Learning to collaborate

Kevin J.S. Zollman

July 3, 2016

Abstract

This paper presents a rudimentary model of collaboration with the aim to understand the conditions under which groups of scientists will endogenously form optimal collaborative groups. By analyzing the model with computer simulations, I uncover three lessons for collaborative groups. First, in reducing the cost borne by scientists from collaborating, one benefits the members of the group. Second, increasing the number of potential collaborative partners benefits all those involved in a collaborative group. Finally and counter-intuitively, this model suggests that groups do better when scientists avoid experimenting with new collaborative interactions.

Collaboration involves tackling problems together. Different individuals might bring diverse perspectives to a problem, and by working together they come to solution that none would have reached alone. Collaborating comes with a cost, however. One must expend effort communicate one's approach to another. Collaborations require agreement about the strategies for tackling the problem, and one's collaborator might be difficult or unhelpful. As groups grow the possibility of other epistemic pathologies like group-think and collective ignorance arise.

Whether collaboration is helpful or harmful depends on how these costs and benefits are weighed. Much of the research on collaboration is focused on enumerating, theorizing, and comparing these various costs and benefits. Scholars normally focus on the benefits or harms that collaboration has on those directly involved in the collaborative effort (cf. Kerr and Tindale, 2004). But collaborations also create what economists call "externalities" for those outside of the collaboration. By collaborating with me, you work less with others. If I collaborate with you, you come to learn (at least in part) how I view the world. You might then share this – for better or worse – with future collaborators.

Dealing with externalities can be a tricky matter. Externalities create difficult social dilemmas, like the Prisoner's dilemma and tragedy of the commons, where the best thing for the group is inconsistent with individuals' self interest. Even in situations where optimality is consistent with individual choice, externalities can make achieving this good outcome complicated or nearly impossible.

In this paper, I utilize a rudimentary model to determine what types of collaborative exchanges would be optimal and under what conditions we should expect groups of scientists to endogenously form optimal collaborative groups. While this model does not represent every collaborative exchange, but it does represent some important aspects of collaboration. By analyzing this model, I find a few lessons for those who wish to maximize the benefits of collaboration. First, by reducing the cost borne by scientists from collaborating, one benefits the members of the group, albeit in a particular way. Second, enlarging the group of potential collaborative partners benefits all those involved in a collaborative group. Finally, and perhaps most counter-intuitively, groups do better when scientists have high inertia in choosing collaborative partners – i.e. groups do better when individual scientists don't try out new collaborative interactions too often.

The model presented in this paper fits within a larger literature of epistemic network models (for an overview, see Zollman, 2013). Collaboration is modeled as creating a social network. This model, like others in this literature, are highly idealized. As a result, the results are suggestive rather than definitive.

1 Modeling collaboration

While collaboration has many facets, the model I present focuses on collaborations with a one-way exchange of assistance. You come to collaborate with me and I help you to solve one of your problems, but you don't help me to solve any of mine. One can model other forms of collaboration as well. For instance, both Jackson and Wolinsky (1996) and Bala and Goyal (2000) develop models of collaborations that feature two-way collaboration.

The interaction I model is minimally collaboration. The exchange of information and assistance goes only one-way. As a result, this model lies at the border of what we might call "collaboration" and "consultation." By focusing on one-way collaboration, I do not mean to imply other forms of collaboration are non-existent or uninteresting. But already in this limited situation several interesting complexities arise.

In this paper I will combine, re-interpret, and extend two different models developed by Bala and Goyal (2000) and Hong and Page (2004). We begin with a fixed group of n *actors*. These actors represent any single individual or cooperatively acting group that is attempting to solve a problem. Actors might be an individual scientist interacting with other individual scientists. Or actors might be a scientific lab interacting with other scientific labs. The model need not be restricted to science, many other groups will fit the assumptions. Henceforth, I will refer to the individuals as "scientists," although the reader should remember that the interpretation of the model can be broad.

Each scientist faces a different problem with many solutions. This might be a high-level problem, like developing a novel theory or designing a complex experiment. The problem might be more simple and mundane, like attempting to prove a small theorem or fix a broken piece of machinery. Whatever its level of significance, the problem must have many solutions that range in quality from worthless to exceptional.

Although the problems faced by the scientists differ from one another, we will suppose they occupy the same field – each scientist can provide assistance on every other problem. Although they work in the same area, each scientist approaches problems in her own way. Each has a unique organization of the space of solutions. The scientists begin with a potential solution to the problem and then move to other solutions that occur to them if those solutions appear better.

Suppose, for instance, an ecologist – call him Carlos – confronts an odd behavior in the field. Carlos finds in the organism he studies, males exhibit bright colors. He wants to figure out which of the many potential explanations for this behavior is the correct one (cf. Maynard Smith and Harper, 2003). Potential experiments abound, and Carlos must try to figure out the one that would be best for uncovering this phenomena.

Carlos begins by considering experimental design u. He forms a judgment – I'll suppose a correct one – about how effective design u is likely to be at discovering the underlying mechanism of interest. While considering design u, Carlos can also imagine designs w and x as possible. Another ecologist, call her Julie, might have a different conceptual organization of the space of possible experiments. When Julie begins by considering u she conceives of designs y and z as possible. Carlos, if left to work on the problem by himself, would overlook y and z, while Julie would miss w and x.

Collaboration in this model is represented by scientists teaching one another their conceptual schemes. So, if Carlos goes to Julie for help, she can provide at least some information about how she conceives of the problem that might bring designs y and/or z to Carlos' attention.

In this model, gaining access to every additional conceptual scheme is as good as the last – there is no decreasing marginal returns. To put it another way: learning two conceptual schemes is twice as good as learning one, and learning three is three times as good as one, etc.[1] Finally, each person's conceptual scheme is as good as any other's – everyone is equally smart. This allows me to represent the benefit of collaboration with a variable, b. Gaining access to one collaborative scheme improves one's ability to solve the problem by degree b, two schemes are worth $2b$, etc. For simplicity I will normalize the expected quality of one's solution without collaboration to zero and normalize the benefit for each additional collaboration b to 1.

But, as we all know, collaboration comes at a cost. Learning another's conceptual scheme takes time; time that might be better spent working in isolation. Worse yet, the new conceptual scheme might be misleading causing one to waste time on inferior solutions. Whatever the source, I will represent this cost by c. We henceforth assume that there is always a cost, by stipulating that $c > 0$.

Each scientist chooses a group of collaborators. Collaborators can, with high fidelity, communicate the relevant parts of all the conceptual schemes they know – both their own and schemes gleaned from others. So, if Carlos collaborates with Julie and Julie collaborates with Shannon, Carlos gains access to both the relevant parts of Julie's conceptual scheme and also to the relevant parts of Shannon's. As mentioned before, all exchanges are one-way. So, if Carlos contacts Julie he gains Julie's conceptual scheme (at a cost to Carlos but at no cost to Julie[2]) but Julie does not gain access to Carlos'.

[1] If there are only finitely many solutions and each individual scientist is good at finding the best possible solution on her own, this assumption cannot hold. At some point there must be decreasing marginal returns. If however, each scientist is ineffective at finding the best solution on her own this assumption is perfectly reasonable. Given the difficulty of most scientific problems, the equal gains assumption is not far off.

[2] One might balk at this assumption, since communicating one's conceptual scheme to another might take time. Since Julie does not gain from her interaction with Carlos,

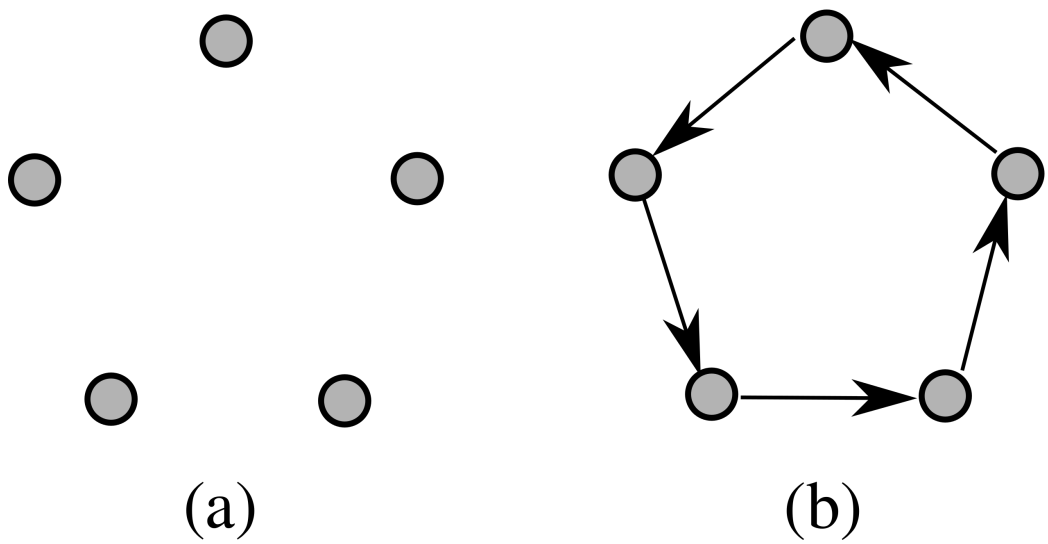

Figure 1: Two graphs: the empty graph (a) and the directed cycle (b).

The model proceeds in two distinct steps. First, everyone chooses which scientists they will contact. Second, schemes are exchanged. This removes complex order effects where it matters whether I contact you before or after you contact another scientist. While this reduces the realism of the model, it greatly increases its tractability.

With these assumptions one can now use pictures to represent the outcome of the choices of each scientist. Scientist i forming a connection to scientist j is represented by an arrow from j to i (representing the direction of the flow of conceptual schemes). A scientist can then be assigned a "payoff" from his choice: the total number of upstream scientists minus the total cost of the connections he forms.

This representation makes clear why models of this form are called "network formation" models. Each scientist chooses whom to link to in a network, and she receives a payoff determined by the entire network's structure. These models are used to represent a multitude of relationships including co-authorship and friendship (Jackson, 2008; Goyal, 2007). But here I will focus on the epistemic interpretation (Zollman, 2013).

Julie would have no reason collaborate with Carlos if it costs her. One might assume, for instance, that part of the cost imposed on Carlos (captured by the variable c) is the amount of remuneration he must give Julie for her time.

There are three distinct metrics for the effectiveness of a collaborative network. The first is called *social optimality*. A pattern of collaboration is socially optimal if it maximizes the sum of the individual payoffs. This represents the situation where all problems are solved to the greatest degree possible given the conceptual schemes available in the community. In this model, only two different states could be socially optimal. When $c \leq n-1$, collaboration is worthwhile so long as doing so will give one access to a rich set of conceptual schemes. In this context the *directed cycle* is optimal (see figure 1(b)). If $c > n-1$ then collaboration is not worthwhile and the only optimal state is one where no one collaborates with anyone (Bala and Goyal, 2000).

This result has important implications. Consider for instance the apparently best case scenario for collaboration, $1 > c$. Here collaboration is valuable even if I only gain access to one additional conceptual scheme. In a large group, one might be inclined to infer that everyone ought to collaborate with everyone. After all any pair of scientists, if they were considered in isolation, ought to collaborate. But, when conceptual schemes are transmitted second hand, this is not true; the optimal structure of collaborative exchanges is sparse. The system would not be improved by encouraging further collaboration.

The reason for this, somewhat counter-intuitive conclusion comes from the presence of externalities mentioned at the outset. If Julie has contacted Shannon and learned her conceptual scheme, then Julie is more valuable as a collaborator for Carlos. If Julie can transmit Shannon's scheme with high fidelity, then Carlos does best by working with Julie and gaining access to two additional conceptual schemes. Carlos has little to gain by also interacting with Shannon, since Julie has given him all the information that Shannon might provide. This will hold true even when second hand transmission of information loses fidelity – up to a point, see Bala and Goyal 2000.

Social optimality in games is a technical concept which imperfectly approximates our intuitive notion of "good for the group." First, if one is modeling scientists researching a problem that has implications for those outside of science, one also might want to consider the impact of the research on society at large. Under the intended interpretation of the model provided here, the payoff to scientists is just the degree to which they solve the problem. If each scientific problem is valuable (and equally valuable) to the world-at-large, then what is socially optimal for the community of scientists reflects the good of the community at large.

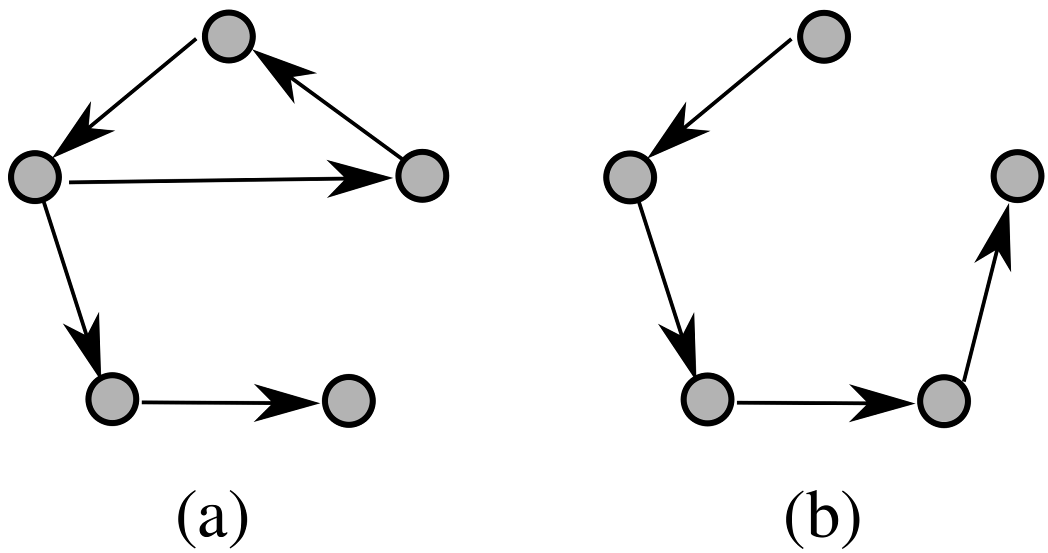

Figure 2: Two non-stable graphs.

One might worry about social optimality because one objects to maximizing the sum or average payoff. This focus is reasonable for groups of scientists since we are primarily interested – as a society – in the total progress of science. In many political settings, in contrast, some advocate for giving special consideration to the worst off. Egalitarian considerations reach the same conclusion in this model since the worst off does best in the state I call "socially optimal." Other measures are possible, too, and some of them will not so neatly coincide with my choice for a measure group effectiveness. However, given the intended interpretation of the model, they cannot be given much defense.

The second metric of evaluation for a pattern of collaboration is stability. A pattern of collaboration is stable, if no individual scientist is able improve his situation by changing his collaborative partners. For instance, the pattern pictured in figure 2(a) is not stable, because the individual in the upper right can do better by connecting to the individual in the bottom right. By doing so, she would increase her payoff from $2 - c$ to $4 - c$.

In this model, socially optimal states are always stable (Bala and Goyal, 2000). When $c < 1$ scientists are willing to collaborate with one other scientist even if the potential collaborator can only provide her with one new conceptual scheme. In these cases, only the directed cycle is stable – in any other situation at least one scientist has a positive incentive to change her pattern of collaboration.[3] When $n-1 > c > 1$, the directed cycle remains stable but the state where no one interacts is also stable (see figure 1(a)). Here collaboration cannot get "off the ground" because no one is willing to take the first step. The first person to collaborate must pay the cost to gain access to one conceptual scheme, and since c is greater than the value of a single scheme, no one wants to pay the cost. We have an example of a sub-optimal, but stable, outcome – a situation familiar to many of us. Turning to the last case, when $c > n-1$, we find the only stable state is the one where no scientist collaborates with any other which is also optimal.

The last metric of evaluation, which will be our primary focus, is one of "learnability." When $n-1 > c > 1$, there is more than one stable state. Even if scientists always find a stable states, one cannot be assured that they will come to land on the socially optimal one. Scientists might not come to find a stable state, but wander around moving from one unstable state to another, perhaps indefinitely. To uncover which states are learnable, one must model of learning.

2 Modeling learning

In order to model learning, I now must extend the model of one-shot collaboration to repeated collaboration. Now scientists are confronted with problems sequentially. Each time a scientist is confronted by a new problem, she can connect to others and collaborate on *that* problem according to the model above. Each new problem differs from the previous one, and each instance of collaboration only provides enough information about a conceptual scheme to solve the problem under consideration. This assumption seems plausible as collaboration is usually problem-focused. The resulting model is tractable because the benefits from each instance of collaboration is independent of the previous collaborations.

Many models of learning in game theory require scientists to form a belief about what others are doing. In games of this type, this requirement amounts to having a probability distribution over all directed graphs with n nodes – a large space even for moderately sized n. Not only would this be difficult to analyze, it is unlikely to provide any real insight into how scientists behave.

[3] When $c = 1$ things are slightly more complex. A scientist has neither a positive nor a negative incentive to change away from some other patterns of interaction.

Instead Huttegger et al. (2014) suggest this game is best analyzed using a learning rule called "probe and adjust" which represents unadorned learning. Each scientist has a default strategy which she usually employs. Occasionally, a scientist experiments – she tries a new strategy at random (a probe). If the new strategy outperforms the default, she adopts it as her new default. On the other hand if the new strategy is inferior to the old one, she returns to the previous default. If the two are tied, she chooses which one will be the new default at random.[4]

This method of strategy revision has a few interesting properties. If one enforced the rule that there must be one round of default play in between single probe events, then the process would be guaranteed to evolve into a stable state and then the default behavior would never change. While mathematically helpful, limiting the system in this way seems implausible – how would scientists ensure they weren't probing simultaneously or immediately after another? If one relaxes this assumption the group can escape from stable states, even optimal ones.

As an example, consider the directed cycle (figure 1(b), the optimal state), where $c = 3$. Suppose the individual in the upper right probes by connecting to the individual in the upper left (figure 2(a)). Not only does this lower her own payoff but her probe lowers the payoff of everyone else except for the person on the lower right. Of particular importance for this example is the individual on the top, whose payoff is now $2 - c = -1$. After one round of probing the individual on the top left will switch back to her default strategy. If the individual on the top probes on the subsequent round by abandoning collaboration altogether, his situation has improved over the previous round – his payoff was -1 (because the other individual was probing) and now it is 0. So, he will stick with this strategy (figure 2(b)). Now the individual on the upper left will prefer to drop his connection, and so on. The system has now abandoned an optimal state.

[4] Whether or not probe and adjust represents *exactly* how scientists are learning is not important. The tractability of probe and adjust makes this model an important starting point from which more general lessons can be drawn. The phenomena illustrated in the next section strike me as sufficiently general as to reoccur in almost any learning rule where scientists are unaware of the other connections in the network – i.e. where scientists do not know the other's patterns of collaboration.

3 Results

The overarching question is, to what extent can scientists come to collaborate efficiently? This is a multifaceted question that is addressed in parts.

First, to what degree is experimentation with collaboration patterns helpful? Is the group improved if people are constantly exploring – perpetually probing – or should they be more set in their ways? The answer to this question is not obvious. On the one hand, high probing rates will help the system escape sub-optimal ones relatively quickly. But, high rates will also increase the probably that the system will leave an optimal state in favor of a sub-optimal one.

A second research question concerns how the cost of collaboration, c, affects the system. If every group always attained optimality, then everyone improves c becomes smaller. But, if we cannot be assured of optimality, the answer is not so straight-forward. While a lower c will make the final payoff in the higher for everyone, it also provides a smaller incentive to find the optimal state. A lower c might cause the community to wander around more. Furthermore, lowering c might also affect the ease with which the system escapes the optimal state.

The last research question relates to group size. Smaller groups are likely to reach optimality more quickly. As Huttegger et al. (2014) point out, the number of non-optimal states grows quickly as new scientists are added, but the set of optimal states grows much more slowly. So, the proportion of states that are optimal approaches zero as the number of scientists approaches infinity. On the other hand, the payoff for scientists *even in non-optimal states* goes up as the set of potential collaborators increases (assuming we hold c constant). How do these two considerations trade off against one another?

3.1 Experimentation rates

Huttegger et al. (2014) prove that when $c < 1$ (i.e. when the only stable state is also the only optimal one), the system will spend most of its time in the optimal state – the directed cycle – in the limit. Stated more formally, if you choose a probability p other than 1, there is a probe probability where the long-run probability the system is in the optimal state is at least as high as p. This proof is achieved by showing that from any state there is a chain of intermediate states where (a) at each stage only one player must change to go from one state to the next, (b) the person who changes does not make herself worse off, and (c) the end of the chain is the directed cycle.

This result is limited. First, we cannot generalize to cases where $c > 1$ – the theorem is clearly no longer true. Second, the result does not systematically answer the question about experimentation rates. Finally, we only know what happens in the infinite limit which can be very different from short-run results.

In order to provide a more general analysis, a simulation study was conducted. Simulations considered groups of size 4, 5, 6, 7, 8, 9, and 10 with probe probabilities ranging from 0.01 to 0.21 in 0.03 increments. c varied from 0.1 to $n + 0.1$ in 0.5 increments.[5] For each setting of the parameters 1,000 simulations were generated where the probe and adjust process continued for 100,000 rounds.[6]

Figure 3 illustrates the conclusions regarding probe probabilities. The top plot shows how well scientists fared when measured by the time they spent in the optimal state.[7] Non-optimal states can nonetheless be better or worse. The bottom plot shows how well the groups fared when measured by normalized payoff (0 is the worst possible payoff for that configuration and 1 is the best possible payoff). Both plots reveal the same general trend: lower probe probabilities are better. (This generalizes simulation results presented in Huttegger et al. 2014 for a small cost version of this game.)

Obviously a probe probability of zero would be inferior because the group would never change. However, low probe probabilities do much better than higher ones. Of those tested, a probe probability of 1% was superior. In communities appropriately modeled in this way, one should be cautious about encouraging scientists to experiment with new patterns of collaboration. These simulations suggest that encouraging exploration will hinder the ability of

[5]The results for these group sizes and probe probabilities show clear patterns that can be projected to larger group sizes and probe probabilities (but not without some care to *smaller* probe probabilities).

[6]100,000 rounds were chosen in order to model a long, but finite time. This will provide a helpful contrast to the limiting results of Huttegger et al. (2014). The results between the limit analysis and this, very long, time differ in significant respects. Since real scientists are operating at much shorter time spans, there might be further differences. The fundamental differences between finite and infinite limit analysis can be identified at this long time scale. This satisfies the duel goals of understanding the system while not risking a time-scale that is "too short."

[7]In general there is also a strong relationship between time spent in optimal state and the first time optimality is obtained.

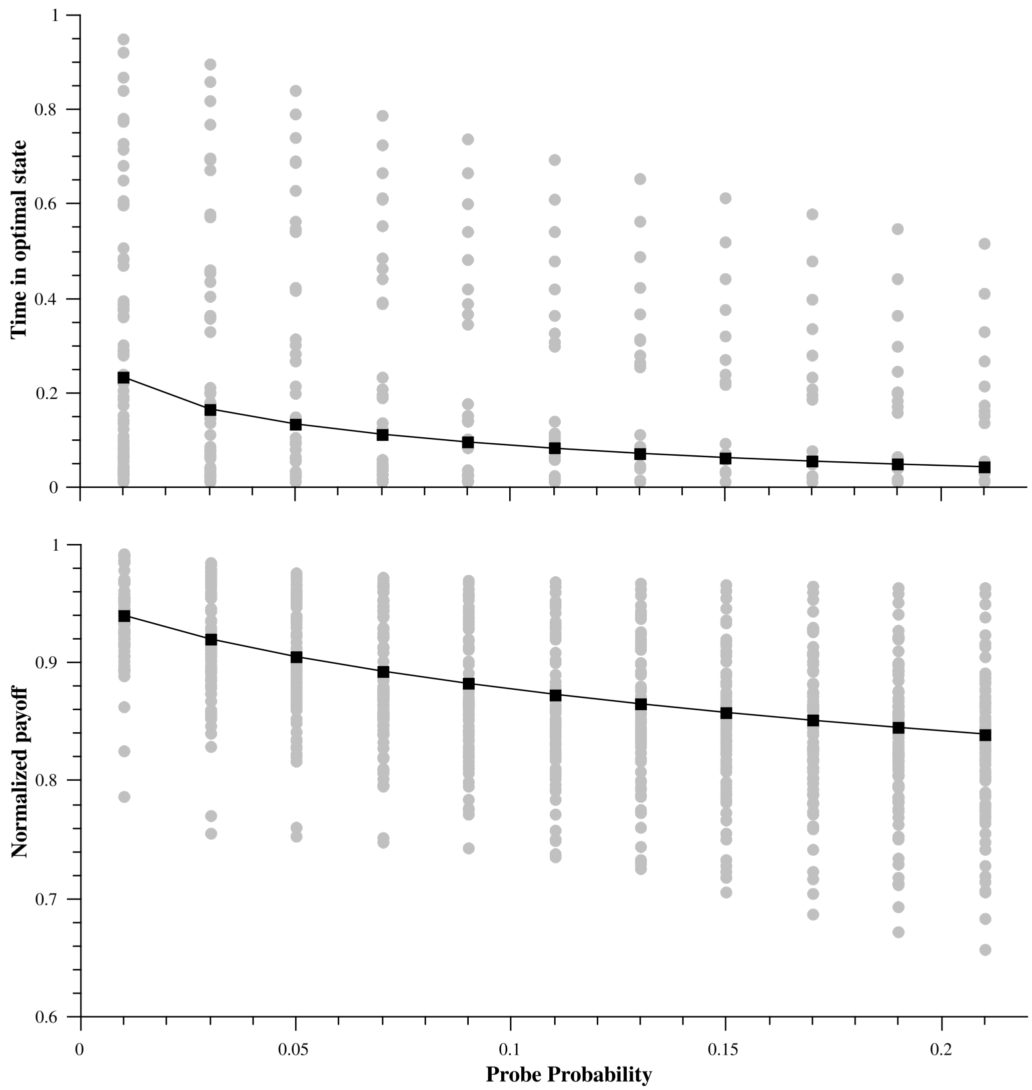

Figure 3: Two plots which relate the probability of a scientist probing to the proportion of time spent in the optimal state (top plot) and to the average normalize payoff (bottom plot). In both plots the gray circles represent the average of 1,000 simulations for each of the parameter settings. The black squares represent the average of all simulations for all parameter settings which use the same probe probability.

the system to find and remain in the optimal state. Beyond that, increased experimentation will harm the communities ability to do well even when not in an optimal state.

3.2 Cost

What effect will the cost of collaboration have on the overall performance of the system? It seems intuitive to strive to minimize the cost of collaboration, but will the simulations endorse this policy? To answer this question I will focus on a representative group size, 7 individuals, all of whom have a low probe probability of 0.01. All the qualitative facts reported here are true of the other group sizes studied.

First let us consider how cost relates to the ability of scientists to find optimality. This is pictured in the top plot of figure 4. Recall the three different regimes. When $c < 1$, the optimal state is the only stable one. Here, however, the system is rarely in the stable state. This is consistent with results reported in (Huttegger et al., 2014). They conjecture that it is difficult to find the optimal state because the search space is so large. Indeed this is likely part of the problem. But the size of the search space cannot be the whole story, because when the cost is higher in this region ($c = 0.6$) the group performs slightly better than when it is lower ($c = 0.1$). Furthermore, comparing the low cost regime to cases where $c > 1$ suggest the situation is rather more complicated.

Turning to the second regime, where $6 > c > 1$, we find a complicated relationship between cost and time in the optimal state. There is an optimal cost between 2 and 3 where the system is best (in this regime) at finding, and remaining in, the optimal state. While the changes in cost do not effect what state is optimal, they do affect the ordering of non-optimal states, and in so doing make the optimal state easier to find. Considering only those simulations that occupied the optimal state for at least one round, the system with a cost of 2.1 was over three times faster at finding the optimal state than the system with a cost of 0.1.

One should be careful about jumping to the conclusion that one should attempt to increase the cost of collaboration. While the identity of the optimal state is not changed as the cost increases, the objective quality of this state changes. When the cost is 0.1, each individual in a seven person group receives a payoff of 5.9 in the optimal state, but when the cost is 2.1, each individual receives a payoff of 3.9. So while the system spends more

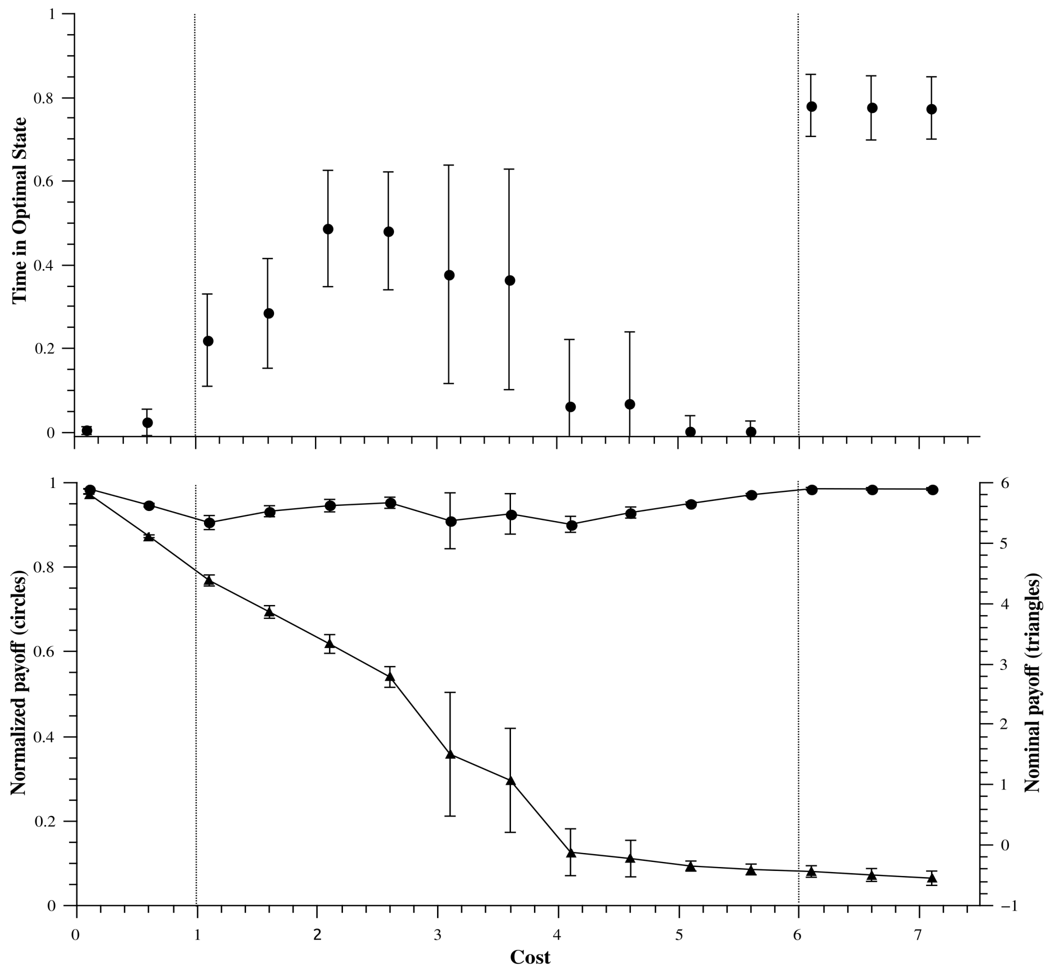

Figure 4: Simulations results for a 7-person group with probe probability set at 0.01. The plots represent the average payoff and standard deviation for different values of cost versus time in the optimal state (top plot) and normalized payoff (bottom plot, circles plotted against the left y-axis) and nominal payoff (bottom plot, triangles plotted against the right y-axis).

time in the optimal state, this is not equivalent to improving how individuals fare.

And in fact, they are not faring better with higher costs. The bottom plot of figure 4 shows both the normalized payoff and the nominal payoff. The later is most important. While increasing the cost increases the time spent in the optimal state, doing so also makes that state worse. The later consideration outweighs the former. While the low cost communities wander around for a significant amount of their time, they nonetheless do well.

Finally, in third regime, where $c > n-1$, scientists are adept at achieving optimality. This is not surprising because in this situation connecting to no one dominates connecting to anyone. Because connecting to no one is one strategy out of 64, individuals must find the dominant strategy which is why they are not in the optimal state the entire time.

These results can underwrite the intuition that collaboration is improved by making collaboration less costly (by perhaps improving mechanisms of communication, or by providing direct incentives designed to compensate for other costs). But, reducing the cost does not increase the chance that the community arranges itself in optimal ways – to the contrary.

3.3 Community size

As community size increases, it will become more difficult for the group to find the optimal state (Huttegger et al., 2014). But, as in the previous section, increasing the population of scientists also improves the nominal payoff of that best case. It therefore is an open question whether increasing the size of the community will be beneficial or harmful to that community.

Figure 5 illustrates the settings of the parameters which maximize the time in optimal state for each group size. While, as expected, the time spent in the optimal state decreases rapidly as the group size grows, the benefit from the larger size swamps the loss. The nominal payoff increases despite a greater fraction of the time is spent in non-optimal states.

There is a harm to increasing the group size; they will spend more time exploring non-optimal states. In this model, that harm is worth incurring because the benefit of increased collaborative possibilities is sufficiently large to outweigh the cost (at least in the best case).

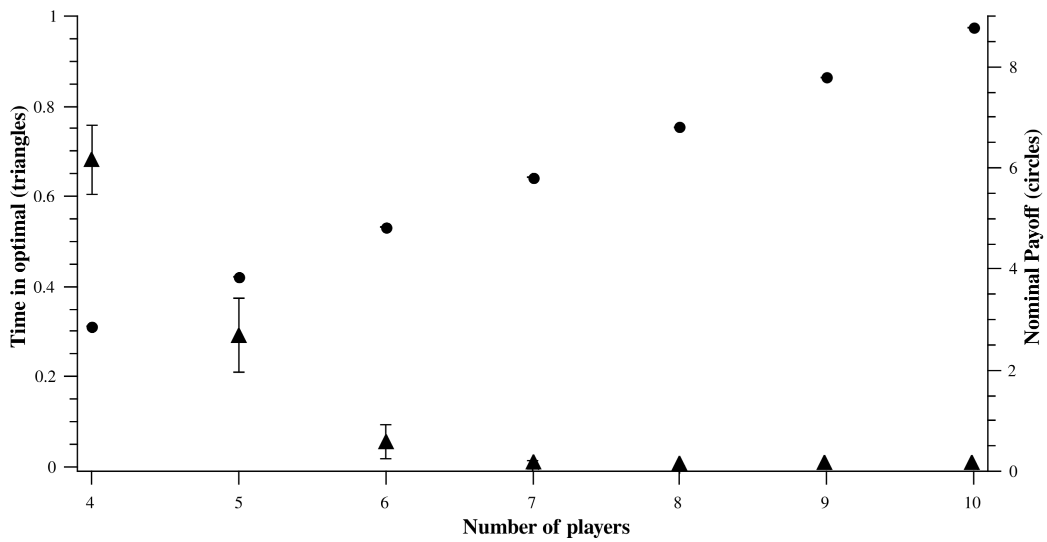

Figure 5: Simulation results comparing the number of players to the time in the optimal state (triangles plotted against left y-axis) and to the nominal payoff (circles plotted against the right y-axis). Each point represents the single probe probability and cost that maximized the nominal payoff relative to other values of those parameters for groups of the same size.

4 Conclusion

Not all collaborative interactions are analogous to this model. I have argued that there are unlikely to be any fully general results about the structure of epistemic networks (Zollman, 2013), and I expect this will be true with collaboration as well. However, I do believe that this model provides an appropriate idealization of some situations of collaboration, and where it does, it offers some clear guidance to how best achieve effective collaboration.

The model underwrites the general belief that productive collaboration is facilitated by (a) increasing the number of individuals with whom collaboration is possible and (b) decreasing the cost to scientists for engaging in collaborative exchanges. These are true despite a reduction in the time spent in optimal states.

More surprisingly, however, the model suggests that encouraging scientists to experiment more with different collaborative arrangements will not be productive. With the increased interest in interdisciplinarity, funding agencies and administrators have been encouraging scientists to find new patterns of collaboration. However, in this model, spuring new collaboration by introducing new scientists to one another or by encouraging scientists to try something new, is counterproductive.

References

Bala, V. and S. Goyal (2000, September). A noncooperative model of network formation. *Econometrica 68*(5), 1129–1181.

Goyal, S. (2007). *Connections: An introduction to the economics of networks.* Princeton: Princeton University Press.

Hong, L. and S. E. Page (2004, November). Groups of diverse problem solvers can outperform groups of high-ability problem solvers. *Proceedings of the National Academy of Sciences of the United States of America 101*(46), 16385–9.

Huttegger, S., B. Skyrms, and K. J. S. Zollman (2014). Probe and Adjust in Information Transfer Games. *Erkenntnis 79*, 835–853.

Jackson, M. O. (2008). *Social and economic networks.* Princeton: Princeton University Press.

Jackson, M. O. and A. Wolinsky (1996). A Strategic Model of Social and Economic Networks. *Journal of Economic Theory 74*(0108), 44–74.

Kerr, N. L. and R. S. Tindale (2004, January). Group Performance and Decision Making. *Annual Review of Psychology 55*, 623–655.

Maynard Smith, J. and D. Harper (2003). *Animal signals.* Oxford: Oxford University Press.

Zollman, K. J. S. (2013). Network Epistemology: Communication in Epistemic Communities. *Philosophy Compass 8*(1), 15–27.