+}\Psi} = \frac{x + e^{-B^+}y - c}{e^{-B^+}y} = re^{B^+} + 1. \tag{52}$$

To apply (51) to core agents and peripherals, write asymptotic pre-cutoff learning and experimentation incentives as $B_{\ell,\tau_\ell} = \lim B^I_{\ell,\tau^I_\ell}$ and $\psi_{\ell,\tau_\ell} = \lim \psi^I_{\ell,\tau^I_\ell}$ and similarly for core agents, substituting "k" for "ℓ"; note that even though $\tau^I_\ell \to \tau_\ell = 0$, this is distinct from, and generally greater than the other limit $B_{\ell,0} = \lim B^I_{\ell,0} = 0$, cf (50). For core agents, $B^- = B_{k,\tau_k} = \Upsilon_k$, $B^+ = \Upsilon_\ell$, and (51) coincides with the limit of (12) as $L \to \infty$. For peripherals, we get an explicit expression of B_{ℓ,τ_ℓ} in $\Upsilon_k, \Upsilon_\ell$ only for $\rho > 0$, (57).

Lemma 11. *Assume $\kappa = \infty$, $\rho < 1$, and any prior $p_0 > \underline{p}$.*

(*a*) *Core agents are initially indifferent*

$$\Psi(\Upsilon_k, \Upsilon_\ell) = P^\emptyset(\Upsilon_k)\left(x + e^{-\Upsilon_\ell}y\right) - c = 0. \tag{53}$$

(*b*) *Pre-cutoff learning of core agents and peripherals coincides: $B_{\ell,\tau_\ell} = \Upsilon_k$.*

Proof. Part (a): For internal cutoffs $\tau^I_k > 0$, the indifference conditions $\psi^I_{k,\tau^I_k} = 0$ converge to (53). By contradiction, assume that $\tau^I_k = 0$ for large I, so that $\Upsilon_k = 0$ and $\psi_{k,\tau_k} = \psi_{k,0} = p_0\left(x + e^{-\Upsilon_\ell}y\right) - c < 0$. Using Lemma 10(b) (immediate learning by both core agents and peripherals) and the greater importance of pre-cutoff learning (52), strict shirking incentives by core agents carry over to peripherals[37]

$$\psi_{\ell,\tau_\ell} = \Psi(B_{\ell,\tau_\ell}, \Upsilon_\ell - B_{\ell,\tau_\ell}) \leq \Psi(0, \Upsilon_\ell) = \psi_{k,\tau_k} < 0.$$

Thus $\tau^I_\ell = 0$ for large I, leading to the contradiction that $\Upsilon_k + \Upsilon_\ell = 0$ and $\psi_{k,0} = \psi_{\ell,0} = p_0\,(x + y) - c = \Psi(0,0) > 0$.

Part (b): The indifference condition of core and peripheral agents imply

$$\Psi(\Upsilon_k, \Upsilon_\ell) = \psi_{k,\tau_k} = 0 = \psi_{\ell,\tau_\ell} = \Psi(B_{\ell,\tau_\ell}, \Upsilon_k + \Upsilon_\ell - B_{\ell,\tau_\ell}) = \Psi(\Upsilon_k - (\Upsilon_k - B_{\ell,\tau_\ell}), \Upsilon_\ell + (\Upsilon_k - B_{\ell,\tau_\ell})).$$

The greater effect of pre-cutoff learning on incentives (52) thus implies $\Upsilon_k - B_{\ell,\tau_\ell} = 0$. □

Lemma 11 establishes two equations for $\Upsilon_k, \Upsilon_\ell$. Below we show they admit a unique solution; a corner solution for $\rho = 0$, and an internal one for $\rho \in (0,1)$.

[37]Note the contrast to the case with bounded core size $\kappa < \infty$ (and $p_0 < \bar{p}$), where peripherals learn slower than core agents, so that $\psi_{k,\tau_k} < \psi_{\ell,\tau_\ell} = 0$.

Case 2a: $\rho = 0$. In this case we get a corner solution for $\Upsilon_k, \Upsilon_\ell$ with $\Upsilon_k/\Upsilon_\ell = 0$. Indeed, using Lemma 11(b), pre-cutoff learning is a vanishing proportion of post-cutoff learning

$$\Upsilon_k = B_{\ell,\tau_\ell} = \lim B^I_{\ell,\tau^I_\ell} < \lim K^I \tau^I_\ell = \lim \frac{K^I}{L^I} L^I \tau^I_\ell \leq \text{“}\frac{\rho}{1-\rho}(\Upsilon_k + \Upsilon_\ell)\text{”}. \tag{54}$$

Since $\rho = 0$, we must have either $\Upsilon_k = 0$ or $\Upsilon_\ell = \infty$ (then the last term "$0 \cdot \infty$" is not well defined), or both.

For pessimistic priors $p_0 < \bar{p}$, core agents' indifference (53) rules out asymptotic learning, so $\Upsilon_\ell < \infty$ and (54) implies $\Upsilon_k = 0$. In turn, aggregate information Υ_ℓ solves $\Psi(0, \Upsilon_\ell) = p_0 \left(x + e^{-\Upsilon_\ell} y\right) - c = 0$. This is the same indifference condition we found in case 1c as $\kappa \to \infty$, so aggregate information is continuous in this limit.

For $p_0 \geq \bar{p}$, Υ_k solves $P^\emptyset(\Upsilon_k) = \bar{p}$ and $\Upsilon_\ell = \infty$.[38] Core agents' indifference (53) clearly requires experimentation until the myopic threshold, $P^\emptyset(\Upsilon_k) \leq \bar{p}$. If, by contradiction, core agents experiment past the myopic threshold, $P^\emptyset(\Upsilon_k) < \bar{p}$, then (53) implies $\Upsilon_\ell < \infty$, and (54) leads to the contradiction that $\Upsilon_k = 0$.

Case 2b: $\rho \in (0, 1)$. In this case we get an internal solution for $\Upsilon_k, \Upsilon_\ell$. We first further operationalize Lemma 11(b) by replacing the upper bound in (54) with an explicit expression of peripherals' pre-cutoff learning B_{ℓ,τ_ℓ} in terms of $\Upsilon_k, \Upsilon_\ell$, (57). To analyze (50) as the integrand $K^I a^I_t$ explodes and the integration domain $[\tau^I_k, \tau^I_\ell]$ vanishes, we rescale time $\alpha^I_t := a^I_{t/I}$. The ODE (13) for core agents' experimentation intensity thus becomes

$$I \frac{\dot{\alpha}^I_t}{1 - \alpha^I_t} = \begin{cases} L^I - 1 & t < I\tau^I_k \\ L^I - 1 - K^I \alpha^I_t & t \in (I\tau^I_k, I\tau^I_\ell) \\ -K^I \alpha^I_t & t > I\tau^I_\ell \end{cases} \tag{55}$$

Recalling the definition of $\rho, \Upsilon_k, \Upsilon_\ell$, as $I \to \infty$, this converges to the solution α_t of

$$\frac{\dot{\alpha}}{1 - \alpha} = \begin{cases} 1 - \rho & t < \Upsilon_k \\ 1 - \rho - \rho\alpha & t \in (\Upsilon_k, \Upsilon_k + \Upsilon_\ell/(1-\rho)) \\ -\rho\alpha & t > \Upsilon_k + \Upsilon_\ell/(1-\rho) \end{cases} \tag{56}$$

Peripherals' pre-cutoff learning (50) then converges to

$$B_{\ell,\tau_\ell} = \rho \left(\Upsilon_k + \int_{\Upsilon_k}^{\Upsilon_k + \Upsilon_\ell/(1-\rho)} \alpha_t dt \right), \tag{57}$$

[38] In the borderline case with $p_0 = \bar{p}$, we get both $\Upsilon_k = 0$ and $\Upsilon_\ell = \infty$.

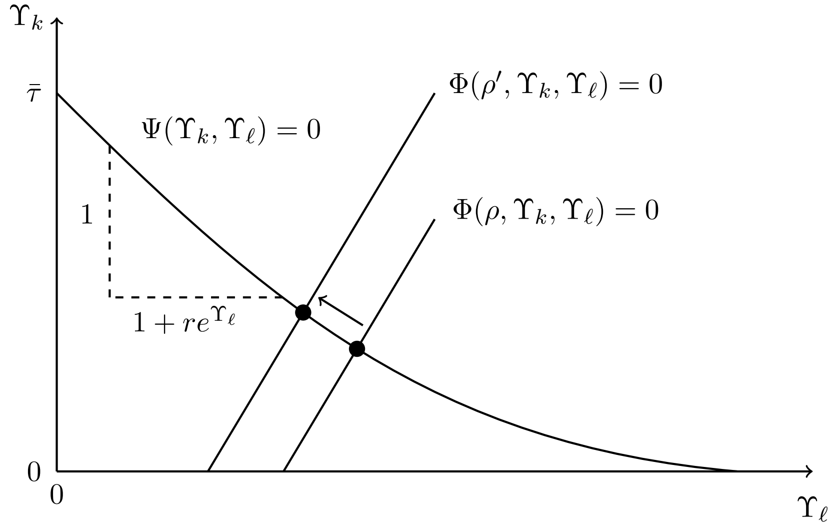

Figure 8: **Solutions of $\Phi(\rho, \Upsilon_k, \Upsilon_\ell) = 0$ and $\Psi(\Upsilon_k, \Upsilon_\ell) = 0$.**

so we can rewrite Lemma 11(b) as

$$\Phi(\rho, \Upsilon_k, \Upsilon_\ell) := \rho \left(\Upsilon_k + \int_{\Upsilon_k}^{\Upsilon_k + \Upsilon_\ell/(1-\rho)} \alpha_t dt \right) - \Upsilon_k = 0. \tag{58}$$

We can now characterize equilibrium learning.

Lemma 12. *Equations (53), (58) admit a unique solution $(\Upsilon_k, \Upsilon_\ell)$. This solution satisfies $0 < \Upsilon_k, \Upsilon_\ell < \infty$, and aggregate information $\Upsilon_k + \Upsilon_\ell$ falls in ρ.*

The proof of Lemma 12 relies on the following generalization of Leibniz's integral rule: For some "cutoff" $s > 0$ and Lipschitz-continuous functions f, g with $g(x_s(s)) \neq 0$, let $x_t = x_t(s)$ be the continuous solution of an ODE with fixed x_0, and

$$\dot{x} = \begin{cases} f(x) & \text{for } t < s \\ g(x) & \text{for } t > s. \end{cases}$$

Lemma 13. *For any $\Delta > 0$*

$$\frac{\partial}{\partial s} \int_s^{s+\Delta} x_t(s) dt = \frac{f(x_s(s))}{g(x_s(s))} \left(x_s(s + \Delta) - x_s(s) \right) \tag{59}$$

Proof of Lemma 12. Equation (58) together with $\Upsilon_k + \Upsilon_\ell > 0$ and the fact that the solution α of (56) is bounded away from zero imply $\Upsilon_k > 0$, and in turn that $0 < \Upsilon_\ell < \infty$. Thus, asymptotic learning fails.

To solve (53), (58), we note that Φ clearly rises in ρ and Υ_ℓ. We show below in (58) that it falls in Υ_k. Hence zero-sets of Φ in $(\Upsilon_\ell, \Upsilon_k)$-space are increasing and shift left when ρ rises to ρ', as illustrated in Figure 8. Recalling from (52) that zero-sets of Ψ are decreasing with slope $-1/(1+re^{-\Upsilon_\ell}) > -1$, equations (53), (58) admit a unique solution $(\Upsilon_k, \Upsilon_\ell)$. A rise in ρ shifts this solution left on the zero-set of Ψ, so $\Upsilon_k + \Upsilon_\ell$ falls.

In fact, the monotonicity of $\Upsilon_k + \Upsilon_\ell$ extends to the boundary points $\rho = 0, 1$: We recall that for $\rho = 0$ all learning is post-cutoff, $\Upsilon_k = 0, \Psi(0, \Upsilon_\ell) = 0$,[39] and anticipate that for $\rho = 1$ all learning is pre-cutoff, $\Upsilon_\ell = 0, \Psi(\Upsilon_k, 0) = 0$, thus attaining the extreme points on the zero set of $\Psi(\Upsilon_k, \Upsilon_\ell) = 0$ as illustrated in Figure 8.

To show that Φ falls in Υ_k, we write $\alpha_* = \alpha_{\Upsilon_k}$ and $\alpha^* = \alpha_{\Upsilon_k + \Upsilon_\ell/(1-\rho)}$, assume that $1 - \rho - \rho\alpha_* \neq 0$, and then argue[40]

$$\frac{\partial \Phi}{\partial \Upsilon_k} = -(1-\rho) + \rho \frac{1-\rho}{1-\rho-\rho\alpha_*}(\alpha^* - \alpha_*) = -(1-\rho)\frac{1-\rho-\rho\alpha^*}{1-\rho-\rho\alpha_*} < 0.$$

The first equality follows from Lemma 13 by substituting $s = \Upsilon_k$ and $\Delta = \Upsilon_\ell/(1-\rho)$ for the integral boundaries, $x_t = \alpha_t$ for the trajectory, and $f(\alpha) = (1-\rho)(1-\alpha)$ for the law-of-motion before $s = \Upsilon_k$ and $g(\alpha) = (1-\rho-\rho\alpha)(1-\alpha)$ after Υ_k.

The middle equality is an elementary algebraic transformation, and the final inequality owes to the fact that $\dot{\alpha}/(1-\alpha) = 1-\rho-\rho\alpha$ cannot switch signs on $[\Upsilon_k, \Upsilon_k + \Upsilon_\ell/(1-\rho)]$, cf (56), so that $\frac{1-\rho-\rho\alpha^*}{1-\rho-\rho\alpha_*} > 0$. □

Proof of Lemma 13. The Leibniz rule evaluates the LHS of (59) "vertically", computing $\frac{\partial}{\partial s} x_t(s) = \lim_{\delta \to 0} \frac{1}{\delta}(x_t(s+\delta) - x_t(s))$ for fixed $t \in [s, s+\Delta]$. Since the ODE $\dot{x} = g(x)$ is autonomous, it is more economical to compare the trajectories $\{x_t(s+\delta)\}_t$ and $\{x_t(s)\}_t$ "horizontally", as illustrated in Figure 9.

Formally, assume first that $f(s)$ and $g(s)$ have the same sign, and for $\delta > 0$ small, let $\delta' > 0$ solve $x_{s+\delta'}(s) = x_{s+\delta}(s+\delta)$. At $s + \delta'$ the original trajectory "merges" with the shifted trajectory and since $\dot{x} = g(x)$ is autonomous we get $x_{s+\delta'+\hat{\delta}}(s) = x_{s+\delta+\hat{\delta}}(s+\delta)$, as illustrated in Figure 9(left). Thus

$$\int_{\delta}^{\delta+\Delta} x_{s+\tilde{\delta}}(s+\delta) d\tilde{\delta} = \int_0^{\Delta} x_{s+\delta+\hat{\delta}}(s+\delta) d\hat{\delta} = \int_0^{\Delta} x_{s+\delta'+\hat{\delta}}(s) d\hat{\delta} = \int_{\delta'}^{\delta'+\Delta} x_{s+\tilde{\delta}}(s) d\tilde{\delta} \quad (60)$$

using the change of variable $\tilde{\delta} = \delta + \hat{\delta}$ in the first equality, and $\tilde{\delta} = \delta' + \hat{\delta}$ in the last. Thus

[39] This assumes $p_0 < \bar{p}$. For $p_0 \geq \bar{p}$, asymptotic information is infinite for $\rho = 0$, and hence trivially greater than the finite learning for $\rho > 0$.

[40] Since $\alpha_t = 1 - \exp(-(1-\rho)t)$ for $t < \Upsilon_k$, there exists at most one value of Υ_k with $1 - \rho - \rho\alpha_{\Upsilon_k} = 0$. Since Φ is continuous in Υ_k and decreasing in Υ_k everywhere else, it decreases everywhere.

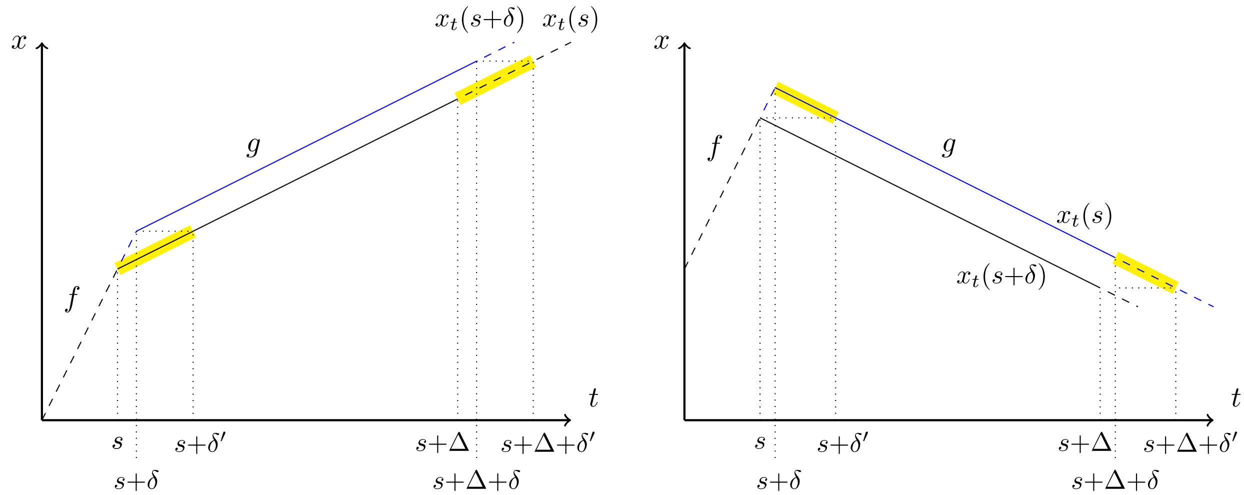

Figure 9: **Proof of Leibniz Rule.** In both figures the difference of between the integral of the upper solid line, $x_t(s+\delta)$ over $t \in [s+\delta, s+\delta+\Delta]$, and the lower solid line, $x_t(s)$ over $t \in [s, s+\Delta]$, equals the difference in the integrals of the shaded lines. E.g. In the left picture this is difference between $x_t(s)$ over $t \in [s+\Delta, s+\delta'+\Delta]$ and $x_t(s)$ over $t \in [s, s+\delta']$, which is the RHS of (61) after substituting $t = s+\tilde{\delta}$.

$$\int_{s+\delta}^{s+\Delta+\delta} x_t(s+\delta)dt - \int_{s}^{s+\Delta} x_t(s)dt = \int_{\delta}^{\delta+\Delta} x_{s+\tilde{\delta}}(s+\delta)d\tilde{\delta} - \int_{0}^{\Delta} x_{s+\tilde{\delta}}(s)d\tilde{\delta}$$
$$= \int_{\delta'}^{\delta'+\Delta} x_{s+\tilde{\delta}}(s)d\tilde{\delta} - \int_{0}^{\Delta} x_{s+\tilde{\delta}}(s)d\tilde{\delta} = \int_{\Delta}^{\Delta+\delta'} x_{s+\tilde{\delta}}(s)d\tilde{\delta} - \int_{0}^{\delta'} x_{s+\tilde{\delta}}(s)d\tilde{\delta} \quad (61)$$

where the first equality uses the change of variables $t = s+\tilde{\delta}$, the second uses (60), and the third cancels identical terms $\int_{\delta'}^{\Delta} x_{s+\tilde{\delta}}(s)d\tilde{\delta}$. In the limit

$$\frac{\partial}{\partial s}\int_{s}^{s+\Delta} x_t(s)dt = \lim_{\delta\to 0}\frac{\delta'}{\delta}\left(x_{s+\Delta}(s) - x_s(s)\right) = \frac{f(x_s(s))}{g(x_s(s))}\left(x_{s+\Delta}(s) - x_s(s)\right),$$

where we used that at first-order $\delta' g(x_s(s)) = \delta f(x_s(s))$.

If f and g have different signs, we let $\delta' > \delta$ solve $x_{s+\delta'}(s+\delta) = x_s(s)$, so $\delta f(s) + (\delta' - \delta)g(s) = 0$, as illustrated in Figure 9(right). Analogous arguments as above then show

$$\frac{\partial}{\partial s}\int_{s}^{s+\Delta} x_t(s)dt = \lim_{\delta\to 0}\frac{\delta'-\delta}{\delta}\left(x_s(s) - x_{s+\Delta}(s)\right) = \frac{f(x_s(s))}{g(x_s(s))}\left(x_{s+\Delta}(s) - x_s(s)\right).$$

as required □

Case 2c: $\rho = 1$. While Lemmas 10 and 11 and most other substantive intermediate results remain true for $\rho = 1$, their proofs divide by $1-\rho$, and sometimes invoke that $L \to \infty$.

Instead of re-proving everything, we provide a separate analysis, solely based on the function Ψ and its derivatives, (51-52), and the ODE (55). Specifically we will show that

$$\Psi(\Upsilon_k, \Upsilon_\ell) \leq \psi_{k,\tau_k} \leq 0 = \psi_{\ell,\tau_\ell} = \Psi(\Upsilon_k + \Upsilon_\ell, 0) \tag{62}$$

Together with (52), this implies $\Upsilon_\ell = 0$, so the inequalities in (62) must hold with equality. In particular $0 = \Psi(\Upsilon_k, 0) = P^{\emptyset}(\Upsilon_k)(x + y) - c$, so total information is as in the clique (or the single-agent problem) $\Upsilon_k + \Upsilon_\ell = \Upsilon_k = \bar{\tau}$.

We now show (62). The middle inequality and equality reflect (the limits of) peripherals' indifference and core agents weak shirking incentives at their respective cutoffs. The first inequality takes the limit of the strict inequality $\Psi(\Upsilon^I_k, \Upsilon^I_\ell) < \psi^I_{k,\tau^I_k}$, which reflects that core agents' observe post-cutoff information Υ^I_ℓ with a delay.

Only the last equality in (62), which states that peripherals' learning is entirely pre-cutoff, requires a novel argument and the assumption $\rho = 1$. Intuitively, information transmission by K^I core agents is infinitely faster than generation by L^I peripherals.

Formally, we will show that peripherals' aggregate post-cutoff learning vanishes

$$\frac{K^I}{I} \int_{I\tau^I_\ell}^{\infty} \alpha^I_t dt \to 0. \tag{63}$$

By (63), peripherals pre-cutoff learning B^I_{ℓ,τ^I_ℓ} converges to total information $\Upsilon_k + \Upsilon_\ell$, implying the last equality in (62).

To see (63) we first argue that $\alpha^I_t \to 0$ for all t. By line one of (55), $\alpha^I_t \leq L^I t/I \leq L^I \bar{\tau}/I \to 0$ for all $t < I\tau^I_k < \bar{\tau}$; at $t > I\tau^I_k$, lines two and three of (55) imply $\dot{\alpha}^I_t < 0$ when $\alpha^I_t \geq L^I/K^I \to (1-\rho)/\rho = 0$. All told, $\alpha^I_t \to 0$ for all t. Turning to the aggregate in (63), line three of (55) states that α^I_t decays exponentially at rate $(1 - \alpha^I_t)K^I/I$. Since this rate converges to 1, we have $\int_{I\tau^I_\ell}^{\infty} \alpha^I_t dt - \alpha^I_{I\tau^I_\ell} \to 0$. Together with $\alpha^I_{I\tau^I_\ell} \to 0$, this implies (63).

C Appendix: Proofs from Section 5

C.1 Proof of Theorem 3

We first show the comparisons between directed and undirected trees

$$\bar{\tau}^{(n+1)} < \vec{\tau}^{(n)} < \bar{\tau}^{(n)} \qquad \text{and} \qquad \bar{V}^{(n+1)} > \vec{V}^{(n)} > \bar{V}^{(n)} \tag{64}$$

and then comment how the same arguments imply the comparison between undirected trees and triangle trees.

Emphasizing the role of degree n and cutoff τ, we write the neighbor's expected time-t

effort in directed and undirected tree as $\vec{a}_t^{(n)}(\tau), \bar{a}_t^{(n)}(\tau)$. These equal 1 for $t < \tau$, and solve (16) and (19) for $t > \tau$.

We first show $\vec{\tau}^{(n)} < \bar{\tau}^{(n)}$. For a given cutoff $\tau > 0$, we have $\vec{a}_t^{(n)}(\tau) > \bar{a}_t^{(n)}(\tau)$ for all $t \geq \tau$: At the cutoff $\vec{a}_\tau^{(n)}(\tau) = 1 - e^{-n\tau} > 1 - e^{-(n-1)\tau} = \bar{a}_\tau^{n}(\tau)$, and this ranking prevails for $t > \tau$ since the RHS of (16) exceeds the RHS of (19). Additionally, $\vec{a}_t^{(n)}(\tau), \bar{a}_t^{(n)}(\tau)$ rise in τ, strictly for $\tau < t$. By Lemma 6(b), for any $\tau \leq \vec{\tau}^{(n)}$

$$0 = \psi_{\vec{\tau}^{(n)}}(\{n\vec{a}_t^{(n)}(\vec{\tau}^{(n)})\}) < \psi_\tau(\{n\bar{a}_t^{(n)}(\tau)\}),$$

so in equilibrium we must instead have $\bar{\tau}^{(n)} > \vec{\tau}^{(n)}$, as desired. Agents in the directed network then also have lower pre-cutoff social learning $n\vec{\tau}^{(n)} < n\bar{\tau}^{(n)}$ and hence higher welfare $\vec{V}^{(n)} = \mathcal{V}(\vec{\tau}^{(n)}, n\vec{\tau}^{(n)}) > \mathcal{V}(\bar{\tau}^{(n)}, n\bar{\tau}^{(n)}) = \bar{V}^{(n)}$ since $\mathcal{V}$ falls in both of its arguments by Lemma 2.

We next show $\bar{\tau}^{(n+1)} < \vec{\tau}^{(n)}$. For a given cutoff τ, the degree difference exactly offsets the difference in the laws-of-motion (16) and (19), so at the level of i's random neighbor j, social learning coincides $\bar{a}_t^{(n+1)}(\tau) = \vec{a}_t^{(n)}(\tau)$. But then total social learning is higher in the undirected network $(n+1)\bar{a}_t^{(n+1)}(\tau) > n\vec{a}_t^{(n)}(\tau)$. By Lemma 6(a,b), for any $\tau \leq \bar{\tau}^{(n+1)}$ we have $0 = \psi_{\bar{\tau}^{(n+1)}}(\{(n+1)\bar{a}_t^{(n+1)}(\bar{\tau}^{(n+1)})\}) < \psi_\tau(\{n\vec{a}_t^{(n)}(\tau)\})$, so in equilibrium we must instead have $\vec{\tau}^{(n)} > \bar{\tau}^{(n+1)}$, as desired.

For the associated welfare ranking we will show more strongly that

$$\bar{\tau}^{(n+1)} < \tau' := \frac{n+1}{n+2}\vec{\tau}^{(n)} \tag{65}$$

for then, by Lemma 2,

$$\bar{V}^{(n+1)} = \mathcal{V}(\bar{\tau}^{(n+1)}, (n+2)\bar{\tau}^{(n+1)} - \bar{\tau}^{(n+1)}) > \mathcal{V}(\bar{\tau}^{(n+1)}, (n+1)\vec{\tau}^{(n)} - \bar{\tau}^{(n+1)}) > \mathcal{V}(\vec{\tau}^{(n)}, n\vec{\tau}^{(n)}) = \vec{V}^{(n)},$$

where the first inequality uses (65), and the second that adding $\vec{\tau}^{(n)} - \bar{\tau}^{(n+1)} > 0$ to the first argument of $\mathcal{V}$ and subtracting it from the second argument decreases $\mathcal{V}$, since $\partial_\tau \mathcal{V} < \partial_B \mathcal{V} < 0$, as shown in (31).

To see (65), we compare a neighbor's expected experimentation $\delta \geq 0$ after the cutoff for the directed n-tree with equilibrium cutoff $\vec{\tau}^{(n)}$, $\vec{a}_\delta := \vec{a}_{\vec{\tau}^{(n)}+\delta}^{(n)}(\vec{\tau}^{(n)})$, and the undirected $(n+1)$-tree with non-equilibrium cutoff τ' from (65), $\bar{a}'_\delta := \bar{a}_{\tau'+\delta}^{(n+1)}(\tau')$. We will show that

$$n\vec{a}_\delta < (n+1)\bar{a}'_\delta. \tag{66}$$

Since pre-cutoff learning coincides, $\vec{p}^{(n)}_{\vec{\tau}^{(n)}} = P^{\emptyset}((n+1)\vec{\tau}^{(n)}) = P^{\emptyset}((n+2)\tau') = \bar{p}^{(n+1)}_{\tau'}$, we get

$$\begin{aligned} 0 &= \psi_{\vec{\tau}^{(n)}}(\{n\vec{a}^{(n)}_t(\vec{\tau}^{(n)})\}) = \vec{p}^{(n)}_{\vec{\tau}^{(n)}}(x + rye^{-\int_0^\infty (r+n\vec{a}_\delta)d\delta}) - c \\ &> \bar{p}^{(n+1)}_{\tau'}(x + rye^{-\int_0^\infty (r+(n+1)\bar{a}'_\delta)d\delta}) - c = \psi_{\tau'}(\{(n+1)\bar{a}^{(n+1)}_t(\tau')\}) \end{aligned}$$

and hence in equilibrium $\bar{\tau}^{(n+1)} < \tau'$, which is (65).

We now argue (66). We start by claiming a partial converse, $\bar{a}'_\delta < \vec{a}_\delta$. This follows by $\bar{a}'_0 = 1 - \exp(-n\tau') < 1 - \exp(-n\vec{\tau}^{(n)}) = \vec{a}_0$ and the fact that $\bar{a}'_\delta$ and $\vec{a}_\delta$ follow the same law-of-motion $\dot{a} = (n-1)a(1-a)$ and hence can't change position. Now (66) follows for $\delta = 0$ as follows

$$n\vec{a}_0 = n(1-\exp(-n\vec{\tau}^{(n)})) < (n+1)(1-\exp(-n\frac{n}{n+1}\vec{\tau}^{(n)})) < (n+1)(1-\exp(-n\tau')) = (n+1)\bar{a}_0 \tag{67}$$

where the first inequality uses that $(1-\exp(-x))/x$ falls in x, and the second inequality that $\frac{n}{n+1}\vec{\tau}^{(n)} < \frac{n+1}{n+2}\vec{\tau}^{(n)} = \tau'$. We show (66) for all $\delta > 0$ by arguing that $n\vec{a}_\delta$ cannot cross $(n+1)\bar{a}_\delta$ from below: Assume $n\vec{a}_\delta = (n+1)\bar{a}_\delta$. Then, using $\vec{a}_\delta > \bar{a}'_\delta$

$$n\dot{\vec{a}}_\delta = n(n-1)\vec{a}_\delta(1-\vec{a}_\delta) < (n+1)(n-1)\bar{a}'_\delta(1-\bar{a}'_\delta) = (n+1)\dot{\bar{a}}'_\delta.$$

We have thus proven the comparison between directed and undirected trees (64).

The comparison between undirected and triangle trees

$$\hat{\tau}^{(n+1)} < \bar{\tau}^{(n)} < \hat{\tau}^{(n)} \qquad \text{and} \qquad \hat{V}^{(n+1)} > \bar{V}^{(n)} > \hat{V}^{(n)}$$

follows analogously: For equal degree n, we use the fact that $\bar{a}^{(n)}_t(\tau) > \hat{a}^{(n)}_t(\tau)$ to argue $\bar{\tau}^{(n)} < \hat{\tau}^{(n)}$ and consequently $\bar{V}^{(n)} > \hat{V}^{(n)}$. For triangle networks with degree $n+1$, we use the fact that $\hat{a}^{(n+1)}_t(\tau) = \bar{a}^{(n)}_t(\tau)$ to argue $\hat{\tau}^{(n+1)} < \bar{\tau}^{(n)}$; the argument that $\hat{V}^{(n+1)} > \bar{V}^{(n)}$ relies again on showing that incentives in the triangle network are strictly positive at $\tau' := \frac{n+1}{n+2}\bar{\tau}^{(n)}$. The only difference is that now $\hat{a}^{(n+1)}_t(\tau) = \bar{a}^{(n)}_t(\tau)$ equals $1 - e^{-(n-1)t}$ for $t = \tau$ and then evolves according to $\dot{a} = (n-2)a(1-a)$, while in the comparison of the undirected $n+1$-tree to the directed n-tree the corresponding values were to $1 - e^{-nt}$ for $t = \tau$ and then $\dot{a} = (n-1)a(1-a)$, so in the analogue of (67) we must now use the fact that $\frac{n-1}{n} < \frac{n+1}{n+2}$.

C.2 Convergence: Proof of Proposition 4

We lead with the proof of part (b) for undirected networks (which are also used in Section 4.2), dropping the "upper bar", say on $\bar{a}_t$, to ease notation. Subsequently, we discuss how to adapt the proof for directed and triangular networks in parts (a) and (c) of Proposition 4.

Notation and conventions. For an arbitrary cutoff $\tau \in [0, \bar{\tau}]$, write the social learning curve in the I-agent network as $\mathcal{B}^I(\tau) = \{\mathcal{B}^I_t(\tau)\}_t$; in the unique equilibrium, $b^I_t = \mathcal{B}^I_t(\tau^I)$ and $\psi_{\tau^I}(b^I) = 0$. Analogously, in the infinite regular n-tree $\mathcal{T}$, define $\mathcal{A}(\tau) = \{\mathcal{A}_t(\tau)\}_t$ as follows: For $t < \tau$, $\mathcal{A}_t(\tau) := 1$; for $t > \tau$, it is the solution of (19), $\dot{a} = (n-2)a(1-a)$ with boundary condition $a_\tau = 1 - e^{-(n-1)\tau}$. In equilibrium, $(\tau^*, a^* = \{a^*_t\})$ uniquely solve $a^* = \mathcal{A}(\tau^*)$ and $\psi_{\tau^*}(na^*) = 0$. Convergence of functions $b^I = \{b^I_t\}$ is always point-wise for all but at most one t, namely the cutoff $t = \bar{\tau}^*$.

We will prove that $\tau^I \to \tau^*$ and that $b^I_t \to na^*_t$ for all $t \neq \tau^*$. We restrict attention to a subsequence where τ^I converges to some τ^∞. The triangle inequality implies that for all I

$$\begin{aligned}|\psi_{\tau^\infty}(n\mathcal{A}^*(\tau^\infty))| \leq & |\psi_{\tau^\infty}(n\mathcal{A}^*(\tau^\infty)) - \psi_{\tau^\infty}(n\mathcal{A}^*(\tau^I))| + \\ & |\psi_{\tau^\infty}(n\mathcal{A}^*(\tau^I)) - \psi_{\tau^\infty}(\mathcal{B}^I(\tau^I))| + \\ & |\psi_{\tau^\infty}(\mathcal{B}^I(\tau^I)) - \psi_{\tau^I}(\mathcal{B}^I(\tau^I))| + |\psi_{\tau^I}(\mathcal{B}^I(\tau^I))|\end{aligned}$$

As $I \to \infty$, the first term vanishes by continuity of $\mathcal{A}^*_t(\tau)$ in τ for all $t \neq \tau^*$, and continuity of $\psi_{\tau^\infty}(b)$ in $b = \{b_t\}$. The second term vanishes by continuity of $\psi_{\tau^\infty}(b)$ in $b = \{b_t\}$ and because for all $t \geq 0$

$$\lim_{I \to \infty} \sup_{\tau \in [0, \bar{\tau}]} |\mathcal{B}^I_t(\tau) - n\mathcal{A}^*_t(\tau)| = 0 \tag{68}$$

as we show below. The third term vanishes because by Lemma 6(b). The fourth term is 0 for all I since τ^I is the equilibrium cutoff of $\mathcal{G}^I$

Thus, $\psi_{\tau^\infty}(n\mathcal{A}^*(\tau^\infty)) = 0$. Since τ^* is the unique solution of this equation, we have $\tau^\infty = \tau^*$. Since the subsequence of τ^I that converges to τ^∞ was arbitrary, the entire sequence τ^I converges to τ^* as desired. The triangle inequality then implies $|b^I_t - na^*_t| = |\mathcal{B}^I_t(\tau^I) - n\mathcal{A}^*_t(\tau^*)| \leq |\mathcal{B}^I_t(\tau^I) - n\mathcal{A}^*_t(\tau^I)| + n|\mathcal{A}^*_t(\tau^I) - \mathcal{A}^*_t(\tau^*)| \to 0$ for all $t \neq \tau^*$.

Proof of (68). The social learning converges for fixed τ. Fix an agent i, and consider times $t > \tau$. Let $G^{I,r}_i$ be the event that, i has n neighbors, $n(n-1)$ second neighbors, ..., $n(n-1)^{r-1}$ agents at distance r, and all of these agents are distinct. For all fixed r and t, $\lim_{I\to\infty} \Pr(G^{I,r}_i) \to 1$. For the upcoming arguments, we state that this convergence also conditional on the event $\{\theta = H, t < T_i, S_i\}$

$$\lim_{I \to \infty} \Pr{}^H(G^{I,r}_i | t < T_i, S_i) \to 1. \tag{69}$$

We will now define upper and lower bounds $\underline{a}^r_t(\tau), \bar{a}^r_t(\tau)$ for the expected effort of i's neighbors j in both the network $a^{I,r}_t(\tau) := E^H[A^I_{j,t} | G^{I,r}_i, t < T_i, S_i]$ and in the infinite tree

$\mathcal{A}_t(\tau)$. We show below that

$$\lim_{r\to\infty} \sup_{\tau\in[0,\bar{\tau}]} |\bar{a}_t^r(\tau) - \underline{a}_t^r(\tau)| = 0. \tag{70}$$

Then, by the triangle inequality

$$\begin{aligned}|\mathcal{B}_{i,t}^I(\tau) - n\mathcal{A}_t(\tau)| &\le |\mathcal{B}_t^I(\tau) - na_t^{I,r}(\tau)| + |nA_t^{I,r}(\tau) - n\mathcal{A}_t^*(\tau)| \\ &\le n(1 - \Pr{}^H(G_i^{I,r}|t < T_i, S_i)) + n|\bar{a}_t^r(\tau) - \underline{a}_t^r(\tau)|\end{aligned}$$

and so (69) and (70) imply

$$\lim_{r\to\infty}\lim_{I\to\infty}\sup_{\tau\in[0,\bar{\tau}]} |\mathcal{B}_{i,t}^I(\tau)-n\mathcal{A}_t^*(\tau)| \le \lim_{r\to\infty}\lim_{I\to\infty} n(1-\Pr{}^H(G_i^{I,r}|t<T_i,S_i)) + \lim_{r\to\infty}\sup_{\tau\in[0,\bar{\tau}]} n|\bar{a}_t^r(\tau)-\underline{a}_t^r(\tau)| = 0$$

which is (68), since the LHS does not depend on r. □

Proof of (70). Construction of the bounds $\underline{a}_t^r, \bar{a}_t^r$ and their convergence. We define the bounds $\underline{a}_t^r, \bar{a}_t^r$ (dropping τ for a moment to ease notation) as i's expectation over neighbor j's effort conditional on pessimistic/optimistic assumptions about successes of distant agents. Specifically, we define expectations $\underline{E}^{-i,r}, \bar{E}^{-i,r}$ over the first success times T_i of all agents k with distance $1, ..., r$ from i, both of which condition on $G_i^{I,r}, \theta = H$ and the fact that i's neighbors j have not seen i succeed. Additionally, $\underline{E}^{-i,r}$ conditions on no "leaf agent" ℓ with distance r from i having observed a success from an "outside" agent at distance $r+1$ from i; conversely, $\bar{E}^{-i,r}$ conditions on *every* "leaf agent" ℓ having observed a success from an "outside" agent. We then set $\underline{a}_t^r := \underline{E}^{-i,r}[A_{j,t}|t < T_j]$ and $\bar{a}_t^r := \bar{E}^{-i,r}[A_{j,t}|t < T_j]$.

We proceed by induction over r. For $r = 1$, this means $\underline{a}_t^1 \equiv 0, \bar{a}_t^1 \equiv 1$ for $t > \tau$. More generally, for $r > 1$, i's neighbor j shirks at $t > \tau$ iff none of his $n-1$ other neighbors $k \in N_j(G)\backslash\{i\}$ have succeeded.

$$1 - \underline{a}_t^r = \frac{\Pr^{-i}(t < T_j, t < T_k \forall k \in N_j\backslash\{i\})}{\Pr^{-i}(t < T_j)} = \frac{\exp\left(-n\tau - (n-1)\int_\tau^t \underline{a}_s^{r-1} ds\right)}{\exp\left(-\tau - \int_\tau^t \underline{a}_s^r ds\right)} \tag{71}$$

The last equality is analogous to the undirected line in Example 4: The denominator follows because the hazard rate of T_j equals 1 before τ and $\underline{a}_s^r$ after. In turn the event in the numerator has hazard rate n when all agents experiment before τ; after τ, having observed no success j shirks, while the expected effort of each of his $n-1$ neighbors k equals $\underline{a}_s^{r-1}$ since the event $G_i^{I,r}$ implies $G_j^{I,r-1}$. We rewrite (71) as an ODE

$$\dot{\underline{a}}^r = ((n-1)\underline{a}^{r-1} - \underline{a}^r)(1 - \underline{a}^r) \tag{72}$$

with initial condition $\underline{a}^r_\tau = 1 - e^{-(n-1)\tau}$. The upper bounds $\bar{a}^r_t$ also obey (72) with anchor $\bar{a}^1_t \equiv 1$.

Since successes outside $G^{I,r}_i$ only affect j's expected effort via the leaf agents, and the solution of (72) is monotone in $\underline{a}^{r-1}$, the so-defined functions indeed bound expected effort, $\underline{a}^r_t < A^{I,r}_t, \mathcal{A}^*_t < \bar{a}^r_t$. Moreover, the monotonicity of (72) together with $\underline{a}^1 \equiv 0$ implies that $\underline{a}^r$ increases in r and so converges to some $\underline{a}^\infty = \{\underline{a}^\infty_t(\tau)\}_t$ which must then solve (19), so $\underline{a}^\infty_t(\tau) = \mathcal{A}^*_t(\tau)$ for all t. Similarly, $\bar{a}^r(\tau) \to \mathcal{A}^*(\tau)$. Since $\underline{a}^r_t(\tau), \bar{a}^r_t(\tau), \mathcal{A}^*_t(\tau)$ are all increasing and equi-Lipschitz in τ, the convergence is uniform in $\tau \in [0, \bar{\tau}]$, so we have proven (70). □

Proof of parts (a) and (c). The only difference is the number of neighbors in (71) and (72). For n-regular directed network, we define $G^{I,r}_i$ as the event that i has n neighbors, n^2 second neighbors, ..., and n^r agents with distance r. Since i's neighbor j has n additional neighbors k, (72) becomes $\dot{\underline{a}}^r = (n\underline{a}^{r-1} - \underline{a}^r)(1 - \underline{a}^r)$ with boundary condition $\underline{a}^r_\tau = 1 - e^{-n\tau}$, so as $r \to \infty$, we obtain (16).

For triangular networks, i's neighbor j shares one more, triangular neighbor j' with i, as well as $n-2$ other neighbors k with distance two from i. Thus, (71) becomes

$$1 - \underline{a}^r_t = \frac{\Pr^i(t < T_j, T_{j'}, t < T_k \forall k \in N_j \backslash \{i, j'\})}{\Pr^i(t < T_j, T_{j'})} = \frac{\exp\left(-n\tau - \int_\tau^t ((n-2)\underline{a}^{r-1}_s + \underline{a}^r_s) ds\right)}{\exp\left(-2\tau - 2\int_\tau^t \underline{a}^r_s ds\right)}.$$

To understand the integral in the numerator, after τ expected effort is 0 for j, $\underline{a}^r_s$ for neighbor j', and $\underline{a}^{r-1}_s$ for each of the $n-2$ second neighbors k. Thus, (72) becomes $\dot{\underline{a}}^r = ((n-2)\underline{a}^{r-1} - \underline{a}^r)(1 - \underline{a}^r)$ with boundary condition $\underline{a}^r_\tau = 1 - e^{-(n-2)\tau}$, so as $r \to \infty$, we obtain (20).