Sharing Rival Information

Nicole Immorlica, Brendan Lucier,*Evan Sadler†

November 17, 2014

Abstract

We study the spread of rival information through social ties. A piece of information loses value as more individuals learn it, but players derive a positive payoff from being first to share the information with friends and neighbors. If there are either very few or very many opportunities to interact with neighbors, sharing occurs robustly in equilibrium. For moderate levels of interaction, there is a non-trivial tradeoff between the benefits of sharing and the lost value from further spread of the information. In general, equilibrium is non-unique, and increasing connectivity in the network may lead to less sharing. In a mean-field approximation to the model, we fully characterize the set of equilibria. As the interaction rate increases, the informed portion of the population in the maximal equilibrium exhibits a discontinuous jump upwards.

1 Introduction

With the advance of technology, from smart phones to the proliferation of social networking sites, sharing information is easier than ever. In principle, any piece of information can find its way to everyone who would benefit from knowing, but incentives often stand in the way. Say you hear about a job opportunity that is a good match for your skills, but the field is highly competitive. If you tell your friends, the low cost of communication works against you as they tell theirs, and so on, and your competition gets stiffer. The prospective employer is probably happy with a larger selection of applicants, but your chances of getting the job have fallen. If you want the job, you have a reason to keep your knowledge of the position to yourself. Awareness of the job opportunity is an example of rival information: information that becomes less valuable as more people learn it.

Information rivalry arises whenever information confers a competitive edge; it appears in contexts ranging from auctions to stock trading to innovation contests. Despite the clear incentive to keep such information secret, people regularly share. Employees learn about jobs

*Microsoft Research, New England

†Stern School of Business, New York University – esadler@stern.nyu.edu

through social ties (Marmaros and Sacerdote, 2002), and investment professionals share tips with each other—both legally and otherwise (Cohen et al., 2008). The DARPA red balloon challenge is a striking recent example of a contest in which word-of-mouth recruiting played a central role. Our reasons for sharing this kind information are often altruistic—we wish our friends to benefit from it—but we may also anticipate some reciprocity through future favors. In some cases, explicit incentives motivate sharing; in the red balloon challenge, the winning team made payments to those who recruited friends who found balloons. In each of these examples, there is a local incentive to share with immediate friends, but we would prefer that the information spreads no further.

We seek an understanding of what inclines individuals to share rival information. What trade offs do people make? How far does the information spread? How can we encourage more sharing? Our focus is a game in which local sharing incentives determine individual behavior within a broader network of friendships and social ties. Two important effects emerge: the gag effect and the blabbermouth effect. The gag effect says that if opportunities to share information are scarce, then individuals are more likely to share when they get the chance. Since it is difficult for a friend to spread the information further, you may as well let them in on the secret. If opportunities to share are common, the blabbermouth effect says individuals are also more likely to share. If word travels fast, a friend will find out anyway, and you may as well be the one to tell them. Each of these effects is always present to some extent, depending on the structure of the network and how often neighbors interact. Importantly, since these effects operate at different extremes, best replies are non-monotonic in others' behavior. When neither one dominates, individuals face a complex trade off.

We first demonstrate the two effects through a pair of examples. The first highlights the blabbermouth effect and the stability of sharing behavior when interactions are frequent. The second highlights the gag effect and shows that the presence of central individuals may lead to less overall sharing: equilibrium sharing behavior is non-monotonic in network connectivity. Theorems 1 and 2 generalize our basic insights to arbitrary networks. We vary a parameter that governs the interaction rate between any two neighbors. As this parameter becomes very small, the gag effect drives individual behavior, and as it becomes large, the blabbermouth effect dominates. At either extreme, players are happy to share information with one another at any opportunity.

In between is where rivalry can stifle the spread of information, and we introduce a mean-field approximation to study this more difficult case. Theorem 4 shows that we can often find an equilibrium in which players share at a positive fraction of their opportunities, even if we cannot sustain complete sharing. The informed portion of the population follows a logistic curve over time, and the partial-sharing equilibrium corresponds to a unique point at which this curve becomes sufficiently steep. This implies that the final informed fraction of the population is insensitive to small changes in our parameters: given a few more or a few less opportunities to share, players adjust their behavior in equilibrium so that the final outcome remains unchanged. Surpassing this aggregate level of dissemination requires reaching critical parameter thresholds at which complete sharing becomes an equilibrium.

Many authors study the spread of information through social networks. The social learning literature is increasingly focused on how network structure influences the spread of information, with numerous papers taking both Bayesian (Acemoglu et al., 2011; Lobel and Sadler, 2014) and more heuristic (Golub and Jackson, 2010, 2012; Jadbabaie et al., 2012) approaches to learning. Several empirical studies of innovation adoption highlight the role of information obtained through social ties (Bandiera and Rasul, 2006; Conley and Udry, 2010). Making information rival introduces a novel strategic element that is absent in much of this literature.

The spread of rival information has only recently piqued the interest of economists. Two efforts stand out for considering rivalry as a novel element in shared information. Banerjee et al. (2012) conducted an experiment in rural Indian villages in which villagers shared information about an opportunity to participate in a study with a limited number of slots. Not only did the villagers share their knowledge of this opportunity with others, there is evidence that they strategically selected with whom to share. Although our game captures fewer details about individual actions than this study, we provide a starting point to think more rigorously about the strategic effects at play.

Closer in spirit to our work, Lee et al. (2011) analyze a game in which players sequentially choose whether to share a piece of rival information with their friends. Our approach differs in two important ways. First, instead of acting in sequence, our players have opportunities to share at randomly determined times; this reflects social interactions that occur in a less structured environment. Second, Lee et al. suppose that the motivation to share information with friends is purely altruistic: a player is happy if her friends know the information regardless of how they obtain it. We assume that a player only benefits from informed friends if *she was the one who informed them.* This captures situations in which expected reciprocity, or an explicit incentive, is responsible for sharing. Our approach facilitates studying a different set of questions, focusing on how interaction patterns influence strategic sharing behavior.

Our findings also contribute to the recent literature on innovation contests. The growth of online platforms like Innocentive has led to more open contests for research and development purposes, and social scientists are exploring various ways to incentivize efficient effort. In these contests, there are at least two types of rival information present. Awareness of the contest is rival because having more competitors reduces one's chance of winning, and knowledge of intermediate findings or dead ends is rival because this could help competitors succeed more quickly. Current research analyzes inefficiencies arising from the second type of rival information (Akcigit and Liu, 2014) and considers centralized disclosure policies to produce better outcomes (Hallac et al., 2014). Our work complements these efforts, suggesting ways to encourage decentralized information disclosure and recruitment through local incentives to share.

2 The Information Sharing Game

There are N players in a population, and a random set is endowed with a piece of rival information; players must decide whether and to what extent they will share the information if they obtain it. We use a directed graph G to represent the network of connections between players. If there is a directed edge from player i to player j, we say that j is a neighbor of i. A player with the information can only share with neighbors.

Players have limited information about the social network. Players know their own degrees—how many neighbors they have—and they share a common prior over possible graphs. This prior is invariant to permutations of player indices, and players receive no information about neighbors' network positions. Hence, all players with a given degree have the same beliefs about their neighbors, and any two neighbors of a particular player are indisitinguishable to that player.

The game takes place in continuous time, lasting from time $t = 0$ until time $t = 1$. Nature first selects k players uniformly at random who are endowed with the information. During play, each edge in the graph G is active at random arrival times that occur with Poisson rate $\lambda > 0$. When the edge from i to j is active, player i has an opportunity to share information with player j if she has information. Each player i chooses a single probability $p_i \in [0, 1]$ with which she shares information at any opportunity. Hence, if player i is informed, any edge leading away from player i transmits information at a Poisson rate λp_i. A strategy profile σ for the players induces a probability distribution over the number who are informed at time $t = 1$.

We have deliberately simplified the set of available strategies. A player makes a single decision about how talkative she will be when she runs into friends, and she never revisits this decision to condition on timing, the identity of the friend, or how many people she has already informed. This represents an uncertain situation in which decisions during play are made quickly with a minimal allocation of mental resources. We can also imagine that the players do not precisely know the start and end points of the game, or when the information will cease to have value, and players do not perfectly recall how they obtained the information. This assumption will simplify our analysis when neither the gag effect nor the blabbermouth effect dominates player decisions. For the extreme cases in which one effect is dominant, our results readily extend with little additional effort.

Payoffs are realized at the conclusion of the game based on how many people are informed and who gave information to whom. The value of the information is $\pi(q)$, where $q \in [0, 1]$ is the informed fraction of the population at the end of the game. We assume $\pi(q) > 0$ and $\pi'(q) < 0$, indicating that information is valuable, but the value declines as more people become informed. In addition to the direct value of the information, an informed player receives a payoff $c\pi(q)$ with $c > 0$ for each neighbor who first learned the information from her. Note that the payoff from informing a neighbor is proportional to benefit the neighbor derives. Importantly, there are no additional payoffs for sharing information with a neighbor who is already informed.

A strategy profile σ, together with the common prior on G and common knowledge of k and λ, induces a distribution over outcomes. We consider Bayes-Nash equilibria in symmetric strategies, meaning that we can express a player's strategy as a function of her degree. We write $\sigma(d) : \mathbb{N} \to [0, 1]$ for a symmetric strategy profile, and we write (p_i, σ_{-i}) for a profile in which player i makes a unilateral deviation to p_i. Let Y_i denote the number of players who first receive the information from player i. If the other players adopt the profile σ, player i's expected payoff from choosing p_i is

$$u_\sigma(p_i) = \mathbb{E}_{(p_i,\sigma_{-i})}\left[\pi(q)\left(1 + cY_i\right)\right].$$

A strategy profile σ is an equilibrium if for each player i and each $p \in [0,1]$ we have

$$u_\sigma\left(\sigma(d_i)\right) \geq u_\sigma(p),$$

where d_i is player i's degree.

3 Examples

A pair of examples offers intuition for equilibrium outcomes, illustrating our two strategic effects. Consider first the game played in a complete network with $N = 3$ and $k = 2$. For notational simplicity, let $\pi_2 = \pi(2/3)$ and $\pi_3 = \pi(1)$ denote the value of the information when two and three players are informed respectively. This network allows the simplest possible demonstration of the blabbermouth effect. There is only one player who lacks the information, and an endowed player's decision to inform will depend on her expectations about the other endowed player's propensity to share.

Two distinct equilibria arise. First, players may refrain from sharing information at all. This is an equilibrium whenever $\pi_2 \geq (1 + c)\pi_3$, or equivalently

$$c \leq \frac{\pi_2 - \pi_3}{\pi_3}.$$

For sufficiently low c, it makes no sense to inform the third player if they are likely to remain uninformed otherwise. In this case, the endowed players can sustain secrecy in equilibrium. Another possibility is that players share as much as possible, choosing $p = 1$. The existence of this equilibrium will depend upon both the incentive to share c and the interaction rate λ because this rate affects the likelihood that the unendowed player learns the information from the other endowed player.

Since the unendowed player has no impact on the outcome, we can condition our analysis on a player endowed with the information. If the other endowed player shares at every chance, the expected utility from choosing sharing propensity p is

$$\pi_2 e^{-\lambda(1+p)} + \pi_3(1 - e^{-\lambda(1+p)})\left(1 + \frac{pc}{1+p}\right).$$

The first derivative of the payoff with respect to p is

$$-\lambda e^{-\lambda(1+p)}(\pi_2 - \pi_3) + \frac{c\pi_3}{1+p}\left(\lambda p e^{-\lambda(1+p)} + \frac{1 - e^{-\lambda(1+p)}}{1+p}\right).$$

Observe that for large λ, this approaches $\frac{c\pi_3}{(1+p)^2} > 0$, so the derivative is always positive for sufficiently high λ. Hence, for arbitrarily small c, we can find a sufficiently high λ such that $p = 1$ is an equilibrium.

Given a high λ and expectations that the other player shares, the unendowed player likely becomes informed regardless of the reply. Choosing $p = 1$ at least ensures an endowed player is first to inform with probability $\frac{1}{2}$. This shows the blabbermouth effect at work: with a high enough λ there is no longer a question whether the third player learns the information. There is only a question of who tells the third player.

We can make a stronger statement. If we believe the other endowed player has sharing propensity $\hat{p}$, then the corresponding first derivative of the payoff function is

$$-\lambda e^{-\lambda(\hat{p}+p)}(\pi_2 - \pi_3) + \frac{c\pi_3}{1+p}\left(\lambda p e^{-\lambda(\hat{p}+p)} + \frac{1 - e^{-\lambda(\hat{p}+p)}}{1+p}\right).$$

For any fixed $\hat{p} > 0$, this is always positive for sufficiently high λ, so choosing $p = 1$ is the unique best reply when λ is large. Alternatively, there is some threshold p^* such that $p = 1$ is a unique best reply whenever $\hat{p} > p^*$, and the threshold p^* converges to zero as λ approaches infinity. This suggests that in environments with frequent opportunities to share information, the full sharing equilibrium is particularly stable. Even a small risk of leaking information in any interaction makes it likely that the third player eventually learns, and the best reply is to share first if possible.

A second example reveals nuanced interactions between network structure and equilibrium strategies. Consider the networks depicted in Figure 1, and suppose one player is endowed with information uniformly at random. Network (a) comprises a ring in which each player has a single neighbor. Fix c and λ, and suppose there are N players in the ring. Given a symmetric strategy profile with common sharing probability p_σ, the number of players who become informed follows a truncated Poisson distribution with parameter λp_σ. For large enough N, the expected informed fraction at the end of the game is negligible regardless of player strategies, so $p_\sigma = 1$ is the unique equilibrium.

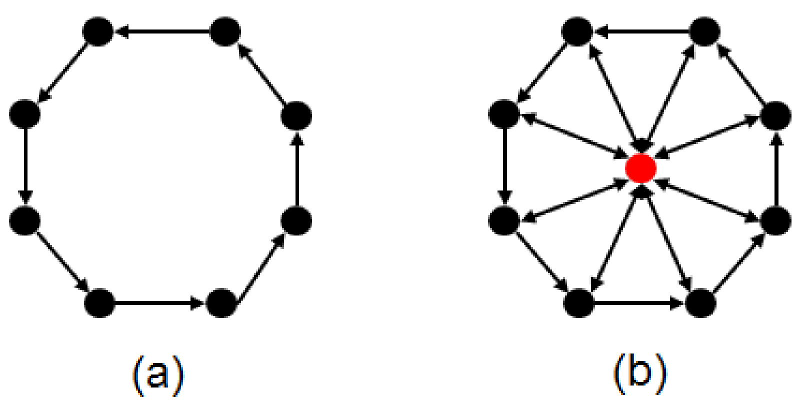

Figure 1: The networks in example 2

In network (b), we have added a central player with undirected edges to all players in the ring. Conditional on becoming informed, the central player expects she will be first to inform a positive fraction of the population if she decides to share information. Consider a fixed, bounded $\pi(q)$. For large N, the payoff from informing many neighbors will dominate the lost value of the information, so the central player must choose $p = 1$. Knowing this, the peripheral players may choose not to share any information because this risks the central player finding out and spreading it to a strictly positive fraction of the population.

To see this, suppose a peripheral player i expects others to choose $p = 0$. Conditional on becoming informed, player i is either endowed with the information or informed by the central player. We use P_i to denote the probability player i is endowed conditional on becoming informed, and note that $P_i > \frac{1}{2}$. Let $\pi_e < \pi(1/N)$ denote the expected value of $\pi(q)$ if the central player is endowed with the information and no peripheral players share. If player i chooses $p = 0$, then conditional on becoming informed, her payoff is

$$P_i\pi(1/N) + (1 - P_i)\pi_e.$$

Now, let $\tilde{\pi}_e$ denote a hypothetical benchmark, representing the expected value of $\pi(q)$ if the central player were endowed with the information at a uniformly drawn random time between 0 and 1. If player i is endowed with the information and chooses some $p_i > 0$, then the uniform distribution first order stochastically dominates the distribution of times at which the central player learns the information, conditional on player i informing the central player. Hence, the value $\tilde{\pi}_e < \pi(1/N)$ is an upper bound on the expected value of the information, conditional on player i informing the central player.

The probability of informing the central player is $1 - e^{-\lambda p_i}$. Conditional on becoming informed, player i's payoff is at most

$$\begin{aligned} P_i &\left[(1 - e^{-\lambda p_i})\tilde{\pi}_e \left(1 + c(2 - e^{-\lambda p_i})\right) + e^{-\lambda p_i}\pi(1/N)\left(1 + c(1 - e^{-\lambda p_i})\right)\right] \\ &+ (1 - P_i)\pi_e(1 + c(1 - e^{-\lambda p_i})). \end{aligned}$$

Here, the first term bounds the payoff conditional on being endowed with the information, and the second term bounds the payoff conditional on being informed by the central player. The payoff from choosing some $p_i > 0$ less the payoff from choosing $p_i = 0$ is then at most

$$P_i(1 - e^{-\lambda p_i})\left[\tilde{\pi}_e\left(1 + c(2 - e^{-\lambda p_i})\right) - \pi(1/N)(1 - ce^{-\lambda p_i})\right] + (1 - P_i)c\pi_e(1 - e^{-\lambda p_i}),$$

and since $P_i > \frac{1}{2}$, this is less than zero if

$$\tilde{\pi}_e\left(1 + c(2 - e^{-\lambda p_i})\right) - \pi(1/N)(1 - ce^{-\lambda p_i}) + c\pi_e < 0.$$

If $c < \frac{\pi(1/N) - \tilde{\pi}_e}{\pi(1/N) + \pi_e + \tilde{\pi}_e}$, this holds for any value of p_i, implying that $p_i = 0$ is the unique best reply.

Adding connectivity or reducing path lengths between players need not increase the spread of information in equilibrium, and it may lead to a decrease. In this example, the possibility of sharing with a highly connected individual discourages the peripheral players from sharing at all when they would otherwise prefer to share with their neighbors. If peripheral players could identify the central player in this example and condition sharing decisions according to which link is active, the presence of the central player would still accomplish little. Equilibrium would involve sharing with peripheral players only, and outcomes would appear similar to those in network (a). This demonstrates the gag effect: players share more freely in a sparse network or with a low interaction rate λ, knowing that information is unlikely to spread too much further.

4 Extreme Equilibria

Many insights from our examples readily extend to general networks. The Theorems in this section deal with extreme cases as the interaction rate λ becomes very small or very large, emphasizing the gag effect and the blabbermouth effect respectively.

For small values of λ, a player knows that if she has an opportunity to share, it is likely the only such opportunity. The gag effect turns the game into a simple individual decision problem, and the direct incentive to share c determines a unique equilibrium.

Theorem 1 (Gag Effect)**.** *Suppose that*

$$c > \frac{\pi\left(\frac{k}{N}\right) - \pi\left(\frac{k+1}{N}\right)}{\pi\left(\frac{k+1}{N}\right)}.$$

There is a threshold $\underline{\lambda}(c)$ *such that* $\sigma(d) = 1$ *for all* d *is the unique equilibrium whenever* $\lambda < \underline{\lambda}(c)$.

Suppose that

$$c < \frac{\pi\left(\frac{k}{N}\right) - \pi\left(\frac{k+1}{N}\right)}{\pi\left(\frac{k+1}{N}\right)}.$$

There is a threshold $\underline{\lambda}(c)$ *such that* $\sigma(d) = 0$ *for all* d *is the unique equilibrium whenever* $\lambda < \underline{\lambda}(c)$.

Proof. Consider a player's expected payoff conditional on becoming informed and conditional on having at least one opportunity to share with an uninformed neighbor. For sufficiently small λ, the conditional probability of having a second opportunity to share, or that a neighbor subsequently has an opportunity to share, is negligible. Moreover, conditional on being informed at all, it is almost certain that the player was endowed with the information. Given an $\epsilon > 0$, we can find λ_ϵ sufficiently small so that the conditional payoff from sharing is at least

$$\pi\left(\frac{k+1}{N}\right)(1+c) - \epsilon$$

whenever $\lambda \leq \lambda_\epsilon$, while the conditional payoff from not sharing is at most $\pi\left(\frac{k}{N}\right)$. If we have

$$c > \frac{\pi\left(\frac{k}{N}\right) - \pi\left(\frac{k+1}{N}\right)}{\pi\left(\frac{k+1}{N}\right)},$$

then we can find ϵ small enough so that the payoff of sharing is strictly larger than the payoff of not sharing, which implies $p = 1$ is the unique optimal action. The second part follows from an analogous argument. □

When the gag effect dominates decision making, we get a strong prediction involving a unique equilibrium in dominant strategies. Players can consider only whether they would like to inform their immediate neighbors, ignoring the possibility of further transmission. This effect generally encourages information sharing, particularly when the network is large relative to the number of endowed players. However, this prediction depends on a sharing incentive that is not too small. If c is negligible in size, we get an equally robust equilibrium with no information sharing.

High interaction rates also encourage information sharing, though for different reasons. If a player expects that everyone will eventually learn the information regardless of her sharing decision, then the game becomes a race to inform her friends first. The blabbermouth effect can operate no matter how small the incentive c is, but we no longer have a unique equilibrium.

Theorem 2 (Blabbermouth Effect)**.** *Suppose $k \geq 2$ and G is biconnected with probability one. For any $c > 0$ and any $p \in (0, 1]$, there exists a threshold $\overline{\lambda}(c, p)$ such that choosing sharing propensity 1 is a best reply for any player whenever $\lambda \geq \overline{\lambda}(c, p)$ and others share at a rate at least p.*

In particular, for any c there exists a threshold $\overline{\lambda}(c) = \overline{\lambda}(c, 1)$ such that $\sigma(d) = 1$ for all d is an equilibrium whenever $\lambda \geq \overline{\lambda}(c)$.

Proof. Fix a a strategy profile in which all players share at a rate at least p, and consider player i with expected payoff $u_\sigma(p_i)$ as a function of p_i. First, we show that for any $\epsilon > 0$ there exists λ_ϵ such that $p_i = 1$ provides a strictly higher payoff than any $p_i \in [0, 1-\epsilon]$. Suppose player i has a neighbor j and player i is endowed with the information. Now consider the following event: player i is first to interact with j, and after this, a chain bypassing player i provides j with the information before i has a second chance. This event has some positive probability $P > 0$; moreover this probability is bounded above 0 for all sufficiently high λ since it depends only on the order of arrivals. If player i adopts $p_i = 1$ instead of some $p_i \in [0, 1-\epsilon]$, then the probability of being first to inform the neighbor increases by at least $\frac{Pk\epsilon}{N} > 0$, corresponding to a payoff increase of at least $c\pi(1)\frac{Pk\epsilon}{N} > 0$ in expectation. For sufficiently large λ, the impact of this strategy change on the expected value of $\pi(q)$ is effectively zero since all players are almost certain to eventually learn the information, proving the claim.

Now, we show that $u'_\sigma(p_i) > 0$ in a neighborhood of 1 for sufficiently large λ, completing the proof. We have

$$u'_\sigma(p_i) = \mathbb{E}_{(p_i, \sigma_{-i})}\left[\frac{\partial \pi(q)}{\partial p_i}(1 + Y_i c) + c\pi(q)\frac{\partial Y_i}{\partial p_i}\right].$$

The argument of the previous paragraph implies that the second term is bounded above zero. The first term is bounded below by

$$(1+d_i c)\left(\pi(1)-\pi(k/N)\right)\frac{\partial}{\partial p_i}\mathbb{P}(q=1\,|\,p_i).$$

As λ grows, this last derivative approaches zero for any positive p_i since the probability $\mathbb{P}(q=1\,|\,p_i)$ converges uniformly to 1. Consequently, the derivative becomes positive in a neighborhood of 1 for sufficiently high λ, completing the proof. □

In contrast with the gag effect, the blabbermouth effect can sustain an equilibrium with complete sharing for arbitrarily low values of the sharing incentive c. The expectation that others will share, a least a little, is important for the conclusion in Theorem 2. If no one else shares, then a player's decision problem boils down to choosing a distribution over how many neighbors to inform. If c is sufficiently low, for instance if

$$\pi\left(\frac{k}{N}\right) \geq \sup_{m\leq N-k}(1+mc)\pi\left(\frac{k+m}{N}\right),$$

then $\sigma(d)=0$ for all d is an equilibrium strategy profile for any value of λ. This highlights the importance of expectations about others' behavior: for a range of parameter values, both no sharing and complete sharing are equilibria.

As λ grows, the set of equilibria collapses to either one or two points, depending on the sharing incentive c, but even when two equilibria exist, we can make a strong case in favor of complete sharing. The first part of the theorem implies that the equilibrium with sharing is robust to small deviations by other players, and the no sharing equilibrium is quite fragile in comparison. If there is even a tiny chance of a leak in any interaction, a sufficiently high value of λ implies that $p=1$ is a unique best reply.

Our assumption of a biconnected network, and at least two initial seeds, is crucial to the existence of a full sharing equilibrium for arbitrarily low values of c. This condition implies that no single player can unilaterally prevent information from reaching some part of the network. Having a path around each player, and ample opportunity for information to flow through that path, is what guarantees that the blabbermouth effect will appear. Taken together, these observations suggest one reason why norms of sharing information can develop in large well-connected groups, even if the direct incentive to share is small.

If we face the problem of designing an explicit sharing incentive, Theorem 2 indicates that a very simple mechanism with small rewards can be effective. This contrasts with the more complicated winning strategy in the DARPA red balloon challenge, which utilized geometrically declining payments to all players in a chain of recruits. Papers studying this and related mechanisms make stronger behavioral predictions, but the results often depend on restrictive assumptions about the network structure (e.g. Babaioff et al., 2012). In particular, if the underlying graph is a uniform tree, and there is a single seed at the root, the more intricate mechanism induces a sharing equilibrium in dominant strategies. However, providing incentives is relatively expensive because individual players can single-handedly determine whether information reaches large parts of the network. By relaxing this assumption on the network, we uncover a novel strategic effect that makes it easier to sustain sharing behavior.

For high and low interaction rates, the set of equilibria is remarkably simple. For intermediate values, the problem becomes far more complex because an individual's sharing decision can have a material impact on the final informed fraction: neighbors will have time to relay information to others, and the probability that these individuals would otherwise remain uninformed is non-trivial. For many parameter values, we cannot sustain full sharing in equilibrium, but we can ask whether some positive level of sharing is sustainable and the extent to which information spreads in these equilibria. The next section introduces a mean-field approximation to better understand these more complex cases.

5 A Mean-Field Approach

The intuition for our mean-field results is easiest to understand visually. If a continuum of players share information at a fixed rate in an infinitely large random graph, we can represent the informed fraction over time using a logistic curve. Equivalently, the final informed fraction as a function of the interaction rate λ, or of the equilibrium sharing propensity p_σ, also follows a logistic curve. Figure 2 shows how different aggregate levels of sharing affect an individual player's incentives to share. If the choices of the other players place the final informed fraction on a steep part of the logistic curve, then one player's incremental decision to share has a relatively large impact on the final informed fraction. Conversely, if the sharing of other players puts us on a flatter part of the logistic curve, an individual's decision to share has relatively little impact on the informed fraction. This also lets us visually represent the gag and blabbermouth effects. The gag effect corresponds to being on the flat part at the lower end of the logistic curve, while the blabbermouth effect corresponds to being on the flat part at the higher end of the curve.

The interaction rate λ determines how far up the curve we go when players choose $p = 1$. For a low interaction rate, we always remain on a flatter part of the curve, and a player's incremental impact on the final informed fraction will never dissuade sharing. As we increase the interaction rate, we approach steeper parts of the curve until we hit the critical threshold shown in Figure 3. Past this point, individual players would unilaterally choose to reduce their own level of sharing in order to prevent the information from spreading. For a range of interaction rates, the highest informed fraction that can occur in equilibrium is exactly the level of this threshold: the equilibrium sharing propensity p_σ drops with λ, and the outcome no longer depends on the interaction rate in this range. As we continue increasing λ, we reach a point of discontinuity where full sharing once again becomes an equilibrium. With a high interaction rate, the expectation that others will share puts us on the flat part of the curve at the high end, so individual decisions no longer have a large impact on the informed fraction.

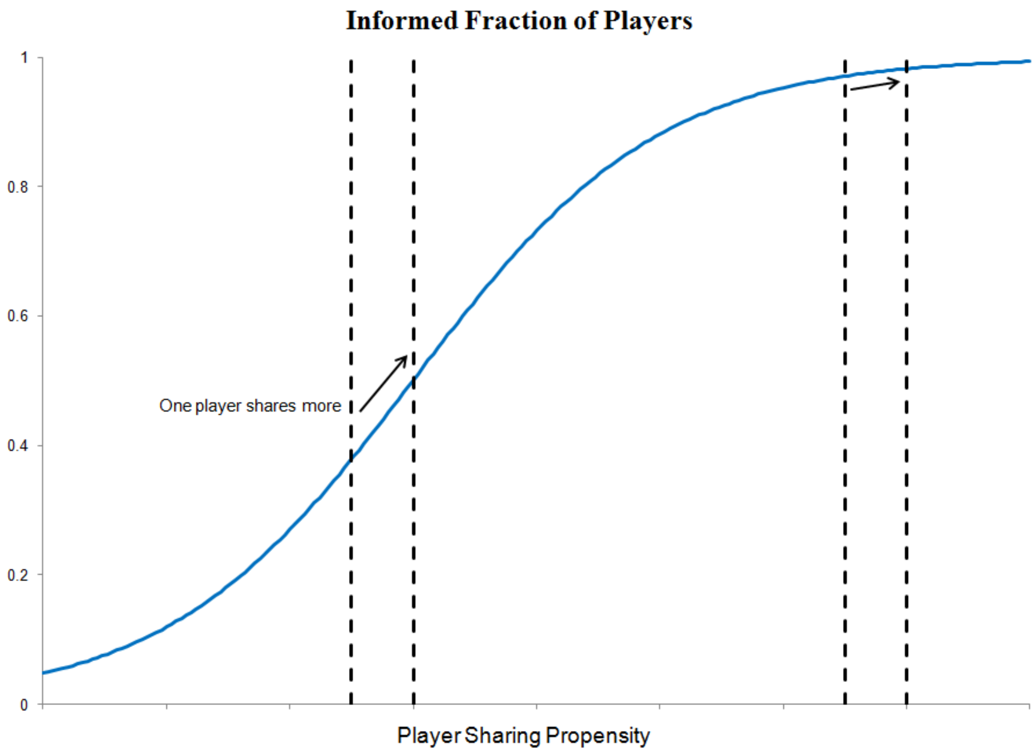

Figure 2: One player's impact on the final informed fraction

5.1 The Mean-Field Game

We still think of a population comprising N players, but we shall represent the population as a unit mass in which the action of one player constitutes the action of a mass $\frac{1}{N}$ of players in the continuum. We can view this either as an approximate way to model a large population or as a simplified representation of the decision problem adopted by boundedly rational players. An initial fraction $q_0 \geq \frac{1}{N}$ hold the information, all players have degree μ, and we imagine the network as a large random regular graph. Time is continuous, and sharing opportunities along each link arrive at random times with Poisson rate $\lambda > 0$. Players choose a fixed probability p with which they share information at each arrival, conditional on being informed. Since all players have the same degree, a single probability p_σ characterizes a symmetric strategy profile. From here on we assume that the value of the information is $\pi(q) = \frac{\beta}{q}$ for some $\beta > 0$.

The game ends after one unit of time. If a player i chooses the equilibrium sharing propensity p_σ, the informed portion of the population evolves as

$$\frac{\mathrm{d}q}{\mathrm{d}t} = \lambda \mu p_\sigma q(1-q).$$

That is, the informed fraction increases at a rate proportional to the arrival rate λ and the network density μ. Informed players account for a fraction q of sharing opportunities, but

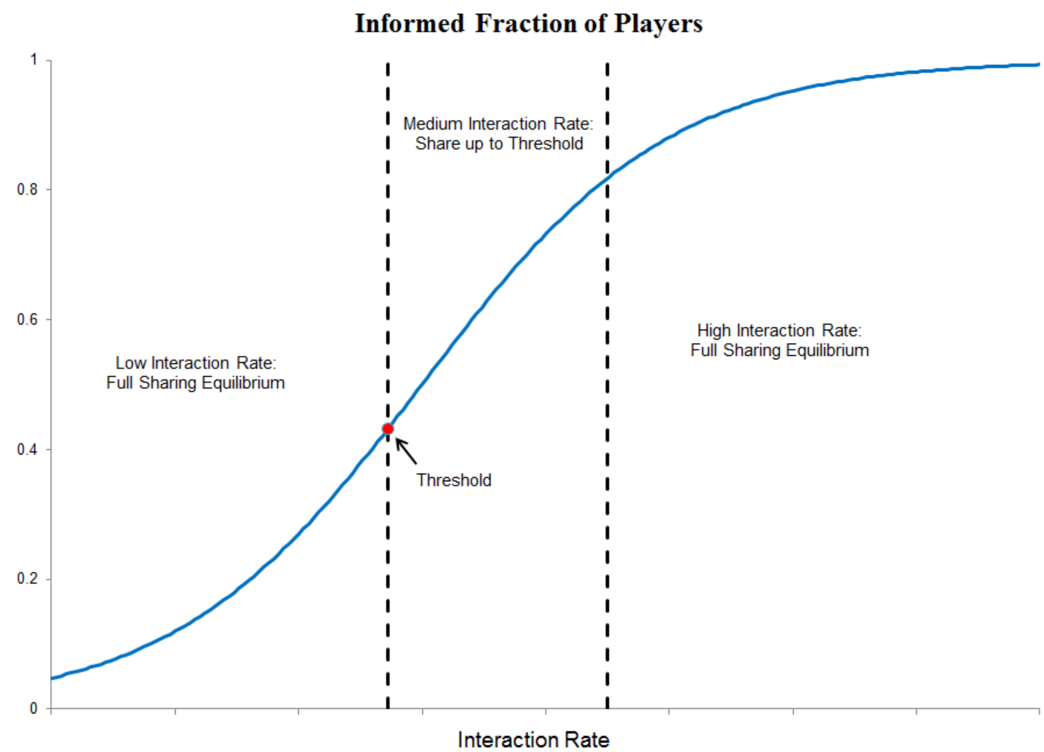

Figure 3: The critical threshold

of the neighbors with whom they share, only a fraction $1 - q$ are uninformed. We recognize a logistic equation, giving the informed fraction at time t as

$$q_t = \frac{q_0 e^{\lambda \mu p_\sigma t}}{1 + q_0 (e^{\lambda \mu p_\sigma t} - 1)}.$$

Hence, the final informed fraction is

$$q_1 = \frac{q_0 e^{\lambda \mu p_\sigma}}{1 + q_0 (e^{\lambda \mu p_\sigma} - 1)}.$$

If player i deviates from p_σ, playing p instead, then the informed fraction evolves as

$$\frac{\mathrm{d}q}{\mathrm{d}t} = \lambda \mu \left(p_\sigma + \frac{1}{N}(p - p_\sigma) \right) q(1 - q),$$

with a corresponding final informed fraction

$$q(p) = \frac{q_0 e^{\lambda \mu \left(p_\sigma + \frac{1}{N}(p - p_\sigma) \right)}}{1 + q_0 (e^{\lambda \mu \left(p_\sigma + \frac{1}{N}(p - p_\sigma) \right)} - 1)}.$$

To complete the mean-field model, we specify the expected payoff of a player conditional on becoming informed. Player i expects a payoff of

$$u(p) = \frac{\beta}{q(p)} \left(1 + c\mu \frac{(q(p) - q_0)p}{q(p)\left((\mu - 1)p_\sigma + p\right)}\right). \tag{1}$$

The value $\frac{(q(p)-q_0)p}{q(p)((\mu-1)p_\sigma+p)}$ represents the probability that a neighbor is first informed by player i, conditional on player i becoming informed. Let A denote the event that the neighbor is first informed by player i and B the event that player i becomes informed. We then have

$$\mathbb{P}(A \,|\, B) = \frac{\mathbb{P}(A, B)}{\mathbb{P}(B)} = \frac{\mathbb{P}(A)}{\mathbb{P}(B)}.$$

Taking the effect of her own choice as given, player i becomes informed with probability $q(p)$. A random neighbor learns the information from another player with probability $q(p) - q_0$, and $\frac{p}{(\mu-1)p_\sigma+p}$ is player i's portion of sharing opportunities with a random neighbor.

We first replicate and enrich the basic findings from the discrete model. The same intuition on extreme equilibria carries over into the mean-field framework. Sharing equilibria are easier to sustain when there are either very many or very few sharing opportunities and when a larger initial fraction q_0 is informed. The following technical lemma provides useful properties that allow us to state more precise comparative statics.

Lemma 1. *Given a strategy profile p_σ, a player's best reply p is increasing in q_0 and c. For sufficiently large N, the best reply is also increasing in N.*

Proof. Consider a monotone transformation of the utility function $u(p)$ defined as

$$v(p) = \frac{N}{(1 - q_0)} \ln u(p).$$

Maximizing u is equivalent to maximizing v, so we can treat v as the player's utility. Let $\alpha = \lambda\mu \left(p_\sigma + \frac{1}{N}(p - p_\sigma)\right)$. We compute

$$\begin{aligned} v'(p) &= \frac{N}{(1 - q_0)} \left[\frac{-q'(p)}{q(p)} + c\mu \frac{q_0 p \frac{q'(p)}{q(p)} + (q(p) - q_0)\frac{(\mu-1)p_\sigma}{(\mu-1)p_\sigma+p}}{q(p)\left((\mu - 1)p_\sigma + p\right) + c\mu p\left(q(p) - q_0\right)}\right] \\ &= \frac{-\lambda\mu}{1 + q_0(e^\alpha - 1)} + c\mu \frac{\lambda\mu p + (e^\alpha - 1)N\frac{(\mu-1)p_\sigma}{(\mu-1)p_\sigma+p}}{e^\alpha\left((\mu - 1)p_\sigma + p\right) + c\mu p(1 - q_0)(e^\alpha - 1)}. \end{aligned} \tag{2}$$

The lemma follows because v' is strictly increasing in q_0 and c, and for sufficiently large N, the derivative v' is increasing in N.

□

The intuition of the gag and blabbermouth effects applies just as before, and we obtain an analog of Theorems 1 and 2.

Theorem 3. *The strategy $p_\sigma = 0$ is an equilibrium if and only if $c \leq \frac{1}{\mu}$. If $c > \frac{1}{\mu}$, there is a threshold $\underline{\lambda}(c)$ such that $p_\sigma = 1$ is the unique equilibrium whenever $\lambda < \underline{\lambda}(c)$. The threshold $\underline{\lambda}$ is increasing in c and in q_0.*

For any $c > 0$ and any $p_\sigma \in (0, 1]$, there exists a threshold $\overline{\lambda}(c, p)$ such that choosing $p = 1$ is a best reply for any player whenever $\lambda \geq \overline{\lambda}(c, p)$ and others share at the rate p_σ. The threshold $\overline{\lambda}$ is decreasing in c and q_0. Moreover, the thresold $\overline{\lambda}$ converges to zero as either μ or N approaches infinity.

Proof. For the first part, taking $p_\sigma = 0$ in equation (2) gives

$$v'(p) = \frac{-\lambda\mu}{1 + q_0(e^\alpha - 1)} + c\mu\frac{\lambda\mu}{e^\alpha + c\mu(1 - q_0)(e^\alpha - 1)}.$$

This is positive if and only if

$$2q_0 - 1 + \frac{2(1 - q_0)}{e^\alpha} \geq \frac{1}{c\mu}.$$

Note the left hand side is at most 1, with equality if and only if $p = 0$. If $c\mu > 1$, the right hand side is strictly less than 1, so $p = 0$ is not a best reply, and $p_\sigma = 0$ is not an equilibrium. Continuity of v' ensures the existence of the threshold $\underline{\lambda}$, and the comparative statics follow immediately from Lemma 1. Alternatively, if $c\mu \leq 1$ the right hand side is at least 1, and the derivative $v'(p)$ is non-positive for all p, implying $p_\sigma = 0$ is an equilibrium.

Taking $p_\sigma = 1$ in equation (2), we have

$$v'(p) = \frac{-\lambda\mu}{1 + q_0(e^\alpha - 1)} + c\mu\frac{\lambda\mu p + (e^\alpha - 1)N\frac{\mu - 1}{\mu - 1 + p}}{e^\alpha(\mu - 1 + p) + c\mu p(1 - q_0)(e^\alpha - 1)}. \tag{3}$$

Note that the negative term in $v'(p)$ converges to zero as the product $\lambda\mu$ approaches infinity, while the positive term converges to

$$\frac{c\mu N\frac{\mu - 1}{\mu - 1 + p}}{\mu - 1 + p + c\mu p(1 - q_0)},$$

which is bounded above zero. Hence, for any fixed c, the profile $p_\sigma = 1$ is an equilibrium for sufficiently high λ or μ. Similarly, for large enough N, the positive term dominates, and again $p_\sigma = 1$ becomes an equilibrium for any fixed c. The comparative statics for the threshold $\overline{\lambda}$ again follow from Lemma 1. This limiting outcome is not specific to $p_\sigma = 1$, and in fact for any fixed $p_\sigma > 0$, the positive term in $v'(p)$ will dominate for sufficiently large λ, μ, or N, completing the proof.

□

The mean-field approach allows us to go a step beyond these results and characterize equilibria with p_σ strictly between zero and one. For intermediate ranges of c and λ, we can find an equilibrium with positive sharing even if there is a non-trivial tradeoff between the benefits of sharing with neighbors and the costs of information spreading further. Outside of these ranges, the set of equilibria collapses to the extreme points.

Theorem 4. *Suppose $q_0 \leq \frac{1}{2}$. There exists a c^* and thresholds $\underline{\lambda}(c)$ and $\overline{\lambda}(c)$ such that for $c \in \left(\frac{\mu}{\mu+N(\mu-1)}, c^*\right)$ and $\lambda \in \left(\underline{\lambda}(c), \overline{\lambda}(c)\right)$ there is a unique equilibrium with a positive level of sharing, taking $p_\sigma = \frac{\underline{\lambda}(c)}{\lambda}$. If $\lambda \leq \underline{\lambda}(c)$ or $\lambda \geq \overline{\lambda}(c)$, then $p_\sigma = 1$ is an equilibrium.*

If $c \leq \frac{\mu}{\mu+N(\mu-1)}$, only extreme equilibria exist; $p_\sigma = 1$ will be an equilibrium for sufficiently high λ. If $c \geq c^*$, only extreme equilibria exist, and $p_\sigma = 1$ is always an equilibrium.

Proof. See Appendix. □

If the benefit from sharing c is very low, sharing can only occur in equilibrium if the blabbermouth effect dominates decision making; if c is very high, then full sharing is always an equilibrium. In between, we find an interesting shift as λ increases. Below $\underline{\lambda}(c)$, the gag effect dominates and we can sustain sharing. This threshold marks where the logistic curve becomes steep enough that individual players will unilaterally reduce their own sharing behavior to keep the information from spreading. Past this point, the equilibrium p_σ shifts downwards to maintain the same level of spread we would have with $p_\sigma = 1$ and $\lambda = \underline{\lambda}(c)$. Eventually, we cross a second threshold, at which point the blabbermouth effect takes over, and full sharing is once again an equilibrium.

6 Final Remarks

Sharing rival information elicits a complex set of tradeoffs depending on the benefits from sharing, the frequency of interactions, and the structure of the social network. We identify two important strategic effects, working at opposite extremes, that influence decision making. When interactions are infrequent, the gag effect reduces the problem to a myopic decision. When interactions are highly frequent, the blabbermouth effect means we cannot keep secrets; the best reply is to share first. These two effects are highly robust to the model specification, and we can apply the intuition behind our results to generalizations both in the discrete and in the mean-field setting. For instance, if we model distinct identifiable groups of players, then depending on expectations and who the initial seeds are, we can find equilibria in which players share with one group but not another.

In the mean-field approximation, we find a sharp discontinuity in equilibrium behavior. As the interaction rate increases, the maximal equilibrium sharing propensity first decreases before jumping upwards. Such threshold phenomena are common in the mathematical literature on random graphs, for instance when looking at the level of connectivity that allows a giant component to emerge. However, rather than emerging from properties of a graph, the threshold phenomenon we find emerges from strategic considerations. The source of the discontinuity is distinct. There is a critical point at which the blabbermouth effect can take hold, but up until that point players have an incentive to reduce their sharing propensity as interactions become more frequent.

Understanding the factors that influence information sharing can inform not only contest design and recruitment efforts but also social policies more broadly. Our findings suggest that technologies that make communication easier can have a discontinuous impact on norms around information sharing. This insight may help make sense of a number of social phenomena such as changing attitudes, particularly on social media, towards privacy.

References

Acemoglu, Daron, Munther Dahleh, Ilan Lobel, and Asuman Ozdaglar (2011), "Bayesian Learning in Social Networks." *The Review of Economic Studies*, 78, 1201–1236.

Akcigit, Ufuk and Qingmin Liu (2014), "The Role of Information in Innovation and Competition." Working Paper.

Babaioff, Moshe, Shahar Dobzinski, Sigal Oren, and Aviv Zohar (2012), "On bitcoin and red balloons." In *ACM Conference on Electronic Commerce (EC'12)*, ACM.

Bandiera, Oriana and Imran Rasul (2006), "Social Networks and Technology Adoption in Northern Mozambique." *The Economic Journal*, 116, 869–902.

Banerjee, Abhijit, Emily Breza, Arun Chandrasekhar, Esther Duflo, and Matthew Jackson (2012), "Come Play With Me: Experimental Evidence of Information Diffusion about Rival Goods." Working Paper.

Cohen, Lauren, Andrea Frazzini, and Christopher Malloy (2008), "The Small World of Investing: Board Connections and Mutual Fund Returns." *Journal of Political Economy*, 116, 951–979.

Conley, Timothy and Christopher Udry (2010), "Learning About a New Technology: Pineapple in Ghana." *American Economic Review*, 100, 35–69.

Golub, Ben and Matthew Jackson (2010), "Naïve Learning in Social Networks and the Wisdom of Crowds." *American Economic Journal: Microeconomics*, 2, 112–149.

Golub, Ben and Matthew Jackson (2012), "How Homophily Affects the Speed of Learning and Best-Response Dynamics." *The Quarterly Journal of Economics*, 127, 1287–1338.

Hallac, Marina, Navin Kartik, and Qingmin Liu (2014), "Contests for Experimentation." Working Paper.

Jadbabaie, Ali, Pooya Molavi, Alvaro Sandroni, and Alireza Tahbaz-Salehi (2012), "Non-Bayesian Social Learning." *Games and Economic Behavior*, 76, 210–225.

Lee, Samuel, Petra Persson, and Nikita Roketskiy (2011), "Circles of Trust." Working Paper.

Lobel, Ilan and Evan Sadler (2014), "Information Diffusion in Networks through Social Learning." Forthcoming in Theoretical Economics.

Marmaros, David and Bruce Sacerdote (2002), "Peer and Social Networks in Job Search." *European Economic Review*, 46, 870–879.

A Appendix

Proof of Theorem 4

We introduce an alternate representation of the players' decision problem. Let $s_\sigma = \lambda p_\sigma$, and a player chooses r such that $p = r p_\sigma$. In equilibrium, the choice $r = 1$ must maximize utility. The parameter λ now corresponds to a constraint on the range of feasible values for s_σ and r: we have $s_\sigma \in [0, \lambda]$ and $r \in \left[0, \frac{\lambda}{s_\sigma}\right]$. We can now rewrite player utility as

$$\tilde{u}(s_\sigma, r) = \frac{\beta}{q(s_\sigma, r)} \left[1 + c\mu \left(1 - \frac{q_0}{q(s_\sigma, r)}\right) \frac{r}{\mu - 1 + r}\right].$$

Define the transformation

$$\tilde{v}(s_\sigma, r) = \frac{N}{\mu s_\sigma (1 - q_0)} \ln \tilde{u}(s_\sigma, r),$$

and note that the utility function $\tilde{v}$ induces the same best replies as the function $\tilde{u}$. We compute the partial derivative

$$\frac{\partial}{\partial r} \tilde{v}(s_\sigma, r) = \frac{-1}{1 + q_0(e^\alpha - 1)} + c \frac{\mu r + \frac{N(\mu-1)}{s_\sigma(\mu-1+r)}(e^\alpha - 1)}{e^\alpha(\mu - 1 + r) + c\mu(1 - q_0) r (e^\alpha - 1)},$$

where $\alpha = \frac{\mu s_\sigma}{N}(N - 1 + r)$. Note this derivative is positive for all sufficiently high s_σ and r. Central to our proof is the behavior of this partial derivative as we change s_σ.

Lemma 2. *Define the function $f_{s_\sigma}(r) = \frac{\partial}{\partial r}\tilde{v}(s_\sigma, r)$. We have for any $h > 0$*

$$f_{s_\sigma + h}(r) \geq f_{s_\sigma}\left(r + \frac{h(N - 1 + r)}{s_\sigma}\right).$$

That is, increasing the population sharing rate s_σ shifts the curve f left and upwards.

Proof. Observe that a small increase $h > 0$ in s_σ leads to an increase in the exponent α, which we could reverse through a corresponding decrease in r by the amount $\frac{h(N-1+r)}{s_\sigma}$. This produces a leftward shift in the curve. To finish the argument, we must show that a decrease in r, holding the exponent α constant, leads to an increase in $f_{s_\sigma}(r)$. That is, ignoring the dependence of α on r, we show

$$\frac{\partial}{\partial r} \frac{\mu r + \frac{N(\mu-1)}{s_\sigma(\mu-1+r)}(e^\alpha - 1)}{e^\alpha(\mu - 1 + r) + c\mu(1 - q_0) r (e^\alpha - 1)} < 0.$$

Applying the quotient rule, this is equivalent to

$$\mu(\mu - 1)e^\alpha - \frac{N(\mu - 1)^2 e^\alpha (e^\alpha - 1)}{s_\sigma(\mu - 1 + r)^2} - \frac{N(\mu - 1)(e^\alpha - 1)\left(e^\alpha + c\mu(1 - q_0)(e^\alpha - 1)\right)}{s_\sigma(\mu - 1 + r)} < 0.$$

Now note that $e^\alpha - 1 \geq \alpha = \frac{\mu s_\sigma(N-1+r)}{N}$, so the above inequality holds if

$$\mu(\mu-1)\left[e^\alpha - \frac{(\mu-1)(N-1+r)e^\alpha}{(\mu-1+r)^2} - \frac{N-1+r}{\mu-1+r}\left(e^\alpha + c\mu(1-q_0)(e^\alpha-1)\right)\right]$$
$$\leq \mu(\mu-1)e^\alpha\left(1 - \frac{N-1+r}{\mu-1+r}\right) < 0,$$

which is implied by $N > \mu$, proving the claim. □

Any equilibrium with $p_\sigma \in (0,1)$ corresponds to a zero at $f_{s_\sigma}(1)$. We compute

$$f_{s_\sigma}(1) = \frac{\partial}{\partial r}\tilde{v}(s_\sigma, 1) = \frac{-1}{1+q_0(e^{\tilde{\alpha}}-1)} + c\frac{1 + \frac{N(\mu-1)}{s_\sigma \mu^2}(e^{\tilde{\alpha}}-1)}{e^{\tilde{\alpha}} + c(1-q_0)(e^{\tilde{\alpha}}-1)},$$

where $\tilde{\alpha} = \mu s_\sigma$. Taking $y = (e^{\tilde{\alpha}} - 1)$, the zeros of this function are the same as the zeros of

$$\frac{-1}{1+q_0 y} + c\frac{1 + \frac{N(\mu-1)y}{\mu \ln(1+y)}}{1 + (1 + c(1-q_0))\, y},$$

which for $y > 0$ are the same as the zeros of

$$w(y) = cq_0 N\frac{\mu-1}{\mu}y^2 + cN\frac{\mu-1}{\mu}y - (1 + c(1-2q_0))\, y\ln(1+y) - (1-c)\ln(1+y).$$

We address each case of the theorem in turn. Suppose $c \leq \frac{\mu}{\mu+N(\mu-1)}$. We then have $w'(0) \leq 0$, $w''(0) < 0$, and $w'''(y) > 0$ for all $y \geq 0$. This implies that w has a single zero with $y > 0$ on the upward sloping part of the curve. Since the location of this zero decreases with c, this cannot be a local maximum; it is a local minimum. Figure 4 shows the functions w and $f_{s_\sigma}(r)$ for a small value of s_σ. As we increase s_σ and shift the curve $f_{s_\sigma}(r)$ left and up, the curve will hit zero at $r = 1$ exactly once on the upward sloping part of the curve. This implies there is no equilibrium with $p_\sigma \in (0,1)$ for this range of c: when λ is low, there will be no sharing, but for sufficiently high λ, the blabbermouth effect creates a full sharing equilibrium.

Now suppose $c > \frac{\mu}{\mu+N(\mu-1)}$. We then have $w'(0) > 0$, $w''(0) < 0$, and $w'''(y) > 0$ for all $y \geq 0$. For sufficiently large c, the function w is always non-negative; let c^* denote the smallest such c. For $c \in \left(\frac{\mu}{\mu+N(\mu-1)}, c^*\right)$, there are then exactly two zeros of w with $y > 0$. That $f_{s_\sigma}(r)$ is positive for sufficiently high r together with Lemma 2 implies that the first zero is a local maximum, and the second is a local minimum. Figure 5 shows the functions w and f_{s_σ} in this case. As we increase s_σ and shift $f_{s_\sigma}(r)$ up and to the left, the first zero crosses $r = 1$ on the downward sloping part of the curve, and the second crosses on the upward sloping part. If λ is too low to reach this first crossing, then $f_{s_\sigma}(r)$ is positive below $r = 1$, and full sharing is an equilibrium. Once λ becomes sufficienty large, this zero corresponds to the only equilibrium with a positive level of sharing as $f_{s_\sigma}(1)$ is negative for slightly higher s_σ. As λ becomes larger still, the blabbermouth effect again creates a full sharing equilibrium.

Finally, by definition, $w(y)$ is non-negative whenever $c \geq c^*$, and $f_{s_\sigma}(r)$ is likewise positive for all $r < 1$, so full sharing is an equilibrium for all values of λ. □

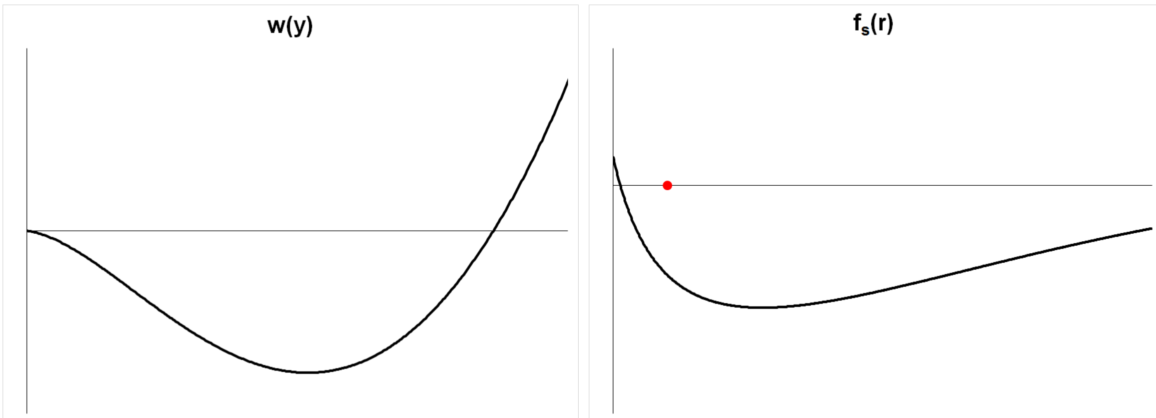

Figure 4: The case $c \leq \frac{\mu}{\mu+N(\mu-1)}$

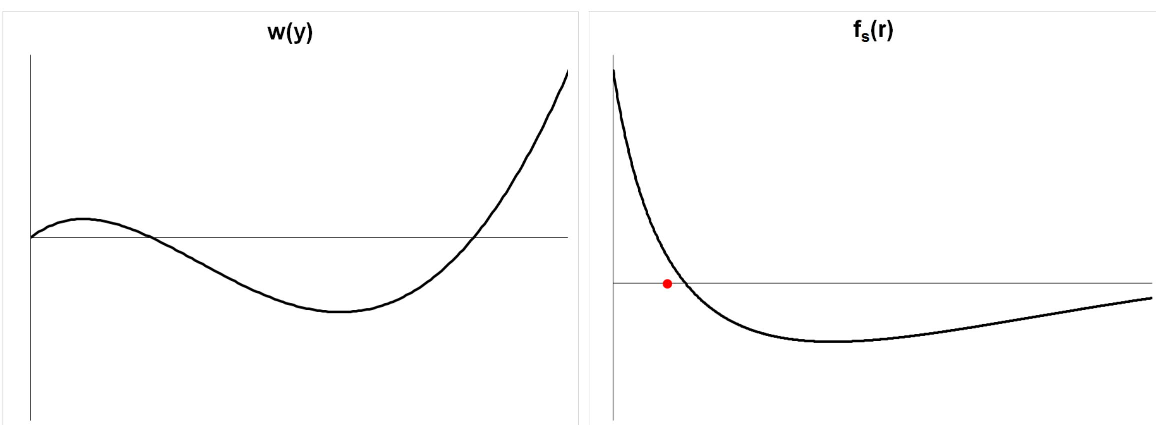

Figure 5: The case $c \in \left(\frac{\mu}{\mu+N(\mu-1)}, c^*\right)$