Structure-sensitive testimonial norms

Benedikt T. A. Höltgen¹

Received: 16 July 2020 / Accepted: 16 March 2021 / Published online: 29 July 2021
© The Author(s) 2021

Abstract

Although there has been a lot of investigation into the influence of the network structure of scientific communities on the one hand and into testimonial norms (TNs) on the other, a discussion of TNs that take the network structure into account has been lacking. In this paper, I introduce two TNs which are sensitive to the local network structure. According to these norms, scientists should give less weight to the results of well-connected colleagues, as compared to less connected ones. I employ an Agent Based Model to test the reliability of the two novel TNs against different versions of conventional, structure-insensitive TNs in networks of varying size and structure. The results of the simulations show that the novel TNs are more reliable. This suggests that it would be beneficial for scientific communities if their members followed such norms. For individual scientists, I show that there are both reasons for and reasons against adopting them.

Keywords Social epistemology · Network epistemology · Testimonial norms · Scientific networks · Zollman effect

1 Introduction

For centuries, the prevalent image of science has been that of solitary men working in secluded studies or laboratories. At least since the seminal work by Thomas Kuhn (1962), philosophers and historians of science have paid more attention to the social and institutional structures that are essential for scientific enterprise. Much has since been learned about these structures, one notable insight being the importance of diversity. While Kuhn emphasized scientists' individual differences in assessing scientific theories and weighting their theoretical values, others have advocated models

✉ Benedikt T. A. Höltgen
b.hoeltgen@campus.lmu.de

¹ Munich Center for Mathematical Philosophy, Ludwig-Maximilians-Universität München, Ludwigstr. 31, 80539 München, Germany

of scientific practice where individuals can share an epistemic strategy yet still scatter among scientific fields and theories. For example, Kitcher (1990) and Strevens (2003) argue that the incentive structures in science contribute to the division of scientific labor. Although Kuhn may well be right that individual differences contribute to a diversification of research agendas, Kitcher's and Strevens' ideas also offer avenues for fruitful investigation into incentive structures and the process of theory choice in scientific communities.

One rather new avenue in this area is that of social epistemology, and the investigation of information networks in particular. In a series of papers culminating in Zollman (2013), Kevin Zollman investigates how scientific networks should be structured in order to provide the optimal basis for a successful division of labor. To this end, he uses computer simulations and mathematical analyzes to investigate how structural properties of these networks influence the communities' convergence to objectively better theories. These modes of investigation are motivated by the observation that since the processes in question display "properties of complex systems, we cannot hope to understand them from mere reflection alone" (Zollman, 2012, 27). In the literature on the division of labor, the question of what exactly is divided is not answered uniformly, and often remains vague (De Langhe, 2014). This also applies to Zollman's work: in his (2007) and (2010), he variously speaks of the choice between treatments, theories, and methodologies. Bedessem (2019) rightly notes that this vagueness can be problematic when it comes to arguments about how the division of labor should be organized, e.g. by providing incentives through funding. In defense of abstract models, one may note that the applicability of such a model to a particular case (e.g. whether it fits the discipline) always needs to be justified individually, so the model itself does not need to commit to one interpretation. In fact, some level of generality can be desirable as it makes the model potentially more widely applicable. More importantly, the potential problems arising for policy advice discussed by Bedessem do not directly affect the present work in the first place, as the choice of a testimonial norm is a choice for individual scientists to make. Therefore, though acknowledging the vagueness, I follow Kummerfeld and Zollman (2016, 1059f) in postulating that the division in question may concern "different theoretical commitments, paradigms, research methodologies, treatments strategies in medicine, and so on". In his work, Zollman (2007, 2010) shows that less connected networks are more resilient to errors and that "regular networks are best when individuals are all equally reliable" (Zollman, 2012, 41). In other words, networks should be structured such that all individuals have the same degree, i.e. the same number of connections to other individuals. However, restricting information (or connectivity) to yield a uniform degree distribution, as Zollman suggests for scientists (e.g. in Zollman 2010), does not seem to be a realistic option.

In reality, scientific networks often fall short of the – prima face plausible – expectation that equally reliable scientists would have equal connectivity. One particularly clear example is the case of mid-twentieth-century psychology. Among the "four founders of behavioral psychology" (Hull, 1990, 368), which was the dominant paradigm at the time, the central figure was undisputedly B.F. Skinner. In his field, he was the most successful in attracting students and followers both through personal

and impersonal channels, as observed by Krantz and Wiggins (1973). Most striking, also in the light of their study, is the comparison with E.C. Tolman.¹ According to one of the study's findings, Skinner was the most successful among the four in getting his own students to work within his theory, while Tolman is tied for the last place (along with Clark Hull). Like D.T. Campbell (1979), one might ask how this squares with the observations that "of all the learning theories of the 1930s, Tolman's can now be seen to have been clearly the best" (p. 187) and that he was the "most personally beloved of all the four theorists" (p. 188).² In a similar vein – though formulated more cautiously – D.L. Hull (1990) "doubt[s] that professional psychologists today would evaluate the quality of the work of these four men in this same order as the degree to which they were successful" (p. 368). Peter Godfrey-Smith (2003) even goes so far as to "conjecture, and many psychologists would agree, that if Tolman had dominated mid-twentieth-century psychology rather than Skinner, it would have been far better for the field" (p. 171). While the specific historical details, let alone the counterfactuals, might be considered debatable, it seems fairly uncontroversial to assert that the field would have benefitted from more psychologists following in the footsteps of the less influential theorists, such as Tolman. Generally speaking, such a principle of giving more weight to less connected scientists seems to be a more realistic method to avoid a concentration on less fruitful theories, as compared to the general restriction of connectivity that Zollman's observations suggest.

In this paper, I investigate the former option, namely the effects of changing the process of information aggregation at the individual level in the framework of Zollman's bandit models. In the cited work, he only considers aggregating information in a way that gives equal weight to all neighbors. The idea pursued in this paper is that if individuals assign lower weights to the beliefs of those neighbors who are particularly well connected, this should cancel out the higher structural influence of the latter. In particular, I propose and investigate two testimonial norms which require to discount the results of neighbors with higher degrees in comparison to neighbors with lower degrees. As Zollman (2012) shows in the context of opinion pooling, the problem with non-regular graphs is that nodes with a higher degree have a higher influence on the aggregated result, which is not desirable if all nodes are assumed to produce equally reliable information. However, social networks are not structured like this in reality. As various empirical investigations into real-life networks of various sorts, including social networks, have shown, individuals' degrees are usually distributed within a wide range of values. My results suggest that the proposed structure-sensitive testimonial norms (SSTNs) are indeed more reliable than conventional ones in terms of the networks' convergence to the objectively better theory, at least in the parameter space of the original model. In general, this comes at the cost of longer convergence times.

¹The other two founders were Clark Hull and Kenneth Spencer.

²Among other indicators, the study of Krantz and Wiggins (1973) suggests that it was at least partly due to Skinner's "crusading zeal" (Godfrey-Smith, 2003, 171).

2 The Zollman model

In this section, I explain the model that Zollman develops in his 2007 and refines in his 2010 paper, drawing on earlier work by Bala and Goyal (1998). For my investigations, I employ the same general model while adding an alternative network structure and additional testimonial norms. In mathematical language, networks are called graphs and I use these two terms interchangeably. The basis of Zollman's model for investigating the effect of different network structures is given by a finite number of scientists (nodes) with connections (edges) between them; there is a path between any pair of scientists, either directly or through other scientists (making the graph a 'connected' one). Following Bala and Goyal (1998), Zollman models the choice between scientific theories (or methodologies) as a bandit problem, in this case involving two-armed bandits as scientists decide between two different theories. More precisely, they can decide between (in the language of bandit problems) using one of two slot machines which each have a specific (unknown) probability of success. He motivates this by observing that "[s]cientists often must choose between different methodologies in approaching a particular problem, and different methods have different intrinsic probabilities of succeeding" (Zollman, 2007, 347). Every time they 'pull the lever' at one of the slot machines (i.e. apply the methodology), they win or lose (the methodology succeeds or fails).

Bandit problems have been investigated in terms of optimal strategies for maximizing successes (or payoff). These usually involve a balance between exploration (getting increasingly reliable estimates for the probabilities of both machines) and exploitation (maximizing the expected number of successes according to current beliefs). While focusing on exploration would lead to pulling both levers equally often, focusing on exploitation leads to only pulling the lever which currently appears to offer the highest rate of success. Zollman has his scientists employ this latter strategy, assuming that they "would like to work only on those theories which can most successfully be applied" (Zollman, 2010, 23), given that "scientists are often rewarded for current successes (whether via tenure, promotion, grants, or awards)" (Zollman, 2007, 346). The process of the model is iterative. Each round, all scientists conduct 1000 experiments (pulls) in the methodology they currently think best and count the number of successes they achieve. Then they update their credences about the success rates of the two methodologies, based on their own new experiments *as well as the new experiments of their neighbors* (i.e. the scientists they are directly connected to). Thus, connections between scientists imply that they communicate their results with each other.³ As scientists are only concerned with exploitation rather than exploration, the risk is that once all adopt the same theory, "they are no longer learning about the other methodology, and unless the right sort of results occur, they will not ever return" (Zollman, 2010, 28).

³Communication always goes both ways, i.e. we are only concerned with *undirected graphs*.

Fig. 1 Three network structures with ten nodes with increasing connectivity: cycle, wheel and complete graph (from Zollman, 2010)

In his refined model reported in the 2010 paper, Zollman represents the scientists' credences through beta distributions. These are characterized by two parameters (α , β) and defined as follows:

$$f(x) = \frac{x^{\alpha-1}(1-x)^{\beta-1}}{B(\alpha, \beta)} \text{ with } B(\alpha, \beta) = \int_0^1 u^{\alpha-1}(1-u)^{\beta-1} du.$$

To compare two credence distributions, the respective expected success rates are computed via the distribution mean, given by $\frac{\alpha}{\alpha+\beta}$. This family of distributions is well suited for modelling Bayesian learning from such trials, as the posterior will always be a beta distribution again.⁴ In particular, if the two parameters of the prior are α and β , then the posterior has parameters $\alpha + s$ and $\beta + p - s$ where p denotes the number of pulls and s the number of successes. In the model, scientists are randomly assigned initial α_i 's and β_i 's ($i \in \{0, 1\}$) between 0 and 4 for the two theories.⁵

Zollman (2007, 2010) runs simulations on three different types of network structures which he calls cycle, wheel and complete graph (Fig. 1). The results of his simulations show that lower connectivity leads to higher reliability, i.e. a higher proportion of simulations converging to the truth.⁶ On the other hand, higher connectivity also leads to faster convergence, i.e. the networks converged after a lower number of rounds. The reason for this "robust trade off between speed and reliability" (Zollman, 2007, 339) is not hard to see: the higher the connectivity, the more information is available to the individual scientist and, thus, the faster she will, on average, receive reliable evidence for the superiority of the objectively superior theory. However, if the scientists get a few bad results with the better theory, this information also spreads further and leads to the above-mentioned scenario where all scientists adopt the inferior theory and the better theory is lost.

Although investigating different network structures is a worthy endeavor, there is another aspect that can be changed. In the remainder of this paper, I report the

⁴Formally speaking, beta distributions are a conjugate prior for the binomial distributions that model multiple pulls from the slot machines.

⁵As Zollman (2010, 30) notes, "the size of the initial α 's and β 's determines the strength of an individual's prior belief".

⁶In Zollman (2010), networks were considered to converge to the truth if all scientists preferred the objectively better theory after 10,000 rounds.

design and results of a series of simulations where the influence of scientists on their neighbors can depend on the local structure of the network.

3 Structure-sensitive testimonial norms

Zollman's models assume that in determining their preferred strategies, scientists give equal weight to all of their neighbors' experiments. This is an assumption about the testimonial norms (TNs) that the scientists adopt. In this section, I introduce two TNs according which require to give less weight to neighbors who have a high degree, i.e. many connections (information channels) to other scientists. I call them the Proportional and the Symmetric TN. Before turning to them, I formulate Zollman's TN in two versions (Additive and Averaging) to make it easier to compare the new ones to the old one. The section closes with a discussion of individual scientists' motivation to adopt such norms.

3.1 Conventional TNs

For every scientist i , let $\mathcal{N}_i$ denote her neighborhood, i.e. $j \in \mathcal{N}_i : \Leftrightarrow i = j \vee \text{"}j \text{ is connected to } i\text{"}$. In order to allow comparisons between the different TNs, I will formulate all of them in terms of weights (which can be interpreted as constituting a quadratic weight matrix). For a given TN and two scientists i and j , let w_{ij} denote the weight that i assigns to j 's experiments for updating her beta distribution. In the case of Zollman's TN, these weights are either zero or one, depending on whether j is in the neighborhood of i . I shall call this the *Additive TN*:

$$\text{Additive TN: } w_{ij}^{Add} = \begin{cases} 1 & \text{if } j \in \mathcal{N}_i \\ 0 & \text{otherwise} \end{cases}$$

In order to make the different TNs as comparable as possible, I formalize them in such a way that the total amount of newly acquired information per round is constant among scientists. Therefore, I normalize the weights such that the weights that every scientist assigns to her neighbors (still including herself) sum to one:

$$\forall i : \sum_j w_{ij} = 1 \quad (1)$$

If we normalize the weights of the Additive TN in this way, this yields what I call the *Averaging TN*:

$$\text{Averaging TN: } w_{ij}^{Avg} = \begin{cases} \frac{1}{|\mathcal{N}_i|} & \text{if } j \in \mathcal{N}_i \\ 0 & \text{otherwise} \end{cases}$$

Now we have $\forall i : \sum_j w_{ij}^{Avg} = |\mathcal{N}_i| \cdot \frac{1}{|\mathcal{N}_i|} = 1$ such that (1) is satisfied.⁷

⁷This generally increases the influence of the priors for each scientist (by dividing the change to the parameters each round by $|\mathcal{N}_i|$). As this increase depends on $|\mathcal{N}_i|$, the effect is stronger for scientists with high connectivity. This, however, makes no significant change to the results (see the differences between the Averaging TN and the Additive TN in Figs. 3 and 4); in the case of the structure-sensitive TNs, multiplying the weights by $|\mathcal{N}_i|$ appears to make them even more reliable.

3.2 SSTNs

The two TNs defined above assign equal weight to all neighbors, which allows scientists with a higher degree to have more influence. If the objectively better theory comes off unusually bad in their experiments, this misleads all neighbors and may make the network converge towards the inferior theory. Therefore, I shall test the hypothesis that networks are more reliable when the influence of such scientists is restricted. As a first approach, I consider a TN requiring that the weights assigned to one's neighbors are inversely proportional to the size of their neighborhood. I call this the *Proportional TN*:

$$\text{Proportional TN: } w_{ij}^{Pro} = \begin{cases} \gamma_i \cdot \frac{1}{|\mathcal{N}_j|} & \text{if } j \in \mathcal{N}_i \\ 0 & \text{otherwise} \end{cases}$$

using the normalization constant $\gamma_i := \frac{1}{\sum_{j \in \mathcal{N}_i} \frac{1}{|\mathcal{N}_j|}}$ to satisfy (1).

However, if we think more rigorously, then all scientists should ideally have equal influence if they are equally reliable (which is an assumption of the model). This means that the different weights that are assigned to a particular scientist j by her neighbors should sum to a fixed value, in this case to one. If the sum of the weights assigned to individual scientists is fixed, then all of them have the same influence. This leads to a second constraint:

$$\forall j : \sum_i w_{ij} = 1 \quad (2)$$

On a first note, it is easy to check that the Proportional TN generally fails to fulfill this constraint.⁸ So how should the weights for a reasonable TN be constructed such that both (1) and (2) are satisfied? Such a TN needs to be symmetric, i.e. satisfy $w_{ij} = w_{ji}$, at least where either i or j is only connected to one other node (given that $|\mathcal{N}_i| = 2$, by (1) and (2), implies $w_{ij} = 1 - w_{ii} = w_{ji}$). Generalizing this property, I will require the norm to be symmetric for all connections between scientists. If we assume that each scientist only knows her own and her neighbors' degrees, rather than the whole network structure, $w_{ij} = w_{ji}$ can only depend on $|\mathcal{N}_i|$ and $|\mathcal{N}_j|$. These constraints do not leave much space for possible TNs;⁹ one (yet not the only) possibility is then given by what I call the *Symmetric TN*:

$$\text{Symmetric TN: } w_{ij}^{Sym} = \begin{cases} \frac{1}{\max\{|\mathcal{N}_i|, |\mathcal{N}_j|\}} & \text{if } j \in \mathcal{N}_i \setminus i \\ 1 - \sum_{j \in \mathcal{N}_i \setminus i} w_{ij}^{Sym} & \text{if } i = j \\ 0 & \text{otherwise} \end{cases}$$

In the case that the Proportional and the Symmetric TN are more reliable than the conventional ones – which is indeed the case, as we will see in the next section –,

⁸This is already the case in a network of three nodes where two nodes have degree one and one node has degree two.

⁹In fact, it can be shown that the maximal weight that two scientists can give each other under these constraints is bounded by $\frac{1}{\max(|\mathcal{N}_i|, |\mathcal{N}_j|) - 1}$.

one might attribute this to the former two being more conservative than the latter two. In other words, they put more weight on one's own results. At least since the work by Weatherall and O'Connor (2020), we know that giving more weight to one's neighbors' current beliefs has a negative influence on the group's convergence rate. To control for this, I introduce another testimonial norm which requires that scientists put as much weight on their own results as with the Symmetric TN but average over their neighbors' results. I call this the *Conservative TN*:

$$\text{Conservative TN: } w_{ij}^{Con} = \begin{cases} \delta & \text{if } i = j \\ \frac{1-\delta}{|\mathcal{N}_j|-1} & \text{if } j \in \mathcal{N}_i \setminus i \\ 0 & \text{otherwise} \end{cases}$$

where δ is defined as the average weight a scientist assigns to her own results if everyone followed the Symmetric TN, i.e. $\delta := \frac{1}{N} \sum_{i=1}^N w_{ii}^{Sym}$ with N denoting the number of scientists.

3.3 Individual rationality

So far, structure-sensitive testimonial norms have been motivated by some form of 'the greater good', by a potential advantage for the whole scientific community. But even if it is beneficial for the community to adopt such a norm, it is not clear whether it is also individually rational to do so (cf. Mayo-Wilson et al. 2011). One indication against it is that, as I am considering undirected graphs, a high out-degree also entails a high in-degree. This means that neighbors with high influence also have access to a large amount of information. It might, therefore, be questioned whether it is desirable to give low weights to such neighbors in actual scientific practice. Such a potential divergence of individual and group rationality has already been discussed in the context of the Zollman effect. In fact, Zollman (2010, 30) observes that "[t]he offered solutions to this problem all turn on individuals being arranged in ways that make each individual look epistemically sub-optimal". But in the case of SSTNs, this is not as clear: giving less weight to those who influence a lot of fellow scientists may not be irrational for the individual. It has already been noted by Kitcher (1990) and Strevens (2003) that the incentive structure in science favors those who go against the big trends and thereby increase their chance of being one of the first on a new avenue of research. If scientists face the choice between a conformist strategy that will have them in good company if they fail and a more risky strategy that increases their chance of being the first to pursue a better methodology, there is not a general answer as to which choice is the rational one. So if this indeed describes the choice between conventional and structure-sensitive testimonial norms, then the latter are not necessarily 'epistemically sub-optimal'. To investigate these two predictions about SSTNs, I also report simulations of mixed networks, comprising scientists that follow different TNs, in the next section.

4 Model and results

Through agent-based models, I compare the Symmetric and the Proportional with the Additive and the Averaging, as well as the Conservative TN. I include the (original) Additive TN to make my results comparable to the results in Zollman (2010); the other four norms all have normalized weights to satisfy Eq. 1 which makes them more easily comparable. In regular graphs, i.e. graphs where all nodes have the same degree, such as the cycle and the complete graph in Zollman (2007, 2010), the new testimonial norms are equivalent to the ‘conventional’ ones (as all neighbors are assigned the same weights). The only other network structure investigated by Zollman (2010) is the wheel. As this is not a realistic network, I also investigate the behavior of the different TNs on graphs that are formed by preferential attachment (Barabási & Albert 1999; see Fig. 2). PA networks are constructed by iteratively adding nodes with a single connection, where the probability of being connected to a specific existing node is proportional to its degree.

Studies of actual social networks show that PA networks are fairly realistic in terms of the distribution of degrees among nodes (Barabási & Albert, 1999). To make the results as robust as possible, each simulation on a PA network uses an individually generated graph, sampled from the probabilistic generative algorithm.

As in Zollman (2010), I have modelled the existence of two rival theories in terms of a two-arm bandit problem where theory 1 has probability of success 0.5 and theory 2 has probability 0.499. A round consists of 1000 experiments (or pulls), where each scientist tests the theory she is currently favoring. The simulation stops after 10,000 rounds or when all scientists have used the same methodology for 200 rounds. I investigate both the percentage of simulations where scientists all adopt the better theory (i.e. theory 1) and at the average number of rounds it takes them to converge. I ran 10,000 simulations for each triple of network structure, size and TN in NetLogo (Tisue & Wilensky, 2004). After presenting the main results obtained in the original setting as well as some robustness results, I report how scientists following different TNs behave when put in a network together.

Fig. 2 Example of a preferential attachment (PA) network with ten nodes

Fig. 3 Probability of converging to the better theory for the different TNs in wheel networks as well as PA networks of different sizes. In all cases, the structure-sensitive TNs are more reliable

4.1 Main results

In accordance with Zollman (2010), reliability increases with the size of the network (Fig. 3). Furthermore, the reliability is higher in PA networks than in wheel networks of the same size for all testimonial norms. As predicted, the structure-sensitive TNs are more reliable in both wheel networks and PA networks. The Averaging and the Additive as well as the Symmetric and the Proportional TN yield quite similar results, respectively. Between the latter, the Symmetric TN is slightly more reliable, at least in PA networks. In line with the conformism results by Weatherall and O'Connor (2020), this difference might relate to the fact that following the Symmetric TN entails giving more weight to one's own results – thereby, in a sense, reducing connectivity –, an effect that is especially strong for nodes with degree one. However, this cannot account for the difference between the structure-sensitive and the conventional norms, as can be seen when considering the Conservative TN. Although this one is more reliable than the other conventional norms, it is still less reliable than the structure-sensitive ones.

As already observed by Zollman (2007, 2010), higher reliability generally comes at the expense of speed; one example is that convergence time increases with network size (Fig. 4). Also, in both wheel and PA networks, the structure-sensitive TNs take longer to converge. On this matter, there is a considerable difference between the two

Fig. 4 Average time needed to converge to the better theory for the different TNs in wheel networks as well as PA networks of different sizes. The structure-sensitive TNs always take longer, with the exception of the pair Proportional vs Conservative TN

structure-sensitive TNs in PA networks, with the Proportional TN allowing a much faster convergence. Surprisingly, the trade-off between reliability and speed cannot be found when comparing the Conservative with the Proportional TN, as the latter outperforms the former on both criteria. In Appendix C, I extend these results to the mixed networks of Section 4.3 and show that the reliability and convergence time of networks increase linearly with the number of scientists following the respective SSTN.

4.2 Robustness and complementary models

In the literature reacting to Zollman's initial results, one line of criticism targets his restricted choice of parameters. In particular, Rosenstock et al. (2017) find that the superiority of the cycle over the fully connected graph (i.e. the Zollman effect) vanishes if the success rates of the two theories differ by a larger margin, the networks are larger, or the scientists conduct a higher number of experiments per round. This lets them conclude that "the effect does not arise for significant portions of the relevant parameter space" (p. 243). As the convergence rates approach 100% for larger differences between success rates as well as for higher numbers of pulls per round, one might expect that this should also affect the results of the present work: If even the Additive TN – the one considered in their experiments – comes close to 100% reliability for complete graphs, then so should the other TNs (which have higher convergence rates, see Fig. 3) on the wheel (which has higher convergence rates than the complete graph, see Zollman's work). The same should apply to PA networks, as they appear to come with even higher convergence rates than the wheel (Fig. 3). There is, however, an important caveat: Rosenstock and colleagues investigated the model of Zollman (2007), while my results apply to the slightly more sophisticated version from Zollman (2010), the one described above.¹⁰ And indeed, the results are more resilient to changes in parameters, at least in the case of more experiments per trial. As the detailed analysis presented in Appendix A shows, the advantage of the structure-sensitive Symmetric TN over the insensitive Conservative TN is – contrary to the Zollman effect – fairly robust to changes in the number of experiments per round, in both wheel and PA networks. It is, however, not robust to changes in the other two parameters that were investigated by Rosenstock et al. (2017): Both for larger differences between the success rates of the two theories and for larger network sizes, the probability of converging to the better theory approximates 100% for all TNs and in both network types. Therefore, the difference in reliability between the TNs becomes negligible.

The model of Zollman (2010) is limited in several ways and various other models have been put forward in the literature to provide additional insights. A number of models that are particularly close to Zollman's have been proposed and discussed in Frey and Šešelja (2018a, b). Here, the central difference to Zollman's models is that the success rates of the two theories change over time: Theory 1 and 2 converge towards success rates of 1 and 0, respectively. Consequently, all scientists eventually

¹⁰To study the differences in detail, the reader is encouraged to compare the descriptions of the implementations in the two papers.

pursue the better theory in every run. In this more optimistic framework, we can only investigate the time until the networks converge (since they always end up with the better theory). In such models, the investigated TNs all perform very similarly, as reported in detail in Appendix B. Even the different variations of this model do not provide substantial insights into differences between the norms. Other complementary models have been proposed in the literature (for some examples, see Appendix B) and an investigation into the performance of SSTNs in such models might provide further insights. One particularly interesting avenue would be to investigate extensions to the more realistic class of directed graphs. Contrary to the undirected models discussed here, a high out-degree would not entail a high in-degree. In other words, higher influence would not need to come with more information about the work of fellow scientists (see Section 3.3). This could mean that structure-sensitive TNs only considering neighbors' out-degrees would perform even better in such a framework.

4.3 Individual rationality in mixed networks

So far, we have only considered networks in which everyone follows the same testimonial norm. While this is useful for comparing the different TNs, networks in which scientists do not all follow the same norm provide a more realistic setting. To investigate SSTNs in such mixed settings, I simulate networks of ten scientists where one group follow an SSTN and the rest follow a conventional norm. I consider two such pairs: in one case, the scientists follow one of Proportional/Averaging TN (which also differ with respect to the weight they give to one's own results) and in the other case, one of Symmetric/Conservative TN (which are directly comparable in this regard). For both pairs, I ran 20,000 simulations for every possible proportion between the considered testimonial norms. Only runs where the network eventually converged to the better theory were considered for the following results.

To analyze how the choice of testimonial norm affects the time it takes individual scientists to settle on the better theory, arguably the most natural measures are the mean and median of the number of rounds it takes them. Therefore, these measures are reported in Fig. 5a and b, for the Symmetric/Conservative and the Proportional/Averaging networks, respectively. While the medians do not show differences in any setting, the means of the SSTNs are higher than the means of their structure-insensitive counterparts, at least in PA networks. This holds for any proportion of SSTN to conventional TN (x-axis). It suggests that while the choice of a structure-sensitive TN is not detrimental for most of the scientists, it can lead to a slower adoption of the better theory when many other scientists are already following it. Note that there is one interesting difference between Fig. 5a and b: While the means in the Symmetric/Conservative networks decrease in PA networks with higher numbers of scientists following the SSTN, the means in the Proportional/Averaging networks increase. The latter is not surprising, as we know that networks where all scientists follow the Averaging TN converge much faster than networks where all scientists follow the Proportional TN (Fig. 4). What is more surprising is that, even though the Conservative TN is generally 'faster' than the Symmetric TN (Figs. 4, 5a), individual scientists settle faster on the better theory when there are more other

Fig. 5 Mean and median time needed to settle on the better theory for individual scientists following either a structure-sensitive or a conventional TN in mixed networks of size ten. In PA networks, the mean time for SSTNs is higher than for their conventional counterparts while there is noticeable difference w.r.t. the median

scientists following the Symmetric TN. In particular, the average time to settle individually is higher in PA networks where all scientists follow the Conservative TN, as compared to PA networks where everyone follows the Symmetric TN (cf. $x=0$ and $x=10$ in Fig. 5a). Further investigations show that only the last scientist takes longer in Symmetric networks as compared to Conservative networks (on average), which leads to longer convergence times for the entire network. When considering networks with five scientists following the Symmetric and five following the Conservative TN, the former group shows a higher variance: While the first, second and third scientist in the former group are slightly faster than their respective counterparts in the second group, the fifth and last scientist of the first group is much slower than his counterpart; in fact, this difference is large enough to outweigh the differences between the first three scientists and to result in a higher mean for the first group (cf. $x=5$ in Fig. 5a).

The results about average settling times can be seen as a reason against taking up structure-sensitive testimonial norms. But as noted in Section 3.3, one particularly important incentive for scientists is to be among the first in pursuing a theory or methodology that later turns out to be superior. Therefore, it is also interesting to see how the choice of testimonial norm affects the probability for this event. For different TNs, Fig. 6 shows the proportion of its followers that were the (sometimes shared) first to have settled on the better theory (where the mixed networks are the same as those considered above). These results show that both in networks with

Fig. 6 Proportion of scientists of each TN that are one of the first to settle on the better theory, where each network comprises ten scientists following either one of Symmetric/Conservative TN or one of Averaging/Proportional TN. In both cases, following the respective SSTN comes with a higher chance to be among the first.

the Symmetric/Conservative TNs and in networks with Proportional/Averaging TNs, scientists following the structure-sensitive norm (i.e. Symmetric and Proportional, respectively) have a distinctively higher chance of being among the first, compared to scientists following the respective structure-insensitive norm.¹¹ This holds for any proportion of SSTN to conventional TN and for both wheel and PA networks.

5 Discussion

Most notably in his 2007 and 2010 papers, Zollman has established that under certain modelling assumptions, less connected networks are more reliable than more connected ones. Furthermore, he infers that it is not the “unequal connectivity” of networks that is problematic and that in moving from cycles to wheels, “[t]he harm done by the individual at the center cannot be simply overcome by removing their centrality” (Zollman, 2007, 343). However, one can mitigate this loss of reliability to some extent. In particular, I have shown that under the modelling assumptions of Zollman (2010), certain structure-sensitive testimonial norms are more reliable (in irregular networks) than their insensitive counterparts. This generally comes at the price of longer convergence times. Surprisingly, however, the conjectured “robust trade off between speed and reliability” (Zollman, 2007, 339) does not apply everywhere. In particular, when structure-insensitive norms put much enough on one’s own results (as for the Conservative TN), they can be both less reliable *and* slower than appropriate structure-sensitive norms. This observation adds to the appeal of the latter.

The empirical question to what extent reliability and connectedness are correlated for individual scientists certainly deserves further investigation. However, among many others, the story of Tolman and Skinner mentioned in the introduction suggests that they do not necessarily go hand in hand. This observation creates a tension

¹¹Note that all values are around 20%, suggesting that on average, 2 of the 10 scientists first settle on the better theory at the same time (e.g. at the start).

with a particular result of the ‘passive learning’ model of Zollman (2012), where scientists are not thought to perform experiments but to repeatedly pool their own credence with their neighbors’. Again, he only considers pooling norms which require scientists to give equal weights to all neighbors. When solving the Markov model for the stationary distribution, Zollman (2012, 41) observes that “more highly connected individuals will receive higher final weights as compared to those who are less connected”. And again, his interpretation is that networks should be regular, since equally reliable scientists should have equal weight on the consensus view. The Symmetric TN satisfies this desideratum also for networks that are not regular: It can easily be proven that with this as the pooling norm, every scientist has equal influence in the stationary solution, and thus on the consensus view, for any given network.¹²

It might be questioned whether structure-sensitive TNs are realistic, whether real world scientists *could* and *would* adopt such norms. On the first point, empirical investigations have shown that individuals are quite good at tracking the structure of social networks around them (Kumbasar et al., 1994; Banerjee et al., 2014). Thus, they should have, at least approximately, the necessary information to adopt structure-sensitive TNs. Further worries about the practicality of SSTNs might be based on the observation that conformism seems to be a natural tendency of us humans – which can also drive scientific practice, as argued by O’Connor and Weatherall (2019), e.g. in the initial rejection and later widespread adoption of small-pox variolation. The descriptive observation of widespread conformism is, however, not relevant for the analysis of whether it is rational for the individual scientist to follow an SSTN. I have tried to address this question of individual rationality by noting that a higher chance to be among the first on a new avenue of research is generally considered a strong incentive for researchers. The simulations of mixed networks show that while following a SSTN indeed increases the probability to be the first scientist working with the theory that turns out to be better, this comes at the risk of a much later adoption when many other scientists are already following it.¹³ So while there is not a clear case that following SSTNs is individually rational for scientists, it is also everything but clear that it is irrational.

As reported above, the main results are not resilient to all changes of the parameter space. While they are robust under changes in the number of experiments per round, the model ceases to be informative when the difference between the success rates of the two theories or the number of scientists in the network increases. This can, however, not be attributed to the testimonial norms. Rather, it is an artifact of the model as the task becomes too easy: the probability of convergence to the better theory approximates 100% for all considered norms, so no comparisons can be made. The SSTNs also cease to be superior to conventional norms when investigated through a family of complementary, more optimistic models that was proposed in Frey and Šešelja (2018a). Hence, the results cannot be considered to be conclusive. In particular, it might be fruitful to investigate the performance of the novel type of

¹²This result follows from the constraint described by Eq. 2.

¹³Note that – using the terms coined by Laudan (1984) – it would be enough for the scientists to adopt such a norm for determining which theory to *pursue* rather than *accept*.

norms in some of the other models that have been put forward in the area of network epistemology. After all, the results reported here, along with the observation that changing TNs is a more realistic strategy than restricting information channels, make structure-sensitive testimonial norms a worthwhile avenue for further research.

Appendix A: Robustness

Here, I report investigations into the robustness of the main results reported in Section 4.1, considering extensions of the parameter space. They are restricted to the Symmetric and the Conservative TN as this is the most interesting and, indeed, decisive comparison. Figure 7 shows results for settings where ten scientists conduct between 1,000 and 10,000 experiments per round. It can be seen that – contrary to the Zollman effect as investigated by Rosenstock et al. (2017) – the advantage of the structure-sensitive Symmetric TN over the insensitive Conservative TN is fairly robust to changes in the number of experiments per round, in both types of networks.

The observed advantage of SSTNs is, however, not robust to changes in the other two parameters that were investigated by Rosenstock et al. (2017). Figure 8 shows results for settings where theory 2 has a probability of success between 0.499 and 0.46, while the probability for theory 1 remains at 0.5. The number of scientists was again fixed at 10 and the number of experiments per round at 1000. Here, the effect quickly vanishes for larger differences between the success rates of the two theories. This is because the scientists' task becomes, in a way, too easy: in all settings, the probability of converging to the better theory approximates 100%. This does not allow to infer anything about possible advantages or disadvantages of either testimonial norm.

Lastly, Rosenstock and colleagues also tested what happens when the number of scientists in the network is higher than in the parameter space considered by Zollman (2007, 2010). Similar to the Zollman effect, the advantage of the Symmetric over the Conservative TN vanishes in such settings. Figure 9 shows results for larger network sizes; again, the task becomes too easy and the probability always approximates 100%. In sum, the results are robust with regard to one of the three considered parameter space extensions.

Fig. 7 Probability of converging to the better theory for the Symmetric and Conservative TN in wheel networks as well as PA networks, for different numbers of experiments per round. The advantage of the Symmetric TN is fairly robust to changes across this dimension

Fig. 8 Probability of converging to the better theory for the Symmetric and Conservative TN in wheel networks as well as PA networks, for different success rates of theory 2. With decreasing success rate (and, thus, an increasing gap to the rate of theory 1, which is 0.5), the probability to converge to the better theory approaches 100% for all settings and both TNs

Fig. 9 Probability of converging to the better theory for the Symmetric and Conservative TN in wheel networks as well as PA networks, for larger numbers of scientists. When the number exceeds around 20, the probability approaches 100% for all settings and both TNs

Appendix B: Complementary models

Here, I report a detailed analysis of SSTNs in the context of a family of models proposed and discussed in Frey and Šešelja (2018a, b). The central difference to Zollman's models consists in the success rates of the two theories changing over time. This is motivated by the assumption that in many contexts, the objectively better theory (theory 1) becomes more and more superior over time. Accordingly, Frey and Šešelja propose to increase the success rate of theory 1 and decrease the success rate of theory 2 by $\frac{|r_{cur} - r_{lim}|}{1000}$ every 100 rounds, where r_{cur} is the current rate and r_{lim} is the end point to which it converges. The latter is fixed at 1 for theory 1 and at 0 for theory 2. This means that the model is more optimistic, as *in every single run, all scientists end up with the better theory* given that after enough time, the two theories' success rates approach 1 and 0, respectively. Therefore, they only consider the convergence time as an evaluation metric. Results from comparing the SSTNs

Fig. 10 Average time needed to converge to the better theory for the different TNs in wheel networks as well as PA networks of different sizes, for the basic model proposed in Frey and Šešelja (2018a). All TNs perform very similarly

with the Conservative TN are depicted in Fig. 10.¹⁴ In this model, the respective performances of the three TNs are almost indistinguishable. When pressed to find something noteworthy, one might say that the Conservative TN is the slowest, at least in wheel networks. However, the main takeaway should be that there are no interesting differences.

Frey and Šešelja also consider further variations of their model. Among others, they introduce notions of ‘critical interaction’ and ‘rational inertia’.¹⁵ Both of them can be seen as features that make the model more descriptively adequate without representing irrational behavior. *Critical interaction* is based on the plausible assumption – defended e.g. in Betz (2012) – that “critique tends to be truth conducive since it allows for false beliefs to be exposed as such” (Frey & Šešelja, 2018a, 12). This is implemented as follows. Critical interaction is triggered whenever a scientist receives a report by a neighbor pursuing the rivalling theory according to which that theory is better than the first scientist thought. On such occasions, the scientist reflects on her methods; this is thought to affect the success rate of the theory on which she is currently working in a ‘truth-conducive’ way. If she is working on the better theory, her success rate increases. If she is working on the worse theory, it decreases. In both cases, this individual update follows the same formula as the regular global update, i.e. the success rate changes by $\frac{|r_{cur} - r_{lim}|}{1000}$. This means that the average convergence times are shorter in this setting, given that the scientists are faster to converge to theory 1. Results for this setting are shown in Fig. 11. Again, there are no interesting differences between the three norms.

The second possible amendment to the model is *rational inertia*. Here, the idea is that scientists don’t necessarily abandon their current theory as soon as the other theory appears to be better. As argued by Kelp and Douven (2012), this is not irrational

¹⁴As the results in this model are evaluated according to a metric that, in a way, mixes the two metrics considered in the main model, both the Symmetric and the Proportional TN are analyzed in this section. As can be seen, there are no qualitative differences between the two in all but the last setting.

¹⁵They also consider something called a ‘theory threshold’ where scientists only switch theories if the other theory seems more successful by a margin of 0.1. This setting seems to be helpful mainly for descriptive, rather than normative, contexts so I am not considering it here.

Fig. 11 Average time needed to converge to the better theory for the different TNs in wheel networks as well as PA networks of different sizes, for the model proposed in Frey and Šešelja (2018a) with critical interaction. All TNs perform very similarly

as it might be possible to improve the theory. Frey and Šešelja (2018a) implement this idea by demanding that a scientist only changes the theory he is working on when the rivaling theory has appeared better to him for ten rounds. The results are shown in Fig. 12. Again, there are no stark differences between the norms. However, the PA setting here provides the strongest trend among the six considered configurations of the general model: The Symmetric TN is the slowest norm in this configuration.

In sum, the different configurations of this model did not provide substantial insights into differences between the norms. Other related models have been suggested in the literature; while the investigation of SSTNs in these models is beyond the scope of the present work, it appears to be a worthwhile endeavor for further research. Mayo-Wilson (2014) considers a different setting which is more complex in the sense that it involves the notion of different disciplines and domain experts, yet simpler in the sense that all researchers pursue only one methodology, so there is no bandit decision problem. It is also in a way more optimistic than the Zollman model as all scientists eventually find the truth in their disciplines. In further research, it could be interesting to construct structure-sensitive testimonial norms for this framework and test them against those previously introduced.¹⁶ Zollman (2015) considers a setting where scientists can differ in their reliability but comes to conclusion that scientists should not bother assessing this reliability; this indicates that SSTNs might give good results in this model already in their current form but it would also need to be investigated in future work. Another limitation of the models used in the present work is that the networks are static. A nice alternative approach is proposed by Alexander (2013) where scientists form additional links with more successful colleagues. Lastly, the performance of SSTNs can be tested in models where scientists need not be honest, as suggested by Holman and Bruner (2015). This may make it even more important to not give too much weight to individuals with high influence which would favor SSTNs even more.

¹⁶His proximitist TN is the only TN in the literature I am aware of that takes some properties of the network structure into account, although in a different, more restricted, way.

Fig. 12 Average time needed to converge to the better theory for the different TNs in wheel networks as well as PA networks of different sizes, for the model proposed in Frey and Šešelja (2018a) with rational inertia. All TNs perform very similarly, although the Symmetric TN takes a little longer in PA networks

Appendix C: Global properties of mixed networks

Here, I report global properties of mixed networks; that is, I extend the results of Section 4.1 to the networks investigated in Section 4.3. We already know that in homogenous networks, the Symmetric TN and the Proportional TN are more reliable and have longer convergence times than the Conservative TN and the Averaging TN, respectively (see Figs. 3 and 4). The question arises how reliability and convergence time change when there are some scientists following one norm and some scientists following the other. We basically interpolate between the case of ten scientists following the Symmetric (Proportional) TN and ten scientists following the Conservative (Averaging) TN. A priori, there are two possibilities how the intermediate networks could behave. One option is that there is some threshold between zero and ten such that there is a sudden jump in reliability (and/or convergence time). This would mean that it is enough for a community that some (rather than all) researchers discount the results of well-connected researchers. The other option is that the interpolation curve is (approximately) linear, i.e. every researcher that follows a SSTN increases the reliability of the network by roughly the same magnitude, irrespective of what the other

Fig. 13 Average time needed to converge to the better theory for mixed networks of size ten. Some networks only have scientists following the Symmetric or Conservative TN (blue), other networks only have scientists following the Proportional or Symmetric TN (green). Increasing the number of scientists following the respective structure-sensitive TNs (from 0 to 10) leads to a roughly linear increase of reliability for both network structures

Fig. 14 Average time needed to converge to the better theory for mixed networks of size ten. Some networks only have scientists following the Symmetric or Conservative TN (blue), other networks only have scientists following the Proportional or Symmetric TN (green). Increasing the number of scientists following the respective structure-sensitive TNs (from 0 to 10) leads to a roughly linear increase of convergence time for both network structures

researchers are doing. This is exactly what the results of the simulations depicted in Fig. 13 show. It is then not surprising that the convergence time also increases linearly with the number of scientists following the respective SSTN (Fig. 14).¹⁷

Acknowledgements I would like to thank the anonymous reviewers as well as the editorial team for their help to improve this paper. Furthermore, I would like to thank Rush Stewart whose seminar on social epistemology at the MCMP inspired this work.

Funding Open Access funding enabled and organized by Projekt DEAL.

Data Availability The simulation data can be shared.

Code Availability The NetLogo code can be shared.

Declarations

Conflict of interests The authors declare that they have no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

¹⁷Note that in Figs. 13 and 14, the ‘border cases’ of 0 and 10 scientists following the respective SSTN correspond to the homogenous networks investigated in Section 4.1.

References

- Alexander, J. M. (2013). Preferential attachment and the search for successful theories. *Philosophy of Science*, 80(5), 769–782.
- Bala, V., & Goyal, S. (1998). Learning from neighbours. *The Review of Economic Studies*, 65(3), 595–621.
- Banerjee, A., Chandrasekhar, A. G., Duflo, E., & Jackson, M.O. (2014). Gossip: Identifying central individuals in a social network. Technical report. National Bureau of Economic Research.
- Barabási, A.-L., & Albert, R. (1999). Emergence of scaling in random networks. *Science*, 286(5439), 509–512.
- Bedessem, B. (2019). The division of cognitive labor: Two missing dimensions of the debate. *European Journal for Philosophy of Science*, 9(1), 3.
- Betz, G. (2012). *Debate dynamics: How controversy improves our beliefs* (Vol. 357). Springer Science & Business Media.
- Campbell, D. T. (1979). A tribal model of the social system vehicle carrying scientific knowledge. *Knowledge*, 1(2), 181–201.
- De Langhe, R. (2014). A unified model of the division of cognitive labor. *Philosophy of Science*, 81(3), 444–459.
- Frey, D., & Šešelja, D. (2018a). Robustness and idealizations in agent-based models of scientific interaction. *The British Journal for the Philosophy of Science*.
- Frey, D., & Šešelja, D. (2018b). What is the epistemic function of highly idealized agent-based models of scientific inquiry? *Philosophy of the Social Sciences*, 48(4), 407–433.
- Godfrey-Smith, P. (2003). *Theory and reality: An introduction to the philosophy of science*. University of Chicago Press.
- Holman, B., & Bruner, J. P. (2015). The problem of intransigently biased agents. *Philosophy of Science*, 82(5), 956–968.
- Hull, D. L. (1990). *Science as a process: An evolutionary account of the social and conceptual development of science*. University of Chicago Press.
- Kelp, C., & Douven, I. (2012). Sustaining a rational disagreement. In *EPSA philosophy of science: Amsterdam 2009* (pp. 101–110). Springer.
- Kitcher, P. (1990). The division of cognitive labor. *The Journal of Philosophy*, 87(1), 5–22.
- Krantz, D. L., & Wiggins, L. (1973). Personal and impersonal channels of recruitment in the growth of theory. *Human Development*, 16(3), 133–156.
- Kuhn, T. S. (1962). *The structure of scientific revolutions*. University of Chicago Press.
- Kumbasar, E., Rommey, A. K., & Batchelder, W.H. (1994). Systematic biases in social perception. *American Journal of Sociology*, 100(2), 477–505.
- Kummerfeld, E., & Zollman, K. J. S. (2016). Conservatism and the scientific state of nature. *The British Journal for the Philosophy of Science*, 67(4), 1057–1076.
- Laudan, L. (1984). *Science and values: The aims of science and their role in scientific debate*. University of California Press.
- Mayo-Wilson, C. (2014). Reliability of testimonial norms in scientific communities. *Synthese*, 191(1), 55–78.
- Mayo-Wilson, C., Zollman, K. J. S., & Danks, D. (2011). The independence thesis: When individual and social epistemology diverge. *Philosophy of Science*, 78(4), 653–677.
- O'Connor, C., & Weatherall, J. O. (2019). *The misinformation age: How false beliefs spread*. Yale University Press.
- Rosenstock, S., Bruner, J., & O'Connor, C. (2017). In epistemic networks, is less really more? *Philosophy of Science*, 84(2), 234–252.
- Strevens, M. (2003). The role of the priority rule in science. *The Journal of Philosophy*, 100(2), 55–79.
- Tisue, S., & Wilensky, U. (2004). Netlogo: A simple environment for modeling complexity. In *International conference on complex systems* (Vol. 21 pp. 16–21). Boston, MA.
- Weatherall, J. O., & O'Connor, C. (2020). Conformity in scientific networks. *Synthese*, 1–22.
- Zollman, K. J. S. (2007). The communication structure of epistemic communities. *Philosophy of Science*, 74(5), 574–587.
- Zollman, K. J. S. (2010). The epistemic benefit of transient diversity. *Erkenntnis*, 72(1), 17.
- Zollman, K. J. S. (2012). Social network structure and the achievement of consensus. *Politics, Philosophy & Economics*, 11(1), 26–44.

- Zollman, K. J. S. (2013). Network epistemology: Communication in epistemic communities. *Philosophy Compass*, 8(1), 15–27.
- Zollman, K. J. S. (2015). Modeling the social consequences of testimonial norms. *Philosophical Studies*, 172(9), 2371–2383.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.