The Impact of Connectivity on the Production and Diffusion of Knowledge

Gustavo Manso † and Farzad Pourbabaee ‡

February 2022

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.

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 adoption of innovation.

While such enhanced knowledge diffusion is in principle 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

We thank Robert M. Anderson, Omer Tamuz and participants at Caltech CSIC seminar series for helpful feedback.

†Haas School of Business, University of California Berkeley. Email: manso@berkeley.edu.

‡California Institute of Technology. Email: far@caltech.edu.

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. Due to free riding, there is over-exploitation and under-exploration relative to the social optimum. Moreover, equilibrium social surplus is non-monotonic in the connection probability between the two agents. For some intermediate levels of connectivity, an increase in connectivity 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. For intermediate initial beliefs, equilibrium is asymmetric, with a given number of players exploring while the remaining players exploiting.

The equilibrium results in Section 3 applies 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. 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, thus 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. 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 tied 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 4, 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. 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) goes 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 features the discontinuous falls appearing in the finite economies. This means the limiting average equilibrium social surplus is weakly increasing and continuous in connectivity index. In addition, for intermediate levels of initial beliefs we identify a connectivity threshold, above which the limit of equilibrium per-capita social surplus remains *constant*. Equivalently, in the limit the social informational gain to having one more agent exploring exactly cancels out the present exploration cost, thus leading to the 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

[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).

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.

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 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.

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) and 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.

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:

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

Therefore, receiving a high payoff of $y = 1$ is perfectly conclusive about the underlying state of the world θ. 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:

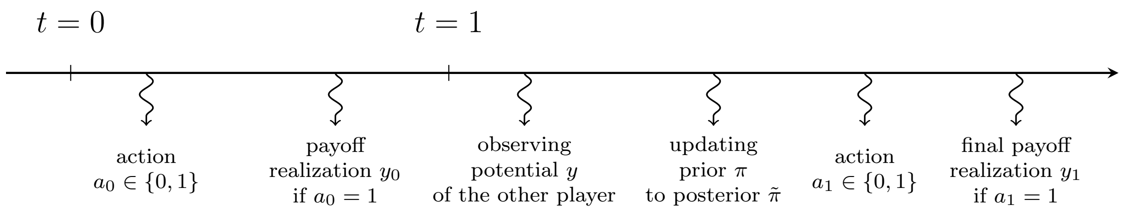

Figure 1: Timeline of the two-period bandit

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 $m \in \{0, 1, 2\}$ signals, out of which $\ell \in \{0, 1, 2\}$ had high realizations (i.e. $y = 1$):

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

The game ends with each agent making a second action choice between the safe or the risky arm. Since each agent always has the safe option at hand, the expected payoff after Bayesian updating is

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

There will be two types of *symmetric* equilibrium: *exploration* equilibrium in which both agents choose the risky arm in the first period, and *exploitation* equilibrium where both agents select the safe arm in the first period. The equilibrium is called *asymmetric* when one agent explores and the other one exploits. 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 δ and $1 - \delta$. This means that our agents are not myopic and they incorporate future gains from current exploration in their decision problem.[2]

[2]This is in contrast to the social learning models of Bala and Goyal (1998) and Sadler (2020) in which

Proposition 2.1. *There exist two thresholds $\underline{\pi} < \bar{\pi}$ such that the exploitation equilibrium appears only on $[0, \underline{\pi}]$, and the exploration equilibrium appears only on $(\bar{\pi}, 1]$. In the intermediate region $(\underline{\pi}, \bar{\pi}]$ the asymmetric equilibrium with only one agent exploring prevails. Closed form expressions for the cutoffs are*

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

The proof is expressed in the appendix.[3] This result shows 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 signal precision (β) on the above thresholds.

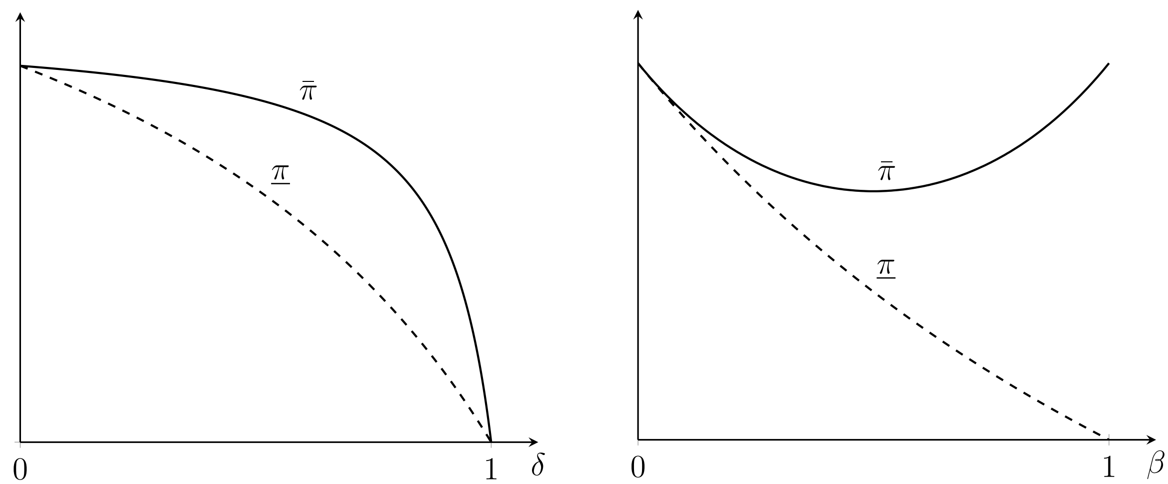

(a) Effect of δ

(b) Effect of β

Figure 2: Comparative statics of thresholds

As it appears from figure 2a higher patience (namely higher δ) is associated with smaller exploration thresholds, thereby increasing the incentives to sacrifice current payoff to learn about the risky arm and recoup 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 upper 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 of such signals. Therefore,

players are myopic. In particular, they collect information from their neighbors just to maximize their *current* period payoff.

[3]Henceforth, the proofs of all claims that are not stated in the main body are relegated to the appendix.

it lowers the individual's incentive to exploit the safe arm, hence shrinking the exploitation region (see $\underline{\pi}$ in figure 2b).

Now suppose the connection between 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 *enlarges* the exploration region.

Proposition 2.2. *In presence of imperfect connections ($p < 1$), there exist two thresholds $\underline{\pi} < \bar{\pi}$ such that the exploitation equilibrium appears only on $[0, \underline{\pi}]$, and the exploration equilibrium appears only on $(\bar{\pi}, 1]$. In the intermediate region $(\underline{\pi}, \bar{\pi}]$ the asymmetric equilibrium with only one agent exploring prevails. 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 takeaway of 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. Because, stronger connections between players increase the free-riding motives, and hence lowers the incentive for the first period exploration, that in turn translates to a higher belief threshold required for exploring the risky arm in the first period.

At this point, it is illuminating to draw the analogy between the connection probability p in the above analysis and the *number* of strategic players in Bolton and Harris (1999). In the multiplayer bandits with *perfect* connections between agents, the other players' experimentation is both a substitute and a complement for current player's exploration incentives. The substitution effect simply arises because of free-riding, that is also present in our model. The complementarity however is created due to the *encouragement effect* of each player's current exploration into the future incentives of other players' exploration, and thereby providing value to the pioneer. The latter effect is absent in our model (much like the exponential bandits in Keller et al. (2005)), because the only way to send encouraging signals to other players is to achieve a conclusive breakthrough, that in turn means the induced experimentation on others will provide no further benefit to the pioneering player.

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

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

The substantial takeaway from the above 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 (e.g. see the x-axis in figure 3a).

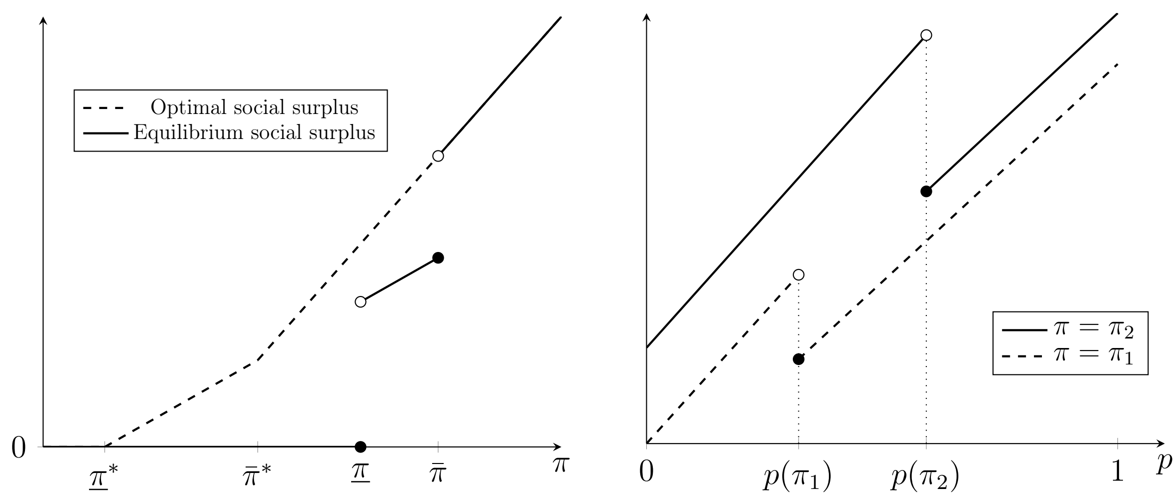

(a) Optimal vs. equilibrium social surplus

(b) Equilibrium social surplus w.r.t p

Figure 3: Social Surplus

Figure 3a shows 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 4.2, 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. Using the expressions for the social surplus in the proof of proposition 2.3, 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 asymmetric equilibrium. This pattern is exhibited 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 2.2 is increasing in p. Let $p(\pi)$ be the level at which $\bar{\pi}(p) = \pi$. For every $p < p(\pi)$, 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 increases the social surplus because it only raises the benefits of information sharing between agents and not alter the free-riding incentives (as one of them is already exploiting the safe arm).

In Section 4.1, we study the average equilibrium social surplus for the economy with many players. There we demonstrate that this pattern of being increasing in connection probability as long as the equilibrium regime does not change, while discontinuously falling at thresholds of regime change is a robust feature of this economy with many players.

3 Equilibria in Large Economy

In this section, we extend the previous two-player model to 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. This group could be the set of her immediate neighbors in the graph of connections (referred to as *local* case), or on the other extreme the set of all agents who belong to her *connected component* (referred to as *global* case).[4] 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 (that is denoted by $\mathcal{C}$ with the size of $M + 1 := |\mathcal{C}|$) in the graph of social connections.[5] In this case, effectively we think of the second period as a collection of several message passing sub-periods through which each agent gets to observe the exploration outcome of every other agent who is connected to her via a path on the graph of connections. Importantly, we further assume the random realization of the connections resolve in the beginning of period two. That is all agents are ex ante similar as of the beginning of the period one. At this stage we rather not make a specific probabilistic structure on the graph of connections (or equivalently the distribution of M), as the following equilibrium results do not depend on the specifics of the underlying random graph nor on the depth of signal observability.

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

[5] 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.

3.1 Symmetric Equilibria

Here, we study two *symmetric* equilibria, exploitation and exploration equilibrium. The exploitation equilibrium is the one in which *all* players choose the safe arm in the first period. It prevails whenever the initial belief falls below the stated $\underline{\pi}$ 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 $w_0(\pi) \geq v_1(\pi)$ whenever $\pi \leq \underline{\pi}$ stated in proposition 2.1. 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 interesting case is the examination of the existence of the exploration equilibrium in which *all* agents choose the risky arm in the first period. For this we form two payoff functions, $w_{n-1}(\pi)$ and $v_n(\pi)$. The former refers to the agent's payoff when she decides 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 (that is when all n players choose the risky arm) and play optimally in the second period. The exploration equilibrium prevails whenever $v_n(\pi) > w_{n-1}(\pi)$.

Suppose all except one individual 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 pertaining agent (whose incentive problem is being studied) observes, which is surely 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 $\ell \geq 1$ 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 everyone else is 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 $q(m) := \mathsf{P}(M = m)$, which is the probability of the randomly picked agent observing the exploration signals of m other players. Then, the above payoff can be written as

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

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

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

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

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

The exploration equilibrium thus appears when the combination of the current 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

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

Theorem 3.1 (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 interesting 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 are associated with higher bars for exploration in the equilibrium.

Remark 3.2. 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 any other choice for the distribution of M. In Section 3.3, we let the connections to follow random Erdos-Renyi graphs, thereby presenting sharper comparative static results for the exploration threshold $\bar{\pi}$.

3.2 Intermediate 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 both pure- and mixed-strategy equilibria in which both types of agents (explorers and exploiters) are present. Let $0 < k < n$, and v_k (resp. w_k) denote the expected payoff of an exploring (resp. 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 $w_k(\pi) \geq v_{k+1}(\pi)$.

Let $q_k(m) := \mathsf{P}_k(M = m)$ denote the probability of observing the first period signals of m out of k exploring individuals. 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. One should take this notation with a grain of salt because all that matters is the distribution of M, which is determined by the subscript of 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 π the second incentive constraint above fails to hold.

Lemma 3.3. *Suppose $\pi > \alpha/(1+\alpha)$, then $w_k(\pi) < v_{k+1}(\pi)$ for every k.*

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

Theorem 3.4 (Asymmetric pure-strategy equilibrium)**.** *The asymmetric equilibrium in which k players explore, where $0 < k < n$, exists if and only if*

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

Using equations in (3.5), the lower bound in (3.6) drops out of the incentive constraint $v_k > w_{k-1}$, and the upper bound from $w_k \geq v_{k+1}$, therefore 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 previous theorem, the asymmetric equilibrium with k agents exploring prevails whenever $\pi_{k-1,n} < \pi \leq \pi_{k,n}$. The full exploitation appears on $\pi \leq \pi_{0,n} \equiv \underline{\pi}$ 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.[6] 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 $\pi_{k-1,n} \leq \pi_{k,n}$. This means that the number of exploring agents in the equilibrium weakly *increases* in π.[7]

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

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

Next, we examine the *symmetric* mixed-strategy equilibria in the intermediate region. Suppose 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

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

Lastly, before stating the next result we define what it means for a random graph to be *exchangeable*, a requirement we need for the next proposition. The random structure of connections is called *exchangeable* if the probability of any event on the graph does not change with relabeling the vertices.

Proposition 3.5 (Symmetric mixed-strategy equilibrium)**.** *In the intermediate region, i.e.* $\pi \in (\underline{\pi}, \bar{\pi}]$*, with exchangeable connections, there is a unique symmetric mixed-strategy equilibrium. Furthermore, the equilibrium probability of exploration* μ *is increasing in* π*.*

3.3 Limits of Equilibria

For the first time in the paper, we make a particular assumption about the random nature of graph connections. Specifically in this section, we assume every pair of agents are connected with probability $p = \lambda/n$. We then study the impact of average degree λ on the exploration threshold $\bar{\pi}$ expressed in theorem 3.1.

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:

$$\begin{aligned}\bar{\pi}_\infty^{\text{local}} := \lim_{n\to\infty} \bar{\pi}_n^{\text{local}} &= \lim_{n\to\infty} \frac{\alpha(1-\delta)}{(1+\alpha)(1-\delta) + \delta\beta \mathsf{E}\left[(1-\beta)^D\right]} \\ &= \frac{\alpha(1-\delta)}{(1+\alpha)(1-\delta) + \delta\beta e^{-\lambda\beta}}.\end{aligned} \tag{3.7}$$

Lemma 3.6. *In the local regime, the exploration threshold* $\bar{\pi}_n$ *is eventually increasing in* n *and converges to* $\bar{\pi}_\infty^{local}$ *in* (3.7)*.*

Proof. To justify $\bar{\pi}_{n+1} > \bar{\pi}_n$, we use equation (3.4) and show that

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

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

Lemma 3.7. *In the local regime, for a fixed $k \in \mathbb{N}$, and large enough n the following ordering holds: $\pi_{k-1,n} \le \pi_{k,n+1} \le \pi_{k,n} \le \pi_{k+1,n+1}$.*

This lemma explains that for large economies with local connections adding one more individual *never* leads to fewer exploring agents in equilibrium. That is the previous agents do not change their exploration decisions as a result of newcomers joining the economy.[8]

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

Proposition 3.8 (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:*

$$\lim_{n\to\infty}\frac{k_n(\pi)}{n} = \kappa(\pi) := \begin{cases} 0 & \pi \le \underline{\pi} \\ \frac{1}{\lambda\beta}\log\frac{\delta\pi\beta}{(1-\delta)(\alpha(1-\pi)-\pi)} & \underline{\pi} < \pi < \bar{\pi}^{local}_{\infty} \\ 1 & \pi \ge \bar{\pi}^{local}_{\infty} \end{cases} \tag{3.8}$$

Figure 4 depicts the limiting fraction of exploring agents $\kappa(\pi)$ as a function of the initial belief π. The function exhibits two kinks at $\underline{\pi}$ and $\bar{\pi}^{\text{local}}_{\infty}$, where there are equilibrium regime changes from full exploitation to the intermediate asymmetric region and then to the full exploration.

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. 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 n-player economy. To study the limiting behavior of the exploration threshold in the global regime, we need an asymptotic result on the limiting distribution of $|\mathcal{C}|$. Let T be the random variable indicating the total number of

[8]It is noteworthy to mention that this conclusion mainly relies on holding the average degree λ constant while increasing the size of the economy.

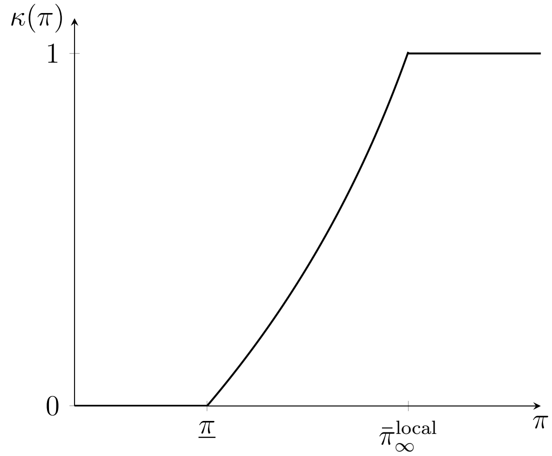

Figure 4: Limiting fraction of explorers

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 $|\mathcal{C}|$ weakly converges to T, and hence the following asymptotic result follows.

Proposition 3.9. *Let $p = \lambda/n$, and T be the total progenies of a Branching process with $\mathsf{Poisson}(\lambda)$ 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 $n \to \infty$:

$$\bar{\pi}_{\infty}^{global} := \lim_{n\to\infty} \bar{\pi}_n^{global} = \frac{\alpha(1-\delta)}{(1+\alpha)(1-\delta) + \delta\beta \mathsf{E}\left[(1-\beta)^{T-1}\right]}. \tag{3.9}$$

The moment generating function for the number of descendants of a Poisson Branching process (T) can be pinned down by the following fixed-point relation – see Section 10.4 of Alon and Spencer (2000). Fix $z \in [0,1]$ and let $X_1 \sim \mathsf{Poisson}(\lambda)$ denote the number of first-generation offspring, then

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

The solutions to the equation $xe^x = y$ are denoted by the Lambert-W function, and based on the above expression one obtains,[9]

$$- \lambda\psi(z)e^{-\lambda\psi(z)} = -\lambda z e^{-\lambda} \Rightarrow \psi(z) = -\frac{1}{\lambda}\mathsf{W}(-\lambda z e^{-\lambda}). \tag{3.10}$$

Rapid fall of exploration in the global regime (small β and $\lambda \approx 1$). As figure 5 shows there is a rapid tightening of the exploration region in the global connections, when λ increases from values just below 1 to the ones just above. Specifically, the marginal impact of increasing λ on the exploration threshold $\overline{\pi}_{\infty}^{\text{global}}$ changes greatly 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.

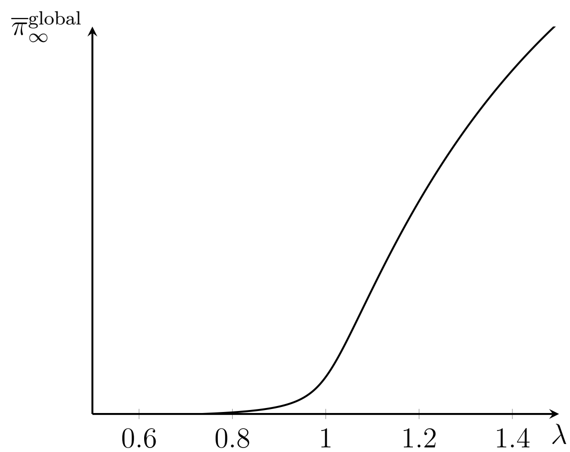

Figure 5: Rapid tightening of the exploration region
$[\delta = 0.15, \alpha = 1, \beta = 0.002]$

We can mathematically justify this sudden fall of exploration incentives by studying the effect of λ on $\overline{\pi}_{\infty}^{\text{global}}$ in equation (3.9). The only place where λ makes an impact is through $\mathsf{E}\left[(1-\beta)^T\right]$ in the denominator. 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 (3.10) 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 w.r.t λ by ψ'_λ and ψ''_λ, respectively:

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

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

At $\lambda = 1$ the above expressions imply

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

Since $\lim_{\beta\to 0}(1-\psi_\lambda) = 0$, then the above ratio explodes as $\beta \to 0$, justifying the rapid change in the sensitivity of exploration threshold w.r.t the average connections at $\lambda = 1$.

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

4 Social Surplus

In this section, we study the properties of the social surplus function in the economy with local connections. We start by studying the large-n limit of the equilibrium average social surplus in three equilibrium regions characterized previously. 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. Lastly, we establish a sufficient condition for the presence of complementarity between the initial belief π and the size of the exploring group k.

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

$$\begin{aligned} u_{k,n}(\pi) &= (1-\delta)k\left(\pi - \alpha(1-\pi)\right) \\ +\delta k\pi\beta + \delta k \sum_{m=0}^{k-1} q_{k-1}(m)\left[\pi(1-\beta)^{m+1} - \alpha(1-\pi)\right]^{+} &+ \delta k \sum_{m=0}^{k-1} q_{k-1}(m)\pi(1-\beta)\left(1-(1-\beta)^m\right) \\ +\delta(n-k)\sum_{m=0}^{k} q_k(m)\left[\pi(1-\beta)^m - \alpha(1-\pi)\right]^{+} &+ \delta(n-k)\sum_{m=0}^{k} q_k(m)\pi\left(1-(1-\beta)^m\right). \end{aligned} \tag{4.1}$$

The first line in $u_{k,n}$ denotes the first period payoff of exploration accrued to k exploring agents who chose the risky arm in the first period. The second line is the discounted second period payoff of this group, consisting of three components: discounted expected payoff when each agent received a conclusive signal in the first period (and optimally chooses the risky arm in the second period); discounted expected payoff when neither the agent nor any of her second period's contacts received a conclusive signal, and thirdly is 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's contacts did. The third line expresses the discounted second period payoff of the remaining $n-k$ exploiting agents who chose the safe arm in the first period that 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 when at least one of them did receive such a conclusive signal.

Leveraging the above representation, the following lemma studies the marginal value of one more explorer in the economy, that is $\Delta u_k := u_{k+1} - u_k$. It will be invoked both to investigate the equilibrium social surplus and the social optimum. We further use the notation $Q_a(b) := \sum_{m \leq b} q_a(m)$ to refer to the cumulative function of q_a, with the additional definition that $Q_0(0) = q_0(0) = 1$.

Lemma 4.1. *The marginal value of one more exploring agent takes the following form:*

(i) On $\frac{\pi}{1-\pi} \leq \alpha$:

$$\Delta u_k(\pi) = (1-\delta)\big(\pi - \alpha(1-\pi)\big) + \delta\pi\beta(1-p\beta)^k \left(1 + (n-1)p - \frac{kp(1-p)\beta}{1-p\beta}\right). \tag{4.2}$$

(ii) On $\frac{\pi}{1-\pi} \geq \frac{\alpha}{(1-\beta)^r}$ *and* $0 \leq k < r$:

$$\Delta u_k(\pi) = (1-\delta)\big(\pi - \alpha(1-\pi)\big).$$

(iii) When $\frac{\alpha}{(1-\beta)^r} \leq \frac{\pi}{1-\pi} \leq \frac{\alpha}{(1-\beta)^{r+1}}$ *and* $k \geq r \geq 0$,

$$\Delta u_k(\pi) = \pi B_k - \alpha(1-\pi)A_k,$$

where

$$\begin{aligned} A_k(\pi) &:= (1-\delta) + \delta(k+1)Q_k(r-1) \\ &+\delta(n-k-1)Q_{k+1}(r) - \delta k Q_{k-1}(r-1) - \delta(n-k)Q_k(r), \end{aligned} \tag{4.3}$$

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

The proof readily follows once we note that the piecewise linear components in (4.1) are positive so long as $m+1 \le r$ in the first one and $m \le r$ in the second one, thus we omit the proof.

4.1 Equilibrium Social Surplus

For $\pi \le \underline{\pi}$ no agent explores the risky arm, and thus the equilibrium social surplus is zero. On the intermediate region, i.e. $\pi \in \left(\underline{\pi}, \bar{\pi}_\infty^{\text{local}}\right)$, k_n number of individuals choose to explore where $k_n/n \to \kappa$ characterized in proposition 3.8. On this region the average equilibrium social surplus is

$$\begin{aligned} \frac{u_{k_n,n}(\pi)}{n} &= (1-\delta)\frac{k_n}{n}\left(\pi - \alpha(1-\pi)\right) + \delta\pi \\ -\delta\frac{k_n}{n}\pi(1-\beta)\mathsf{E}_{k_n-1}^{(n)}\left[(1-\beta)^M\right] &- \delta\frac{n-k_n}{n}\pi\mathsf{E}_{k_n}^{(n)}\left[(1-\beta)^M\right]. \end{aligned} \tag{4.4}$$

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. On a fixed equilibrium region (namely k_n remaining constant) increasing λ positively shifts the distribution of M in the sense of first-order stochastic dominance, and therefore based on the above representation *raises* the equilibrium social surplus. In the next proposition, we prove that on all thresholds, where the economy undergoes an equilibrium regime change, the social surplus falls, thus confirming our intuition from the two-player case.

Proposition 4.2. *The equilibrium social surplus falls discontinuously on 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 have k players exploring. Part (i) of lemma 4.1 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 (4.2) implies that at $p = \lambda_{k,n}(\pi)/n$,

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

which is always negative. Therefore, the equilibrium social surplus evaluated just above $\lambda_{k,n}(\pi)$ is smaller than that just below this threshold. That is we have a discontinuous fall of equilibrium social surplus at $\lambda_{k,n}(\pi)$. □

The graph in figure 6 shows 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. One should take this plot as a counterpart of figure 3b, except that there are more than one discontinuous jumps because of multiple equilibrium regime changes when $n > 2$.

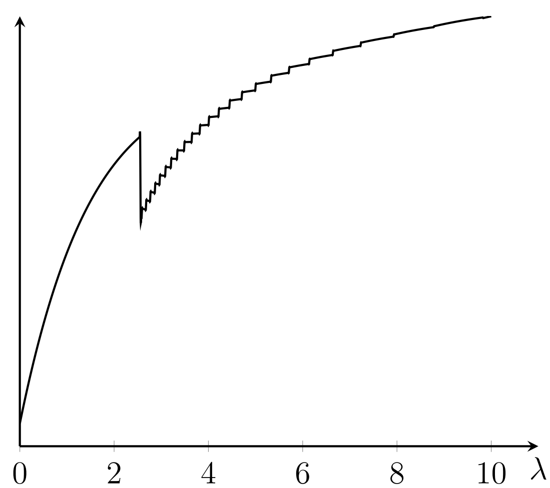

Figure 6: Finite-n equilibrium social surplus

As $n \to \infty$, the fraction k_n/n converges to κ and

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

Therefore, the limit of the average equilibrium social surplus expressed in (4.4), henceforth denoted by $\bar{u}_\infty$, is

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

which after replacing $\kappa(\pi)$ from equation (3.8) simplifies to

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

Lastly, for $\pi \geq \bar{\pi}_\infty^{\text{local}}$ all agents explore the risky arm, and it follows from (4.1) that

$$\begin{aligned} \bar{u}_\infty = \lim_{n\to\infty} \frac{u_{n,n}(\pi)}{n} &= (1-\delta)\left(\pi - \alpha(1-\pi)\right) + \delta\pi\left(1 - (1-\beta)e^{-\lambda\beta}\right) \\ &+ \delta \mathsf{E}_{M\sim\mathsf{Pois}(\lambda)}\left[\left(\pi(1-\beta)^{M+1} - \alpha(1-\pi)\right)^+\right]. \end{aligned} \tag{4.5}$$

It is worth mentioning that in the full exploration region (where $k_n = n$) $M_{k_n-1}^{(n)}$ converges weakly to $\mathsf{Poisson}(\lambda)$, and that is behind the final term in the above expression. Next proposition summarizes the above results on $\bar{u}_\infty$ as a function of both π and λ. Specifically, we are interested in how $\bar{u}_\infty$ changes w.r.t λ.

Proposition 4.3. *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 $\pi \leq \underline{\pi}$, $\bar{u}_\infty(\pi, \lambda) = 0$.

(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 $\lambda \geq \lambda(\pi)$, and follows (4.5) on $\lambda < \lambda(\pi)$.

(iii) For all $\pi \geq \frac{\alpha}{1+\alpha}$, $\bar{u}_\infty$ follows equation (4.5).

Notably, for the intermediate values of π, there is a region where the limit of the average equilibrium social surplus is independent of average degree λ. That is as long as $\pi \in \left(\underline{\pi}, \bar{\pi}_\infty^{\text{local}}\right)$ and $\lambda \geq \lambda(\pi)$ (thus the prevailing equilibrium is asymmetric) the equilibrium social surplus per-capita does not change by increasing or decreasing the connections. This is because increasing λ is associated with more free-riding and thus fewer exploring agents in equilibrium, that in turn lowers the social cost of first period exploration. On the other hand, fewer explorers corresponds to smaller benefits of second period exchange of information among the agents. In the large-n limit these two effects exactly cancel each other, thus leaving the equilibrium per-capita social surplus unaffected by λ. In particular, this is the region where for finite n, the equilibrium social surplus features bounded jumps due to the

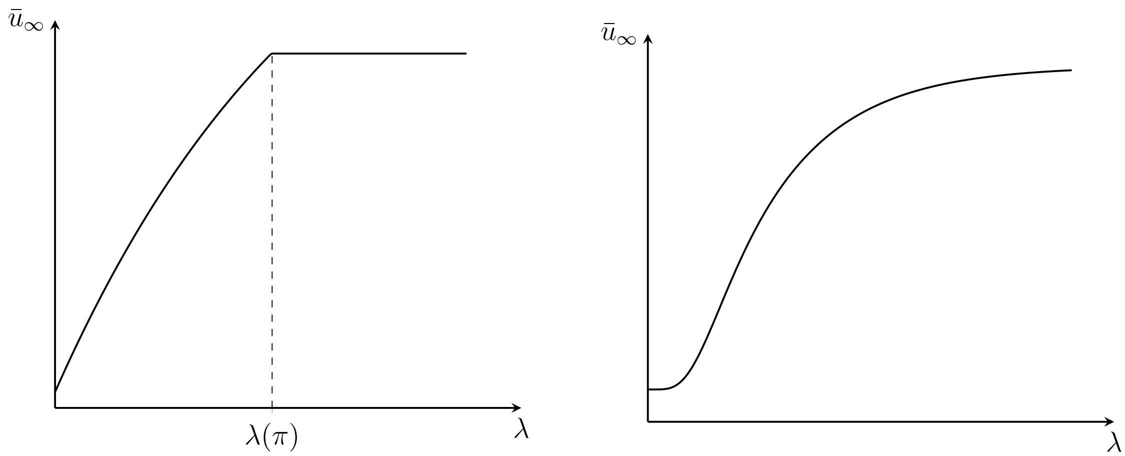

(a) $\underline{\pi} < \pi < \frac{\alpha}{1+\alpha}$ (b) $\pi \geq \frac{\alpha}{1+\alpha}$

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

regime changes in the equilibrium number of explorers (see figure 6). In the limit $n \to \infty$, these jumps in the per-capita equilibrium social surplus vanish and it becomes flat in λ.

4.2 Social Optimum

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

Firstly, we need to know when the marginal impact of one more exploring agent is positive, that is to examine $\Delta u_k = u_{k+1} - u_k$. Lemma 4.1 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 next 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 4.4. *For every $k \geq r \geq 0$:*

$$B_k \geq \max\left\{(1-\beta)^r A_k, (1-\beta)^{r+1} A_k\right\}. \tag{4.6}$$

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

Lemma 4.5. *For every fixed $\pi \leq \alpha/(1+\alpha)$, the marginal value $\Delta u_k(\pi)$ is decreasing in k.*

This result is an immediate consequence of part (i) of lemma 4.1. It stops short at claiming diminishing return for the social surplus w.r.t the number of exploring agents, and only claims that on the region where the initial belief is small. However, together with the lemmas 4.1 and 4.4, they characterize the regions where full exploitation and exploration are socially optimal.

Theorem 4.6 (Social optimum)**.** *The socially optimal outcome is full exploitation iff $\pi \leq \underline{\pi}^*$, and full exploration iff $\pi \geq \overline{\pi}^*$. Furthermore, on $[0, \underline{\pi}^*]$ the social surplus is decreasing in k ($\Delta u_k \leq 0$), and on $[\overline{\pi}^*, 1]$ it is increasing in k ($\Delta u_k \geq 0$). The cutoff points are*

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

Proof. First, we justify the lower cutoff rule for the optimality of full exploitation. Part (i) of lemma 4.1 implies $\Delta u_0 \leq 0$ on $\pi \leq \underline{\pi}^*$ and $\Delta u_0 > 0$ on $\underline{\pi}^* < \pi \leq \alpha/(1+\alpha)$. In addition, since Δu_k is decreasing on $[0, \alpha/(1+\alpha)]$ (because of lemma 4.5), then $\Delta u_k \leq \Delta u_0 = u_1 \leq u_0 = 0$ on $\pi \leq \underline{\pi}^*$, implying the optimality of full exploitation on this region. When $\pi/(1-\pi) \in [\alpha, \alpha/(1-\beta)]$, part (iii) in lemma 4.1 (with $r = 0$) says that $\Delta u_0(\pi) = \pi B_0 - \alpha(1-\pi)A_0$. If $A_0 \geq 0$, then $B_0 \geq 0$ due to lemma 4.4 and hence

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

Alternatively, if $A_0 < 0$, then again because of lemma 4.4,

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

Lastly, when $\pi/(1-\pi) > \alpha/(1-\beta)$ part (ii) of lemma 4.1 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 iff $\pi \leq \underline{\pi}^*$.

Next, we establish the optimality of full exploration above $\overline{\pi}^*$. On the region $\pi/(1-\pi) \leq \alpha$, part (i) of lemma 4.1 shows that $\Delta u_{n-1}(\pi) < 0$ on $\pi < \overline{\pi}^*$ and $\Delta u_{n-1}(\pi) \geq 0$ on $[\bar{\pi}^*, \alpha/(1+\alpha)]$. Also, lemma 4.5 results in $\Delta u_k(\pi) \geq \Delta u_{n-1}(\pi) \geq 0$ for $\pi \in [\bar{\pi}^*, \alpha/(1+\alpha)]$. Therefore, establishing that the social surplus is increasing in k, i.e. $\Delta u_k(\pi) \geq 0$ for every $\pi > \alpha/(1+\alpha)$, concludes the proof of the theorem. For every $\pi > \alpha/(1+\alpha)$, there exists r such that $\alpha/(1-\beta)^r \leq \pi/(1-\pi) \leq \alpha/(1-\beta)^{r+1}$. If $k < r$, then part (ii) states that Δu_k is positive. Alternatively, suppose $k \geq r$. Then, if $A_k \geq 0$, from lemma 4.4 it falls out that

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

and if $A_k < 0$, then again from lemma 4.4 one obtains

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

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

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

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

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

Effect of λ on the optimal exploration cutoff. The optimal exploration cutoff $\overline{\pi}^*_\infty$ 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 (that is Δu_{n-1}), and specifically its contribution in the positive externality of community exploration on an agent whose exploration failed in the first period, which is the marginal of the last term in the second line of the surplus function (4.1), namely

$$\delta\pi(1-\beta)\left[n\sum_{m=0}^{n-1} q_{n-1}(m)\big(1-(1-\beta)^m\big) - (n-1)\sum_{m=0}^{n-2} q_{n-2}(m)\big(1-(1-\beta)^m\big)\right]. \tag{4.7}$$

We employ an intuitive coupling argument to further highlight the above difference and its reaction 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 (4.7) 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 thus simplifies 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 every fixed $\pi \leq \alpha/(1+\alpha)$ the marginal impact of the exploration of the n-th individual (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 draws the large-n limits of the equilibrium and optimum exploration cutoffs as a function of λ in the local economy.

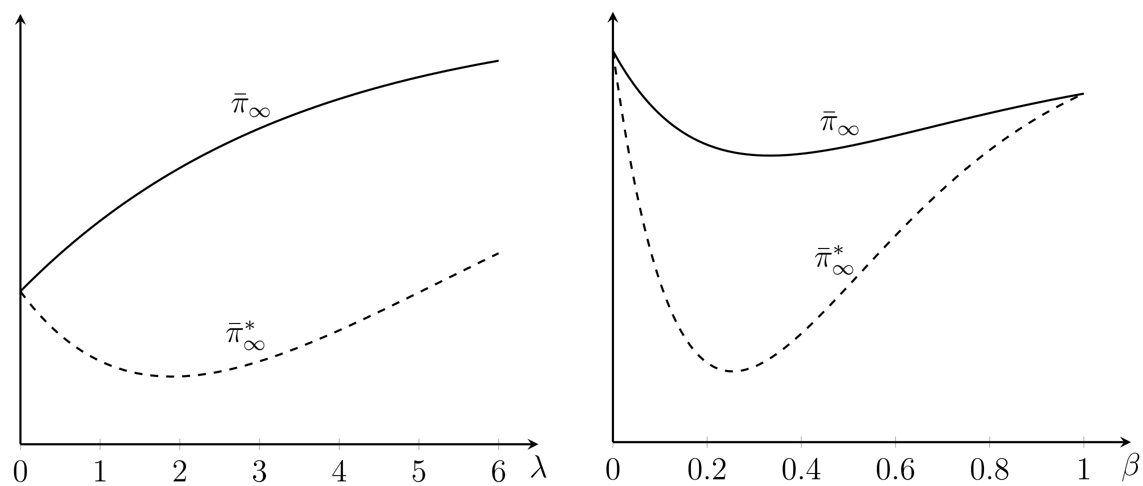

(a) Effect of λ
[$\beta = 0.3$]

(b) Effect of β
[$\lambda = 3$]

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

In figure 8b, we draw 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.

4.3 Asymptotic Complementarity

A natural question one might have is to know under what circumstances the number of exploring agents (k) and the initial belief (π) act as complements in the social surplus function. Here we define and further establish the notion of *asymptotic complementarity* between these two variables.

Definition 4.7. The social surplus function features asymptotic complementarity between k and π, when for every $k \in \mathbb{N}$ and $\pi' < \pi''$ in $[0, 1]$:

$$\liminf_{n\to\infty} \min_{0\le k<n} \left\{ \big(u_{k+1}(\pi'') - u_k(\pi'')\big) - \big(u_{k+1}(\pi') - u_k(\pi')\big) \right\} \ge 0. \tag{4.8}$$

Next proposition establishes that when it comes to the complementarity between k and π the discount factor δ and the average connections λ play a substitutable role.

Proposition 4.8. *For sufficiently small δ (specifically $\delta \le \frac{1}{\lambda+2}$), or equivalently sufficiently sparse connections, the social surplus function features asymptotic complementarity.*

Proof. Since $\Delta u_k(\pi)$ is continuous in π and differentiable except at finitely many (kink) points, then condition (4.8) is equivalent to

$$\liminf_{n\to\infty} \min_{k<n} \frac{\mathrm{d}}{\mathrm{d}\pi}\Delta u_k(\pi) \ge 0.$$

Using lemma 4.1 we verify that for large n the above condition holds. For the region (i) in lemma 4.1 we have,

$$\begin{aligned}\frac{\mathrm{d}}{\mathrm{d}\pi}\Delta u_k(\pi) &= (1-\delta)(1+\alpha) + \delta\beta(1-p\beta)^k \left(1 + (n-1)p - \frac{kp(1-p)\beta}{1-p\beta}\right) \\ &\ge (1-\delta)(1+\alpha) + \delta(1-\lambda\beta/n)^{n-2}\big(\lambda(1-\beta) + 1 - \lambda/n\big),\end{aligned}$$

and consequently,

$$\liminf_{n\to\infty} \frac{\mathrm{d}}{\mathrm{d}\pi}\Delta u_k(\pi) \ge (1-\delta)(1+\alpha) + \delta e^{-\lambda\beta}\big(\lambda(1-\beta)+1\big) > 0.$$

On the region (ii), $\mathrm{d}\Delta u_k(\pi)/\mathrm{d}\pi = (1-\delta)(1+\alpha) > 0$. Lastly on the region (iii), due to the lemma 4.4

$$\frac{\mathrm{d}}{\mathrm{d}\pi}\Delta u_k(\pi) = B_k + \alpha A_k \ge \big((1-\beta)^r + \alpha\big)A_k. \tag{4.9}$$

Given the definition of A_k in (4.3), and using the the fact that $\forall a \ge b$, $Q_{a+1}(b+1) = pQ_a(b) + (1-p)Q_a(b+1)$, one obtains the following equivalent expression for A_k:

$$A_k = 1 - \delta - \delta kpq_{k-1}(r-1) - \delta(n-k)pq_k(r) - \delta(1-p)q_k(r).$$

Since $q_{k-1}(r-1)$ and $q_k(r)$ are less than or equal to 1, then $A_k \ge 1 - \delta(\lambda+2)$ which is nonnegative, and hence (4.9) implies $\min_k \mathrm{d}\Delta u_k(\pi)/\mathrm{d}\pi \geq 0$, thereby concluding the proof. $\square$

The characterization of complementarity in the previous result is rather sharp as depicted in the simulation of figure 9. Essentially the main obstacle behind the positivity of $\mathrm{d}\Delta u_k(\pi)/\mathrm{d}\pi$, as can be verified in the above proof, is related to the region where $k \geq r$ and $\pi/(1-\pi) \in [\alpha/(1-\beta)^r, \alpha/(1-\beta)^{r+1}]$. To simulate the graph shown in figure 9, we picked $r = \lfloor k\lambda/n \rfloor$, that is closest to the peak of the Binomial probabilities $q_k(r)$ and $q_{k-1}(r-1)$, and hence minimizes A_k the most. Then, for each $\lambda \in (\delta^{-1}-3, \delta^{-1}-1)$ we find $\min_k \mathrm{d}\Delta u_k(\pi)/\mathrm{d}\pi$ with the above informed guess for r to ease and speed up the computation.

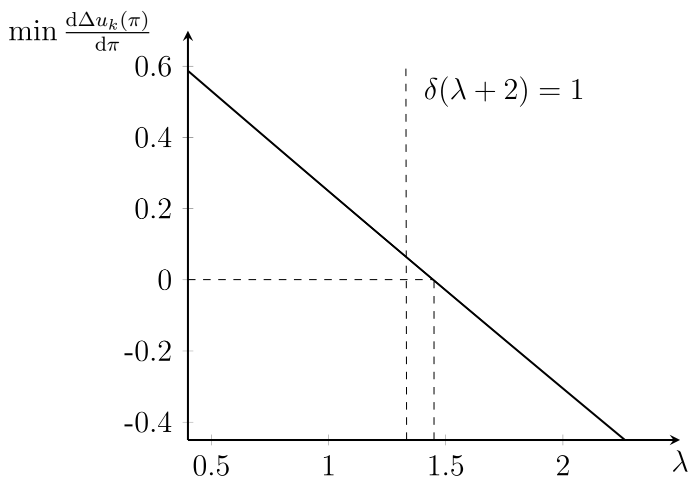

Figure 9: Simulation of $\min \frac{\mathrm{d}\Delta u_k(\pi)}{\mathrm{d}\pi}$ around $\lambda = \delta^{-1} - 2$
$[\delta = 0.3, \alpha = 1, \beta = 0.05, n = 500]$

5 Conclusion and Additional Discussion

We highlighted the tension between information diffusion and production in societies and organizations. Connectivity naturally enhances knowledge diffusion but may induce free riding and homogenization, undermining knowledge production and social welfare.

Our model is stylized and some of its aspects may exacerbate results. For example, all an agent needs in order to mimic other agents is knowledge. In practice, however, resources, which are unevenly distributed, may play a key role. Due to resource constraints an agent may not be able to mimic another agent even if they have the same knowledge. In addition, our agents have homogeneous preference, increasing the incentives for free riding and homogenization. Relaxing these and other assumptions are fruitful avenues for further research.

Finally, we only considered connectivity in terms of a basic random matching structure in which all agents are equally likely to meet other agents. In practice, however, societies and organizations have network structures with differently connected agents. Investigating the effects of such particular structures on knowledge production and diffusion seems promising.

A Proofs

A.1 Proof of Proposition 2.1

Let M and L be the random variables encoding the total number of signals each agent observes in the second period (other than herself) and the number of successful high signals among them, with respective realizations of m and ℓ.

Exploitation equilibrium. We first show that the exploitation equilibrium prevails only when $\pi \leq \underline{\pi}$. Suppose agent j decides to choose the safe arm in the first period. Then agent i's expected payoff from exploration is

$$\begin{aligned} v_1(\pi) &:= (1-\delta)\left(\pi - \alpha(1-\pi)\right) + \delta \sum_{\ell=0}^{1} \mathsf{P}(L=\ell \big| M=1) \mathsf{E}\left[\theta - \alpha(1-\theta) \big| L=\ell, M=1\right]^+ \\ &= (1-\delta)\left(\pi - \alpha(1-\pi)\right) + \delta \sum_{\ell=0}^{1} \mathsf{E}\left[\theta - \alpha(1-\theta) 1_{\{L=\ell\}} \big| M=1\right]^+ \\ &= (1-\delta)\left(\pi - \alpha(1-\pi)\right) + \delta \left\{\left[\pi(1-\beta) - \alpha(1-\pi)\right]^+ + \pi\beta\right\}. \end{aligned}$$

The above expression coined as $v_1(\pi)$ is the exploration payoff when only one agent is exploring and the other agent is inactive. This function has to be weighed against $w_0(\pi)$, namely the expected payoff of agent i when neither of the agents are exploring, where $w_0(\pi) = \delta\left[\pi - \alpha(1-\pi)\right]^+$. Exploitation equilibrium thus prevails whenever $v_1(\pi) \leq w_0(\pi)$:

$$\begin{aligned} v_1(\pi) \leq w_0(\pi) \Leftrightarrow & \overbrace{(1-\delta)\left(\pi - \alpha(1-\pi)\right)}^{\text{current risky payoff}} + \overbrace{\delta\pi\beta}^{\substack{\text{exploration gain} \\ \text{upon conclusive signals}}} \\ &\leq \delta\left[\pi - \alpha(1-\pi)\right]^+ - \delta\left[\pi(1-\beta) - \alpha(1-\pi)\right]^+ \\ &= \text{opportunity cost absent of conclusive signals} \end{aligned}$$

Denote the *lhs* of the above inequality, which is the benefit of exploration, by $B(\pi)$ and the *rhs*, which is the opportunity cost of exploration, by $C(\pi)$. The cost component features two kinks at ξ_0 and ξ_1, that are respectively:

$$\frac{\xi_0}{1-\xi_0} = \alpha, \quad \frac{\xi_1}{1-\xi_1} = \frac{\alpha}{1-\beta}.$$

We next show $B(\pi) > C(\pi)$ at three corner points $\pi \in \{\xi_0, \xi_1, 1\}$, and further $C(0) > B(0)$. These together will prove that there exists $\underline{\pi} < \xi_0$ (expressed in the proposition), below which the exploitation equilibrium prevails.

$$\begin{aligned}
\pi = 0 : B(0) = -(1-\delta)\alpha < 0 = C(0) \\
\pi = \xi_0 : B(\xi_0) = \delta \xi_0 \beta > 0 = C(\xi_0) \\
\pi = 1 : B(1) = 1 - \delta + \delta\beta > \delta\beta = C(1)
\end{aligned}$$

Lastly, with some minor algebraic work, one can show for any combination of parameters at $\pi = \xi_1$,

$$B(\xi_1) > C(\xi_1) \Leftrightarrow \frac{(1-\delta)\alpha\beta}{\alpha + 1 - \beta} > 0,$$

which always holds. Therefore, there exists $\underline{\pi} \in (0, \xi_0)$, only below which the exploitation equilibrium prevails. At $\underline{\pi}$, $B(\underline{\pi}) = C(\underline{\pi})$, that yields the expression in the proposition for $\underline{\pi}$.

Exploration equilibrium. Now we assume agent i believes agent j explores the risky arm in the first period, and then we study her incentive to explore as well. Let $w_1(\pi)$ be her payoff when she chooses to exploit, that is when only one agent is exploring (in this case the opponent j):

$$\begin{aligned}
w_1(\pi) &:= \delta \sum_{\ell=0}^{1} \mathsf{P}(L = \ell \big| M = 1) \mathsf{E}\left[\theta - \alpha(1-\theta) \big| L = \ell, M = 1\right]^+ \\
&= \delta \left\{ [\pi(1-\beta) - \alpha(1-\pi)]^+ + \pi\beta \right\}.
\end{aligned}$$

Alternatively, if agent i explores, that is when two agents are exploring, then her expected payoff would be

$$\begin{aligned}
v_2(\pi) &:= (1-\delta)\left(\pi - \alpha(1-\pi)\right) + \delta \sum_{\ell=0}^{2} \mathsf{P}(L = \ell \big| M = 2) \mathsf{E}\left[\theta - \alpha(1-\theta) \big| L = \ell, M = 2\right]^+ \\
&= (1-\delta)\left(\pi - \alpha(1-\pi)\right) + \delta \left\{ \left[\pi(1-\beta)^2 - \alpha(1-\pi)\right]^+ + 2\pi\beta(1-\beta) + \pi\beta^2 \right\}.
\end{aligned}$$

Agent i selects the risky arm and the exploration equilibrium prevails if $v_2(\pi) > w_1(\pi)$. Let

ξ_1 and ξ_2 be the respective kink points of w_1 and v_2:

$$\frac{\xi_1}{1-\xi_1} = \frac{\alpha}{1-\beta}, \quad \frac{\xi_2}{1-\xi_2} = \frac{\alpha}{(1-\beta)^2}.$$

Then, the exploration incentive condition is expressed by

$$\begin{aligned} v_2(\pi) > w_1(\pi) \Leftrightarrow & \overbrace{(1-\delta)\left(\pi - \alpha(1-\pi)\right)}^{\text{current risky payoff}} + \overbrace{\delta\pi\beta(1-\beta)}^{\substack{\text{exploration gain} \\ \text{upon conclusive signals}}} \\ > \delta & \left\{ [\pi(1-\beta) - \alpha(1-\pi)]^+ - \left[\pi(1-\beta)^2 - \alpha(1-\pi)\right]^+ \right\} \\ & = \text{opportunity cost absent of conclusive signals.} \end{aligned}$$

Analogous to the previous case, denote the *lhs* by $B(\pi)$ and the *rhs* by $C(\pi)$. We show at three corner points $\pi \in \{\xi_1, \xi_2, 1\}$, $B(\pi) > C(\pi)$, and $B(0) < C(0)$, therefore, there exists a unique $\bar{\pi}$ above which $v_2(\pi) > w_1(\pi)$ and the exploration equilibrium prevails.

$$\begin{aligned} \pi = 0 : B(0) &= -(1-\delta)\alpha < 0 = C(0) \\ \pi = \xi_1 : B(\xi_1) &= (1-\delta)\left(\xi_1 - \alpha(1-\xi_1)\right) + \delta\xi_1\beta(1-\beta) \\ &= \frac{\alpha\beta(1-\delta\beta)}{\alpha+1-\beta} > 0 = C(0) \\ \pi = 1 : B(1) &= (1-\delta) + \delta\beta(1-\beta) > \delta\beta(1-\beta) = C(1) \end{aligned}$$

With some algebraic work, one can also show at $\pi = \xi_2$,

$$B(\xi_2) > C(\xi_2) \Leftrightarrow 1 - \delta - \delta(1-\beta)^2 > (1-2\delta)(1-\beta)^2,$$

which is always true. Therefore, there exists a unique $\bar{\pi}$ at which $B(\bar{\pi}) = C(\bar{\pi})$, and for all $\pi > \bar{\pi}$ the exploration equilibrium prevails. Solving the previous equality leads to the expression for $\bar{\pi}$ in the proposition.

Asymmetric equilibrium. This is the pure-strategy equilibrium in which only one agent explores. Suppose agent i exploits and agent j explores. From the previous analysis (for exploration equilibrium) agent i is best-responding by exploitation if $\pi \leq \bar{\pi}$. And from the analysis for exploitation equilibrium agent j is best-responding by exploration if $\pi > \underline{\pi}$. □

A.2 Proof of Proposition 2.2

The imperfect connection does not impact the determination of the exploitation equilibrium, because the other player is not exploring, thus having a perfect or imperfect access to her experimentation outcome will not change the incentive problem of the current player. Therefore, we only study the conditions for the existence of the exploration equilibrium. Assume player j is choosing the risky arm in the first period. Recall that M and L are random variables respectively representing the number of signals agent i observes in the second period (other than herself) and the number of successful ones among them. Then, agent i's expected payoff from choosing the safe arm is

$$\begin{aligned} w_1(\pi) &= \delta \sum_{m=0}^{1} \sum_{\ell=0}^{m} \mathsf{P}\left(M=m, L=\ell\right) \mathsf{E}\left[\theta - \alpha(1-\theta) \middle| M=m, L=\ell\right]^+ \\ &= \delta \sum_{m=0}^{1} \sum_{\ell=0}^{m} \mathsf{P}(M=m) \mathsf{E}\left[\theta - \alpha(1-\theta) 1_{\{L=\ell\}} \middle| M=m\right]. \end{aligned}$$

Next, we express agent i's expected payoff from exploration. In this case, agent i can benefit from the outcome of her first period experimentation as well, while evaluating her choice in the second period. Therefore, we further condition her second period expected payoff on the y value she observed in the first period:

$$\begin{aligned} v_2(\pi) &= (1-\delta)\left(\pi - \alpha(1-\pi)\right) + \\ &\quad \delta \sum_{m=0}^{1} \sum_{\ell=0}^{1} \sum_{y\in\{-\alpha,1\}} \mathsf{P}\left(M=m, L=\ell, y_0=y\right) \mathsf{E}\left[\theta - \alpha(1-\theta) \middle| M=m, L=\ell, y_0=y\right]^+ \\ &= (1-\delta)\left(\pi - \alpha(1-\pi)\right) + \\ &\quad \delta \sum_{m=0}^{1} \sum_{\ell=0}^{1} \sum_{y\in\{-\alpha,1\}} \mathsf{P}\left(M=m\right) \mathsf{E}\left[\theta - \alpha(1-\theta) 1_{\{L=\ell\}} 1_{\{y_0=y\}} \middle| M=m\right]^+. \end{aligned}$$

The exploration equilibrium prevails when $v_2(\pi) > w_1(\pi)$, that is equivalent to

$$\begin{aligned} &\overbrace{(1-\delta)\left(\pi - \alpha(1-\pi)\right)}^{\text{current risky payoff}} + \overbrace{\delta\pi\beta\left[(1-p) + p(1-\beta)\right]}^{\substack{\text{exploration gain} \\ \text{upon conclusive signals}}} > \\ &\quad \delta(1-p)\left\{\left[\pi - \alpha(1-\pi)\right]^+ - \left[\pi(1-\beta) - \alpha(1-\pi)\right]^+\right\} \\ &\quad + \delta p \left\{\left[\pi(1-\beta) - \alpha(1-\pi)\right]^+ - \left[\pi(1-\beta)^2 - \alpha(1-\pi)\right]^+\right\} \\ &\quad = \text{opportunity cost absent of conclusive signals}. \end{aligned} \tag{A.1}$$

The opportunity cost function is piecewise linear and increasing in π, with three breaking points at $\{\xi_r : r = 0, 1, 2\}$, where $\frac{\xi_r}{1-\xi_r} = \frac{\alpha}{(1-\beta)^r}$. In what follows we examine (A.1) over four intervals of initial beliefs:

(i) $\pi \leq \xi_0$: on this region the opportunity cost is zero and (A.1) reduces to

$$\frac{\alpha(1-\delta)}{(1+\alpha)(1-\delta)+\delta\beta(1-p\beta)} < \pi. \tag{A.2}$$

Note that the above lower bound is always less than $\alpha/(1+\alpha)$, leaving us with a non-empty region for exploration equilibrium.

(ii) $\xi_0 \leq \pi \leq \xi_1$: on this region only the first component of the opportunity cost is nonzero, and (A.1) simplifies to

$$(1-2\delta+\delta p+\delta\beta(1-p\beta))\,\pi > (1-2\delta+\delta p)\alpha(1-\pi). \tag{A.3}$$

If the coefficient of π in the above inequality is nonnegative, then it reduces to

$$\frac{\pi}{1-\pi} > \frac{\alpha(1-2\delta+\delta p)}{1-2\delta+\delta p+\delta\beta(1-p\beta)},$$

which always holds on $\pi \in [\xi_0, \xi_1]$, as the *rhs* above is smaller than $\alpha = \xi_0/(1-\xi_0)$. Alternatively, if the coefficient of π is negative, then (A.3) becomes equivalent to

$$\frac{\pi}{1-\pi} < \frac{\alpha(1-2\delta+\delta p)}{1-2\delta+\delta p+\delta\beta(1-p\beta)},$$

which again always holds on $\pi \in [\xi_0, \xi_1]$, because in this case the *rhs* above is larger than $\frac{\alpha}{1-\beta} = \frac{\xi_1}{1-\xi_1}$. Therefore, exploration equilibrium appears entirely on this region.

(iii) $\xi_1 \leq \pi \leq \xi_2$: on this region the first three components of the opportunity cost term in (A.1) become active, and the inequality reduces to

$$\left(1-\delta-\delta p(1-\beta)^2\right)\pi > (1-\delta-\delta p)\alpha(1-\pi).$$

With similar reasoning as in the previous region, one can show regardless of the sign of the coefficient of π in the above inequality, it always holds on the region $\pi \in [\xi_1, \xi_2]$, therefore exploration equilibrium appears entirely on this set as well.

(iv) $\xi_2 \leq \pi$: all four components of the opportunity cost are active on this region, hence the incentive inequality (A.1) boils down to

$$(1-\delta)\left(\pi-\alpha(1-\pi)\right)+\delta\pi\beta-\delta p\pi\beta^2>\delta(1-p)\pi\beta+\delta p\pi(1-\beta)\beta,$$

which always holds.

Therefore, the only restriction for the existence of the exploration equilibrium is (A.2), above that threshold, such equilibrium always exists.

Lastly, the previous analysis for exploitation and exploration equilibria implies that when $\pi\in(\underline{\pi},\bar{\pi}]$ the only equilibrium that survives is the asymmetric one, in which only one player explores. □

A.3 Proof of Proposition 2.3

Let $u_k(\pi)$ be the social surplus function when k players are exploring in the first period, then:

$$\begin{aligned}
u_0(\pi) &= 2\delta\left[\pi-\alpha(1-\pi)\right]^+,\\
u_1(\pi) &= (1-\delta)\left(\pi-\alpha(1-\pi)\right)+\delta\left\{\pi\beta+\left[\pi(1-\beta)-\alpha(1-\pi)\right]^+\right\}\\
&\quad+\delta p\left\{\pi\beta+\left[\pi(1-\beta)-\alpha(1-\pi)\right]^+\right\}+\delta(1-p)\left[\pi-\alpha(1-\pi)\right]^+,\\
u_2(\pi) &= 2(1-\delta)\left(\pi-\alpha(1-\pi)\right)+2\delta p\left\{\pi\beta^2+2\pi\beta(1-\beta)+\left[\pi(1-\beta)^2-\alpha(1-\pi)\right]^+\right\},\\
&\quad+2\delta(1-p)\left\{\pi\beta+\left[\pi(1-\beta)-\alpha(1-\pi)\right]^+\right\}.
\end{aligned}$$

Some straightforward analysis shows that $u_0(\pi)\geq u_1(\pi)\Leftrightarrow\pi\leq\underline{\pi}^*$ and $u_2(\pi)\geq u_1(\pi)\Leftrightarrow\pi\geq\bar{\pi}^*$, thereby establishing the proof. □

A.4 Proof of Theorem 3.1

If $\alpha(1-\pi)/\pi\leq 1$, define $\bar{m}:=\max\left\{0\leq m\leq n:(1-\beta)^m\geq\frac{\alpha(1-\pi)}{\pi}\right\}$, 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

$$\begin{aligned}
&(1-\delta)\left(\pi-\alpha(1-\pi)\right)+\delta\pi\beta\sum_{m=0}^{n-1}q(m)(1-\beta)^m\\
&>\delta\pi\beta\sum_{0\leq m<\bar{m}}q(m)(1-\beta)^m+\delta q_{\bar{m}}\left[\pi(1-\beta)^{\bar{m}}-\alpha(1-\pi)\right]^+,
\end{aligned}\tag{A.4}$$

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 $\pi/(1-\pi) \leq \alpha$. In this case equation (A.4) 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

$$\frac{\pi}{1-\pi} \in \left[\frac{\alpha}{(1-\beta)^{\bar{m}}}, \frac{\alpha}{(1-\beta)^{\bar{m}+1}}\right], \tag{A.5}$$

equation (A.4) holds consistently. On the above region, (A.4) is equivalent to

$$\pi\Big\{1 - \delta - \delta(1-\beta)^{\bar{m}} q_{\bar{m}} + \delta\beta\mathsf{E}\left[(1-\beta)^M; M \geq \bar{m}\right]\Big\} > \alpha(1-\pi)\left(1 - \delta - \delta q_{\bar{m}}\right). \tag{A.6}$$

If the coefficient of π in (A.6) is positive, then it becomes equivalent to

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

which always holds on (A.5) because $\frac{\alpha}{(1-\beta)^{\bar{m}}}$ is greater than the *rhs* above. Alternatively, if the coefficient of π in (A.6) is negative, then it becomes equivalent to

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

which again always holds on (A.5), 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.5), 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 $\pi > \bar{\pi}$. □

A.5 Proof of Lemma 3.3

Let us look at the difference

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

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

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

which is always positive. ◻

A.6 Proof of Proposition 3.5

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 $j \in \{1, \ldots, n-1\}$. 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 $M \overset{d}{=} X_1 + \ldots + X_k$. Using this coupling approach one can write the increment of Γ as

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

in that we use the notation $X_1^{k-1} := \{X_1, \ldots, X_{k-1}\}$. The inner expectation above is equal to

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

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

Proving the proposition, we first show in the intermediate region there exists a unique μ satisfying $v(\pi;\mu) = w(\pi;\mu)$. This constraint is equivalent to

$$\begin{aligned} w(\pi;\mu) &= \delta\pi \sum_{k=0}^{n-1} \binom{n-1}{k} \mu^k (1-\mu)^{n-1-k} \mathsf{E}_k\left[1 - (1-\beta)^M\right] \\ &= v(\pi;\mu) = (1-\delta)\left(\pi - \alpha(1-\pi)\right) + \delta\pi\beta \\ &\qquad + \delta\pi(1-\beta) \sum_{k=0}^{n-1} \binom{n-1}{k} \mu^k (1-\mu)^{n-k} \mathsf{E}_{k-1}\left[1 - (1-\beta)^M\right], \end{aligned}$$

that in turn holds iff

$$\frac{(1-\delta)\left(\pi - \alpha(1-\pi)\right)}{\delta\pi} = \sum_{k=0}^{n-1} \binom{n-1}{k} \mu^k (1-\mu)^{n-1-k} \Gamma_k = \mathsf{E}_{k\sim\mathsf{Bin}(n-1,\mu)}\Gamma_k. \tag{A.7}$$

Since $v(\pi;0) = v_1(\pi) > w(\pi) = w(\pi;0)$ and $v(\pi;1) = v_n(\pi) \leq w_{n-1}(\pi) = w(\pi;1)$ on the intermediate region, then there exists $\mu^* \in (0,1]$ satisfying (A.7). In addition, $\mathsf{Bin}(n-1,\mu)$

increases in the FOSD sense w.r.t μ. Due to the previous lemma, Γ is increasing in k, therefore, the *rhs* of (A.7) becomes increasing in μ, and this establishes the uniqueness of μ^* satisfying (A.7). Lastly, since the *lhs* is increasing π, the equilibrium point μ^* increases in π. □

A.7 Proof of Lemma 3.7

Since $p_n = \frac{\lambda}{n} \geq \frac{\lambda}{n+1} = p_{n+1}$, then a coupling argument shows that on a same probability space $M_k^{(n)} \geq M_k^{(n+1)}$, therefore $\mathsf{E}\left[(1-\beta)^{M_k^{(n)}}\right] \leq \mathsf{E}\left[(1-\beta)^{M_k^{(n+1)}}\right]$, that in turn implies $\pi_{k,n+1} \leq \pi_{k,n}$. Next, note that with local connections,

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

Additionally,

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

Since for large n, one can readily show

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

then it holds that

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

and $\pi_{k+1,n+1} \geq \pi_{k,n}$ for large enough n. Similarly, $\pi_{k,n+1} \geq \pi_{k-1,n}$, thus concluding the proof. □

A.8 Proof of Proposition 3.8

For every $\pi \leq \underline{\pi}$, the full exploitation equilibrium prevails, thus $k_n(\pi) = 0$. Also, for every $\pi \geq \bar{\pi}_\infty^{\text{local}}$, due to lemma 3.6, it follows that $\pi > \bar{\pi}_n$ for large enough n, hence $k_n(\pi) = n$. Thus, 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 iff

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

Therefore,

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

□

A.9 Proof of Proposition 3.9

First, we show how the size of the connected component in a random Erdos-Renyi graph with parameters $(n, p=\lambda/n)$ can be approximated with the descendants of a Branching process with $\mathsf{Bin}(n,p)$ offspring distribution, denoted 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:

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

Next, we see how the total number of the progenies of a Binomial Branching process with parameters (n,p) can be approximated by the Branching process with $\mathsf{Poisson}(np)$ offspring distribution, denoted 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

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

Subsequently, the last two relations give us

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

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 $\mathsf{P}_{\lambda_n}(T \geq k) \to \mathsf{P}_\lambda(T \geq k)$. Using this and the above inequality one reaches the conclusion that for every $k \in \mathbb{N}$,

$$\lim_{n\to\infty} \mathsf{P}(|\mathcal{C}| \geq k) = \mathsf{P}_\lambda(T \geq k). \tag{A.8}$$

Let $\bar{\mathbb{N}} = \mathbb{N} \cup \{\infty\}$, then $|\mathcal{C}|$ and T are $\bar{\mathbb{N}}$-valued random variables, which is a discrete metric space. Therefore, the limiting result in (A.8) 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 proposition 3.9. Part (ii) immediately 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

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

supporting equation (3.9). □

A.10 Proof of Lemma 4.4

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

$$\begin{aligned}(1-\delta)\big(1-(1-\beta)^r\big) \geq & -\delta k\Big[Q_{k-1}(r-1)(1-\beta)^r + \sum_{m=r}^{k-1} q_{k-1}(m)(1-\beta)^{m+1}\Big] \\ & - \delta(n-k)\Big[Q_k(r)(1-\beta)^r + \sum_{m=r+1}^{k} q_k(m)(1-\beta)^m\Big] \\ & + \delta(k+1)\Big[Q_k(r-1)(1-\beta)^r + \sum_{m=r}^{k} q_k(m)(1-\beta)^{m+1}\Big] \\ & + \delta(n-k-1)\Big[Q_{k+1}(r)(1-\beta)^r + \sum_{m=r+1}^{k+1} q_{k+1}(m)(1-\beta)^m\Big].\end{aligned} \tag{A.9}$$

The *lhs* of the above inequality is nonnegative, thus to justify that $B_k \geq (1-\beta)^r A_k$ it is enough to show that the following equivalent representation for the *rhs* is negative. In that, we use the notation E_k to express the expectation w.r.t to the distribution $\mathsf{Bin}(k,p)$, and the random variable M follows the corresponding distribution in the subscript of E.[10]

$$\begin{aligned} rhs \text{ of (A.9)} &= -\delta k \mathsf{E}_{k-1}\left[(1-\beta)^{(M+1)\vee r}\right] + \delta(k+1)\mathsf{E}_k\left[(1-\beta)^{(M+1)\vee r}\right] \\ &\quad - \delta(n-k)\mathsf{E}_k\left[(1-\beta)^{M\vee r}\right] + \delta(n-k-1)\mathsf{E}_{k+1}\left[(1-\beta)^{M\vee r}\right] \end{aligned}$$

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 $\mathsf{Bin}(k+1,p) \succeq \mathsf{Bin}(k,p)$, yields the following upper bound:

$$\begin{aligned} rhs \text{ of (A.9)} &\leq \delta \mathsf{E}_k\left[(1-\beta)^{(M+1)\vee r}\right] - \delta \mathsf{E}_{k+1}\left[(1-\beta)^{M\vee r}\right] \\ &= \delta \mathsf{E}_k\left[(1-\beta)^{(M+1)\vee r}\right] - \delta p \mathsf{E}_k\left[(1-\beta)^{(M+1)\vee r}\right] - \delta(1-p)\mathsf{E}_k\left[(1-\beta)^{M\vee r}\right] \\ &= \delta(1-p)\mathsf{E}_k\left[(1-\beta)^{(M+1)\vee r} - (1-\beta)^{M\vee r}\right] \leq 0. \end{aligned}$$

For the second part of the inequality, namely $B_k \geq (1-\beta)^{r+1} A_k$, one arrives to the following equivalent condition:

$$\begin{aligned} (1-\delta)\big(1-(1-\beta)^{r+1}\big) \geq \delta(1-\beta)\Bigg\{ & -k\Big[Q_{k-1}(r-1)(1-\beta)^r + \sum_{m=r}^{k-1} q_{k-1}(m)(1-\beta)^m\Big] \\ & - (n-k)\Big[Q_k(r)(1-\beta)^r + \sum_{m=r+1}^{k} q_k(m)(1-\beta)^{m-1}\Big] \\ & + (k+1)\Big[Q_k(r-1)(1-\beta)^r + \sum_{m=r}^{k} q_k(m)(1-\beta)^m\Big] \\ & + (n-k-1)\Big[Q_{k+1}(r)(1-\beta)^r + \sum_{m=r+1}^{k+1} q_{k+1}(m)(1-\beta)^{m-1}\Big]\Bigg\}. \end{aligned} \tag{A.10}$$

The *lhs* to (A.10) is nonneagtive, thus it is enough to show the *rhs* is negative to justify $B_k \geq (1-\beta)^{r+1} A_k$. For that, we appeal to the following equivalent representation:

$$\begin{aligned} rhs \text{ of (A.10)} = \delta(1-\beta)\Big\{ & -k\mathsf{E}_{k-1}\left[(1-\beta)^{M\vee r}\right] + (k+1)\mathsf{E}_k\left[(1-\beta)^{M\vee r}\right] \\ & - (n-k)\mathsf{E}_k\left[(1-\beta)^{(M-1)\vee r}\right] + (n-k-1)\mathsf{E}_{k+1}\left[(1-\beta)^{(M-1)\vee r}\right]\Big\}. \end{aligned}$$

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

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

$$\begin{aligned}
rhs \text{ of (A.10)} &\leq \delta(1-\beta)\Big(\mathsf{E}_k\left[(1-\beta)^{M\vee r}\right] - \mathsf{E}_{k+1}\left[(1-\beta)^{(M-1)\vee r}\right]\Big) \\
&= \delta(1-\beta)\Big(\mathsf{E}_k\left[(1-\beta)^{M\vee r}\right] - p\mathsf{E}_k\left[(1-\beta)^{M\vee r}\right] - (1-p)\mathsf{E}_k\left[(1-\beta)^{(M-1)\vee r}\right]\Big) \\
&= \delta(1-\beta)(1-p)\mathsf{E}_k\left[(1-\beta)^{M\vee r} - (1-\beta)^{(M-1)\vee r}\right] \leq 0.
\end{aligned}$$

Therefore, both inequalities were proved, and thus the claim (4.6) in the lemma is established.

□

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

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

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

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

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

[9] 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.

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

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

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

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

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

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

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

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

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

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

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

[21] ——— (2009). "Informational Externalities and Emergence of Consensus," *Games and Economic Behavior*, 66(2): 979–994.

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

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

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

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