The Supply Side: Are Biased Information Sources Necessary for Belief Polarisation?

Jens Koed Madsen
(j.madsen2@lse.ac.uk)

Department of Psychological and Behavioural Science, London School of Economics and Political Science, Houghton Street, WC2A 2AE, London, UK.

Abstract

Do biased minds or sources cause issue polarisation? Demand-side accounts emphasise motivated reasoning where people selectively attend to, and favourably evaluate, attitude-consistent information. Supply-side accounts emphasise the information environment, such as ideologically committed sources. We present an agent-based model to test if motivated reasoning mechanisms can create polarisation without biased information sources. Our simulations suggest that populations converge to consensus, regardless of motivated reasoning strength. Polarisation only emerges given ideological broadcasters. Critically, pre-polarised populations converge without persistent biased sources, indicating that motivated reasoning does not maintain existing polarisation. We also find that selective exposure counterintuitively *reduces* polarisation from consensus by symmetrically filtering all extreme sources. These results support a supply-side thesis: within the logic of the model, motivated reasoning mechanisms were not sufficient for belief polarisation without biased information sources. Here, motivated reasoning and echo chambers are amplifiers, not generators, of issue polarisation.

Keywords: polarisation; motivated reasoning; selective exposure; agent-based modelling; information environment

Introduction

Political polarisation happens when beliefs move in opposing directions while living in the same information environment. It has become a core feature of contemporary democracies (Iyengar et al., 2019; Mason, 2018; Webster & Abramowitz, 2017). Understanding its causes has significant implications for democratic functioning, as polarised populations struggle to find common ground on policy issues and may become susceptible to misinformation (Cinelli et al., 2020). Further, it seems to have increased in places like the USA over the past years (Young et al., 2026). Two broad approaches dominate. Demand-side accounts emphasise cognitive biases in how people process information. Supply-side accounts focus on the information environment itself.

The demand-side perspective is often rooted in motivated reasoning literature (Kunda, 1990; Taber & Lodge, 2006), which holds that people are biased information processors who evaluate attitude-consistent evidence more favourably (biased assimilation), generate more counterarguments against attitude-inconsistent evidence (disconfirmation bias), and preferentially seek out confirming information (confirmation bias or selective exposure). According to this view, even balanced information exposure should produce polarisation because cognitive biases transform neutral information into belief-reinforcing experiences.

However, recent work challenged this account. Coppock (2023) has a Bayesian reinterpretation of biased assimilation and disconfirmation bias, showing that identical behavioural patterns could emerge from rational agents with different priors rather than motivated cognition. Young, Madsen, and de Wit (2025) show experimentally that belief polarisation can occur through rational Bayesian updating when people disagree about source independence, without requiring motivated reasoning. Tappin, Pennycook, and Rand (2020) argue that paradigmatic study designs in the motivated reasoning literature often undermine causal inference. These critiques suggest that the evidence for demand-side mechanisms may be weaker than commonly assumed. Jern, Chang, and Kemp (2014) further show that Bayesian agents can polarise on identical evidence given different priors, suggesting prior informational asymmetries may underlie apparent demand-side effects.

Meanwhile, a growing literature emphasises the role of information supply. Studies of misinformation dynamics treat information ecosystems as complex systems where feedback loops, network effects, and the actions of influential sources play critical roles (Bak-Coleman et al., 2025). Rosenstock et al. (2017) show that in epistemic network models, the benefits of reduced connectivity are not robust, occurring only when learning is especially difficult. This suggests that we should be cautious about attributing polarisation to network structure alone. Madsen et al (2018) show that even rational Bayesian agents can form persistent echo chambers through confirmatory information search, where agents progressively narrow the range of beliefs they accept as their confidence increases.

This paper contributes to this debate by testing if demand-side factors (motivated reasoning) can create polarisation independently of supply-side factors (biased information sources). We develop an agent-based model that implements multiple motivated reasoning mechanisms, including a strong form of selective exposure that lacks an obvious Bayesian interpretation (Coppock, 2023). We explore if motivated reasoning, even in a strong form, can create polarisation from initial consensus without biased information sources?

In G. Boleda, I. Dautriche, M. Maldonado, C. Saldana and T. Torrent (Eds.), *Proceedings of the 48th Annual Conference of the Cognitive Science Society*

Motivated reasoning

Kunda's (1990) leading theory identifies motivated reasoning as occurring when directional goals influence the cognitive strategies people employ. Three mechanisms have been proposed where motivated reasoning affects belief formation.

Biased assimilation is the tendency to evaluate attitude-consistent evidence more favourably than inconsistent evidence. Lord et al (1979) showed this when proponents and opponents of capital punishment rated the same mixed evidence differently, depending on each side finding studies supporting their position more convincing. **Disconfirmation bias** describes the asymmetric scrutiny applied to congenial versus uncongenial information where people generate more counterarguments against attitude-inconsistent evidence (Taber & Lodge, 2006). Finally, **selective exposure** (also called confirmation bias in information search) refers to the preferential seeking of attitude-consistent information. Unlike the other two mechanisms, selective exposure operates on information acquisition rather than processing.

Bayesian Reinterpretations and the Selectivity of the Critique

Coppock (2023) shows that both biased assimilation and disconfirmation bias can be reinterpreted within a Bayesian framework. If people believe that sources aligned with their priors are more likely to be reliable (which seems reasonable if priors are based on previously accumulated evidence), then discounting counter-attitudinal information emerges from rational updating. Similarly, spending more cognitive effort criticizing low-quality evidence is not irrational but rather an efficient allocation of limited cognitive resources.

Coppock acknowledges that selective exposure is harder to rationalize: active avoidance of potentially informative signals is difficult to justify as optimal information gathering.

Complex systems

Recent papers emphasise the need for complex systems approaches to understand polarisation and misinformation dynamics (Bak-Coleman et al., 2025). Individual-level randomized controlled trials cannot capture emergent properties of information ecosystems, including feedback loops between content creators and audiences, network spillover effects, and hysteresis (path-dependence). Agent-based models provide a complementary method to bridge individual-level functions and population-level outcomes (Epstein & Axtell, 1996).

Our model follows this line, allowing us to test whether micro-level cognitive mechanisms (motivated reasoning) can produce macro-level outcomes (population polarisation) under different information environment configurations.

The model

We developed an agent-based model in NetLogo 7.0.2 that incorporates motivated reasoning mechanisms, network dynamics, and asymmetric information sources. Each model element is grounded in the above literature. The model allows us to systematically vary supply-side factors (presence of biased broadcasters) and demand-side factors (motivated reasoning mechanisms) to test their independent and joint contributions to polarisation.

Agents and beliefs

The model contains 200 agents, each holding a continuous belief between 0 and 1 about some proposition. Following the convention in polarisation research, we interpret 0 and 1 as opposing ideological poles (e.g., liberal vs. conservative positions), with 0.5 representing a centrist or undecided position. This continuous representation allows us to capture gradual belief change and measure polarisation as increased variance in the belief distribution.

Broadcasters (biased information sources)

Following partisan media literature (Levendusky, 2013), we model 'broadcasters' as ideologically positioned information sources. When broadcasters are enabled, a 'left' broadcaster at position 0.1 and a 'right' broadcaster at position 0.9 each influence a subset of the population. Unlike regular agents, broadcasters never update their positions as they represent ideologically committed information sources such as partisan media outlets or political elites.

Broadcaster exposure reflects the empirically observed pattern of partisan news consumption where agents in the lower half of the initial belief distribution are exposed to the left broadcaster while agents in the upper half to the right broadcaster. This captures the fact that people's media diets correlate with their political dispositions, even if weakly so at initial consensus. Broadcaster reach (the proportion exposed, set to 50%) and strength (influence per interaction, set to 0.15) are parameterized to allow sensitivity analysis.

Broadcasters represent the *supply side* of polarisation. They are external to the agent population and persist regardless of the agents' responses. Disabling broadcasters allows us to isolate demand-side mechanisms and test whether motivated reasoning alone can generate polarisation.

Peer influence and network structure

Agents are connected in a social network with initial degree $k = 4$ and maximum degree 12. Networks can be random (Erdős-Rényi), clustered by initial belief, or structured as two communities with limited bridging connections. Each time step, agents receive messages from their network neighbours and update beliefs accordingly.

The basic peer influence mechanism follows a weighted averaging model where agents move their beliefs toward the positions of those they communicate with, with the magnitude proportional to a peer strength parameter (set to 0.1). Formally, for each message received from a peer at position b_peer, the agent updates: b_new = b_old + peer_strength × (b_peer - b_old). This captures the tendency for social influence to produce convergence in connected networks (DeGroot, 1974; Friedkin & Johnsen, 1990).

Motivated reasoning mechanisms

We implement three forms of motivated reasoning, each representing different mechanisms proposed in the psychological literature:

Linear discounting describes the biased assimilation. The influence of a message is reduced proportional to the distance between the agent's belief and the source's position. Formally, effective_strength = base_strength × (1 - mr_strength × |b_agent - b_source|), where mr_strength (set to 0.8) controls the magnitude of discounting. An agent at 0.5 receiving a message from a source at 0.9 discount it by 32% (0.8 × 0.4). This captures Taber and Lodge's (2006) finding that people rate counter-attitudinal arguments as less persuasive, though as Coppock (2023) notes, this pattern is also consistent with Bayesian reasoning about source quality.

Threshold acceptance with repulsion captures the backfire effect sometimes shown in the motivated reasoning. Beyond a threshold distance (set to 0.3), agents may move *away* from the source (negative influence). This represents the strongest form of biased assimilation, though empirical support for backfire effects is limited (Wood & Porter, 2019).

The previous functions introduce motivated reasoning in information *processing*. **Selective exposure** implements confirmation bias in information *search*. To do so, agents probabilistically ignore messages from sources with distant beliefs. The probability of attending to a message decreases with belief distance: P(attend) = max(min_attend, 1 - se_strength × distance × 2). With se_strength = 0.8 and min_attend = 0.1, an agent at 0.5 has only a 36% chance of attending to a message from a source at 0.1 or 0.9, but near-certainty of attending to messages from nearby peers.

Selective exposure is the strongest motivated reasoning function we test. Coppock (2023) notes it cannot have an obvious Bayesian interpretation, as rational agents should attend all informative signals, even if they subsequently discount them for being weak. Ignoring information entirely based on its expected content is a departure from optimal information gathering and so is difficult to rationalise. By implementing selective exposure, we test whether even a strong form of motivated reasoning can create polarisation.

This distance-based implementation is symmetric with respect to direction, as agents filter based on extremity, not attitude-consistency per se. As such, our selective exposure mechanism may be closer to Coppock's rational quality heuristic than to directionally motivated avoidance, making it a conservative test of the motivated reasoning account. A directional alternative (e.g., log-surprisal; cf. Jern et al., 2014) would be undefined at consensus, since directional asymmetry presupposes the polarisation under test.

Social network mechanisms

Following homophily literature (McPherson, Smith-Lovin, & Cook, 2001; Centola, 2011), agents can adapt their networks by severing ties agents that are dissimilar and forming ties with similar agents. When enabled, agents probabilistically drop connections to neighbours whose beliefs differ by more than a tolerance threshold (set to 0.3) and seek new connections with belief-similar agents. This process occurs each time step with a rewiring probability of 0.1.

This function allows echo chambers to form endogenously because of belief dynamics. The question is whether echo chambers alone (segregated networks of like-minded individuals) can create polarisation, or whether they merely amplify pre-existing differences driven by other factors.

Outcome measures

Our primary outcome measure is *belief variance*, which is the variance of beliefs across the population. Polarisation from consensus means variance increases from its initial (near-zero) level. We report the *variance ratio* (final/initial variance) as our main metric. Values greater than 1 indicate polarisation, while values near 0 indicate convergence. We also track *between-group difference* (mean of high group minus mean of low group), *attend rate* (proportion of messages attended to under selective exposure), and *mean network degree* (to assess echo chamber formation).

Simulated experiments and results

We ran systematic simulations that vary the presence of motivated reasoning mechanisms, broadcasters, and network adaptation. Each condition was run for 500 steps with 30 replications to ensure stable estimates. Parameters were set to strongly favour polarisation when possible: mr_strength = 0.8, se_strength = 0.8, broadcaster_strength = 0.15, peer_strength = 0.1. This parameterisation tests the motivated reasoning account. If polarisation cannot emerge under conditions that favour demand-side mechanisms, it suggests fundamental limitations of the account (given the interpretation of the theories that we implemented, of course).

Table 1 presents the results from a 2×2 factorial design crossing selective exposure (on/off) with broadcasters (on/off), starting from consensus. This design allows us to isolate the independent effects of the supply-side factor (broadcasters) and the demand-side factor (selective exposure), as well as their interaction.

Table 1: Factorial Results: Selective Exposure × Broadcasters (consensus start, N=200, 30 reps)

Condition	SE	BC	Final Var.	Var. Ratio	Polarised?
Baseline	OFF	OFF	~0	0.00	No
SE Only	ON	OFF	~0	0.00	No
BC Only	OFF	ON	0.0072	54.51	Yes
SE + BC	ON	ON	0.0010	7.74	Yes (reduced)

The results suggest polarisation does not emerge without broadcasters (rows 1-2), even if selective exposure is enabled. The variance ratio is effectively zero, meaning that agents converge to consensus through peer influence. With broadcasters (rows 3-4), polarisation emerges robustly. The variance ratio of 54.51 indicates substantial divergence in the population. That is, selective exposure on its own does not create polarisation. Even with 64% of messages from distant sources being ignored, convergence occurs with absent

broadcasters. This challenges the demand-side claim that motivated reasoning is sufficient for polarisation.

Figure 1 illustrates the findings through belief trajectories over the 500 steps of the simulation. The contrast between conditions show that all agents converge toward 0.5 regardless of selective exposure without broadcasters (Panels A and B). With broadcasters (Panels C and D), the population splits into two groups moving toward their respective broadcasters.

Figure 1: *Belief trajectories under four conditions.*

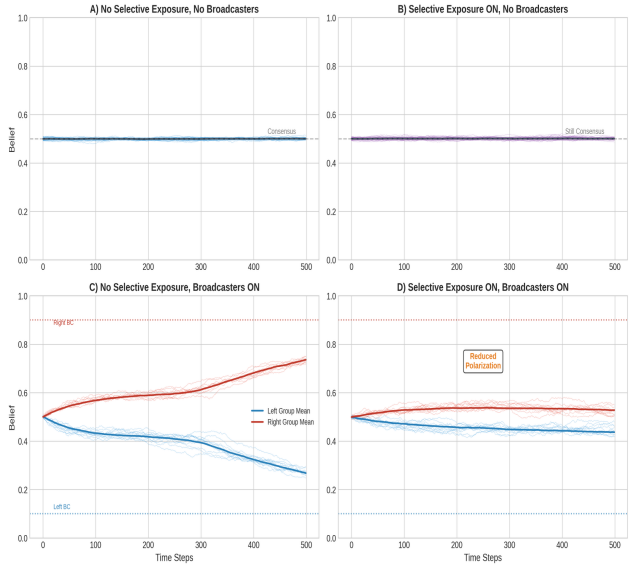

Panels A and B show convergence to consensus without broadcasters, regardless of selective exposure. Panels C and D show polarisation with broadcasters, with Panel D demonstrating that selective exposure actually reduces the magnitude of polarisation from consensus.

Surprisingly, selective exposure even reduces broadcaster-driven polarisation when starting from consensus (variance ratio 7.74 vs. 54.51, representing an 86% reduction). This result challenges the intuition that motivated reasoning should amplify polarisation. By tracing the information flow in the model, this mechanism becomes clear. At consensus, both broadcasters (at 0.1 and 0.9) are equidistant from agents at 0.5, so both are filtered equally (attend rate ~36%). Meanwhile, peer influence remains strong (~95-100% attend rate) because peers are nearby in belief space.

This reveals a crucial insight. Selective exposure does not just filter the 'other side'. It filters *all* extreme sources, including one's 'own' broadcaster. When everyone starts at the centre, everyone is equally distant from both poles, so selective exposure reduces influence from both directions equally. Therefore, peer influence (which has higher attend rates due to smaller belief distances) dominates, pulling everyone toward the mean rather than toward the extremes.

To test if the results depend on the consensus starting point, we ran simulations starting from a bimodal distribution with groups centred at 0.27 and 0.72. This represents a moderately polarised population where each side has substantial "distant" information from the other that could trigger motivated reasoning mechanisms. That is, they start out disagreeing with each other on some issue.

Table 2 presents the results. The findings remain. Without broadcasters, groups converge completely regardless of the new starting conditions. Even beginning with substantial separation, final variance collapses to effectively zero (mean ratio < 0.001). Motivated reasoning does not maintain pre-existing moderate polarisation, let alone create it, within the context of the implemented model.

This is interesting because the bimodal condition should be favourable to motivated reasoning accounts. Instead, we observe convergence at rates nearly identical to the consensus condition. Notably, in the bimodal condition the filter operates asymmetrically in practice (an agent at 0.27 filters the right broadcaster more than the left), yet convergence still occurs.

*Table 2: Bimodal vs. Consensus Starting Conditions**

Initial Condition	BC	Final Var.	Var. Ratio	Outcome
Consensus	OFF	~0	0.00	Convergence
Consensus	ON	0.0016	12.19	Polarisation
Bi-modal	OFF	~0	<0.001	Convergence
Bi-modal	ON	0.0015	0.20	Maintained*

** Consensus = 0.49 and 0.51; Bi-modal = 0.27 and 0.72*

The mechanism is straightforward. Peer influence operates through local averaging, which mathematically guarantees convergence toward the population mean in the absence of persistent extreme sources (DeGroot, 1974). Motivated reasoning can modulate the *rate* of this convergence but cannot reverse its direction. No matter how strongly agents discount distant information, some influence still occurs with each interaction. Over time, this inexorable averaging overwhelms initial differences and leads to convergence.

Network adaptation (echo chamber formation) amplifies polarisation when combined with broadcasters but cannot create it independently. Table 3 shows results from combining motivated reasoning mechanisms and network adaptation with and without broadcasters.

Table 3: The Echo Chamber effect

Condition	**BC**	**Net. Adapt.**	**Var. Ratio**	**Mean Degree**
Combined MR, No BC	OFF	ON	~0	4.0
BC Only	ON	OFF	54.51	4.0
BC + Net. Adapt.	ON	ON	~180	2.1

Network adaptation provides a ~3.3× amplification of broadcaster-driven polarisation (variance ratio increases from 54.51 to ~180). The mechanism suggests that, as agents begin to diverge due to broadcaster influence, they sever ties with increasingly dissimilar others, which creates segregated communities that reinforce within-group similarity while eliminating cross-cutting exposure.

Without broadcasters, network adaptation does not prevent convergence. Echo chambers form around the consensus. The network fragments into clusters, but all clusters converge to similar beliefs because there is no external force pushing

them apart. This suggests that echo chambers *amplify* but do not *generate* polarisation. Figure 2 visualises the network dynamics under these conditions, showing how network structure and belief distribution co-evolve.

Figure 2: Network structure and belief distribution

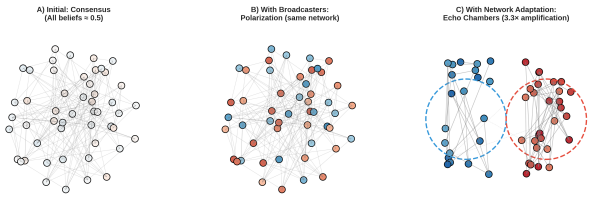

(A) Initial state at consensus with random network. (B) Polarisation with broadcasters but unchanged network topology. (C) Echo chamber formation with network adaptation, showing segregated communities and 3.3× amplification of polarisation.

We conducted robustness tests to ensure our findings are not artifacts of parameter choices. **Selective exposure strength** Even setting se_strength to 1.0 (maximum filtering) and an exponential distance model (which produces stronger filtering at moderate distances), polarisation did not emerge without broadcasters. Varying min_attend from 0.05 to 0.2 did not change the result. **Combined mechanisms** Enabling all motivated reasoning mechanisms simultaneously (selective exposure, linear discounting, threshold repulsion, and network adaptation) still does not produce polarisation without broadcasters. The mechanisms interact but cannot substitute for an asymmetric information supply. **Zealots** Replacing broadcasters with 'zealots' (agents who never change their beliefs) at extreme positions produces similar results. Zealots at 10-20% of the population can substitute for broadcasters in creating polarisation. This suggests the key factor is the presence of persistent, extreme-positioned sources, whether external (broadcasters) or internal (zealots). The supply-side thesis thus generalises beyond external media: polarisation requires persistent asymmetric signals, whether they originate outside or within the population.

Discussion

Our findings support a 'supply-side' thesis of polarisation. Within the logic of the model and its assumptions, demand-side mechanisms (motivated reasoning and echo chambers) were not sufficient to generate or maintain polarisation. In the model, polarisation only emerged when persistent, ideologically separated information sources were present.

This reframes the motivated reasoning debate. The question is not whether motivated reasoning *exists*, as the behavioural patterns documented in the literature are real. But whether it is *sufficient* to explain polarisation. Our model suggests it is not. Even the strongest form of motivated reasoning (selective exposure) did not create divergence from consensus without an asymmetric information supply.

Selective exposure and the Coppock critique

The counterintuitive finding that selective exposure *reduces* polarisation from consensus has important implications for the debate between motivated reasoning and Bayesian accounts. If selective exposure were a genuinely *motivated* bias driven by directional goals to arrive at preferred conclusions, it should selectively filter attitude-inconsistent information more than attitude-consistent information. This asymmetric filtering should, according to the motivated reasoning account, amplify polarisation.

Instead, we observe that selective exposure operates as a symmetric, distance-based filter. From consensus, agents filter both extreme sources equally because both are equally distant. This is in line with Coppock's (2023) Bayesian reinterpretation. If selective exposure reflects a rational heuristic (avoiding sources perceived as low-quality based on their extremity), it should operate symmetrically with respect to distance, not asymmetrically with respect to direction.

Our strongest evidence against the demand-side account comes not from the consensus-start simulations but from the bimodal positions. A defender of motivated reasoning might object that our main results are methodologically biased. By starting everyone at 0.5, we deny motivated reasoning mechanisms any "distant" information to work with. This objection misunderstands what the theory predicts.

If motivated reasoning were sufficient for polarisation, it should at minimum be able to maintain pre-existing differences. The bimodal condition tests this. Agents begin in separated groups with substantial cross-group distance. Under the motivated reasoning account, each group should resist influence from the other through selective attention, biased evaluation, and defensive processing. This does not occur in our simulations. Groups still converge regardless of initial separation.

This finding has important theoretical implications. It suggests that under weighted averaging dynamics—where social influence tends toward the mean of interacting beliefs (DeGroot, 1974)—information processing biases cannot reverse the direction of convergence. Motivated reasoning can slow convergence but cannot reverse it. Creating or maintaining polarisation requires a persistent force that continually pushes groups apart. In our model, that force is the supply side in the form of ideologically committed sources that never moderate their positions. In the real world, analogous forces include partisan media, political elites, and algorithmically amplified extreme content.

Relation to prior work

Our results align with several recent findings. Young et al.'s (2025) demonstration that polarisation can arise from rational updating about source independence suggests that apparent motivated reasoning may sometimes reflect justified differential trust in sources, which can be seen as a supply-side rather than demand-side factor. Our model provides a complementary demonstration: explicitly implementing non-Bayesian selective exposure, the supply side was still required for polarisation to emerge.

Rosenstock et al.'s (2017) finding that the benefits of reduced connectivity are not robust, but occurs in narrow parameter ranges where learning is especially difficult is also consistent with our results. Network structure alone did not reliably produce divergent outcomes. In our model, echo chambers without broadcasters produce convergence to consensus. Polarisation requires an asymmetric signal where the echo chamber structure alone is insufficient.

Policy implications If the supply side is required for polarisation to emerge under these dynamics, interventions should focus on information sources rather than (or in addition to) individual biases. Debiasing interventions that target motivated reasoning may have limited effects if the information environment remains polarised. Platform design changes that reduce amplification of extreme content, or policies that encourage source diversity, may be more effective than literacy interventions aimed at changing how individuals process information. Platform-level inoculation, for instance, can dampen misinformation spread in rational-agent networks (Pilditch, Roozenbeek, Madsen, & van der Linden, 2022).

Limitations

Naturally, the results should be taken with a significant pinch of salt, as they stem directly from the way motivated reasoning, network adaptation, and supply-side elements were built into the model. The model has several simplifying assumptions. Beliefs are unidimensional, while real political attitudes are multidimensional. Agents are homogeneous in their cognitive parameters, while real populations vary in susceptibility to different biases. The network is simplified compared to real social networks with their heavy-tailed degree distributions and community structure.

Our peer influence mechanism implements weighted averaging, which mathematically guarantees convergence in connected networks (DeGroot, 1974). Alternative baseline dynamics such as bounded confidence models where agents only influence those within a similarity threshold (Deffuant et al., 2000; Hegselmann & Krause, 2002) can produce persistent disagreement without external sources. Our results therefore show that motivated reasoning did not *overcome* convergent dynamics, not that it cannot generate polarisation under all possible interaction rules. Testing motivated reasoning within bounded confidence frameworks remains an important direction for future work. Our model tests whether motivated reasoning is *sufficient* for polarisation, not whether it is *necessary*. Our results do not rule out the possibility that motivated reasoning contributes to real-world polarisation in combination with biased information supply. What they suggest is that motivated reasoning was not sufficient on its own within this model.

One could argue that our bimodal starting conditions (0.27 and 0.72) underestimate initial real-world polarisation. We tested this with a uniform distribution over [0.15, 0.85] and found identical results: only sources that continuously inject extreme signals prevent convergence. Our model treats broadcasters as exogenous, but partisan media may emerge endogenously from audience demand and platform incentives, which represents a dynamic our framework cannot address.

Our findings may initially appear to contrast with Madsen et al.'s (2018) demonstration that Bayesian agents form echo chambers without broadcasters; their Bayesian-updating agents preserve initial heterogeneity through confirmatory search, while our weighted-averaging dynamics drive convergence.

Future work could explore heterogeneous agents, more realistic networks, and broadcaster emergence dynamics. Empirically, populations with similar cognitive profiles but different media environments should show different polarisation trajectories if the supply-side account is correct.

Concluding remarks

Our ABM suggests that motivated reasoning mechanisms—including strong selective exposure, aggressive discounting, belief repulsion, and endogenous echo chamber formation—are not sufficient to create or maintain polarisation without biased information sources. This supports a 'supply-side' thesis: within the convergent dynamics of this model, motivated reasoning mechanisms were not sufficient to generate or maintain polarisation. Echo chambers and cognitive biases are *amplifiers*, not generators. The counterintuitive finding that selective exposure reduces polarisation from consensus supports Coppock's (2023) argument that these mechanisms may reflect rational heuristics rather than motivated biases.

For understanding and addressing political polarisation, these results suggest shifting attention from individual minds to information ecosystems. The question may not be 'why are people so biased?' but rather 'why does the information environment present them with such systematically different pictures of reality?'

References

Bak-Coleman, J. B., Lewandowsky, S., Lorenz-Spreen, P., Narayanan, A., Orben, A., & Oswald, L. (2025). Moving towards informative and actionable social media research. *arXiv preprint arXiv:2505.09254*.

Centola, D. (2011). An experimental study of homophily in the adoption of health behavior. *Science*, 334(6060), 1269-1272.

Cinelli, M., et al. (2020). The COVID-19 social media infodemic. *Scientific Reports*, 10, 16598.

Coppock, A. (2023). *Persuasion in Parallel*. University of Chicago Press.

Deffuant, G., Neau, D., Amblard, F., & Weisbuch, G. (2000). Mixing beliefs among interacting agents. *Advances in Complex Systems*, 3, 87-98.

DeGroot, M. H. (1974). Reaching a consensus. *Journal of the American Statistical Association*, 69(345), 118-121.

Epstein, J. M., & Axtell, R. (1996). *Growing Artificial Societies*. MIT Press.

Friedkin, N. E., & Johnsen, E. C. (1990). Social influence and opinions. *Journal of Mathematical Sociology*, 15(3-4), 193-206.

Hegselmann, R., & Krause, U. (2002). Opinion dynamics and bounded confidence. *Journal of Artificial Societies and Social Simulation*, 5(3).

Iyengar, S., et al. (2019). The origins and consequences of affective polarisation in the United States. *Annual Review of Political Science*, 22, 129-146.

Jern, A., Chang, K.-M. K., & Kemp, C. (2014). Belief polarization is not always irrational. *Psychological Review*, 121(2), 206-224.

Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin*, 108(3), 480-498.

Levendusky, M. (2013). *How Partisan Media Polarise America*. University of Chicago Press.

Lord, C. G., Ross, L., & Lepper, M. R. (1979). Biased assimilation and attitude polarisation. *Journal of Personality and Social Psychology*, 37(11), 2098-2109.

Madsen, J. K., Bailey, R., & Pilditch, T. D. (2018). Large networks of rational agents form persistent echo chambers. *Scientific Reports*, 8, 12391.

Mason, L. (2018). *Uncivil Agreement: How Politics Became Our Identity*. University of Chicago Press.

McPherson, M., Smith-Lovin, L., & Cook, J. M. (2001). Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27, 415-444.

Pilditch, T. D., Roozenbeek, J., Madsen, J. K., & van der Linden, S. (2022). Psychological inoculation can reduce susceptibility to misinformation in large rational agent networks. *Royal Society Open Science*, 9(8), 211953.

Rosenstock, S., Bruner, J., & O'Connor, C. (2017). In epistemic networks, is less really more?. *Philosophy of Science*, *84*(2), 234-252.

Taber, C. S., & Lodge, M. (2006). Motivated skepticism in the evaluation of political beliefs. *American Journal of Political Scienc*e, 50(3), 755-769.

Tappin, B. M., Pennycook, G., & Rand, D. G. (2020). Thinking clearly about causal inferences of politically motivated reasoning. *Current Opinion in Behavioral Sciences*, 34, 81-87.

Webster, S. W., & Abramowitz, A. I. (2017). The ideological foundations of affective polarisation in the U.S. electorate. *American Politics Research*, 45(4), 621-647.

Wood, T., & Porter, E. (2019). The elusive backfire effect: Mass attitudes' steadfast factual adherence. *Political Behavior*, 41, 135-163.

Young, D. J., Ackland, J. A., Kapounek, A., Madsen, J. K., Greening, L. J. & De-Wit, L. (2026) A new approach to measuring mass issue polarisation using *k*-means clustering, applied to a longitudinal analysis of US polarisation (1988-2024) and mapping patterns and predictors across the world, *Royal Society Open Science*

Young, D. J., Madsen, J. K., & de Wit, L. H. (2025). Belief polarisation can be caused by disagreements over source independence. *Cognition*, 259, 106126.