Withholding Knowledge

Jingyi Wu*

November 1, 2022
Draft—Please do not cite or circulate without permission.

Abstract

If scientists are purely motivated by the truth of their findings, are they incentivized to share their evidence with each other? I use computer simulations of two paradigmatic models of scientific inquiry to argue that there are distinct epistemic advantages in unilaterally withholding evidence from the outside, even compared to a mutually sharing community in many cases. I further analyze the sharing and withholding dynamics from a game theoretic perspective by constructing epistemic games from simulation results.

1 Introduction

The communist norm in science mandates that scientists should share their work as widely as possible (Merton 1973). It is a core norm in science, but not a given. While academic scientists largely adhere to the communist norm, industrial scientists do not typically share proprietary findings.[1] There seem to be community-level benefits for sharing one's work. For instance, scientific discoveries could be made faster, so the public as well as the rest of the scientific community can reap the benefits earlier (Strevens 2017; Heesen 2017). But do *individual* scientists, or *subgroups* of scientists, have an incentive to share?

Previous work in the philosophy of science answers this question by appealing to credit incentives by which scientists may be instrumentally motivated. Strevens (2017) appeals to the priority rule, which stipulates that credit—a proxy for recognition for one's scientific work—will be allocated only to the scientist who first makes a discovery. He argues that because of the priority

*Department of Logic and Philosophy of Science, UC Irvine. Jingyi.Wu@uci.edu.

[1] In fact, publicly-funded scientists in the US *must* share their research by 2025 (Brainard and Kaiser 2022). Many European funders have similar requirements.

rule, individual scientists would have credit incentives to withhold their evidence from others during the process of inquiry. In so doing, they increase their own chances of making a credit-worthy finding first.[2] However, Heesen (2017) uses a multi-stage game theoretic model to argue that in many cases, it is still rational for credit-maximizing scientists to share, because scientists can publish and claim credit for intermediate results.

Credit is only one of many incentives facing scientists. They are also, or at least nominally, motivated to find the truth (see, e.g., Bright (2017) and Zollman (2018)). Curiously, this epistemic incentive underlying scientific sharing has been relatively under-explored. The received position seems to be that a scientist's decision to share or not during the learning process does not impact the likelihood of them eventually making a discovery. Not only do Bright and Heesen (2021) recently explicitly claim that the communist norm is non-epistemic just in this sense, but both Strevens (2017) and Heesen (2017)'s models also implicitly assume that the probability that a scientist will eventually make a discovery remains unchanged regardless of whether they share during the learning process. It is unclear what grounds this position, especially given the growing literature from network epistemology that shows how our social connections and evidence-sharing dynamics can significantly shape knowledge production (Zollman 2010; O'Connor and Weatherall 2018; Fazelpour and Steel 2022; Wu 2022).

In this paper, I focus on the epistemic incentives underlying scientific sharing. Specifically, I ask, if scientists are purely motivated by the truth or epistemic significance of their own findings, are they incentivized to share evidence? I investigate this question by simulating models of an epistemic community with two subgroups, one capable of withholding evidence from out-group members, and one adhering to the communist norm by sharing evidence.[3] I find that the subgroup that withholds ends up with several epistemic advantages, both compared to the subgroup that shares and to an entirely "communist" community where everyone shares regardless of subgroup membership. This suggests that a truth-seeking scientist may be incentivized to withhold their evidence.[4]

I construct two models from different paradigms for this purpose—a generalized multi-armed

[2]Partha and David (1994) make a similar argument more informally.

[3]Note that my models have two *subgroups*, whereas previous models (e.g. Strevens (2017) and Heesen (2017)) typically have two *agents*. But this is not a major departure—results presented here still hold when there are two agents, one withholding and one sharing.

[4]Note that in my models the scientists are motivated by *themselves* finding the truth. So the epistemic benefits may translate into practical benefits for the scientists too—such as recognition of or financial gains from their work.

bandit model with a network structure based on Zollman (2010) and an NK landscape model with a network structure based on Lazer and Friedman (2007). These two models represent two different types of scientific inquiry. In one model agents figure out which of the two probabilistic epistemic options are better, representing, e.g., clinical doctors finding out the efficacy of two different drugs by conducting trials. In the other, agents search in a vast epistemic landscape with multiple "peaks," representing, e.g., researchers adopting different approaches to solve a problem. Together these models represent a wide array of possible scientific problems.

Results from the bandit model show that members of the "withhold" group reach the true belief more frequently and faster, and select the epistemically better action more frequently during the learning process, as compared to the "share" group. In the NK landscape model, members of the share group end up with worse solutions than both the withhold group and a generally communist community. Moreover, the withhold group gains epistemic advantage even compared to the communist community in terms of arriving at the true belief more frequently (in the bandit model) and ending up with epistemically better solutions (in the NK landscape model in most cases).[5]

These results are troubling, especially given that proprietary industrial scientists (see, e.g. DeAngelis (2003), Michaels (2008), and McGarity and Wagner (2010)) and scientists working on classified research (see, e.g. Galison (2004)) routinely withhold their evidence from others. Academic scientists, on the other hand, largely conform to the communist norm (Louis et al. 2002; Macfarlane and Cheng 2008). My results suggest that even if these scientists are not as instrumentally motivated by credit as academic scientists, they may still have epistemic incentives to withhold evidence. Moreover, while not explicitly modeled in this paper, industrial scientists likely have further financial incentives to withhold evidence. Consequently, academic scientists suffer epistemically.

I will then observe that, based on simulation results, this share-withhold dynamics gives rise to what one might call an Epistemic Weak Prisoner's Dilemma. Each subgroup receives the highest epistemic payoff when they withhold while the other subgroup shares. Their payoffs are the worst when the other group withholds, and their payoff is intermediate when both subgroups share. I will discuss features of this game and explore strategies that may shift communities into mutual sharing.

[5] These results extend and provide robustness checks for recent results from a simpler multi-armed bandit model, built for another context (Wu 2022).

This paper is organized as follows. In §2, I introduce the generalized bandit model and present my simulation results. In §3, I introduce the NK landscape model and discuss my results. I also note common features of the two models that lead to robust qualitative findings and explain why certain results are not as robust. In §4, I show that a weak prisoner's dilemma may represent the epistemic dynamics of scientific sharing. §5 concludes. This paper is supplemented by two technical appendices. Appendix A provides definitions of the end states in the generalized bandit model. Appendix B describes how the solution space of the NK landscape model is generated.

2 The Generalized Bandit Model

In this section, I construct a model where a group of agents is tasked with a learning problem. There are two subgroups: the withhold group, whose members only share evidence with in-group neighbors and withhold from out-group neighbors, and the share group, whose members share evidence with every neighbor. The goal of this model is to explore the epistemic consequences of this asymmetric evidence sharing dynamics across subgroups.

Let us motivate the model by considering a toy example. Suppose that a group of clinical scientists is tasked to figure out which of the two drugs, A or B, are more effective at treating a disease. Each scientist tests the drugs on patients, receives evidence from their tests, and updates their beliefs from the evidence they and others receive. The scientists have limited resources, so they each only test one drug per round. Moreover, they want to minimize patient suffering, so they always assign the drug that they currently think is better. As the scientists keep getting evidence for the drugs, eventually they should reach stable beliefs about which drug is better.

This example is well modeled by what is called a two-armed bandit problem with a network structure. The name "bandit problem" comes from applying the model to a gambling situation, where a gambler aims at maximizing their profits when playing with a multi-armed "bandit" (or slot) machine. Let us first consider the problem for one agent, before thinking about it in a group setting. Every round, the agent selects between two options, A and B, tests their choice a fixed number of times n, and gets evidence about how many times their tests succeed. Each option has a fixed probability of success. Unbeknownst to the agent, I set the success rate of A, P_A, to be .5, and the success rate of B, P_B, to be lower than .5. This means that A is objectively the better choice.

The agent, however, is uncertain about the success rate of either option, and their credence for each is represented by a beta distribution with two parameters, α and β.[6] Details about the beta distribution do not matter for our purpose. What matters is that in the context of Bayesian learning, we can interpret α as the agent's estimate of the "successes" of the arm, and β as their estimate of the "failures." For instance, suppose the agent starts with $\alpha_A = 1$ and $\beta_A = 3$ for option A, tests this option 10 times, and receives 8 successes and 2 failures, then, after Bayesian updating, their posterior will still be a beta distribution, with parameters $\alpha_A = 1 + 8 = 9$ and $\beta_A = 3 + 2 = 5$. If this arm continues to be tested, over time, $\frac{\alpha_A}{\alpha_A+\beta_A}$ will approach the true success rate of A. Every round, the agent selects the option that they currently think has a better chance at succeeding,[7] tests it n times, and updates on their evidence in the (myopically) Bayesian way described above. After sufficient rounds, the agent's belief about which option is better stabilizes.

Now suppose a group of agents, connected via a network structure, solve the same bandit problem together. Each agent starts with four randomly assigned parameters $\alpha_A, \beta_A, \alpha_B, \beta_B$,[8] representing their own credence for the two options. The network structure determines who they are connected to, or who their "neighbors" are. Typically, as in the case of Zollman (2010), every round, after collecting their evidence, agents share the evidence with all their neighbors and update on their neighbors' evidence in the Bayesian way.

In my model, however, not everyone shares evidence with all their neighbors. I divide the community into two subgroups: the withhold group, whose members share their evidence with in-group neighbors but withhold evidence from out-group neighbors, and the share group, whose members share their evidence with all their neighbors. All agents then update their beliefs based on all the evidence they receive (including their own evidence) in the Bayesian way. This modification creates an asymmetry along group membership in the evidence-sharing dynamics. To go back to the previous example of clinical trials of drugs, my model may represent a situation where a community

[6] A beta distribution is a function of the following sort:

$$f(x) = \frac{x^{\alpha-1}(1-x)^{\beta-1}}{B(\alpha,\beta)}$$

where $B(\alpha,\beta) = \int_0^1 u^{\alpha-1}(1-u)^{\beta-1}du$ and $\alpha, \beta > 0$.

[7] If the mean of their beta distribution for A is higher than the mean of their beta distribution for B, i.e. $\frac{\alpha_A}{\alpha_A+\beta_A} > \frac{\alpha_B}{\alpha_B+\beta_B}$, then they choose A; they choose B otherwise.

[8] These are real numbers between 0 and 4 (exclusive).

of scientists is trying to figure out the efficacy of two different drugs, but within this community, there is a subgroup of industry-affiliated scientists who withhold their evidence from the outside, even though the rest of the community continues to share their evidence widely.[9]

I run each simulation for $10,000$ rounds. At the end of the simulation, there are typically three stable end states: (1) community convergence to the true belief, where everyone thinks that A is better; (2) community convergence to the false belief, where everyone thinks that B is better; and (3) polarization, where the withhold group thinks that A is better, but the share group thinks that B is better. An overwhelming majority of simulations end and stay in these states.[10] Detailed definitions of the end states are available in Appendix A.

Polarization is only possible in this model because one group withholds evidence from another. In this state, the withhold group succeeds in learning, but the share group fails in learning. This state is stable because here, the share group has already settled on B and they do not receive evidence about A anymore, even though the withhold group continues to test A. It is important to note that a polarized state in the other direction, i.e. the share group succeeds but the withhold group fails, is not stable. This is because the withhold group would continue to receive evidence for action A from the share group, and since A is in fact better, the evidence they receive would prompt them to eventually switch.

It is important to note that this model is a generalization of a previous model (Wu 2022) involving a simpler bandit problem. There are two subgroups in Wu (2022)'s model, one ignoring evidence from out-group neighbors, and one updating on evidence shared by all their neighbors. The asymmetric evidence *updating* dynamics in Wu (2022)'s model is structurally equivalent to the asymmetric evidence *sharing* dynamics in this model, with the ignored group in Wu (2022) and the withhold group in this model occupying structurally the same position. This model uses a more complex bandit problem where agents start with not knowing the success rates of either epistemic option and there are infinitely many possible success rates for both options, whereas in Wu (2022)'s model, agents start with knowing the success rate of one option and there are only two possible success rates for the other. As we will see, my results on the epistemic advantage of the withhold

[9]Note that industrial scientists mutually share with each other. I briefly explore what would happen if there are multiple mutually withholding "industry" groups in §4.

[10]Simulations with learning problems that are "hard" (e.g. where P_A and P_B are close or when n is small) are more likely to not finish in these end states than problems that are "easy," though for each set of parameter values, more than 97% of simulations reach one of these end states.

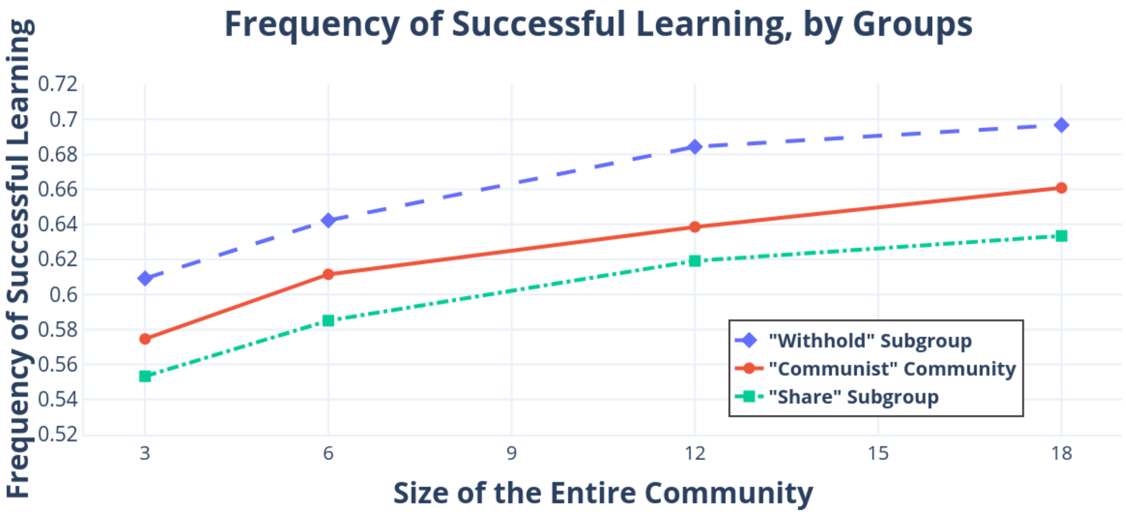

Figure 1: $P_B = .49$, $n = 10$, $d = \frac{1}{3}$, complete network structure, $10,000$ simulations.

group provide a robustness check in a more generalized setting for the results in Wu (2022).[11]

2.1 Results and Discussions

For each set of parameter values, I run the model for $10,000$ simulations. The parameters tested include:

- Total number of agents (k): $3, 6, 12, 18$;
- Proportion of the withhold group in the population (d): $\frac{1}{6}, \frac{1}{3}, \frac{1}{2}, \frac{2}{3}$;
- Number of tests per round (n): $1, 10, 100, 1000$;
- Success rate of B (P_B): $.4, .45, .49$;
- Network Structure: complete.[12]

In addition, I run a "communist" model, where everyone shares with everyone, for all the parameter values. This is equivalent to models from Zollman (2010), and it provides a comparison to my results. Unless otherwise noted, results reported here are robust across all parameter values.

I find first that the withhold group succeeds in learning more frequently than the share group (Figure 1). I measure the frequency of successful learning for a subgroup by calculating the proportion of simulations that the subgroup succeeds in learning out of all simulations where the community ends in one of the three end states. In this model, the withhold group succeeds in learning more frequently

[11] Wu (2022) uses her results on the epistemic advantage of the ignored group to justify a standpoint epistemology thesis that marginalized group can sometimes have better knowledge.

[12] A network structure is complete when every agent is connected to every other agent. I also consider the directed Erdős-Rényi random networks of size 18 for robustness checks.

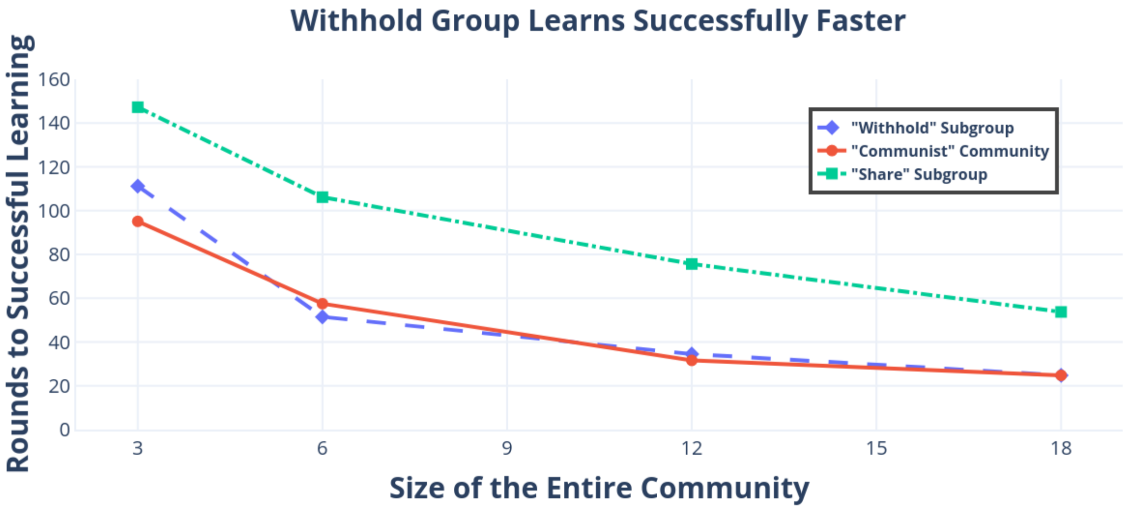

Figure 2: $P_B = .45$, $n = 1$, $d = \frac{2}{3}$, complete network structure, $10,000$ simulations.

because polarization is an end state. In this state, the withhold group succeeds in learning but the share group fails.

Moreover, and perhaps surprisingly, the withhold group succeeds in learning even more frequently than a communist community where everyone shares with everyone (Figure 1).[13] The reason is that in simulations that start with initially promising evidence for B and unpromising evidence for A,[14] members of the communist community may settle on B as the better option. But in my model, since the share group is effectively an isolated communist community of a smaller size, it takes them longer to reach a stable state. This in turn means that evidence from both actions circulates in the network for longer, and the community, including the withhold group, may revert to choosing A. This explanation is related to what is often called the Zollman effect (Zollman 2007), which states that a more sparsely-connected network more frequently succeeds in learning. Here, the sparsity of the network slows down learning, so the community spends more time exploring different options before settling down. This period of time where members of the community test different options is called a period of transient diversity (Zollman 2010; Wu and O'Connor 2021). When transient diversity lasts longer, it brings epistemic benefits to some or all community members.

It may be instructive to think about the withhold group's epistemic advantage through a trade-off between exploration and exploitation inherent in the bandit model. Due to the probabilistic nature

[13]This result holds for more than 92.2% of parameter combinations. It is more likely to fail when the learning is "easy," i.e. when P_B is low and n is high. When $P_B = .49$, this result holds 100%. This pattern of robustness levels is consistent with Rosenstock et al. (2017)'s findings on the Zollman effect.

[14]These simulation runs are possible because of the probabilistic nature of epistemic options in the model.

of both arms of the bandit, agents typically have to test each arm sufficiently many times to form a reliable estimate of its success rate. This gives rise to a dilemma for individual agents in the model—do they keep exploiting the option that they currently think is best, or do they explore the other option that seems inferior? By *not* sharing their evidence with out-group members, the withhold group in a sense takes the best of both worlds. They continue to exploit the option they currently think is better, while benefiting from the exploration of the share group.[15] This is related to the "free rider problem" (Kummerfeld and Zollman 2015), which describes situations where it is rational for individual agents to leave the exploring to others in the community. In my model, the free rider problem is even more insidious, since here the free riders, i.e. the withhold group, reap even more epistemic benefits than the rest of the community.

Furthermore, the withhold group takes a shorter time to succeed in learning than the share group in simulations that end in community convergence to the true belief.[16] The withhold group's speed to successful learning is comparable to the communist community's (Figure 2). The time a subgroup takes to succeed in learning is measured as follows. For every agent in the subgroup, I record the first round after which the agent's credence satisfies the success conditions (see Appendix A for details). I then take the average rounds out of all agents in the subgroup and all simulations that end in community convergence to the true belief. Here, the withhold group succeeds in learning faster because they update on more pieces of evidence than the share group. However, in simulations that end in polarization, the withhold group takes longer to learn the truth. This is because in these cases, the share group converges to believing that B is superior early on, but the withhold group slowly tests A until they have sufficient evidence for its superiority. This means that the withhold group consistently has epistemic advantage over the share group in all situations—when they both succeed in learning, the withhold group succeeds more quickly; when the share group fails in learning, the withhold group still may succeed, albeit slowly.[17]

The withhold group on average selects the epistemically better action more frequently during

[15]Note that the share group explores more not because of differences in their behavioral rules, but because they may be testing different options from the withhold group since they do not have access to the withhold group's evidence.

[16]This result holds for more than 92.8% of parameter combinations. It is more likely to fail when the learning is "hard," i.e. when P_B is closer to P_A. In this situation, it is difficult to distinguish P_A and P_B, so it is more likely to have a small number of simulations that take a long time to finish, thus skewing the average rounds to successful learning.

[17]Contrast this with the Zollman (2007, 2010)'s finding that a less connected community succeeds in learning more frequently but less quickly. The withhold group in my model truly takes the best of both worlds.

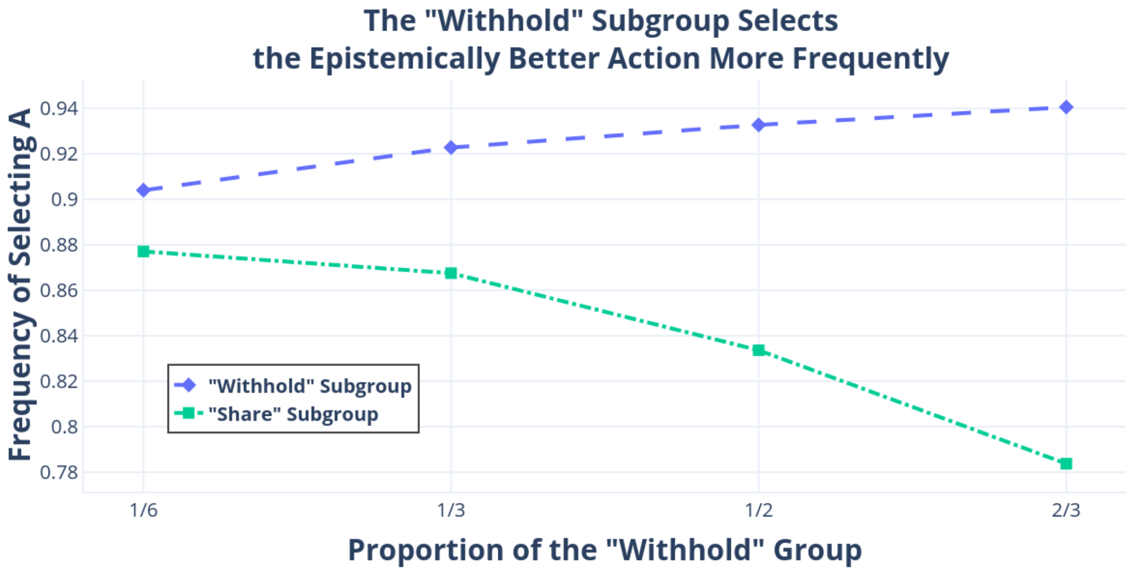

Figure 3: $P_B = .4$, $n = 1$, $k = 6$, complete network structure, $10,000$ simulations.

the learning process than the share group (Figure 3). For this, I count the number of rounds that an agent in the said subgroup selects action A, then divide it by all agents in the subgroup and all simulations.[18] As the withhold group gets larger, it selects the epistemically better action more frequently, and the share group selects it less frequently. The other epistemic advantages for the withhold group similarly depend on its proportion in the community. As the proportion of the withhold group grows, the withhold group's epistemic advantage tends to increase in terms of the frequency and speed of successful learning.

To further test for robustness, I simulate this model and Zollman (2010) with directed Erdős-Rényi random networks and find qualitatively similar results.[19]

[18]Including simulations that do not end in the three end states.

[19]Such a network is generated at the start of each simulation in the following way. First, every agent is linked to themselves. Then, for every agent X and every other agent Y, the probability that X is connected to Y is a fixed number b $(0 < b < 1)$. I ran the simulations for communities of size 18 for all other parameters listed, with probability of connection $b = .6, .7, .8, .9$, for $10,000$ simulations each. Note that this network is directed because X could be connected to Y without Y be connected to X. I expect a model with undirected network structures to produce similar results. Furthermore, I only consider connected* networks in this case, defined in the following way. A network is connected* if (1) there exists a path from any sharing agent to any arbitrary agent in the network; and (2) there exists a path from any withholding agent to any arbitrary withholding agent. Moreover, there exists a path from agent Y to agent Z if there are agents $A_0, ..., A_n$ such that $A_0 = Y$, $A_n = Z$, and A_i shares evidence with A_{i+1}, for $0 \leq i < n$. This definition is analogous to the definition of connected* in Wu (2022).

3 The NK Landscape Model

Next, I consider an influential epistemic landscape model—the NK landscape model.[20] My model again consists of two subgroups—a withhold group and a share group—searching in the same epistemic landscape for solutions to a problem. In what follows, I first introduce the general idea behind an epistemic landscape model. I then introduce the solution space of the NK landscape model, the network structure, and agents' behavioral rules in my model. After that, I compare and contrast this model with the generalized bandit model in §2, before discussing the results.

We can think of an epistemic landscape as containing a large number of research approaches to a particular topic of inquiry. Each research approach is a point on the landscape and has a score associated with it, representing its "epistemic significance." Following Alexander et al. (2015, p. 426), I interpret research approaches in a broad way—they have a number of components, including research questions, methods, skills, instruments, background assumptions and theories, etc. We can then introduce a group of agents searching the landscape, i.e. choosing research approaches and solving problems. This models important aspects of scientific problem solving—scientists constantly communicate with others in their community and decide whether they should stick with the research approach they currently have or try new ones, either by exploring on their own or adopting an approach from the community.

The NK landscape is a sophisticated multi-dimensional landscape with multiple "peaks."[21] The solution space is N-dimensional, with binary strings (0s and 1s) of length N as its points. At the start of each simulation, an algorithm with parameter K ($1 < K < N - 1$) is used to randomly assign scores between 0 and 1 to each string. A full description of the algorithm is available in Appendix B. Roughly speaking, the parameter K determines how "rugged" the solution space is and how correlated "nearby" scores are. When $K = 0$, the solution space is smooth with one single peak. When $K = N-1$, the solution space becomes totally chaotic, where the score of every point is totally independent of adjacent points. When $1 < K < N - 1$, the landscape is rugged, with multiple local optima, and with some correlation between adjacent solutions. In this case, some peaks in the landscape may

[20] While the NK landscape model is influential in other fields, it is under-explored in philosophy (with the exception of Alexander et al. (2015)).

[21] C.f. lower-dimensional epistemic landscape models, e.g. Weisberg and Muldoon (2009) and Hong and Page (2004). The NK landscape model was originally developed in biology to model "synergies" among genes, see Kauffman and Levin (1987) and Kauffman and Weinberger (1989).

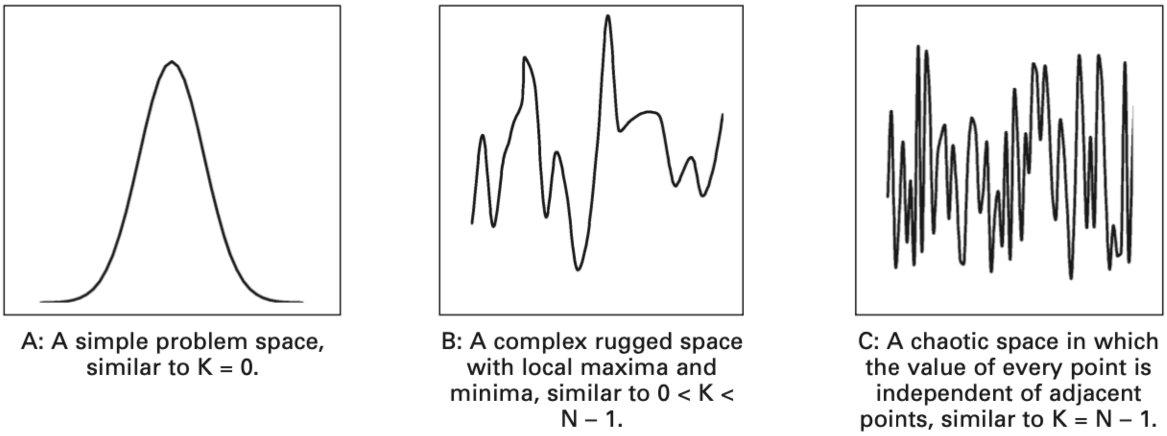

Figure 4: Stylized Representation of the Solution Space (Lazer and Friedman 2007).

only be reached via some solutions but not others. As K grows, the landscape becomes increasingly harder for agents to search. Though it is impossible to sketch a higher dimensional space, Figure 4 provides a stylized representation of the solution space as we vary K. [22] In this paper, we focus on the $1 < K < N - 1$ regime. In the context of scientific problem solving, we can think of these binary strings as different research approaches to a topic of scientific inquiry, and the N digits of each string as possible components of these approaches, e.g. research questions, lab instruments, standards of induction, etc. The epistemic significance of an approach, then, depends on how well its components work together.

I now turn to the network structure and behavioral rules of my model. The model I construct is a variation of Lazer and Friedman (2007)'s, with the addition of asymmetric evidence sharing.

At the start of each simulation, we have 100 agents, connected to each other via a directed Erdős-Rényi random network with a probability of connection b, as generated by the algorithm in Footnote 19.[23] This means that each agent on average has $100 \cdot b$ neighbors. I divide the agents into two subgroups, the withhold group and the share group. Every agent is randomly assigned a starting solution, in the form of a binary string of length N.

Every V rounds, the epistemic community goes through social learning.[24] When this happens,

[22]Note that though the figure is two dimensional, it does not represent the problem where $N = 2$. Rather, it provides a stylized representation of the complexity of the problem as K increases, for an unspecified $N > 2$.

[23]I use directed networks to allow for situations where an agent is aware of a solution of another agent, but not vice versa, even if they are in the same group. The results I report does not depend on the directedness of the network structure.

[24]I start the simulation at round 1, and require that the community go through social learning when the current round is divisible by V. This ensures that agents go through a few rounds of local search before starting social learning, unless $V = 1$.

each member of the withhold group shares their solution and its score with their in-group neighbors, and each member of the share group shares their solution and its score with all their neighbors. Then, every agent in the community looks at all solutions shared with them, and chooses the solution with the highest score to copy, if it is higher than their own.[25] If their own score is higher than all those shared with them, then the agent goes through a local search. This means that they randomly choose a bit in their binary string to alter (1 to 0 or 0 to 1), and, if the altered string has a higher score, they switch to that solution. Otherwise, they maintain their current location. In other rounds (rounds indivisible by V), each agent conducts a local search to try to improve their score. In scientific problem solving, local searches represent situations where each scientist or lab tries to improve their own research approach by making changes to one component of their approach. Note that when $V = 1$, agents undergo social learning every round. I simulate the model for 200 rounds. As we will see, in most simulations, the community reaches stable behaviors much more quickly.

In the NK landscape model, we have a community of agents searching in a vast epistemic landscape with multiple "peaks," so depending on how agents search, they may fail to ever discover the global optimum. There are two characteristics that distinguish the NK landscape model from the bandit model. First, instead of the two options considered in the bandit model, in the NK landscape model we have a myriad of solutions.[26] Second, in the NK landscape model, each solution's epistemic significance is readily known to an agent who chooses it, whereas in the bandit model both options are probabilistic, making their success rates harder to estimate. The exploration and exploitation trade-off for agents in this model is thus the following: do they exploit the epistemic significance of their current solution, or do they keep exploring the landscape in the hope of finding a better solution? Many authors (e.g. March (1991), Lazer and Friedman (2007), and Fang et al. (2010)) find results similar to the Zollman effect in the NK landscape model, and these results are empirically confirmed to some extent (Mason and Watts 2012; Derex et al. 2018). That is, they find that in the NK landscape model too, a less connected community may end up with better solutions. This is because in more connected networks, agents may quickly settle onto a local optimum, thus failing to explore better alternatives elsewhere. In less connected networks, agents

[25] If there are multiple solutions with the highest score, they randomly select one to copy.

[26] In the $N = 20$ case considered in this paper, there are $1,048,576$ solutions.

retain a diverse range of solutions for longer, and thus may ultimately discover better solutions.

Simulating the asymmetric evidence-sharing dynamics in the NK landscape model, therefore, serves as an important robustness check. Moreover, we may think that many of the real scientific problems involve *a wide range of probabilistic options* (Wu and O'Connor 2021), not just a wide range of certain options, as in the case of the NK landscape model, or a limited range of probabilistic options, as in the case of a two-armed bandit model. A replication of my qualitative results here may thus increase our confidence that similar results would hold in a model that combines features of the bandit model and NK landscape model.[27] As I show shortly, the qualitative results exhibit similar patterns across the two modeling paradigms, though some results are more robust than others.

3.1 Results and Discussions

I run the model $1,000$ times for each parameter combination:

- N: 20;[28]
- K: $5, 10, 15$;[29]
- V: $1, 3, 5$;
- Proportion of the withhold group (d): $.2, .4, .6, .8$;
- Probability of connection in directed Erdős-Rényi random networks (b): $.4, .6, .8, 1$.[30]

In addition, for each parameter combination, I run the model of a communist community where everyone shares with everyone.[31] As noted before, each community has 100 agents. The results I present are robust across parameter values unless otherwise noted.

The first result is that the withhold group ends up having better scores on average than the share group. This result always holds because the withhold group has the option of switching to the better solutions from the share group, but not vice versa. Because of this, the withhold group always ends up with at least as good a score as the share group at the end of each simulation.

[27] Such models may take the form of a bandit model with sufficiently many arms, or an epistemic landscape model where the score of each solution is probabilistic. I leave the detailed modeling work to further research.

[28] N here is different from n in §2.1. n is the number of tests per round in the bandit model.

[29] K here is different from k in §2.1. k is the number of agents in the bandit model.

[30] When $p = 1$, the network structure is complete.

[31] This is equivalent to Lazer and Friedman (2007). For ease of comparison, the two communities always search in the same solution space, and have exactly the same random network structure and initial solutions, all of them randomly generated for each simulation.

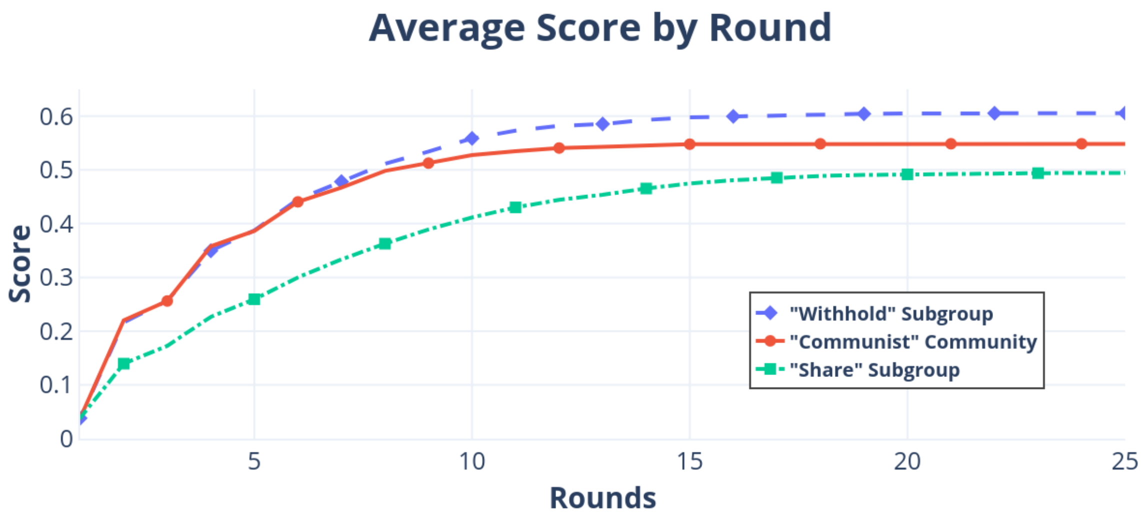

Figure 5: $N = 20, K = 5, V = 1, d = .8, b = .6$. $1,000$ simulation runs. All three groups reach stable scores after 25 rounds.

For almost all ($> 97\%$) parameter combinations, the share group ends up having worse scores on average than the communist community. This is because the share group is smaller in size than the communist community and thus suffers from two epistemically detrimental consequences. First, the initial solutions present in the share group are not as diverse as in the communist community, so high scoring solutions may be "farther away" from agents in the share group. Second, agents in the share group on average have fewer neighbors than agents in the communist community. So when an agent in the share group and an agent in the communist community both have the best solution among their neighbors, it takes longer for the neighbors of the former agent to conduct a thorough local search for better solutions. In the face of unpromising results in early local search, then, these agents may give up and switch to another solution from elsewhere in the community. Figure 5 shows the average scores per round for different groups.

For most sets of parameter values ($> 56\%$) the withhold group ends up with better scores on average than the communist community. When this happens, the withhold group benefits from the extra exploration done by the share group, while the communist community converges to a suboptimal solution too quickly. This is a similar mechanism for epistemic advantage as that identified in the bandit model, but the withhold group's advantage here is not as robust.

To see why, let us consider a simplified case. Suppose that we have four solutions, A, B, C, and D. Suppose further that $A < B < C < D$ in epistemic significance, but D is only available when we explore from B, and C is only available when we explore from A. Now suppose that A is the current

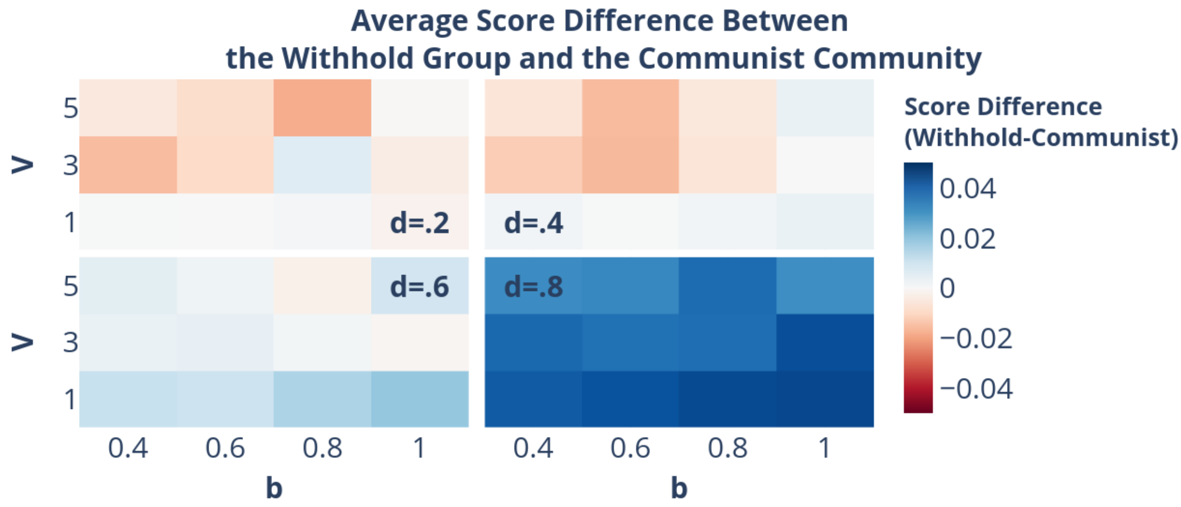

Figure 6: $N = 20$, $K = 10$, $1,000$ simulation runs. Gray-scale version available upon request.

best solution among the share group, and B is the current best solution among the withhold group. Then in cases where the share group discovers C before the withhold group discovers D, the withhold group would switch to C prematurely, without ever discovering D. Curiously, when a communist community encounters this case, they would not make the same mistake, since agents would choose B over A first, and the community would, after sufficient local searches, settle on D. Qualitatively similar cases can arise in the NK landscape model, and in these situations, the withhold group learns worse than the communist community. The same situation does not apply to the two-armed bandit problems.

For this reason, when d is small, i.e. when the withhold group is in the minority, the withhold group is more likely to lose its epistemic advantage compared to the communist community. The withhold group does not have enough agents to do as thorough a local search around the most promising solution as the communist community, given the vast number of possible local variations on a solution. They may instead give up on their local search if the more populous share group reaches a seemingly promising solution in the meantime. As d increases, the withhold group becomes more efficient at local search, and it gains epistemic advantage.[32]

Moreover, as V increases, i.e. when social learning becomes less frequent, the withhold group tends to have less of an epistemic advantage. To understand this, we observe that when $V = 1$, because of the quick social learning, the communist community is more likely to prematurely settle

[32]The withhold group does better than the communist community 100% when $d = .8$.

down to a local optimum, without ever exploring other regions of the landscape. In contrast, in the withhold-share community, diverse solutions are still present even when social learning is fast, because the withhold group does not share. In cases where solutions among the share group are initially unpromising, its members can keep exploring them, and the withhold group benefits if the share group's exploration proves fruitful. However, when V increases, both communities know more about the landscape terrain before commencing social learning. In this case, the communist community simply has more agents conducting local searches around promising solutions after social learning, making it more likely that it would find something.

Finally, as K increases, the withhold group tends to lose its epistemic advantage over the communist community as well. The reason for this is slightly different from before. When K grows, the landscape becomes increasingly chaotic, and the epistemic communities do not typically learn very well at all.[33] In these "hostile" landscapes, it is again better to have more agents conduct targeted local exploration around one solution, as in the case of the communist community, than dividing up labor such that neither subgroup has enough agents to do thorough local searches, as in the case of the withhold-share community. One might find this result surprising because previously Rosenstock et al. (2017) finds that the Zollman effect is the most robust in hard bandit problems, but here it seems that the opposite holds in the NK landscape problem. This conclusion may be too quick, because here what is holding the withhold-share community back is the insufficient *local* exploration caused by the size of the subgroups, not a lack of global diverse solutions. This nonetheless points to a significant difference between the bandit model and the NK landscape problem.

Figure 6 shows the degree of the withhold group's epistemic (dis-)advantage over the communist community as we vary several parameters.

4 Epistemic Prisoner's Dilemma

Results from §2-§3 suggest that a subgroup that withholds evidence can have epistemic advantages over the rest of the epistemic community that shares, and over a communist community in many cases. If scientists are purely motivated by the truth of their own findings, the withhold-share dynamics gives

[33] When $K = 15$, all communities and subgroups' average final scores do not exceed $.4$.

rise to what one might call an Epistemic Weak Prisoner's Dilemma (EWPD).[34]

The game is set up as follows. We have two groups who are the primary players of the game. Group members internally share their evidence with each other. On top of that, each group has two strategies, sharing evidence with out-group members (Share), and withholding evidence from out-group members (Withhold). Each group's payoff is epistemic, corresponding to the probability that the group converges to the truth in that situation. Specifically, the epistemic payoffs in Tables 1 and 2 are calculated from the frequencies of successful learning measure from the bandit model in §2.[35] By assuming that members of the community are already sorted into different groups or coalitions, this game is an instance of a cooperative game (see §9 of Ross (2021)). Because of this assumption, the game may be interpreted to represent certain groups (e.g. industrial scientists) better than others (e.g. academic scientists), as I will discuss shortly.

Symmetric EWPD		Group 2 (1/2)	
		Share	Withhold
Group 1 (1/2)	Share	.58, .58	.56, .63
	Withhold	.63, .56	.56, .56

Table 1: $k = 12$, $n = 1$, $P_B = .49$. $10,000$ simulation runs. Groups 1 and 2 are both $1/2$ the size of the community.

Asymmetric EWPD		Group 2 (1/3)	
		Share	Withhold
Group 1 (2/3)	Share	.58, .58	.57, .61
	Withhold	.67, .55	.57, .55

Table 2: $k = 12$, $n = 1$, $P_B = .49$. $10,000$ simulation runs. Group 1 is $2/3$ the size of the community, and Group 2 is $1/3$.

In the EWPD, each group receives the highest payoff when they withhold evidence from the other group while the other group shares with them. Their payoffs are the worst when the other group withholds evidence from them (no matter what they themselves choose to do), and their payoff is intermediate when both groups share.

There are two features of this game worth noting. First, the game is symmetric only when each group is exactly half of the entire community in size (Table 1); otherwise, the game is asymmetric

[34] In so doing, I provide new examples of how game theory can be used as a tool for the social epistemologists, see Zollman (2021).

[35] For Share-Share, I calculate the frequency of successful learning for the communist community. For Withhold-Withhold, I calculate the same frequency for an isolated communist community of a smaller size.

(Table 2). In the latter case, the group that is the majority in size would gain more when they withhold evidence while the other group shares, and lose less when both subgroups withhold. This may seem counter-intuitive, given that in a public goods game, where players decide whether to contribute to a public pot which will then grow and be distributed to everyone, players have the incentive to be in the very few that withhold. However, in the EWPD, members of each group still share with each other, and this mechanism is absent in a public goods game. Furthermore, social knowledge is unlike public goods. It is not the case that the more people share, the better the group is at learning—the Zollman effect tells us otherwise.

The second feature is that this prisoner's dilemma is *weak*, since the situation where both groups withhold is a weak Nash equilibrium, not a strict one. The situations where one group switches their strategy (top right and bottom left) are two additional weak Nash equilibria of this game, and they are both Pareto improvements from the situation where both groups withhold, since one group would gain and no groups lose.[36] Given this feature, we might imagine a new game where both groups are slightly altruistic, and their payoff is calculated by .9 of their own original epistemic payoff, plus .1 of the other group's original epistemic payoff in the same situation. Then, both situations where one group withholds and the other shares are (strict) Nash equilibria in this game, and Withhold-Withhold is no longer Nash (see Table 3). One might call this new game an Epistemic Hawk-Dove. Furthermore, we might imagine another game where only one group is slightly altruistic. The result is an Asymmetric Epistemic Hawk-Dove (Table 4), and the situation where the altruistic group shares and the other self-interested group withholds is the only strict Nash equilibrium.[37]

Epistemic Hawk-Dove		Group 2 (1/2)	
		Share	Withhold
Group 1 (1/2)	Share	.58, .58	.567, .623
	Withhold	.623, .567	.56, .56

Table 3: New game constructed from Table 1. Both groups are slightly altruistic.

One might take this Asymmetric Epistemic Hawk-Dove to explain our current situation where slightly altruistic academic scientists generally share their evidence while self-interested industrial scientists withhold.[38] However, the strength of this explanation depends on whether we can

[36]The EWPD can also be seen as a weak Hawk-Dove, with Withhold-Withhold as an additional weak Nash equilibrium.

[37]The other Nash equilibrium, where the altruistic group withholds and the self interested group shares, is weak.

[38]For those who are skeptical that academic scientists are altruistic, we can instead introduce academic scientists' credit

Asymmetric Epistemic Hawk-Dove		Group 2 (1/2)	
		Share	Withhold
Group 1 (1/2)	Share	.58, .58	.567, .63
	Withhold	.623, .56	.56, .56

Table 4: New game constructed from Table 1. Group 1 is slightly altruistic.

successfully interpret academic scientists and industrial scientists as group agents able to play the two strategies. While the industrial scientists have centralized decision-making power, it is difficult to coordinate decentralized academic scientists to play the Withhold strategy, especially given that they have additional credit incentive to publish.[39] The assumption that academic scientists can be thought of as a group agent in a cooperative game may not be realistic. Instead, our current situation may be closer to one where academic scientists are stuck with playing the Share strategy, and the industrial scientists, having centralized decision-making power, choose to Withhold to maximize their epistemic payoffs.

Given that industrial scientists can be interpreted as group agents, are there ways in which industrial scientists may come to share from a game theoretic perspective? There is a large literature on how cooperative behaviors could evolve in a prisoner's dilemma (See, e.g., Chapter 6 of O'Connor (2020) for an overview). The key is that when players' actions become correlated—when it is likely that sharers only interact with sharers and withholders with withholders—cooperative behaviors would be expected to evolve over time. One way to achieve correlated interaction is when players keep track of the actions of other players, and reciprocate their actions accordingly (See, e.g., Trivers (1971), Skyrms (2001), and Binmore (2005)). For instance, industrial scientists may start by sharing their evidence, then on successive rounds choose the action that their partner performed before (a strategy called Tit-For-Tat). Other ways to achieve correlated interaction include secret handshakes (sending a signal to identify cooperators before playing, see Robson (1990)), network reciprocity (players disproportionately interact with certain "neighbors" and adopt a reciprocal

incentive to the EWPD by adding a small payoff to one group in situations that they share, representing the recognition they receive when they publish. The result would be an asymmetric Hawk-Dove as well. Furthermore, we can introduce industrial scientists' financial incentive to the game by adding a small payoff to the second group in situations that they withhold. The two additions together would produce a game where the situation where the first group shares and the second group withholds is the only Nash equilibrium, strict or weak. Thanks to Hannah Rubin for thinking through this point with me.

[39]Thanks to Cailin O'Connor for raising this point. The case for classified scientists is a bit complicated. We might think that they are the opposite of academic scientists, since they have centralized decision-making power, but are stuck with the Withhold strategy. Analyzing how this complicates sharing dynamics is beyond the scope of this paper.

strategy, see Alexander (2007)), etc. Under these mechanisms, cooperative behaviors are expected to evolve. But note that this analysis only outlines how mutual sharing can be evolutionarily advantageous between an industry group and other similarly agential groups (e.g. other industry groups). When interacting with a decentralized group that is stuck with the Share strategy (e.g. academic scientists), the best action for an industry group is still Withhold.

Another possibility comes from the observation that in my models, the withholding agents are mutually sharing, but in reality, industrial groups may not share with each other. One might think that if there are multiple mutually withholding industrial groups, then even though they each learn better than the sharing academic scientists, they could reach even better beliefs if the whole community comes to share. This reasoning is (surprisingly) not supported by simulation results. Suppose that we have an epistemic community with three equal-sized subgroups, X that shares with everyone, and Y and Z that only share with in-group members and no one else. My simulation results show that in many cases Y and Z reach the true belief even more frequently than in the situation where X, Y, and Z are all mutually sharing.[40] In other words, the epistemic advantage that Y and Z gain from free-riding off of X's exploration is strong enough that they can even outperform a communist community of a larger size. Furthermore, even if Y and Z *would* perform better if the whole community is communist, they would perform even better if they mutually share but continue to withhold from X, especially when X is stuck with the Share strategy.

Finally, throughout the section, I relied on simulation results from the bandit model in §2 to construct the payoff tables, but we know very little about the problem space of our scientific inquiries (Wu and O'Connor 2021). As §3 shows, for a considerable portion of the parameter space, especially when the withhold group is in the minority, it does not end up scoring better than the communist community in the NK landscape problem. If industrial scientists constitute a minority of the scientific community, then since they may not have enough powers to conduct thorough local searches, perhaps it would be rational for them to share. Then, the base game may instead be more like a stag hunt where mutual sharing offers the best payoffs. However, industrial science is rather sizable. For instance, in the US, industry conducts 75% and funds 72% of research and experimental development in 2019, according to the National Science Foundation (Boroush and Guci 2022). In this case, industrial

[40] I tested this situation in the bandit model with 18 total agents, where $P_B = .45$ and $.49$, and $n = 1$.

scientists would still gain an epistemic advantage if they withhold. Moreover, there may be other advantages associated with withholding than the ones modeled here, such as reaching conclusions faster so they can reap financial benefits sooner. These additional advantages may shift the game back to a prisoner's dilemma, even if the problem space is more like an NK landscape.

5 Conclusion

In this paper, I use network models to show that, scientists who withhold evidence can obtain *epistemic* advantages. My results suggest that industrial scientists and classified researchers may have additional epistemic incentives to withhold their work, even if the academic credit economy does not apply to them in the same way. I further use my modeling results to construct epistemic games that illustrate the underlying sharing dynamics. Before closing, I will briefly discuss some implications for other models of scientific sharing, connect my models to topics in social epistemology, note limitations of my analysis, and suggest directions for future work.

To start, my modeling results may complicate previous credit-based models of scientific sharing. Both Strevens (2017) and Heesen (2017)'s models assume that the probability that a scientist would make a discovery remains fixed regardless of whether they share during the process of discovery. But my results show that the probability of a scientist making the *right* discovery (as in §2) or an *epistemically more significant* discovery (as in §3) does depend on whether they share. In the context of Heesen (2017)'s multi-stage model, if agents consistently share throughout the process of discovery, then they may risk prematurely settling on a worse theory, without ever discovering better ones. Furthermore, in the situation where one agent consistently withholds while the rest of the community shares, then even though the sharing agents can publish and claim credit for intermediate results as Heesen (2017) argues, the withholding agent may reach the final stage much faster, likely with epistemically more significant findings, which could bring them even more credit.

Moreover, my models offer new support for the independence thesis in social epistemology, which states that individual and group rationality can come apart (Mayo-Wilson et al. 2011).[41] Indeed, in my models, a subgroup may be epistemically better off if they withhold evidence, but the epistemically optimal situation for the entire community is when everyone shares. The action that maximizes a

[41] See Bradley (2022) for an analysis of different versions of the independence thesis.

subgroup's epistemic payoff and the action that maximizes the community's epistemic payoff come apart.

My analysis admittedly has a few limitations. First, agents in both models all choose the best available epistemic option when they undergo social learning. However, recent work (e.g. Kummerfeld and Zollman (2015) and Wu (manuscript)) shows that alternative behavioral rules can significantly change epistemic outcomes. Combining these alternative behavioral rules with the asymmetric evidence-sharing dynamics is a worthwhile direction for future research. Moreover, I have only tested a small range of cases where there are multiple withhold groups. It is worthwhile to conduct a fuller analysis with more than two groups, since my limited testing shows that the dynamics may be more complicated than expected.

Furthermore, while it is clear from my analysis that epistemic consideration alone does not incentivize a group to share, I have said very little about what to do as a result. My analysis may support policies that require industrial scientists to share for the benefit of the community, or mechanisms that make academic scientists more agential as a group, or recommendations that ensure a small amount of exploration in the community to reduce the epistemic drawbacks of mutual sharing. We need models that combine instrumental and epistemic incentives of scientific sharing before fuller recommendations are made.

Finally, one may find the differential robustness levels between the bandit models and the NK landscape model independently interesting from a philosophy of modeling perspective. It is worthwhile to explore whether there is a similar difference in robustness for other mechanisms that lead to transient diversity (e.g. those reviewed in Wu and O'Connor (2021)). This can offer insight into how the exploration and exploitation trade-offs differ in these models, and how results obtained from either model can be used to represent and explain reality.

Appendix A End States of the Generalized Bandit Model

After $10,000$ rounds, the community typically settles into one of three states, defined below:

- **Community convergence to the true belief** is reached when, for each agent, the following three conditions are satisfied: (S1) $\frac{\alpha_A}{\alpha_A+\beta_A} > \frac{\alpha_B}{\alpha_B+\beta_B}$; (S2) $\frac{\alpha_A}{\alpha_A+\beta_A} > P_B$; (S3) $\frac{\alpha_B}{\alpha_B+\beta_B} < P_A$. The first condition ensures that all individual agents think that option A performs better than option B. The third condition ensures that agents' actions are, for the most part, stable, because agents continue to receive evidence about A.[42] The second condition is included for consistency reasons—we will need it as a stability check in simulations that lead to polarization.[43]

- **Community convergence to the false belief** is reached when, for each agent, the following two conditions are satisfied: (F1) $\frac{\alpha_A}{\alpha_A+\beta_A} < \frac{\alpha_B}{\alpha_B+\beta_B}$; (F2) $\frac{\alpha_A}{\alpha_A+\beta_A} < P_B$. Condition (F1) says that all agents think that the worse action, B, is better. (F2) is a stability condition similar to (S3).

- **Polarization** is reached when, for every member of the withhold group, the success conditions are satisfied, and for every member of the share group, the failure conditions are satisfied. Condition (S2) is needed because the withhold group continues to receive evidence about B from the share group.

Appendix B The Solution Space of the NK Landscape Model

I now describe how scores are assigned to solutions in the NK landscape model. I run this algorithm at the start of each simulation. So even for the same value of N and K, solution spaces are different in each simulation.

I start with randomly selecting a K-tuple, $J = (j_1, ..., j_K)$, where $j_1, ..., j_K$ are integers between 1 and N inclusive, without repeating. I then generate a valuation function f that takes binary strings of length $K+1$ to $(0,1)$ by assigning each possible binary combination a random number from $(0,1)$. For instance, if $K = 2$, then $f(0,0,1) = .549$, $f(0,1,0) = .235$, $f(1,1,0) = .652$, etc. would be a valuation function. I randomly generate J and f at the start of every simulation, and they remain fixed throughout the simulation.

Remember that each solution S is a binary string of length N, i.e. $S = (s_1, s_2, ..., s_N)$, where $s_i \in \{0,1\}$ for $1 \leq i \leq N$. Each S can now be associated with a score, via the following function

$$F(S) = \frac{1}{N}(\sum_{i=1}^{N} f(s_i, s_{i+j_1}, ..., s_{i+j_K})).$$

[42] Because of the probabilistic nature of this model, this condition does not ensure that individual actions are always stable when it is satisfied.

[43] As a technical aside, the "success conditions" listed here are not the analogue of the success condition discussed in Wu (2022), in the following sense. In Wu (2022), each agent's credence is given by a single number between 0 and 1, representing their credence in the statement "B is better than A." Then, an agent reaches the success condition if this number is $> .99$ (in that model, B is in fact better than A). For the success condition in our current model to be analogous to the one before, we need to calculate the probability that one beta distribution performs better than the other beta distribution and ask whether this probability is $> .99$. This amounts to calculating the following formula

$$Pr(p_B > p_A) = \int_0^1 \int_{P_A}^1 \frac{{p_A}^{\alpha_A - 1}(1-p_A)^{\beta_A - 1}}{B(\alpha_A, \beta_A)} \frac{p_B^{\alpha_B - 1}(1-p_B)^{\beta_B - 1}}{B(\alpha_B, \beta_B)} dp_B dp_A.$$

It turns out that even after simplification, this formula is still too computationally taxing. So I adopted the current success conditions. See Miller (2015) for details on Bayesian A/B testing.

Here, for cyclicity, we set $s_l = s_{l-N}$ for $l > N$. What this function does is to calculate the score of a point as a sum of the "value" of each bit of the string, while taking into consideration its interaction with K other bits. This allows for correlations among nearby solutions to be reflected in the final score. The score is then normalized (multiplied by $\frac{1}{N}$).

To give an example of this, suppose that $N = 3$, $K = 1$, and $J = (1)$. Suppose further that $f(0,0) = .235$, $f(0,1) = .367$, $f(1,0) = .785$, and $f(1,1) = .954$. Then the score of the binary string $S = (0,1,1)$ is

$$F(S) = \frac{1}{3}(\sum_{i=1}^{N} f(s_i, s_{i+1})) = \frac{1}{3}(f(0,1) + f(1,1) + f(1,0)) = \frac{1}{3}(.367 + .954 + .785) = .702.$$

In general, every landscape generated is going to have a different maximum score. For ease of comparison across simulations, I further divide the score of every point by the maximum score in the landscape to normalize it, i.e. $\tilde{F}(S) = F(S)/\text{MaxScore}$. This way, the highest score in every landscape generated is 1.

Finally, because of how the NK landscape model is generated, the distribution of scores is similar to that of a normal distribution, where the majority solutions are moderately good at solving a problem. This can be unrealistic since usually haphazard solutions to a problem do not perform well. I thus follow Lazer and Friedman (2007) and apply the following transformation $\hat{F}(S) = (\tilde{F}(S))^\wedge 8$ to every score, such that most solutions receive a low score, and high scores are distinguished. Note that none of the qualitative results I present depend on this transformation, since the transformation preserves the order among the scores. $\hat{F}(S)$ is what I mean by score throughout §3.

References

Alexander, Jason McKenzie (2007). *The structural evolution of morality*. Cambridge University Press.

Alexander, Jason McKenzie, Johannes Himmelreich, and Christopher Thompson (2015). "Epistemic landscapes, optimal search, and the division of cognitive labor". In: *Philosophy of Science* 82.3, pp. 424–453.

Binmore, Ken (2005). *Natural justice*. Oxford university press.

Boroush, Mark and Ledia Guci (2022). "Research and Development: US Trends and International Comparisons." In: *National Science Foundation*. URL: https://ncses.nsf.gov/pubs/nsb20225/.

Bradley, Clara (2022). "The Many Independence Theses". In:

Brainard, Jeffrey and Jocelyn Kaiser (2022). "US to require free access to papers on all research it funds". In: *Science* 377.6610, pp. 1026–1027.

Bright, Liam Kofi (2017). "On fraud". In: *Philosophical Studies* 174.2, pp. 291–310.

Bright, Liam Kofi and Remco Heesen (2021). "To be Scientific is to be Communist". In:

DeAngelis, Tori (2003). "Does industry funding deserve a bad rap?" In: *Monitor on Psychology* 34.7. URL: https://www.apa.org/monitor/julaug03/industry.

Derex, Maxime, Charles Perreault, and Robert Boyd (2018). "Divide and conquer: intermediate levels of population fragmentation maximize cultural accumulation". In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 373.1743, p. 20170062.

Fang, Christina, Jeho Lee, and Melissa A Schilling (2010). "Balancing exploration and exploitation through structural design: The isolation of subgroups and organizational learning". In: *Organization Science* 21.3, pp. 625–642.

Fazelpour, Sina and Daniel Steel (2022). "Diversity, Trust, and Conformity: A Simulation Study". In: *Philosophy of Science* 89.2, pp. 209–231.

Galison, Peter (2004). "Removing knowledge". In: *Critical Inquiry* 31.1, pp. 229–243.

Heesen, Remco (2017). "Communism and the Incentive to Share in Science". In: *Philosophy of Science* 84.4, pp. 698–716.

Hong, Lu and Scott E Page (2004). "Groups of diverse problem solvers can outperform groups of high-ability problem solvers". In: *Proceedings of the National Academy of Sciences* 101.46, pp. 16385–16389.

Kauffman, Stuart and Simon Levin (1987). "Towards a general theory of adaptive walks on rugged landscapes". In: *Journal of theoretical Biology* 128.1, pp. 11–45.

Kauffman, Stuart and Edward Weinberger (1989). "The NK model of rugged fitness landscapes and its application to maturation of the immune response". In: *Journal of theoretical biology* 141.2, pp. 211–245.

Kummerfeld, Erich and Kevin JS Zollman (2015). "Conservatism and the scientific state of nature". In: *The British Journal for the Philosophy of Science* 67.4, pp. 1057–1076.

Lazer, David and Allan Friedman (2007). "The network structure of exploration and exploitation". In: *Administrative science quarterly* 52.4, pp. 667–694.

Louis, Karen Seashore, Lisa M Jones, and Eric G Campbell (2002). "Macroscope: Sharing in science". In: *American Scientist* 90.4, pp. 304–307.

Macfarlane, Bruce and Ming Cheng (2008). "Communism, universalism and disinterestedness: Re-examining contemporary support among academics for Merton's scientific norms". In: *Journal of Academic Ethics* 6.1, pp. 67–78.

March, James G (1991). "Exploration and exploitation in organizational learning". In: *Organization science* 2.1, pp. 71–87.

Mason, Winter and Duncan J Watts (2012). "Collaborative learning in networks". In: *Proceedings of the National Academy of Sciences* 109.3, pp. 764–769.

Mayo-Wilson, Conor, Kevin JS Zollman, and David Danks (2011). "The Independence Thesis: When Individual and Social Epistemology Diverge". In: *Philosophy of Science* 78.4, pp. 653–677.

McGarity, Thomas O and Wendy E Wagner (2010). *Bending science: How special interests corrupt public health research*. Harvard University Press.

Merton, Robert K (1973). *The Sociology of Science: Theoretical and Empirical Investigations*. University of Chicago Press.

Michaels, David (2008). *Doubt is their product: how industry's assault on science threatens your health*. Oxford University Press.

Miller, Evan (2015). *Formulas for Bayesian A/B Testing*. https://www.evanmiller.org/bayesian-ab-testing.html. Accessed: 2021-10-08.

O'Connor, Cailin (2020). *Games in the Philosophy of Biology*. Cambridge University Press.

O'Connor, Cailin and James Owen Weatherall (2018). "Scientific Polarization". In: *European Journal for Philosophy of Science* 8.3, pp. 855–875. DOI: 10.1007/s13194-018-0213-9.

Partha, Dasgupta and Paul A David (1994). "Toward a new economics of science". In: *Research policy* 23.5, pp. 487–521.

Robson, Arthur J (1990). "Efficiency in evolutionary games: Darwin, Nash and the secret handshake". In: *Journal of theoretical Biology* 144.3, pp. 379–396.

Rosenstock, Sarita, Justin Bruner, and Cailin O'Connor (2017). "In Epistemic Networks, is Less Really More?" In: *Philosophy of Science* 84.2, pp. 234–252.

Ross, Don (2021). "Game Theory". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Fall 2021. Metaphysics Research Lab, Stanford University.

Skyrms, Brian (2001). "The Stag Hunt". In: *Proceedings and Addresses of the American Philosophical Association*. Vol. 75. 2, pp. 31–41.

Strevens, Michael (2017). "Scientific Sharing: Communism and the Social Contract". In: *Scientific collaboration and collective knowledge*, pp. 1–50.

Trivers, Robert L (1971). "The evolution of reciprocal altruism". In: *The Quarterly review of biology* 46.1, pp. 35–57.

Weisberg, Michael and Ryan Muldoon (2009). "Epistemic landscapes and the division of cognitive labor". In: *Philosophy of science* 76.2, pp. 225–252.

Wu, Jingyi (2022). "Epistemic Advantage on the Margin: A Network Standpoint Epistemology". In: *Philosophy and Phenomenological Research*.

— (manuscript). "Better than Best: Epistemic Landscapes and Diversity of Practice in Science". In:

Wu, Jingyi and Cailin O'Connor (2021). "How Should We Promote Transient Diversity in Science?" In: URL: http://philsci-archive.pitt.edu/19428/.

Zollman, Kevin JS (2007). "The communication structure of epistemic communities". In: *Philosophy of science* 74.5, pp. 574–587.

— (2010). "The epistemic benefit of transient diversity". In: *Erkenntnis* 72.1, p. 17.

— (2018). "The credit economy and the economic rationality of science". In: *The Journal of Philosophy* 115.1, pp. 5–33.

— (2021). "The theory of games as a tool for the social epistemologist". In: *Philosophical Studies* 178.4, pp. 1381–1401.