The Impact of Connectivity on the Production and Diffusion of Knowledge

Gustavo Manso [†] and Farzad Pourbabaee [‡]

February 2024

Abstract

We study a social bandit problem featuring production and diffusion of knowledge. While higher connectivity enhances knowledge diffusion, it may reduce knowledge production as agents shy away from experimentation with new ideas and free-ride on the observation of other agents. As a result, under some conditions, greater connectivity can lead to homogeneity and lower social welfare.

[†]Haas School of Business, University of California Berkeley: manso@berkeley.edu.
[‡]California Institute of Technology: far@caltech.edu.

Contents

1 Introduction

Advances in travel and communication technologies have cleared the way for more connected organizations and societies. In well-connected structures, new ideas spread quickly leading to rapid innovation adoption.

While such enhanced knowledge diffusion is beneficial, it may come at the cost of reduced knowledge production. When an organization or society is well-connected, agents may shy away from experimentation with new ideas, since they can easily see the results of the experimentation efforts of other agents and adapt their actions accordingly. Because of this free riding, more connected organizations or societies may become homogeneous, converging on an inferior technology, and having lower overall welfare than less connected organizations or societies.

We study this tension between knowledge diffusion and knowledge production in a simple two-period social bandit model. In each period, each agent has the choice between exploiting a safe well-known action or exploring a risky novel action. At the end of the first period, each agent observes the outcome of a randomly selected group of agents. We show that in equilibrium social welfare is not necessarily increasing in connectivity between agents. That is, in a better connected society or organization, in which each agent is likely to meet with a greater number of agents, the costs of free riding on knowledge production may dominate the benefits of connectivity on knowledge diffusion, leading to lower social surplus.

We begin our analysis in Section 2 with a two-player economy. In this economy, equilibrium features three different regions based on initial beliefs: (i) both agents exploit; (ii) one agent exploits, while the other explores; (iii) both agents explore (Propositions 1 and 2). Due to free riding, there is over-exploitation and under-exploration relative to the social optimum (Proposition 3). Moreover, equilibrium social surplus is non-monotonic in the connection probability between the two agents. For some intermediate levels of connectivity, an increase in the connection probability leads to lower equilibrium social surplus.

In Section 3, we study the equilibrium in the multi-agent economy, and we show it resembles the equilibrium in a two-agent economy. In the sense that, for high (low) initial beliefs about the risky action, all players explore (exploit) in equilibrium (Theorem 1). For intermediate initial beliefs, equilibrium is asymmetric, with a given number of players exploring while the remaining players exploiting (Theorem 2).

The equilibrium results in Section 3 apply to any ensemble of random networks of connections. In particular, we apply them to economies with *local* and *global* connections, where every pair of agents are connected to each other independently with the same probability across all pairs. In the local case, each agent only observes the experimentation outcomes of her immediate neighbors, whereas in the global case her observable circle includes the entire set of agents who are connected to her. Thanks to the tractable results on Binomial processes, we provide asymptotic equilibrium analysis for local economies as the number of agents grows to infinity (Proposition 4). We find closed-form representation for the asymptotic *fraction* of exploring agents in the equilibrium, which turns out to be increasing in the initial belief and agents' patience. Importantly, it is *inversely* related to the average degree of connections, thereby confirming the free riding channel.

In the global case, we establish a rapid tightening of the exploration region when the number of agents an individual is expected to observe rises just above 1 (see Section 4.2). This effect is more significant for radical innovation, when the probability of success of the risky action is small. The intuition is that self-exploration is more beneficial to an agent when its expected future informational gain dominates the present cost of first period exploration. The informational gain is linked to the probability of making a breakthrough (individual success) and receiving failure signals from all other contacts (group failure). As the average degree of neighbors rises above one, the size of the giant connected component in the graph of connections becomes proportional to the number of agents, and hence the probability of group failures (with many members) rapidly falls, lowering the informational benefit to private exploration and thus significantly tightening the exploration region.

In Section 5, we investigate the equilibrium social surplus and compare it to the social optimum. As in the two-player economy, equilibrium social surplus is not increasing in the connectivity of the economy (Proposition 6). Relative to the social optimum, the equilibrium outcome exhibits over-exploitation and under-exploration (Theorem 4).

Higher connectivity exacerbates free riding. Since an agent observes the experimentation efforts of other agents, she may shy away from exploration herself, reducing the social surplus. Specifically, increasing the average degree of connections, *weakly* decreases the number of exploring agents. This number remains constant with respect to the connectivity index, and undergoes discrete drops (of size 1) at separated thresholds as a result of equilibrium regime change in the asymmetric region. On the intervals where the equilibrium number of exploring agents is constant, increasing connectivity enhances knowledge diffusion without affecting the free riding incentives, and hence increases the social surplus. However, at the thresholds where the economy goes through equilibrium regime change (by losing one previously exploring agent) the social surplus falls. Therefore, in the finite economy, the overall look of the social surplus with respect to the connectivity features increasing intervals with discontinuous falls on the thresholds.

In the economy with local connections, where the average degree of peers is constant, the size of these discontinuous drops remains *bounded* as the number of players (n) grows to infinity. Therefore, in the per-capita analysis they decay like $O(1/n)$ and the limit of per-capita equilibrium social surplus no longer exhibits the discontinuous falls appearing in the finite economies. This means the limiting average equilibrium social surplus is weakly increasing and continuous in connectivity index (Proposition 7). Furthermore, for intermediate levels of initial beliefs, we identify a connectivity threshold above which the limit of the per-capita equilibrium social surplus remains *constant*. Put differently, in the limit, the social informational gain of having one more agent exploring precisely offsets the current exploration cost, resulting in constancy with respect to the connectivity index.

Related Literature. In his seminal work Rothschild (1974) studies the *single-agent* experimentation problem in the two-armed bandit environment, and shows that with positive probability the agent settles on the sub-optimal arm. The literature on *multi-agent* strategic experimentation starts with the work of Bolton and Harris (1999) and Keller et al. (2005).[1] In both studies, players are completely connected to each other, that is each player can observe the experimentation outcome of *all* other players. Our paper interpolates the two ends of the experimentation spectrum, since we consider agents who are neither completely connected nor completely isolated from each other. By doing so, we are able to uncover the non-monotonicity of equilibrium social surplus with respect to the connectivity.

Bala and Goyal (1998), Gale and Kariv (2003) and Sadler (2020) study the social learning dynamics of *myopic* agents who are connected in networks and collect information from their neighbors to maximize their *short-run* payoff. Our two-period experimentation framework is a first stab to depart from these works by letting agents to have long-run incentives in their strategic interactions. In a recent work Board and Meyer-ter Vehn (2023) show that the social surplus is single-peaked in network density, namely after a certain connectivity level, social experimentation crowds out private learning and thus lowers the social surplus.

Issues such as long-run social conformity and information aggregation in the context of multi-agent strategic experimentation, when agents observe the actions and not the

[1] A non-exhaustive list of related papers in strategic bandits includes Heidhues et al. (2015), Keller and Rady (2015), Bonatti and Hörner (2017), and Pourbabaee (2020).

payoffs of others, are studied in Chamley and Gale (1994), Aoyagi (1998), Rosenberg et al. (2007), Rosenberg et al. (2009) and Camargo (2014). Aside from the observability of payoffs (rendering tractable equilibrium analysis) our paper differs from these studies in that it mainly focuses on the impact of connectivity on equilibrium strategies and social welfare rather than focusing on the long-run conformity of actions and/or social learning.

Our paper is also related to the broader literature of games with information sharing and externality. For example, Duffie et al. (2009) studies a continuum economy where individuals are initially endowed with informative signals and incur costly search to meet and share their information. Wolitzky (2018) investigates a social learning framework and innovation adoption where agents learn from a random sample of past outcomes, in that they arrive continuously over time and make once-and-for-all action. Also, in a Poisson news settings Frick and Ishii (2020) studies how the arrival rate of public signal (that depends on the mass of current adopters) could impact the adoption of innovation in the economy. Somewhat related to the tension between private experimentation and information sharing, Gordon et al. (2021) show that increasing the *observation delay* among members of a team (that work on a mutual project) increases the individual efforts.

Lastly, the analysis of our paper on how connectivity impacts exploration incentives has implications for designing optimal policies to motivate innovation and exploration in networked economies (e.g., Manso, 2011; Kerr et al., 2014).

2 Two-Player Economy

In this section, we propose a very simple model that aims to capture the essence of equilibrium forces and provide some intuition for the general case of $n > 2$ agents. We respectively study the perfect and imperfect connections (between the two agents), followed by the analysis of social surplus. Propositions 1-3 in this section are followed from the subsequent general case in Sections 3 and 5. Therefore, we will not provide separate proofs for them.

2.1 Perfect Connections

There are two agents i and j, and the game consists of two periods, i.e., $t \in \{0, 1\}$. Every agent faces a binary action choice in each period. Specifically, she can choose a safe action ($a = 0$ that is exploiting the status quo) with a normalized payoff of 0, or take a risky action ($a = 1$ exploring the other alternative). In the latter case, the return is a binary random variable, i.e., $y \in \{-\alpha, 1\}$ (with $\alpha \in (0,1)$) conditioned on the hidden state of the world $\theta \in \{0,1\}$, with the following conditional structure:

$$\mathsf{P}\left(y=1 \mid \theta=1\right) = \beta \in (0,1), \text{ and } \mathsf{P}\left(y=1 \mid \theta=0\right) = 0\,.$$

Therefore, receiving a high payoff of $y = 1$ is perfectly conclusive about the underlying state θ, i.e., the probability of receiving high payoff in the low state is zero. Let $\pi = \mathsf{P}(\theta = 1)$ be the initial prior of both players. The following timeline elucidates the order of events in this two-period economy:

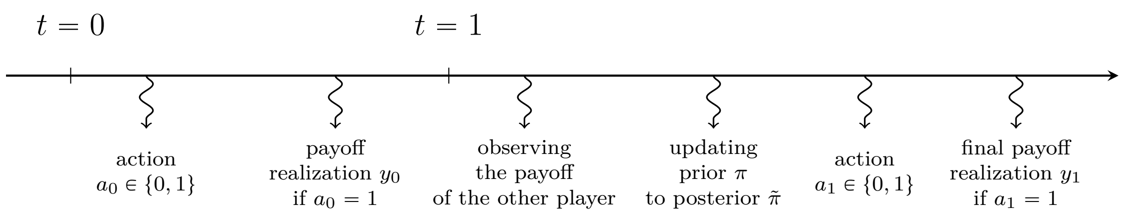

Figure 1: Timeline of the Two-Period Bandit

At the beginning of the second period, agent i gets to observe the outcome of agent j's experimentation, if j chose to pick the risky arm in the first period. This communication step among players is the main point of analysis throughout the paper. After that, she updates her prior about θ given $y_0(i)$ and $y_0(j)$, leading to the posterior $\tilde{\pi}$. Let $\pi_{\ell,m}$ denote the posterior when agent i observes $m \in \{0,1,2\}$ payoff signals, out of which $\ell \in \{0,1,2\}$ had high realizations (i.e., $y = 1$). Observe that it must be that $\ell \leqslant m$. Then,

$$\pi_{\ell,m} = 1_{\{\ell \geqslant 1\}} + \frac{\pi(1-\beta)^m}{1-\pi+\pi(1-\beta)^m}\, 1_{\{\ell=0\}}\,.$$

In particular, before the realizations of $y_0(i)$ and $y_0(j)$, $\tilde{\pi}$ is a random variable whose ex post values are denoted by $\pi_{\ell,m}$.

The game ends with each agent making a second action choice between the safe and the risky arm. Since each agent always has the safe option at hand, the expected payoff *after* Bayesian updating — that is when the agent holds posterior $\tilde{\pi}$ — is

$$\big(\tilde{\pi} - \alpha(1-\tilde{\pi})\big)^+ := \max\{\tilde{\pi} - \alpha(1-\tilde{\pi}), 0\}\,.$$

Therefore, in contrast with the strategic choice of period one (i.e., a_0), the decision in the second period a_1 follows a simple threshold rule: choose the risky arm if and only if $\tilde{\pi} > \frac{\alpha}{1+\alpha}$ (we choose to break the tie in favor of the safe arm).

There will be two types of *symmetric* equilibrium: *exploration* equilibrium in which both agents choose the risky arm in the first period (i.e., $a_0(i) = a_0(j) = 1$), and *exploitation* equilibrium where both agents select the safe arm in the first period (i.e., $a_0(i) = a_0(j) = 0$). The equilibrium is called *asymmetric* when one agent explores and the other one exploits (thus $a_0(i) \neq a_0(j)$). Let $\delta \in [0, 1]$ be the time discount factor, that is each agent values the payoffs in the first and second periods with the respective weights of $1-\delta$ and δ. This means that our agents are not myopic and they incorporate future gains from current exploration in their decision problem.[2]

Proposition 1. *There exist two thresholds $\underline{\pi} < \bar{\pi}$ such that the exploitation equilibrium appears on $[0, \underline{\pi}]$, and the exploration equilibrium appears on $(\bar{\pi}, 1]$. In the intermediate region $(\underline{\pi}, \bar{\pi}]$ the equilibrium is asymmetric with only one agent exploring. Closed form expressions for the cutoffs are*

$$\underline{\pi} = \frac{\alpha(1-\delta)}{(1+\alpha)(1-\delta)+\delta\beta}\,, \quad \bar{\pi} = \frac{\alpha(1-\delta)}{(1+\alpha)(1-\delta)+\delta\beta(1-\beta)}\,. \tag{2.1}$$

This result shows that the equilibrium number of explorers is weakly increasing in the initial belief. Two important comparative statics about the exploration incentives are the effect of patience (δ) and the signal precision (β) on the above thresholds.

As it appears from Figure 2a higher patience (namely higher δ) is associated with smaller exploration thresholds, thereby increasing the incentives to sacrifice the present payoff to learn about the risky arm and gain the benefits in the next period. Specifically, higher patience enlarges the exploration equilibrium region and shrinks the exploitation region.

Higher uncertainty about the risky arm (namely β closer to 1/2) is associated with higher gains from exploration, and hence lower exploration threshold. Figure 2b confirms this intuition. In addition, higher β increases the exploration gain upon receiving conclusive signals about θ more so than it raises the opportunity cost of exploration absent such signals. Therefore, it lowers the individual's incentive to exploit the safe arm, thereby shrinking the exploitation region (see $\underline{\pi}$ in Figure 2b).

[2]This is in contrast to the social learning models of Bala and Goyal (1998) and Sadler (2020) in which players are myopic. In particular, they collect information from their neighbors just to maximize their *present* period payoff.

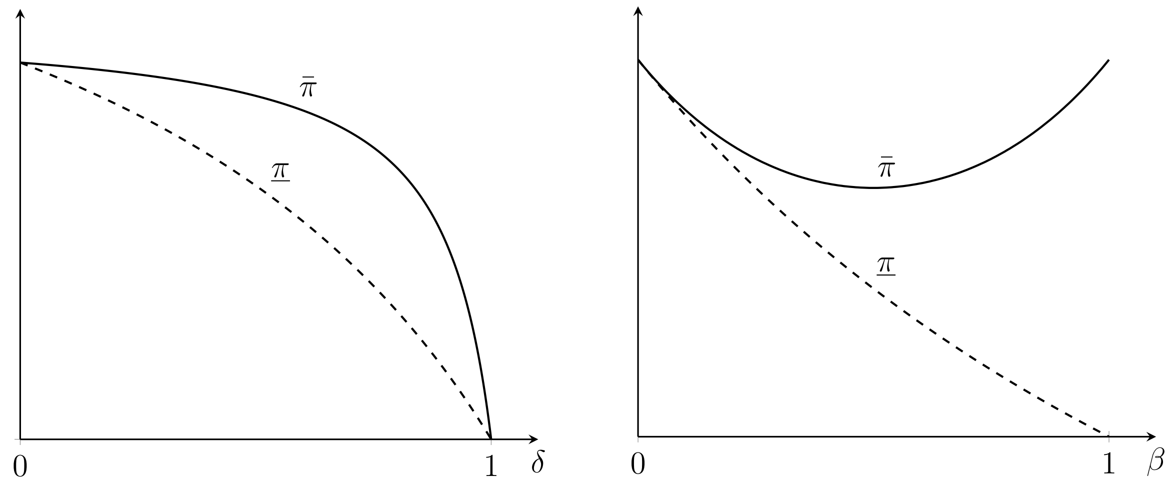

(a) Effect of δ

(b) Effect of β

Figure 2: Comparative statics of Thresholds

2.2 Imperfect Connections

Suppose the connection between the two players is *imperfect*. That is each agent gets to observe the outcome of the other agent's first period experimentation with probability p. In the next proposition, we show such imperfect communication will not impact the exploitation region and *expands* the exploration region.

Proposition 2. *In the presence of imperfect connections ($p < 1$), there exist two thresholds $\underline{\pi} < \bar{\pi}$ such that the exploitation equilibrium appears on $[0, \underline{\pi}]$, and the exploration equilibrium appears on $(\bar{\pi}, 1]$. In the intermediate region $(\underline{\pi}, \bar{\pi}]$ the equilibrium is asymmetric with only one agent exploring. Closed form expressions for the cutoffs are*

$$\underline{\pi} = \frac{\alpha(1-\delta)}{(1+\alpha)(1-\delta)+\delta\beta}\,, \quad \bar{\pi} = \frac{\alpha(1-\delta)}{(1+\alpha)(1-\delta)+\delta\beta(1-p\beta)}\,.$$

The important message behind this result is that $\mathrm{d}\bar{\pi}/\mathrm{d}p > 0$, therefore in this two-player economy weaker ties between agents correspond to higher levels of exploration. This is because stronger connections between players increase the free riding motives, and hence lowers the incentive for the first period exploration. This in turn translates to a higher belief threshold required for exploring the risky arm in the first period.

2.3 Social Surplus

So far we have analyzed the equilibrium response in the two-player bandit game with imperfect connections. One may wonder how the equilibrium response compares to the socially optimum behavior. For that, we subsequently investigate when the "benevolent" planner prescribes the exploitation or exploration by both agents.

Proposition 3. *The socially optimal outcome is for both players to exploit the safe arm whenever $\pi \leqslant \underline{\pi}^*$, and to jointly explore the risky arm on $\pi \geqslant \bar{\pi}^*$, where*

$$\underline{\pi}^* = \frac{\alpha(1-\delta)}{(1+\alpha)(1-\delta)+\delta\beta(1+p)}\,, \quad \bar{\pi}^* = \frac{\alpha(1-\delta)}{(1+\alpha)(1-\delta)+\delta\beta\,(1+p(1-2\beta))}\,.$$

The substantial lesson from this proposition is that the equilibrium outcome features over-exploitation ($\underline{\pi} > \underline{\pi}^*$) and under-exploration ($\bar{\pi} > \bar{\pi}^*$) relative to the social optimum (see the x-axis in Figure 3a).

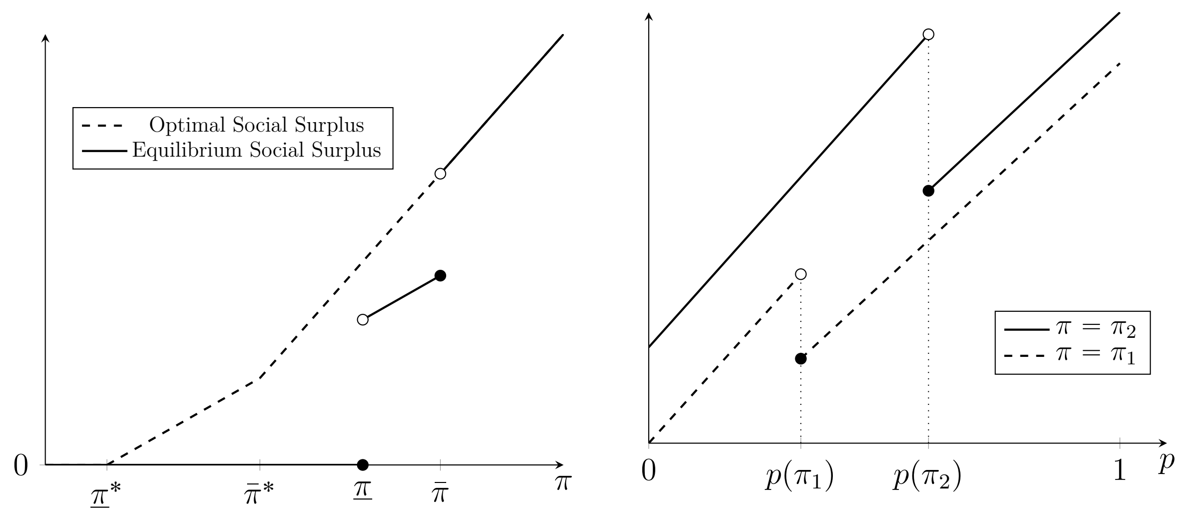

(a) Optimal vs. Equilibrium Social Surplus

(b) Equilibrium Social Surplus with respect to the Connection Probability p

Figure 3: Social Surplus

Figure 3a depicts the equilibrium and optimal social surplus in the two-player economy as a function of the initial belief π. Importantly, because of the inherent externality in this economy, the equilibrium social surplus is discontinuous at $\underline{\pi}$ and $\bar{\pi}$, where it undergoes equilibrium regime changes. As we will see in Section 5.2, the discontinuities in the average equilibrium social surplus remain bounded in large economies with local connections, therefore, they disappear as the number of individuals gets large.

We wrap up this section by investigating the effect of the connection probability p on the equilibrium social surplus. By examining the social surplus function in this two-player economy, one can readily show that it is *increasing* in p in each equilibrium region, and undergoes a single drop when there is a regime change from full exploration to the intermediate region of asymmetric equilibrium. This pattern is represented in Figure 3b, where the dependency of the equilibrium social surplus on p is plotted for two fixed levels of initial beliefs $\pi_2 > \pi_1$. Specifically, the exploration threshold $\bar{\pi}(p)$ found in Proposition 2 is increasing in p. Let $p(\pi)$ be the level at which $\bar{\pi}(p) = \pi$. On the region where $p < p(\pi)$, the full exploration equilibrium prevails and the social surplus *increases* by strengthening the connections until p surpasses $p(\pi)$, at which the equilibrium number of explorers drops from two to one. This creates the discontinuous fall in the equilibrium social surplus. Thereafter, raising the connection probability again *increases* the social surplus because it only raises the benefits of information sharing between the agents and does not alter the free riding incentives (as one of them is already exploiting the safe arm).

In Section 5.1, we study the average equilibrium social surplus for the economy with many players. There, we demonstrate that this pattern of being increasing in the connection probability on a fixed equilibrium regime, while discontinuously falling at thresholds of regime change is a robust feature of this economy with many players.

3 Equilibria with Multiple Agents

In this section, we study the general case of an economy consisting of n individuals, where each player in the second period observes the exploration outcome of a randomly selected group of individuals whose cardinality is denoted by the random variable M. It is crucial in our subsequent analysis to assume that the realization of this random group is not available to the agent in the first period, and she only knows the *distribution* of its size.

All agents are ex ante similar as of the beginning of the period one. In particular, the distribution of the size of the observable group in the second period is the same across all agents, and is denoted by the random variable M. In addition, all agents share a common initial belief π.

The random group of contacts, realized in the second period, for each agent could be her immediate neighbors in the graph of connections (referred to as the **local** case), or, at the other extreme, the set of all individuals who belong to her *connected component* (referred to as the **global** case).[3] In the latter case, each agent not only observes the signals of her immediate neighbors in the second period but also the signals of members in her connected component (denoted by $\mathcal{C}$ with the size of $M+1 := |\mathcal{C}|$) in the graph of social connections.[4] In this case, effectively we consider the second period as a collection of several message-passing sub-periods, during which each agent gets to observe the exploration outcome of every other agent connected to her via a path on the graph of connections. Importantly, as mentioned earlier, we assume the random realization of the connections is resolved in the beginning of period two.

At this stage we would rather not impose a specific probabilistic structure on the graph of connections (or, equivalently the distribution of M), as the following equilibrium results do not hinge on the specifics of the underlying random graph nor on the depth of the signal observability.

In Sections 3.1 and 3.2 below, we examine the pure strategy equilibria as a function of the initial belief π. Specifically, we demonstrate that for large (respectively, small) levels of the initial belief, the pure strategy equilibrium is symmetric and characterized by full exploration (respectively, full exploitation). In contrast, at intermediate belief levels, the equilibrium is asymmetric, involving both groups of explorers and exploiters. We then study the symmetric mixed-strategy equilibrium in Section 3.3.

3.1 Symmetric Pure-Strategy Equilibria

Here, we study two *symmetric pure-strategy* equilibria: exploitation and exploration equilibrium. The exploitation equilibrium occurs when *all* players choose the safe arm in the first period, and it prevails whenever the initial belief falls below the threshold $\underline{\pi}$ stated in equation (2.1). One can readily confirm this by comparing an individual's payoff from exploitation when everyone else is also exploiting (denoted by $w_0(\pi)$) with her exploration payoff when she is the only explorer (denoted by $v_1(\pi)$). This analysis implies that $w_0(\pi) \geqslant v_1(\pi)$ whenever $\pi \leqslant \underline{\pi}$. That is the condition for exploitation equilibrium remains the same as before (in spite of having more than two players and presence of imperfect connections).

The more intriguing case is the examination of the existence of the exploration equilibrium in which *all* agents choose the risky arm in the first period. Toward this, we define two payoff

[3] Our equilibrium analysis encompasses these two cases as well as all intermediate ones.

[4] The connected component of each player includes herself as well. Hence, in the later case, M denotes the number of *other* players connected to the current agent.

functions: $w_{n-1}(\pi)$ and $v_n(\pi)$. The former refers to the agent's payoff when she chooses to exploit in the first period (and optimally act in the second period), while all $n-1$ remaining agents are exploring in the first period. The latter is her payoff from exploration in the first period (when all n players choose the risky arm) and subsequently play optimally in the second period. The exploration equilibrium occurs whenever $v_n(\pi) > w_{n-1}(\pi)$.

Suppose all individuals except one are exploring in the first period. Let L be the random variable indicating the number of successful high outcomes (i.e., $y = 1$) that the agent under study observes, which is certainly less than or equal to M (size of her second period contacts). Let $\pi_{\ell,m} := \mathsf{P}\left(\theta = 1 \middle| L = \ell, M = m\right)$, then $\pi_{\ell,m} = 1$ whenever $\ell \geqslant 1$ and if $\ell = 0$:

$$\frac{\pi_{0,m}}{1-\pi_{0,m}} = \frac{\pi}{1-\pi}(1-\beta)^m .$$

In the second period she chooses the risky arm if $\pi_{\ell,m} > \alpha/(1+\alpha)$, leading to the payoff

$$\mathsf{E}\left[\theta-\alpha(1-\theta)\middle|L = \ell, M = m\right]^+ = \left[\pi_{\ell,m}-\alpha(1-\pi_{\ell,m})\right]^+ .$$

When all others are exploring in the first period, her payoff from exploitation is

$$\begin{aligned} w_{n-1}(\pi) &= \delta \sum_{m=0}^{n-1} \sum_{\ell=0}^{m} \mathsf{P}\left(L = \ell, M = m\right) \mathsf{E}\left[\theta-\alpha(1-\theta)\middle|L = \ell, M = m\right]^+ \\ &= \delta \sum_{m=0}^{n-1} \sum_{\ell=0}^{m} \mathsf{E}\left[\theta-\alpha(1-\theta); L = \ell, M = m\right]^+ . \end{aligned} \tag{3.1}$$

Let $q(m) := \mathsf{P}(M = m)$, representing the probability of a randomly selected agent observing the exploration outcomes of m other players. Then, the above payoff can be written as

$$\begin{aligned} w_{n-1}(\pi) &= \delta \sum_{m=0}^{n-1} \sum_{\ell=0}^{m} q(m) \left[\pi \binom{m}{\ell} \beta^\ell (1-\beta)^{m-\ell} - \alpha(1-\pi) 1_{\{\ell=0\}}\right]^+ \\ &= \delta \sum_{m=0}^{n-1} \underbrace{q(m)\left[\pi(1-\beta)^m-\alpha(1-\pi)\right]^+}_{\mathsf{E}[\theta-\alpha(1-\theta);M=m,L=0]^+} + \underbrace{\delta\pi \sum_{m=0}^{n-1} q(m)\left(1-(1-\beta)^m\right)}_{\delta\mathsf{E}[\theta-\alpha(1-\theta);L>0]} . \end{aligned}$$

Hence the exploitation payoff is decomposed into two components: the expected payoff when no conclusive signal is observed ($L = 0$) and the expected payoff when at least one out of M contacts received a successful outcome ($L > 0$).

Now, suppose the agent chooses to explore in the first period, and denote her random realization of the risky arm by $y_0 \in \{-\alpha, 1\}$. Then, her expected payoff from exploration is

$$\begin{aligned} v_n(\pi) &= (1-\delta)\,(\pi-\alpha(1-\pi)) \\ +\delta \sum_{m=0}^{n-1}\sum_{\ell=0}^{m}\sum_{y\in\{-\alpha,1\}} &\mathsf{P}\,(M=m, L=\ell, y_0=y)\,\mathsf{E}\left[\theta-\alpha(1-\theta)\middle|M=m, L=\ell, y_0=y\right]^+ . \end{aligned} \tag{3.2}$$

The second term, representing the discounted expected payoff in period two, decomposes into two sums:

$$\begin{aligned} \text{discounted expected payoff} &= \delta\sum_{m,\ell}\mathsf{E}\,[\theta-\alpha(1-\theta); M=m, L=\ell, y_0=1]^+ \\ &+\delta\sum_{m,\ell}\mathsf{E}\,[\theta-\alpha(1-\theta); M=m, L=\ell, y_0=-\alpha]^+ . \end{aligned}$$

The first one is the expected payoff in the second period when a success was achieved in the first period, and the second one is the same quantity conditioned on receiving a low output.

Further rearrangements imply that the sum above is equal to

$$\begin{aligned} &\delta\pi\beta+\sum_{m,\ell} q(m)\left[\pi\binom{m}{\ell}\beta^\ell(1-\beta)^{m+1-\ell}-\alpha(1-\pi)1_{\{\ell=0\}}\right]^+ \\ = &\underbrace{\delta\pi\beta}_{\delta\mathsf{E}[\theta-\alpha(1-\theta);y_0=1]} +\delta\sum_{m=0}^{n-1}\underbrace{q(m)\left[\pi(1-\beta)^{m+1}-\alpha(1-\pi)\right]^+}_{\mathsf{E}[\theta-\alpha(1-\theta);M=m,L=0,y_0=0]^+}+\underbrace{\delta\pi(1-\beta)\sum_{m=0}^{n-1}q(m)\,(1-(1-\beta)^m)}_{\delta\mathsf{E}[\theta-\alpha(1-\theta);L>0,y_0=-\alpha]} . \end{aligned}$$

Formally, the exploration equilibrium occurs when $v_n(\pi) > w_{n-1}(\pi)$ — in other words, when the combination of the present payoff from exploration and the discounted exploration gain in presence of conclusive signals ($L > 0$ or $y_0 = 1$) in the second period exceeds the discounted opportunity cost of exploration in the absence of such signals ($L = 0$ and $y_0 = -\alpha$), that is when

$$\begin{aligned} &\overbrace{(1-\delta)\,(\pi-\alpha(1-\pi))}^{\text{present payoff}} + \\ &\delta\Big\{\mathsf{E}\,[\theta-\alpha(1-\theta); y_0=1]+\mathsf{E}\,[\theta-\alpha(1-\theta); L>0, y_0=-\alpha]-\mathsf{E}\,[\theta-\alpha(1-\theta); L>0]\Big\} \\ &> \delta\sum_{m=0}^{n-1}\left(\mathsf{E}\,[\theta-\alpha(1-\theta); M=m, L=0]^+-\mathsf{E}\,[\theta-\alpha(1-\theta); M=m, L=0, y_0=-\alpha]^+\right) \\ &= \text{discounted opportunity cost of exploration absent conclusive signals.} \end{aligned} \tag{3.3}$$

Theorem 1 (Exploration equilibrium)**.** *Let M be the size of the random group of contacts* *in the second period. Then, the exploration equilibrium appears on $\pi > \bar{\pi}$, where*

$$\bar{\pi} = \frac{\alpha(1-\delta)}{(1+\alpha)(1-\delta)+\delta\beta\mathsf{E}\left[(1-\beta)^{M}\right]}\,. \tag{3.4}$$

The important comparative static is the effect of the sparsity of connections on the exploration threshold. Since $x \mapsto (1-\beta)^x$ is a decreasing function, if the distribution of M positively shifts in the sense of first-order stochastic dominance, then the exploration threshold rises, equivalently the exploration region tightens. That is denser connections raise the the free riding incentives and thus are associated with higher thresholds for exploration in the equilibrium.

Remark 1. Note that in the case of local connections $M = D$, which is the degree of a randomly picked agent. And in the global connections scenario $M = |\mathcal{C}|-1$, where $\mathcal{C}$ is the connected component of a randomly chosen individual in the graph of social connections. The result of the previous theorem applies to these two important cases, as well as to any other choice for the distribution of M. In Section 4, we apply this result to random Erdos-Renyi graphs, thereby presenting sharper comparative statics for the exploration threshold $\bar{\pi}$.

3.2 Asymmetric Pure-Strategy Equilibria

In the previous section, we studied the equilibria in which *all* agents were either exploring or exploiting, and thus choosing symmetric equilibrium strategies. In this part, we focus on the equilibria in the intermediate region, where $\pi \in (\underline{\pi}, \bar{\pi}]$. Specifically, we study pure-strategy equilibria in which both types of agents (explorers and exploiters) are present. Let $0 < k < n$, and v_k (respectively, w_k) denote the expected payoff of an exploring (respectively, exploiting) agent when there are a *total* of k individuals exploring in the economy. This will be an equilibrium outcome if the exploring agents have no incentive to revert to exploitation, equivalently $v_k(\pi) > w_{k-1}(\pi)$, and when the exploiting agents find it costly to explore, namely $w_k(\pi) \geqslant v_{k+1}(\pi)$.

Let $q_k(m) := \mathsf{P}_k(M = m)$ denote the probability of observing the exploration outcomes of m individuals out of the k who explored in the first period. Following the recipe of equations (3.1) and (3.2), the payoff functions take the following forms:

$$\begin{aligned} w_k(\pi) &= \delta \mathsf{E}_k\left[\left(\pi(1-\beta)^M-\alpha(1-\pi)\right)^+\right]+\delta\pi\mathsf{E}_k\left[1-(1-\beta)^M\right], \\ v_k(\pi) &= (1-\delta)\big(\pi-\alpha(1-\pi)\big)+\delta\pi\beta+\delta\mathsf{E}_{k-1}\left[\left(\pi(1-\beta)^{M+1}-\alpha(1-\pi)\right)^+\right] \\ &\quad +\delta\pi(1-\beta)\mathsf{E}_{k-1}\left[1-(1-\beta)^M\right]. \end{aligned} \tag{3.5}$$

Note that above, we used the random variable M repeatedly in all expectation operators. However, it is important to interpret this notation with caution, as what truly matters is the distribution of M, which is determined by the subscript of the outer expectation symbol E. For instance, when E_k is used, it means that $\mathsf{P}(M=m)=\mathsf{P}_k(M=m)=q_k(m)$.

As a first step toward analyzing such equilibria, we show that for large values of π, the second incentive constraint (i.e., $w_k(\pi)\geqslant v_{k+1}(\pi)$) fails to hold.

Lemma 1. *Suppose $\pi>\alpha/(1+\alpha)$, then $w_k(\pi)<v_{k+1}(\pi)$ for every k.*

This lemma establishes that a pure-strategy equilibrium with non-zero number of exploiters cannot exist when $\pi>\alpha/(1+\alpha)$. In this region, only the full exploration equilibrium is sustainable. Therefore, to identify intermediate equilibria, whether pure or mixed, we need to focus exclusively on the region $\pi\leqslant\alpha/(1+\alpha)$. On this region all terms that include $(\cdot)^+$ inside the expectation operators in (3.5) are zero, and the following theorem results.

Theorem 2 (Asymmetric pure-strategy equilibrium)**.** *The asymmetric equilibrium in which k players explore, where $0<k<n$, exists if and only if*

$$\frac{\alpha(1-\delta)}{(1+\alpha)(1-\delta)+\delta\beta\mathsf{E}_{k-1}\left[(1-\beta)^M\right]}<\pi\leqslant\frac{\alpha(1-\delta)}{(1+\alpha)(1-\delta)+\delta\beta\mathsf{E}_{k}\left[(1-\beta)^M\right]}. \tag{3.6}$$

Using the expressions in (3.5), the left inequality in (3.6) is implied by the incentive constraint $v_k>w_{k-1}$, and the right inequality results from $w_k\geqslant v_{k+1}$, thus we omit the formal proof. Henceforth, in an economy of n agents we define the threshold $\pi_{k,n}$ as

$$\pi_{k,n}:=\frac{\alpha(1-\delta)}{(1+\alpha)(1-\delta)+\delta\beta\mathsf{E}_{k}\left[(1-\beta)^M\right]}.$$

As a result of the previous theorem, the asymmetric equilibrium with k agents exploring occurs whenever $\pi_{k-1,n}<\pi\leqslant\pi_{k,n}$. The full exploitation appears on $\pi\leqslant\pi_{0,n}\equiv\underline{\pi}$ and the full exploration appears on $\pi>\pi_{n-1,n}\equiv\bar{\pi}_n$. Furthermore, let $M_k^{(n)}$ be the random variable standing for the number of second period contacts of an individual in an economy that has

n agents, among them k are exploring the risky arm in the first period.[5] Then a simple stochastic dominance analysis implies that the distribution of $M_k^{(n)}$ first-order stochastically dominates that of $M_{k-1}^{(n)}$, and hence $\pi_{k-1,n} \leqslant \pi_{k,n}$. This means that the number of exploring agents in the equilibrium weakly *increases* in π.[6]

3.3 Symmetric Mixed-Strategy Equilibrium

In the previous two sections, we characterized all possible pure-strategy equilibria as a function of the initial belief π. Specifically it is shown in Section 3.1, that at the two ends, namely, on the intervals $[0, \underline{\pi}]$ and $(\bar{\pi}, 1]$, the unique equilibrium is, respectively, full exploitation and full exploration. Based on the results of Section 3.2, we can conclude that in the intermediate region $(\underline{\pi}, \bar{\pi}]$, the unique pure-strategy equilibrium is asymmetric, featuring both types of explorers and exploiters. This leaves open the possibility of existence of mixed-strategy equilibria in the intermediate region.

Hence in this section, we study the *symmetric* mixed-strategy equilibrium in the intermediate region. Suppose that each agent explores the risky arm with probability μ. This will be a mixed-strategy equilibrium if the expected payoff from exploitation, namely

$$w(\pi; \mu) = \sum_{k=0}^{n-1} \binom{n-1}{k} \mu^k (1-\mu)^{n-1-k} w_k(\pi)\,,$$

matches the expected payoff from exploration, that is

$$v(\pi; \mu) = \sum_{k=0}^{n-1} \binom{n-1}{k} \mu^k (1-\mu)^{n-1-k} v_{k+1}(\pi)\,.$$

Before presenting the equilibrium result, we provide a definition for a random graph to be *exchangeable*, which is a requirement for the next proposition

Definition 1 (Exchangeability)**.** The random structure of connections is called *exchangeable* if the probability of any event on the graph does not change with relabeling the vertices.

Theorem 3 (Symmetric mixed-strategy equilibrium)**.** *In the intermediate region, i.e.,* $\pi \in (\underline{\pi}, \bar{\pi}]$*, when connections are exchangeable, there exists a unique symmetric mixed-strategy*

[5]Depending on the context, we either use $M_k^{(n)}$ or explicitly specify the indices on the expectation operator, that is e.g., $\mathsf{E}_k^{(n)}$.

[6]The term 'weakly' is used because over each interval $(\pi_{k-1,n}, \pi_{k,n}]$ the equilibrium number of explorers is constant.

equilibrium. Moreover, the equilibrium probability of exploration μ is increasing in π, and decreasing in n.

The above theorem has two substantial implications: first, the uniqueness of the symmetric mixed-strategy equilibrium, and second, that the equilibrium probability of exploration is a decreasing function of the number of agents, thus affirming the presence of the free riding force in the setting of mixed-strategies.

4 Large n Limits of Equilibria

For the first time in this paper, we introduce a specific assumption regarding the random nature of graph connections. In particular, we assume that every pair of agents is connected with a probability $p = \lambda/n$, where λ represents the average number of immediate neighbors of a randomly chosen agent. It readily follows that the induced random graph model is exchangeable. We will first examine the case of local connections, where $M = D$ (i.e., the degree of a certain vertex in the graph of connections). Then, we will study global connections, where $M = |\mathcal{C}|-1$ (i.e., the size of the connected component minus one). In both cases, we focus on the limiting behavior as $n \to \infty$, and study the effect of λ on the equilibrium outcomes.

4.1 Local Connections

Recall that in the local regime $M = D$, the degree of a randomly drawn agent, that has the Binomial distribution $\mathsf{Bin}(n-1, p)$. For a constant λ, the Binomial distribution converges *weakly* to $\mathsf{Poisson}(\lambda)$, and therefore in the local regime the limit of exploration threshold is:

$$\begin{aligned} \bar{\pi}_\infty^{\text{local}} := \lim_{n\to\infty} \bar{\pi}_n^{\text{local}} &= \lim_{n\to\infty} \frac{\alpha(1-\delta)}{(1+\alpha)(1-\delta)+\delta\beta\mathsf{E}\left[(1-\beta)^D\right]} \\ &= \frac{\alpha(1-\delta)}{(1+\alpha)(1-\delta)+\delta\beta e^{-\lambda\beta}}\,, \end{aligned} \tag{4.1}$$

where in the last equality we applied the moment generating function of Poisson distribution, because as $n \to \infty$, D converges in distribution to $\mathsf{Poisson}(\lambda)$.

Lemma 2. *In the local regime, the exploration threshold $\bar{\pi}_n$ is eventually increasing in n and converges to $\bar{\pi}_\infty^{local}$ in (4.1).*

Proof. To justify $\bar{\pi}_{n+1} > \bar{\pi}_n$, we employ equation (3.4) along with the notation in footnote 5, and show that

$$\mathsf{E}_{n-1}^{(n)}\left[(1-\beta)^M\right] > \mathsf{E}_{n}^{(n+1)}\left[(1-\beta)^M\right].$$

This is indeed true because $\mathsf{E}_{n-1}^{(n)}\left[(1-\beta)^M\right] = \left(1-\frac{\lambda\beta}{n}\right)^{n-1}$, which is eventually decreasing in n (as $x \mapsto (x-1)\log(1-\lambda\beta/x)$ has negative derivative with respect to x for large x). $\square$

Based on this lemma, we can conclude that as the number of agents increases, the free riding incentives become stronger, albeit at a diminishing rate.

Lemma 3. *In the local regime, for a fixed $k \in \mathbb{N}$, and large enough n the following ordering holds: $\pi_{k-1,n} \leqslant \pi_{k,n+1} \leqslant \pi_{k,n} \leqslant \pi_{k+1,n+1}$.*

This lemma demonstrates that for large economies with local connections adding one more individual *never* results in fewer exploring agents in the equilibrium. Put differently, the existing agents do not switch to the exploitation status when new individuals join the economy.[7]

Let k_n denote the equilibrium number of exploring agents. The following proposition shows that in an economy with local connections, the equilibrium fraction of explorers k_n/n converges as n grows. The proof relies on using the incentive condition (3.6) to establish matching upper and lower bounds for k_n.

Proposition 4 (Limiting fraction of explorers)**.** *Let $k_n(\pi)$ be the equilibrium number of exploring agents in an economy of n individuals with local connections, then:*

$$\lim_{n\to\infty}\frac{k_n(\pi)}{n} = \kappa(\pi) := \begin{cases} 0 & \pi \leqslant \underline{\pi} \\ \frac{1}{\lambda\beta}\log\frac{\delta\pi\beta}{(1-\delta)(\alpha(1-\pi)-\pi)} & \underline{\pi} < \pi < \bar{\pi}_\infty^{local} \\ 1 & \pi \geqslant \bar{\pi}_\infty^{local} \end{cases} \tag{4.2}$$

Figure 4 illustrates the limiting fraction of exploring agents, i.e., $\kappa(\pi)$, as a function of the initial belief π. The function exhibits two kinks at $\underline{\pi}$ and $\bar{\pi}_\infty^{\text{local}}$, where there are equilibrium regime changes from full exploitation to the intermediate asymmetric region and then to the full exploration. Importantly, the graph exhibits convexity, indicating that the equilibrium fraction of explorers grows at an increasing rate as the initial belief rises.

[7]It is important to note that this conclusion primarily rests on keeping the average degree λ constant while expanding the size of the economy.

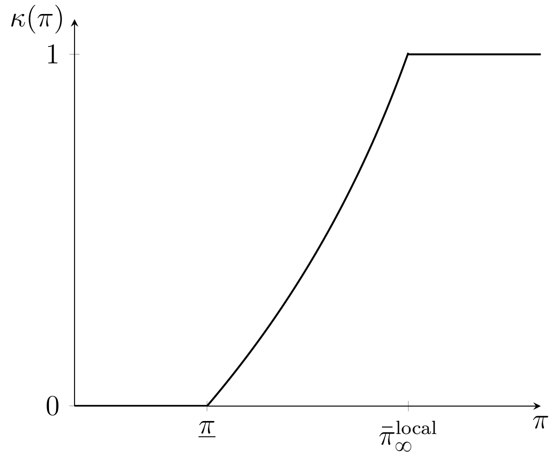

Figure 4: Limiting fraction of explorers

4.2 Global Connections

The analysis in the global regime (where $M = |\mathcal{C}|-1$) is rather intricate. In this regime, an agent meets all members of her connected component in the second period. Hence, in this setting, there is a greater potential for knowledge diffusion due to the increased number of connections, while simultaneously, there are stronger free riding incentives, leading to lower levels of exploration.

One can readily see (via a coupling argument, e.g., Theorem 2.1 in Bollobás (2001)) that the distribution of the size of the connected component $|\mathcal{C}|$ is first-order stochastically increasing in λ, and since $x \mapsto (1-\beta)^x$ is a decreasing function, then $\bar{\pi}$ becomes increasing in λ, thereby confirming the free riding force in the global case. To study the limiting behavior of the exploration threshold in this regime, we need an asymptotic result on the limiting distribution of $|\mathcal{C}|$. Let T be the random variable indicating the total number of descendants of a Branching process with $\mathsf{Poisson}(\lambda)$ offspring distribution. With the help of few lemmas from the literature of Erdos-Renyi random graphs, we show that $|\mathcal{C}|$ weakly converges to T, and hence the following asymptotic result follows.

Proposition 5. *Let $p = \lambda/n$, and T be the number of total progenies in a Branching process with a $\mathsf{Poisson}(\lambda)$ offspring distribution, then*

(i) $|\mathcal{C}|$ *converges in distribution to T, where* $\mathsf{P}(T = k) = \frac{e^{-\lambda k}(\lambda k)^{k-1}}{k!}$*, and*

(ii) as $n \to \infty$:

$$\bar{\pi}_{\infty}^{global} := \lim_{n\to\infty} \bar{\pi}_{n}^{global} = \frac{\alpha(1-\delta)}{(1+\alpha)(1-\delta)+\delta\beta\mathsf{E}\left[(1-\beta)^{T-1}\right]}\,. \tag{4.3}$$

To further analyze $\bar{\pi}_{\infty}^{\text{global}}$ derived from equation (4.3), one needs to determine the moment generating function (henceforth, MGF) of T, i.e., the number of the descendants in a Poisson Branching process. Following from Section 10.4 of Alon and Spencer (2000), the MGF of T is determined by a fixed-point relation. Specifically, fix $z \in [0,1]$ and let $X_1 \sim \mathsf{Poisson}(\lambda)$ denote the number of first-generation offspring, then the MGF of T is pinned down by

$$\begin{aligned}\psi(z) := \mathsf{E}\left[z^T\right] &= \mathsf{E}\Big[\mathsf{E}\left[z^T \middle| X_1\right]\Big] = \sum_{k=0}^{\infty} \mathsf{P}\left(X_1 = k\right) z\psi(z)^k \\ &= \sum_{k=0}^{\infty} \frac{e^{-\lambda}\lambda^k}{k!} z\psi(z)^k = ze^{\lambda(\psi(z)-1)}\,.\end{aligned}$$

The solutions to the equation $xe^x = y$ are denoted by the Lambert-W function, and based on the above expression one obtains:[8]

$$-\lambda\psi(z)e^{-\lambda\psi(z)} = -\lambda ze^{-\lambda} \Rightarrow \psi(z) = -\frac{1}{\lambda}\mathsf{W}(-\lambda ze^{-\lambda})\,. \tag{4.4}$$

Rapid fall of exploration in the global regime (small β and $\lambda \approx 1$). As Figure 5 illustrates, there is a rapid tightening of the exploration region in the case of global connections as λ increases from values slightly below 1 to values just above. More precisely, the marginal effect of increasing λ on the exploration threshold $\bar{\pi}_{\infty}^{\text{global}}$ undergoes a significant shift at $\lambda = 1$. This effect is more significant when β is close to zero, which is the most relevant region in the innovation and entrepreneurship research, when the probability of success is extremely small.

We can mathematically justify this sudden fall of exploration incentives by studying the effect of λ on $\bar{\pi}_{\infty}^{\text{global}}$ in equation (4.3). The only place where λ has an impact is through the MGF expression in the denominator, i.e., $\mathsf{E}\left[(1-\beta)^T\right]$. Therefore, we examine the change in the derivative of this component near $\lambda = 1$, and specifically its second derivative at this point. Let $z := 1-\beta$, then (4.4) implies that $\mathsf{E}\left[(1-\beta)^T\right] = \psi_\lambda(z)$. Dropping z from ψ's argument, we denote the first and second derivatives of ψ with respect to λ by ψ'_λ and ψ''_λ,

[8] We pick the solution branch of the Lambert-W function that guarantees $\psi(z) \leqslant 1$. For further details about this function see Corless et al. (1996).

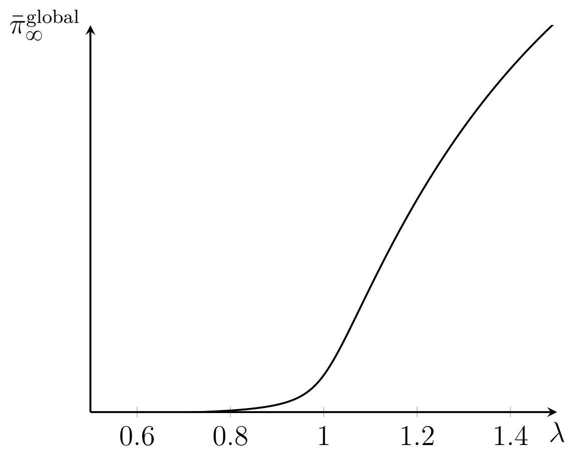

Figure 5: Rapid tightening of the exploration region
$[\delta = 0.15, \alpha = 1, \beta = 0.002]$

respectively:

$$\psi'_\lambda = \psi\big(\psi - 1 + \lambda\psi'_\lambda\big)\,,$$
$$\psi''_\lambda = \psi'_\lambda\big(3\psi - 1 + \lambda\psi'_\lambda\big)\,.$$

At $\lambda = 1$ the above expressions imply that

$$\psi''_1 = \frac{\psi_1(1 - 2\psi_1)}{1 - \psi_1}\,.$$

Since $\lim_{\beta\to 0}(1-\psi_\lambda) = 0$, then the above ratio explodes as $\beta \to 0$, thereby highlighting the rapid change in the sensitivity of the exploration threshold $\bar{\pi}_\infty^{\text{global}}$ with respect to the average number of connections at $\lambda = 1$.

The underlying intuition behind this rapid tightening is that the informational gain appearing in the incentive problem of a potential explorer is linked to the probability of making a breakthrough (individual success) while simultaneously receiving failure signals from all other contacts (group failure). As λ increases just above 1, the size of the giant connected component (and therefore, with high probability, the size of a randomly selected component) becomes proportional to the number of agents, leading to a rapid reduction in the probability of group failures. This decreases the informational benefit to private exploration and thus significantly tightens the exploration region.

5 Social Surplus

In this section, we examine the properties of the social surplus function in an economy with local connections and homogeneous connection probability (namely the one studied in Section 4.1). We begin by analyzing the limit of the equilibrium average social surplus as n grows large, focusing on the three equilibrium regions characterized in Section 3. Then, we study the social optimum, and we demonstrate that similar to the two-player case, over-exploitation and under-exploration are robust features of this economy in spite of the large number of players. We further determine the regions where the social surplus is monotone (increasing or decreasing) with respect to the number of exploring agents.

Suppose that out of n players, k agents choose the risky arm in the first period, and let the resulting social surplus be denoted as $u_{k,n}(\pi)$. Further, in the local regime, let $q_a(b) = \binom{a}{b} p^b (1-p)^{a-b}$ represent the probability of meeting b agents out of a specific set of a individuals in the second period, then

$$
\begin{gathered}
u_{k,n}(\pi) = (1-\delta)k\left(\pi-\alpha(1-\pi)\right) \\
+\delta k\pi\beta+\delta k\sum_{m=0}^{k-1} q_{k-1}(m)\left[\pi(1-\beta)^{m+1}-\alpha(1-\pi)\right]^{+}+\delta k\sum_{m=0}^{k-1} q_{k-1}(m)\pi(1-\beta)\left(1-(1-\beta)^{m}\right) \\
+\delta(n-k)\sum_{m=0}^{k} q_k(m)\left[\pi(1-\beta)^{m}-\alpha(1-\pi)\right]^{+}+\delta(n-k)\sum_{m=0}^{k} q_k(m)\pi\left(1-(1-\beta)^{m}\right).
\end{gathered}
\tag{5.1}
$$

The first line in $u_{k,n}$ denotes the first period payoff of exploration earned by the k exploring agents who chose the risky arm in the first period. The second line represents the discounted second period payoff of this group, consisting of three components: the discounted expected payoff when each agent received a conclusive signal in the first period (and optimally chose the risky arm in the second period); the discounted expected payoff when neither the agent nor any of her second-period contacts received a conclusive signal; and lastly, the discounted expected payoff when the individual herself did not receive a high output in the first period, but at least one of her second-period contacts did. The third line represents the discounted second period payoff of the remaining $n-k$ exploiting agents who chose the safe arm in the first period. This is composed of two components: their payoff when none of their contacts in the exploring group received a high output in the first period, and their payoff when at least one of them did receive such a conclusive signal.

Leveraging the above representation, the following lemma examines the marginal value of adding one more explorer to the economy, denoted by $\Delta u_k := u_{k+1} - u_k$. This will prove useful later when we investigate the equilibrium social surplus and the social optimum. We further use the notation $Q_a(b) := \sum_{m \leqslant b} q_a(m)$ to refer to the cumulative function of q_a, with the additional definition that $Q_0(0) = q_0(0) = 1$.

Lemma 4. *The marginal value of one more exploring agent takes the following form:*

(i) On $\frac{\pi}{1-\pi} \leqslant \alpha$, *it holds that*

$$\Delta u_k(\pi) = (1-\delta)\big(\pi-\alpha(1-\pi)\big)+\delta\pi\beta(1-p\beta)^k\left(1+(n-1)p-\frac{kp(1-p)\beta}{1-p\beta}\right). \tag{5.2}$$

(ii) On $\frac{\pi}{1-\pi} \geqslant \frac{\alpha}{(1-\beta)^r}$ *and* $0 \leqslant k < r$, *one has* $\Delta u_k(\pi) = (1-\delta)\big(\pi-\alpha(1-\pi)\big)$.

(iii) On $\frac{\alpha}{(1-\beta)^r} \leqslant \frac{\pi}{1-\pi} \leqslant \frac{\alpha}{(1-\beta)^{r+1}}$ *and* $k \geqslant r \geqslant 0$, *it holds that* $\Delta u_k(\pi) = \pi B_k - \alpha(1-\pi)A_k$, *where*

$$\begin{aligned} A_k(\pi) &:= (1-\delta)+\delta(k+1)Q_k(r-1) \\ &\quad+\delta(n-k-1)Q_{k+1}(r)-\delta k Q_{k-1}(r-1)-\delta(n-k)Q_k(r), \end{aligned} \tag{5.3}$$

$$\begin{aligned} B_k(\pi) := (1-\delta)-\delta(k+1)(1-\beta)\sum_{m=r}^{k} q_k(m)(1-\beta)^m-\delta(n-k-1)\sum_{m=r+1}^{k+1} q_{k+1}(m)(1-\beta)^m \\ +\delta k(1-\beta)\sum_{m=r}^{k-1} q_{k-1}(m)(1-\beta)^m+\delta(n-k)\sum_{m=r+1}^{k} q_k(m)(1-\beta)^m. \end{aligned}$$

The proof follows easily when we observe that the piecewise linear components in (5.1) are positive as long as $m+1 \leqslant r$ in the first component and $m \leqslant r$ in the second component. Hence, we omit the proof.

5.1 Equilibrium Social Surplus

For $\pi \leqslant \underline{\pi}$ no agent explores the risky arm, and thus the equilibrium social surplus is zero. On the intermediate region, i.e., $\pi \in \big(\underline{\pi}, \bar{\pi}_\infty^{\text{local}}\big)$, k_n number of individuals choose to explore where $k_n/n \to \kappa$ characterized in Proposition 4. On this region the average equilibrium social surplus is

$$\begin{aligned} \frac{u_{k_n,n}(\pi)}{n} &= (1-\delta)\,\frac{k_n}{n}\,(\pi-\alpha(1-\pi))+\delta\pi \\ &\quad -\delta\,\frac{k_n}{n}\,\pi(1-\beta)\mathsf{E}_{k_n-1}^{(n)}\big[(1-\beta)^M\big]-\delta\,\frac{n-k_n}{n}\,\pi\,\mathsf{E}_{k_n}^{(n)}\big[(1-\beta)^M\big]. \end{aligned} \tag{5.4}$$

Figure 6 shows the average equilibrium social surplus for a finite n and some intermediate $\pi \in (\underline{\pi}, \alpha/(1+\alpha))$, where the asymmetric equilibrium prevails. In a fixed equilibrium region (with k_n remaining constant), an increase in λ positively shifts the distribution of M in the sense of first-order stochastic dominance. Consequently, based on the above representation, it *increases* the equilibrium social surplus. In the next proposition, we prove that at all thresholds, where the economy undergoes an equilibrium regime change, the social surplus falls, thus affirming our intuition from the two-player case.

Proposition 6. *The equilibrium social surplus falls discontinuously at every λ where the economy undergoes an equilibrium regime change.*

Proof. Suppose initially at $\lambda = \lambda_0$, the common belief falls in the interval $(\pi_{k,n}, \pi_{k+1,n}]$, and thus there are $k+1$ agents exploring in the equilibrium. Since the belief cutoffs (i.e., $\pi_{k,n}$'s) are increasing in λ, there will be a point $\lambda_{k,n}(\pi) > \lambda_0$ at which $\pi = \pi_{k,n}$ and the prevailing equilibrium will exhibit k players exploring. Part (i) of Lemma 4 implies that the change in the equilibrium social surplus when $\lambda = \lambda_{k,n}(\pi)$ is $u_{k,n}(\pi) - u_{k+1,n}(\pi) = -\Delta u_k(\pi)$. Letting $\pi = \pi_{k,n}$ in expression (5.2) implies that at $p = \lambda_{k,n}(\pi)/n$, one has

$$u_{k,n}(\pi_{k,n}) - u_{k+1,n}(\pi_{k,n}) = \frac{-\alpha p\beta\delta(1-\delta)(1-p\beta)^{k-1}\Big((n-1)(1-p)\beta - k(1-p)\beta\Big)}{(1+\alpha)(1-\delta)+\delta\beta(1-p\beta)^k},$$

that is always negative. As a result, the equilibrium social surplus, when evaluated just above $\lambda_{k,n}(\pi)$, is smaller than the surplus just below this threshold. Consequently, we observe a discontinuous decline in equilibrium social surplus at $\lambda_{k,n}(\pi)$. $\square$

The graph in Figure 6 shows that the equilibrium social surplus is increasing in λ on each equilibrium region, and features discontinuous jumps at critical λ's supporting equilibrium regime change. The largest (and the first) one corresponds to the equilibrium regime change from the full exploration (i.e., $k_n = n$) to the intermediate region, that is when π drops below $\bar{\pi}^{\text{local}}$ as λ increases. This plot can be considered as a counterpart of Figure 3b, with the distinction that it exhibits multiple discontinuous jumps due to the presence of multiple equilibrium regime changes when $n > 2$.

We continue by studying the limit of the average equilibrium social surplus, namely

$$\bar{u}_\infty(\pi) := \lim_{n\to\infty} \frac{u_{k_n,n}(\pi)}{n}.$$

Observe that as $n \to \infty$, the fraction k_n/n converges to κ and

$$\lim_{n\to\infty} \mathsf{E}^{(n)}_{k_n-1}\left[(1-\beta)^M\right] = \lim_{n\to\infty} \mathsf{E}^{(n)}_{k_n}\left[(1-\beta)^M\right] = e^{-\lambda\beta\kappa(\pi)}.$$

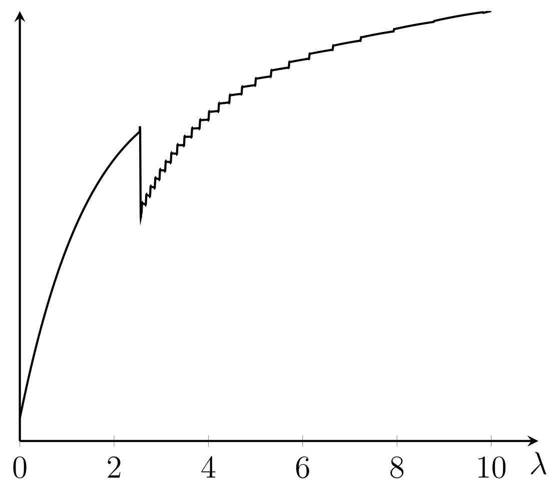

Figure 6: Finite-n equilibrium social surplus

Therefore, applying the expression in (5.4) results in:

$$\lim_{n\to\infty} \frac{u_{k_n,n}(\pi)}{n} = (1-\delta)\kappa(\pi)\left(\pi - \alpha(1-\pi)\right) + \delta\pi + \delta\pi e^{-\lambda\beta\kappa(\pi)}(\kappa(\pi)\beta - 1)\,,$$

which after replacing $\kappa(\pi)$ from (4.2) simplifies to

$$\bar{u}_\infty(\pi) = (1-\delta)\beta^{-1}\left(\pi - \alpha(1-\pi)\right) + \delta\pi\,,\ \text{ for every } \pi \in \left(\underline{\pi}, \bar{\pi}_\infty^{\text{local}}\right)\,.$$

Finally, for $\pi \geqslant \bar{\pi}_\infty^{\text{local}}$ all agents explore the risky arm, and it follows from (5.1) that

$$\begin{aligned}\bar{u}_\infty &= (1-\delta)\left(\pi - \alpha(1-\pi)\right) + \delta\pi\left(1 - (1-\beta)e^{-\lambda\beta}\right)\\ &\quad + \delta \mathsf{E}_{M\sim\mathsf{Pois}(\lambda)}\left[\left(\pi(1-\beta)^{M+1} - \alpha(1-\pi)\right)^+\right].\end{aligned} \tag{5.5}$$

It is worth mentioning that in the full exploration region (where $k_n = n$) $M^{(n)}_{k_n-1}$ converges weakly to $\mathsf{Poisson}(\lambda)$, and this underlies the final term in the above expression. The following proposition summarizes the results regarding $\bar{u}_\infty$ as a function of both π and λ. Additionally, Figure 7 illustrates the asymptotic average equilibrium social surplus as a function of λ for two different levels of initial belief π.

Proposition 7. *The large-n limit of the average equilibrium social surplus $\bar{u}_\infty$ is weakly increasing in λ for every fixed π. In addition,*

(i) for $\pi \leqslant \underline{\pi}$, $\bar{u}_\infty(\pi, \lambda) = 0$.

(ii) For every $\pi \in \left(\underline{\pi}, \frac{\alpha}{1+\alpha}\right)$, there exists a threshold $\lambda(\pi)$ such that $\bar{u}_\infty(\pi, \lambda)$ is constant in λ on $\lambda \geqslant \lambda(\pi)$, and it follows (5.5) on $\lambda < \lambda(\pi)$.

(iii) For all $\pi \geqslant \frac{\alpha}{1+\alpha}$, $\bar{u}_\infty$ follows equation (5.5).

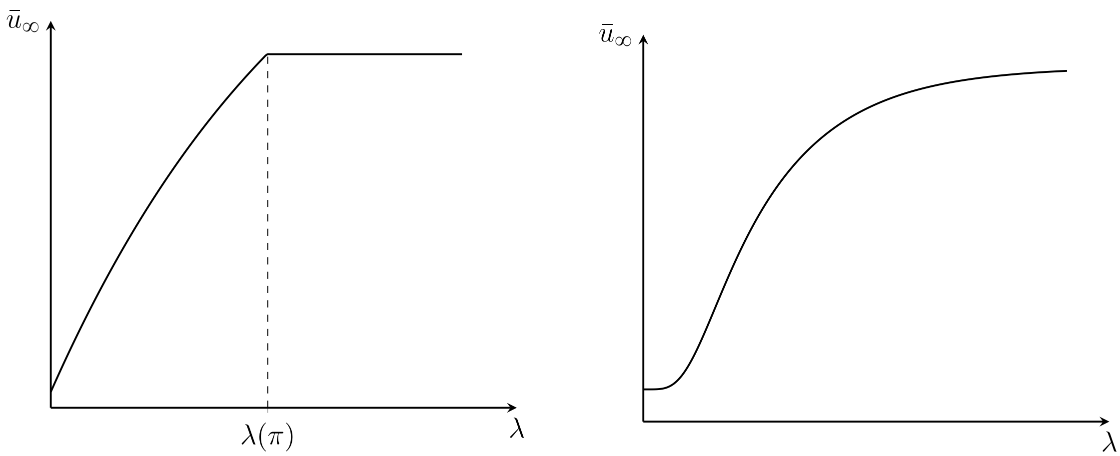

(a) $\underline{\pi} < \pi < \frac{\alpha}{1+\alpha}$

(b) $\pi \geqslant \frac{\alpha}{1+\alpha}$

Figure 7: Effect of λ on $\bar{u}_\infty$

Notably, for intermediate values of π, there exists a region where the limit of the average equilibrium social surplus remains unaffected by the average degree λ (see Panel 7a). That is as long as $\pi \in \left(\underline{\pi}, \bar{\pi}_\infty^{\text{local}}\right)$ and $\lambda \geqslant \lambda(\pi)$ (thus the prevailing equilibrium is asymmetric) the limiting equilibrium social surplus per-capita does not change by increasing or decreasing the connections. This is so because an increase in λ leads to more free riding and consequently fewer exploring agents in the equilibrium, which in turn reduces the social cost of first period exploration. On the other hand, fewer explorers correspond to smaller benefits from second period information exchange among the agents. In the large-n limit, these two effects precisely offset each other, leaving the per-capita equilibrium social surplus unaffected by λ. In particular, this is the region where for finite n, the equilibrium social surplus exhibits bounded jumps due to the regime changes in the equilibrium number of explorers (see Figure 6). In the limit $n \to \infty$, the discontinuous jumps in the per-capita equilibrium social surplus disappear, and it becomes constant with respect to λ (as plotted in Figure 7a).

5.2 Social Optimum

One should expect that the behavior of the social optimum mirrors the pattern in the two-player case. That is exploration (respectively, exploitation) becomes the social optimum when the initial common belief is larger (respectively, smaller) than some threshold. However, the justification of this result in the multi-agent economy follows after a long line of analysis.

Firstly, we need to determine when the marginal impact of adding one more exploring agent is positive, which involves examining $\Delta u_k = u_{k+1} - u_k$. Lemma 4 decomposed Δu_k into two components A and B. The former captures all terms including α and the latter accounts for the β-effect. The following lemma is the cornerstone of the social optimum analysis and its proof largely relies on the first-order stochastic dominance relation for Binomial distributions asserting that $\mathsf{Bin}(k+1, p) \succeq \mathsf{Bin}(k, p)$.

Lemma 5. *For every $k \geqslant r \geqslant 0$, it holds that*

$$B_k \geqslant \max\left\{(1-\beta)^r A_k, (1-\beta)^{r+1} A_k\right\}. \tag{5.6}$$

The previous lemma gives us a tight grip for Δu_k on $[\alpha/(1+\alpha), 1]$. For $\pi \leqslant \alpha/(1+\alpha)$ we need an additional result.

Lemma 6. *For every fixed $\pi \leqslant \alpha/(1+\alpha)$, the marginal value $\Delta u_k(\pi)$ is decreasing in k.*

This result is an immediate consequence of part (i) of Lemma 4. It stops short of asserting diminishing returns for the social surplus with respect to the number of exploring agents and instead only makes this claim for the region where the initial belief is small. However, together with the Lemmas 4 and 5, they characterize the regions where full exploitation and exploration are socially optimal.

Theorem 4 (Social optimum)**.** *The socially optimal outcome is full exploitation if and only if $\pi \leqslant \underline{\pi}^*$, and full exploration if and only if $\pi \geqslant \bar{\pi}^*$. Moreover, on $[0, \underline{\pi}^*]$ the social surplus is decreasing in k ($\Delta u_k \leqslant 0$), and on $[\bar{\pi}^*, 1]$ it is increasing in k ($\Delta u_k \geqslant 0$). The cutoff points are respectively:*

$$\underline{\pi}^* = \frac{\alpha(1-\delta)}{(1+\alpha)(1-\delta)+\delta\beta+\delta(n-1)p\beta},$$

$$\bar{\pi}^* = \frac{\alpha(1-\delta)}{(1+\alpha)(1-\delta)+\delta\beta(1-p\beta)^{n-2}\big(np(1-\beta)+1-p\big)}.$$

Recall that $p = \lambda/n$. Thus, one can find the limit of the lower (respectively, upper) cutoff point for the optimality of full exploitation (respectively, full exploration) as $n \to \infty$:

$$\underline{\pi}^*_\infty := \lim_{n\to\infty} \underline{\pi}^* = \frac{\alpha(1-\delta)}{(1+\alpha)(1-\delta)+\delta\beta(\lambda+1)}\,,$$
$$\bar{\pi}^*_\infty := \lim_{n\to\infty} \bar{\pi}^* = \frac{\alpha(1-\delta)}{(1+\alpha)(1-\delta)+\delta\beta e^{-\lambda\beta}\big(\lambda(1-\beta)+1\big)}\,.$$

Recall that the equilibrium exploitation threshold in the multi-agent economy matches the one in the two-player case, namely $\underline{\pi}$ in Proposition 2. Comparing $\underline{\pi}^*_\infty$ from the above expression with $\underline{\pi}$ confirms the over-exploitation in the equilibrium relative to the social optimum. Additionally, comparing the above $\bar{\pi}^*_\infty$ with the asymptotic equilibrium exploration threshold in the local case, expressed in (4.1), indicates that there is under-exploration relative to the social optimum.

Effect of λ on the optimal exploration cutoff. The optimal exploration cutoff $\bar{\pi}^*_\infty$ initially decreases in λ and then increases. To better understand the reason behind this fall and the subsequent rise, we examine the marginal impact of the n-th exploring agent on the social surplus (i.e., Δu_{n-1}). Specifically, we examine its contribution to the positive externality of community exploration for an agent whose exploration failed in the first period. This is reflected in the marginal change of the last term in the second line of the surplus function (5.1), namely

$$\delta\pi(1-\beta)\left[n\sum_{m=0}^{n-1} q_{n-1}(m)\big(1-(1-\beta)^m\big)-(n-1)\sum_{m=0}^{n-2} q_{n-2}(m)\big(1-(1-\beta)^m\big)\right]. \tag{5.7}$$

We employ an intuitive coupling argument to further highlight the above marginal change and its response to λ. Suppose in the high state of the world an agent who had picked the risky arm failed in the first period, that happens with probability $\pi(1-\beta)$. Let $X \sim \mathsf{Bin}(n-2, \lambda/n)$ be the number of his contacts in the second period (excluding himself and the candidate n-th individual). Then, setting the base event probability $\pi(1-\beta)$ aside, the difference in the bracket in (5.7) is approximately equal to

$$n\Big(\mathsf{E}_{X,Z}\left[1-(1-\beta)^{X+Z}\right]-\mathsf{E}_X\left[1-(1-\beta)^X\right]\Big)\,,$$

where Z is a $\mathsf{Bernoulli}(\lambda/n)$ random variable representing the exploration outcome of the n-th agent in the first period. The above expression is further simplified to

$$n\ \mathsf{E}_Z\left[1-(1-\beta)^Z\right]\ \mathsf{E}_X\left[(1-\beta)^X\right] = n\ \frac{\lambda}{n}\beta\ (1-\lambda\beta/n)^{n-2} \to \lambda\beta e^{-\lambda\beta}\,.$$

This representation tells us that the positive externality of the n-th agent's exploration is proportional to the average number of her meetings with her immediate neighbors when she had experienced a success, i.e., $\lambda\beta$, and the expected probability of group failure among the remaining $n-2$ exploring agents, i.e., $(1-\lambda\beta/n)^{n-2}$. Therefore, for any fixed $\pi \leqslant \alpha/(1+\alpha)$, the marginal impact of the exploration of the n-th individual (namely, Δu_{n-1}) is initially increasing in λ and then decreasing. This translates to an opposite response for the full exploration optimal cutoff. Figure 8a depicts the large-n limits of the equilibrium and optimum exploration cutoffs as a function of λ in the local economy.

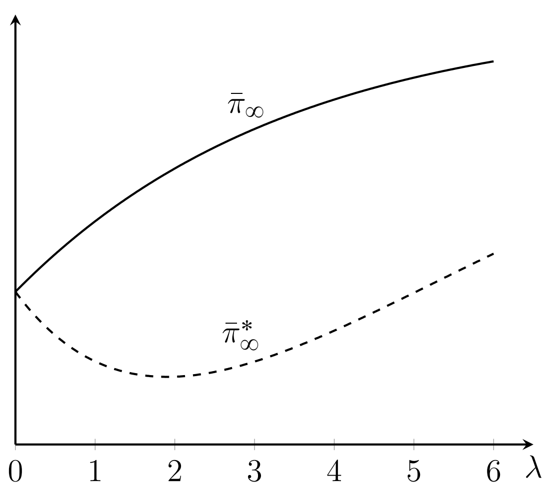

(a) Effect of λ
[$\beta = 0.3$]

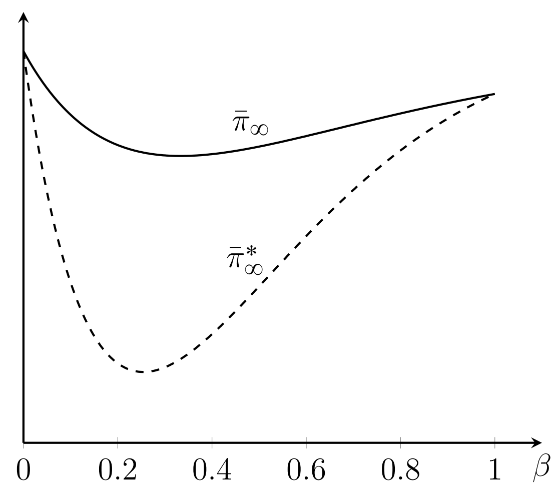

(b) Effect of β
[$\lambda = 3$]

Figure 8: Equilibrium and optimum full exploration threshold
[$\delta = 0.15, \alpha = 1$]

In Figure 8b, we plotted the equilibrium and optimum exploration thresholds as a function of β. Both graphs highlight the idea that higher levels of uncertainty about the risky arm (meaning intermediate values of β) are associated with more exploration. However, this effect is relatively dampened in the equilibrium compared to the optimum.

6 Conclusion and Additional Discussion

The tension between information diffusion and production in organizations and societies represents a complex and evolving challenge. On one hand, greater connectivity has allowed for unprecedented access to knowledge and ideas fostering informed decision-making and quicker adoption of innovation. At the same time, this proliferation of information can undermine knowledge production. In a better connected organization or society, individuals are tempted to rely on knowledge that is shared through the network instead of experimenting with new ideas. Greater connectivity and knowledge diffusion can thus lead to lower knowledge production, reducing overall social welfare.

Our analysis begins with a two-player economy, where equilibrium reveals three distinct regions based on initial beliefs, each with varying degrees of knowledge exploitation and exploration. Free riding leads to over-exploitation and under-exploration relative to the social optimum, and equilibrium social surplus exhibits non-monotonic behavior concerning connection probabilities.

We then analyze multi-agent economies and explore different network structures. With local connections each agent only observes the experimentation outcomes of her immediate neighbors, whereas with global connections each agent's observable circle includes the entire set of agents who are (directly or indirectly) connected to her. In both structures, the tension between information sharing and private exploration remains a significant factor. In particular, higher connectivity can exacerbate free riding and reduce social welfare. Additionally, we discuss the asymptotic effects of connectivity on equilibrium and highlight how the size of the connected component in the network significantly impacts exploration behavior.

Although we showed that the tension between information diffusion and production is prevalent in different settings, interesting variations remain to be studied. For example, what happens when agents have different preferences or endowments? What happens when societies and organizations have network structures with differently connected agents? We hope to address these issues in future work.

A Proofs

A.1 Proof of Theorem 1

If $\alpha(1-\pi)/\pi \leqslant 1$, define $\bar{m} := \max\left\{0 \leqslant m \leqslant n : (1-\beta)^m \geqslant \frac{\alpha(1-\pi)}{\pi}\right\}$, otherwise let $\bar{m} = 0$. Then, after few steps of algebraic manipulations, the condition for $v_n(\pi) > w_{n-1}(\pi)$ laid out in (3.3) reduces to

$$\begin{aligned} &(1-\delta)\left(\pi-\alpha(1-\pi)\right)+\delta\pi\beta\sum_{m=0}^{n-1} q(m)(1-\beta)^m \\ &> \delta\pi\beta\sum_{0\leqslant m<\bar{m}} q(m)(1-\beta)^m+\delta q_{\bar{m}}\left[\pi(1-\beta)^{\bar{m}}-\alpha(1-\pi)\right]^+ , \end{aligned} \tag{A.1}$$

with the interpretation of each component given in (3.3). If $\bar{m} = 0$ the *rhs* in the above inequality is zero, which will be the case when $\pi/(1-\pi) \leqslant \alpha$. In this case, equation (A.1) is equivalent to $\pi > \bar{\pi}$, which as it will turn out is the only restricting condition for the existence of the exploration equilibrium.

Next, we show for every $\bar{m} > 0$ and for every

$$\frac{\pi}{1-\pi} \in \left[\frac{\alpha}{(1-\beta)^{\bar{m}}}, \frac{\alpha}{(1-\beta)^{\bar{m}+1}}\right], \tag{A.2}$$

equation (A.1) holds. On the above region, (A.1) is equivalent to

$$\pi\Big\{1-\delta-\delta(1-\beta)^{\bar{m}}q_{\bar{m}}+\delta\beta\mathsf{E}\left[(1-\beta)^M; M \geqslant \bar{m}\right]\Big\} > \alpha(1-\pi)\left(1-\delta-\delta q_{\bar{m}}\right) . \tag{A.3}$$

If the coefficient of π in (A.3) is positive, then this inequality becomes equivalent to

$$\frac{\pi}{1-\pi} > \frac{\alpha\left(1-\delta-\delta q_{\bar{m}}\right)}{1-\delta-\delta(1-\beta)^{\bar{m}}q_{\bar{m}}+\delta\beta\mathsf{E}\left[(1-\beta)^M; M \geqslant \bar{m}\right]},$$

which always holds on (A.2) because $\frac{\alpha}{(1-\beta)^{\bar{m}}}$ is greater than the *rhs* above. Alternatively, if the coefficient of π in (A.3) is negative, then it becomes equivalent to

$$\frac{\pi}{1-\pi} < \frac{\alpha\left(1-\delta-\delta q_{\bar{m}}\right)}{1-\delta-\delta(1-\beta)^{\bar{m}}q_{\bar{m}}+\delta\beta\mathsf{E}\left[(1-\beta)^M; M \geqslant \bar{m}\right]},$$

which again always holds on (A.2), because it can be readily shown that the *rhs* above is smaller than $\frac{\alpha}{(1-\beta)^{\bar{m}+1}}$. Therefore, the exploration equilibrium appears on every region of type (A.2), and the only constraint restricting the existence of such equilibrium appears on the region $\pi \in \left[0, \frac{\alpha}{1+\alpha}\right]$, which is nothing but $\pi > \bar{\pi}$. $\square$

A.2 Proof of Lemma 1

Let us look at the difference

$$v_{k+1}(\pi) - w_k(\pi) = (1-\delta)\big(\pi - \alpha(1-\pi)\big) + \delta\pi\beta \mathsf{E}_k\left[(1-\beta)^M\right]$$
$$+\delta \mathsf{E}_k\left[\left(\pi(1-\beta)^{M+1} - \alpha(1-\pi)\right)^+\right] - \delta \mathsf{E}_k\left[\left(\pi(1-\beta)^M - \alpha(1-\pi)\right)^+\right].$$

Since $\pi/(1-\pi) > \alpha$, then $\bar{m} := \max\left\{0 \leqslant m \leqslant n : (1-\beta)^m \geqslant \frac{\alpha(1-\pi)}{\pi}\right\}$ exists and $\bar{m} \geqslant 0$. Thus the above difference can be reduced to

$$v_{k+1}(\pi) - w_k(\pi) = (1-\delta)\big(\pi - \alpha(1-\pi)\big) + \delta\alpha(1-\pi) q_k(\bar{m}) + \delta\pi\beta \mathsf{E}_k\left[(1-\beta)^M; M \geqslant \bar{m}\right],$$

which is always positive. $\square$

A.3 Proof of Theorem 3

We need the next lemma to prove the proposition.

Lemma A.1. *In exchangeable random graphs, the mapping* $k \mapsto \Gamma_k := \mathsf{E}_{k-1}\left[(1-\beta)^{M+1}\right] - \mathsf{E}_k\left[(1-\beta)^M\right]$ *is increasing.*

Proof. Let us pick a vertex i uniformly at random and label the other vertices by $j \in \{1, \dots, n-1\}$. Let X_j be the indicator random variable which is one when i is connected via a path to j. Then, when k agents are exploring $M \overset{d}{=} X_1 + \cdots + X_k$. Using this coupling approach one can express the increments of Γ as

$$\Gamma_{k+1} - \Gamma_k =$$
$$\mathsf{E}\left[(1-\beta)^{X_1+\ldots+X_{k-1}} \mathsf{E}\left[(1-\beta)^{X_k+1} - (1-\beta)^{X_k+X_{k+1}} - (1-\beta) + (1-\beta)^{X_k} \middle| X_1^{k-1}\right]\right],$$

where we use the notation $X_1^{k-1} := \{X_1, \dots, X_{k-1}\}$. The inner expectation above equals

$$\beta\left[\mathsf{P}\left(X_k = 0, X_{k+1} = 1 \middle| X_1^{k-1}\right) - (1-\beta)\mathsf{P}\left(X_k = 1, X_{k+1} = 0 \middle| X_1^{k-1}\right)\right].$$

The two probabilities above are equal to each other because of exchangeability. Hence, Γ has positive increments, and is therefore increasing in k. $\square$

To prove the proposition, we first demonstrate that in the intermediate region there exists a unique μ satisfying $v(\pi;\mu) = w(\pi;\mu)$. Observe that

$$w(\pi;\mu) = \delta\pi \sum_{k=0}^{n-1} \binom{n-1}{k} \mu^k (1-\mu)^{n-1-k} \mathsf{E}_k \left[1-(1-\beta)^M\right] ,$$

$$v(\pi;\mu) = (1-\delta)\left(\pi-\alpha(1-\pi)\right)+\delta\pi\beta$$

$$+\delta\pi(1-\beta) \sum_{k=0}^{n-1} \binom{n-1}{k} \mu^k (1-\mu)^{n-k} \mathsf{E}_{k-1} \left[1-(1-\beta)^M\right] .$$

Hence $v(\pi;\mu) = w(\pi;\mu)$ is equivalent to

$$\frac{(1-\delta)\left(\pi-\alpha(1-\pi)\right)}{\delta\pi} = \sum_{k=0}^{n-1} \binom{n-1}{k} \mu^k (1-\mu)^{n-1-k} \Gamma_k = \mathsf{E}_{k\sim\mathsf{Bin}(n-1,\mu)} \Gamma_k . \tag{A.4}$$

Since in the intermediate region $v(\pi;0) = v_1(\pi) > w(\pi) = w(\pi;0)$ and $v(\pi;1) = v_n(\pi) \leqslant w_{n-1}(\pi) = w(\pi;1)$, then there exists $\mu^* \in (0,1]$ satisfying (A.4). In addition, $\mathsf{Bin}(n-1,\mu)$ increases in the sense of first-order stochastic dominance with respect to μ. By previous lemma, Γ is increasing in k, therefore, the *rhs* of (A.4) becomes increasing in μ, and this establishes the uniqueness of μ^*.

Next, since the *lhs* of (A.4) is increasing π, and its *rhs* is increasing in μ, then μ^* must be increasing in π. Lastly, observe that for a fixed μ, one has $\mathsf{Bin}(n,\mu) \geq \mathsf{Bin}(n-1,\mu)$ in the sense of first-order stochastic dominance. Also, the function Γ_k is increasing by Lemma A.1. Therefore, the *rhs* of (A.4) is increasing in n (as well as μ). Since the *lhs* of (A.4) is unaffected by n, then μ^* will become decreasing in n. $\square$

A.4 Proof of Lemma 3

Since $p_n = \frac{\lambda}{n} \geqslant \frac{\lambda}{n+1} = p_{n+1}$, then a coupling argument shows that on a same probability space $M_k^{(n)} \geqslant M_k^{(n+1)}$, and therefore $\mathsf{E}\left[(1-\beta)^{M_k^{(n)}}\right] \leqslant \mathsf{E}\left[(1-\beta)^{M_k^{(n+1)}}\right]$. That in turn implies

$\pi_{k,n+1} \leqslant \pi_{k,n}$. Next, observe that with local connections,

$$\mathsf{E}_{k+1}^{(n+1)}\left[(1-\beta)^M\right] = \left(1-\frac{\lambda\beta}{n+1}\right)^{k+1} \leqslant \left(1-\frac{\lambda\beta}{n+1}\right)\left(1+\frac{\lambda\beta}{n+1}\right)^{-k}$$
$$\leqslant \left(1-\frac{\lambda\beta}{n+1}\right)\left(1+\frac{k\lambda\beta}{n+1}\right)^{-1}.$$

Moreover,

$$\mathsf{E}_{k}^{(n)}\left[(1-\beta)^M\right] = \left(1-\frac{\lambda\beta}{n}\right)^{k} \geqslant 1-\frac{k\lambda\beta}{n}.$$

For large n, one can readily show that

$$\left(1-\frac{\lambda\beta}{n+1}\right) \leqslant \left(1-\frac{k\lambda\beta}{n}\right)\left(1+\frac{k\lambda\beta}{n+1}\right).$$

therefore,

$$\mathsf{E}_{k+1}^{(n+1)}\left[(1-\beta)^M\right] \leqslant \mathsf{E}_{k}^{(n)}\left[(1-\beta)^M\right].$$

Hence $\pi_{k+1,n+1} \geqslant \pi_{k,n}$ for large enough n. Similarly, one can show $\pi_{k,n+1} \geqslant \pi_{k-1,n}$, thereby concluding the proof. □

A.5 Proof of Proposition 4

For every $\pi \leqslant \underline{\pi}$, the full exploitation equilibrium prevails, thus $k_n(\pi) = 0$. Also, for every $\pi \geqslant \bar{\pi}_\infty^{\text{local}}$, due to Lemma 2, it follows that $\pi > \bar{\pi}_n$ for large enough n, hence $k_n(\pi) = n$. Therefore, it remains to examine the limiting behavior of $k_n(\pi)/n$ on the intermediate region $(\underline{\pi}, \bar{\pi}_\infty^{\text{local}})$, where asymmetric equilibria prevail. According to equation (3.6) there will be k_n agents exploring in the equilibrium if and only if

$$\mathsf{E}_{k_n-1}^{(n)}\left[(1-\beta)^M\right] < \frac{(1-\delta)\left(\alpha(1-\pi)-\pi\right)}{\delta\beta} \leqslant \mathsf{E}_{k_n}^{(n)}\left[(1-\beta)^M\right]$$
$$\Leftrightarrow (k_n-1) < \frac{\log\left((1-\delta)\left(\alpha(1-\pi)-\pi\right)/\delta\pi\beta\right)}{\log\left(1-\lambda\beta/n\right)} \leqslant k_n.$$

Therefore,

$$\lim_{n\to\infty}\frac{k_n}{n} = \frac{\log\left((1-\delta)\left(\alpha(1-\pi)-\pi\right)/\delta\pi\beta\right)}{\lim_{n\to\infty} n\log\left(1-\lambda\beta/n\right)} = \frac{1}{\lambda\beta}\log\frac{\delta\pi\beta}{(1-\delta)\left(\alpha(1-\pi)-\pi\right)}.$$

□

A.6 Proof of Proposition 5

First, we show how the size of the connected component $|\mathcal{C}|$ in a random Erdos-Renyi graph with parameters $(n, p = \lambda/n)$ can be approximated with the number of the descendants in a Branching process with $\mathsf{Bin}(n,p)$ offspring distribution, which we denote it by B. We use $\mathsf{P}_{n,p}$ to refer to the distribution of B. Theorem 4.2 and 4.3 of Van Der Hofstad (2016) jointly state that:

$$\mathsf{P}_{n-k,p}\left(B \geqslant k\right) \leqslant \mathsf{P}\left(|\mathcal{C}| \geqslant k\right) \leqslant \mathsf{P}_{n,p}\left(B \geqslant k\right) .$$

Next, we examine how the total number of progenies in a Binomial Branching process with parameters (n,p) can be approximated by the Branching process with a $\mathsf{Poisson}(np)$ offspring distribution, that we denote it by T. We use P_λ to refer to the distribution of the Branching process with $\mathsf{Poisson}(\lambda)$ offspring distribution. Let $\lambda = np$ and fix $k \in \mathbb{N}$. Then, Theorem 3.20 in Van Der Hofstad (2016) implies that

$$\left|\mathsf{P}_{n,p}(B \geqslant k) - \mathsf{P}_\lambda(T \geqslant k)\right| \leqslant \frac{\lambda^2 k}{n} .$$

Subsequently, the last two relations result in

$$\mathsf{P}_{\lambda(1-kn^{-1})}(T \geqslant \ell) - \frac{\lambda^2 (n-k)\ell}{n^2} \leqslant \mathsf{P}\left(|\mathcal{C}| \geqslant k\right) \leqslant \mathsf{P}_\lambda(T \geqslant k) + \frac{\lambda^2 k}{n} .$$

For a fixed $k \in \mathbb{N}$, let $\lambda_n := \lambda(1 - kn^{-1})$. Then, using the method of characteristic functions, one can show P_{λ_n} weakly converges to P_λ as $n \to \infty$ (see Theorem 5.3 in Kallenberg (2002)). This in turn means, $\mathsf{P}_{\lambda_n}(T < k) \to \mathsf{P}_\lambda(T < k)$, and hence $\mathsf{P}_{\lambda_n}(T \geqslant k) \to \mathsf{P}_\lambda(T \geqslant k)$. Using this and the above inequality one reaches the conclusion that for every $k \in \mathbb{N}$,

$$\lim_{n \to \infty} \mathsf{P}(|\mathcal{C}| \geqslant k) = \mathsf{P}_\lambda(T \geqslant k) . \tag{A.5}$$

Let $\bar{\mathbb{N}} = \mathbb{N} \cup \{\infty\}$, then $|\mathcal{C}|$ and T are $\bar{\mathbb{N}}$-valued random variables, where $\bar{\mathbb{N}}$ is a discrete metric space. Therefore, the limiting result in (A.5) implies the weak convergence of $|\mathcal{C}|$ to T. Lastly, the distribution of the descendants of a Poisson Branching process is known to follow the Borel distribution (see Theorem 3.16 of Van Der Hofstad (2016)). This concludes the justification of part (i) of the proposition. Part (ii) quickly follows because every function on $\bar{\mathbb{N}}$ is continuous. In particular, $x \mapsto (1-\beta)^{x-1}$ is bounded and continuous, therefore because of the weak convergence established in the previous part one has:

$$\lim_{n\to\infty} \mathsf{E}\left[(1-\beta)^{|\mathcal{C}|-1}\right] = \mathsf{E}\left[(1-\beta)^{T-1}\right] ,$$

justifying equation (4.3). $\square$

A.7 Proof of Lemma 5

We separately show B_k is larger than both of the arguments of the max operator. First, observe that $B_k \geqslant (1-\beta)^r A_k$ if and only if

$$\begin{aligned}(1-\delta)\big(1-(1-\beta)^r\big) \geqslant{} & -\delta k\Big[Q_{k-1}(r-1)(1-\beta)^r+\sum_{m=r}^{k-1} q_{k-1}(m)(1-\beta)^{m+1}\Big] \\ & -\delta(n-k)\Big[Q_k(r)(1-\beta)^r+\sum_{m=r+1}^{k} q_k(m)(1-\beta)^m\Big] \\ & +\delta(k+1)\Big[Q_k(r-1)(1-\beta)^r+\sum_{m=r}^{k} q_k(m)(1-\beta)^{m+1}\Big] \\ & +\delta(n-k-1)\Big[Q_{k+1}(r)(1-\beta)^r+\sum_{m=r+1}^{k+1} q_{k+1}(m)(1-\beta)^m\Big] .\end{aligned} \tag{A.6}$$

The *lhs* of the above inequality is nonnegative. Therefore, to establish that $B_k \geqslant (1-\beta)^r A_k$, it suffices to show that the following equivalent representation for the *rhs* is negative. In that, we use the notation E_k to express the expectation with respect to the Binomial distribution $\mathsf{Bin}(k,p)$, and the random variable M follows the corresponding distribution in the subscript of E.[9]

$$\begin{aligned}\textit{rhs} \text{ of (A.6)} ={} & -\delta k\mathsf{E}_{k-1}\left[(1-\beta)^{(M+1)\vee r}\right]+\delta(k+1)\mathsf{E}_k\left[(1-\beta)^{(M+1)\vee r}\right] \\ & -\delta(n-k)\mathsf{E}_k\left[(1-\beta)^{M\vee r}\right]+\delta(n-k-1)\mathsf{E}_{k+1}\left[(1-\beta)^{M\vee r}\right] .\end{aligned}$$

Note that each of the functions inside the expectation operators is decreasing in M, therefore, using the first-order stochastic dominance for the first and second lines, respectively $\mathsf{Bin}(k,p) \succeq \mathsf{Bin}(k-1,p)$ and $\mathsf{Bin}(k+1,p) \succeq \mathsf{Bin}(k,p)$, results in the following upper bound:

$$\begin{aligned}\textit{rhs} \text{ of (A.6)} &\leqslant \delta\mathsf{E}_k\left[(1-\beta)^{(M+1)\vee r}\right]-\delta\mathsf{E}_{k+1}\left[(1-\beta)^{M\vee r}\right] \\ &= \delta\mathsf{E}_k\left[(1-\beta)^{(M+1)\vee r}\right]-\delta p\mathsf{E}_k\left[(1-\beta)^{(M+1)\vee r}\right]-\delta(1-p)\mathsf{E}_k\left[(1-\beta)^{M\vee r}\right] \\ &= \delta(1-p)\mathsf{E}_k\left[(1-\beta)^{(M+1)\vee r}-(1-\beta)^{M\vee r}\right] \leqslant 0 .\end{aligned}$$

[9] This means the distribution of M varies across terms.

For the second part of the inequality, namely $B_k \geqslant (1-\beta)^{r+1}A_k$, we arrive at the following equivalent condition:

$$
\begin{aligned}
(1-\delta)\big(1-(1-\beta)^{r+1}\big) \geqslant \delta(1-\beta)\Bigg\{&-k\Big[Q_{k-1}(r-1)(1-\beta)^r+\sum_{m=r}^{k-1} q_{k-1}(m)(1-\beta)^m\Big] \\
&-(n-k)\Big[Q_k(r)(1-\beta)^r+\sum_{m=r+1}^{k} q_k(m)(1-\beta)^{m-1}\Big] \\
&+(k+1)\Big[Q_k(r-1)(1-\beta)^r+\sum_{m=r}^{k} q_k(m)(1-\beta)^m\Big] \\
&+(n-k-1)\Big[Q_{k+1}(r)(1-\beta)^r+\sum_{m=r+1}^{k+1} q_{k+1}(m)(1-\beta)^{m-1}\Big]\Bigg\}.
\end{aligned}
\tag{A.7}
$$

The *lhs* in (A.7) is nonneagtive, thus it is enough to show that the *rhs* is negative to justify $B_k \geqslant (1-\beta)^{r+1}A_k$. Toward that, we appeal to the following equivalent representation:

$$
\begin{aligned}
\textit{rhs} \text{ of (A.7)} = \delta(1-\beta)\Big\{&-k\mathsf{E}_{k-1}\left[(1-\beta)^{M\vee r}\right]+(k+1)\mathsf{E}_k\left[(1-\beta)^{M\vee r}\right] \\
&-(n-k)\mathsf{E}_k\left[(1-\beta)^{(M-1)\vee r}\right]+(n-k-1)\mathsf{E}_{k+1}\left[(1-\beta)^{(M-1)\vee r}\right]\Big\}.
\end{aligned}
$$

Using the first-order stochastic dominance once again yields the following upper bound:

$$
\begin{aligned}
\textit{rhs} \text{ of (A.7)} &\leqslant \delta(1-\beta)\Big(\mathsf{E}_k\left[(1-\beta)^{M\vee r}\right]-\mathsf{E}_{k+1}\left[(1-\beta)^{(M-1)\vee r}\right]\Big) \\
&= \delta(1-\beta)\Big(\mathsf{E}_k\left[(1-\beta)^{M\vee r}\right]-p\mathsf{E}_k\left[(1-\beta)^{M\vee r}\right]-(1-p)\mathsf{E}_k\left[(1-\beta)^{(M-1)\vee r}\right]\Big) \\
&= \delta(1-\beta)(1-p)\mathsf{E}_k\left[(1-\beta)^{M\vee r}-(1-\beta)^{(M-1)\vee r}\right] \leqslant 0\,.
\end{aligned}
$$

Therefore, both inequalities were proved, and thus the claim (5.6) in the lemma is established.

□

A.8 Proof of Theorem 4

First, we justify the lower cutoff rule for the optimality of full exploitation. Part (i) of Lemma 4 implies that $\Delta u_0 \leqslant 0$ on $\pi \leqslant \underline{\pi}^*$ and $\Delta u_0 > 0$ on $\underline{\pi}^* < \pi \leqslant \alpha/(1+\alpha)$. In addition, since Δu_k is decreasing on $[0, \alpha/(1+\alpha)]$ (because of Lemma 6), then $\Delta u_k \leqslant \Delta u_0 \leqslant 0$ on $\pi \leqslant \underline{\pi}^*$, implying the optimality of full exploitation on this region. When $\pi/(1-\pi) \in [\alpha, \alpha/(1-\beta)]$, part (iii) in Lemma 4 (with $r = 0$) states that $\Delta u_0(\pi) = \pi B_0 - \alpha(1-\pi)A_0$. If $A_0 \geqslant 0$, then

$B_0 \geqslant 0$ due to Lemma 5 and hence

$$\Delta u_0(\pi) = \pi B_0 - \alpha(1-\pi)A_0 \geqslant \big(\pi - \alpha(1-\pi)\big)A_0 \geqslant 0\,.$$

Alternatively, if $A_0 < 0$, then again because of Lemma 5, it holds that

$$\Delta u_0(\pi) = \pi B_0 - \alpha(1-\pi)A_0 \geqslant \big(\pi(1-\beta) - \alpha(1-\pi)\big)A_0 \geqslant 0\,.$$

Lastly, when $\pi/(1-\pi) > \alpha/(1-\beta)$ part (ii) of Lemma 4 implies $\Delta u_0(\pi) > 0$. We can now conclude that $u_1(\pi) > u_0(\pi)$ for all $\pi > \underline{\pi}^*$, and therefore full exploitation becomes optimal if and only if $\pi \leqslant \underline{\pi}^*$.

Next, we establish the optimality of full exploration above $\bar{\pi}^*$. On the region where $\pi/(1-\pi) \leqslant \alpha$, part (i) of Lemma 4 shows that $\Delta u_{n-1}(\pi) < 0$ on $\pi < \bar{\pi}^*$ and $\Delta u_{n-1}(\pi) \geqslant 0$ on $[\bar{\pi}^*, \alpha/(1+\alpha)]$. Also, Lemma 6 results in $\Delta u_k(\pi) \geqslant \Delta u_{n-1}(\pi) \geqslant 0$ for $\pi \in [\bar{\pi}^*, \alpha/(1+\alpha)]$. Hence, it follows that full exploration is the social optimum if and only if $\pi \geqslant \bar{\pi}^*$.

Thus it remains to show that the social surplus is increasing in k, i.e., $\Delta u_k(\pi) \geqslant 0$, for every $\pi > \alpha/(1+\alpha)$. For every $\pi > \alpha/(1+\alpha)$, there exists r such that $\alpha/(1-\beta)^r \leqslant \pi/(1-\pi) \leqslant \alpha/(1-\beta)^{r+1}$. If $k < r$, then part (ii) states that Δu_k is positive. Alternatively, suppose $k \geqslant r$. Then, if $A_k \geqslant 0$, from Lemma 5 it follows that

$$\Delta u_k(\pi) = \pi B_k - \alpha(1-\pi)A_k \geqslant \big(\pi(1-\beta)^r - \alpha(1-\pi)\big)A_k \geqslant 0\,,$$

and if $A_k < 0$, then one obtains

$$\Delta u_k(\pi) = \pi B_k - \alpha(1-\pi)A_k \geqslant \big(\pi(1-\beta)^{r+1} - \alpha(1-\pi)\big)A_k \geqslant 0\,.$$

This justifies that u_k is increasing on $[\bar{\pi}^*, 1]$, and hence concludes the proof. □

References

[1] Noga Alon and Joel H. Spencer (2000). *The Probabilistic Method*, Wiley, 2nd edition.

[2] Masaki Aoyagi (1998). “Mutual Observability and the Convergence of Actions in a Multi-Person Two-Armed Bandit Model,” *Journal of Economic Theory*, 82(2): 405–424.

[3] Venkatesh Bala and Sanjeev Goyal (1998). "Learning from Neighbours," *The Review of Economic Studies*, 65(3): 595–621.

[4] Simon Board and Moritz Meyer-ter Vehn (2023). "Experimentation in Networks," *Working Paper*.

[5] Béla Bollobás (2001). *Random Graphs*, Cambridge University Press.

[6] Patrick Bolton and Christopher Harris (1999). "Strategic Experimentation," *Econometrica*, 67(2): 349–374.

[7] Alessandro Bonatti and Johannes Hörner (2017). "Learning to Disagree in a Game of Experimentation," *Journal of Economic Theory*, 169: 234–269.

[8] Braz Camargo (2014). "Learning in Society," *Games and Economic Behavior*, 87: 381–396.

[9] Christophe Chamley and Douglas Gale (1994). "Information Revelation and Strategic Delay in a Model of Investment," *Econometrica*, 62(5): 1065–1085.

[10] Robert M. Corless, Gaston H. Gonnet, David EG. Hare, David J. Jeffrey, and Donald E. Knuth (1996). "On the Lambert W function," *Advances in Computational Mathematics*, 5(1): 329–359.

[11] Darrell Duffie, Semyon Malamud, and Gustavo Manso (2009). "Information Percolation with Equilibrium Search Dynamics," *Econometrica*, 77(5): 1513–1574.

[12] Mira Frick and Yuhta Ishii (2020). "Innovation Adoption by Forward-Looking Social Learners," *Working Paper*.

[13] Douglas Gale and Shachar Kariv (2003). "Bayesian Learning in Social Networks," *Games and Economic Behavior*, 45(2): 329–346.

[14] Sidartha Gordon, Chantal Marlats, and Lucie Ménager (2021). "Observation Delays in Teams and Effort Cycles," *Games and Economic Behavior*, 130: 276–298.

[15] Paul Heidhues, Sven Rady, and Philipp Strack (2015). "Strategic Experimentation with Private Payoffs," *Journal of Economic Theory*, 159: 531–551.

[16] Olav Kallenberg (2002). *Foundations of Modern Probability*, Springer, 2nd edition.

[17] Godfrey Keller and Sven Rady (2015). "Breakdowns," *Theoretical Economics*, 10(1): 175–202.

[18] Godfrey Keller, Sven Rady, and Martin Cripps (2005). "Strategic Experimentation with Exponential Bandits," *Econometrica*, 73(1): 39–68.

[19] William R. Kerr, Ramana Nanda, and Matthew Rhodes-Kropf (2014). "Entrepreneurship as Experimentation," *Journal of Economic Perspectives*, 28(3): 25–48.

[20] Gustavo Manso (2011). "Motivating Innovation," *The Journal of Finance*, 66(5): 1823–1860.

[21] Farzad Pourbabaee (2020). "Robust Experimentation in the Continuous Time Bandit Problem," *Economic Theory*: 1–31.

[22] Dinah Rosenberg, Eilon Solan, and Nicolas Vieille (2007). "Social Learning in One-Arm Bandit Problems," *Econometrica*, 75(6): 1591–1611.

[23] ——— (2009). "Informational Externalities and Emergence of Consensus," *Games and Economic Behavior*, 66(2): 979–994.

[24] Michael Rothschild (1974). "A Two-Armed Bandit Theory of Market Pricing," *Journal of Economic Theory*, 9(2): 185–202.

[25] Evan Sadler (2020). "Innovation Adoption and Collective Experimentation," *Games and Economic Behavior*, 120: 121–131.

[26] Remco Van Der Hofstad (2016). *Random Graphs and Complex Networks*, volume 1, Cambridge University Press.

[27] Alexander Wolitzky (2018). "Learning from Others' Outcomes," *American Economic Review*, 108(10): 2763–2801.