Ignorance in Social Networks

Discounting Delays and Shape Matters

Abstract

Knowledge is the proper basis for action (Williamson, 2000). But disinformation causes ignorance, whether through error (false belief) or omission (agnosticism). We use philosophical simulations (Mayo-Wilson and Zollman, 2021) to study how ignorance persists in networks of inquiring rational agents. Following Zollman (2007), we simulate communities of agents who generate evidence, share it with their neighbours, and then update their beliefs. After recapping previous results, we report two novel findings. First, in variations on the mistrust models developed by O'Connor and Weatherall (2018), discounting evidence received from neighbours delays convergence to the truth. Second, ignorance persists differentially in networks of different shapes, even when they have the same overall connectivity. These results shed light on the structural causes of ignorance that can be exploited by those engaged in disinformation campaigns; and they constrain the space of knowledge-conducive responses.

1. Introduction

Actions undertaken by social groups are best when underpinned by relevant knowledge. Appropriate responses to climate change, or the coronavirus pandemic, for example, require informed (coordinated) democratic action on the part of whole societies. And these actions may be legitimated by the inquiries of e.g. certain scientific communities. Yet mis- or even disinformation may disrupt knowledge from taking hold within groups of individuals, preventing timely and appropriate action. Indeed, knowledge is hard won, and ignorance may persist or prevail even in their absence.

The current paper investigates ignorance in social networks. We begin by describing our approach, detailing key aspects of our models and methods. We then summarize previous results (surrounding the Zollman effect and polarization) before presenting two novel findings of our own: first, that discounting good evidence delays the emergence of a correct, knowledgeable consensus within communities of rational inquirers; and second, that when it comes to the persistence of ignorance, social network shape, or topology, matters. We close with some concluding remarks that contextualize our findings, indicating their relevance to issues of trust and disinformation.

2. Models and Methods

We model the social situations that interest us as involving networks of individual rational agents, conducting experiments, sharing their findings, and updating their beliefs in light of the evidence at their disposal (cf. Zollman, 2007). Allow us to unpack this.

2.1 Social Groups as Networks of Individuals

We model social groups as networks of individuals. These networks can be represented as *graphs*: (sets of) *nodes* connected by *edges*. Although the terminology here derives from network theory, the approach should be familiar to those acquainted with (the semantics of) predicate logic: a given graph represents the extension of a relation amongst (or, on a domain consisting of) some individuals. Indeed, multi-edge graphs are possible - these represent the extensions of a number of different relations simultaneously (on the domain in question). Moreover, labels can be attached to the nodes: this allows for the representation of properties of the individuals. Here we focus on the simple case of unlabelled, single-edge graphs.

Note that, while it is common to present the semantics of predicate logic - and network theory - in the language of set theory, so that domains are sets of individuals, this is not essential: a plural framework can be employed instead. We will not make much of this here: but it is worth remarking that we are not assuming that social groups are *sets* of individuals; for sets are themselves individual entities over and above, and distinct from, their members. (For example, even the singleton set {i} comprising just the individual i is distinct from i itself.) We are rather concerned with e.g. scientific communities, or democratically self-governing peoples, construed as the individuals themselves - taken collectively.[1]

Indeed, there has been much interesting work recently on the metaphysics of social groups. For instance, Uzquiano (2004) argued that the Supreme Court of the United States must be construed as a *sui generis* entity distinct from its justices (yet not a set or mereological sum of them either). While we are sympathetic to there being such entities (e.g. corporations, and other institutional organizations), here we more closely follow the work of Ritchie (2020), for whom 'social structures are central to the nature of all social groups' (2020: 402), and where structures in turn are 'complexes… of relations' (2020: 405). Thus, we crucially take the social groups with which we are concerned, to be not mere pluralities (or sets!) of individuals, but *structured* pluralities, or *networks* of individuals - which can be appropriately represented by graphs.

2.2 Individuals as Rational Agents

At the nodes of our networks are rational agents. As *agents*, they take actions. These actions in turn modify their environment; thus, the effects of their actions can be observed, yielding new evidence. In our models, the actions are limited to two: agents perform either action A or action B. In particular, we can think of the agents as (scientists) conducting experiments, such as trialling a drug on a certain number *n* of experimental patients.

Insofar as they are *rational*, our agents' beliefs are based on the evidence at their disposal, and the actions they take are appropriate given their attitudes. Thus, they perform action A (administer a familiar drug) if they take it to be more likely to yield a desirable result (e.g. help more of their experimental patients), and B if they think (as is in fact the case) B is better (in this same sense).[2] They observe the results of their experiments (e.g. 6 out of 10 patients recovered), and communicate their findings (by testimony) to the people with whom they are connected (by edges) in the network - their *neighbours*. Finally, they update their beliefs in light of the evidence at their disposal (that which they themselves generate, as well as that which is communicated to them by their neighbours).

What exactly is rational belief updating? We consider two variants. In both, our agents are modelled as having a *credence*, or degree of belief (between 0 and 1), in the (true) proposition that action (or drug) *B is better* (than A).[3] A is assumed to have a constant (and known), 0.5 probability (i.e. 50% chance) of helping a given patient to recover from their illness; B has a 0.5 + (positive) epsilon probability of doing so. The proposition that B is better just is the proposition that epsilon is positive

---

[1] It is common to distinguish 'collective' from 'distributive' plural predications: the latter are equivalent to conjunctions of individual predications, while the former are not; thus, the claim that some people carried a piano is likely meant collectively – unless the people in question are e.g. participants in a strongman competition.

[2] More sophisticated models of the practical rationality involved in deciding what action to take are of course possible. Here it is the agents' beliefs alone that determine what they do; but we might expect their desires, preferences, or utilities to play a role as well - as in standard decision theory. For instance, O'Connor and Weatherall (2019) discuss 'conformity' models in which scientific agents are motivated not only by the desire to ascertain the truth about the matter they are investigating, but also by a desire to conform with what their peers are doing. We do not here explore such variations on the basic models under discussion.

[3] Note that this is already a robust theoretical assumption: some theorists regard (outright) belief as a basic, all-or-nothing affair; that belief comes in degrees, as assumed in the models we employ, is therefore contentious.

– which in our models it always is (so that our target proposition is true); but this is not initially known to the agents in the models.

In the models originally employed by Zollman (2007), agents are initially assigned a credence drawn from a random uniform distribution (in line with the *subjective* Bayesian assumption that any credence is rational prior to the receipt of evidence). They then update their beliefs in light of the evidence available at any given stage using

Bayes' rule: $P_f(h) = P_i(h/e) = P_i(e/h) \times P_i(h)/P_i(e)$

(where P_f is the final probability, after updating, and P_i the initial probability, prior to updating; h is some hypothesis, and e is a given piece of evidence).

Indeed, O'Connor and Weatherall (2019) argue (based on diachronic Dutch book considerations) that this is the unique rational method for updating beliefs over time.[4] Nevertheless, they themselves subsequently suggest that it might sometimes be rational to be more cautious in dealing with the available evidence. Thus, in their variant 'mistrust' models, agents are uncertain of the evidence with which they are presented, and update using

Jeffrey's rule: $P_f(h) = P_f(e) \times P_i(h/e) + P_f(\text{not } e) \times P_i(h/\text{not } e)$.

It is worth noting that Bayes' rule is the special case of Jeffrey's rule in which $P_f(e) = 1$ and $P_f(\text{not } e) = 0$. Thus, the use of Jeffrey's rule might be thought to simply generalize the use of Bayes' rule. However, the matter is not quite so simple: as Williamson (2000) has noted, Jeffrey's rule cannot be immediately operationalized; in order to put it to use, we need to know what $P_f(e)$ is - and nothing in the rule itself determines this.

In order to address this issue, and put Jeffrey's rule into practice in their models, O'Connor and Weatherall make an assumption of *homophily*: people trust others more when they are more alike.[5] Moreover, they assume a particular version of this idea - namely, that people trust each other more when their *beliefs* (on the question whether B) are more similar.[6] Finally, they implement this (attitudinal) version of the assumption in a *specific mathematical formula*:

$P_f(E) = 1 - \min(\{1, d \cdot m\}) \cdot (1 - P_i(E))$.[7]

Here d is the distance (i.e. the absolute value of the difference) between the credences of the recipient and the provider of the evidence supplied through testimony. Thus, agents do not discount the

---

[4] Such arguments are contentious - see Williamson (2000).

[5] Note that it is possible to implement Jeffrey's rule without the assumption of homophily. For instance, in our 'testimonials' models (AUTHORS, IN PREPARATION), we consider the case in which agents set the final probability of the evidence they receive to the frequency of truth-telling in their network and/or neighbourhood.

[6] It is also possible to implement the homophily assumption without assuming that the relevant respect of similarity is in the credence given to the proposition that B is better. For instance, in our 'epistemic injustice' models (AUTHORS, IN PREPARATION), we label nodes to indicate different social groups. We might then set one (relatively low, possibly null) discount rate for evidence received from those in the same group, and one or more other (higher) rates for those from different groups. This would embody a homophily assumption not based on differences in credences (which, after all, might be difficult to determine). An alternative approach would simply set a single discount rate applied by all agents determined by the group label of the speaker, thereby abandoning the homophily assumption entirely.

[7] In fact, O'Connor and Weatherall (2018) use two different rules: they also run simulations using a rule - namely, $P_f(E) = \max(\{1 - d \cdot m \cdot (1 - P_i(E)), 0\})$ - that allows for 'anti-updating' (whereby the final credence in E is less than the initial credence in that same proposition). Here we discuss only simulations with no anti-updating - i.e. those using the rule given in the main text.

evidence they themselves produce at all, since d = 0: accordingly, the final probability of the evidence they provide themselves with is 1; that is, they treat this evidence as certain. But for evidence supplied by other agents, as d increases, the product of d and the 'mistrust multiplier' m increases until it reaches (and then exceeds, but is replaced by the smaller value) 1: at this point $P_f(E) = P_i(E)$; in other words, the evidence provided is simply ignored (and treated as precisely as certain as it was prior to its being given).

2.3 The Method of Philosophical Simulation

Like the researchers discussed above, we run computer simulations based on our models of the social situations we wish to investigate. Our (python) code first builds a network (of a specified kind and size - see below) and assigns initial credences to the agents at its nodes (drawn from a random uniform distribution in all of our simulations to date). It then simulates that group of agents conducting experiments, sharing results, and updating credences in the manner specified in the model – with the values of certain key parameters (e.g. epsilon, number of trials) set. It runs until all agents have credence less than 0.5 or over 0.99 in B – or else, in some mistrust models (those leading to polarization – see below), the differences in credences preclude any further dynamics (because any new evidence produced by B believers is ignored - i.e. completely discounted - by the A believers).

Some might worry that our simulations will not be appropriately realistic - that is, that they may not have 'external validity'. In particular, it might be thought that the simulations we run make erroneous assumptions about either: (i) the nature of the individuals that constitute our social groups (specifically, it might be thought that people are not rational though we assume they are); or (ii) the character of the networks that bind them into communities; or both. If that is right, then (it is argued) we will not be able to draw (legitimate) conclusions - and in this way learn - about the situations we are modelling from findings about our simulations.

In response to this objection, and by way of clarification of our method, we must, unfortunately, be brief. We begin with point (i). First, while many irrationalities have been suggested (Kahneman, 2011), there is some debate around whether reliance on such heuristics is genuinely irrational (Mousavi and Gigerenzer, 2017). Second, we are idealizing: just as physicists model situations as involving no air resistance, or a frictionless plane, so we too imagine our social groups as comprising agents without irrationalities. This simplification allows us to approximate the truth before refining our models. Third, relatedly, this allows us to identify causes and provide explanations. Our aim is not (yet, at least) to make predictions about how social groups will behave when investigating a given issue under particular circumstances. (Thus, in a sense we require less of the external validity of our simulations than some might expect.) More specifically, fourth, we wish to uncover *structural* causes of ignorance (e.g. due to misinformation): if we assume (even if falsely) that agents are perfectly rational, and we nevertheless discover that communities of these agents remain ignorant of the truth in our simulations (for an extended period of time), then it will be apparent that (at least some) ignorance has a cause other than the psychological failings of the agents constituting them. These causes ultimately derive from the (social) structures within which the agents are embedded. And finally, fifth, we are also engaged in an interpretive, and normative, (humanities) investigation (not just a descriptive, scientific one): we are aiming to assess (a) whether our models of rational agents are appropriate, or *apt*, in the sense that they capture what we have in mind when we think of (the ideal of) rationality (this is the interpretive element), and (b) whether behaving in the ways our agents do really is rational (especially in light of the foreseeable social consequences) - that is, whether agents really *ought* to behave in those ways (this is the normative aspect). For all of these reasons, we believe that much can be learned about real social situations, even through simulations based on somewhat simplified models.[8]

Turning to (ii), we are aware of the concern – which is why, in addition to looking at highly controlled, artificial and relatively small networks, we also run our simulations on networks drawn from real world data sets (AUTHORS, UNDER REVIEW). Thus, unlike (at least some) previous efforts in 'computational philosophy' (Grim and Singer, 2022), our framework allows us to *scale* our simulations to run on large data sets - and we exploit the opportunities this affords. Moreover, whatever our findings in relation to smaller simulations on artificially generated networks, we hope to extrapolate from them to explain findings based on real world data sets with the structural features, whatever they might be, that actual social networks possess. For instance, we may find cleaner patterns within more controlled environments that may nevertheless be discerned - albeit with some noise - in more complex, realistic settings; though they might have been difficult to detect, or recognize, if we worked exclusively with real-world data sets. Nevertheless, we here focus on simulations run on small, artificial networks.

3. Results

Researchers employing philosophical simulations built on the models described above have obtained a number of results. We begin by sketching previous results, before reporting our own findings.

3.1 The Zollman Effect

Zollman (2007) found (amongst other things) that simulations (using Bayesian updating) converged to the true belief that B is better more quickly (i.e. in fewer steps) in complete networks (see figure 1) in which each agent is connected (by an edge) to every other than in cycle networks (see figure 2) where each agent is connected to exactly two neighbours (with the network as a whole forming a ring).[9] But he also found that the complete networks were more likely than cycles to converge to the false belief that A is better. In his own words, his results suggested 'two interesting conclusions. First, in some contexts, a community of scientists is, as a whole, more reliable when its members are less aware of their colleagues' experimental results. Second, there is a robust tradeoff between the reliability of a community and the speed with which it reaches a correct conclusion.' (2007: 574)

[8] Mayo-Wilson and Zollman (2021) defend 'simulation as a core philosophical method' (2021: 3647), offering two central arguments, and responding to objections. The second of their arguments is irrelevant for our purposes (since it concerns benefits to the researchers, rather than the research), but the first depends on the crucial claim that computer simulations can play many of the same roles as thought experiments do - indeed, they argue that the former can be better at doing so in relation to five of the six aims they identify for the latter. We are inclined to agree with much of what they say, and note that one of the aims they mention is to '[d]istinguish explanatory reasons and identify those causes that explain a phenomenon' (2021: 3650) - which is, at least roughly, one of the objectives we have given above in relation to our simulations. The only aim that Mayo-Wilson and Zollman recognize in the case of thought experiments, but do not argue is well served by simulations is that of '[e]licit[ing] normative intuitions' (2021: 3650) - but that is something that we do hope our simulations can help us to do, as hinted above. Of course, we agree that 'simulations almost never answer philosophical questions by themselves' (2021: 3649) - but we think they can play a role, within a given dialectical context, in drawing out normative intuitions that can help to settle an issue (such as what rationality requires).

[9] Zollman looked at various other small networks as well. Some of these will be discussed below.

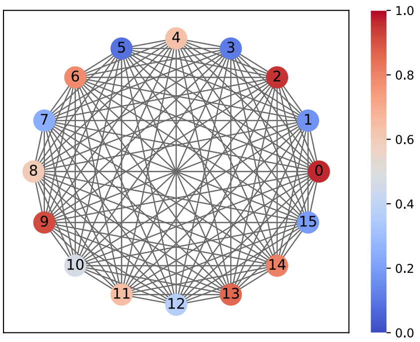

Fig 1: Complete network (size 16)

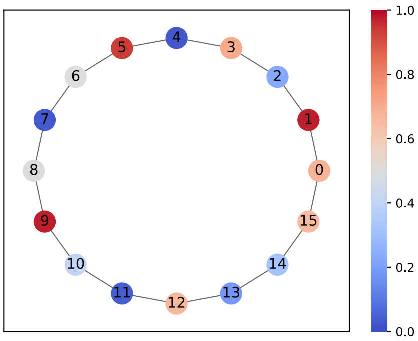

Fig 2: Cycle network (size 16)

O'Connor and Weatherall introduced the term 'Zollman effect' as a name for the 'phenomenon' in which 'scientists improve their beliefs by failing to communicate' (2019: 63) – i.e. the first of Zollman's findings (in his own report above). Interestingly, however, Rosenthal, Bruner, and O'Connor (2017) speak of 'the general effect noted by Zollman—that sparser networks are more reliable than well-connected ones.' Now, one network is sparser than another, when it has lower *density* – where this is the number of actual connections (or edges) in the network divided by the number of potential connections (which itself is half the number of nodes n multiplied by n-1). (Thus, a complete network has the highest possible density: 1.) And yet O'Connor and Weatherall (2018) suggest (in footnote 28 on p.198) that the Zollman effect concerns 'social structure': and while density is determined by structure, there is more to the latter than the former. We will return to this below.

In any case, Rosenstock, Bruner, and O'Connor report 'results replicating Zollman's (2007) simulations with a wider parameter space' (2017: 241) than Zollman himself employed. They 'find that parameters for which there is a notable benefit to decreased network connectivity occupy a relatively small niche of the total space.' (2017: 241) In particular: 'the difference in likelihood of successful convergence between the two network configurations [cycle and complete] decreases' (2017: 241) as epsilon increases; 'the difference in rates of successful convergence for the cycle and complete networks is more significant for lower [number] *n* [of trials]' (2017: 243); and crucially, '[t]he Zollman effect is strongest for smaller networks, but as network size increases it drops off' (2017: 244).

3.2 Polarization

O'Connor and Weatherall (2018) ran simulations in complete networks using Jeffrey's rule for updating, with the evidence provided by others being discounted increasingly with larger differences in credences (as described above). (As we have seen, this assumes homophily: we trust others more when they are more like us – and in particular, when their opinions are similar to our own.) They found that their simulations often ended in a state they refer to as 'polarization'. Of course, polarization can be - and has been - understood in a number of ways. O'Connor and Weatherall themselves say that 'polarization involves the emergence of two subgroups, one whose members all have credence > .99, and the rest with a variety of stable, low credences, such that they prefer the worse theory [that A is better]' (2018: 866).[10]

[10] They go on to say: '(More precisely, a stable outcome is one in which every agent either (a) has credence > .99 or else (b) has credence <= .5 such that their distance to all agents whose credence is > .99 satisfies m ∗ d >= 1.) Because the agents with low credences are outside the "realm of influence" of those testing the informative theory, they do not update their beliefs.' (2018: 866)

Reporting their findings on polarization (so understood), O'Connor and Weatherall say: 'In our models,… over all parameter values, we found that only 10% of trials ended in false consensus, 40% in true consensus, and 50% in polarization. These values should not be taken too seriously, since parameter choices influence where and when polarization happens, but the point is that adding evidential assessments based on shared belief dependably generates stable polarization.' (2018: 866) This requires unpacking, as well as qualification.

We begin with the qualification. In the (no anti-updating) formula (given above) that O'Connor and Weatherall (2018) use to operationalize Jeffrey's rule in the context of their homophily assumption, polarization is only a possible outcome when the mistrust multiplier is sufficiently large that min({1, d . m}) is 1. Since d itself is a real number between 0 and 1, this is only the case when $m > 1$. Thus, adding evidential assessments based on shared belief dependably generates stable polarization in this case only. So much for the qualification; now for the unpacking. Unfortunately, O'Connor and Weatherall do not themselves provide the information needed to unpack their findings in the case where mistrust can lead to polarization. Elsewhere [AUTHORS, UNDER REVIEW], we report the results of the simulations we ran (on complete networks) to uncover the requisite details. In brief, we found that ignorance – whether through error (i.e. false belief) or omission (i.e. agnosticism due to polarization) - was more likely to prevail in the long-run in O'Connor and Weatherall's polarization models (with $m>1$) than in Zollman's original models (keeping other parameter settings fixed). Here we focus on the case where the mistrust multiplier $m \leq 1$.

3.3 Discounts Delay

What happens in O'Connor and Weatherall-style models when polarization is not a possible outcome? To find out, we ran simulations on complete networks, using O'Connor and Weatherall's (no anti-updating) operationalization of Jeffrey's rule with mistrust $m = 1$. Since polarization is not a possible outcome in this case, we did not see the increased ignorance in the long-run due to agnosticism that we saw for $m > 1$. But when it comes to some of the issues that must be resolved through collective decision-making, such as efforts within democracies to avoid catastrophic climate change, timeliness is important; and so we looked at how long ignorance persisted in these models when it was eventually overcome. More specifically, we compared the mean number of steps in simulations that converged on the correct consensus that B is better in these models and in Zollman's original models. To ensure the robustness of our findings, we also considered ways of amplifying or dampening the effect of the distance between credences that, unlike in O'Connor and Weatherall's rule, cannot result in polarization. Thus, rather than multiplying the initial probability of the negation of the evidence ($1 - P_i(E)$) by min({1, d . m}) before setting the final probability of the evidence as 1 minus the result, we explored the effects of multiplying by the square, and by the square root, of d instead.[11] Table 1 below shows our findings.

Size	Epsilon	Trials	Model	Total (count)	Steps (mean)	U value	P value
16	0.001	16	Zollman	558	5121		
			Square(d)	94	7,867	36,844.5	0.00
			Mistrust m=1	430	14,649	193,902	0.00

[11] When $m = 1$, O'Connor and Weatherall's no anti-updating rule is equivalent to $P_f(E) = 1 - d \cdot (1 - P_i(E))$. We also considered the d-squared rule, $P_f(E) = 1 - d^2 \cdot (1 - P_i(E))$, and the root-d rule, $P_f(E) = 1 - \sqrt{d} \cdot (1 - P_i(E))$. Since $d \leq 1$, squaring dampens the effect of distance in discounting the evidence, while taking the root amplifies it.

			Root(d)	91	19,215	47,956	0.00
	0.001	64	Z	577	1,287		
			S	95	1,955	38,336	0.00
			M	421	3,233	202,012	0.00
			R	98	7,499	52,553	0.00
	0.01	16	Z	93	53		
			S	95	78	5,506	0.00
			M	430	129	31,830.5	0.00
			R	96	228	8,306	0.00
	0.01	64	Z	99	14		
			S	95	19	6,103	0.00
			M	435	59	34,939	0.00
			R	99	9,962	8,908.5	0.00
64	0.001	16	Z	597	1,710		
			S				
			M	451	7,197	234,197	0.00
			R	100	15,738	57,935.5	0.00
	0.001	64	Z	595	419		
			S	97	688	41,121.5	0.00
			M	456	1,575	244,823.5	0.00
			R	97	2,987	56,180.5	0.00
	0.01	16	Z	98	16		
			S	98	31	7,251.5	0.00
			M	452	61	39,782.5	0.00
			R	99	89	9.475	0.00
	0.01	64	Z	100	5		
			S	98	6	6,292.5	0.00
			M	453	13	38,155.5	0.00
			R	98	22	9,326.5	0.00

Table 1: Comparing O'Connor and Weatherall-style simulations with Zollman's original models. Jeffrey's rule is employed and evidence is discounted in line with a homophily assumption in the former. In the latter, Bayes' rule is used, and evidence is fully trusted. We used a Mann-Whitney U-test ($p < 0.05$) to show that ignorance persists for longer (i.e. more steps) in homophily models, even when it is eventually overcome. Discounting good evidence delays convergence to the truth.

As can be seen, agents' discounting of the evidence they receive in these models leads to significant delays in arriving at the correct consensus that B is better. Moreover, the mean number of steps required to converge to the truth increases as the discounting effect of the distance between credences becomes more pronounced. It should, of course, be remembered that, in these models, all agents are reliable truth-tellers, reporting the results of their experiments accurately, so that the evidence received is factive (i.e. true), even if (due to its probabilistic nature) potentially misleading as regards the underlying issue under investigation. It is therefore possible that in other contexts, where information sources are less reliable, discounting what is presented as 'evidence' may yield benefits. Nevertheless, the current findings provide a kind of benchmark result: discounting good evidence delays convergence to the truth; ignorance persists longer under mistrust than full trust.

3.4 Shape Matters

We reported (above) the finding, due to Zollman, that sparser (small) networks converge to the truth more reliably than denser ones: that is, for a given size of (small) network, the proportion of simulations that achieve consensus in the long run that B is better than A decreases with the density of the network; or, in other words, better connected (small) networks are more error prone.

For a given network size, the density of the network determines the average neighbourhood size in that network. Moreover, how much evidence a node receives is controlled by (the number of trials conducted by each agent and) the size of its neighbourhood. (It is not strictly *determined* by this, since how many of its neighbours take action B can vary while the size of its neighbourhood does not.) So network size and density together control the quantity of evidence each node receives on average at any given simulation step. Thus, what Zollman found was that, for a given size of (small) network, there being more information on average for each node yields a consensus more quickly, but – somewhat surprisingly – with a greater chance of error.

But there is more to network structure than density alone. Each network has a specific shape (or topology). Above we saw the complete (figure 1) and cycle (figure 2) networks. (Clearly, the complete network is more densely connected than the cycle.) A star network (figure 3) consists of a central node with all other nodes connected only to (themselves and) it. Networks of other kinds are generated by stochastic processes, with different instances of the kind having somewhat different shapes. For instance, a Barabasi-Albert network (figure 4) is generated by 'preferential attachment': as each new node is added to the network it forms a fixed number of attachments to previous nodes (1 in the network shown in figure 4), with the probability of forming an edge connecting it to a given node determined by the size of the neighbourhood of that other node; in other words, the more edges a node has, the more likely it is for the next new node to attach to it. Random networks (not pictured) are formed by making edges between any two nodes with a given, fixed probability; and Watts-Strogatz networks (not pictured) begin as rings, or cycles, of a given thicknesses (with each node connected to e.g. 2 neighbours, or 4, or…), with edges then rewired, with a certain probability, to connect to a randomly chosen node.

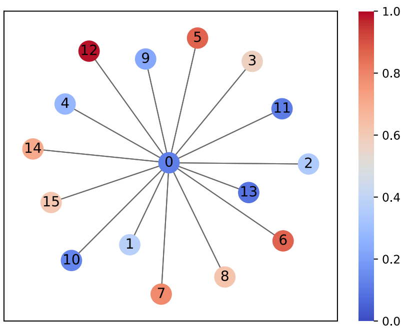

Figure 3: Star network (size 16)

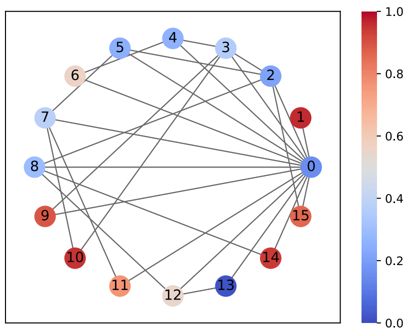

Figure 4: Barabasi-Albert network (size 16)

Do other structural features of networks, besides density, influence the knowledge or ignorance of the group? We compared simulations run (using Bayes' rule) on networks of different kinds, but with similar densities, and found that, even when it was eventually eradicated (with all the nodes converging on the correct consensus that B is better), ignorance persisted significantly longer in networks of some kinds as opposed to others. For example, in networks of size 64, cycles and stars, as well as certain Barabasi-Albert (attachments: 1) and Watts-Strogatz (knn: 2) networks, the density of connections is approximately 0.03 (i.e. roughly 3% of possible connections are actualized). And yet (with epsilon 0.001) the number of steps required for convergence to B differed significantly between the simulations run on these networks. Stars were the fastest. Cycles were next, though much slower. Barabasi-Albert networks slower still; and Watts-Strogatz networks were the slowest. We tested each pairwise comparison for significance (p<0.005) using a Mann-Whitney U-test – our findings are in Table 2 below.[12]

	Star (mean steps: 68,818)	Cycle (mean steps: 109,127)	Barabasi-Albert (mean steps: 119, 976)	Watts-Strogatz (mean steps: 120, 596)
Star (n=600)	-	U=294,660.5 p=0.00	U=43,898.5 p=0.00	U=31282.0 p=0.00
Cycle (n=600)		-	U=134,036.0 p=0.00	U=98,518.5 p=0.00
Barabasi-Albert (n=500)			-	U=107,933.5 p=0.04
Watts-Strogatz (n=400)				-

Table 2: Comparing the number of steps required for convergence to B in various kinds of (size 64) networks with similar densities (all approximately 0.03). In these simulations, each B node conducted 64 trials, and the value of epsilon was 0.001. Here we report the number of simulations of each kind, as well as the U- and p-values (<0.05) for each Mann-Whitney U-test performed.

These considerations show that the persistence of ignorance in a social network of rational (Bayesian) agents is influenced not only by the density of connections in that network, but also by the kind of network it is: the number of steps needed for convergence to the truth is significantly different between networks of different kinds with similar densities.

Interestingly, Zollman claimed that, in the small networks he tested, 'the in-network degree variance is not correlated with success' (2007: 482): that is, while the average neighbourhood size within a network affects reliability, variation in neighbourhood size does not. However, we found that variation in the size of the nodes' neighbourhoods can make a difference to the persistence of ignorance. We used a normalized measure we called 'moment_sd': it is the standard deviation in neighbourhood size divided by the mean neighbourhood size. Figure 5 (below) shows that (with epsilon 0.001 and 16 trials) steps increase as this measure increases in Barabasi-Albert and Watts-Strogatz networks with density

[12] With the parameter setting above, but with 16 (rather than 64) trials, we again found significant differences between networks kinds, except between Barabasi-Albert and Watts-Strogatz; the same was true for size 16 networks of these kinds, with both 16 and 64 trials. Increasing the numbers of attachments and nearest neighbours, however, we did see significant differences between Barabasi-Albert and Watts-Strogatz networks with both 16 and 64 trials, for: size 16, density 0.53; size 64, density 0.06; and size 64, density 0.12.

0.53.[13] While more work is needed to understand the way(s) in which it does so, it is clear that network shape matters.

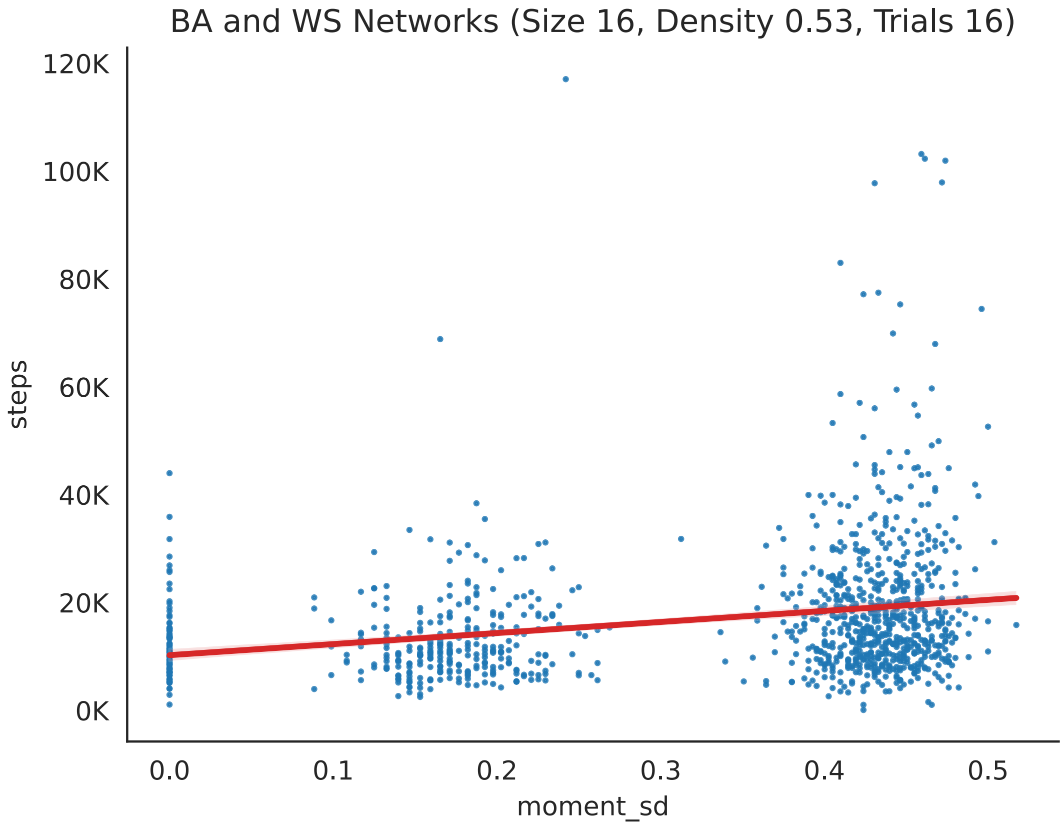

Figure 5: The effect of moment_sd on steps. R-squared: 0.064014. Slope (CI: 95%): 20513.379535 +/- 4644.725867. Number of simulations: 1100.

4. Concluding Remarks

We have explored ignorance in social networks, discussing computer simulations that employ simple models of inquiring communities. We avoided imbuing the agents in these models with psychologically motivated departures from rationality, in the hope of uncovering other (social and environmental) reasons why knowledge may fail to emerge in a timely fashion. Building on previous work, our results show: first, that when all agents truthfully report their own observations, discounting the evidence they provide delays any eventual correct consensus; and second, that the shape (and not just the connectivity) of the network formed by a community of agents influences how long ignorance is likely to persist within it.

These findings may be disconcerting. The first suggests that it will not do to respond to the threat of mis- and disinformation simply by advocating greater sceptical caution in incorporating the evidence available within one's social network: trustworthy information will not be efficiently assimilated if distrust runs rampant. The second finding may even suggest strategies for disinformants - such as

[13] The trend seen here was less pronounced, but still present, in these networks when trials increased to 64; and in size 64 networks with densities 0.12 and 0.06. In the lowest density networks (i.e. Barabasi-Alberts with 1 attachment and Watts-Strogatz, knn: 2) there was no significant correlation.

finding ways to occupy the central node in a star network (O'Connor and Weatherall, 2019). Yet this cuts equally in the opposite direction: those wishing to ensure the dissemination of trustworthy information may equally accommodate the point – whether they be social or other (e.g. news) media organizations, or governments or charities that provide oversight. In any case, the results themselves provide a starting point upon which further research on strategies for coping with our informational environment may build, and that constrain viable approaches to combatting mis- and disinformation.

References

AUTHORS, (IN PREPARATION).

AUTHORS, (UNDER REVIEW).

Grim, P. & Singer, D. (2022). Computational Philosophy. In E. N. Zalta & U. Nodelman (Eds.), *The Stanford Encyclopedia of Philosophy.* Metaphysics Research Lab, Stanford University, Fall 2022 edition.

Kahneman, D. (2011). *Thinking, fast and slow*. Macmillan.

Mayo-Wilson, C. & Zollman, K. (2021). The computational philosophy: simulation as a core philosophical method. *Synthese*, 199(1), 3647–3673.

Mousavi, S., & Gigerenzer, G. (2017). Heuristics are tools for uncertainty. *Homo Oeconomicus*, *34*, 361-379.

O'Connor, C. & Weatherall, J. O. (2018). Scientific polarization. *European Journal for Philosophy of Science*, 8(3), 855–875.

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.

Williamson, T. (2000). *Knowledge and its Limits*. Oxford University Press.

Zollman, K. J. S. (2007). The communication structure of epistemic communities. *Philosophy of Science*, 74(5), 574–587.