Cambridge Default Journal (2025), **XX**, 1–36

CAMBRIDGE
UNIVERSITY PRESS

Polarisation through Homophily. A Generative Algorithm

D.M. Mayerhoffer,[*†‡] M.A. Schulz,[¶§] S. Scheller,[§] and J. Schulz[|0‡]

†Amsterdam Institute for Social Science Research & Data Science Centre, University of Amsterdam
‡Contribution co-funded by the German Federal Ministry of Research, Technology and Space (grant: DATI Pilot, e-lernen)
¶Institute of Philosophy, FAU Erlangen-Nürnberg
§Institute for Political Science, University of Bamberg
|0Economics Institute, University of Bamberg
*Corresponding author. Email: d.m.mayerhoffer@uva.nl

Abstract

This paper presents a generative algorithm for simulating network polarisation based on attitudinal homophily, i.e., the tendency to connect to others with similar attitudes as oneself. To do so, it applies the notion of preferential attachment to node properties other than degree, aiding intuitive communication within and beyond the network science community. The algorithm works with one or more flexibly weighted attitude dimensions, heterogeneous populations. The generated networks commonly share features of real-world social networks such as (weak) small-worldiness. They can contribute to how-possibly explanations of stylised empirical facts, especially when empirical network data is missing. An example on polarisation in Germany on migration attitudes is used to illustrate this. We find that homophily can generate patterns resembling empirically reported polarisation and individual perceptions, such as viewing one's own opinion as moderate, over-estimating the actual level of societal polarisation, and dwindling open-mindedness towards others with different attitudes.

1. Introduction

Polarisation threatens the foundations of liberal democracy when political disagreement turns into mutual hostility. Social networks play a central role in this process; indeed, they are arguably what is getting polarised. Hence, network science provides tools to study how polarisation emerges and persists. Specifically, homophily — the tendency to link to others with similar views — is a likely candidate mechanism for generating polarised network structures (McPherson, Smith-Lovin, and Cook 2001).

While existing research has produced insights into belief and affective polarisation (Iyengar and Westwood 2015), we know far less about the generative mechanisms through which segregated network structures that feature polarisation between communities emerge from an initial non-polarised attribute distribution: Classical opinion dynamics models such as Degroot (1974), Hegselmann, Krause, et al. (2002), or Deffuant et al. (2002) specify how individuals update their beliefs given an existing set of neighbours but do not address how polarised network structures arise in the first place. Co-evolutionary models allow networks to rewire as opinions diverge (Holme and Newman 2006; Nardini, Kozma, and Barrat 2008), yet they typically generate segregation as a consequence of opinion dynamics rather than modelling attitudinal similarity as the primary mechanism of tie formation. Moreover, empirical study of these mechanisms is often hampered by a lack of empirical network data. Mechanisms that feature homophily to generate synthetic networks may help bridge these gaps.

To specifically tackle homophilic linkage, one must first set the node property on which the homophily mechanism. Should operate. Prima facie, there appear three good candidates for this: I) Group identity to model homophilic intra-group vs heterophilic inter-group linkage, whereby association between attitude(s) and group identity can be strong (e.g., partisanship) or weaker (e.g., socioeconomic stratum). Kim and Leskovec (2011) and Karimi et al. (2018) propose mechanisms for this. II) The node's location in the network for preferred linkage between geographically close nodes, as done by Santos, Lelkes, and Levin (2021) Segregation in attitudes will occur only indirectly (or as a correlate of) homophily in group identity or geographic location. III) To model homophily in attitude(s) themselves, one must, of course identify the relevant attitude. Baumann et al. (2020) do this with a power-law decay in linkage probability for growing attitude dissimilarity, and get temporal networks to simulate interactions on social media. Notably, a single target for the homophily mechanism is not always sufficient because empirical research shows that political attitudes are increasingly multidimensional and structured by varying issue salience (Mason 2018; Hare 2022). Finally, one must also set the functional form that governs the homophilic linkage, i.e. the decay and linkage probability with growing dissimilarity in attitudes or node attribute. All functional forms currently discussed in the literature allow for a flexible homophily strength and setting this parameter is also important for generating homophilic networks. Depending on the application case, it might be necessary to calibrate the generative function as well as the strength parameter theoretically and/or empirically.

To complement existing approaches, we introduce a parsimonious general-purpose generative algorithm that simulates attitudinal network formation through homophily, flexibly incorporating continuous, weighted attitude distributions as base for this homophily without a need for a prior group allocations of individual nodes. Its core idea resembles preferential attachment (Albert and Barabási 2002), but applied to any node property, namely the similarity of nodes with regard to this property. Thereby, the homophilic neighbour selection is rooted in discrete choice. Furthermore, the formal framework explicitly and flexibly accommodates multiple attitude dimensions with different relative importance in tie formation, heterogeneous populations and directed/undirected links. It can yield graphs that resemble selected properties of empirically observed polarised networks.

Hence, the proposed algorithm offers a conceptual and practical basis for how-possibly explanations of stylised empirical facts, e.g., how observed attitudinal clustering and segregation can plausibly arise from similarity-based tie formation, how segregation on one dimension emerges as a byproduct of homophily on others, or how beliefs or attitudes and ties co-evolve over time. The algorithm may be particularly useful when empirical networks are unknown or partially observed,where it generates realistic surrogate networks for simulating processes such as social sampling, belief updating, and other dynamic interactions. We demonstrate the use of the algorithm to explain misperceptions about polarisation in an exemplary case of polarised attitudes towards migration in Germany. To this end, we generate a network based on attitudinal survey data, where no information on social ties is present in the survey.

The paper proceeds as follows. Section 2 clarifies key concepts surrounding polarisation. Section 3 introduces the model framework, including the homophilic linkage mechanism. Section 4 outlines how the model supports how-possibly explanations and how its outputs can be compared with stylised facts and other empirical benchmarks, before exploring its use in simulating belief-updating. Section 5 presents an empirical case using migration attitudes in Germany. Section 6 concludes with a discussion of contributions and future directions both on the exemplary case and the algorithm in general.

2. Key Concepts

The concept of polarisation has no doubt made its way into the established vocabulary of describing and discussing politics, especially so in the US context: A Google search for the term "polarization" on the New York Times website returns 2,170 results for the year of 2024 alone (an average of six articles per day). As is perhaps unsurprising for a concept that has such currency both in public discourse and in several academic disciplines (political science, sociology, psychology, philosophy, media studies), it is not always clear whether the term is used consistently, or indeed what is denoted by it in the first place (Lelkes 2016). It is not, of course, the aim of this paper to contribute to explicating this concept and arbitrate among divergent uses. Still, it will be helpful to give some indication of the sort of phenomena which we understand may be studied using the proposed generative algorithm. To this end, we will in this section first suggest a core sense of polarisation operative across different contexts, and then briefly sketch what the term is usually further associated with when it is wielded as a timely diagnosis of putative problematic trends in Western democracies.

At its core, polarisation is a collective-level concept that picks out a typically bimodal distribution of relevant features within a population which combines *inter-group heterogeneity* with *intra-group homogeneity* (Mehlhaff 2024). So for this concept to be applicable, there must be at least one dimension on which individuals vary and individuals need to be attributable to groups (which may emerge from the distribution itself or may be identified independently, e.g. through party affiliation). When we say that the population is polarised with respect to this dimension, we imply two things: First, the groups significantly lie apart – for instance, Democrats have very different attitudes towards immigration than Republicans do. And second, members of either group are very much alike. This does not necessarily follow from the first: In theory, Republicans and Democrats could form non-overlapping sets with widely diverging means but still be spread out wide, such that there are significant disagreements among members of either group. Yet when we speak of polarisation, we typically have in mind that individuals gravitate towards tight-knit clusters on opposing ends of a spectrum.

In addition, it is usually taken for granted that the dimensions along which individuals vary are broadly speaking attitudinal or attitude-dependent features. For instance, we would not ordinarily hold that a given society is strongly polarised along the dimension of biological sex, even though it might be correct that biological males and females strongly resemble one another in sex-constitutive features (intra-group homogeneity) and both groups are in these respects quite unlike one another (inter-group heterogeneity). Rather, discussion of polarisation revolves around features such as normative beliefs or preferences (e.g., immigration policy, redistributive taxation, morality of abortion), empirical beliefs (e.g., dangerousness of vaccines, existence of climate change, theism), affective dispositions (e.g., hatred or fear of the outgroup, warm feelings towards the ingroup), and resultant behaviours (e.g., voting, social interaction).

Social psychologists have long studied subsets of these phenomena under the rubrics of *belief polarisation* and *affective polarisation*. In this literature, belief polarisation is understood to denote the phenomenon that deliberation in a like-minded group reliably causes the members' beliefs to become more extreme (Moscovici and Zavalloni 1969; Myers and Bishop 1970; Baron et al. 1996; Sunstein 2002, 2009), where 'more extreme' is sometimes best construed as an increase in confidence in an existing belief and sometimes as the acquisition of a new belief that sits further along on the relevant attitudinal dimension, in the direction congruent with the group's established identity (Talisse 2021, 213–15). Information-based accounts of the underlying mechanism have highlighted how either kind of shift can be subjectively epistemically rational for group members given their exposure to corroborating evidence to the exclusion of countervailing arguments (Burnstein and Vinokur 1977; Begby 2022). Such epistemic dynamics have garnered particular interest with regard to 'echo chambers' or 'filter bubbles' (Pariser 2011; Steppat, Castro Herrero, and Esser 2022; Hobolt, Lawall, and Tilley 2024; Iyengar, Konitzer, and Tedin 2018), especially on social media, which have in turn been conjectured to drive polarisation (Sunstein 2018; Vaidhyanathan 2018; Bruns 2019; Nyhan et al. 2023). Alternatively, the underlying mechanism has been theorised to accord with Social Identity Theory (Tajfel and Turner 1986; Hornsey 2008), whereby the mere fact of salient group membership is thought to propel an interest in ingroup cohesion and further distinction from the outgroup, such that members align their beliefs to those of others merely upon learning that the belief is characteristic of group members (Myers et al. 1980; Abrams et al. 1990; Lee 2007; Anderson 2021).

In turn, a concern with social identity is shared by scholars who study *affective polarisation*, typically associated with *partisanship* within the purview of political science (Huddy, Mason, and Aarøe 2015): Whereas Social Identity Theory generally expects group membership to cause a favourable evaluation of the ingroup as compared to the outgroup (Billig and Tajfel 1973), this literature is specifically interested in how political identity causes value-laden affective dispositions (such as fear or loathing of the outgroup) and discusses these under the rubric of polarisation (Iyengar and Westwood 2015; Iyengar et al. 2019).[1]

This richer notion of polarisation as affect-laden and intertwined with partisan identities brings us closer to the term's mainstream use as a proliferating anxious diagnosis in political discourse (Iyengar 2016; Finkel et al. 2020; Sides, Tesler, and Vavreck 2018; Mason 2018). Here, it tends to be invoked when we observe a society-wide pattern of (increasing) inter-group heterogeneity and intra-group homogeneity across an (increasing) range of attitudinal and behavioural dimensions, which become associated with and more tightly integrated into a salient social — specifically, political — identity, a pattern commonly described as partisan sorting. In such accounts, attitudes or aggregate identities that were formerly intercombinable become increasingly perceived as incompatible, increasing the number of features people readily understand as signifiers of group membership. Accordingly, it has been associated with polarisation that attitudes on long-standing policy controversies become increasingly unidimensional (Hare 2022): One's views on immigration strongly predict their views on tax policy and on abortion rights. It is in this sense in which Bakker and Lelkes (2024, 419) hold that "polarisation ultimately refers to a collapse of dimensionality and the flattening of conflict". Yet beliefs beyond policy preferences equally seem to get implicated: For instance, your belief about whether a woman can have a penis might predict your belief about whether the sun was shining during President Trump's first inauguration, even though the two issues likely appear utterly unrelated to a naive outside observer (Schulz and Scheller 2024). In general terms, increasing pressures of

1. Indeed, Bakker and Lelkes (2024) admonish that the study of affective polarisation is squarely focussed on partisan identity and hardly at all on affect.

intra-group homogeneity have plausibly been conjectured to co-opt contingent markers of group affiliation into constituting and maintaining the group identity (Mason and Wronski 2018) – leading Talisse (2020, 84) to contend: "It's no stretch to say that today politics *simply is* our lifestyle." Such enriched and salient identities in turn impact preferences for social interaction beyond the realm of politics: Party affiliation has in the US become a major factor in the selection of potential romantic partners (Huber and Malhotra 2017; Iyengar, Konitzer, and Tedin 2018); formerly apolitical fora of civil society such as churches or social clubs increasingly overtly identify as liberal or conservative (Talisse 2020); and certain professions skew along partisan lines (Bonica, Chilton, and Sen 2016).

To summarise: When we talk about polarisation, we generally have in mind an attitude distribution that combines inter-group heterogeneity with intra-group homogeneity. When we worry about polarisation, we usually have in mind that polarisation clusters across dimensions into purified and affectively opposed identities.

3. Model Framework

3.1 Overview

The parsimonious generative model presented in the following builds upon existing frameworks for network formation based on income homophily, as introduced in previous research (Schulz, Mayerhoffer, and Gebhard 2022; Mayerhoffer and Schulz 2022; Schulz and Mayerhoffer 2023). This model adapts the generative algorithm to provide a setting in which possible patterns of polarisation and social division can be explored under the assumption that network linkages depend on attitudinal homophily. The network generation algorithm lets agents choose their own neighbours based on a straightforward choice rule: Stochastically prefer nodes whose attitude(s) is (are) closer to yours. The strength of this stochastic preference can be flexibly adapted through a single parameter, the homophily strength ρ. This epistemically, ontologically and mathematically undemanding makeup renders the algorithm is versatile: it can be applied to any attributes or sets of attributes that can be quantified on a monotonous (preferably but not necessarily linear) scale. It is even possible to assess homophily based on a combination of attitude and demographics, such as gender or age. Furthermore, it can specify the number of links formed by a node (and hence its minimum degree), e.g. to capture different layers of social interaction (MacCarron, Kaski, and Dunbar 2016).

3.2 Agent Attitudes

The model simulates a population of N agents, each with a unique *attitude score* A_i for agent i, which is drawn from a continuous or discrete distribution representing the range of possible attitudes within a population.[2] The attitude score A_i can reflect a single attitude dimension (e.g., political ideology) or can be extended to a vector of multiple attitude dimensions $\mathbf{A}_i = (A_{i,1}, A_{i,2}, \ldots, A_{i,m})$ if linkage based on their combination is of interest. For each agent, $\mathbf{A}_i$ should contain identical attitude dimensions, e.g., if migration attitudes are part of the model, all agents should be initialised with information on their migration attitude. Furthermore, attitudes need to be on (approximately) cardinal scale so that the homophilic linkage mechanism (cf. Section 3.3) can be applied to them. It is advisable to normalise all attitude dimensions, dividing all nodes' attitude values by the mean (of the respective dimension). This ensures that the segregation in the simulated network depends only on the homophily (introduced below), and is not impacted by attitude scale.

Apart from these formal requirements, the notion of attitudes, as a technical one, may be operationalised in any way fit for the study in question. Attitudes can be taken from empirical data,

2. The model has several features of an Agent-Based Model. The reason why we do not dub it as such is that one core feature is missing for the baseline presented here: While agents base their decisions whom to link to on others' attitudes, they do not directly react to the choices made by others within the link formation itself. Only if one uses the generated network to model some form of interaction or even simple mutual observation (as presented in Section 5), this last commonly cited feature of an Agent-Based Model is fulfilled, too.

other simulation models, or theoretically inferred distributions. Moreover, "attitude (dimension)" may refer to any node property, including demographics, on which one intends agents to form their links.[3] (For the sake of simplicity, the remainder of this paper will refer to the relevant properties as attitudes.) If, for example, one would like to model homophily with regard to both age and beliefs about social security, $\mathbf{A}_i$ will consist of an age variable and one or multiple measures of the desired level of social security either as multiple variables (e.g., one on the level of unemployment benefits and another on social insurance contributions) or as a single measure combining them.

Any other basic agent characteristics concern their linking behaviour (see below). In addition to these basic properties, additional ones can be introduced, for example if nuanced interaction is to be modelled (cf. Section 4.5).

3.3 Link Formation Based on Attitudinal Homophily

The actual connection between agents follows a discrete choice mechanism, as outlined in Manski and McFadden (1981). In a first step, an agent i calculates a weight $w_{i,j}$ for potential linkage with each other agent j. Based on these weights, the agent then draws d_i link-neighbours.

In the single-attitude case, the weight $w_{i,j}$ is given by:

$$w_{i,j} = \frac{1}{\exp(\rho_i \cdot |A_i - A_j|)} \tag{1}$$

where $\rho_i \in \mathbb{R}_0^+$ is the homophily parameter, determining the rate of decay in link probability with increasing attitudinal distance. Trivially, $w_{i,j} = w_{j,i}$ if $\rho_i = \rho_j$, and hence homogenous ρ implies fully symmetrical link weights.[4]

The key idea of the weighting function is to penalise differences in attitudes between i and j (shown in the denominator). The exponential nature of this penalty ensures that small differences receive only mild punishments, whereas large ones result in particularly low weights. As such, the weighting emphasises the homophilic tendency to pick others with close attitudes, making them much more likely as link-neighbours. Thereby, $\forall i \rho_i = 0$ implies equal weights for all agents. For an increasing ρ, the weights of those others with great attitude differences shrink quickly (cf. Schulz, Mayerhoffer, and Gebhard 2022, Appendix A for an analytical proof of this mechanism). Figure 1 showcases the effect of homophily in an example case.

In the multi-attitude case, the weight equation adapts as follows:

$$w_{i,j} = \frac{1}{\exp\left(\rho \cdot \sum_{k=1}^{m} \omega_{i,k} |A_{i,k} - A_{j,k}|\right)} \tag{2}$$

where the homophily parameter $\rho_i \in \mathbb{R}_0^+$ works as outlined above, and where $\omega_{i,k}$ is a weighting factor that determines the relative importance of attitude dimension k in agent i's overall homophily calculation. This weighted multi-dimensional approach allows researchers to explore how alternative weights assigned to different attitude dimensions are associated with clustering and polarisation in the generated networks, adjusting the values of $\omega_{i,k}$ to emphasize or downplay specific dimensions. Like for the one-dimensional case, $\rho_i = \rho_j$ and $\forall k(\omega_{i,k} = \omega_{i,k})$ imply $w_{i,j} = w_{j,i}$.

Link weights are straightforwardly translated to probabilities when divided by the sum of weights that the agents assign to all others. Based on these probabilities, i establishes their desired number

3. One could even encode group membership as "attitude (dimension)" to model group-based homophily. Likewise, this is possible for proximity to other nodes to model homophily based on proximity in the network in some form of iterative process.

4. The algorithm is also straightforwardly applicable to j's attributes without direct comparison between i and j. This can be used to, for example, model people with extreme attitudes as especially (un)attractive link neighbours, or to cover linkage on demographic characteristics (such as a propensity to listen to those living in specific places).

of links d_i, using a multinomial sampling approach. If links are directed, it is more intuitive to think of agents forming their in-links (from the other agent to oneself) because the notion of the communication structure detailed below is one of individuals choosing whom to listen to rather than whom to address (consider, e.g., following others on social media).[5] In that case, d_i specifies i's in-degree. For undirected links, in contrast, d_i specifies i's minimum degree for other agents can still select i as their link-neighbour if i has not selected them. Even after i has selected j as link-neighbour, i remains in the set of j's potential choices with no change to $w_{j,i}$. Should j also pick i, there still is simply one undirected link created[6] and neither of the agents picks an additional link-neighbour. Consequently, the order in which agents are picked for creating links is irrelevant to the resulting network (safe for the impact of randomness).

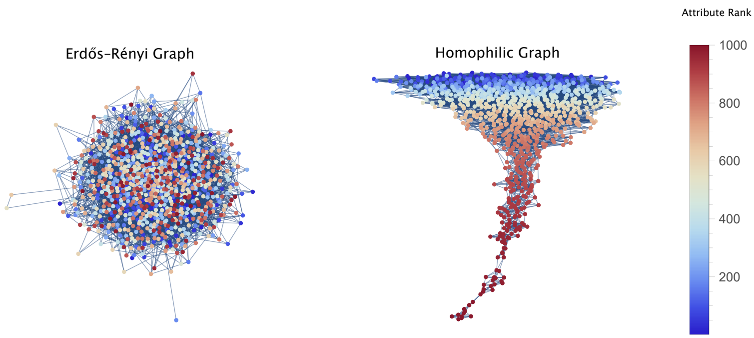

Figure 1. Illustration of network segregation under attribute homophily. 1000 nodes, single log-normally distributed attribute with $\sigma = 1\,\forall i d_i = 5 \wedge \rho_i = 10$, undirected graph. Visualisation uses Mathematica 13's "Spring Electrical Embedding" routine.

3.4 Communication and Sensing

Once the network is established, there are several ways to model agents' perceptions and interactions: (i) simple sensing of immediate or higher-degree neighbors' attitudes; (ii) social learning with updating one's own attitude based on network information and (iii) political action based on social influence. Section 4.5 details these options and reviews several well-studied processes in the network literature.

3.5 Accounting for (Random) Variation

Since attitudes, ω, ρ, and chance interact, drawing appropriately qualified inferences (see Section 4.1 for details) from the networks generated by the algorithm requires preparation in setup. First, when working with theoretical input distributions, one might want to experiment with different distributional shapes. Second, exploring the empirically plausible combinations of attitude weights ω is necessary. Third, this also goes for values of ρ, where, as a rule of thumb, one might want to depart from a random network ($\rho = 0$), and increase ρ until the network reaches a level of segregation which is inadequate for the case studied (e.g., the network is no longer connected or the components get too small). Generally, the space to be explored depends on how much is known about the inputs. If, for example, attributes or their weights are known from empirical research (cf. Section 5 for an example), there is no need to vary these parameters, except for intentionally exploring counterfactuals.

5. Note, however, that formation of out-links is also possible and, of course, technically mostly equivalent.
6. In a model variant, one could, of course, add higher weights to links that are established by both ends.

Finally, to account for the impact of chance, one should have multiple networks simulated - do so-called Monte Carlo runs - for each combination of input values. Since the algorithm is computationally efficient, one might easily opt for a high number, such as at least 50, or even more if opting for heterogenous, randomly distributed values of ρ or ω. However, should one intend to run computationally demanding calculations on the resulting networks (such as calculating the clustering coefficient on larger graphs), working with fewer Monte Carlo runs might also be feasible, especially for homogeneous ρ and ω, where results tend to be fairly stable.

4. Application of the Network Simulation in Your Workflow

The generative algorithm proposed above alone does not answer any research questions, but the resulting networks can contribute to such answers in various ways, which are discussed in this section. Figure 2 provides an overview on how the different use cases relate to the model mechanism.

Besides its use for analytical work or theory building, it can bridge gaps in empirical data in two forms: First, for surveys that collect attitudes but no or incomplete network data, a network may be generated, for example to assess whether correlations (such as between partisanship and theoretically unrelated social attitudes) may be caused by network formation. Similarly, it allows for the study of echo chamber effects in social media environments where platform-level network data are unavailable. Second, simulation of homophilic linkage can offer a longitudinal lense when network data is only known for one cross-section but there are changes in issue salience, for example during the 2015 refugee inflow in the EU, which might impact the co-evolution of network and attitudes.

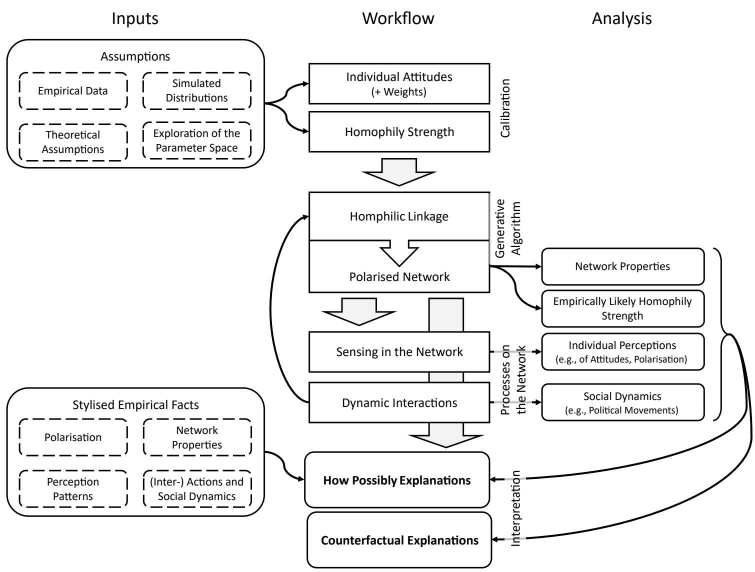

Figure 2. Workflow and use-cases of the homophily algorithm

4.1 *How-Possibly Explanations of Stylised Empirical Facts*

In the social sciences, stylised facts refer to simplified empirical regularities that capture broad, recurrent patterns across cases, abstracted from the noise and complexity of individual observations. Initially introduced by Kaldor (1961) in macroeconomics, the notion of stylised facts intends to allow theorists to construct models based on broad tendencies rather than the precise — and often imprecise — details of specific datasets. Kaldor saw stylising facts as a necessary abstraction to enable theory-building in the face of messy real-world data. Stylised facts typically represent unexplained empirical associations rather than fully elaborated causal claims (Hirschman 2016). For example, the persistent clustering of political attitudes within social networks, or the alignment of demographic characteristics and ideological positions, are stylised facts if they are consistently observed but not yet understood in terms of underlying mechanisms. In this sense, stylised facts are not endpoints of inquiry but starting points: empirical phenomena that call for theoretical explanation.

One response to stylised facts is through the use of how-possibly explanations. Unlike how-actually explanations, which aim to identify the true mechanism driving a phenomenon in a specific real-world case, how-possibly explanations offer candidate mechanisms that could plausibly account for the phenomenon. They help address the question: How could this empirical pattern conceivably come about? (Epstein 1999; Reutlinger, Hangleiter, and Hartmann 2018) In this context, simulation models are valuable tools for constructing how-possibly explanations. They develop "parallel realities" (Sugden 2009) where specific mechanisms operate in silico to approximate stylised empirical patterns observed in real-world data. Such "surrogate systems" allow to study a target system where direct observation of the latter is not possible (Mäki 2009). Our homophily-based generative algorithm is intended to serve this purpose. By modelling tie formation based on attitudinal similarity, it can generate network structures resembling stylised patterns of attitudinal clustering, segregation by political ideology, and associations between demographic and attitudinal characteristics. Correspondence between simulated and observed patterns would be compatible with homophily as a candidate explanatory mechanism and could motivate further investigation. It would not, by itself, identify homophily as the actual cause of these patterns in any particular empirical case.

4.2 Evaluation of the Generated Graphs

The networks generated by the algorithm can straightforwardly be evaluated through standard graph-theoretic metrics, i.e. clustering coefficient (to measure the degree of local cohesion within the network), degree distributions (to assess whether the network exhibits features such as right-skewness or hubs), modularity (to capture the extent of community formation and segmentation), and assortativity, particularly with respect to attitudes (to detect the emergence of polarisation), as well as connectivity and path length measures (to see how extreme segregation is). Of course, one can also use these basic properties to calculate compound ones, e.g., to assess small-world properties (Watts and Strogatz 1998). Since the algorithm explicitly works on node attributes, the resulting graphs are attributed (Bothorel et al. 2015), and hence further case-specific methods should be considered.

These metrics may be complemented by visual inspections to gauge face validity. While a detailed calibration to empirical data is addressed in the following Subsection 4.3, in some cases a mere eyeballing might be enough to validate the graphs as internally coherent and externally in line with plausible social configurations.

Notably, for many input parameters, the algorithm produces graphs that share key properties with Random Geometric Graphs, a well-studied class of spatial networks (Newman and Ziff 2001; Dall and Christensen 2002). Like these, the networks generated by the present algorithm typically exhibit high clustering, right-skewed degree distributions, and a natural tendency towards local connection, all of which emerge from the underlying similarity-based linking rule in the simulation. Firstly, this geometric structure aligns the generative model with observed topologies of real-world social and attitudinal networks, thereby reinforcing the plausibility of the simulated outcomes.[7] Secondly, methods for studying Random Geometric Graphs are good candidates if one desires a deeper analysis of the graph structure.

7. Note that the similarities, by themselves, establish that the model captures the actual formation process of empirical networks. They merely show that the outcome can be approximated by the generative algorithm.

If one does not know the features of an empirical network (e.g., due to lack of data), one can use the model to see which graphs are compatible with the input attitudes. With that, one can benchmark empirical networks of schools, workplaces, or online communities when the real-world network topology is only partially known. Thereby, sensitivity analyses for a wide range of ρ and ω are necessary. The more information one has about the configuration of the corresponding empirical network, the easier it becomes to narrow down these values, as will be discussed next.

4.3 Gauging Empirically Likely Homophily Levels and Predicting the Effects of Changes in Homophily

Moving beyond face validity and sensitivity analysis, the model allows more rigorous validation strategies by comparing the simulated networks' structural properties, attitude compositions, or behavioural patterns to empirical data or expectations derived from theory.[8] Networks generated under empirically plausible homophily parameters often exhibit weak small-world properties (Watts and Strogatz 1998), high clustering, and distinct attitudinal clusters — features that are consistent with observed social and political network structures. Knowledge about how attitudes are distributed in empirical networks can help narrow down plausible values of ρ and ω.

This approach fits into what Squazzoni (2012, ch. 4.3) labels "multi-level empirical validation", which combines two logically distinct but substantively interdependent tasks: The first is empirical specification, the calibration of model inputs — here, the attitudinal distributions. The second is empirical validation of model outputs against stylised facts and known empirical patterns. Importantly, this process retains the explanatory orientation of the model. While not aiming to pinpoint the exact generative process in any specific empirical case (as in a how-actually explanation), the model's ability to replicate empirically grounded regularities supports its role as a how-possibly explanation. In this context, empirical homophily levels inferred from empirical data (e.g., attitudinal agreement among ego-alter pairs) can inform the model's parameter space. Conversely, the model can be used to identify ranges of homophily parameters under which it generates levels of segregation or polarisation comparable to observed benchmarks, offering a tool for counterfactual or policy-relevant scenarios.

Additionally, the model can be used to explore the effects of hypothetical changes in homophily – for instance, reductions brought about by deliberate mixing policies or increases due to rising affective polarisation, or whether increased ideological hostility would produce pronounced community fragmentation under current network conditions These experiments can suggest possible nonlinearities, tipping points, or emergent dynamics that would be difficult to anticipate through intuition alone.

4.4 Investigating the Complex Relationship Between Attitude Dimensions and Homophily

The model allows for the investigation of interdependencies between different attitude dimensions by simulating how homophily on one dimension may lead to segregation on another — even if the latter is not directly involved in link formation. One may simply achieve this by setting the homophily weight for the dependent dimension k to zero, i.e., $\forall i(\omega_{i,k} = 0)$. If the resulting network nonetheless exhibits significant clustering or assortativity on that dimension, this may suggest the opportunity to investigate whether segregation could be related to structural homophily on the other included dimension. This approach might often be interesting to assess the impact of homophily in demographic features on attitudinal segregation, or vice versa.

Such an analysis can compare associations already present in the input data with those observed after network generation. It thereby allows researchers to examine whether the modelled linking

8. For the sake of simplicity and because we expect most of our readers to have this background, we will in the following only talk about empirical data. Yet, our explanations also cover an application of the simulation model to theory building (Taghikhah, Filatova, and Voinov 2021; Schwaninger and Grösser 2008).

process is compatible with an amplification of cross-dimensional associations. In particular, the model can identify parameter settings under which associations resembling observed attitudinal alignments could arise from homophilic network formation. Further empirical evidence would be required to assess whether this possibility applies to a particular real-world case.

This approach may be especially useful when the interaction between different homophily levels and attitude-dimension weights is analytically intractable. The simulation framework can map regions of the parameter space in which co-segregation occurs and thereby help formulate candidate how-possibly accounts of affective polarisation, ideological bundling, or demographic echo chambers outlined in Section 2. In this sense, the model is a tool for exploring whether network formation could contribute to observed attitudinal alignments and for distinguishing, within the simulation, patterns introduced by the structural assumptions from patterns already present in the input data. Furthermore, simulated networks can aid testing whether racial or socioeconomic homophily alone could generate apparent ideological "bundling", even if ideology itself plays no direct role in tie formation, or whether homophily on lifestyle or consumption patterns could indirectly induce political polarisation.

4.5 Simulation of Processes on the Network

Simulated networks generated by the model can serve as surrogates for real-world social structures in cases where empirical network data are unavailable or incomplete (Mäki 2009). In such scenarios, the simulated graphs provide a plausible basis for studying how beliefs or attitudes evolve through social influence processes embedded in homophilically structured environments. Moreover, these simulations can be used to explore counterfactuals or to evaluate interventions aimed at reducing polarisation, such as changes in network openness or the introduction of bridging agents, potentially in preparation for further theoretical or empirical analysis. The surrogate network approach thus offers a flexible framework for studying belief dynamics in a controlled yet socially plausible setting, grounded in empirically-informed network structures.

4.5.1 Basic Sensing as Social Sampling

The most basic set of processes to model on the network is social sampling (Galesic, Olsson, and Rieskamp 2012; Dasgupta, Kumar, and Sivakumar 2012), that is agents sensing others' attitudes. For example, they can situate their own attitude in their immediate social context of their ego network, compare their link-neighbours' attitudes to each other, or compute summary statistics of attitudes. Furthermore, simple inductive statistics, such as correlations between link neighbours' attitude dimensions, might be of interest. To assess the representativeness of the individual social samples, one can compare these local samples to the corresponding characteristics of the population as a whole.

Based on their sensing, agents can form beliefs about central tendencies in the population but also about how attitudes are distributed, and about properties of the network, which they live on. For the purpose of investigating polarisation in networks, individual beliefs about this polarisation might be the prime characteristic; similarly, access to their link-neighbours' node properties would allow agents to gauge, for example, density or proportion of closed triads. Such basic sensing would allow to construct metrics of individually perceived attitude distributions and perceived polarisation. Comparing these perceptions to the actual values allows identifying where social devices go unnoticed, or where people mistakenly believe that their society is polarised with respect to salient issues like climate-change or migration.

All the above can be taken only from the agents' own link-neighbours, but including neighbours of neighbours or mediate connections selected according to other criteria is equally possible. Likewise, agents can be made to retrieve characteristics other than attitudes, such as node properties or the aforementioned perceptions. The latter can be used to model meta-perceptions or iterative updating, as detailed in the following.

4.5.2 Dynamic Interaction

The iterative updating of beliefs about attitudes described above actually constitutes a first form of dynamic interaction: A simulation on the set network would run for multiple time steps and in each step, agents would formulate their new belief about attitudes, taking into account visible attitudes as well as a combination of one's own and others' beliefs from the previous step in a Bayesian updating process. Again, taking the ego-network, or including higher-degree neighbours, is the most straightforward way to define an agent's vision but other approaches are possible as well. For example, one could have each agent i select a random other j to retrieve information from, and each agent on the shortest path from j to i would also gain this information; such a mechanism would put more central nodes in an epistemically superior position.

One could also have agents utilise the information they receive over time to form more nuanced beliefs about the network topology and their own position within it. While the static version of this sensing described above would limit agents to extrapolation from their immediate context, learning about information passing through them or observing their neighbours' belief changes might also allow for insights into clusters that the assessing agent is not part of.

Probably the most obvious use-case for the attitudinal networks generated by the presented algorithm is to explore updating of the attitudes themselves, given initial polarisation. To that end, one could simply run Degroot (1974) on the network. Alternatively, one might use the network as a starting point for bounded confidence dynamics (Meylahn and Searle 2024; Hegselmann, Krause, et al. 2002; Deffuant et al. 2002): Agents in the Deffuant-Weisbuch Model could select communication partners only from link-neighbours; those in the Hegselmann-Krause Model could have the size and shape of their confidence interval determined by the attitudes in their ego-network. By applying such updating mechanisms to the simulated networks, one can examine the consequences of different homophily strengths or network topologies for belief convergence, polarisation, or fragmentation. For instance, higher homophily may not only increase attitudinal clustering in the initial network but also hinder cross-cutting influence, thus stabilising or even intensifying initial divides.

Obviously, one can also have the network adapt over time to changes in attitudes to showcase feedbacks between opinion and network polarisation. Agents can drop links with a likelihood inferred from inverse weights and if they dropped a link form a new one according to the algorithm. This way, it is also possible to model polarisation from scratch, by having agents start from a random network. Likewise, one may want to study how attitudes and the network co-evolve over time, following e.g. (Henry, Mitsche, and Prałat 2016).

Building on belief formation, one could also use the network to study individuals taking action in a polarised world. For example, to explain the fates of revolutions and street protests, the extant theoretical and simulation literature focuses mainly on individual decisions or attitude patterns (Kuran 1989; Klein and Marx 2017; Asgharpourmasouleh et al. 2019). Running such simulations on homophilic networks can help assess and highlight the possible role of polarisation in this context: When agents see that their peers share their attitudes or deem change attainable, they might feel motivated to take action, even if they are actually in the minority.

5. Exemplary Analysis: Migration Attitudes, Polarisation, and Perceptions in Germany

The purpose of this section is to illustrate potential use cases of the proposed generative algorithm. To do so, we generate polarised networks based on attitudes towards migration in Germany, and showcase to what extent they can explain stylised empirical facts on polarisation, and especially on differences between actual polarisation and what people believe it to be.[9]

9. The algorithm was implemented in a network model in Julia, which is available at Anonymised (2025).

5.1 Polarisation of Attitudes towards Migration in Germany and Its Perception

5.1.1 Relevant Stylised Empirical Facts

Attitudes toward migration have become highly polarised in many societies, which is to say opinions split into opposing camps with little middle ground (Alesina and Tabellini 2024). Polarisation on immigration is evident in anti-immigrant movements and contributes to the rise of far-right anti-immigration parties. Germany, with its recently emboldened right-wing *AfD* party or the xenophobic *PEGIDA* movement, following the large influx of refugees around 2015-16, is an example of this (Ueffing, Rowe, and Mulder 2015; Asgharpourmasouleh et al. 2019). Windzio (2025) empirically identifies these two distinct clusters in German opinion data, and shows that similarity ties in polarising attitudes are strong within each cluster but weak between clusters. Hence, pro-immigration Germans tend to resemble each other (intra-group homogeneity) and differ starkly from those who are anti-immigration (inter-group heterogeneity), forming two coherent network components.

According to the literature, communication plays a key role in polarisation. Coverage of the topic in traditional media (Schneider-Strawczynski and Valette 2025) and social media interactions (Nasuto and Rowe 2024; Esteve-Del-Valle 2022) drive attitude development; due to its emotional dimension, the topic is likely to produce stronger narratives and information received will be more uniform on an intrapersonal level, yet less uniform on an interpersonal one; this, in turn, causes a more polarised attitudinal landscape. Polarisation is also salient for many people, with perceptions of polarisation being empirically upwards biased (Robinson et al. 1995; Chambers, Baron, and Inman 2006), and this bias being stronger for people with more extreme attitudes of their own (Westfall et al. 2015). Relatedly, people perceive attitudes held by typical supporters of political camps to be more extreme than what they actually are (**Graham2012 VanBoven2012**). Moreover, it is this perception rather than actual polarisation being related to (a lack of) trust and voting propensity (Enders and Armaly 2019), or hostility between groups (Moore-Berg et al. 2020). Furthermore, people believe their own attitude to be in line with the social norm (Dippel, Hetzer, and Burger 2022).

Beyond the clustering of attitudes and the media dynamics described above, segregation and polarisation have behavioural consequences for how citizens process information. A more segregated attitudinal environment reduces exposure to discrepant views and lowers willingness to engage with them. As a result, agents display reduced open-mindedness and greater resistance to updating when encountering counter-attitudinal signals (Turner 2023; Wollebæk et al. 2019). In our terms, this manifests as narrower confidence bounds over others' positions as well as over parties' programmatic locations–particularly for out-group targets.

In sum, we derive four stylised facts on the perception of polarisation which we will investigate using the network model:

1. People tend to perceive their own attitude as moderate, i.e., close to the mean and median of the attitude landscape.
2. People tend to overestimate the level of polarisation, i.e., they perceive their own social context to be more different from randomly chosen other social contexts than is the case in comparison with society-wide estimates.
3. People tend to perceive attitudes as more extreme than they actually are.
4. Segregation and polarisation cause lower levels of open-mindedness; hence, confidence intervals are small towards both individuals and political parties.

5.1.2 Data and Variable Construction

To model polarisation on migration attitudes representatively for Germany, this study draws on data from the German General Social Survey (ALLBUS) by GESIS-Leibniz-Institut Für Sozialwissenschaften (2024). The study does not collect any network data for it samples representatively from the German populace. The 2021 wave, used in this analysis, employed a mixed-mode self-administered design (paper- or web-based) and covered a broad range of sociopolitical topics, including attitudes toward immigration and migration policy. For model calibration, we use all of these variables that use a Likert scale (and can thus be merged into a composite measure), shown in Table 1. Since questions measuring attitudes toward different migrant groups (mi05 to mi11) were only administered to respondents in a subsample, we only include this subsample in our analysis.

Table 1. Overview of ALLBUS 2021 variables used in the analysis, with English item translations and response scales.

Variable	Question (English Translation)	Scale / Range
mi05	Immigration of people fleeing war should be...	
mi06	Immigration of people persecuted for political reasons should be...	
mi07	Immigration of economic migrants should be...	1 = Fully permitted
mi08	Immigration of workers from Eastern EU countries should be...	2 = Restricted
mi09	Immigration of workers from other EU countries should be...	3 = Completely prohibited
mi10	Immigration of non-EU workers should be...	
mi11	Immigration of spouses and children of migrants should be...	
mp16inv	Refugees are a risk for the welfare state	7-Point Likert scale:
mp17inv	Refugees are a risk for public safety	1 = Strongly disagree to
mp18inv	Refugees are a risk for social cohesion	7 = Strongly agree
mp19inv	Refugees are a risk for the economy	[inverted for consistency]
pa01	In politics, where would you place yourself on a scale from left to right?	10-point Likert scale: 1 = Far left to 10 = Far right

We construct a composite migration scepticism index based on all migration-related variables in Table 1, i.e., all but *pa01*. This index is simply computed as the unweighted arithmetic mean of the recoded responses across all eleven items for each respondent (ignoring missing values). It is normalised, such that the population mean index is 1; values below 1 thus indicate above-average openness to migration, while values above 1 indicate below-average openness.

Figure 3 shows the distribution of the composite migration scepticism index, for each category on the self-description on a left-right scale (variable *pa01*). The latter will be used to create a benchmark for polarisation of migration attitudes. There are 3011 respondents in the subsample relevant for the present study who answered the political self-perception question, and the questions on migration issues.

To assess whether the visual differences across political self-placement categories are systematic, we conduct several complementary significance tests. The results are reported in Appendix 1. A linear trend regression of the composite migration scepticism index on left-right self-placement yields a positive and highly significant slope of 0.054 per scale point ($p < 0.001$, $R^2 = 0.193$). This is corroborated by a Spearman rank correlation of $\rho = 0.448$ ($p < 0.001$), indicating a clear monotonic association. Finally, both a one-way ANOVA ($F = 89.998$, $p < 0.001$) and a Kruskal-Wallis test ($H = 632.104$, $p < 0.001$) reject equality across the ten self-placement categories. Hence, political self-placement is systematically associated with migration scepticism, although the association is far from deterministic, which motivates using the continuous migration scepticism index as the network-forming attitude.

5.1.3 Perceived Polarisation

To assess the level of polarisation in the population, we calculate Ashman's D of the composite migration scepticism index. The measure is commonly used to assess the degree of separation between two overlapping distributions, and it has been proposed as a robust indicator of bimodality

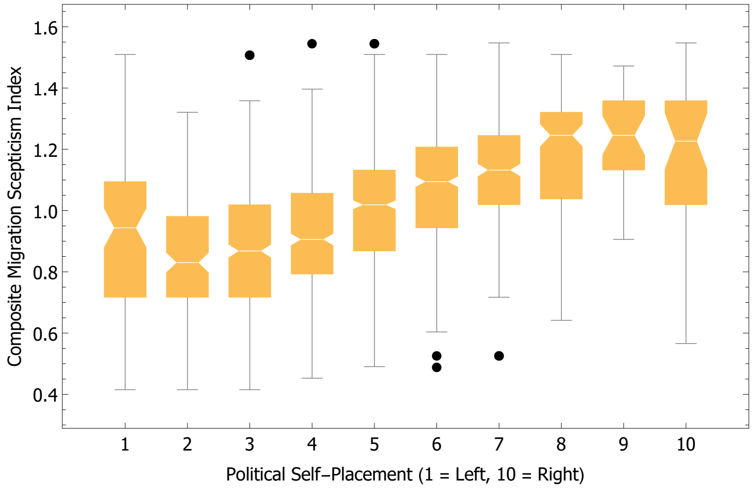

Figure 3. Distribution of migration scepticism per self-identified left-right position in the ALLBUS 2021 wave.

— a statistical property closely linked to polarisation (Tang et al. 2021), as outlined in Section 2: Bimodality captures several senses in which polarisation is typically conceptualised, including the fragmentation of a population into subgroups, the distinctness of those subgroups, and the distance between their respective positions (cf. above Section 2). Hence, we apply Ashman's D as a compact measure of the extent to which a single attitudinal dimension, such as migration scepticism, exhibits a polarised distribution. It computes the (dis)similarity of two groups as follows:

$$D = \frac{|\mu_1 - \mu_2|}{\sqrt{2 \cdot (\sigma_1^2 + \sigma_2^2)}}, \tag{3}$$

where μ_i refers to the mean score of a group i and σ_i to its standard deviation. The intuition behind this index is straightforward: Measured polarisation increases, whenever the absolute distance in the average attitude of the two groups in question increases ($|\mu_1 - \mu_2|$) or whenever the two groups get more homogeneous (i.e., the dispersion measured by σ_1 or σ_2 decreases). In our case, we derive the groups from the self-determined left-right score. This is a salient dimension of political belonging, and Windzio (2025) finds an association with attitudes towards migration in Germany, making overall political views a suitable category for reference groups. Since it is a matter of interpretation, where along this 10-point scale to split between the right-wing and left-wing cluster, we compute Ashman's D for all possible splits, and Table 2 shows the results.

5.2 Model

The descriptions for this example present a shortened overview of which specific values for attitudes, ω and ρ are used. The reader should plug them into the general framework introduced in Section 3.

5.2.1 Calibration and Network Setup

We initialise the model based on the combined migration (policy) attitude score and the left-right self-description score from the ALLBUS study 2021, as described in Section 5.1.2. We initialise each respondent with their composite migration scepticism index and political self-description on a left-right scale. This results in a population of N = 3011 nodes, based on the representative sample of Germany.

Table 2. Ashman's D benchmark values for different political self-placement splits based on pa01, measuring separation of migration scepticism index distributions.

Left Cluster (pa01)	Left Size	Right Cluster (pa01)	Right Size	Ashman's D
1	85	2–10	2926	0.0805
1–2	245	3–10	2766	0.1544
1–3	716	4–10	2295	0.1847
1–4	1168	5–10	1843	0.1973
1–5	1968	6–10	1043	0.1888
1–6	2512	7–10	499	0.2071
1–7	2798	8–10	213	0.2360
1–8	2946	9–10	65	0.2373
1–9	2977	10	34	0.1871

Despite the left-right self-description being an agent attitude, this is merely for evaluation/validation purposes. Agents perform their homophilic linkage based solely on the composite migration scepticism index. Technically, this can either be understood as two attitude dimensions $\mathbf{A} = migindex, pa01$ with $\omega_{pa01} = 0$, or as the migration scepticism index being a single attitude of interest.

The network is simulated for various homophily levels, and we use select levels to showcase behavioural patterns: $\rho = 0$ (random network) for a benchmark without endogenous polarisation, $\rho = 4$ for mild homophily, $\rho = 8$ for moderate homophily, and $\rho = 14$ for strong homophily. This parameter range is informed by a previous application in the extant literature studying inequality perceptions. In their application, Schulz, Mayerhoffer, and Gebhard (2022) report that values of ρ between 4 and 14 were needed for the model to generate misperceptions comparable to the empirical benchmark, with $\rho = 8$ corresponding to the average level of perceived inequality. To account for random variations in network generation, we simulate 100 Monte Carlo runs for each ρ value. Since the distribution of attitudes stems directly from the empirical input study, it is identical across all 100 Monte Carlo runs per ρ, and across all homophily levels. The figures in Section 5.3 generally pool observations from all 100 Monte Carlo runs, such that each respondent from the survey appears 100 times, each with different network calibrations. The exemplary network graphs in Figure 4 that are based on a single random run for illustrative purposes constitute the only exception from this.

5.2.2 Agents' Observations

Once the network is generated, agents make their observations. They can sense all attitudes in their own ego-network (including their own attitude). With this, they estimate their own migration attitude rank relative to their link neighbours as well as the dispersion of migration attitudes, operationalised as standard deviation in their ego-network. Furthermore, agents can sense the ego-network of one randomly selected other agent. We assume random choice of alters as a lower-bound benchmark for perceived polarisation. In empirical reality, this choice will likely not be random, as e.g. social media algorithms are typically deliberately designed to expose users to the "other side" to increase engagement Munn (2020). Comparing this other group to their own, they calculate Ashman's D.

Note that there is no bias in agents' observations. Firstly, agents sense the true attitudes of their link-neighbours and the other randomly chosen ego-network; this could be understood to be the result of clear and unambiguous communication of attitudes. Secondly, agents correctly calculate the relevant distributional properties, i.e., they are not constrained by biases in information processing. Both these assumptions are likely violated in actual debates. However, the present study deliberately idealises here because its goal is to explore possible implications of structurally selective information availability in a polarised environment. Namely, agents take their own lived experience in their immediate social context – modelled via ego-networks – to be (representative) samples that allow inferences on the whole population.

5.3 Results

5.3.1 Polarisation of the Networks

We first discuss the resulting network topology of our generating mechanism. Under the parameter settings examined here, even mild levels of homophily cause high segregation: Figure 4 illustrates that in the network for $\rho = 4$, connections between agents at the lower and upper end of the attitude spectrum are rare, notably notably distinguishing its community structure from the random network's. Yet, ties with others who have a somewhat different attitude are still fairly common. This results from the exponential form of the weighting function in Equation 1, which makes linkage to others with far distant attitudes especially unlikely, keeping the overall relative weights of those with moderately different attitudes constant. In the public forum, such connections from both extremes to moderates might create intermediaries who function as bridges between the extreme clusters. Conversely, for higher values of the homophily parameter ρ, agents keep more strictly to their like-minded peers, and these bridges largely disappear, making communication between extremes harder or requiring more intermediaries, as evident in the increased average shortest path length. We discuss these findings in more detail below.

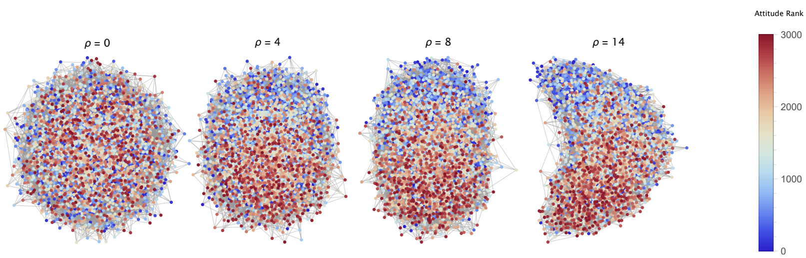

Figure 4. Exemplary graphs for different homophily levels (all taken from the first Monte Carlo Run with random seed 1). Visualisations based on Mathematica 13's "Spring Electrical Embedding" routine. Polarisation increases in ρ, as illustrated by communities growing more homogenous in attitude rank (visualised by colour).

5.3.2 Degree distributions and the relationship between attitudes and degree

Across all values of ρ, the simulated networks depicted in Figure 5 display a highly homogeneous degree distribution. For $\rho = 0$, where attitudes do not influence the probability of forming ties, the degree distribution is well approximated by a truncated Gaussian with small variance. This shape arises directly from the network generation procedure: each node deterministically selects at least five neighbours, which imposes a positive lower bound on degrees, and additional connections are formed when other nodes select it. Because these incoming ties represent the sum of many small stochastic contributions, their distribution is approximately Gaussian by the central limit theorem. The combination of a deterministic minimum and Gaussian noise naturally yields the observed truncated Gaussian degree distribution. Degrees thus have a characteristic scale and the network should consequently be interpreted as a network of close social contacts, as discussed in MacCarron,

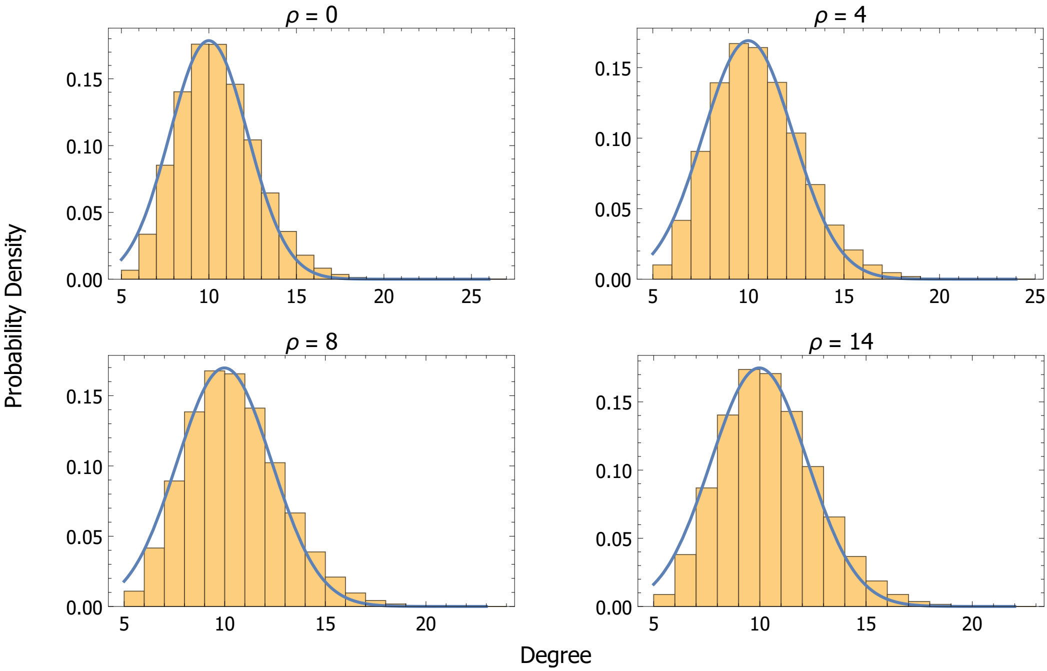

Figure 5. Degree distribution for varying homophily levels $\rho \in \{0, 4, 8, 14\}$ with a superimposed Gaussian fit. In all cases, degree distributions of the simulated networks are well approximated by a Gaussian density truncated at a minimum of 5.

Kaski, and Dunbar (2016) and McPherson, Smith-Lovin, and Brashears (2006), rather than of weak ties on social networks that typically exhibit fat-tailed degree distributions (see e.g. Aparicio, Villazón-Terrazas, and Álvarez 2015, for one of many examples).

To verify that this is indeed the generating mechanism for the truncated Gaussian degree distribution in the model, Figure 6 plots the binned degrees for 10 intervals of the attitude distribution against the midpoints of these attitude bins. As expected, degree is independent of attitudes at $\rho = 0$, when homophily is inactive. For $\rho > 0$, degrees become weakly associated with attitudes, but this systematic effect is not large enough to disrupt the overall Gaussian shape of the degree distribution. Appendix 4 formalises this point by decomposing expected degree into the fixed own-choice component, the expected number of incoming selections, and a correction for reciprocal selections. The derivation shows that the attitude distribution enters through the incoming-selection component: nodes in denser parts of the attitude distribution tend to receive more selections from others, while the fixed-choice component keeps the aggregate degree distribution concentrated around a characteristic scale. The relationship for $\rho > 0$ exhibits an inverted U-shape pattern, reflecting the fact that most of the attitude density is close to moderate positions. Thus, agents holding these moderate positions have a wider choice of others with similar or even identical attitudes to link to, and selections as link-neighbours are less likely to be mutual, implying a higher degree for these agents.[10]

10. Appendix 2 and Appendix 3 test whether this result is specific to the empirical ALLBUS attitude distribution. It is not. We repeat the analysis for artificial attitude distributions with substantially different shapes, including high-variance lognormal distributions and strongly U-shaped, bimodal Beta distributions such as Beta(0.5, 0.5). Across these cases, the aggregate degree distribution remains well approximated by a Gaussian density truncated at the minimum degree, while the conditional relationship between attitude position and degree changes with the shape of the input distribution. This supports the interpretation that the truncated-Gaussian degree distribution is a robust implication of the fixed-choice homophilic generator, whereas the attitude distribution primarily governs where, in attitude space, degree is somewhat higher or lower.

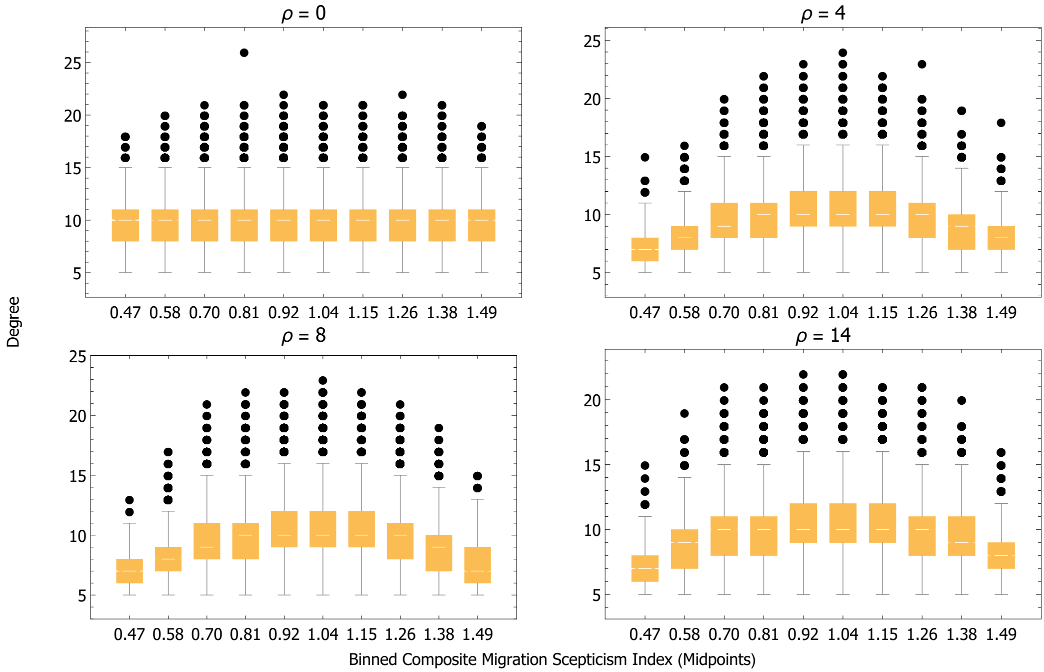

Figure 6. Binned degrees for 10 equally spaced intervals of the attitude distribution against the midpoints of these intervals. Degree is independent of attitudes for $\rho = 0$ and exhibits an inverted U-shape relationship for $\rho > 0$.

5.3.3 Small-world structure

Following established practice in the analysis of small-world networks (e.g. Humphries and Gurney 2008), we benchmark clustering against an Erdős–Rényi random graph $G(n, p)$ with the same number of nodes n and the same mean degree $k = p \cdot n$. For this benchmark, the expected clustering coefficient is

$$C_{\mathrm{ER}} = \frac{k}{n}. \tag{4}$$

This expression implicitly assumes a Poisson-like degree distribution, where the variance of degree is approximately equal to the mean. In such networks, a fraction of nodes will, purely by chance, attain substantially higher degrees than average; these high-degree nodes generate many connected triples and thus contribute disproportionately to triangle formation.

As Figure 5 shows, our simulated networks at $\rho = 0$ deviate from this assumption. The degree variance is much smaller than the mean, and neither very high-degree nor very low-degree nodes are present. Consequently, the number of connected triples is smaller than in an Erdős–Rényi graph with identical (n, k), and the observed clustering lies below C_{ER}.

We quantify the deviation from the Erdős–Rényi expectation using the normalized clustering coefficient

$$\Gamma = \frac{C}{C_{\mathrm{ER}}}. \tag{5}$$

In contrast to clustering, the average path length is not highly sensitive to the variance of the degree distribution (Chung and Lu 2002). Following Schulz, Mayerhoffer, and Gebhard (2022), we use the approximation

$$L_{ER} = \frac{\ln(n) - \gamma}{\ln(k)} + \frac{1}{2}, \qquad \gamma = 0.577\ldots \tag{6}$$

for the expected average path length of an Erdős–Rényi graph with n nodes and mean degree k.

Because the simulated networks at $\rho = 0$ have strongly homogeneous degrees and exhibit no community structure, their average path length closely matches L_{ER}, even though their clustering is lower than C_{ER}. We define the normalized path length as

$$\Lambda = \frac{L}{L_{ER}}. \tag{7}$$

Following Humphries and Gurney (2008), we summarise small-world structure using the small-world index

$$\Phi = \frac{\Gamma}{\Lambda}. \tag{8}$$

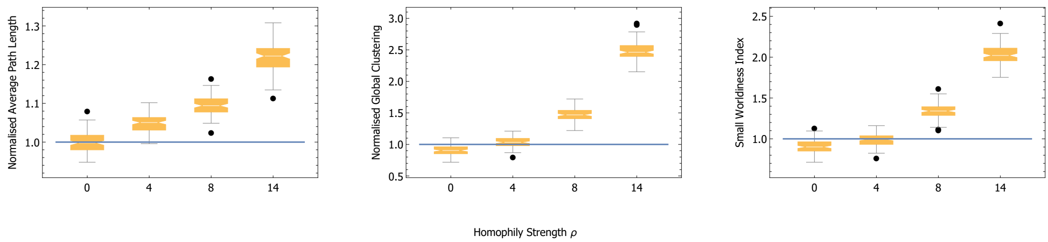

Figure 7. Box-whisker charts for three statistics determining small-worldiness, i.e., Λ, Γ and Φ. Vertical axes display values against the theoretical Erdős–Rényi benchmarks. Our simulated network shows weak small-worldiness, as increasing homophily strength increases average path lengths slightly above the benchmark but clustering increases much faster in ρ.

Figure 7 summarises these statistics for varying degrees of homophily. As expected, the average path length at $\rho = 0$ corresponds closely to the Erdős–Rényi benchmark but increases slightly for increasing ρ. For clustering, at $\rho = 0$, the small degree variance reduces clustering relative to C_{ER}, resulting in $\Gamma < 1$ and low small-worldiness. For $\rho \geq 4$, clustering increases while path lengths remain close to the Erdős–Rényi expectation, i.e. $\Lambda \approx 1$, producing $\Gamma > 1$ and indicating weak to moderate small-world structure for higher values of ρ.

In sum, this indicates that the resulting networks mimic one important feature of real-world social networks even for very homogeneous attitude distributions, at least approximately: small-worldiness, i.e., low average path lengths coupled with strong clustering (Watts and Strogatz 1998). This resemblance provides some face-validity support for using the generated networks in a how-possibly explanation of the stylised facts for perceptions of polarisation, as discussed below.

5.3.4 Perception of Own Attitude Position

The first stylised fact on perceptions identified in Section 5.1.1 concerns the perceived own relative attitude. To this end, agents simply sort all attitudes in their ego-network and then check their own rank.

For the random network, this results in self-perceptions close to the true values, as shown in Figure 8.[11] However, even for mild homophily, the simulations display a marked tendency towards the median, producing a pattern consistent with the first stylised fact.

11. The differences from the theoretical expectation of each decile containing 10% of the population are not fully matched due to the non-continuous nature of the constructed composite migration scepticism index, causing some agents with the same score to be put in different deciles.

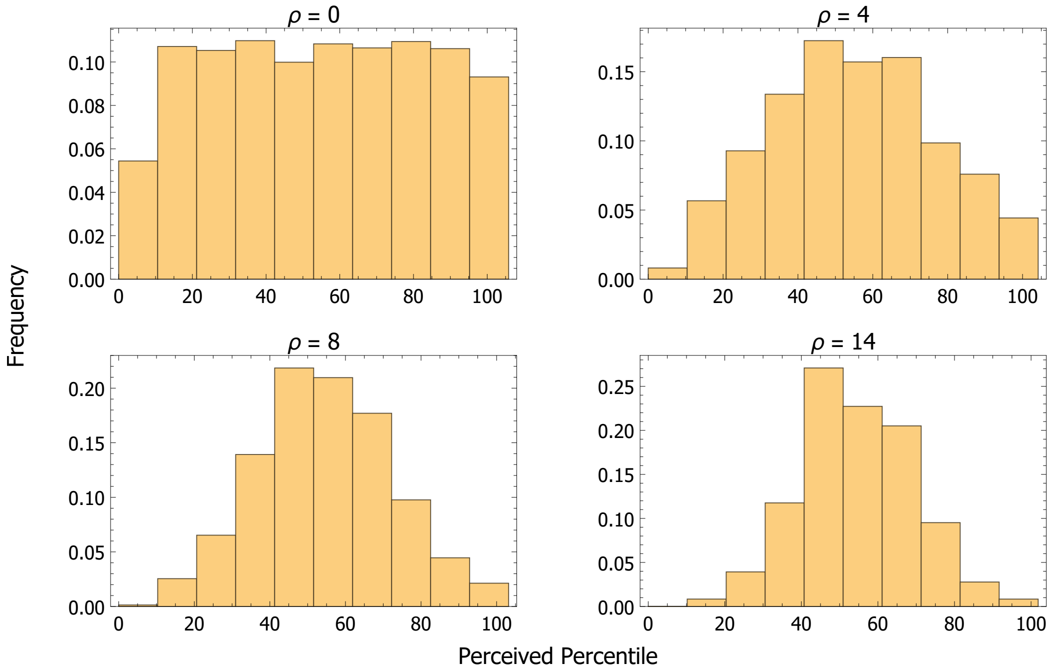

Figure 8. Perceptions of own attitude position (decile), compared to one's ego network for different homophily levels. (Trivially, accurate perception referencing the whole population would be uniformly distributed.)

5.3.5 Perceived Polarisation

The second stylised fact concerns perceptions of polarisation. We assume agents to perceive polarisation by comparing their own ego-network to the one of another randomly selected agent. This can be understood as people gaining insights from incidental interactions with others outside of their peer group. One might, for instance, think of such encounters as meeting strangers on the train or engaging another bubble in a (possibly controversial) social media debate. Indeed, especially for the case of social media, we treat random choice of alters as a possible lower-bound benchmark for perceived polarisation, as discussed above. Based on the two groups, agents compute Ashman's D with the formula introduced in Equation 3. Hence, like for the self-perceptions above, the sensing and processing are unbiased, but the sample which agents use is biased.

Figure 9 shows the individual perceptions, comparing them to the left-right benchmark values from Table 2. While people almost underestimate polarisation in the case of a random network, homophily causes agents in the simulation to perceive polarisation to be (sometimes drastically) higher than it actually is, in line with the stylised empirical fact.

The perception patterns follow a U-shape, meaning that those with attitudes further from the mean not only contribute more to the actual polarisation but also to its overestimation. Within the model, the mechanism driving this shape is the combination of homophily in the generation of ego-networks with randomly selected ego-networks as a comparison group. Hence, the difference between ego and alter tends to be larger for those with extreme attitudes. For similar reasons, the dispersion of perceptions in the population as a whole and for a given attitude rises in the homophily level: For higher ρ, ego-networks grow more homogenous, making it more influential which other group is selected for comparison.

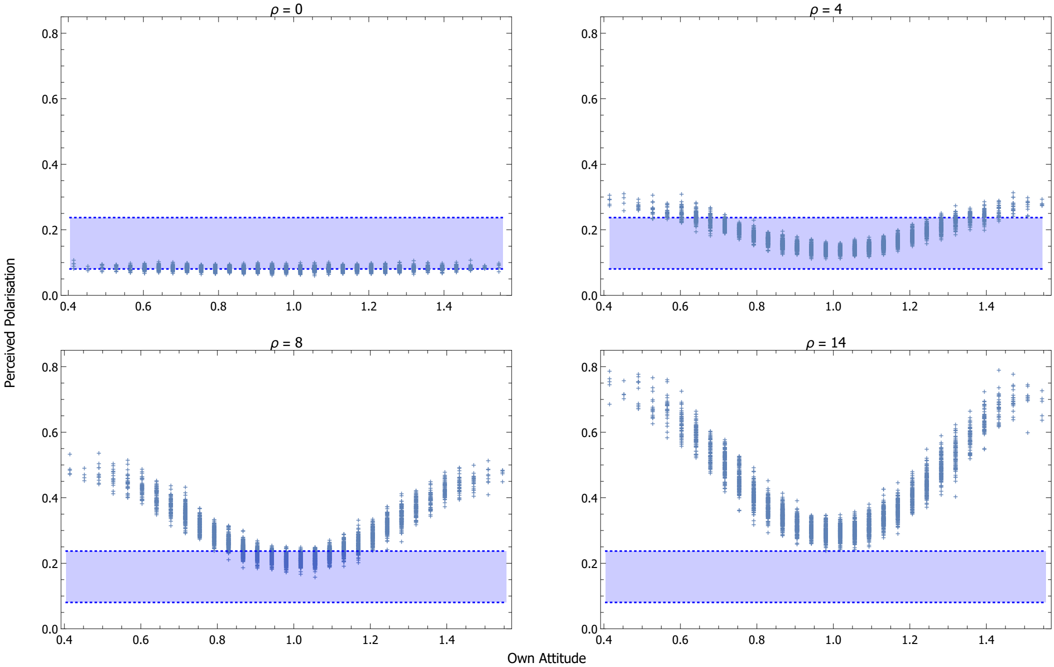

Figure 9. Perceived polarisation (Ashman's D) for different homophily levels. The marked blue area shows the range of objective polarisation, based on the left-right scale.

5.3.6 *Perceived Extreme Attitudes and Homogeneity of Ego-Networks*

The third stylised fact to be investigated here concerns the perception of out-group members' attitudes to be more extreme than they actually are. Such perception depends on the quality of information that one might have about these attitudes as well as the threshold for identifying an attitude as extreme. Trivially, agents in the model are less likely to acquire knowledge about people with less similar attitudes because of the homophilic linkage. Yet, since the network does not model all interaction channels, the actual quality of information that people get depends on what these other interaction channels look like. If random encounters, as simulated for the prior stylised fact, are sufficiently frequent, the available information can become almost accurate.

To determine a threshold to distinguish extreme from non-extreme settings, agents could refer to some fixed number, e.g., below 0.5 and above 1.5 on the composite migration scepticism index. Such fixed threshold would mean that one considers certain (expressions of) attitudes extreme, regardless of the discoursive context. Namely, it could happen that everyone in a society holds a view considered to be extreme in the same direction. Alternatively, agents might set the threshold as relative to what they deem normal in the population, e.g., below or above one standard deviation of the the composite migration scepticism index. For such relative setting, they could either rely on the true values in the population (as, for example, communicated through representative polls), or on their ego-network. Figure 10 displays the standard deviation of ego-networks. As expected, the ego-networks in homophilic networks show less plurality than the population as a whole. Thus, the simulations are consistent with the stylised fact if – but only if – agents base their threshold for considering attitudes extreme on observations in their ego-network.

A related implication concerns the idea that polarisation inhibits openness and keeps one from taking attitudes that differ from one's own seriously. To test this fact, we draw on bounded confidence approaches and take individuals to establish confidence intervals around their own attitude:

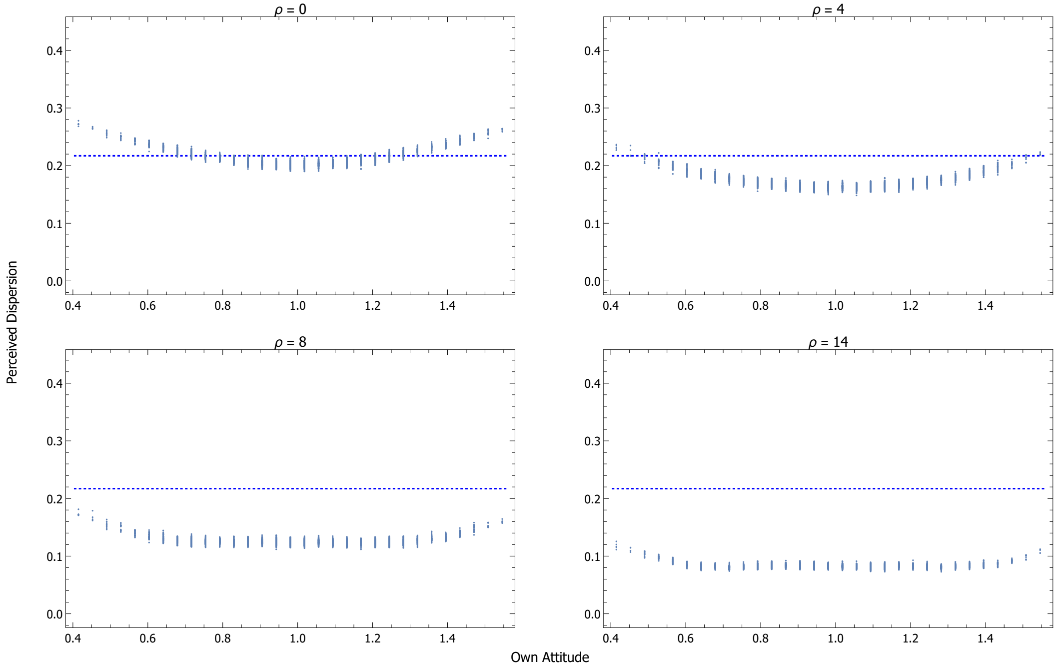

Figure 10. Perceived dispersion in attitudes compared to actual dispersion, measured by the standard deviation of attitudes. Population standard deviation as reference.

people/attitudes within this confidence interval are taken seriously by the agent in question, the ones outside it are not. To set the size of their confidence interval, an agent could refer to what the peers in their ego-network believe and assume this to be the range of normal/acceptable beliefs in society. The population standard deviation corresponds to a confidence interval within which social discourse can reach consensus — for example, in the Hegselmann-Krause bounded confidence model. However, if agents instead were to determine their confidence interval based on their ego-network, communication dynamics would solidify polarisation (Hegselmann, Krause, et al. 2002). To get a more detailed picture, future studies could run the actual confidence dynamics interaction on the different graphs; yet even the static results presented here suggest one possible pathway through which homophily could narrow exposure to differing attitudes, inhibit fruitful debate and hence possibly contribute to polarisation of networks as well as attitudes.

5.4 *Summary of the Case Study*

The exemplary application investigated perceptions of polarisation in homophilically structured networks. Holding the empirically observed distribution of migration attitudes in Germany fixed, the simulations generate patterns consistent with the empirical observations that people tend to perceive their own attitudes as moderate, and overestimate the level of societal polarisation. This positions exponential homophily in attitudes as a possible mechanism driving these stylised facts. The tendency to perceive others' attitudes as more extreme than they actually are requires additional assumptions on how these others are perceived, and on how thresholds for identifying attitudes as extreme are set. The model therefore identifies possible structural pathways through which homophilic linkage could produce systematically unrepresentative perceptions, while also identifying conditions under which this mechanism alone is insufficient.

For the mechanism to work as a how-possibly explanation, the synthetic network structures need to resemble ones from the real world, which is to be evaluated in future empirical research. However, this does not mean that real world network formation has to explicitly feature homophily in attitudes. Rather, there could also be homophily in attributes positively associated with similarity in migration attitudes, or even heterophily in ones negatively associated with it. Likely, the mechanism is much weaker in reality than in our simulations, and contributes to the stylised facts together with other factors.

Finally, the foci on perceptions and on network generation mean that several other aspects of polarisation are not modelled here. Thus, various stylised facts are not explained in the model. For example, the literature identifies traditional media, social-media interaction, emotional framing, and political communication as possible contributors to polarisation. In order to explore them, one would need to the production, selection, or diffusion of media content, and more generally processes of communication, on the network. Because attitudes remain fixed, the developed model does not examine whether communication among close contacts produces attitude convergence, affective polarisation, or political extremism. Likewise, the application does not model any affective or behavioural consequences of perceived polarisation. To account propose a mechanism governing these phenomena, one may only use the present synthetically generated networks as a starting point.

6. Discussion

This paper presented a generative algorithm for simulating network polarisation based on attitudinal homophily applicable to both undirected and directed link formation. In essence, it applies the notion of preferential attachment to node properties other than degree, aiding intuitive communication within and beyond the network science community. This proposed algorithm can work with single or multiple attitude dimensions with flexible weights; it allows for heterogeneity in the simulated population regarding these weights as well as the homophily strength, and for the specification of a minimum degree. In the simulations examined, the generated networks display selected properties also associated with Random Geometric Graphs. Under the parameter settings considered, they can also exhibit features found in some real-world social networks, such as (weak) small-worldiness with high clustering and local cohesion. The algorithm's output networks can be used to explore the model-implied structural consequences of homophily, to generate how-possibly explanations of stylised empirical facts about polarisation in attitudes, and to simulate social influence processes.

The exemplary application of the algorithm to migration attitudes and polarised networks in Germany illustrates how stylised empirical facts can be compared with outcomes generated by a candidate mechanism. Simulations based on real-world attitude distributions produce patterns that resemble empirical observatins of polarisation and perception bias: model agents in homophilic networks overestimate polarisation, misperceive their own attitudinal position, and experience reduced attitudinal diversity in their ego-networks. Hence, individual perceptions of societal attitudes may arise without biased cognition solely from the structure of social interaction.[12] If future empirical work supports the relevance of the proposed mechanism, the resulting synthetic network could help to identify promising targets for educational interventions with potentially significant spill-over effects.

Generally, one of the key contributions of the algorithm is to establish a how-possibly result: Simple homophilic preferences in link formation can generate segregated network structures which, in turn, yield biased individual samples and thus misperceptions. Hence, under certain conditions present in the simulation model, complex cognitive, psychological, or ideological assumptions are neither necessary for polarised networks to emerge, nor for biased individual assumptions about attitudes and polarisation. Therefore, the algorithm can help efforts to complement the analytical focus on individual traits by one on the structure in which attitudes form, circulate, and solidify.

12. We do not claim that empirical misperceptions arise solely from network structure. Indeed, studying how cognitive biases operate on the homophilic network might be a helpful next step.

The presented straightforward algorithm offers a way for capturing homophily in attitudes, has specific features. These may constitute an asset when studying some aspects of polarisation phenomena and/or can complement other mechanisms well, while, in the context of other phenomena, they are a limitation and/or require it to be complemented by other mechanisms. First, it captures homophily in attitudes, not any other node properties like group membership or link-neighbours. Likewise, homophily is isolated from other tie-formation mechanisms. Yet, in reality, people's link formation is likely captured by an aggregate function, of which attitude homophily is only one component. Second and relatedly, the model primarily generates network segregation, not the full phenomenon of polarisation. The algorithm's simplicity makes application versatile: it holds other mechanisms fixed in order to isolate the effect of homophily, and it offers a clean baseline for extensions, in terms of dynamic rewiring, endogenously shifting attitudes, or more nuanced sensing and cognition. However, one has to add these features ex post, as they are not included in the proposed algorithm. Third, the homophily mechanism uses a fixed exponential decay function. To operationalise attitudinal distance for this, one must assume the attitude(s) in question to be on meaningful cardinal (and comparable) scale(s). That appears to be common practice but can nonetheless be a sword of Damocles hanging above any findings drawn from the generated networks. Moreover, while the functional form is founded in discrete choice, it has very specific implications; namely, I) every agent can potentially select every other agent as link neighbour, but II) as attitudes become dissimilar, linkage probability declines rather steeply. There are various plausible alternatives (some of which have been explored in the literature, as outlined in Section 2) such as a sharp similarity threshold beyond which there is no linkage possible, a tolerance plateau within which differences matter little, or a slower power-law decline of link probability. Fourth, multidimensional formulation uses a weighted sum of absolute differences inside the exponential. Hence, each attribute contributes independently and additively to total distance. This excludes, for example, multiplicative effects, or any on the above-mentioned alternatives to the exponential linkage operating on one attitude dimension (e.g., there is a sharp threshold for one dimension prohibiting connection even for high similarity in all other dimensions, but if the threshold is met, this first dimension does not matter much). Besides this conceptual feature, cross-dimensional segregation may partly reflect correlations already present in the input data. Lastly, the parsimonious baseline version of the algorithm presented here has some shortcomings that one can straightforwardly counteract with minor adjustments: I) The network is represented as a single layer of largely equivalent ties, even in the multi-attitude case. As an alternative, one might want to create a multi-layer network with separate layers for each attitude. II) The fixed number of selected neighbours shapes the resulting network, but one might, of course, vary that number either randomly, based on attitudes or other known (or presumed) node attributes. III) Raw attributes do not matter for the algorithm. However, the functional form is straightforwardly applicable to these (by dropping references to self in Equations 1 and 2) and for example capture tendencies to listen to populists with more extreme attitudes. In sum, the generative algorithm discussed here is one way of synthesising polarised networks based on homophily in attitudes, and may especially be helpful when combined with other mechanisms.

References

Abrams, Dominic, Margaret Wetherell, Sandra Cochrane, Michael A. Hogg, and John C. Turner. 1990. Knowing what to think by knowing who you are: Self-categorization and the nature of norm formation, conformity and group polarization [in en]. _Eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.2044-8309.1990.tb00892.x, British Journal of Social Psychology 29 (2): 97–119. ISSN: 2044-8309, accessed April 25, 2025. https://doi.org/10.1111/j.2044-8309.1990.tb00892.x. https://onlinelibrary.wiley.com/doi/abs/10.1111/j.2044-8309.1990.tb00892.x.

Albert, Réka, and Albert-László Barabási. 2002. Statistical mechanics of complex networks. Reviews of Modern Physics 74, no. 1 (January): 47–97. ISSN: 1539-0756. https://doi.org/10.1103/revmodphys.74.47. http://dx.doi.org/10.1103/RevModPhys.74.47.

Alesina, Alberto, and Marco Tabellini. 2024. The political effects of immigration: culture or economics? Journal of Economic Literature 62, no. 1 (March): 5–46. ISSN: 0022-0515. https://doi.org/10.1257/jel.20221643. http://dx.doi.org/10.1257/jel.20221643.

Anderson, Elizabeth. 2021. Epistemic Bubbles and Authoritarian Politics [in en]. In Political Epistemology, edited by Elizabeth Edenberg and Michael Hannon, 11–30. Oxford: Oxford University Press. ISBN: 978-0-19-289333-8. https://doi.org/10.1093/oso/9780192893338.003.0002.

Anonymised. 2025. Homophilic linkage (model code). https://anonymous.4open.science/r/Homophilic-Linkage-041225/. V. 1.03. https://anonymous.4open.science/r/Homophilic-Linkage-041225/.

Aparicio, Sofía, Javier Villazón-Terrazas, and Gonzalo Álvarez. 2015. A model for scale-free networks: application to Twitter. Entropy 17 (8): 5848–5867.

Asgharpourmasouleh, Ahmadreza, Masoud Fattahzadeh, Daniel Mayerhoffer, and Jan Lorenz. 2019. On the fate of protests: dynamics of social activation and topic selection online and in the streets. In Computational conflict research, 141–164. Springer International Publishing, November. ISBN: 9783030293338. https://doi.org/10.1007/978-3-030-29333-8_7. http://dx.doi.org/10.1007/978-3-030-29333-8_7.

Bakker, Bert N., and Yphtach Lelkes. 2024. Putting the affect into affective polarisation. Publisher: Routledge _eprint: https://doi.org/10.1080/02699931.2024.2362366, Cognition and Emotion 38, no. 4 (May): 418–436. ISSN: 0269-9931, accessed April 23, 2025. https://doi.org/10.1080/02699931.2024.2362366. https://doi.org/10.1080/02699931.2024.2362366.

Baron, Robert S., Sieg I. Hoppe, Chuan Feng Kao, Bethany Brunsman, Barbara Linneweh, and Diane Rogers. 1996. Social Corroboration and Opinion Extremity. Journal of Experimental Social Psychology 32, no. 6 (November): 537–560. ISSN: 0022-1031, accessed April 25, 2025. https://doi.org/10.1006/jesp.1996.0024. https://www.sciencedirect.com/science/article/pii/S0022103196900244.

Baumann, Fabian, Philipp Lorenz-Spreen, Igor M. Sokolov, and Michele Starnini. 2020. Modeling echo chambers and polarization dynamics in social networks. Physical Review Letters 124 (4): 048301. https://doi.org/10.1103/PhysRevLett.124.048301.

Begby, Endre. 2022. From Belief Polarization to Echo Chambers: A Rationalizing Account [in en]. Publisher: Cambridge University Press, Episteme, 1–21. ISSN: 1742-3600, 1750-0117, accessed July 27, 2022. https://doi.org/10.1017/epi.2022.14. https://www.cambridge.org/core/journals/episteme/article/from-belief-polarization-to-echo-chambers-a-rationalizing-account/FCC9AB2691CD160F53BB904C128DD81D.

Billig, Michael, and Henri Tajfel. 1973. Social categorization and similarity in intergroup behaviour [in en]. _Eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/ejsp.2420030103, European Journal of Social Psychology 3 (1): 27–52. ISSN: 1099-0992, accessed April 27, 2025. https://doi.org/10.1002/ejsp.2420030103. https://onlinelibrary.wiley.com/doi/abs/10.1002/ejsp.2420030103.

Bonica, Adam, Adam S. Chilton, and Maya Sen. 2016. The Political Ideologies of American Lawyers. Journal of Legal Analysis 8, no. 2 (December): 277–335. ISSN: 2161-7201, accessed April 27, 2025. https://doi.org/10.1093/jla/lav011. https://doi.org/10.1093/jla/lav011.

Bothorel, Cecile, Juan David Cruz, Mattei Magnani, and Barbora Micenková. 2015. Clustering attributed graphs: models, measures and methods. Network Science 3, no. 3 (March): 408–444. ISSN: 2050-1250. https://doi.org/10.1017/nws.2015.9. http://dx.doi.org/10.1017/nws.2015.9.

Bruns, Axel. 2019. Are Filter Bubbles Real? [In en]. Google-Books-ID: rcOrDwAAQBAJ. Cambridge: Polity, August. ISBN: 978-1-5095-3646-7.

Burnstein, Eugene, and Amiram Vinokur. 1977. Persuasive argumentation and social comparison as determinants of attitude polarization. Journal of Experimental Social Psychology 13, no. 4 (July): 315–332. ISSN: 0022-1031, accessed April 25, 2025. https://doi.org/10.1016/0022-1031(77)90002-6. https://www.sciencedirect.com/science/article/pii/0022103177900026.

Chambers, John R., Robert S. Baron, and Mary L. Inman. 2006. Misperceptions in intergroup conflict: disagreeing about what we disagree about. Psychological Science 17, no. 1 (January): 38–45.

Chung, Fan, and Linyuan Lu. 2002. The average distances in random graphs with given expected degrees. Proceedings of the National Academ 99 (25): 15879–15882.

Dall, Jesper, and Michael Christensen. 2002. Random geometric graphs. Physical Review E 66, no. 1 (July). ISSN: 1095-3787. https://doi.org/10.1103/physreve.66.016121. http://dx.doi.org/10.1103/PhysRevE.66.016121.

Dasgupta, Anirban, Ravi Kumar, and D. Sivakumar. 2012. Social sampling. In Proceedings of the 18th acm sigkdd international conference 235–243. KDD '12. New York, NY, USA: Association for Computing Machinery. ISBN: 9781450314626. https://doi.org/10.1145/2339530.2339572.

Deffuant, Guillaume, Frédéric Amblard, Gérard Weisbuch, and Thierry Faure. 2002. How can extremism prevail? a study based on the relative agreement interaction model. Journal of Artificial Societies and Social Simulation 5 (4).

Degroot, Morris H. 1974. Reaching a consensus. Journal of the American Statistical Association 69, no. 345 (March): 118–121. ISSN: 1537-274X. https://doi.org/10.1080/01621459.1974.10480137. http://dx.doi.org/10.1080/01621459.1974.10480137.

Dippel, Alina S., Lukas Hetzer, and Axel M. Burger. 2022. Links oder rechts? die ideologische selbstverortung von wähler:innen und ihre wahrnehmung von parteien in deutschland [in de]. easy social sciences, https://doi.org/10.15464/EASY.2022.04. https://www.ssoar.info/ssoar/handle/document/79901.

Enders, Adam M, and Miles T Armaly. 2019. The differential effects of actual and perceived polarization. Political Behavior 41:815–839.

Epstein, Joshua M. 1999. Agent-based computational models and generative social science. Complexity 4 (5): 41–60.

Esteve-Del-Valle, Marc. 2022. Homophily and polarization in twitter political networks: a cross-country analysis. Media and Communicatio 10, no. 2 (April). ISSN: 2183-2439. https://doi.org/10.17645/mac.v10i2.4948. http://dx.doi.org/10.17645/mac.v10i2.4948.

Finkel, Eli J., Christopher A. Bail, Mina Cikara, Peter H. Ditto, Shanto Iyengar, Samara Klar, Lilliana Mason, et al. 2020. Political sectarianism in America. MAG ID: 3095417376, Science 370 (6516): 533–536. https://doi.org/10.1126/science.abe1715.

Galesic, Mirta, Henrik Olsson, and Jörg Rieskamp. 2012. Social sampling explains apparent biases in judgments of social environments. Psychological Science 23, no. 12 (October): 1515–1523. ISSN: 1467-9280. https://doi.org/10.1177/0956797612445313. http://dx.doi.org/10.1177/0956797612445313.

GESIS-Leibniz-Institut Für Sozialwissenschaften. 2024. Allgemeine bevölkerungsumfrage der sozialwissenschaften allbus - kumulation 19 [in de]. https://doi.org/10.4232/1.14333. https://search.gesis.org/research_data/ZA5284?doi=10.4232/1.14333.

Hare, Christopher. 2022. Constrained Citizens? Ideological Structure and Conflict Extension in the US Electorate, 1980–2016 [in en]. British Journal of Political Science 52, no. 4 (October): 1602–1621. ISSN: 0007-1234, 1469-2112, accessed April 27, 2025. https://doi.org/10.1017/S000712342100051X. https://www.cambridge.org/core/product/identifier/S000712342100051X/type/journal_article.

Hegselmann, Rainer, Ulrich Krause, et al. 2002. Opinion dynamics and bounded confidence models, analysis, and simulation. Journal of Artificial Societies and Social Simulation 5 (3): 1–33.

Henry, Adam Douglas, Dieter Mitsche, and Paweł Prałat. 2016. Homophily, influence and the decay of segregation in self-organizing networks. Network Science 4, no. 1 (March): 81–116. ISSN: 2050-1250. https://doi.org/10.1017/nws.2016.1. http://dx.doi.org/10.1017/nws.2016.1.

Hirschman, Daniel. 2016. Stylized facts in the social sciences. Sociological Science 3:604–626. ISSN: 2330-6696. https://doi.org/10.15195/v3.a26. http://dx.doi.org/10.15195/v3.a26.

Hobolt, Sara B., Katharina Lawall, and James Tilley. 2024. The Polarizing Effect of Partisan Echo Chambers [in en]. American Political Science Review 118, no. 3 (August): 1464–1479. ISSN: 0003-0554, 1537-5943, accessed April 23, 2025. https://doi.org/10.1017/S0003055423001211. https://www.cambridge.org/core/journals/american-political-science-review/article/polarizing-effect-of-partisan-echo-chambers/5044B63A13A458A97CA747E9DCA07228.

Holme, Petter, and M. E. J. Newman. 2006. Nonequilibrium phase transition in the coevolution of networks and opinions. Physical Review E 74, no. 5 (November). ISSN: 1550-2376. https://doi.org/10.1103/physreve.74.056108. http://dx.doi.org/10.1103/PhysRevE.74.056108.

Hornsey, Matthew J. 2008. Social Identity Theory and Self-categorization Theory: A Historical Review [in en]. _Eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1751-9004.2007.00066.x, Social and Personality Psychology Compass 2 (1): 204–222. ISSN: 1751-9004, accessed December 4, 2022. https://doi.org/10.1111/j.1751-9004.2007.00066.x. https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1751-9004.2007.00066.x.

Huber, Gregory A., and Neil Malhotra. 2017. Political Homophily in Social Relationships: Evidence from Online Dating Behavior. Publisher: [The University of Chicago Press, Southern Political Science Association], The Journal of Politics 79 (1): 269–283. ISSN: 0022-3816, accessed April 27, 2025. https://www.jstor.org/stable/26550847.

Huddy, Leonie, Lilliana Mason, and Lene Aarøe. 2015. Expressive Partisanship: Campaign Involvement, Political Emotion, and Partisan Identity [in en]. American Political Science Review 109, no. 1 (February): 1–17. ISSN: 0003-0554, 1537-5943, accessed April 27, 2025. https://doi.org/10.1017/S0003055414000604. https://www.cambridge.org/core/journals/american-political-science-review/article/abs/expressive-partisanship-campaign-involvement-political-emotion-and-partisan-identity/7D2A2C87FBEBBE5DABAAF9658B3162AA.

Humphries, Mark D, and Kevin Gurney. 2008. Network 'small-world-ness': a quantitative method for determining canonical network equivalence. PloS one 3 (4): e0002051.

Iyengar, Shanto. 2016. E Pluribus Pluribus , or Divided We Stand. Public Opinion Quarterly 80, no. S1 (January): 219–224. ISSN: 0033-362X, accessed April 27, 2025. https://doi.org/10.1093/poq/nfv084. https://doi.org/10.1093/poq/nfv084.

Iyengar, Shanto, Tobias Konitzer, and Kent Tedin. 2018. The Home as a Political Fortress: Family Agreement in an Era of Polarization. Publisher: The University of Chicago Press, The Journal of Politics 80, no. 4 (October): 1326–1338. ISSN: 0022-3816, accessed April 27, 2025. https://doi.org/10.1086/698929. https://www.journals.uchicago.edu/doi/abs/10.1086/698929.

Iyengar, Shanto, Yphtach Lelkes, Matthew Levendusky, Neil Malhotra, and Sean J. Westwood. 2019. The Origins and Consequences of Affective Polarization in the United States [in en]. Publisher: Annual Reviews, Annual Review of Political Science 22, no. Volume 22, 2019 (May): 129–146. ISSN: 1094-2939, 1545-1577, accessed April 23, 2025. https://doi.org/10.1146/annurev-polisci-051117-073034. https://www.annualreviews.org/content/journals/10.1146/annurev-polisci-051117-073034.

Iyengar, Shanto, and Sean J. Westwood. 2015. Fear and Loathing across Party Lines: New Evidence on Group Polarization [in en]. _Eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/ajps.12152, American Journal of Political Science 59 (3): 690–707. ISSN: 1540-5907, accessed April 25, 2025. https://doi.org/10.1111/ajps.12152. https://onlinelibrary.wiley.com/doi/abs/10.1111/ajps.12152.

Kaldor, Nicholas. 1961. Capital accumulation and economic growth. In The theory of capital, 177–222. Palgrave Macmillan UK. ISBN: 9781349084524. https://doi.org/10.1007/978-1-349-08452-4_10. http://dx.doi.org/10.1007/978-1-349-08452-4_10.

Karimi, Fariba, Mathieu Génois, Claudia Wagner, Philipp Singer, and Markus Strohmaier. 2018. Homophily influences ranking of minorities in social networks. Scientific Reports 8, no. 1 (July). ISSN: 2045-2322. https://doi.org/10.1038/s41598-018-29405-7. http://dx.doi.org/10.1038/s41598-018-29405-7.

Kim, Myunghwan, and Jure Leskovec. 2011. Modeling social networks with node attributes using the multiplicative attribute graph model. https://doi.org/10.48550/ARXIV.1106.5053. https://arxiv.org/abs/1106.5053.

Klein, Dominik, and Johannes Marx. 2017. Wenn du gehst, geh ich auch. die rolle von informationskaskaden bei der entstehung von massenbewegungen. Politische Vierteljahresschrift 58 (4): 560–592. https://doi.org/10.5771/0032-3470-2017-4-560.

Kuran, Timur. 1989. Sparks and prairie fires: a theory of unanticipated political revolution. Public choice 61 (1): 41–74.

Lee, Eun-Ju. 2007. Deindividuation Effects on Group Polarization in Computer-Mediated Communication: The Role of Group Identification, Public-Self-Awareness, and Perceived Argument Quality [in en]. Journal of Communication 57, no. 2 (June): 385–403. ISSN: 0021-9916, 1460-2466, accessed April 25, 2025. https://doi.org/10.1111/j.1460-2466.2007.00348.x. https://academic.oup.com/joc/article/57/2/385-403/4102671.

Lelkes, Yphtach. 2016. Mass Polarization: Manifestations and Measurements. MAG ID: 2308341436, Public Opinion Quarterly 80, no. 1 (January): 392–410. https://doi.org/10.1093/poq/nfw005.

MacCarron, P., K. Kaski, and R. Dunbar. 2016. Calling Dunbar's numbers. Social Networks 47:151–155.

Mäki, Uskali. 2009. MISSing the world. Models as isolations and credible surrogate systems. Erkenntnis 70:29–43.

Manski, Charles F, and Daniel McFadden. 1981. Structural analysis of discrete data with econometric applications. MIT press Cambridge, MA.

Mason, Lilliana. 2018. Uncivil Agreement: How Politics Became Our Identity [in en]. Google-Books-ID: R29RDwAAQBAJ. Chicago: University of Chicago Press, April. ISBN: 978-0-226-52468-9.

Mason, Lilliana, and Julie Wronski. 2018. One Tribe to Bind Them All: How Our Social Group Attachments Strengthen Partisanship [in en]. _Eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/pops.12485, Political Psychology 39 (S1): 257–277. ISSN: 1467-9221, accessed April 27, 2025. https://doi.org/10.1111/pops.12485. https://onlinelibrary.wiley.com/doi/abs/10.1111/pops.12485.

Mayerhoffer, Daniel M., and Jan Schulz. 2022. Perception and privilege. Applied Network Science 7 (1).

McPherson, Miller, Lynn Smith-Lovin, and Matthew E Brashears. 2006. Social isolation in america: changes in core discussion networks over two decades. American Sociological Review 71 (3): 353–375.

McPherson, Miller, Lynn Smith-Lovin, and James M Cook. 2001. Birds of a feather: homophily in social networks. Annual Review of Sociology 27, no. 1 (August): 415–444. ISSN: 1545-2115. https://doi.org/10.1146/annurev.soc.27.1.415. http://dx.doi.org/10.1146/annurev.soc.27.1.415.

Mehlhaff, Isaac D. 2024. A Group-Based Approach to Measuring Polarization [in en]. American Political Science Review 118, no. 3 (August): 1518–1526. ISSN: 0003-0554, 1537-5943, accessed April 23, 2025. https://doi.org/10.1017/S0003055423001041. https://www.cambridge.org/core/journals/american-political-science-review/article/groupbased-approach-to-measuring-polarization/7979CFDC243FB1E21FE7AE3724B42EEA.

Meylahn, Benedikt V, and Christa Searle. 2024. Opinion dynamics beyond social influence. Network Science 12 (4): 339–365. https://doi.org/10.1017/nws.2024.14.

Moore-Berg, Samantha L., Lee-Or Ankori-Karlinsky, Boaz Hameiri, and Emile Bruneau. 2020. Exaggerated meta-perceptions predict intergroup hostility between american political partisans. Proceedings of the National Academy of Sciences 117 (26): 14864–14872.

Moscovici, Serge, and Marisa Zavalloni. 1969. The group as a polarizer of attitudes. Place: US Publisher: American Psychological Association, Journal of Personality and Social Psychology 12 (2): 125–135. ISSN: 1939-1315. https://doi.org/10.1037/h0027568.

Munn, Luke. 2020. Angry by design: toxic communication and technical architectures. Humanities and Social Sciences Communications 7 (1): 1–11.

Myers, David G., and George D. Bishop. 1970. Discussion Effects on Racial Attitudes. Publisher: American Association for the Advancement of Science, Science 169, no. 3947 (August): 778–779. Accessed April 25, 2025. https://doi.org/10.1126/science.169.3947.778. https://www.science.org/doi/abs/10.1126/science.169.3947.778.

Myers, David G., James B. Bruggink, Robert C. Kersting, and Barbara A. Schlosser. 1980. Does Learning Others' Opinions Change One's Opinions? [In EN]. Publisher: SAGE Publications Inc, Personality and Social Psychology Bulletin 6, no. 2 (June): 253–260. ISSN: 0146-1672, accessed April 25, 2025. https://doi.org/10.1177/014616728062011. https://doi.org/10.1177/014616728062011.

Nardini, Cecilia, Balázs Kozma, and Alain Barrat. 2008. Who's talking first? consensus or lack thereof in coevolving opinion formation models. Physical Review Letters 100, no. 15 (April). ISSN: 1079-7114. https://doi.org/10.1103/physrevlett.100.158701. http://dx.doi.org/10.1103/PhysRevLett.100.158701.

Nasuto, Andrea, and Francisco Rowe. 2024. Understanding anti-immigration sentiment spreading on twitter. Edited by Claire Seungeun Lee. PLOS ONE 19, no. 9 (September): e0307917. ISSN: 1932-6203. https://doi.org/10.1371/journal.pone.0307917. http://dx.doi.org/10.1371/journal.pone.0307917.

Newman, M. E. J., and R. M. Ziff. 2001. Fast monte carlo algorithm for site or bond percolation. Physical Review E 64, no. 1 (June). ISSN: 1095-3787. https://doi.org/10.1103/physreve.64.016706. http://dx.doi.org/10.1103/PhysRevE.64.016706.

Nyhan, Brendan, Jaime Settle, Emily Thorson, Magdalena Wojcieszak, Pablo Barberá, Annie Y. Chen, Hunt Allcott, et al. 2023. Like-minded sources on Facebook are prevalent but not polarizing [in en]. Publisher: Nature Publishing Group, Nature 620, no. 7972 (August): 137–144. ISSN: 1476-4687, accessed April 25, 2025. https://doi.org/10.1038/s41586-023-06297-w. https://www.nature.com/articles/s41586-023-06297-w.

Pariser, Eli. 2011. The Filter Bubble: What The Internet Is Hiding From You [in en]. London: Penguin, May. ISBN: 978-0-14-196992-3.

Reutlinger, Alexander, Dominik Hangleiter, and Stephan Hartmann. 2018. Understanding (with) toy models. The British Journal for the Phi 69, no. 4 (December): 1069–1099.

Robinson, Robert J., Dacher Keltner, Andrew Ward, and Lee Ross. 1995. Actual versus assumed differences in construal: "naive realism" in intergroup perception and conflict. Journal of Personality and Social Psychology 68, no. 3 (March): 404–417. ISSN: 0022-3514. https://doi.org/10.1037/0022-3514.68.3.404. http://dx.doi.org/10.1037/0022-3514.68.3.404.

Santos, Fernando P., Yphtach Lelkes, and Simon A. Levin. 2021. Link recommendation algorithms and dynamics of polarization in online social networks. Proceedings of the National Academy of Sciences 118 (50): e2102141118. https://doi.org/10.1073/pnas.2102141118.

Schneider-Strawczynski, Sarah, and Jérôme Valette. 2025. Media coverage of immigration and the polarization of attitudes. American Economic Journal: Applied Economics 17, no. 1 (January): 337–368. ISSN: 1945-7790. https://doi.org/10.1257/app.20230300. http://dx.doi.org/10.1257/app.20230300.

Schulz, Jan, and Daniel M. Mayerhoffer. 2023. A network approach to expenditure cascades. Review of Behavioral Economics 10 (3): 229–262. ISSN: 2326-6201.

Schulz, Jan, Daniel M. Mayerhoffer, and Anna Gebhard. 2022. A network-based explanation of inequality perceptions. Social Networks 70:306–324.

Schulz, Moritz A., and Simon Scheller. 2024. Being a believer: social identity in post-truth political discourse. Inquiry: An Interdisciplinary Jour https://doi.org/10.1080/0020174x.2024.2312201.

Schwaninger, Markus, and Stefan Grösser. 2008. System dynamics as model-based theory building. Systems Research and Behavioral Science 25, no. 4 (July): 447–465. ISSN: 1099-1743. https://doi.org/10.1002/sres.914. http://dx.doi.org/10.1002/sres.914.

Sides, John, Michael Tesler, and Lynn Vavreck. 2018. Identity Crisis: The 2016 Presidential Campaign and the Battle for the Meaning of Ame [in English]. Princeton: Princeton University Press, October. ISBN: 978-0-691-17419-8.

Squazzoni, Flaminio. 2012. Agent-based computational sociology. John Wiley & Sons.

Steppat, Desiree, Laia Castro Herrero, and Frank Esser. 2022. Selective exposure in different political information environments – How media fragmentation and polarization shape congruent news use [in en]. Publisher: SAGE Publications Ltd, European Journal of Communication 37, no. 1 (February): 82–102. ISSN: 0267-3231, accessed November 21, 2022. https://doi.org/10.1177/02673231211012141. https://doi.org/10.1177/02673231211012141.

Sugden, Robert. 2009. Credible worlds, capacities and mechanisms. Erkenntnis 70 (1): 3–27.

Sunstein, Cass R. 2002. The Law of Group Polarization. MAG ID: 3125714960, Journal of Political Philosophy 10 (2): 175–195. https://doi.org/10.1111/1467-9760.00148.

———. 2009. Going to Extremes: How Like Minds Unite and Divide [in en]. Oxford: Oxford University Press, May. ISBN: 978-0-19-979314-3.

———. 2018. Republic: Divided Democracy in the Age of Social Media [in en]. Publisher: Princeton University Press, 1–328. Accessed April 25, 2025. https://www.torrossa.com/en/resources/an/5559626.

Taghikhah, Firouzeh, Tatiana Filatova, and Alexey Voinov. 2021. Where does theory have it right? a comparison of theory-driven and empirical agent based models. Journal of Artificial Societies and Social Simulation 24 (2). ISSN: 1460-7425. https://doi.org/10.18564/jasss.4573. http://dx.doi.org/10.18564/jasss.4573.

Tajfel, Henri, and John C. Turner. 1986. The Social Identity Theory of Intergroup Behavior. In Psychology of intergroup relations, edited by Stephen Worchel and William G. Austin, 7–24. 00754. Chicago: Nelson.

Talisse, Robert B. 2020. Overdoing Democracy: Why We Must Put Politics in its Place [in eng]. New York: Oxford University Press. ISBN: 978-0-19-092419-5, accessed July 21, 2022. https://doi.org/10.1093/oso/9780190924195.001.0001. https://oxford.universitypressscholarship.com/view/10.1093/oso/9780190924195.001.0001/oso-9780190924195.

———. 2021. Problems of Polarization [in en]. In Political Epistemology, edited by Elizabeth Edenberg and Michael Hannon, 209–225. Oxford: Oxford University Press. ISBN: 978-0-19-289333-8. https://doi.org/10.1093/oso/9780192893338.003.0012.

Tang, Tanzhe, Amineh Ghorbani, Flaminio Squazzoni, and Caspar G. Chorus. 2021. Together alone: a group-based polarization measurement. Quality amp; Quantity 56, no. 5 (December): 3587–3619. ISSN: 1573-7845. https://doi.org/10.1007/s11135-021-01271-y. http://dx.doi.org/10.1007/s11135-021-01271-y.

Turner, Cody. 2023. Online echo chambers, online epistemic bubbles, and open-mindedness. Episteme (November): 1–26. ISSN: 1750-0117. https://doi.org/10.1017/epi.2023.52. http://dx.doi.org/10.1017/epi.2023.52.

Ueffing, Philipp, Francisco Rowe, and Clara H. Mulder. 2015. Differences in attitudes towards immigration between australia and germany: the role of immigration policy. Comparative Population Studies 40, no. 4 (December). ISSN: 1869-8980. https://doi.org/10.12765/cpos-2015-18. http://dx.doi.org/10.12765/CPoS-2015-18.

Vaidhyanathan, Siva. 2018. Antisocial Media: How Facebook Disconnects Us and Undermines Democracy [in en]. Google-Books-ID: h05WDwAAQBAJ. Oxford University Press. ISBN: 978-0-19-084116-4.

Watts, Duncan J, and Steven H Strogatz. 1998. Collective dynamics of 'small-world' networks. Nature 393 (6684): 440–442.

Westfall, Jacob, Leaf Van Boven, John R. Chambers, and Charles M. Judd. 2015. Perceiving political polarization in the united states: party identity strength and attitude extremity exacerbate the perceived partisan divide. Perspectives on Psychological Science 10, no. 2 (March): 145–158.

Windzio, Michael. 2025. Limits of rationality? moral foundations and polarization in attitudes towards diversity in germany. Rationality and Society (March). ISSN: 1461-7358. https://doi.org/10.1177/10434631251327035. http://dx.doi.org/10.1177/10434631251327035.

Wollebæk, Dag, Rune Karlsen, Kari Steen-Johnsen, and Bernard Enjolras. 2019. Anger, fear, and echo chambers: the emotional basis for online behavior. Social Media + Society 5, no. 2 (April). ISSN: 2056-3051. https://doi.org/10.1177/2056305119829859. http://dx.doi.org/10.1177/2056305119829859.

Appendix 1. Association between left-right self-placement and migration scepticism

This appendix reports additional tests for the association between political self-placement and the composite migration scepticism index used in the empirical application. The tests complement the graphical evidence in the main text and assess whether migration scepticism differs systematically across left-right self-placement categories.

Table 3. Tests of the association between left-right self-placement and migration scepticism

Test	Statistic	Degrees of freedom	p-value	Effect size / note
Linear trend: migration scepticism on left-right placement	Slope = 0.0540; t = 26.8523	df = 3009	< 0.001	R^2 = 0.1933
Spearman rank correlation	ρ = 0.4476; t = 27.4565	df = 3009	< 0.001	Monotonic association
One-way ANOVA across left-right categories	F = 89.9981	df = 9, 3001	< 0.001	η^2 = 0.2125
Kruskal-Wallis test across left-right categories	H = 632.1041	df = 9	< 0.001	ϵ^2 = 0.2076

Notes: The dependent variable is the composite migration scepticism index. Left-right self-placement ranges from 1 (left) to 10 (right). The table reports complementary parametric and rank-based tests. The linear trend test treats left-right self-placement as an ordered variable; the ANOVA and Kruskal-Wallis tests compare the ten self-placement categories.

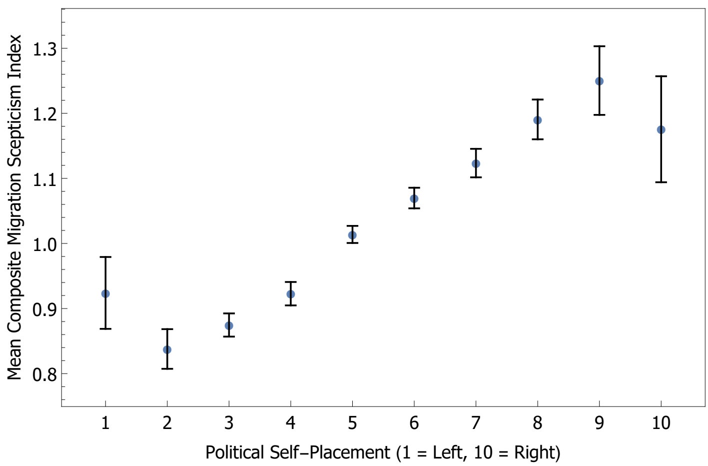

Note: Points show category-specific means of the composite migration scepticism index. Vertical bars show 95 percent confidence intervals. Political self-placement ranges from 1 (left) to 10 (right). The figure shows that migration scepticism increases systematically with rightward self-placement, while the association remains imperfect.

Figure 11. Mean composite migration scepticism by political self-placement

Appendix 2. Degree distributions and the degree-attitude relationship for lognormal attitudes

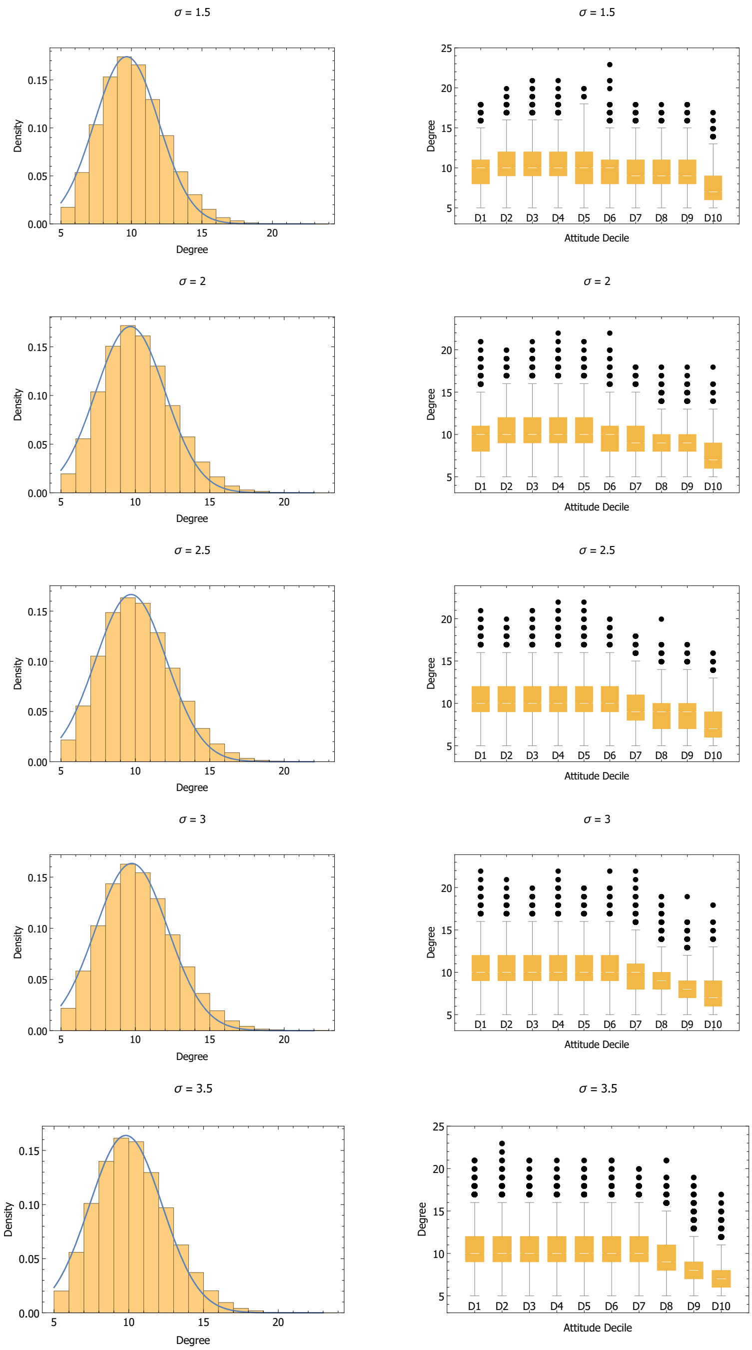

Figure 12. Degree distributions and degree by attitude decile under lognormal attitude distributions and varying dispersion parameters $\sigma \in \{1.5, 2, 2.5, 3, 3.5\}$. 100 Monte Carlo runs each. All simulations use $\rho = 8$.

Appendix 3. Degree distributions and the degree-attitude relationship for beta, uniform, and exponential attitudes

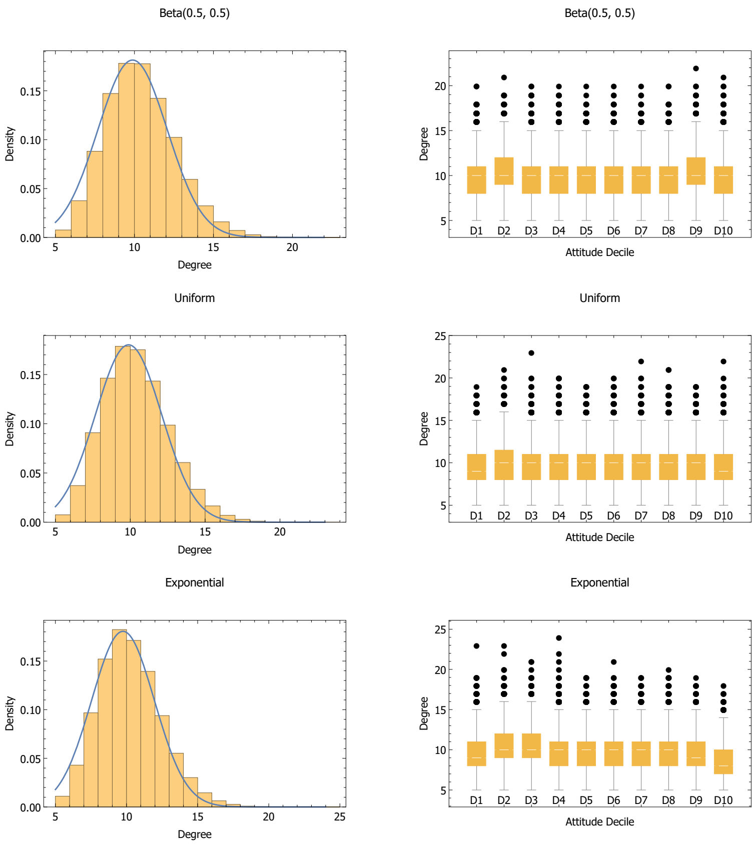

Figure 13. Degree distributions and degree by attitude decile under beta, uniform, and exponential attitude distributions. 100 Monte Carlo runs each. All simulations use $\rho = 8$.

Appendix 4. Approximation of expected degrees as a function of attitude

Formally, let p_{ij} denote the normalised probability weight with which agent i selects agent j. In the one-dimensional case, this is

$$p_{ij} = \frac{\exp(-\rho|A_i - A_j|)}{\sum_{\ell \neq i} \exp(-\rho|A_i - A_\ell|)}. \tag{9}$$

Let Y_{ij} be an indicator that takes the value one if agent i selects agent j, and zero otherwise. If agent i forms d_i links and $d_i \ll N$, then

$$\Pr(Y_{ij} = 1) \approx d_i p_{ij}. \tag{10}$$

The expected number of incoming selections received by node j is therefore

$$\begin{aligned} E\left[k_j^{\text{in,sel}}\right] &= E\left[\sum_{i \neq j} Y_{ij}\right] \\ &= \sum_{i \neq j} E[Y_{ij}] \\ &\approx \sum_{i \neq j} d_i p_{ij} \\ &= \sum_{i \neq j} d_i \frac{\exp(-\rho|A_i - A_j|)}{\sum_{\ell \neq i} \exp(-\rho|A_i - A_\ell|)}. \end{aligned} \tag{11}$$

For the undirected simple graph, total degree also includes the neighbours chosen by j itself. Moreover, reciprocal selections have to be counted only once, because a mutual selection creates a single undirected edge. Let Z_{ij} indicate whether an undirected edge between i and j exists. Then

$$Z_{ij} = \mathbf{1}\{Y_{ij} = 1 \text{ or } Y_{ji} = 1\}. \tag{12}$$

We can express this as a sum, i.e.,

$$Z_{ij} = Y_{ij} + Y_{ji} - Y_{ij} Y_{ji}. \tag{13}$$

Hence, the expected undirected degree of node j is

$$\begin{aligned} E[k_j] &= \sum_{i \neq j} E[Z_{ij}] \\ &= \sum_{i \neq j} E[Y_{ij}] + \sum_{i \neq j} E[Y_{ji}] - \sum_{i \neq j} E[Y_{ij} Y_{ji}] \\ &\approx \sum_{i \neq j} d_i p_{ij} + \sum_{i \neq j} d_j p_{ji} - \sum_{i \neq j} d_i d_j p_{ij} p_{ji}. \end{aligned} \tag{14}$$

Since $\sum_{i \neq j} p_{ji} = 1$, the first summand equals d_j. Thus,

$$E[k_j] \approx d_j + \sum_{i \neq j} d_i p_{ij} - \sum_{i \neq j} d_i d_j p_{ij} p_{ji}. \tag{15}$$

If reciprocal selections are rare, as should be the case for a small minimum degree compared to the number of agents N, the final correction term is small. In that case, the expected degree can be approximated by

$$E[k_j] \approx d_j + \sum_{i \neq j} d_i \frac{\exp(-\rho |A_i - A_j|)}{\sum_{\ell \neq i} \exp(-\rho |A_i - A_\ell|)}. \tag{16}$$

Equation (16) separates the two components of degree: the fixed own-choice component d_j, and the incoming-selection component generated by other agents' homophilic choices. It also shows where the attitude distribution enters the degree-generating process. For $\rho > 0$, incoming selection probabilities vary with the position of A_j in the attitude distribution. Nodes in denser regions of attitude space tend to be close to more potential choosers and therefore receive more incoming selections. At the same time, dense regions also contain more competing targets, because each chooser normalises over all possible neighbours. The attitude distribution therefore affects the conditional degree-attitude relationship but, as our numerical simulations demonstrate, not sufficiently strongly to affect the approximation of the degree distribution as (truncated) Gaussian.