Essays on the Source Credibility Evaluation and Political Belief Formation under Social Contexts

by

Minsu Jang

Bachelor of Arts in Political Science, Dongguk University, 2011

Master of Arts in Political Science, Dongguk University, 2016

Master of Arts in Political Science, University of Pittsburgh, 2018

Submitted to the Graduate Faculty of

the Dietrich School of Arts and Sciences in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

University of Pittsburgh

2024

UNIVERSITY OF PITTSBURGH

DIETRICH SCHOOL OF ARTS AND SCIENCES

This dissertation was presented

by

Minsu Jang

It was defended on

May 00th, 2024

and approved by

Jonathan Woon, Department of Political Science

Michael Colaresi, Department of Political Science

Steven E. Finkel, Department of Political Science

Sera Linardi, Graduate School of Public and International Affairs

Essays on the Source Credibility Evaluation and Political Belief Formation under Social Contexts

Minsu Jang, PhD

University of Pittsburgh, 2024

The abstract of the document. This document is a sample file for the creation of ETD's at Pitt through LaTeX.

Table of Contents

List of Tables

List of Figures

1.0 Introduction

2.0 Why We Need A Computational Model of Political Learning

2.1 Introduction

Understanding how individuals acquire and process political information is crucial for comprehending the dynamics of political belief formation, its reinforcement, and belief polarization in democracies. Scholars have delved into the intricacies of the political learning process, aiming to uncover the mechanisms underlying how citizens form and update their political beliefs. This study contributes to this body of literature by proposing a series of agent-based models that address existing loopholes and offer theoretical insights into the comprehensive political learning process.

The existing literature on political learning predominantly focuses on two main themes: individual-level mechanisms and the effects of social networks. At the individual level, studies have explored how citizens process political information, demonstrating that people's beliefs are shaped by a combination of their prior convictions and new information obtained through social interactions. The echo-chamber effect, where individuals reinforce their attitudes by gravitating towards like-minded sources, is a prominent example illustrating this process. Additionally, the role of cognitive biases in information evaluation further influences individuals' information acquisition strategies and belief updating processes.

However, existing approaches to modeling political learning often face methodological and theoretical challenges, including the difficulty of capturing the cyclic nature of belief reinforcement and the intricate dynamics of social networks. To address these challenges, this project employs agent-based modeling as a powerful tool for theorizing the political learning process. Agent-based modeling allows for the simulation of complex individual behaviors and interactions within a system, making it well-suited for capturing the nuanced dynamics of political learning in real-world contexts.

Through a thorough literature review, this chapter outlines why a computational modeling approach is employed to theorize political belief formation and its reinforcement mechanism. I begin with a detailed explanation of existing approaches in the political information processing literature and identify the gaps that have been understudied. Subsequently, I illustrate why agent-based modeling is useful for filling in these gaps, emphasizing its ability to capture the complex, dynamic nature of political learning processes. Additionally, I highlight that this approach is not entirely new in the literature by outlining related computational models. Finally, I provide a brief introduction to the models proposed in this dissertation project in the following chapters.

2.2 Existing Approaches to Political Learning & Loopholes

The existing literature on the political learning process can be broadly categorized into two main themes. The first approach focuses on understanding the mechanisms through which individuals process political information at the micro-level. Studies within this approach have explored how people's political beliefs are shaped by a combination of their prior beliefs and the new information they receive through social interactions [115, 98, 100, 152, 47]. An exemplar of this process is the echo-chamber effect, wherein individuals reinforce their attitudes when consistently exposed to messages that align with their preexisting beliefs. Conversely, exposure to counter-attitudinal information can both weaken [47, 100] or strengthen [24, 52] their predispositions, depending on the strength of the listener's directional motives.

Another strand of research within the individual-level political learning literature investigates the reverse causal direction: how individuals' prior beliefs influence their patterns of political information acquisition. Studies employing this approach also confirm that individuals tend to rely more on information sources that confirms their existing attitudes [93, 64]. Moreover, some research suggests that certain individuals actively seek out counter-attitudinal sources to engage in argument refutation [159, 99, 163].

Importantly, synthesizing findings from both strands of research suggests that the political learning process operates cyclically, reinforcing preexisting beliefs over time. Figure 1 illustrates this comprehensive political learning mechanism, incorporating insights from prior literature. Studies emphasizing the role of cognitive biases have revealed that individuals often exhibit biased assessments when evaluating the credibility of information sources, in-

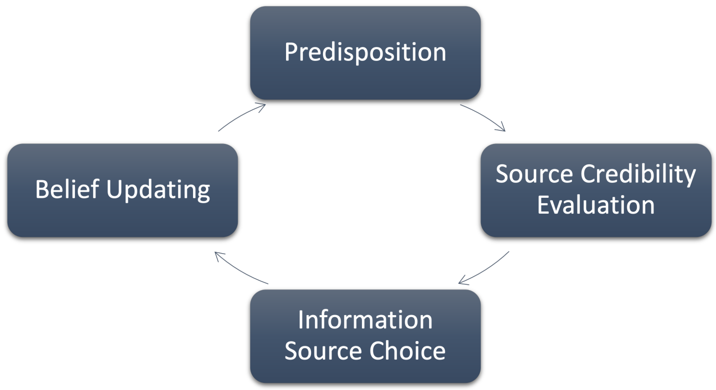

Figure 1: The cyclic nature of the micro-level political learning.

fluencing their subsequent information acquisition decisions. As individuals update and form new opinions, these updated beliefs serve as the prior for future learning steps, contributing to the continuation of the cycle of political belief reinforcement.

The cyclic nature of this mechanism presents both empirical and theoretical challenges in constructing comprehensive understandings of citizens' political learning processes. The reinforcement learning described above introduces endogeneity, making it difficult to establish one-directional causal path and develop effective empirical analysis strategies. Consequently, grappling with this inherent complexity in the political learning process has prompted scholars to approach research by deconstructing the mechanism and concentrating on individual steps within the cycle.

The second theme explored in existing literature on the political learning process revolves around the effects of social networks. In the process of deliberating with fellow citizens, individuals not only absorb information from better-informed sources but also function as conduits, disseminating this information to others. For example, friendship networks often establish social norms among their members, fostering consensus of opinion within the network [107, 146, 72, 38]. Similarly, active social interactions are expected to facilitate exposure to a variety of viewpoints from different groups.

Another intriguing aspect of social networks, as briefly discussed previously, is their self-selective nature. As described in Figure 1, citizens can select information sources based on their assessments of credibility. In other words, individuals can choose whom they communicate with among their neighbors. Consequently, cognitive biases can influence citizens' preferences for information channels, consequently impacting the fragmentation of communication channels.

2.3 Agent-Based Models as Alternatives

In this study, I employ a computational modeling approach to conceptualize the political learning process within a social context. As discussed in the preceding section, the political learning process exhibits a cyclic nature, reinforcing an agent's preexisting attitudes. Agent-Based Modeling (ABM) is a valuable tool for investigating such complex individual behaviors. ABM, a specific approach within computational modeling, focuses on representing individual entities, or agents, and modeling their interactions within a system.

In ABMs, agents' behavior is defined by sets of rules, which can range from simple to highly intricate [59]. The recursive nature of the cyclic mechanism depicted in Figure 1 implies that past behaviors determine future actions [109, 143]. Moreover, its self-reinforcing aspect introduces self-adaptivity, enabling citizens to swiftly adjust strategies in response to changes in their counterparts' behavior to enhance their learning about the world [128]. ABMs excel in exploring such complex behavioral rules by accommodating multiple processes simultaneously [41]. While existing literature has examined the political learning process by dissecting the cycle into discrete components due to methodological concerns outlined in the preceding section, ABMs offer avenues to indirectly delve into the comprehensive mechanisms of political reinforcement learning more realistically through simulation outcomes [137, 134].

Another useful aspect of ABMs is their ability to incorporate agent heterogeneity [41, 109, 134]. This heterogeneity not only includes diverse belief distributions among agents but also variations in behavioral rules based on defined characteristics. Previous studies suggest that individuals can have different incentives in the learning process, and their behavior may be contingent on various factors. The inclusion of heterogeneity can thus enrich the model, leading to unexpected emergent patterns stemming from dynamics of micro-level behavior [136, 144].

Furthermore, researchers employing ABMs can exploit their methodological strength by introducing dynamics between agents. In repeated political learning processes, citizens mutually influence each other. This interdependence and adaptability in the political learning process make ABMs particularly useful for theoretically modeling mechanisms and exploring outcomes. ABMs also allow agents to interact with collective entities such as organizations [4, 66] or macro-level systems they reside in [7, 4]. Thus, ABMs provide an effective exploration of environments that accompany the dynamics of components that construct the system.

Researchers can also incorporate spatial attributes into ABMs, where agents exhibit social dynamics. These spaces can be geographic or social environments. For instance, [42] replicated the spatio-temporal demographic history in Long House Valley in northern Arizona. Additionally, modelers can include networks that determine social interaction channels [104, 134]. As previously discussed, the existing literature on political learning emphasizes the social network effect. With an ABM, one can establish an initial network structure to trace the exchange of information and observe changes in communication channel selection over time as agents build beliefs about the world and assess the credibility of communication partners.

Lastly, ABMs offer the capability to track agents' behavioral changes over longer time spans. The primary challenge in empirically studying reinforcement political learning lies in the difficulty of repeatedly observing each individual's behavior. Moreover, in real-world contexts, citizens consolidate their beliefs through social interactions and select information sources based on their learning outcomes. ABMs enable us to trace the temporal process of belief formation and the evolution of social interaction patterns over time. Furthermore, the tractability of ABMs facilitates the examination of how information disseminates within social networks and how competing pieces of information interact.

2.4 Related Models: Opinion Dynamics and Social Learning

2.4.1 Opinion Dynamics and Social Learning

Numerous attempts have been made to theorize and explore opinion formation dynamics using ABMs. One of the most well-known classical ABMs of opinion dynamics is the Degrootian approach [44]. In this model, agents update their beliefs by interacting with others in an attempt to reach a consensus. Belief updating is achieved through a weighted averaging function that combines an agent's prior belief with messages from other agents. These weights are exogenously given and invariant. The advantage of this model lies in its straightforward comprehension using matrix theory and Markov chains [118]. Some studies have modified the model by treating opinions as discrete variables [157, 151, 12].

The strength of the Degrootian approach stems from its ability to provide a useful framework for explaining how social consensus is built. These models, in particular, have proven effective in illustrating how agents adopt other opinions over a shorter time span. However, given the prevalence of social disagreements and belief segregation, other works have attempted to model the social learning process with more constraints to achieve social agreement. One such constraint is bounded confidence, which posits that an agent only accepts another agent's message if it falls within its acceptability boundary [43, 82, 169]. This approach facilitates theoretical exploration into how fractionalized post-communication opinion distributions are shaped given the acceptability of opinions.

[117] adopts Bayesian inference for belief updating in social learning. Under this framework, each individual agent holds an opinion represented by a normal distribution, where the mean indicates the agent's estimation about the state they are trying to guess, and the standard deviation denotes the uncertainty the agent holds. The model implies that uncertainty can serve as a threshold that determines the number of opinion clusters in the end: the lower the uncertainty each agent possesses, the more fractionalized the post-communication opinion distribution becomes. While Martins' model assumes that uncertainty is not shared and each agent expects others to hold the same level of uncertainty, [2] extends the model further by relaxing the "unshared uncertainty assumption." By allowing agents to infer each

other's uncertainty levels, their model illustrates that mistrust plays a fundamental role in driving polarization, although its predecessor only indicates its impact on the speed at which agents converge to the steady-state (i.e., social consensus or polarization).

2.4.2 Source Credibility in Social Communication

The aforementioned models shed light on how individuals access information originating from various sources. In the Degrootian framework, for example, the weight assigned to each source signifies the degree to which an information receiver values messages from the speaker. Similarly, in bounded confidence models, the confidence bound indicates the level of acceptability an agent maintains towards different messages. Moreover, the concepts of uncertainty and (mis)trust [117, 2] elucidate how each agent discerns one source from another in the information processing procedure. Likewise, [46] model opinion formation through social communication, categorizing statements into discrete classifications based on the listener's perceived agreeability towards each statement.

Understanding the credibility assessment mechanism is crucial for comprehending how people select information sources and its consequences on social dimensions, such as belief distribution and social connections. Previous studies on echo chambers have emphasized cognitive bias as the essential driver of segregated communication channel emergence, thereby deepening the echo chamber [45, 145, 112, 150]: as people tend to dislike hearing arguments that contradict their pre-existing beliefs, they naturally communicate more often with like-minded neighbors, ultimately creating segregated communication networks through repeated interactions. Conversely, [67] argue that cognitive bias is not the sole factor creating echo chambers. Instead, the study suggests that echo chambers could arise from lateral transmission of information combined with limited access to information, and the presence of communication partners with strong knowledge can amplify echo-chamber effects. Nevertheless, their work still emphasizes that cognitive bias can consolidate echo-chambers.

2.4.3 Diverse Agent Types and Social Dynamics

Under social dynamic environments of opinion formation, one important factor determining the flows of information is the definition of agents' behavior. Specifically, understanding how interactions among diverse types of agents impact micro and macro-level opinion formation is of particular interest. Several endeavors have been made to incorporate special types of agents into existing models, including experts, extremists, and strategic politicians. For example, [3] introduce informed agents into social networks and discovered that better-informed agents could steer social beliefs in desirable directions even through microscopic interactions with neighboring agents. Conversely, as previously mentioned, [67] illustrate that experts could act as gatekeepers, creating segregated opinion channels and fostering local filter bubbles. Additionally, [142] introduce media as a special category of elite information providers, playing pivotal roles in shaping the opinion formation of other ordinary agents.

Another strand of research explores the impact of extremists on social opinion formation. Overall, these studies indicate that agents with strong opinions often hinder the building of social consensus. For instance, [2] demonstrate that the inclusion of extreme agents tends to fragment agents' belief distribution, even in environments where agents have high levels of social trust, which indicates the openness to diverse opinions, including counterarguments to their prior beliefs. Similarly, [142] show that biased media can foster polarization even when agents engage in town meetings with all other agents. They found that moderate-sized town meetings with polarized media initially lead to polarized divergence, followed by eventual convergence towards an extreme media-defined pole of opinion. Likewise, [140] highlight that stronger arguments mainly contribute to polarization[1], while strong informational biases are necessary for bipolarization in the absence of homophily.

Finally, [113] introduce a politician aiming to win elections by persuading electorates with micro-targeted messages. The model simulates micro-targeted political campaigns by incorporating voters, politicians, and persuasive connections between them. Voters hold beliefs and care rankings about political issues and perceive politicians' credibility, enabling

[1]They define polarization as "a uniform change of the opinion of the whole group towards the same direction."

them to revise their beliefs when contacted. Politicians, in turn, aim to persuade voters about their positions on various issues by contacting them during the campaign and transmitting persuasive attempts, prompting voters to update their beliefs accordingly. In their following work, [139] extend their model of micro-targeted campaign strategies by introducing additional heterogeneity and sophistication to voter representation, allowing voters to consider multiple political issues and rank their importance. This extension enables the simulation of various voter types, ranging from single-issue voters to highly engaged ones. The simulation results demonstrate the effectiveness and efficiency of micro-targeted campaigns in navigating complex voter landscapes, with micro-targeted campaigns consistently outperforming stochastic counterparts, particularly in heterogeneous electoral environments.

2.5 Brief Introduction to This Study's Models

In this dissertation project, a series of agent-based models is proposed to address the aforementioned loopholes in existing approaches and to theoretically model the political learning process comprehensively. First, the Baseline Model outlined in the next chapter explores the individual-level mechanisms of political learning illustrated in Figure 1. Unlike previous studies that have dissected the reinforcement political learning mechanism step-by-step, the baseline model combines the process into one cyclic algorithm. The *Citizen* agents in the model aim to learn the state of the world by drawing messages from elite *Information Provider* agents. When they sample messages, they 1) assess the credibility of each source, and 2) determine their sampling strategy probabilistically reflecting the source credibility assessment results. The behavioral rules of agents are defined employing preexisting methods from the literature, particularly biasedly evaluating source credibility, as suggested by cognitive bias literature. Ultimately, the baseline model chapter examines how citizens' biased assimilation of political information affects macro-level outcomes including 1) social political belief distributions, 2) choice of information outlets, and 3) how quickly agents' beliefs reach the steady-state.

The Social Network Model extends the Baseline Model by introducing social dynamics between *Citizen* agents. While the baseline model restricts social interaction to dyadic conversations between a citizen and each information provider, the Social Network Model aims to explore how diverse social network structures affect macro-level belief distribution. This chapter investigates social dynamics under three types of social network structures. First, it compares the baseline simulation results with a model where each citizen is randomly matched with two other agents regardless of their types. Then, simulations are run under a homophilous environment, where each agent has a group identity conditional on its initial belief and is more likely to be connected with in-group members. The final variation allows citizens to communicate with both citizens and information providers, combining the network structures of the baseline model and the random 2-neighbor matching. The goal of this chapter is to compare how social belief distribution changes as the degree of homophily and accuracy setups of elite information providers vary.

The final model introduces strategic agents into both the baseline and social network models. In this chapter, a strategic information provider replaces an information provider and competes with an unbiased optimal information provider. The strategic agent aims to disrupt citizens' social learning about the true state of the world by sending jamming messages with respect to the optimal messages provided by its unbiased rival. The outcomes of interest in this chapter are the social belief distribution and the speed of belief stabilization. Simulations are run under both the baseline model and social network environments. This chapter aims to theorize how adversarial disinformation providers behave and the social consequences of adversarial messages. To model this behavior, the micro-targeted messaging approach is adopted: the disruptive message provider can partially observe the belief distribution of citizens with clustered points of view and tailors messages to members of each cluster to distract targeted citizens from discerning the true state of the world.

3.0 Baseline Model: Credibility Learning, Information Source Choice, and Political Belief Formation

3.1 Introduction

Social scientists have extensively researched how individuals form their attitudes towards real-world political issues, with the political learning process being particularly complex. When making political decisions, individuals take into account various factors, including their pre-existing beliefs, the opinions of others, the credibility of information sources, and which sources to trust. These factors are closely intertwined, as research has demonstrated that an individual's perception of an information sender's credibility is linked to their updated beliefs about the world. Specifically, the more a receiver trusts the source, the more their updated beliefs align with the information sender's message.

Political predisposition is a crucial aspect of political learning, and prior research has shown that a person's pre-existing beliefs can influence their selection of information sources. According to [159], individuals who possess a deeper understanding of an issue are more inclined to search for arguments that contradict their pre-existing beliefs, while less informed individuals are more likely to seek out information that aligns with their existing attitudes. Moreover, [152] found that individuals with stronger political predispositions are more likely to choose political media outlets that share their viewpoints.

Recent work by [108] suggests that people learn about the status of the world and the credibility of the information source simultaneously after receiving a message from the sender. Furthermore, his model illustrates that the receiver's directional motives can distort their beliefs about the message's credibility. Although Little's model focuses on the individual level, this study expands the scope to include social contexts. The model proposed here demonstrates how 1) political predisposition affects the learning of information source credibility, and 2) how this mechanism reinforces beliefs about the world over time. In other words, this study addresses the question of how learning about source credibility and beliefs about the world affect the evolution of media consumption in a competitive political news media market. Additionally, this model delineates how the evolution of media consumption affects public knowledge about the state of the world.

My model integrates existing models of political learning with reinforcement learning from the machine learning literature. Reinforcement learning describes how agents can learn to make optimal decisions through repeated experiences. One well-known example of reinforcement learning is the Multi-armed Bandit problem (MAB problem). In this problem, an agent must choose between alternatives to maximize their gains. For instance, imagine a gambler standing in front of two slot machines. The gambler must decide which machine to play first, how many times to play each machine, and whether to continue playing the same machine or switch to the alternative. In reinforcement learning, each agent's current action not only determines their current payoff but also affects future states, such as the choice of machine and the rewards that follow those choices over time.

Reinforcement learning can be applied to the scenario where people gather political information and update their political beliefs. For instance, suppose a citizen wants to learn which candidate will perform better in dealing with economic issues in an election. The person needs to 1) decide from where to acquire relevant information, 2) evaluate how trustworthy the information collected from each source is, and 3) choose whether to use the new information to update their beliefs about the candidates. Thus, the political learning process is working as a cyclic mechanism, rather than a clear one-direction path. Since this mechanism is well-captured in the reinforcement learning procedure, it is worth combining the pre-existing knowledge of political information processing with reinforcement learning models.

The shared characteristics between the political learning process and the reinforcement learning mechanism raise both theoretical and methodological concerns. Theoretically, the cyclic nature of political learning makes it challenging to clearly identify the causal direction. Additionally, the cyclic causal path introduces endogeneity issues in real-world observations. Thus, this study adopts a computational modeling approach using agent-based simulation—an effective tool for modeling cyclic mechanisms and generating plausible observations in such environments.

My model proposes that individuals' perceptions of message bias and message precision are critical factors in evaluating the credibility of an information source. Therefore, the agent must learn both bias and precision by sampling messages from the outlet. To begin, I endogenize the sampling mechanism with credibility assessment strategies: people first assess the credibility of information sources and then decide how many messages they want to listen to from that source. In other words, the number of samples drawn from a source depends on how much the agent believes the outlet is superior to other alternatives. I then compare how different credibility assessment methods affect information sampling patterns and people's beliefs about the world.

My theoretical model provides two valuable perspectives. First, it extends our understanding of biased information processing from the individual level to the social context. While previous studies have explored how an individual's preexisting beliefs and slanted media outlets can affect their political decisions, my model takes it one step further by examining how media consumption networks evolve over time based on an agent's credibility evaluations at each point in time. This provides a broader view of information consumption patterns and the resulting belief distributions among all agents. Moreover, my model can also be applied to the consumption of fake news, providing insights into when and how fake news providers may be more likely to survive within media diets and the impact this has on the overall distribution of beliefs among agents. The abstract nature of the model allows for analysis of a wide range of scenarios and circumstances, providing valuable insights into the dynamics of information consumption in competitive political news media markets.

3.2 Related Literature

The preexisting literature on political learning has explored two mechanisms: 1) how people assimilate new information, and 2) how preexisting attitudes affect the choice of information sources. In this section, I summarize the findings of preexisting studies illustrating these two mechanisms. Additionally, as my model attempts to endogenize the choice of information source based on previous experience, I review existing models of reinforcement learning that delineate the optimal arm choice mechanism.

3.2.1 Information Acquisition and its Effect on Political Attitude

People learn about the status of the world based on the information they acquire from various sources. They combine the new information with their prior knowledge about the true state of the world, so the messages they receive from sources can either confirm their preexisting beliefs or provide counterarguments. People's reaction to confirming messages is straightforward - they strengthen their prior beliefs. Previous studies found that people who selectively exposed themselves to politically congruent media sources uniformly strengthen their predispositions [115, 98, 100, 152, 47]. However, some scholars have raised concerns about the echo chambers created by biased exposure to like-minded messages, as it can deepen political polarization [90, 155, 154, 68, 63, 65].

The story becomes more complicated when new information is not consistent with prior beliefs. [103] argue that exposure to contradictory information can persuade partisans to vote for the rival party. Similarly, other studies found that counterargument exposure led people to deviate from their own party [47] and weaken their prior beliefs [100]. However, some people use counterattitudinal information to refute rival arguments and strengthen their prior beliefs eventually, leading to the "backlash effect" [79, 122, 16, 58]. They argue that when people encounter new information inconsistent with their prior beliefs, they start to discredit it [48, 158] and believe that they have refuted it [106]. Recent studies also found that exposure to out-partisan media sources activated people's directional motives, resulting in the strengthening of their predispositions [24, 52]. However, the opposite effect was observed when subjects engaged in person-to-person conversations with out-partisans: the post-communication attitude became weaker than the pre-communication attitude.

Democratic citizens are exposed to new political information every day, such as news about the economy, policy debates, or politicians' behavior. Many political scientists agree that citizens learn new information by combining their preexisting beliefs with new signals acquired from information sources [172, 14, 71, 159, 84]. Bayesian learning provides a mathematical formulation of such belief updating that citizens face every day. As Bayesian updating illustrates, after looking at new information, people evaluate the newly acquired information and combine it with their preexisting beliefs.

Although Bayesian updating is a useful tool to understand political learning among citizens, it has been criticized for being an unrealistic model that does not account for cognitive biases prevalent in political information processing literature. Studies have shown that people often process new information in a biased manner that reflects their preexisting beliefs [101, 159, 158, 133]. For example, [27] explains that party identification creates a "perceptual screen" that allows only messages from one's own party. Similarly, [172] argues that citizens filter out information that contradicts their preexisting beliefs.

More recent works attempt to incorporate cognitive biases into the Bayesian updating framework. For instance, [71] proposes a Bayesian interpretation of biased assimilation in political learning, which is discussed in a later section of this paper. In the model, a citizen evaluates the persuasiveness of a signal message based on the comparison between their posterior beliefs and the signal. [108] also illustrates how directional motives can distort a citizen's evaluation of the credibility of an information source. Interestingly, these attempts still suggest that information receivers are not "unmoved by it" [71]. As [84] found, even though people are not perfect Bayesians, they still learn in the appropriate direction, even with directional motives.

3.2.2 Political Predisposition Affecting Information Source Choices

Another body of literature exploring the relationship between political information exposure and political attitudes seeks to answer the question of how political predisposition affects people's information source choices. There are two interesting patterns of information source choices: congruent message pursuit and counterargument exposure. The first pattern, congruent message pursuit, is based on the argument that people tend to avoid counter-attitudinal messages and want to expose themselves to messages that are consistent with their preexisting attitudes [93, 64]. Previous studies have shown that the strength of party identity is positively correlated with politically motivated selective exposure: Stronger partisans are more likely to pursue attitude-consistent information [152, 153, 89, 99, 13, 37].

The more interesting information exposure pattern is counterargument exposure. [130] found that cross-cutting exposure was more common in media usage than interpersonal

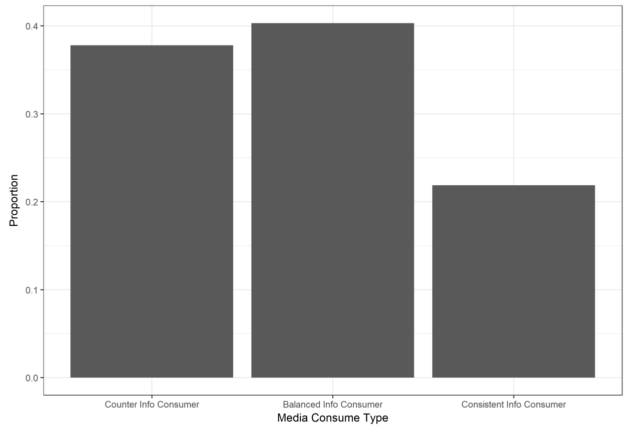

Figure 2: Media Consumption Type Distribution (NAES 2004 & 2008).

communication. Figure 2 illustrates the distribution of media outlet choice patterns using the National Annenberg Election Survey (NAES) conducted in 2004 and 2008. Approximately 40.3 percent of the respondents are balanced media consumers who consume an equal number of congruent and counterargument media outlets. The second most frequent category is, interestingly, counterargument consumers. They consist of about 37.8 percent of the sample, while only 21.9 percent of the respondents are categorized as consistent media consumers.

There are various reasons why people may be exposed to counter-attitudinal information. One reason is that people may unintentionally choose counterargument sources because they are not skilled in assessing the biases of media outlets [164, 40].

Another reason for counterargument exposure is intentional. People who have moderate levels of party attachment (i.e., weak partisans) are more likely to be open to counterarguments. Moreover, people with stronger motivations, such as strong party attachment or higher levels of anxiety, may seek out counter-attitudinal information to defend their own argument from rival arguments [159, 99, 163].

3.2.3 Reinforcement Learning And Optimal Arm Choice

As previously explained, reinforcement learning illustrates how agents might learn to make optimal decisions from repeated experiences [156]. The optimal action refers to the policy that can maximize the utility. Since the action taken in the previous time step affects my action and rewards, the model satisfies some Markov Chain properties [32].

The multi-armed bandit (MAB) problem is a special case of reinforcement learning that posits the situation of how limited resources will be allocated among alternatives. In the decision making of allocation, people face trade-offs between the expected gains of relying on the channel of which performance is already known and trying a new alternative so they can learn the alternative's rewarding parameter. The former refers to exploitation, while the latter is called exploration.

One simple example is illustrated in [17]. In this example, suppose that a person makes a choice between 2 coins. Let's call each coin's Bernoulli parameter θ_1 and θ_2, respectively. In addition, it is known that coin 1 is fair (i.e., $\theta_1 = \frac{1}{2}$), and coin 2 is either 1 or 0 (i.e., it is either two-headed or two-tailed), with probability π and $1-\pi$. The player aims to maximize the expected number of heads in n flips. Here, the gambler has to make a choice between 1) a myopic strategy that flips coin 1 because it is already known (i.e., exploitation), or 2) taking the risk of exploring coin 2 to learn the probability parameter (i.e., exploration). Since its first appearance in the literature [161], many attempts have been made to explore this problem using various types of reward distributions [22, 166, 62, 6, 1, 74, 75, 73, 165, 8].

While the canonical form of MAB problems assumes that the reward function parameter for at least one machine is known while the others require exploration, another type of conflict a person may face in MAB problems is when there is no prior knowledge on reward distributions. Under this scenario, all actions taken by the player are purely exploratory. In addition, the player's rewards are data that are used to learn each machine's reward function. [33] provides a clear example of pure exploration under a clinical trial. In this example, there are two drugs to treat a disease, but the probability of curing the disease is unknown: call it θ_1 and θ_2 for drug 1 and drug 2, respectively. The experimenter faces the problem of finding the more efficient drug with the minimum number of trials. Thus, under pure exploration, the experimenter needs to rely on an inference strategy to build a rule for trial-stopping points where sufficient accuracy is detected and no more experimentation is needed [5, 114, 23].

Reflecting Whittle's articulation [170] that such resource allocation conflicts are evident in all human actions, MAB problems have been applied in various fields. MAB problems can be used in studies exploring communication within social networks and social learning [173]. In deliberation networks, for instance, people need to make choices about which channel they want to communicate through, or whether they want to pass on the argument they received from the previous counterpart to another neighbor. Despite the MAB's general applicability in studies on communication and learning, few attempts have been made so far. [11] use a two-armed bandit problem to model social learning where each agent can share their experience of trials and rewards with their neighbor. They found convergence of payoffs and choice of actions in the long run. Similar effects were also found in [76, 26, 167].

Thus, reinforcement learning can be applied to the scenario where people gather political information and update their political beliefs. As previously mentioned, in the political learning process, people should decide 1) which source to rely on to acquire useful information, and 2) how much to reflect newly acquired information to update their beliefs. Furthermore, throughout the learning process, agents accumulate knowledge on the credibility of each source. In the next section, I propose an original model of reinforcement political learning that combines the aforementioned three bodies of literature.

3.3 Baseline Model: Credibility Learning and Bayesian Belief Update

This paper employs a generative social science approach to illustrate the reinforcement learning procedure over time. Here, I argue that the mechanism of individuals' political learning is cyclic: The predisposition affects how people acquire information; and people learn about the world in a Bayesian way. Therefore, it is crucial to accumulate the observation of each individual's behavior over time. Agent-based models are well-suited for this purpose since they generate both individual-level information acquisition history and macro-level phenomena that result from individual agents' interactions [59].

An agent-based model provides an appropriate platform to delineate the mechanism that traces 1) how an agent learns the credibility of information sources reflecting its own prior belief, 2) how the updated credibility affects its information acquisition methods, and 3) how the information-acquiring pattern determines the entire society's belief distribution about the state of the world. Previous literature treated the information consumption network structure as a given and fixed parameter, partly due to difficulties in observing network evolution. We can only observe the network at present, and the history of evolution can only be traced if studied over time. However, agent-based modeling can address this problem due to the aforementioned characteristics of the method.

3.3.1 Agents

First, I begin with defining the attributes of agents residing in this model. Let $n \in 1, N$ represent the N agents in a social network g, where $(i, j) \in g$ indicates that agents i and j are connected in the network. The agents in this model are attempting to learn about the true state of the world θ.

There are two types of agents. The first type is the *Information Providers*, which represent the political information sources such as politicians, media outlets, or social influencers. They can directly observe the true state of the world with noise. The noise consists of their own bias δ_i, representing the bias of *Information Provider* i, and different levels of precision parameter σ_i. The belief of *Information Provider* i about θ is normally distributed $\theta_i \sim N(\theta + \delta_i, \sigma_i^2)$. They send out messages to *Citizen* agents so they can learn the true state of the world. In other words, the messages that *Information Provider* i sends to *Citizen* j, x_{ij}, are also drawn from a normal distribution $x_{ij} \sim N(\theta + \delta_i, \sigma_i^2)$. However, *Information Providers* do not receive messages from other agents.

The second type of agent is *Citizens*, who update their beliefs β_j by receiving messages from *Information Providers*. They cannot see where the true state of the world is, and can only communicate with *Information Providers*. The behavioral rules for *Citizens* are defined as follows.

3.3.2 Belief Updating of *Citizens*

As previously explained, *Citizen* agents aim to learn the belief about the true state of the world (θ) by sampling messages from *Information Providers*. Also, during the learning process, *Citizens* have to update their beliefs about the credibility of information sources. Thus, at each step, *Citizens* have to choose which beliefs they want to update before beginning communication —beliefs about the true state of the world or credibility of a source. When they are learning their beliefs on the state of the world (θ), they update their beliefs in Bayesian ways. The posterior beliefs on the state of the world conditional on the sampling result at time t is $\mu^{\theta}_{i,t}|x_{i,t} \sim N(\widehat{\mu^{\theta}_{i,t}}, \sigma^2_{\theta_{i,t}})$; where $x_{j,t}$ is the sampling result that a *Citizen* agent acquired from *Information Provider* j at step t,

$$\widehat{\mu^{\theta}_{i,t}} = \mu^{\theta}_{i,t-1} + (x_{i,t} - \mu^{\theta}_{i,t-1}) \left(\frac{\sigma^2_{\theta_{i,t-1}}}{\sigma^2_{\theta_{i,t-1}} + \sigma^2_{x_{j,t}}} \right) \tag{1}$$

and

$$\sigma^2_{\theta_{i,t}} = \frac{\sigma^2_{\theta_{i,t-1}} \sigma^2_{x_{j,t}}}{\sigma^2_{\theta_{i,t-1}} + \sigma^2_{x_{j,t}}} \tag{2}$$

where $\sigma^2_{x_{j,t}}$ is the precision of new messages. The credibility assessment methods are described in the following paragraphs.

Figure 3 illustrates how *Citizens* behave in this model. In this study, they make a decision randomly at each time step t. Let $c \in 1, 0$, and $c = 1$ is the decision for learning credibility, and $c = 0$ means the agent learns about the state of the world. The probability of choosing credibility decreases as time goes on, assuming that agents want to minimize cognitive efforts if they have enough previous knowledge about the credibility of an information source. Thus, I set the probability of choosing credibility learning to be the reciprocal of time t: $Pr(c) = \frac{1}{t}$.

Each *Citizen* samples messages from *Information Providers* using the credit, R. Each sampling costs 1 credit, and they can sample R number of messages in total at each step t. If they choose to update the credibility of each information source (i.e., $c = 1$), they assign the resources equally to each information source. For instance, if $R = 10$ and there are two *Information Providers*, they sample 5 messages from *Information Provider* A and 5 messages from the other. Based on the sampled messages, they compare the credibility of each source and conclude the better-performing outlet. Here, I employ two ways of

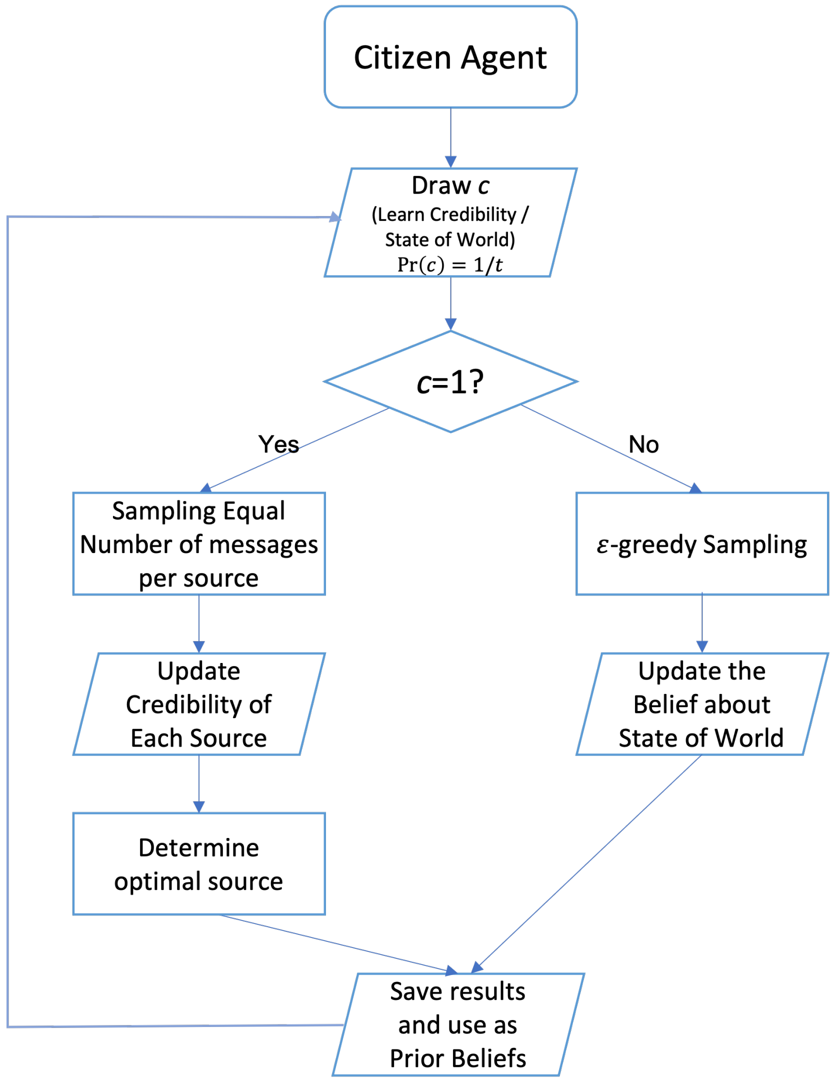

Figure 3: Citizen Agent Behavior Algorithm.

credibility assessment from existing literature, which are based on motivated reasoning in source credibility evaluation: z-Statistics Comparison [71] and δ Comparison [108].

3.3.2.1 Source Credibility and Learning Strategies

As aforementioned, I argue that people learn the state of the world by communicating with available information sources. This process consists of two sub-steps: 1) how plausible the new information is and 2) updating the belief based on the credibility evaluation. As the information acquisition process is represented as drawing numbers in this model, at each step of interaction, *Citizens* can infer the communication partner's belief parameters from the samples. From the sampling theory, the sample accuracy consists of precision and bias. Thus, *Citizens* evaluate each *Information Provider's* credibility based on the sample accuracy.

The precision is a straightforward indicator for the credibility judgment. Concise sample distribution is preferred because 1) the information source has concrete beliefs on the state of the world, and 2) the samples from future interactions are also expected to be concise.

The bias, on the other hand, is the more complicated attribute in the credibility inference process. From the motivated reasoning literature, it is known that people tend to conclude that attitude-congruent information is more believable than counterarguments. Reflecting preexisting literature, this model is built upon two key assumptions. First, the bias is subjectively evaluated. As *Citizens* have initial beliefs about the state of the world, they conclude that an information source is less biased when it sends out messages that are more similar to their prior beliefs. Second, *Citizens* always learn from communication partners' messages assuming that they can be better informed from the social interaction.

This study employs two credibility judgment strategies from existing studies for *Citizens'* optimal outlet choices. The first strategy is called δ *Comparison*. Under this strategy, *Citizens* focuses solely on the biases of messages samples from each information outlet: they prefer attitude-congruent messages and conclude that the *Information Provider*, whose messages are more similar to their pre-existing beliefs about the world, is more accurate and reliable. In other words, they believe the more similar message outlet is the best outlet and rely more on the outlet for the future belief updating process. The second strategy, contrarily, *Citizens* are concerned not only with the perceived bias of the sample messages, but also with how precise the messages are. Thus, they utilize both sample bias and precision for the source optimality evaluation. Details about each strategy are described as follows. Again, *Citizens*' credibility judgment is incurred from their preferences on each *Information Providers*. The information provider, who is believed more credibile, is preferred in the future sampling process.

δ ***Comparison Strategy***

Previous research in political communication has shown that United States citizens believe in the existence of pervasive bias in political information outlets [164, 168, 89, 129, 60]. Moreover, individuals' political predispositions significantly affect how they evaluate news quality [138, 108]. [108] provides insights into how people process political information while assuming the possibility of bias in the signal.

Little's model assumes that the signal x is composed of three elements: the true state of the world (θ), the signal's bias (δ), and the idiosyncratic error (ε): $x = \theta - \delta + \epsilon$. Consequently, assuming that the signal can be biased, the agent becomes interested not only in learning the state of the world but also in detecting the bias (δ) in the signal.

Assuming that the three elements of the signal are independent and normally distributed, we can obtain the mean of the bias, μ^{δ}, using the following equation:

$$\mu^{\delta}_{j,t}|x_{j,t} = \frac{\mu^{\delta}_{j,t-1}(\sigma^2_{\theta_{i,t-1}} + \sigma^2_{\varepsilon}) - (x_{j,t} - \mu^{\theta} i, t-1)\sigma^2 \delta_{j,t-1}}{\sigma^2_{\theta_{i,t-1}} + \sigma^2_{\delta_{j,t-1}} + \sigma^2_{\varepsilon}} \tag{3}$$

Additionally, we can obtain the posterior variance of δ_j beliefs using the following equation:

$$\bar{\sigma}^2_{\delta_j,t} = \frac{\sigma^2_{\delta_{j,t-1}}\sigma^2_{\varepsilon} + \sigma^2_{\delta_{j,t-1}}\sigma^2_{\theta_{i,t-1}}}{\sigma^2_{\theta_{i,t-1}} + \sigma^2_{\delta_{j,t-1}} + \sigma^2_{\varepsilon}} \tag{4}$$

Here, δ_j refers to the bias of *Information Provider* j and x_j is the message obtained from agent j. $\bar{\sigma}^2_{\delta_j}$ is the posterior variance of δ_j beliefs, and σ^2_{θ}, $\sigma^2_{\delta_j}$, and σ^2_{ε} are the prior variance beliefs on the state of the world, *Information Provider* j's bias, and the error term, respectively.

This means that when an agent receives a signal that is far from its prior belief about the state of the world, it will conclude that the information contains a significant amount of bias. Similarly, when the information receiver's prior on the state of the world is precise (i.e., σ_θ is low), it perceives greater bias from the new information.

When there are multiple information sources available to the *Citizens*, they can use this method to compare the biasedness of each *Information Provider* with absolute values of $\mu^B_{\delta_j}|x_j$. Since the bias can have either negative or positive values, the credibility assessment depends on the biasedness in general rather than focusing on the bias itself. If $|\mu^B_{\delta_A}| > |\mu^B_{\delta_B}|$, then *Information Provider B* is a better source. Finally, they can keep track of updated $\mu^B_{\delta_j}$ and $\bar{\sigma}^2_{\delta_j}$, which enables the Bayesian learning of the biases of information sources in future learning processes.

z-Statistics Comparison Strategy

[71] propose a model of Bayesian biased assimilation, in which people selectively accept information that confirms and strengthens preexisting beliefs. In their model, people assimilate new information after considering 1) how much it is different from (or similar to) their own prior beliefs, and 2) how precise the new signal is. In this model, after receiving a message, they estimate the probability that the posterior beliefs on the state of the world are equal to or greater than the signal, $1 - \Phi(z)$,

$$z = \frac{|x_{j,t} - \widehat{\mu^{\theta}_{i,t}}|}{\sqrt{(\sigma^2_{x_{j,t}} + \sigma^2_{\theta_{i,t}})}}\left(\frac{\sigma_{\theta_{i,t}}}{\sigma_{x_{j,t}}}\right) \tag{5}$$

where $x_{j,t}$ is the message sent from *Information Provider j*, $\sigma_{x_{j,t}}$ is the standard deviation of the message sent from j at time t. Thus, when z is too large, the signal becomes unconvincing. This model implies that people are more likely to accept messages with lower z values.[1]

[1]In the original equation [71], it does not take the absolute value for the difference between the new message $x_{j,t}$ and the posterior beliefs $\widehat{\mu^{\theta}_{i,t}}$. This study, however, takes the absolute value for this term to represent the degree of deviation from the preexisting beliefs better.

In this study, when learning the credibility of *Information Providers*, *Citizens* compare the z statistics from each source and conclude that the source with lower z is better performing. They also become more likely to sample more messages from the source when learning the state of the world.

3.3.2.2 Learning θ with Endogenous Sampling Strategy

The ε-greedy strategy is a popular approach to balance exploration and exploitation in multi-armed bandit problems, where one needs to decide between different actions with uncertain rewards [165, 162, 149]. It involves exploiting the best-performing option most of the time, while occasionally exploring with a globally controlled probability parameter (ε) that determines the choice of actions.

In this study, I adopt an endogenous sampling strategy that is based on the ε-greedy method. Before each attempt to sample a message, a *Citizen* randomly draws a number from a uniform distribution. If the number equals or exceeds ε, the citizen draws a message from the more credible outlet. Otherwise, the agent samples a message from the less preferred one. Thus, ε represents how much an agent is open to getting information from a sub-optimal source. The following pseudo-code describes the algorithm:

```
epsilon = 0.05
p = random.uniform()
if p < epsilon: explore less accurate outlet
else: exploit the best outlet
```

In this example, the probability that an agent samples a message from the optimal outlet is 95%, while only 5% are assigned to the less accurate information source.

I argue that this public information-acquiring process also resembles how people consume news in real life. Figure 2 shows that people often expose themselves to media outlets that hold different perspectives. Additionally, previous studies provide abundant evidence of selective exposure to like-minded information sources. Further, information receivers rely on the perceived source credibility when they have to learn new topics [31]. Therefore, in the social learning process, people tend to gather information from more credible sources, but they do not strictly filter out the messages from less preferred outlets.

3.3.2.3 Adjustment of the Model with Sampling Results

This study adopts a quasi-Bayesian approach to define belief updating. Under the Bayesian learning model with normal distributions, an agent holds a belief that is defined as a random normal distribution, and the signal also takes the form of a normal distribution. As explained earlier, here I attempt to model belief formation with a sampling-based approach. An agent samples some pieces of information from the other agent and uses the sample for belief updating. Therefore, this study argues that biased information processing occurs at the credibility assessment stage, which defines the sampling patterns. In addition, since the agent draws samples from the information source, the sample means and sample standard deviations are used to infer the information source's belief distribution and as the Bayesian signal, consisting of the mean and variance.

Thus, reflecting that an agent samples R number of messages at each step of learning, the aforementioned equations can be adjusted with sampling results by substituting the message sent by *Information Provider* j at time t, $x_{j,t}$, into the sample mean, $\overline{x_{j,t}}$, and the precision of the message, $\sigma^2_{x_{j,t}}$, into the sample variance $s^2_{x_{j,t}}$:

$$\widehat{\mu^{\theta}_{i,t}} = \mu^{\theta}_{i,t-1} + (\overline{x_{i,t}} - \mu^{\theta}_{i,t-1}) \left(\frac{\sigma^2_{\theta_{i,t-1}}}{\sigma^2_{\theta_{i,t-1}} + s^2_{x_{j,t}}} \right) \tag{1-A}$$

$$\sigma^2_{\theta_{i,t}} = \frac{\sigma^2_{\theta_{i,t-1}} s^2_{x_{j,t}}}{\sigma^2_{\theta_{i,t-1}} + s^2_{x_{j,t}}} \tag{2-A}$$

$$z = \frac{|\overline{x_{j,t}} - \widehat{\mu^{\theta}_{i,t}}|}{\sqrt{(s^2_{x_{j,t}} + \sigma^2_{\theta_{i,t}})}} (\frac{\sigma_{\theta_{i,t}}}{s_{x_{j,t}}}) \tag{3-A}$$

$$\mu^{\delta}_{j,t}|x_{j,t} = \frac{\mu^{\delta}_{j,t-1}(\sigma^2_{\theta_{i,t-1}} + \sigma^2_{\varepsilon}) - (\overline{x_{j,t}} - \mu^{\theta}_{i,t-1})\sigma^2_{\delta_{j,t-1}}}{\sigma^2_{\theta_{i,t-1}} + \sigma^2_{\delta_{j,t-1}} + \sigma^2_{\varepsilon}} \tag{4-A}$$

Table 1: Hypothetical Scenarios for Simulations

Citizen θ Beliefs	Information Provider's Message Crafting Parameters		
	Unbiased / Precise Vs Biased / Precise	Biased / Precise Vs Biased / Imprecise	Unbiased / Precise Vs Unbiased / Imprecise
Consensual	Scenario A-1	Scenario B-1	Scenario C-1
Flat	Scenario A-2	Scenario B-2	Scenario C-2
Polarized	Scenario A-3	Scenario B-3	Scenario C-3

3.4 Agent-Based Simulation

3.4.1 Simulation Settings

I use an Agent-based simulation to investigate how individual agents' endogenous information sampling affects macro-level social learning outcomes. Specifically, I examine how different credibility comparison rules impact *Citizens*' belief distributions, outlet choices, and the time taken for *Citizens* to stop learning the state of the world.[2] The objective of this chapter is to compare which credibility assessment rule performs better at 1) exploiting more accurate information outlets, 2) creating consensus among *Citizens*, and 3) reducing the time taken to reach a public consensus.

Simulations are conducted under 3× 3 hypothetical scenarios (as shown in Table 1). According to [50], the persuasion process is affected by the attributes of both speakers and information receivers. One aspect that characterizes the persuasion process is how the speakers craft their messages [56, 57, 69, 70]. Therefore, the first dimension of the scenarios focuses on the message-composing parameters of *Information Providers*. Similarly, the prior attitudes of the information receivers also influence social learning [55, 50, 20]. People who have extreme attitudes and more confidence in their beliefs are more difficult to persuade

[2]For computational purposes, I consider *Citizens* to have stopped learning the state of the world if their average percent changes in $\mu^{\theta}_{i,t-1}$ and $\mu^{\theta}_{i,t}$ fall below 1%.

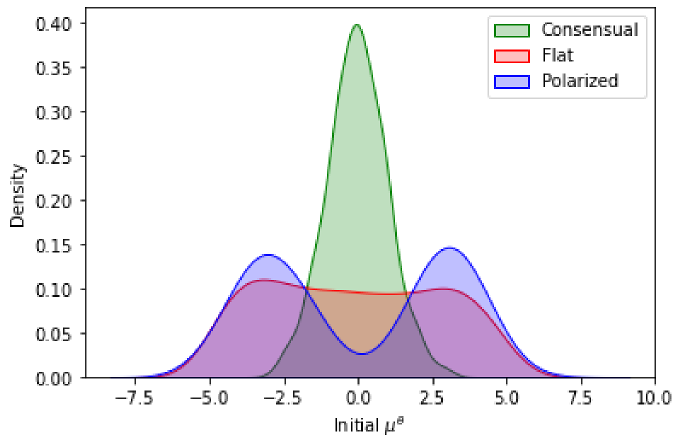

Figure 4: Three Initial μ^θ Distributions of *Citizens*

[87, 54]. To replicate various social circumstances and how *Citizens'* attitudes are initially distributed, I employ three cases: Consensual, Polarized, and Flat distributions. Figure 4 shows how the initial μ^θ beliefs of each type of *Citizens* are distributed. The standard normal distribution is used for the consensual distribution. In contrast, the polarized distribution combines two normal distributions, each characterized by a standard deviation of 1 and means of -3 and 3, respectively. The flat distribution is derived from a uniform distribution ranging from -5 to 5.

The agent-based simulation consists of 500 *Citizens* and 2 *Information Providers*. *Citizens* are the only type of agents who can communicate with *Information Providers* by sampling messages from them. Communication between agents of the same type is not allowed for the baseline model; there is no communication between *Information Providers* and between *Citizens*. At each communication step, *Citizens* sample 20 messages. The true state of the world is fixed at 0 ($\theta = 0$).

The simulation includes nine hypothetical scenarios, which are divided into three groups. The first group (Scenario A-1 to A-3) explores the effect of one unbiased and one biased *Information Providers* on *Citizens*. The precision parameters of both *Information Providers* are the same. The second group (Scenario B-1 to B-2) involves two biased *Information*

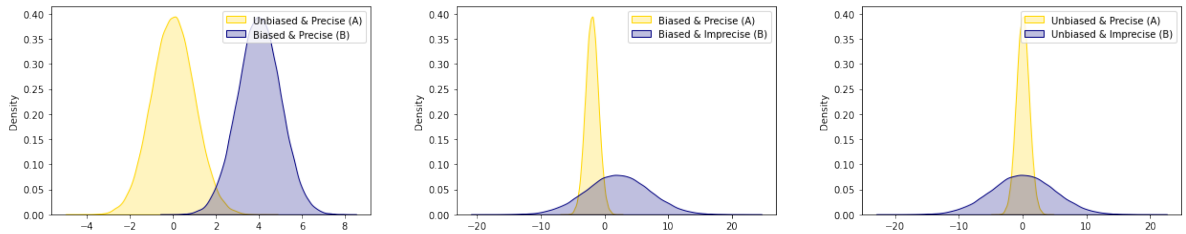

(a) Scenario A (b) Scenario B (c) Scenario C

Figure 5: *Information Providers*' Message Crafting Parameter Settings

Providers, with only one of them sending out precise messages. The third group (Scenario C-1 to C-2) has two unbiased *Information Providers* with different levels of precision. Figure 5 visualizes the probability density functions of the *Information Providers'* message crafting parameters.

For each model configuration, the simulation is conducted 100 times[3], with a maximum of 10,000 steps per simulation[4]. Should agents engage in interactions exceeding 10,000 steps, the program terminates the simulation and reports the outcomes at the last reached step.

The simulation programs were implemented using Python and the `Mesa` library [**?**]. The simulations were run on the H2P cluster[5] at the University of Pittsburgh Center for Research Computing.

3.4.2 Conjectures

This study investigates how different credibility assessment rules affect the social learning process and proposes three sets of hypotheses that focus on three outcomes of interest: 1) exploitation of the better source, 2) μ^θ distribution, and 3) time to reach equilibrium. The first set of hypotheses examines the relationship between credibility assessment mechanisms

[3]Given the nature of the model involving frequent random samplings, running simulations 100 times per model configuration allows for a substantial number of observations.

[4]Throughout the simulation process, some model configurations required more than 13,000 steps, leading to computational breakdowns due to memory constraints. As a result, a maximum step limit was imposed for computational efficiency.

[5]H2P cluster is supported by NSF award number OAC-2117681.

and optimal outlet exploitation.

Equation 3-A and Equation 4-A illustrate how agents learn the credibility of each information source. The numerator of Equation 3-A indicates that the difference between the prior beliefs and the sampled message has a positive correlation because the posterior beliefs given the sampled message are the result of biased assimilation [71]. Similarly, as citizens compare the absolute value of $\mu^{\delta}_{j,t}|x_{j,t}$ for each *Information Provider*, this implies that an information receiver's perception of biasedness becomes minimal when the difference between the messages obtained from the information source and its prior beliefs are equal. Therefore, when the variances of each *Information Provider*'s message are equal, it is expected that the optimal arm selection is solely based on the difference between an agent's prior and sampled messages, resulting in the same exploitation patterns for both credibility assessment rules. Moreover, since *Citizens* sample messages similarly, there should be no significant differences in the final belief distributions and the time required to reach a stalemate under this circumstance.

C1. If both *Information Providers'* messages are equally precise, *Citizens'* better-arm-exploitation patterns will be the same regardless of the choice of credibility assessment rules.

C2. If both *Information Providers'* messages are equally precise, *Citizens'* belief distributions should be the same regardless of the choice of credibility assessment rules.

C3. If both *Information Providers'* messages are equally precise, the number of steps until *Citizens* no longer learn should be the same regardless of the choice of credibility assessment rules.

As shown in Equation 3-A, the z-statistics comparison rule uses sampling variance to infer the credibility of the signal, while the δ comparison focuses solely on the biasedness of the signal. Thus, when there is an *Information Provider*, *Citizens* will do better at identifying the better source if they employ the z-statistics comparison rule. In addition, since *Citizens* exploit the better outlet more efficiently when they use the z-statistics strategy, the time taken to reach a stable belief state will be much shorter when they use the z-statistics comparison rule.

C4. If the variances of messages from each source vary, the z-statistics comparison rule will make *Citizens* exploit the better source more effectively compared to when they use the δ comparison strategy.

C5. If the variances of messages from each source vary, the z-statistics comparison rule will make *Citizens* stop learning sooner compared to when they use the δ comparison strategy.

Finally, unless the averages of messages from each source are the same, the δ comparison rule will produce polarization, while the z-statistics comparison rule creates homogeneous beliefs distributions because it only relies on the biasedness of messages. Similarly, if both *Information Providers* send messages with equal precision but different mean values, the z-statistics comparison rule will also produce polarized μ^θ belief distributions.

C6. If the messages from each *Information Provider* have different mean values and *Citizens* use the δ comparison rule, the posterior belief distribution will ultimately become polarized.

C7. The z-statistics comparison rule will result in polarized *Citizens'* μ^θ distributions if *Information Providers* send out messages with equal precision but different mean values.

3.4.3 Simulation Results: Baseline Model

3.4.3.1 Unbiased & Precise vs. Biased & Precise Information Providers

The first set of scenarios examines the effect of source credibility evaluation choices under the environment where one unbiased and a biased *Information Provider* sending out equally precise messages (Scenario A-1 to A-3). Figure 6 summarizes the simulation results. As previously explained, it is expected that both rules will result in the same exploitation pattern (C1), *Citizens'* learning speed should be equal (C3), and their posterior beliefs should be equally polarized (C3, C6, and C7).

Figure 6b delineates the better performing outlet exploitation patterns of *Citizens*. Interestingly, the results indicate that the δ comparison rule performs better than the z-statistics comparison rule in identifying the better outlet in all sub-scenarios. It is interesting to note that as the level of polarization increases, from consensual distribution to polarized

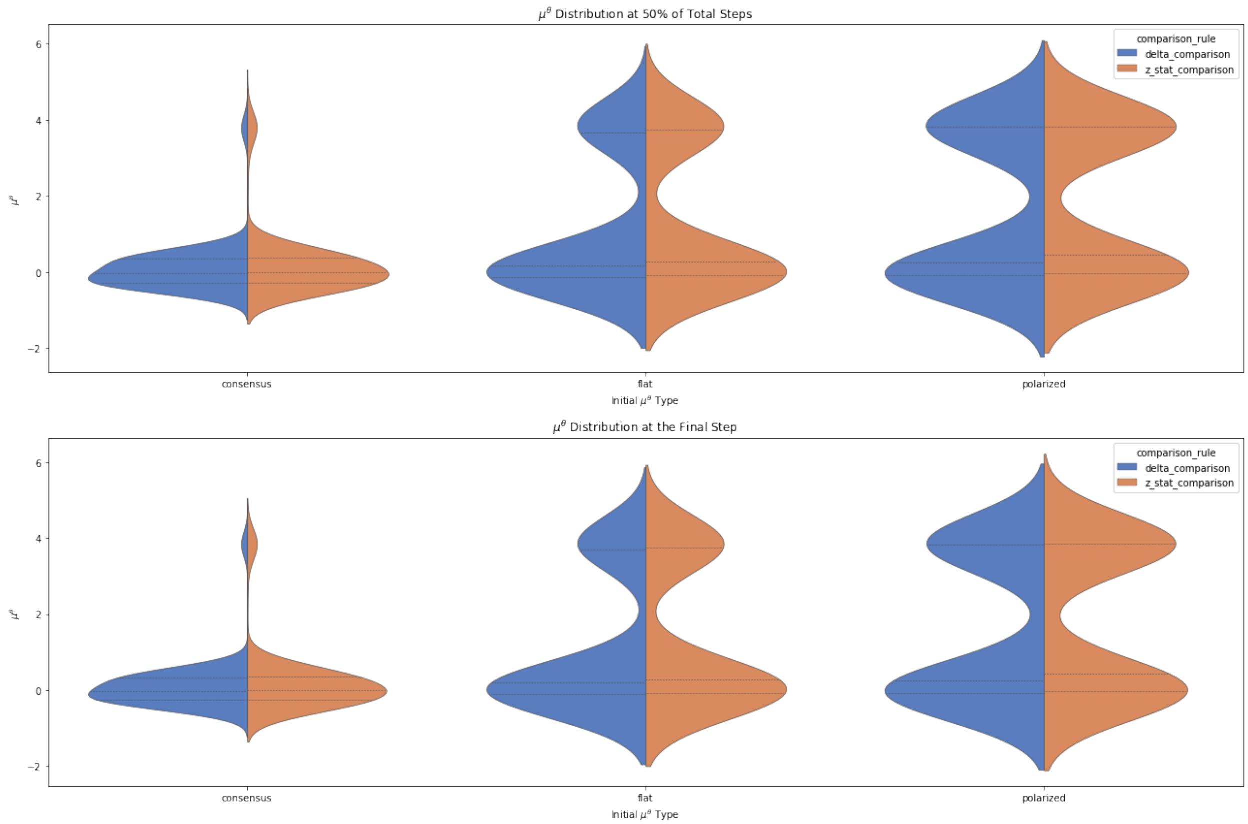

(a) μ^θ at 50% & Final Steps

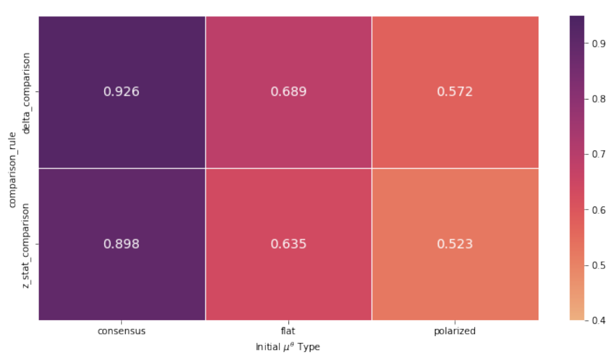

(b) IP A choice Ratio

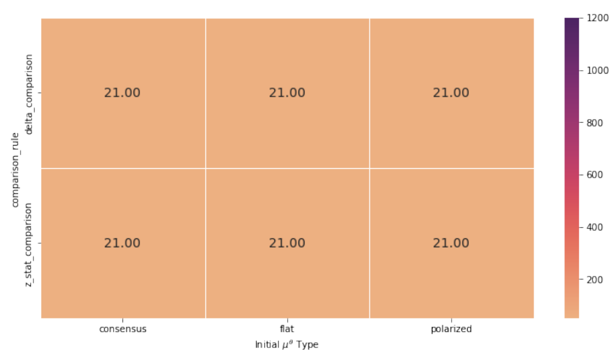

(c) Total Steps Run

Figure 6: Baseline Model Simulation Results: Scenario A

distribution, the exploitation patterns are also affected: Citizens were most successful in identifying the optimal information providers under the consensual scenario, while they were least efficient under polarized one. Under the consensual μ^{θ} distribution, 92.6% of messages were sampled from the unbiased source (*Information Provider A*) using δ comparison rule. In contrast, 89.9% of messages were sampled from it when using the z-statistics comparison strategy. When *Citizens* have flat initial beliefs, the δ comparison rule led 68.9% of total messages to be drawn from the unbiased outlet, while the ratio was 63.5% using the z-statistics comparison. Under the polarized environment, similarly, 57.2% of messages were exploited from *Information Provider A* with the δ comparison rule, while only 52.3% were sampled when *Citizens* compared z-statistics. Pair-wise comparisons of the differences in proportions between each assessment rule reveal that the differences are all statistically significant at the 99.9% confidence level. This finding indicates that when *Citizens* observe messages of equal precision from each source, their credibility evaluation is predominantly influenced by their predispositions.

Figure 6c shows consistent results with C3. Even though the proportion differences shown in Figure 6b were statistically significant, the difference was not substantially large enough: Differences were approximately 5% point or less. Thus, the same steps were run in all scenarios.

More interestingly, both credibility assessment rules produced highly similar posterior belief distributions (See Figure 6a). Unless *Citizens'* initial μ^{θ} is strongly consensual, both mechanisms produced polarized circumstances in the end (C2, C6, and C7). As it was shown in Figure 6b, the level of polarization in posterior beliefs is correlated with the degree of priors' polarization: it is most polarized when the initial beliefs are polarized, while the consensual scenario results in being least polarized. As previously explained, when *Citizens* are unable to discern differences in message precision from each source, the subjectively evaluated bias becomes the only indicator for inferring source credibility. Thus, the level of polarization in their information consumption patterns and posterior belief distributions is contingent upon how citizens' prior beliefs are distributed.

3.4.3.2 Biased & Precise vs. Biased & Imprecise Information Providers

Here, both *Information Providers* are biased, but one of them sends out imprecise messages: $m_A \sim N(-2, 1^2)$ and $m_B \sim N(2, 5^2)$ (Scenario B-1 to B-3). Reflecting the differences in signal variances, I expect that *Citizens* should exploit the more precise outlet more if they rely on the z-statistics comparison rule (C4). Further, due to better exploitation, *Citizens* are expected to stop learning the state of the world if they use the z-statistics comparison rule (C5). Finally, the δ comparison rule will lead *Citizens* to become polarized.

Again, as C4 conjectures, Figure 7b shows that *Citizens* exploit the better arm more frequently when they rely on the z-statistics comparison rule. The proportion differences between the two comparison rules were approximately 10% points in all subcategories of scenarios. Under the z-statistics comparison mechanism, *Citizens* exploited the better outlet in 87% of total messages when they already had consensus, 91% with a flat prior μ^θ distribution, and 92.5% in a polarized environment. These proportions were significantly different from those obtained with the δ comparison rule at $p < 0.001$, as confirmed by statistical testing of pair-wise comparison of proportion differences. It is more interesting to note that citizens become more successful exploiting the objectively optimal outlet (i.e., IP A) when their initial beliefs are polarized compared to when they were consensual.

Figure 7c indicates that the z-statistics comparison strategy required fewer time steps to reach the equilibria than the δ comparison rule, except in the polarized environment where both strategies exceeded the maximum number of interaction time steps. Results indicate that *Citizens* learned the more credible source early, and the better optimal arm detection led to quicker convergence among the agents (C5).

Both methods resulted in highly similar μ^θ distributions in the end (See Figure 7a). Further, the most interesting finding is that citizens' post-communication beliefs become most consensual under the polarized initial belief scenario: As indicated in Figure 7b, citizens under this setting were able to identify the objectively better outlet in the long run, which results in their posterior beliefs being concentrated on -2 and being unimodal. Under the consensual settings, on the other hand, they eventually become polarized with the modes at 0 and -2, which reflects the lower efficiency in identifying more precise information outlet.

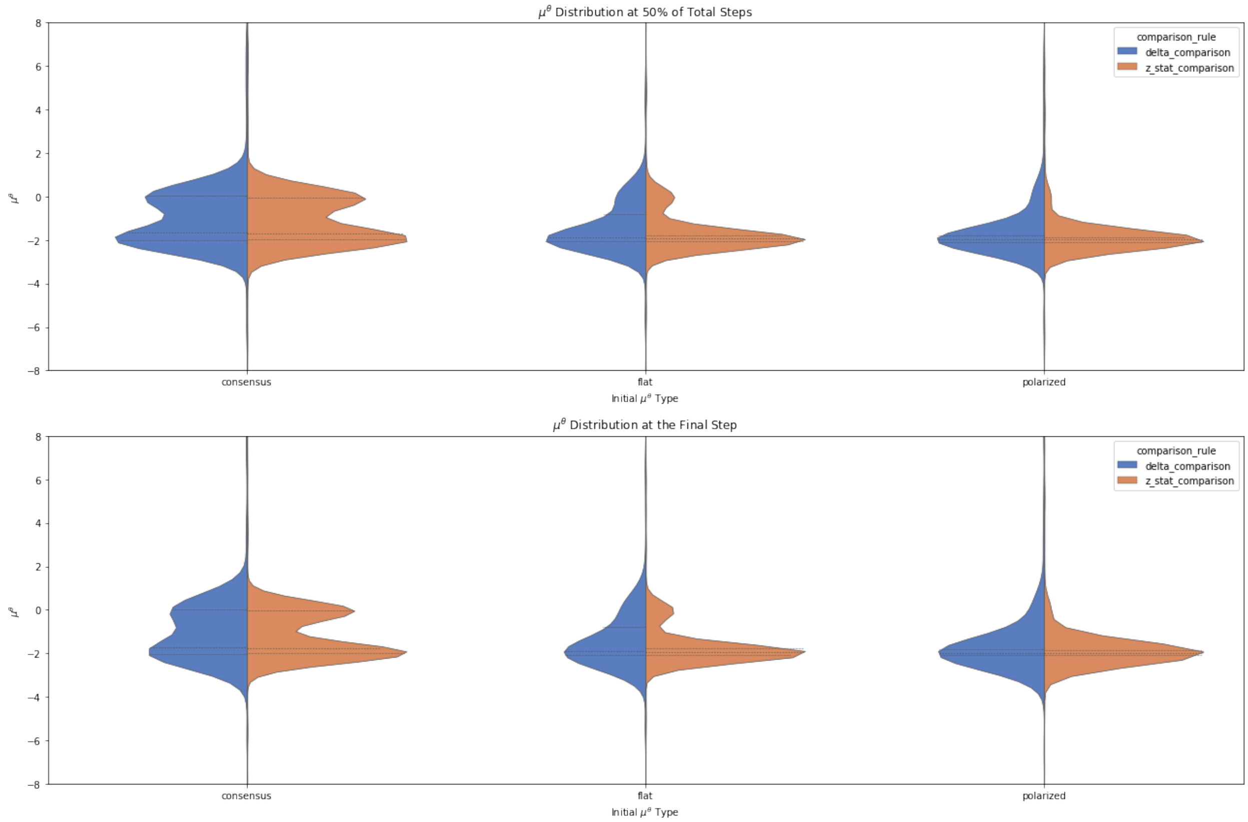

(a) μ^θ at 50% & Final Steps

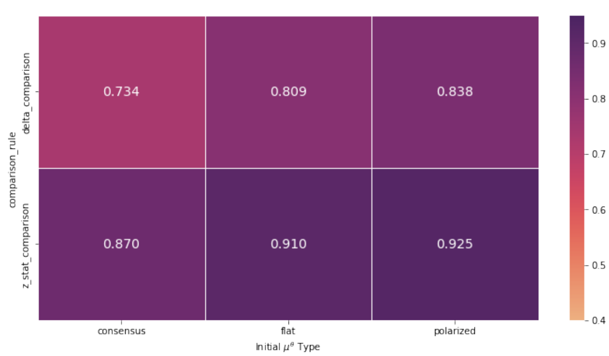

(b) IP A choice Ratio

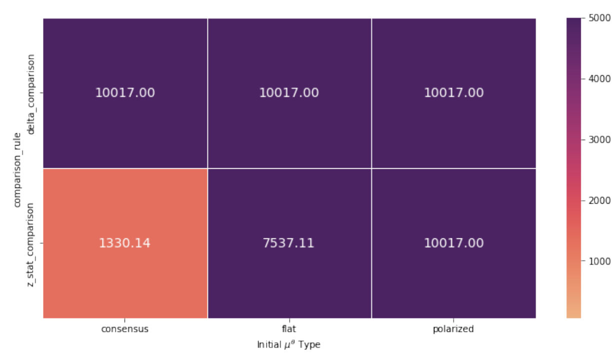

(c) Total Steps Run

Figure 7: Baseline Model Simulation Results: Scenario B

Interestingly, the δ learning process replicates the z-statistics comparison strategy even though it does not directly utilize the sample variance of messages drawn from each *Information Provider*. This finding contradicts the previously conjectures that the δ comparison strategy would result in polarized information source reliance and belief distributions due to a sole focus on perceived message biasedness (C3 and C6).

To investigate this further, I ran another set of simulations with two biased and precise information providers (Figure 12a). Figure 6a and Figure 12b indicate that the δ comparison strategy polarizes belief distributions only if messages from each source are equally precise, confirming C3 but partially confirming C6. Similarly, Figure 12c shows that when there is not an objectively better information source, information choices become polarized. At the same time, *Citizens* successfully identified the precise outlet even with the δ comparison strategy in Figure 7b.

In summary, although *Citizens* do not directly use sample variance information in credibility assessment procedures, they can slowly but successfully learn which source is more optimal than the other even if they rely on the δ comparison rule.

3.4.3.3 Unbiased & Precise vs. Unbiased & Imprecise Information Providers

The final scenarios investigate how *Citizens* learn about the state of the world and the credibility of information sources when both *Information Providers* are unbiased, but the variances of their messages differ (Scenario C-1 to C-3): $m_A \sim N(0, 1^2)$ and $m_B \sim N(0, 5^2)$ (see Figure 5c). Since both messages are unbiased, the accuracy should be assessed by comparing their variances. As previously conjectured, C4 predicts that the z-statistics comparison strategy should outperform the δ comparison strategy in identifying the best arm. It is expected that *Citizens* will be divided into two groups, each equally likely to choose either outlet because they do not rely on the sample message variances.

Figure 9b partially confirms these expectations. Under all sub-scenarios, the z-statistics comparison strategy performed better than the δ comparison strategy. Approximately 92% of messages were collected from the better outlet using the z-statistics comparison rule, while they sampled 88% of messages from the optimal source when comparing δ. Statistical

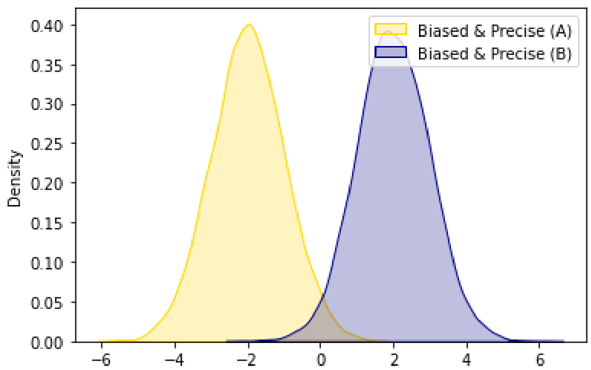

(a) Message Crafting Settings

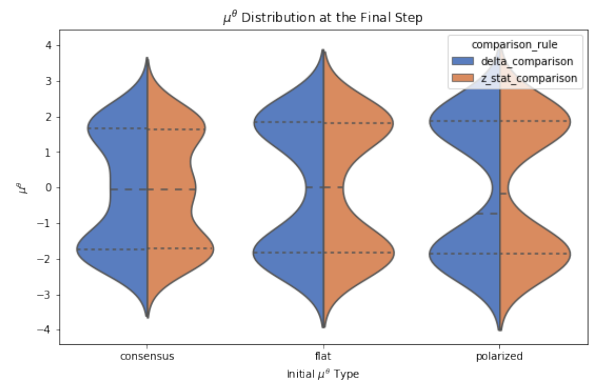

(b) Final *Citizens'* μ^θ

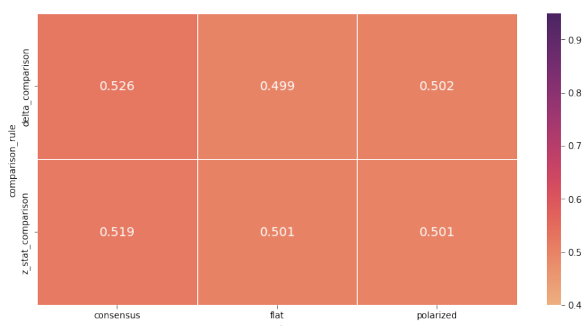

(c) *Info. Provider* A Choices

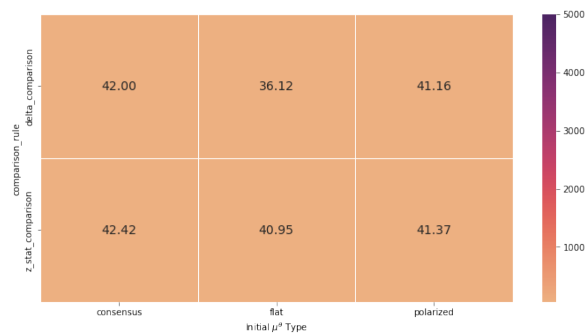

(d) Avg. Steps Run

Figure 8: Both Biased & Precise *Information Providers*

testing of the differences in proportions confirmed that the rates are significantly different from each other. As previously observed in the other scenarios, even though *Citizens* using the δ comparison rule performed less well compared to the z-statistics comparison, they successfully exploited the better outlet.

Figure 9c confirms the expectation that *Citizens* will stop learning sooner when they compare z-statistics to evaluate the credibility under this type of scenario (C5). While the average number of time steps run until agents stopped learning the state of the world was below 30 steps, it took much longer time to reach the stable status. Figure 9b and Figure 9c imply that *Citizens* slowly learn the credibility of information source objectively when they use δ comparison strategy. Finally, since they both successfully identify the optimal arm, the posterior μ^θ are precisely distributed at the true state of the world (i.e., $\theta = 0$).

3.5 Implication and Discussions

Agent-based simulation results confirmed some of my theoretical expectations. First, as predicted by C1 to C3, when *Citizens* perceive the differences in message tones (i.e., the sample average of received messages) with the same degree of precision, both credibility assessment mechanisms produce highly similar learning patterns. As their behaviors are solely determined by the differences of means, they tend to create a polarized environment shortly after agents begin interactions. Their choices of preferable information sources also become similar.

When agents perceive the differences in the sample variances of messages from each source, the z-statistics comparison policy performs much better at identifying better sources compared to when they rely on the δ comparison strategy. As they exploit better sources more frequently by comparing z-statistics, *Citizens'* posterior μ^θ beliefs reach an agreement sooner. Finally, under these circumstances, *Citizens'* posterior μ^θ belief distributions become consensual because of the high reliance on the better information source.

An interesting finding from the simulation results is that even though the δ comparison strategy does not take variances of sampled messages into account to calculate δ values,

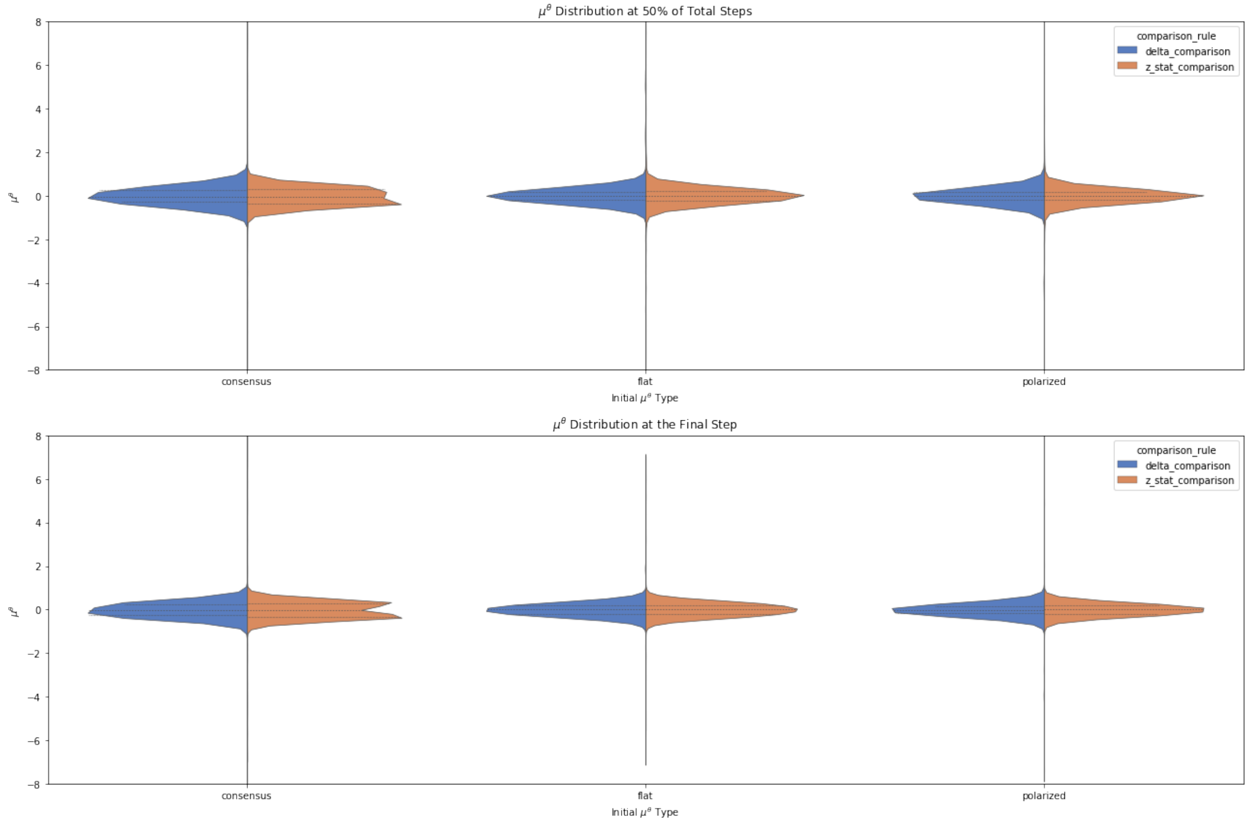

(a) μ^θ at 50% & Final Steps

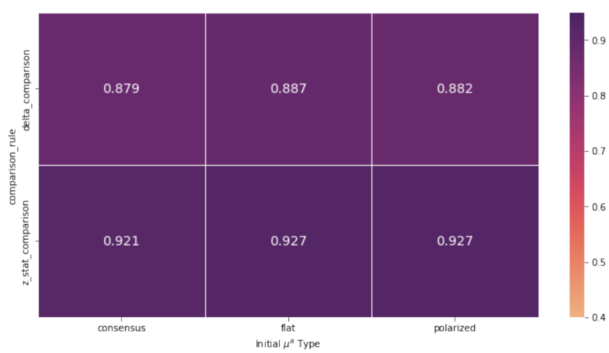

(b) IP A choice Ratio

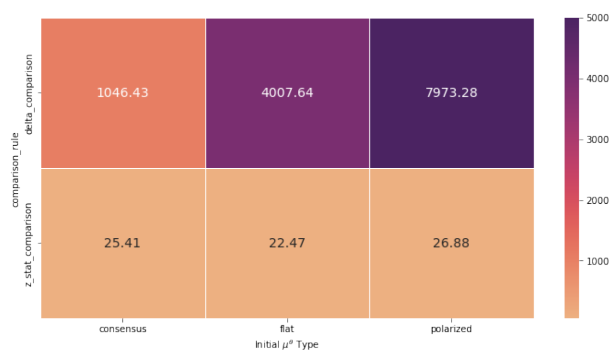

(c) Total Steps Run

Figure 9: Baseline Model Simulation Results: Scenario C

it successfully replicates the performance of the z-statistics comparison policy, which directly utilizes the sample variance (See Figure 7b and Figure 9b). However, as the upper rows of Figure 9c and Figure 9c indicate, the credibility learning process took much longer when agents employ the δ comparison rule. The posterior belief distributions also become homogeneous because of the successful detection of optimal information sources.

[84] argues that people learn as "cautious Bayesian." He found that people's beliefs do not get polarized even though they hold cognitive biases. This is supported by the agent-based simulation results presented in this study, which indicate that even if agents assimilate information in biased ways, they can still learn the objective credibility of information sources using the Bayesian learning rule.

[50] proposes Generalizing Persuasion (GP) framework. This framework defines the information receivers' persuasion mechanism as the function of 1) their cognitive efforts to assimilate information, 2) directional motives, and 3) prior attitudes. This study's model of social learning echoes the GP framework, as the *Citizens* learn the credibility of information sources and determine their reliance on each outlet based on their pre-existing attitudes and directional motives. The simulation results suggest that agents can still reach agreements even when reflecting these aspects moderately by exploiting objective optimal outlets.

GP framework also emphasizes the role of speakers. Two factors should be considered in this approach: The types of speakers and how speakers craft messages. Previous studies found that people acquire political information both from the better informed [91, 120, 88, 29] and peers who are easily accessible [116, 147, 127]. Further, reflecting the information senders also have motives, better understanding in message crafting efforts are needed. However, little is known about how information providers tune their messages. In the next paper, the author proposes another model that introduces social networks where both *Citizens* and *Information Providers* can send out messages to others.

4.0 Social Network Model: Biased Credibility Assessment, Social Network, and Political Belief Distribution

4.1 Introduction

In the first chapter, I delved into the consequences of biased source credibility assessment on citizens' political belief distribution, considering dyadic interactions between citizens and information sources. However, the intricate landscape of contemporary political discourse entails a multifaceted process for the formation of political beliefs, shaped by individual-level cognitive biases, diverse information sources, and the dynamic interplay within social networks. The rise of digital media, marked by both a significant increase in available information sources and the segregation of information consumption networks through the "filter bubble," has fundamentally reshaped how individuals engage with political information. In response to this evolving landscape, this chapter seeks to expand the individual-level mechanism established in the baseline model into social contexts, by placing agents in the complexity of social interaction.

More specifically, this section aims to delve into the dynamics of social learning, extending the exploration from the baseline model to incorporate the influential factor of social interactions. While the baseline model provided valuable insights into the interplay between information providers and citizens, it inherently operated in a more isolated learning environment. Acknowledging the real-world prevalence of social connections shaping individuals' perceptions and belief systems, this chapter introduces the Social Network Model to capture the complexities of information dissemination in a socially interconnected society.

The Social Network Model aspires to unravel the intricate ways in which citizens' beliefs evolve, influenced not only by information providers but also through between-citizen interactions within their social networks. Leveraging the structured connectivity of these networks, we scrutinize how citizens' social interactions contribute to or counteract belief polarization and misinformation spread. This chapter commences by offering a detailed exploration of the variations in social network structures employed in our simulations, emphasizing the

Random 2 Matching, Group ID Matching, and Extended Network scenarios.

Embarking on this exploration, I recognize the significance of comprehending how diverse social structures shape the dynamics of information diffusion. The Random 2 Matching scenario allows for random communication among citizens, offering insights into consensus-building and information sharing in a more unpredictable network. On the other hand, the Group ID Matching scenario reflects a homophilic social network where citizens connect based on shared initial beliefs, providing a perspective on information dissemination within like-minded clusters. Finally, the Extended Network scenario introduces a broader interaction landscape, encompassing interactions not only with fellow citizens but also information providers, providing a nuanced understanding of the interplay between various actors.

Through the lens of the Social Network Model, this chapter extends the examination of credibility assessment mechanisms, incorporating not only individual-level assessments but also the collective impact of social interactions on belief dynamics. The results presented in this section aim to unravel how citizens navigate and interpret information in a socially interconnected world, shedding light on the critical role of social networks in shaping belief systems and mitigating the adverse effects of biased information.

4.2 Related Literature

In the previous chapter, the baseline model focused solely on how information receivers react to information delivered by information providers. Specifically, the previous model theorized how citizens assess the credibility of political information sources with cognitive bias. Additionally, under the baseline model, citizens interact only with information providers. In the social learning environment, however, all citizens can simultaneously be the information consumer and the outlet delivering it to fellow citizens. Thus, communication channels among citizens are essential aspects that determine individuals' political belief updates and social dissemination of beliefs.

Existing research has emphasized the importance of social interaction in citizens' political attitude changes. When people interact with better-informed individuals, for instance, they tend to become more interested in politics, which leads them to become politically more active [120, 95, 96, 97, 94, 28]. In other words, communication with the well-informed group can improve the less-informed public's political knowledge, and improved political knowledge leads them to be more willing to share it. This finding implies that interaction with the better-informed can promote the long-term dissemination of political knowledge through second-hand delivery from the previously less-informed population. Furthermore, since the better-informed are more likely to be open to political belief updating [172, 53, 121, 35, 92], political communication can create an open-minded social environment.

Deliberation literature also highlights the importance of political conversations as a fundamental element of political persuasion. Specifically, it examines the effects of citizens' exposure to diverse political arguments. From a normative and empirical perspective, communication between groups is considered beneficial for democracy because exposure to opposing viewpoints can: 1) promote interpersonal deliberation [124, 80], 2) improve awareness of different perspectives [130, 131, 132, 61], and 3) foster a cooperative environment between groups [123, 135, 18, 30].

While exposure to counter-attitudinal information is assumed to lead to political persuasion, research shows that it often has the opposite effect. In fact, people from rival groups are often not persuaded, and conflicting messages can sometimes cancel each other out, failing to affect attitude change [148]. In some cases, exposure to contradictory information can even reinforce preexisting attitudes [79, 122, 16, 58, 10, 9]. This backlash effect is influenced by various factors, including the receiver's motivation to assimilate new information in a biased manner [50], social pressure [107, 146], or group identity [111].

Social interactions with neighbors can indirectly influence a citizen's attitude formation. As briefly discussed previously, consensus within close friendship networks can shape social norms, influencing in-group members to conform [107, 146, 72, 38]. Such social norm perceptions further depend on the frequencies of messages people encounter within social networks. [105] found that if a person is overly confident that their own opinion is consistent with the majority's, they may resist changing their prior attitude because they believe they are complying with the social consensus, even though their beliefs about the social majority are incorrect globally.

Another interesting aspect of the social communication channel for political information exchange is that it can be both self-selective and given. For example, party identity is often formed through socialization in the early stages of life when social network choice is less selective. However, as individuals get older, they have the opportunity to choose their communication counterparts, such as friends, spouses, and life partners, who dominate daily conversations. In these network choices, individuals tend to seek out like-minded individuals [81, 110, 34]. People tend to limit their communication to like-minded counterparts [131], and political ideology-driven homophily can also be observed on friendship networks on social media [37, 21].

Self-selection in communication partners becomes more salient when a listener can clearly observe the speaker's identity cues, such as party or media outlet names. [49] argues that citizens rely on well-informed discussants with compatible political orientations, as this shortcut can reduce the cost of analyzing political information. Subsequent studies have confirmed that like-minded political experts are crucial sources of information for citizens [39, 111, 19, 126, 29]. Furthermore, a politician is often perceived as less credible when the citizen and the politician are not from the same party [102, 57, 78, 171]. Ideological selective exposure also illustrates how political predisposition affects news media choices [152, 153, 89, 99, 13, 37].

The intricate intersection between the self-selective nature of the communication network and political belief formation through social interaction again illustrates the endogenous cyclic mechanism emphasized in the previous chapter: a citizen prefers talking to like-minded neighbors, reinforcing existing attitudes. This section proposes a social network model, comparing how different types of social networks affect the social belief distribution among citizens, assuming citizens behave as they do in the baseline model.

4.3 Social Network Model: Incidental Social Learning & Biased Assimilation of Political Information

4.3.1 Model Outline

This chapter delves into the impact of social interactions on the macro-level outcomes with comparison to the baseline model. Social Network Model is also built upon the baseline model. While the baseline model only includes dyadic links between each *Information Provider* and *Citizen*, the Social Network Model, on the other hand, *Citizens* can also communicate with each other. Except for the network structure that agents resides in, the behavioral rules of agents remain fixed.

As the Baseline Model is, two types of agents reside in the model: *Information Provider* and *Citizen*. *Information Providers* can directly observe the true state of the world (θ) with their own bias δ_A. *Information Provider* A's belief about θ is normally distributed: $\theta_A \sim N(\theta + \delta_A, \sigma_A^2)$. They send out messages to *Citizen* agents which perfectly reflect *Information Provider* A's beliefs about the state of the world with no distortion: $x_{Aj} \sim N(\theta + \delta_A, \sigma_A^2)$. Again, they do not update their beliefs about the state of the world unless it is changed. This model has two *Information Providers* available: IP_A and IP_B.

The first variation in social network structure is described in Figure 10b (Random 2 Matching). Here, each citizen agent is randomly connected with two other agents, regardless of the type. Thus, a citizen can communicate with 1) both information providers, 2) one information provider and one citizen, or 3) two fellow citizens. When choosing communication partners, a citizen randomly chooses two agents with equal probabilities: $Pr(E_{i,j}) = \frac{1}{N}$, where $Pr(E_{i,j})$ is the probability of having an edge from citizen i to other agent j and N is the number of entire agents (i.e., 499 citizens + 2 information providers).

The Group ID Matching network, on the other hand, represents the homophilic social network. Each agent is categorized into two groups based on its initial belief about the world status: If their initial belief is greater than 0, they are categorized as Group 1; otherwise, Group 2 (Figure 10c). The group identities are assigned for all agents regardless of their types. They are connected to two agents reflecting their group ID: the probability of being linked with the in-group member is 90%, but they can be matched with the out-group members with the probability of 10%.

Finally, citizens are connected with two information providers, and with two citizens under the Extended Network (Figure 10d). At each time step t, a citizen decides the communication partner type (i.e., whether citizen or information provider). For instance, if citizen i chooses to communicate with fellow citizens, it allocates resources to sample messages from two connected citizen neighbors.

As this model allows *Citizens* to share their beliefs with others, they also send out messages that perfectly represent a *Citizen*'s θ beliefs, which are drawn from a normal distribution: $x_{ij} \sim N(\theta_i, \sigma_i^2)$, where x_{ij} is the message x sent from *Citizen* i to another *Citizen*, j, and θ_i is *Citizen* i's beliefs about the actual state of the world.

All agents are placed in a directed network space: If *Citizen 1* is connected with *Citizen 2* (i.e., $(C_1, C_2) \in E_{C_1}$), A can sample messages from B; but it does not guarantee that B will listen to the messages coming from A reciprocally (i.e., $(C_2, C_1) \notin E_{C_2}$), where E_{C_i} represents network edges that *Citizen* C_i has with other agents in the network g. In other words, E_{C_i} represents a set of communication channels that *Citizen* i can use to learn the state of the world (θ) in this model. After sampling messages from the conversation partners, they update their θ beliefs using the Bayesian updating rule explained in the previous chapter:

$$\mu_{i,t}^{\theta} | x_{i,t} \sim N(\widehat{\mu_{i,t}^{\theta}}, \sigma_{\theta_{i,t}}^2) \tag{6}$$

$$\widehat{\mu_{i,t}^{\theta}} = \mu_{i,t-1}^{\theta} + (x_{i,t} - \mu_{i,t-1}^{\theta}) \left(\frac{\sigma_{\theta_{i,t-1}}^2}{\sigma_{\theta_{i,t-1}}^2 + \sigma_{x_{j,t}}^2} \right) \tag{7}$$

$$\sigma_{\theta_{i,t}}^2 = \frac{\sigma_{\theta_{i,t-1}}^2 \sigma_{x_{j,t}}^2}{\sigma_{\theta_{i,t-1}}^2 + \sigma_{x_{j,t}}^2} \tag{8}$$

Again, agent-based simulations are implemented under 3×3 hypothetical scenarios (see Table 1). There are 2 *Information Providers* and 500 *Citizens* in this model. As set in the last chapter, information providers are never affected by other agents. Also, each citizen has $R = 20$, which means they sample 20 messages at each communication step. Finally, the true state of the world is also fixed at 0 ($\theta = 0$).

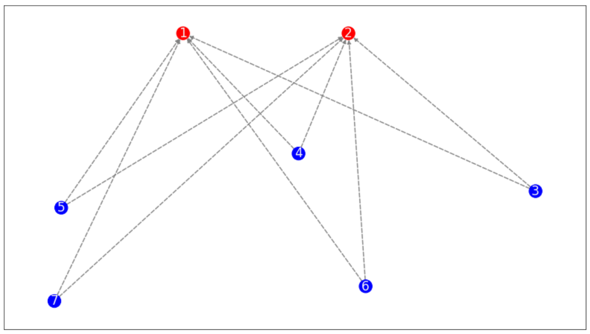

(a) Baseline Model Network

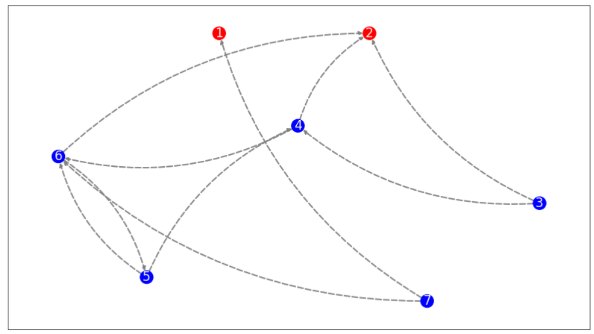

(b) Random 2 Matching

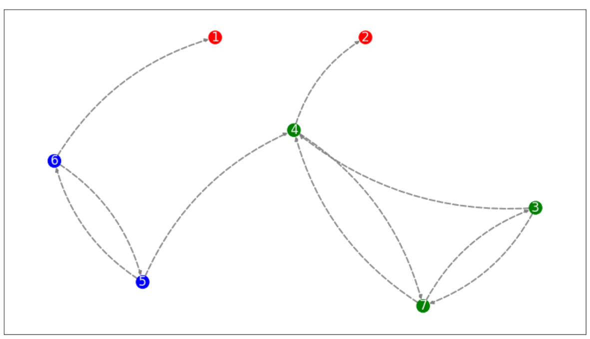

(c) Group ID Matching

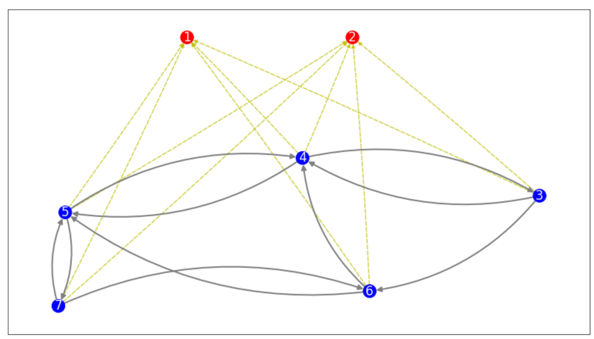

(d) Extended Network

Figure 10: Network Structures

4.3.2 Simulation Results

4.3.2.1 Unbiased & Precise vs. Biased & Precise Information Providers

First, I begin with the analysis of the scenarios where there is a competition between a unbiased and precise information provider and a precise but biased information provider. In addition, as the model does not guarantee all information providers are always connected with citizens, I focus on the analysis of post-communication belief distribution and the pace of social belief formation.

Figure 11 illustrates the citizens' post-communication belief distributions per network settings.[1] Overall, with comparison to the result from baseline model, the figures reveal that the social interaction among the citizens reduces the belief polarization. The result under the Random 2 Matching network is intuitively straightforward: As they are matched with random two neighbors regardless of the partner's prior belief status, the network should facilitate the consensus building as it will average beliefs of all citizens.

The findings from the latter two network structures, on the other hand, are particularly intriguing because they are expected to result in the belief polarization. The Group ID Matching, for instance, represents homophily among citizens as they are more likely to be connected with like-minded agents with a probability of 90%. Similarly, as observed from the baseline results, the extended network is anticipated to result in belief polarization because information providers consistently send precise messages that attract citizens. Figure 11b and Figure 11c depict more homogeneous posterior belief distributions, centered around 0. These findings reflect the second-hand learning through the social interactions. If a citizen is linked with the optimal (i.e., unbiased) information provider, it rapidly adopts beliefs similar to the provider's, effectively serving as a quasi-information provider, spreading objectively unbiased messages to others.

Figure 12 displays the average number of steps taken by citizens to reach equilibrium. Although the pace slows down due to citizen-to-citizen communication compared to the baseline model without such exchanges, it remains reasonably fast when compared to results

[1] The extreme values are dropped out for a visualization purposes. There are only few cases of having extreme μ^{θ} values throughout the simulations. The total number of dropped observations is only 6.

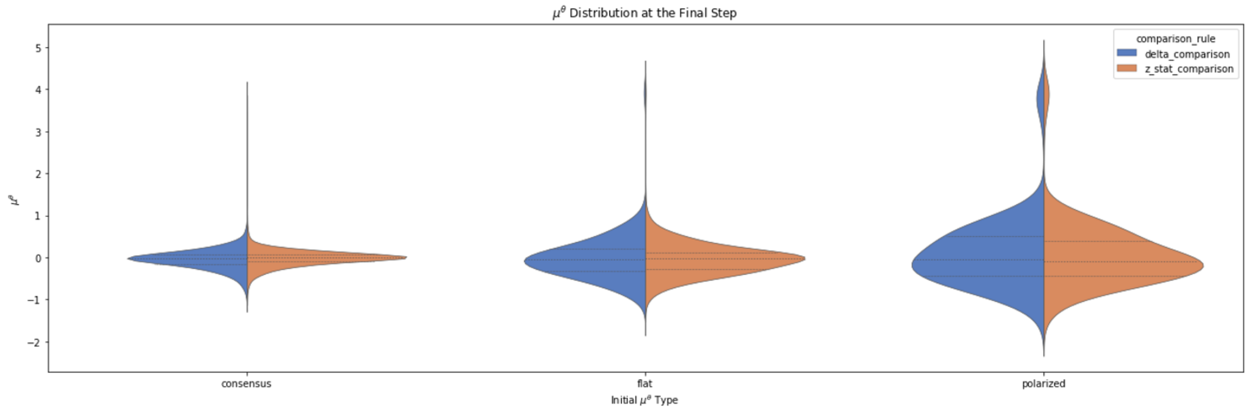

(a) Random 2 Matching

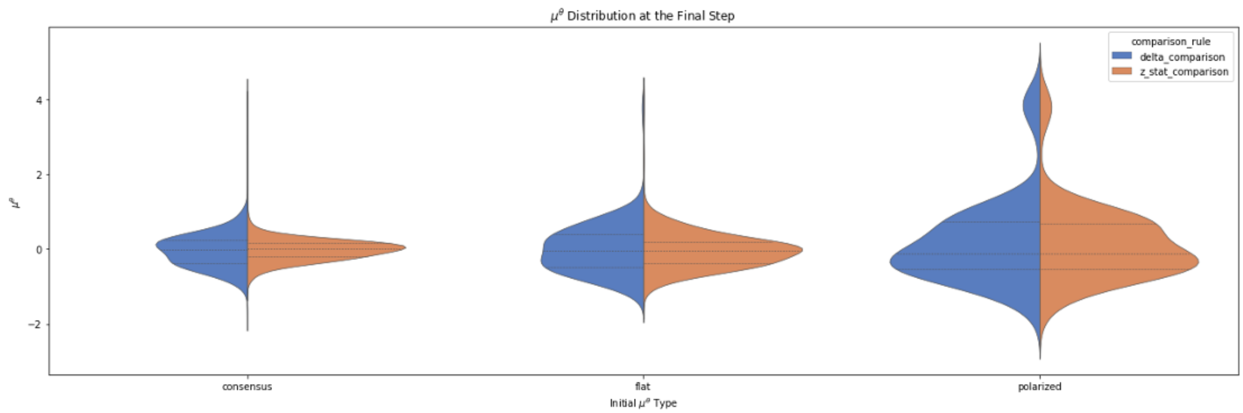

(b) Group ID Matching

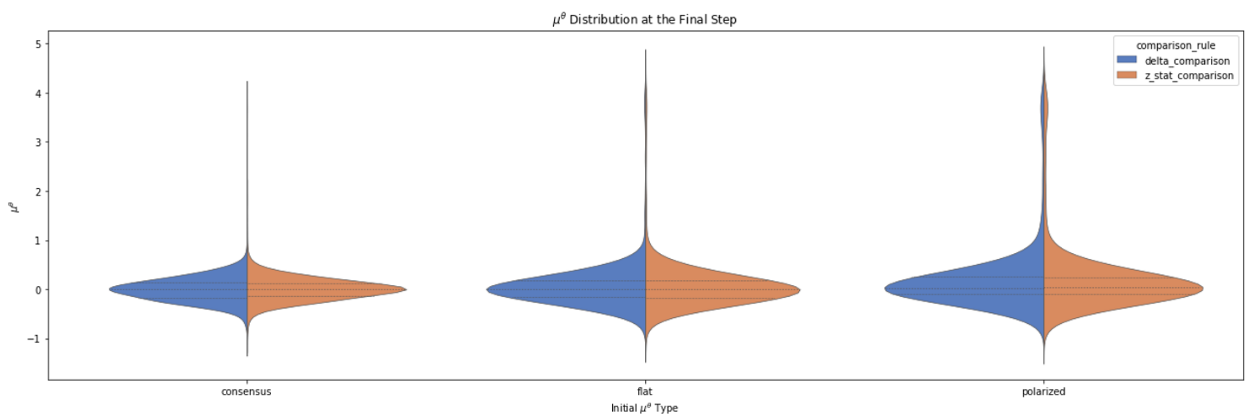

(c) Extended Network

Figure 11: μ^θ at Final Steps: Scenario A

from other baseline scenarios. On average, the Random 2 Matching scenario has the fastest pace, while the Extended Network scenario has the slowest pace: specifically, 104 steps for Random 2 Matching, 109 steps for Group ID Matching, and 114 steps for the Extended Network. Notably, the δ comparison strategy performs just as effectively as the z-statistics comparison rule.

4.3.2.2 Biased & Precise vs. Biased & Imprecise Information Providers

The second set of scenarios explore the environment where both information providers are equally biased in opposite directions, but only one of them are sending out precise messages to citizens. Figure 13 summarizes the citizens' post-communication belief distributions. Interestingly, even though the posterior beliefs are distributed differently, all simulation results reveal that citizens are not affected by the imprecise information provider.

Under Random 2 Matching network, the belief distributions are highly similar to the baseline model's result (Figure 13a). When citizen's initial beliefs are distributed homogeneously (i.e., the consensual distribution), they ended up with having slightly polarized belief distribution with two mods around -2, where the precise information provider's μ^{θ} belief is located, and 0, where is the citizens' initial consensus point. Other initial beliefs, on the other hand, they tend to have consensual belief distribution in the end.

Further, it is interesting to note that the Group ID Matching network lead to consensus building among citizens even though they are placed under homopilic environment. As previously mentioned, the network is expected to create belief polarization among citizens due to the nature of network structure. However, Figure 13b illustrates the most homogeneous posterior belief distributions compared to the other network structures. Again, all initial belief distributions of citizens are finally transformed into the consensual post-communication belief distributions.

The posterior belief distribution under the Extended Network combines findings from the previous results. When citizens' initial beliefs are homogeneously distributed, they tend to stay at the initial consensus point even after the social interactions. Under the polarized initial belief scenario, their post-communication beliefs are significantly affected by the pre-

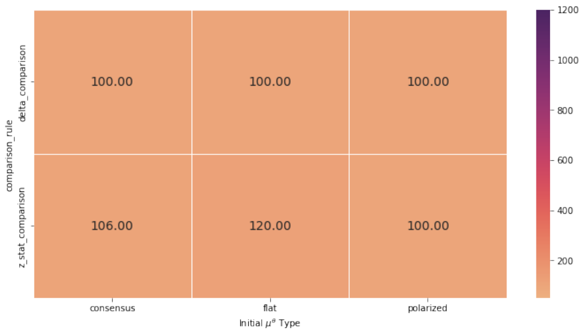

(a) Random 2 Matching

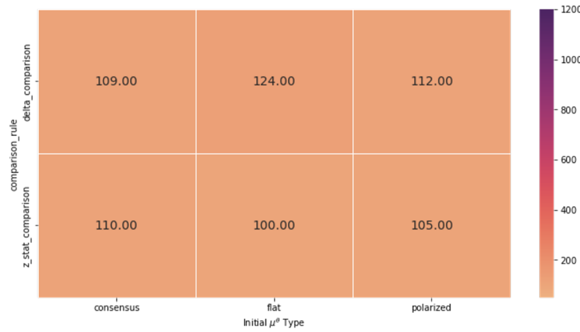

(b) Group ID Matching

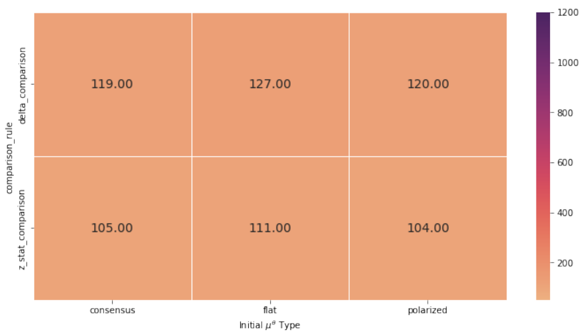

(c) Extended Network

Figure 12: The Average Steps Run Until The Equilibrium

cise information provider's messages. The flat initial belief scenario results in the loose belief polarization with two modes arount -2 and 0.

Figure 14 shows the average number of steps taken by citizens to reach equilibrium. The baseline model had a very slow belief convergence: overall, the average number of steps required was 8144.54. The social interactions, on the contrary, significantly expedites the social belief formation. When citizens are placed under the random 2 matching network and they use the δ comparison strategy for credibility evaluation, the communication lasts much longer than other scenarios. On average, the group ID matching has the fastest pace, while the random 2 matching scenario has the slowest pace: approximately 5251 steps for Random 2 Matching, 106 steps for Group ID Matching, and 117 steps for the Extended Network.

4.3.2.3 Unbiased & Precise vs. Unbiased & Imprecise Information Providers

Finally, when both information providers are unbiased, the post-communication belief distributions become consensual (Figure 15), as I found from the baseline model. Again, the random 2 matching network (Figure 15a produces highly similar results with the baseline model. The group ID matching network also produces highly consensual belief distribution, except the polarized initial belief distribution. Under the polarized initial belief scenarios, citizens' belief distribution become slightly polarized with the mods at -0.5 at +0.1 respectively. This reflects that citizens from each group forms own group consensus, but they are still pushed towards 0 due to the unbiased information providers. Finally, citizens under extended network structure reach on a social consensus at 0, with more dense distribution (Figure 15c). While previous two networks produces relatively long tails in belief distributions, the post-communication beliefs under the extended network are densely distributed with short tails.

Figure 16 shows the average number of steps taken by citizens to reach equilibrium. The baseline model had a very slow belief convergence: overall, the average number of steps required was 2183.85. The social interactions, on the contrary, significantly expedites the social belief formation. When citizens are placed under the random 2 matching network and they use the δ comparison strategy for credibility evaluation, the communication lasts much

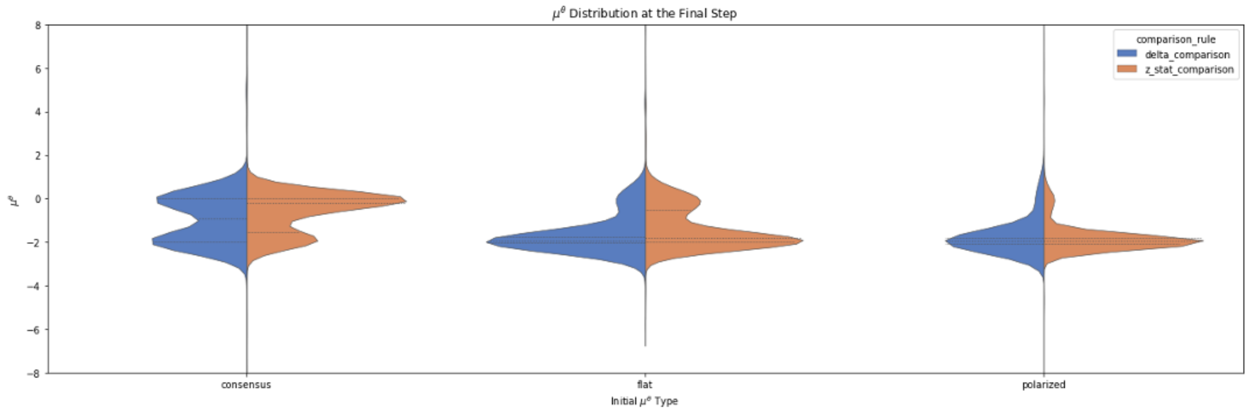

(a) Random 2 Matching

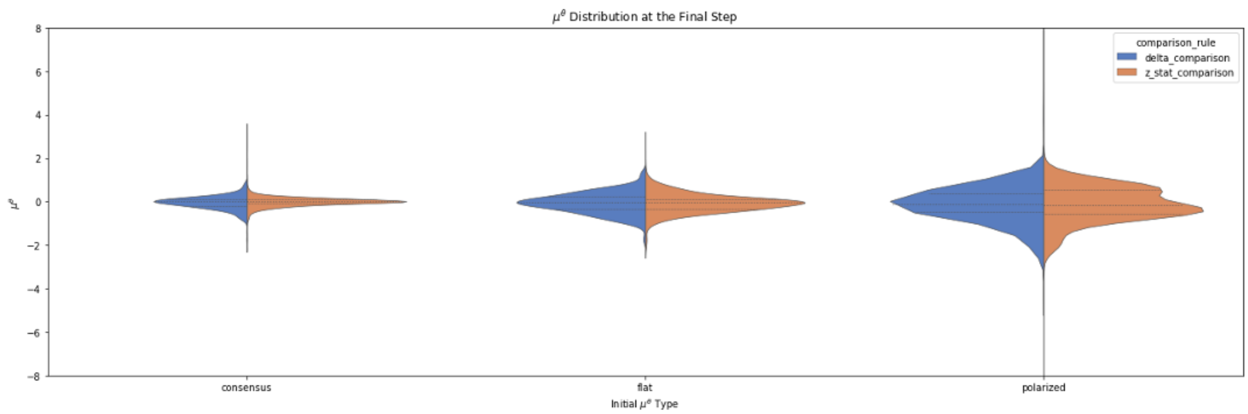

(b) Group ID Matching

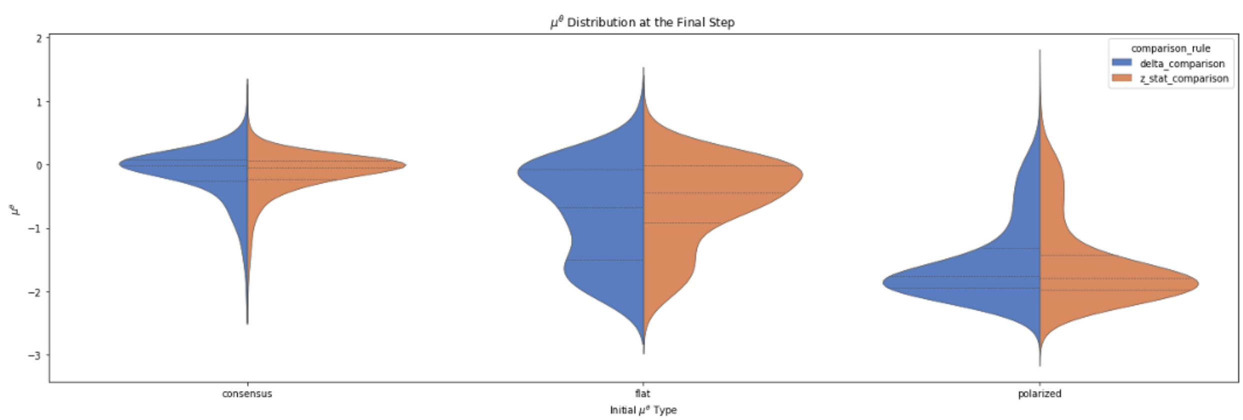

(c) Extended Network

Figure 13: μ^θ at Final Steps: Scenario B

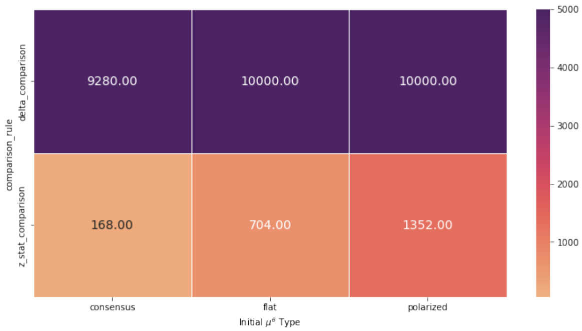

(a) Random 2 Matching

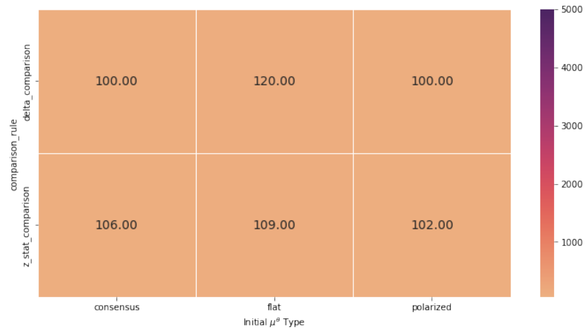

(b) Group ID Matching

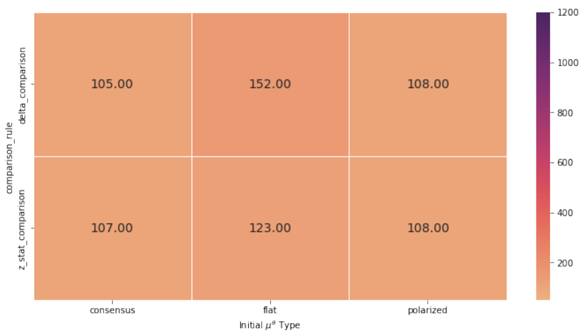

(c) Extended Network

Figure 14: The Average Steps Run Until The Equilibrium: Scenario B

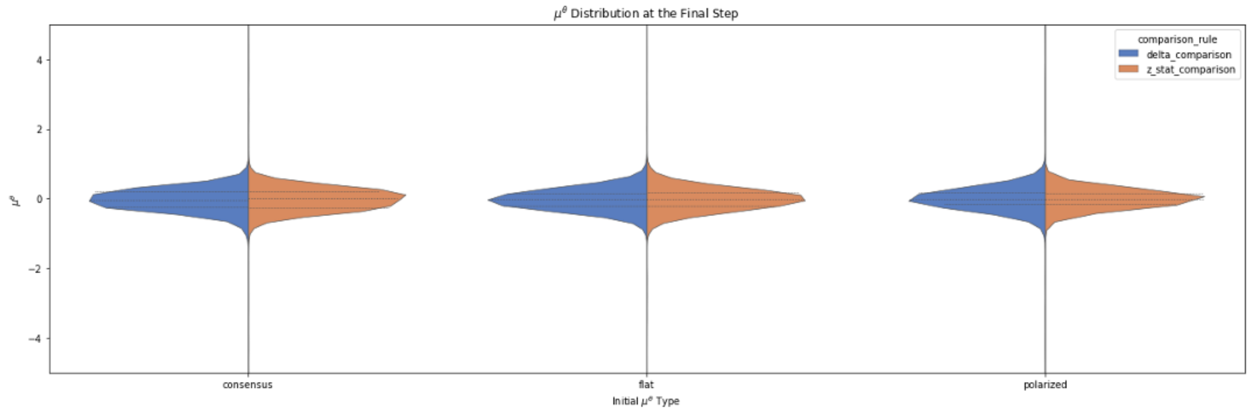

(a) Random 2 Matching

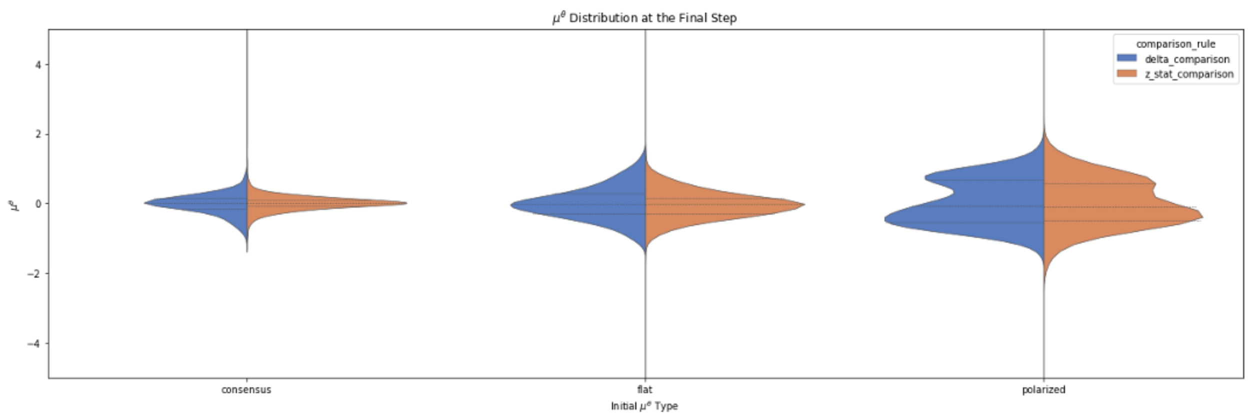

(b) Group ID Matching

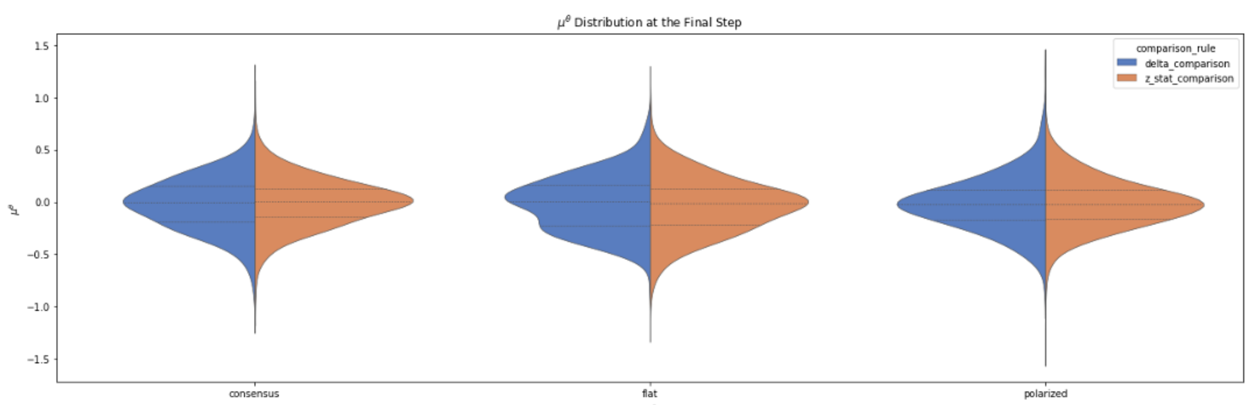

(c) Extended Network

Figure 15: μ^{θ} at Final Steps: Scenario C

longer than other scenarios. On average, the extended network has the fastest pace, while the random 2 matching scenario has the slowest pace: approximately 1977 steps for Random 2 Matching, 118 steps for Group ID Matching, and 107 steps for the Extended Network.

4.3.2.4 Second-handed Learning or the Artifact of Belief Distribution?

To examine the factors contributing to consensus-building across various social network structures, additional analyses were conducted, presenting two variations and a step-by-step exploration of the Social Network Model. Appendix A provides detailed descriptions of two extensions of the Group ID Matching network, introducing variations in the tendencies of homophily. The results from these extensions consistently align with the main model findings, emphasizing that social interaction inevitably leads to the formation of social consensus (For a comprehensive explanation, please refer to the Appendix).

Another factor that may influence the simulation results is the inherent nature of the belief distributions themselves. Given that citizens' initial belief distributions in all sub-scenarios are centered at 0, it's plausible to consider whether social interaction naturally leads to consensus building by converging to the distributional centroid, regardless of the presence of elite information providers. However, this conjecture is contradicted by the findings discussed earlier, as they ultimately result in different belief distributions depending on the information providers' belief distributions and the social network structures.

To further investigate this, additional simulations were conducted with two equally biased information providers, each with messages of equal precision but in opposite directions: $x_{Aj} \sim N(-2, 1^2)$ and $x_{Bj} \sim N(2, 1^2)$. If the convergence to consensus were solely a product of natural convergence, citizens should end up with consensual belief distributions regardless of the network types. Therefore, simulations were run under Random 2 Matching and Group ID Matching networks, as belief distributions would be polarized under these scenarios if second-handed learning existed.

However, Figure 17 contradicts the conjecture that social consensus would be achieved regardless of social network structures if it were solely an artifact of belief distributions. Interestingly, citizens' posterior beliefs tend to polarize under the Random 2 Matching net-

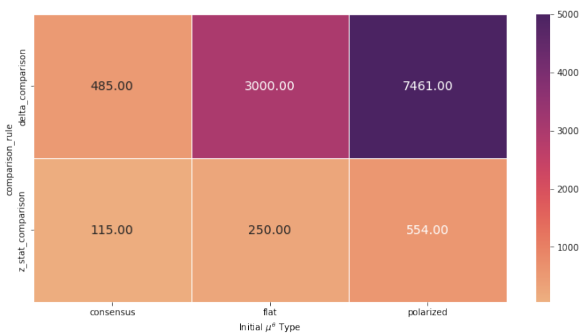

(a) Random 2 Matching

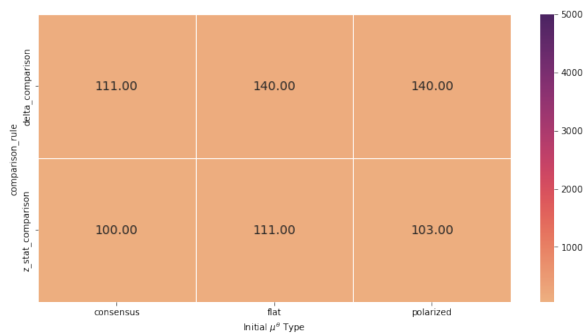

(b) Group ID Matching

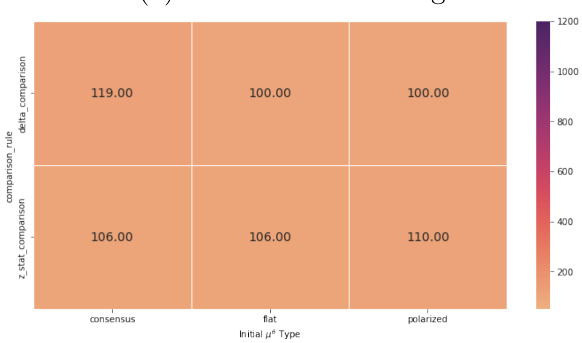

(c) Extended Network

Figure 16: The Average Steps Run Until The Equilibrium: Scenario C

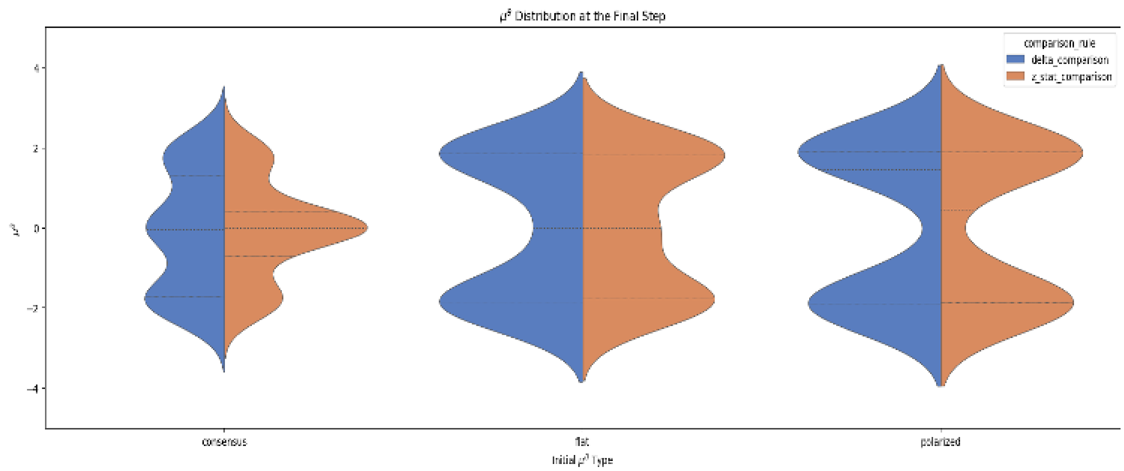

(a) Random 2 Matching

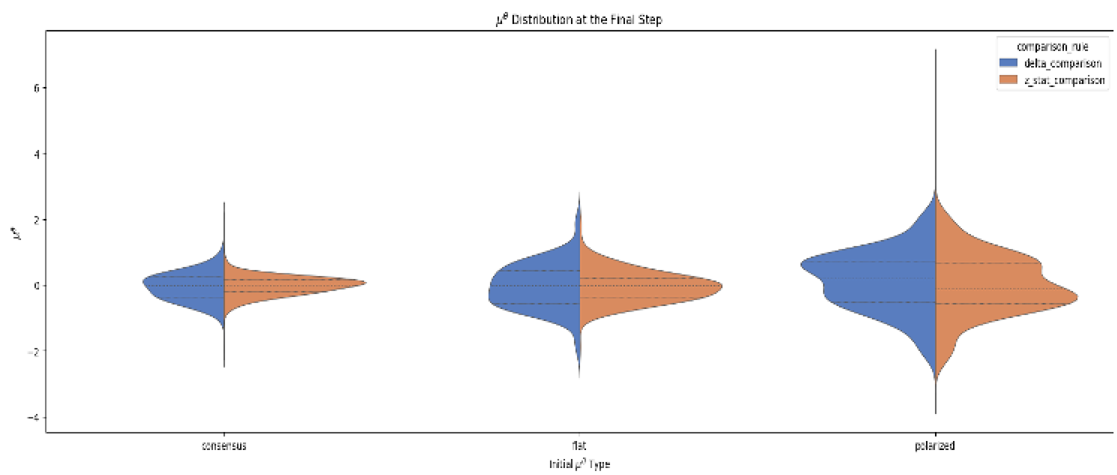

(b) Group ID Matching

Figure 17: μ^θ at Final Steps: Both Biased and Precise

work (see Figure 17a), while Group ID Matching does not lead to belief polarization (see Figure 17b). Nonetheless, it is noteworthy that the belief distributions become less densely distributed compared to Scenario B (see Figure 13b), where both information providers are biased but have different precision levels.

Figure 18 delves into the process of how agents achieved social consensus. The figure summarizes simulation outcomes from a simplified model using a Group ID Matching network, featuring 100 citizens (Blue and Red nodes) and 2 information providers (Green nodes). In this illustrative scenario, citizens start with polarized initial beliefs, and both information providers transmit equally precise messages — one biased and the other unbiased (Scenario A-3). The horizontal locations of nodes in Figure 18 corresponds to each agent's μ^{θ} values, with citizens employing a z-statistics learning strategy for source credibility assessment. Social interactions were implemented for a limited duration of 50 steps.

As depicted in Figure 18a, the social network among agents are highly segregated: Most communication channels exist between in-group members. There are two main mechanisms leading to the social consensus. The first dynamic is in-group convergence, whereby citizens connected to in-group members quickly converge, forming a consensus within their respective groups. Simultaneously, as previously mentioned, citizens connected to unbiased information providers swiftly adopt beliefs similar to those of the unbiased provider shortly after social interactions commence. Subsequently, they disseminate precise messages within their in-group, expediting the process of social consensus across the group membership. Over time, the number of citizens holding biased beliefs tends to decrease (See Figure 18b and 18c). These patterns vividly illustrate the positive impact of unbiased information reproduction through indirect (or second-handed) learning from citizens.

4.4 Implication and Discussions

This study delved into the the dynamics of social learning, focusing on the interplay between information providers, citizens, and their social networks. The model developed here combine individual-level source credibility assessment with information dissemination

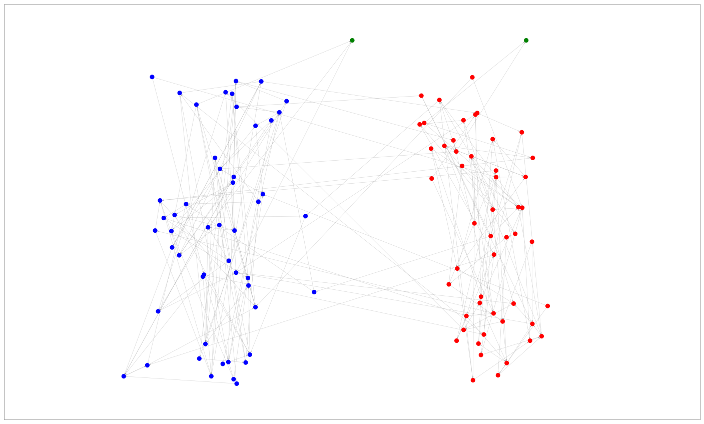

(a) Initial Step $t = 0$

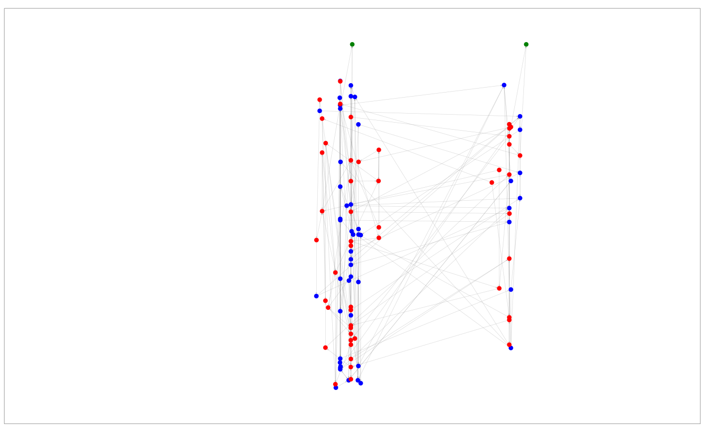

(b) Step $t = 25$

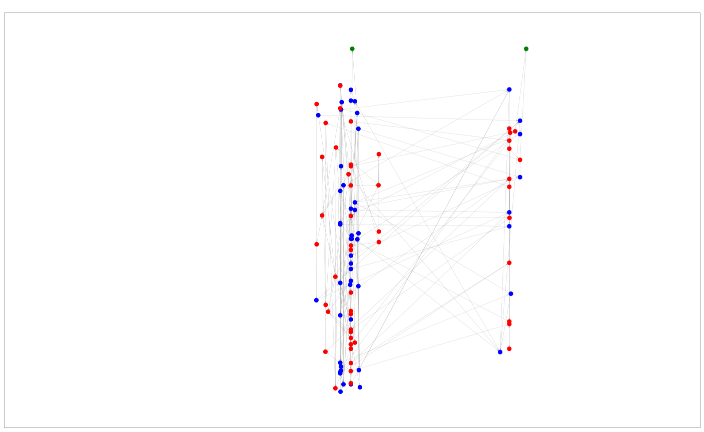

(c) Step $t = 50$

Figure 18: Step-by-step Illustration of Second-handed Learning: Scenario A-3

through social networks, offering insights into how citizens' beliefs are shaped, altered, and polarized amidst biased or unbiased information sources and diverse social interactions.

The baseline model's simulation results reveal that signal precision plays a pivotal role in individual-level source credibility evaluation. Particularly, the results imply that it is crucial for elite information providers to maintain consistent signals to citizens to attract citizens into their ideal points. Further, the baseline model implies that citizens become susceptible to misleading information when they cannot observe the differences in signal precision.

The social network model, on the other hand, highlights the importance of the between-citizen communication in building the social consensus among citizens. In addition, with the comparison of the average number of steps run until the models reach to the equilibria, I found that the social interaction significantly expedites the consensus building among citizens: the social interaction among citizens drastrically reduces the time required to reach the equilibria. There are still some exceptions that failed to create consensual post-communication belief distributions as found in Figure 13a and Figure 15b, but they still illustrates the tendency forcing citizens' beliefs centered around the true state of the world.

The findings from the social network models have significant implications for moderating ideological polarization and correcting misbeliefs. Scenario A, where unbiased and biased information providers are competing, represents the environment where the experts attempts to provide unbiased truth to citizens, but the adversarial disinformation provider attempts to draw citizens' beliefs away from the true state of the world. Figure 11 demonstrates that communication between citizens results in more homogeneous belief distributions, indicating a moderation of ideological polarization. Moreover, the results imply that the impact of misinformation on citizens' belief formation can be substantially reduced through active interpersonal interactions. Likewise, scenario B (with two biased information providers) represents the social learning under polarized media environment. Findings from scenario B reveals that the social interaction among the citizens significantly expedites the social belief formation.

Finally, baseline and social network model assume the information providers send out messages that perfectly represent their sincere beliefs. However, it is essential to consider the strategic motives of information senders, as they might want to influence citizens more significantly through message crafting efforts. For example, politicians have incentives to manipulate their messages to secure electoral victories. Similarly, adversarial trolls exploit these benefits by diverting citizens' attention from the true state of the world. In the next chapter, I will explore how speakers compose messages, specifically focusing on disinformation provider intentionally manipulating messages. This investigation aims to understand how the strategic message composition of elite information providers affects citizens' information source choice and belief formation.

5.0 Disruptive Jamming Model: Social Learning with A Disruptive Disinformation Provider

5.1 Introduction

As previously outlined, [50] emphasizes the role of speakers in the persuasion process. One important aspect of the speaker's effect on citizens' social learning procedure is how speakers formulate their messages. Politicians, for instance, aim to persuade voters with their messages, employing strategies such as advertising, credit claiming, and position-taking to secure reelection [119]. Similarly, presidents attempt to sway voters by priming the salience of issues they support [51]. Experts strive to lead public opinion to agreements by informing the true state of the world, while media outlets aim to provide correct information to the public. These goals are directly related to the utility function of elite speakers, driving their message composition strategies.

Elite speakers, as rational actors, have incentives to manipulate their messages to maximize their utility, particularly in political contexts. However, little is known about how information providers formulate their messages [50]. Politicians seek to move public opinion in the direction of their preferences through priming and framing, with messages often converging towards the political center to attract the most votes under a two-party system [49]. In the realm of media, circulation has become a crucial revenue source, incentivizing news media companies to tune news reports to attract more subscribers [15]. Studies have shown that consumer preferences strongly predict media slant [70].

However, message manipulation by elite speakers is constrained by two factors: their sincere ideals and reputation concerns. Political information reflects the preferences of its source, limiting the extent to which messages can deviate from their original ideals. Previous research on media bias has highlighted the significant impact of editorial board preferences or ownership on biases in news reports [86, 25, 141, 77]. Reputation concerns also play a role, as less precise signals are perceived as less convincing by information receivers [71]. To maintain credibility, speakers must compose consistent messages over time and conform to their audience's prior expectations [69].

In an era dominated by social media and online communication, the spread of misinformation and disinformation has become a significant concern. The ease with which false or misleading information can be disseminated to mass audiences poses a threat to the integrity of public discourse and democratic processes. Understanding the dynamics of how misinformation spreads and influences public opinion is thus of paramount importance in combating its adverse effects.

This paper proposes an interesting disinformation provider into the agent-based model proposed in the previous chapters. More specifically, I adopt the disruptive adversarial agent from [36], who aims to distract citizens' learning about the state of the world. This disruption idea reflects previous literature from various perspectives, including cheap-talk literature focusing on the jamming effects [125] and micro-targeting strategies in political communication [113]. Combining previous literature, the disruptive agent in this model aims to jam the spread of true information with micro-targeted messages, constrained by their sincerely held beliefs and reputation concerns.

The objectives of this essay are twofold: first, to investigate how disruptive messaging affects citizens' beliefs in different social contexts, and second, to explore the role of social networks in mitigating or exacerbating the impact of disruptive messaging. By systematically varying factors such as the credibility assessment strategies employed by citizens, the degree of belief polarization, and the structure of social networks, this chapter uncovers the nuanced dynamics that shape the evolution of political beliefs amidst disruptive messaging campaigns. Through a novel model of disruption and thorough interpretation of simulation results, this study contributes to a deeper understanding of the complex interplay between disruptive messaging, social networks, and belief formation processes, ultimately informing efforts to safeguard the integrity of public discourse in the digital age.

The structure of this study is as follows. First, I introduce the utility function of the disruptive jammer and its message composition strategy. Then introduces the behavioral rules how disruptive jammer works in the simulation setup. The simulations are implemented in the same environment from the previous two chapters, except it replaces information providers with an unbiased expert and a disruptive jammer. The simulation results are thoroughly explained with the comparison with other models including the ones in previous chapters. In the final section, I summarize the simulation results and their implications.

5.2 Model

5.2.1 Overview

In this chapter, I propose a model that explores how the behavior of disinformation providers affects the distribution of social beliefs. In previous chapters, biased information providers were defined as those who consistently spread their skewed perspectives to citizens. In this study, however, I introduce a distinct type of information provider that deliberately manipulates messages to maximize its utility after social learning processes. Specifically, I present a disinformation provider solely interested in disrupting citizens' understanding of the true state of the world within the social system, drawing from [36].

This model is built upon two main assumptions grounded in findings from preceding chapters. Firstly, all (dis)information providers in this model transmit messages of equal precision to citizens. Simulation results from earlier chapters suggest that it is consistently advantageous for information providers to send highly precise messages to persuade citizens effectively. Given that the goal of an information provider in this model is to steer the social belief distribution in its desired direction, all information providers are incentivized to finely tune their message parameters, as precision has been demonstrated to be effective in social persuasion.

Furthermore, (dis)information providers in this model assume that they are in competition with one another, without direct communication among citizens. Results from the social network model indicate the presence of second-hand learning effects, resulting in an increase in the number of semi-elite information providers within social networks. Moreover, substantial effort is required for information providers to observe the network structures, making learning about the network structure costly. Leveraging the effectiveness of second-hand learning observed in the social network model, an information provider can wield significant persuasion power by convincing citizens that it is a superior outlet compared to others.

Once again, as the objective of this chapter is to investigate the impact of strategically crafted disinformation within social networks by comparing results with previous findings, other elements of model environments such as social network structures, citizens' behavioral rules, and initial simulation setups remain unchanged.

5.2.2 Information Providers: Expert and Disruptive Jammer

In this model, there are two subclasses of information providers: an Expert and a Disruptive Jammer. The Expert, as delineated in previous chapters, represents unbiased and precise information providers. Experts can be likened to simplified versions of journalists or academic experts who directly observe the objective state of the world and deliver this observation to citizens without bias. Their goal is to transmit unbiased truth to citizens and aid in the formation of beliefs consistent with the true state of the world.

A disruptive jammer, on the other hand, represents the strategic disinformation provider, who aims to disturb citizens' learning about the truth. While misleading information providers in previous chapters disseminate incorrect messages due to their biased observation of the world, a Disruptive Jammer in this model is motivated to intentionally inject fabricated messages into the social network. Their utility function is defined by how much they push citizens' post-communication beliefs away from the true state of the world [36]. Additionally, I introduce constraint terms to the utility function, penalizing the jammer for composing messages deviating from their original beliefs. Thus, the jammer j's post-communication utility at time t, incurred from setting the message parameter sent to citizen i, $x_{j,i}^{(t-1)} \sim N(m^{(t-1)}, \sigma_j^2)$, at time $t-1$, is represented as:

$$U_j^{(t)}|x_{j,i}^{(t-1)} = \sum_{i=1}^{N} \left[((\mu_{i,t}^{\theta}|x_{j,i}^{(t-1)}) - \theta)^2 - (m^{(t-1)} - \mu_j^{\theta})^2 \right] \tag{9}$$

Here, $U_j^{(t)}|x_{j,i}^{(t-1)}$ denotes the jammer j's utility after sending messages from $x_{j,i}^{(t-1)}$, $\mu_{i,t}^{\theta}|x_{j,i}^{(t-1)}$ represents citizen i's belief about the world at t after observing the jammer's message composed at $t-1$, θ signifies the true state of the world, and μ_j^{θ} reflects the jammer's belief about the world.

To tailor disruptive messaging parameters, at the initial stage, Disruptive Jammers explore citizens' belief distribution. They surveil citizens' beliefs with a clustered view. Clustered categorization of citizens' attitudes is a familiar concept; for instance, in an election survey, respondents are often categorized based on their strength of party identification. Moreover, previous studies have demonstrated the effectiveness of micro-targeted messages in social persuasion [85, 83, 113, 139, 160]. During surveillance, Disruptive Jammers obtain categorized information about citizens' belief locations with k clusters: they only learn the cluster membership of citizens and the cluster's belief centroid locations. Although Disruptive Jammers lack information about the network structure, surveillance is repeated at certain time points; in this paper, surveillance occurs during the initialization stage and every fifth step throughout the simulation. Between the surveillance steps, the jammer cannot change the disruptive messaging setup until they update the surveillance result. In addition to the citizens' belief distribution, the jammer is aware of the expert's beliefs about the world.

Estimation of a cluster's centroid relies on its members' μ^θ values, with the centroid representing the grand mean of their μ^θ. Since Disruptive Jammers cannot directly observe the variance of each citizen's belief distribution, they estimate the clusters' variance using the variance of its members' μ^θ values: $\mu_k^{(t)} \sim N(\overline{\mu_{k_i}^\theta}, Var(\mu_{k_i}^\theta))$, where $\overline{\mu_{k_i}^\theta}$ denotes the mean of cluster k members' beliefs about the world status and $Var(\mu_{k_i}^\theta)$ represents the variance of the cluster members' $\mu_{k_i}^\theta$.

After the initial surveillance, a jammer pretends as if its $\mu_j^\theta = \mu_{k,t=0}^\theta$ for a citizen who is a member of cluster k. As citizens are not aware of the jammer's actual belief distribution, through this deceptive behavior, the jammer aims to make citizens believe the jammer is optimal outlet compared to the expert: under the z-statistics comparison strategy, $z_{i,j} < z_{i,o}$ and $\delta_{i,j} < \delta_{i,o}$ under the δ comparision rule, where $z_{i,j}$ and $\delta_{i,j}$ are the citizen i's credibility perception on the jammer j, and $z_{i,o}$ and $\delta_{i,o}$ are the vice versa. Immediately after the deception stage, the jammer changes the message parameters to the jamming setup.

Observing citizens' belief distribution with clustered views, the jammer tunes message parameters tailored to each cluster and its members. The jammer's expected utility for cluster k given its tailored message to members of k is:

$$E(U_{j,k}^{(t)}|x_{j,k}^{(t-1)}) = \Big[(E(\mu_{k,t}^{\theta}|x_{j,k}^{(t-1)}) - \theta)^2 - (m_{j,k}^{(t-1)} - \mu_j^{\theta})^2\Big] \tag{10}$$

Here, $m_{j,k}^{(t)}$ represents the mean of the messaging parameter targeted to cluster k at time t. As all information providers in this model are assumed to provide equally precise messages, citizens' beliefs given the jamming message are dominated by citizens' biased perceptions. Thus, the expected utility function becomes:

$$E(U_{j,k}^{(t)}|m_{j,k}^{(t-1)}) = \Big[(E(\mu_{k,t}^{\theta}|m_{j,k}^{(t-1)}) - \theta)^2 - (m_{j,k}^{(t-1)} - \mu_j^{\theta})^2\Big] \tag{11}$$

Moreover, during the deceptive messaging stage, citizens are expected to believe the jammer is the optimal information source. The expected θ beliefs of citizens who are exposed to messages with the mean of $m_{j,k}^{(t-1)}$ composed at $t-1$ are calculated as:

$$\begin{aligned} E(\mu_{k,t}^{\theta}|m_{j,k}^{(t-1)}) &= \mu_{k,t-1}^{\theta} + (\alpha m_{j,k}^{(t-1)} + (1-\alpha)\overline{x_{o,k}} - \mu_{i,t-1}^{\theta}) \left(\frac{\sigma_{k,t-1}^2}{\sigma_{k,t-1}^2 + s_{x_{t-1}}^2}\right) \\ &= \mu_{k,t-1}^{\theta} + (\alpha m_{j,k}^{(t-1)} - \overline{x_{o,k}}) \left(\frac{\sigma_{k,t-1}^2}{\sigma_{k,t-1}^2 + \alpha^2 - (1-\alpha)^2}\right) \\ &= \mu_{k,t-1}^{\theta} + (\alpha \overline{x_{j,t}} - \overline{x_{o,t}}) \left(\frac{\sigma_{k,t-1}^2}{\sigma_{k,t-1}^2 + 2\alpha - 1}\right) \end{aligned} \tag{12}$$

Here, $\alpha = 0.95$ represents the exploitation probability given by ε-greedy sampling strategy, ensuring the jammer is the optimal outlet for the members of cluster k. Additionally, the expert's message remains time-invariant as they always convey unbiased true state of the world. The jammer also knows that its competitor, the expert, is the unbiased and precise information provider. Thus, the expected mean of the expert's message is held at 0: $\overline{x_{o,t}} = 0$. Finally, the sample variance of entire messages drawn by a citizen at time $t-1$ is held at 1, $s_{x_{t-1}}^2 = 1$, as precise information providers have the variance set at 1 in previous chapters.

The optimal disruptive message for the cluster k, $m_{j,k}^{(t-1)'}$, can be derived from the first derivative of the equation 12:

$$m_{j,k}^{(t-1)'} = \frac{V\alpha\theta - \mu_{k,t-1}^{\theta} - V\alpha(1-V)\mu_j^{\theta}}{V\alpha - 1} \tag{13}$$

where, $V = \frac{\sigma_{k,t-1}^2}{\sigma_{k,t-1}^2 + 2\alpha - 1}$.

5.3 Simulation Results

This chapter delves into the social consequences of the jammer's disruption on the social learning process. To facilitate clearer comparisons, I conduct simulations under identical environments to those in the previous chapters, except for the changes in the types of information providers. Moreover, a noteworthy aspect of the disruptive jammer proposed in this chapter is their capacity to observe citizens' beliefs in clustered views. Essentially, as they allocate more resources to the surveillance stage, they should be capable of crafting finely tuned messages for targeted audiences. Thus, this section also investigates the effects of surveillance by varying the number of clusters they can observe. The number of clusters is varied by the powers of two, ranging from 2^0 to 2^8, and 500, considering the total number of citizens: $k \in \{1, 2, 4, 8, 16, 32, 64, 128, 256, 500\}$.

Additionally, the simulations are conducted under three sub-scenarios that vary the initial belief distributions of citizens: consensual, flat, and polarized distributions. Since the (dis)information providers in this model deliver equally precise messages and the jammer's messages are not primarily affected by its world status beliefs, simulations are not executed under scenarios that vary information providers' messaging parameters, as in the previous chapters.

5.3.1 Baseline Model: Introducing A Disruptive Jammer

To begin, I compare the pure Baseline model with the baseline model featuring a disruptive jammer. Once again, citizens in this setup are not allowed to communicate with each other, but they learn about the state of the world by sampling messages from information providers: the unbiased expert and the disruptive jammer.

Figure 19 illustrates citizens' post-communication belief distribution. Each subplot in Figure 19 describes the simulation results based on different initial belief settings of citizens. The vertical axis shows the distribution of $\mu^{\theta}_{i,T}$ at the time when citizens no longer update their beliefs ($t = T$). The horizontal axis represents the disruptive jammer's surveillance ability, indicating how many clusters the jammer can perceive. As expected from the disruptive message optimization function, Figure 19 clearly demonstrates the jamming effect. Particularly, as the number of clusters the jammer perceives increases, the jamming effects become more pronounced.

Moreover, two interesting findings emerge from the results in Figure 19. First, when the jammer is only capable of detecting the global mean of citizens' belief distributions (i.e., $k = 1$), the macro-level distribution of citizens' beliefs about the world tends to be double-peaked but within a narrow interval. This minor polarization is more distinct under the consensual initial setting scenario (on the left end of Figure 19a). This result suggests that because the jammer only observes the global mean, their messaging parameter serves as if citizens are observing two information providers competing with equally precise but with different message means.

The second, and more intriguing, finding is the presence of a clear ceiling effect regarding the jammer's surveillance ability: there is a threshold point beyond which the effectiveness of jamming messages does not drastically increase. In Figure 19, the tipping point is at a surveillance level of 8; after surpassing this point, the posterior belief distributions look highly similar to each other. This observation is interesting because it aligns with the political science tradition of categorizing American voters into 7 groups depending on their attachment to each party: Strong Democrat, Democrat, Democrat-leaner, independent, Republican-leaner, Republican, and Strong Republican. This implies that these categories can serve as the optimal surveillance ability level in crafting micro-targeted messages in political information outlet competition.

Figure 20 illustrates the ratio of how many messages were sampled from the jammer in the last time step compared to the unbiased expert. The jammer successfully attracts citizens, leading them to believe it outperforms the unbiased expert. Likewise, it supports the observation of the ceiling effect of the surveillance ability across all sub-scenarios. Once the surveillance ability exceeds 8, the exploitation rate does not increase steeply. The successful attraction results from the deceptive message sending after the initialization stage: the jammer successfully misleads citizens by pretending its belief is located at each cluster's centroid.

Another interesting finding is that disruption outperformed when citizens employ the δ

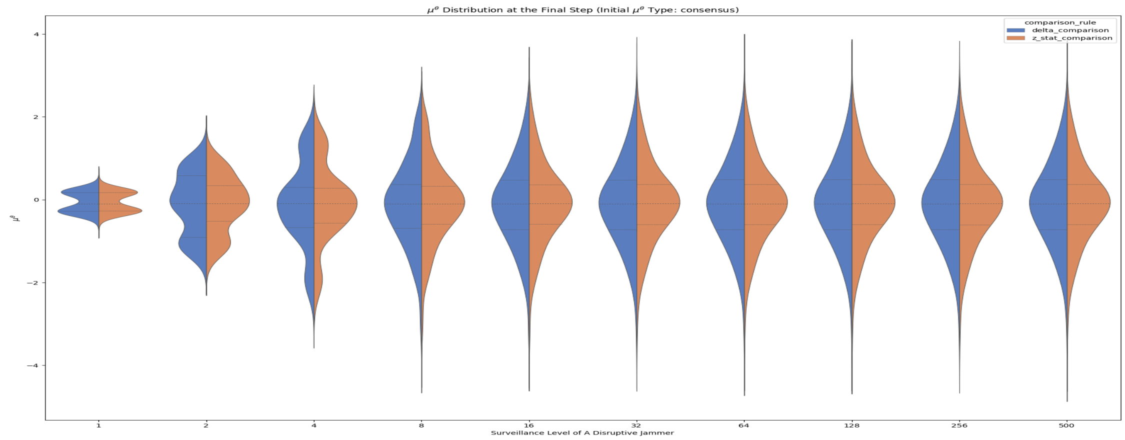

(a) Consensual Initial Beliefs

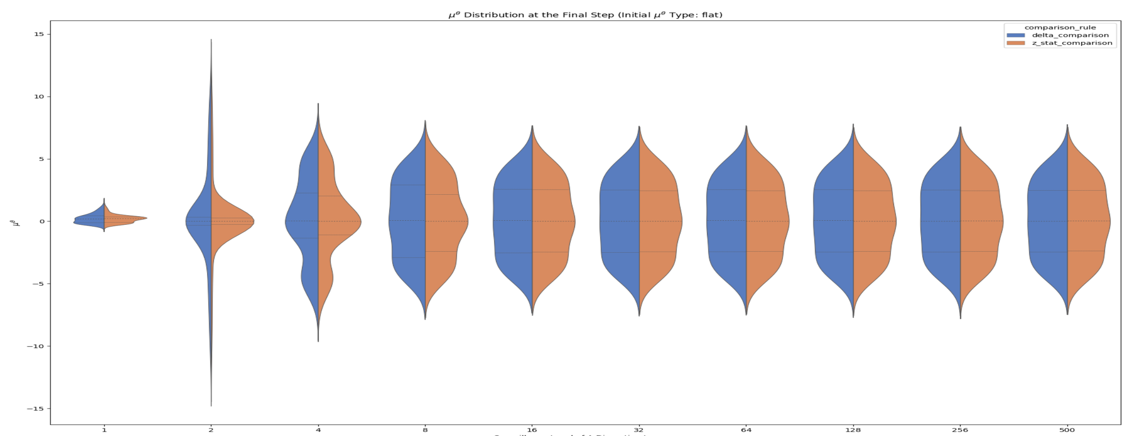

(b) Flat Initial Beliefs

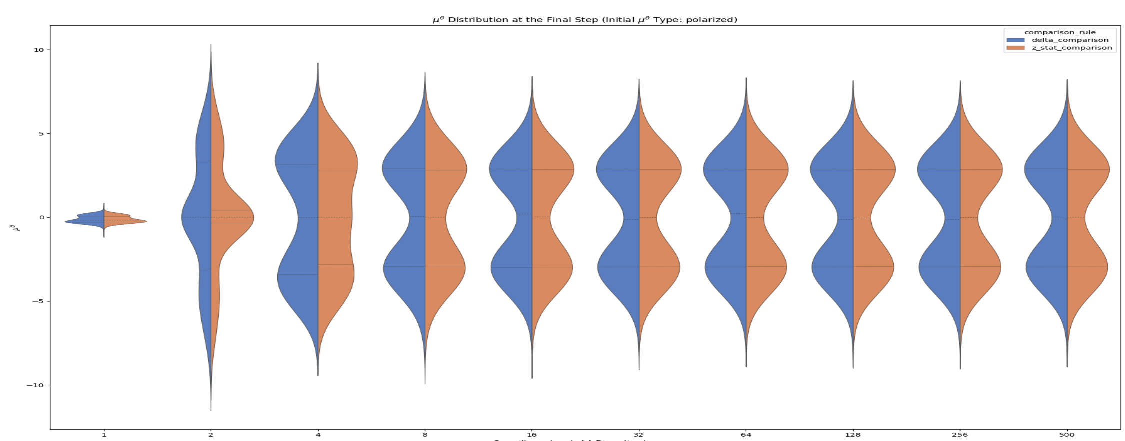

(c) Polarized Initial Beliefs

Figure 19: μ^θ at Final Steps: Baseline Model

comparison strategy (see the upper row of Figure 20a - 20c) for source credibility evaluation compared to when they use the z-statistics comparison strategy (see the bottom row). This result stems from the nature of each credibility assessment rule: under the δ comparison rule, credibility perception is dominated by the similarity of messages with respect to preexisting beliefs, while the z-statistics comparison strategy strikes a balance between bias detection and message precision.

Finally, Figure 21 summarizes the average number of steps required to reach the 'steady-state' where citizens do not update their beliefs further. Under all subscenarios, the simulation lasts approximately 100 steps and reaches the steady-state relatively quickly. This result is closely related to the findings from the exploitation pattern: as the jammer successfully attracts citizens relatively early in the simulation steps, they swiftly form static beliefs fixed at their initial values, even though the unbiased expert continues to share unbiased information. Due to the fast belief consolidation, there is no evidence illustrating the effectiveness of jamming messages in terms of determining the speed of belief consolidation.

5.3.2 Social Network Model: Random 2 Matching

Now, I examine the effects of jamming messages on social belief formation under a scenario where citizens can communicate with fellow citizens. The first network I explore is the Random 2 Matching network, where each citizen is connected with 2 other agents regardless of their types. In contrast to the baseline model where elite-level information providers are connected with every citizen, the communication partner selection in this model is entirely random. Since communication between citizens and information providers is not guaranteed in this model, the outcomes of interest are limited to the social belief distribution and the speed of belief consolidation.

Figure 22 summarizes the post-communication belief distribution under the Random 2 Matching social network. Consistent with findings from the social network model, opening communication channels with other agents with diverse beliefs fosters social consensus building even in the presence of the adversarial disruptive agent. While disruptive messages lead to the consolidation of initial beliefs even after the social learning process, the Random 2

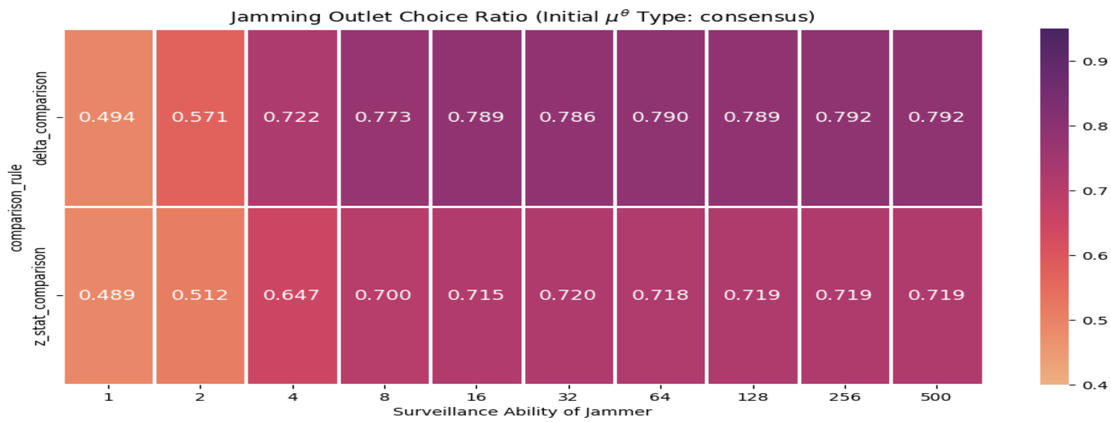

(a) Consensual Initial Beliefs

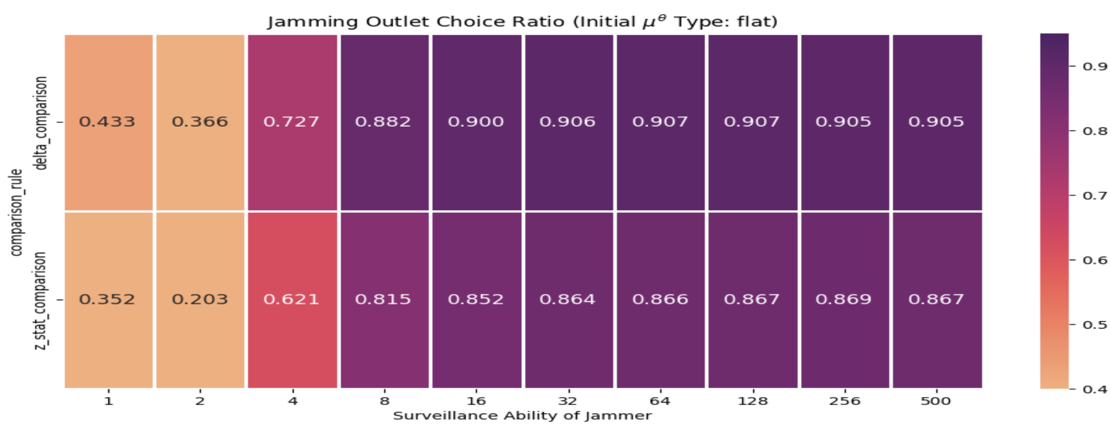

(b) Flat Initial Beliefs

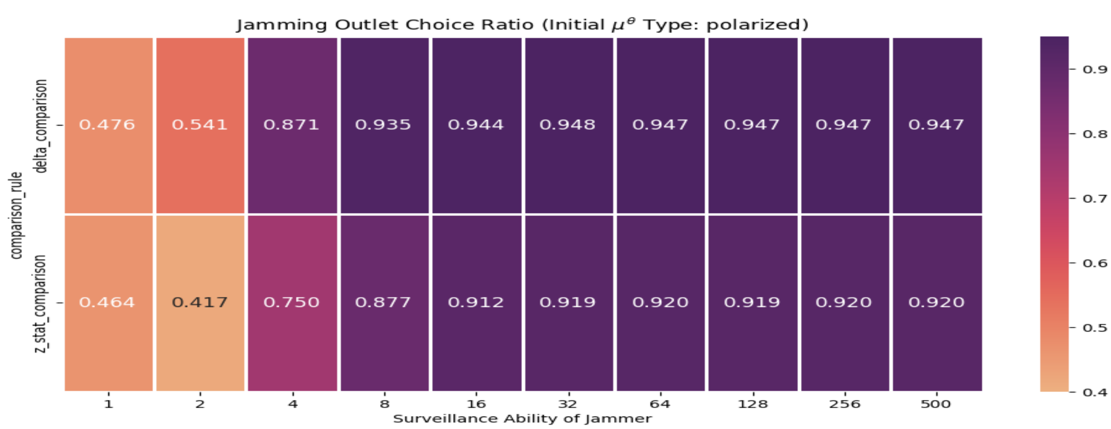

(c) Polarized Initial Beliefs

Figure 20: Jamming Outlet Choice Ratio: Baseline Model

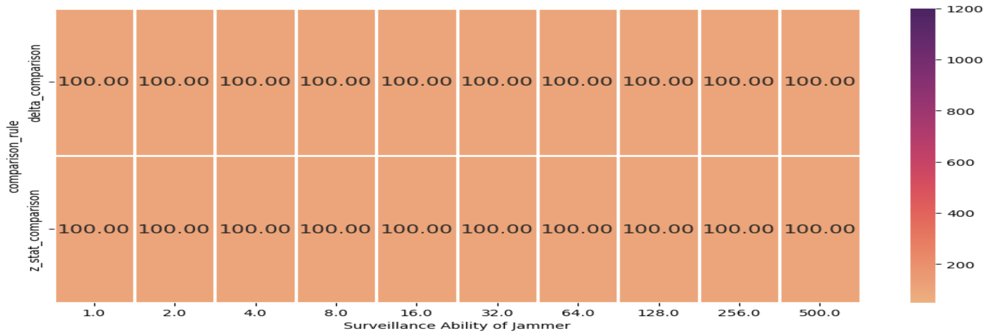

(a) Consensual Initial Beliefs

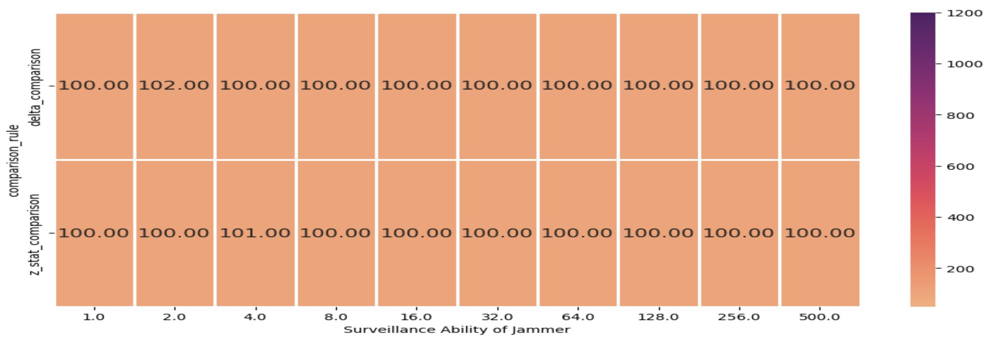

(b) Flat Initial Beliefs

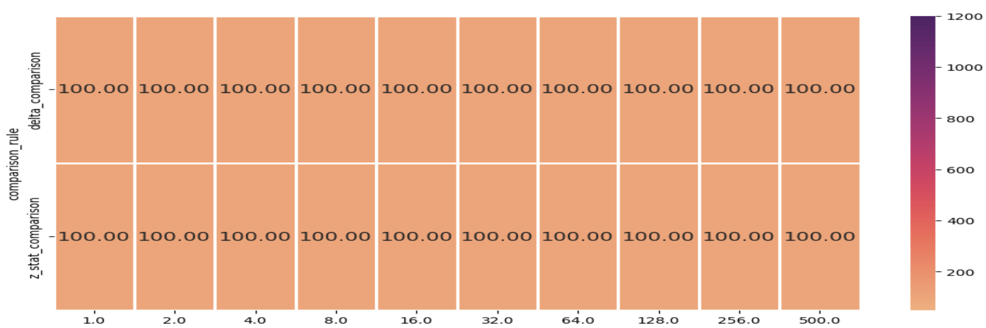

(c) Polarized Initial Beliefs

Figure 21: Total Steps Run: Baseline Model

Matching network facilitates the mixing of diverse perspectives and converges them toward the true state of the world in all scenarios.

An interesting finding is that the jamming effect is more evident when citizens rely on the δ comparison strategy than the z-statistics method. The blue density plots in Figure 22 illustrate the posterior belief distribution at the end of social interaction when citizens rely on the δ comparison rule. It is noteworthy that these density plots tend to have longer and thicker tails than their comparison sets in orange, indicating that jamming was slightly more effective under this condition than the z-statistics comparison rule. As previously explained, this difference arises from the nature of the method, which focuses solely on the message receivers' biased perception of information sources.

Finally, surveillance ability does not affect the shape of the posterior belief distribution, as jamming is not effective under this environment.

Figure 23 illustrates the average number of total steps run until citizens reach the steady-state. Once again, citizens' beliefs are consolidated relatively quickly after approximately 110 steps. The number of steps run under the social network model was similar to this, implying that the jamming messages do not significantly affect the increase or decrease in the belief consolidation speed. Likewise, the surveillance ability does not influence the changes in speed.

5.3.3 Social Network Model: Group ID Matching

I also conducted simulations under the Group ID matching network, where all agents are assigned group identities based on their initial beliefs, and there are higher chances of being connected with in-group members. Previous chapters revealed that even though the network is designed to deepen homophily, the presence of elite information providers can foster consensus building due to interactions between in-group consensus and second-hand learning from optimal information sources. Similar patterns are observed in Figure 24: despite the network structure representing homophily, it does not hinder social consensus building, even when disruptive messages are consistently delivered to citizens. Additionally, as observed in the random 2 matching network, the belief distribution becomes slightly wider

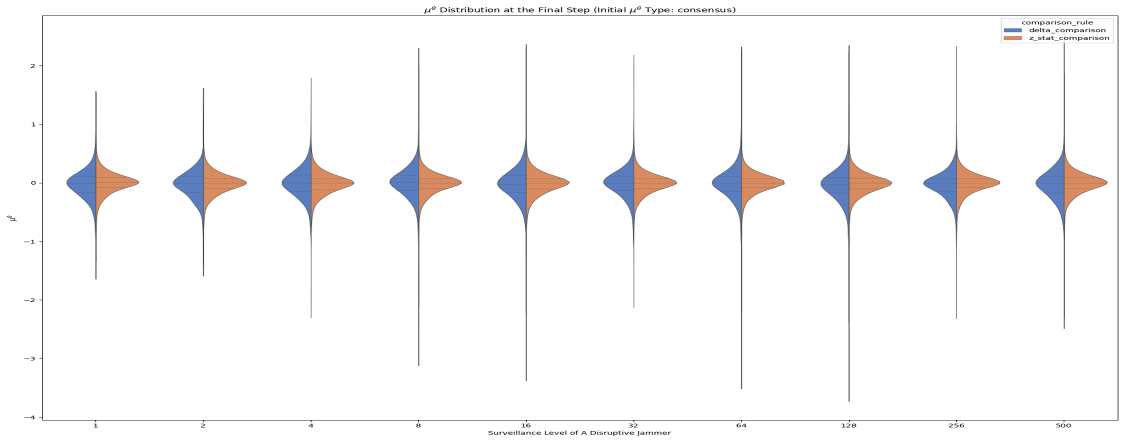

(a) Consensual Initial Beliefs

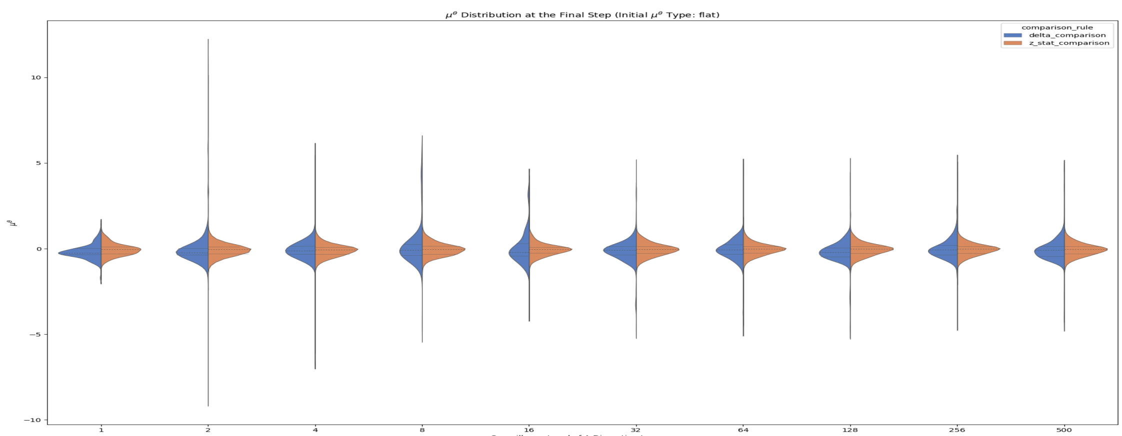

(b) Flat Initial Beliefs

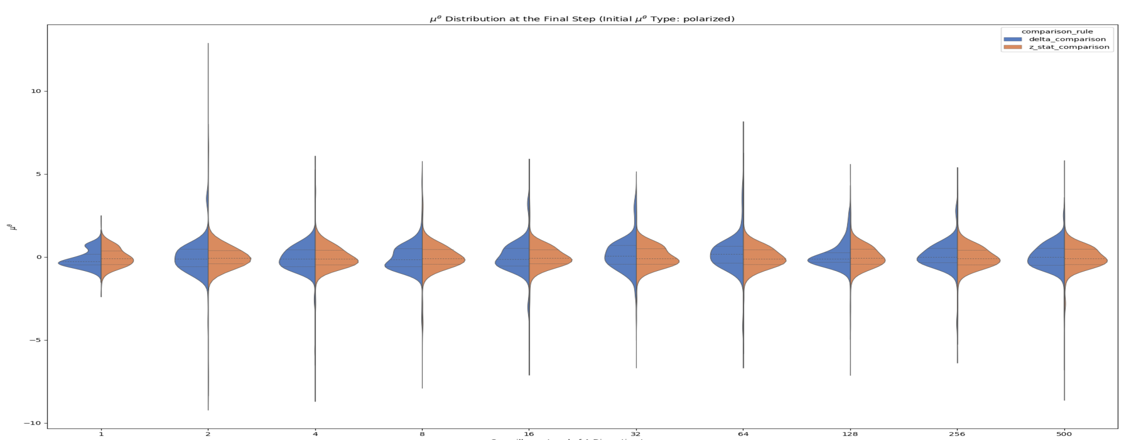

(c) Polarized Initial Beliefs

Figure 22: μ^{θ} at Final Steps: Random 2 Matching

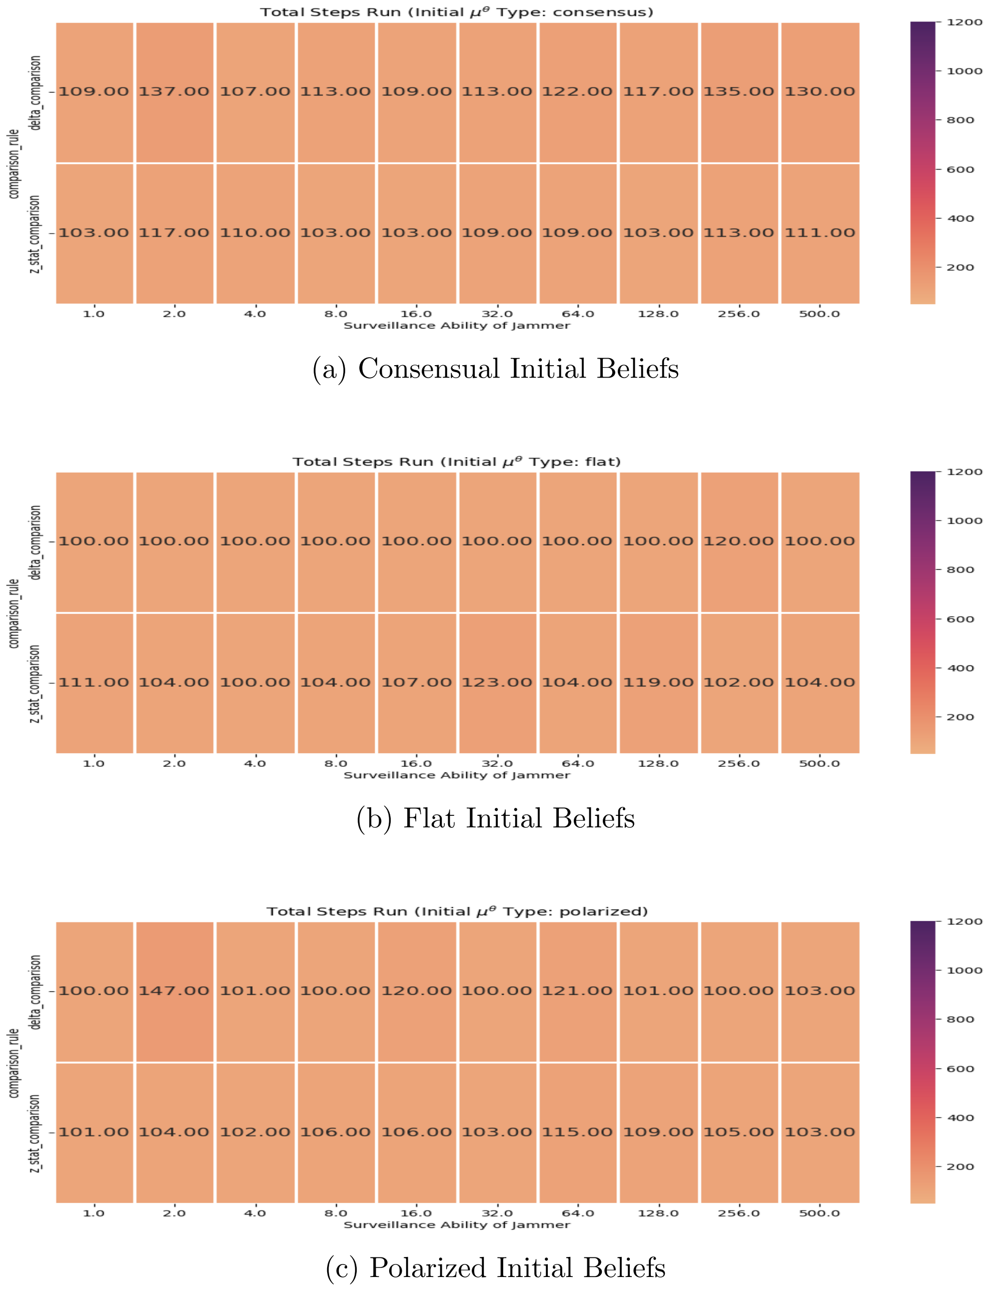

(a) Consensual Initial Beliefs

(b) Flat Initial Beliefs

(c) Polarized Initial Beliefs

Figure 23: Total Steps Run: Random 2 Matching

when citizens employ the δ comparison strategy. Furthermore, the jammer's surveillance ability does not lead to changes in citizens' posterior belief distribution.

Figure 25 compares the average number of steps run depending on the surveillance ability of the jammer and the initial belief distribution of citizens. Once again, even under the homophilous social network, citizens reach social consensus quickly. Comparing to the social network model's simulation result, the jamming messages do not significantly affect the speed of consensus building. Likewise, as there is no jamming effect, the surveillance ability does not lead to changes in speed.

An additional supplementary analysis is conducted in this section. Previously in the second study, the finding of social consensus building in the group ID matching was interpreted as the interaction effect of in-group consensus building and second-handed learning. Especially notable was that social agreements were built even with polarized initial beliefs, expected to deepen belief polarization. Linking these observations from the previous chapter with this section's results, further exploration is undertaken to determine if the results change with greater initial belief polarization. While the polarized initial belief had two modes at -3 and +3, here the gap between two clusters is expanded by three times at -9 and +9. As the space between two peaks increases, it should provide a larger room for the jammer to tailor more diverse and appealing messages to citizens. Thus, combined with the homophilous nature of the network building mechanism, significant jamming effects deterring social consensus building are expected.

Figure 26 illustrates the simulation results with greater polarized initial beliefs, supporting the conjecture with increased numbers of clusters. Additionally, similar to the jamming baseline model, the jamming effect becomes more evident when the jammer's surveillance ability is over 8. While the jamming baseline led citizens to stick to their initial beliefs, it creates a belief distribution with three clusters peaked around -10, 0, and +10, respectively. This result implies that the jammer can exploit benefits from greater levels of belief disparities between citizens and clustered network setup, leading to the fractionalization of the belief distribution in the end. Likewise, the increase in the efficiency of jamming messages is not linear: it shows similar levels of jamming effect when the surveillance ability goes over 8.

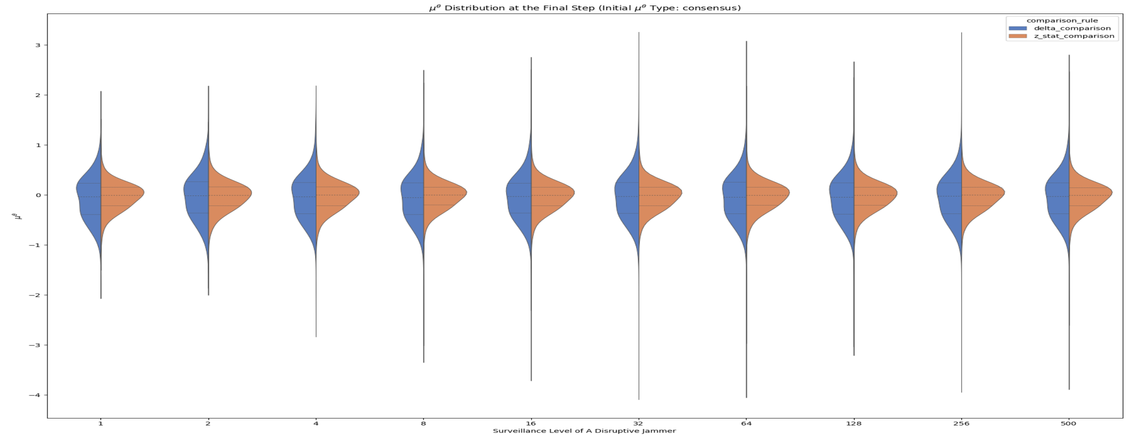

(a) Consensual Initial Beliefs

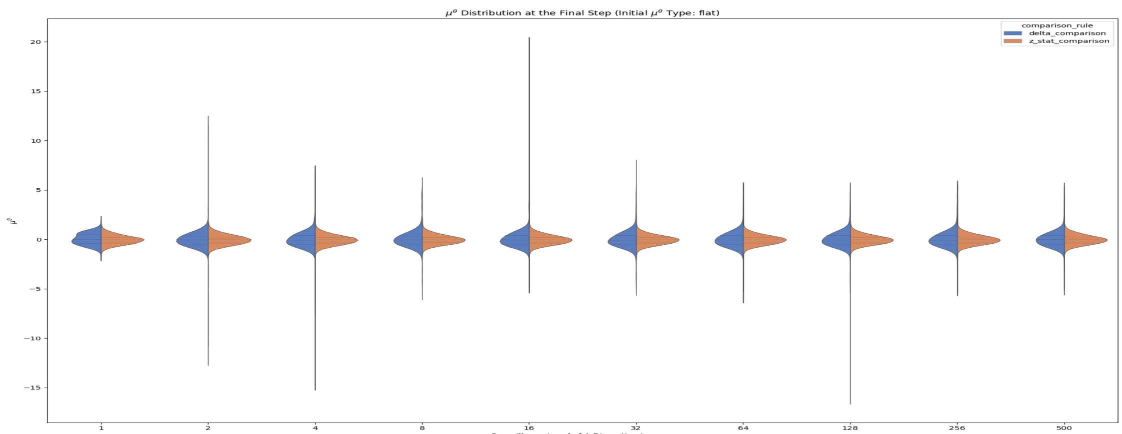

(b) Flat Initial Beliefs

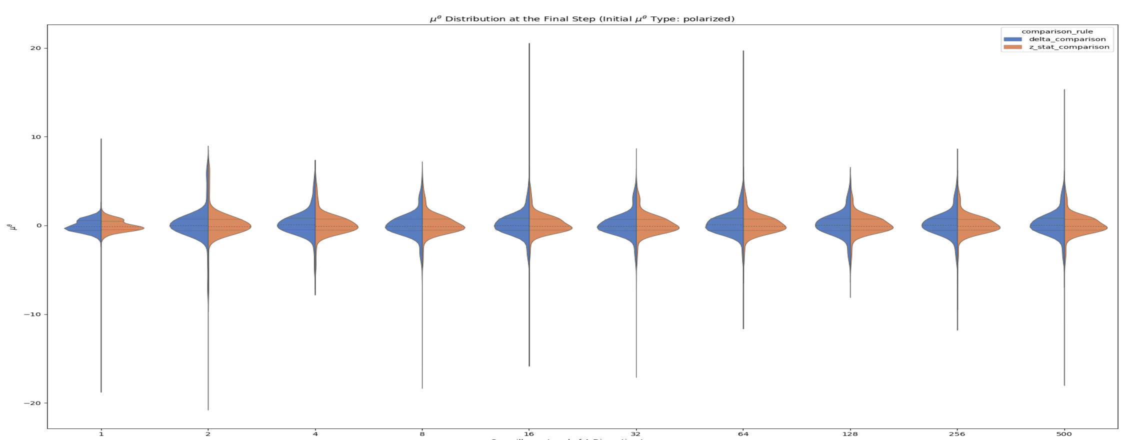

(c) Polarized Initial Beliefs

Figure 24: μ^θ at Final Steps: Group ID Matching

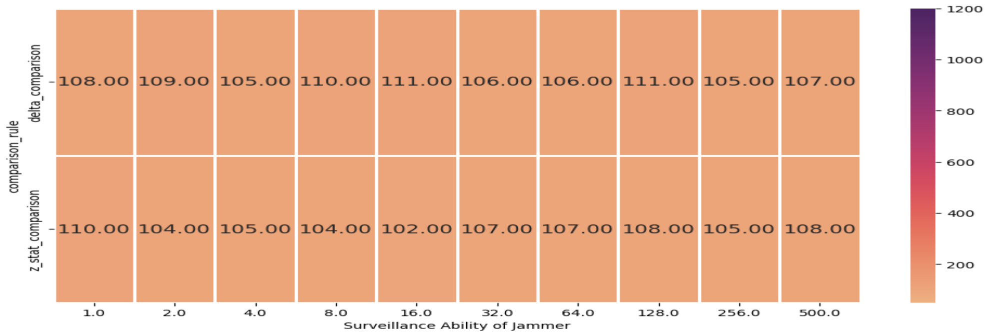

(a) Consensual Initial Beliefs

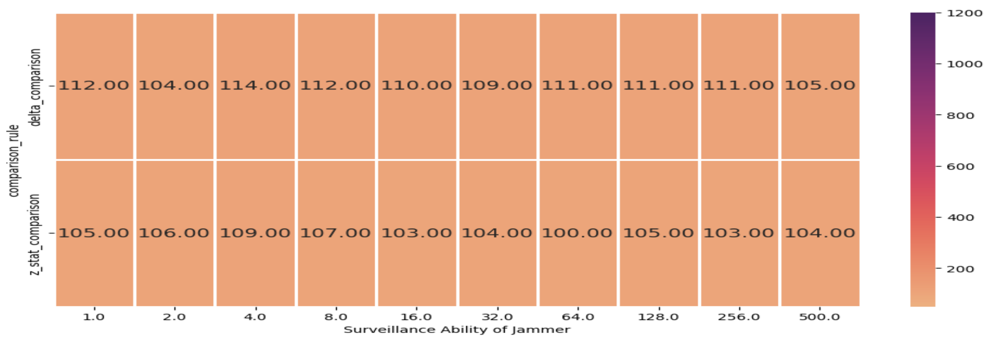

(b) Flat Initial Beliefs

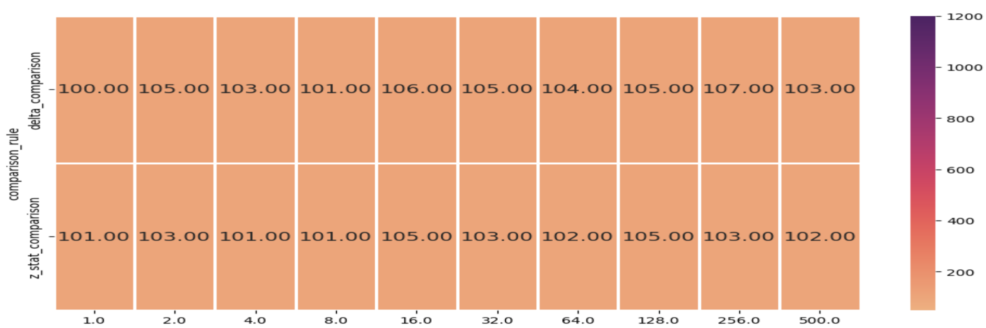

(c) Polarized Initial Beliefs

Figure 25: Total Steps Run: Group ID Matching

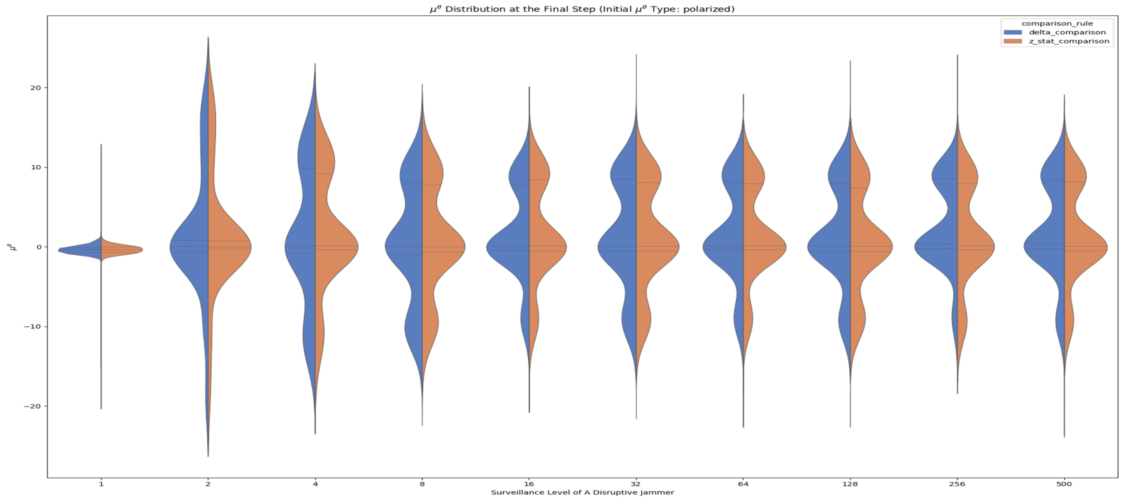

Figure 26: μ^θ at Final Steps: Group ID Matching with Further Initial Belief Polarization

5.3.4 Social Network Model: Extended Network

Finally, I examine the jamming effect under the extended network, which combines the random 2 matching and the baseline model networks. In this network, each citizen is connected with 4 agents: two elite-level information providers and two randomly matched citizens. Under this structure, at each time step, citizens randomly choose the type of communication partner and sample messages from the chosen type of neighbors. Thus, in this part, I compare the simulation outcomes, including the social belief distribution, the ratio of citizens' choices of the disruptive jammer, and the number of steps required to reach the steady-state.

Figure 27 shows the social belief distributions under each subscenario. Consistent with previous jamming social network models, this model also creates consensual posterior belief distributions. Likewise, as jamming messages are suppressed by consensus-building forces, the jammer's surveillance ability has no effect on social belief formation. This result highlights the importance of social support in political information processes. When citizens are isolated and only able to communicate with elite-level information providers, they become vulnerable to being attracted by disinformation, deterring social agreement. But when citizens can communicate with other fellow citizens, social agreement becomes inevitable, and the expert's messages become salient, correcting the disinformation.

The aforementioned superiority of expert messages in this model is supported by Figure 28. It summarizes the proportion of messages sampled from the jammer at the final interaction step (i.e., $t = T$). None of the subscenarios indicate that more than 50% of messages are drawn from the disruptive jammer, indicating that most citizens found experts' messages more truthful. An interesting observation is that citizens become least susceptible to disinformation when their initial beliefs are flat. Likewise, Figure 27b shows the shortest distribution tails compared to the other scenarios. This implies that the jammer found it more difficult to distract citizens, and the social force to create agreement is greater than in other scenarios.

The average number of steps run is depicted in Figure 29. Simulations under the extended network last longer than the other jamming models. This reflects two aspects. First, it is the product of the behavioral rule and the network structures. As a citizen is connected with both types of agents, in each step, the agent picks its communication partner type, which slightly slows down overall learning speed by complicating interactions at each time step. Additionally, it implies that even though jamming messages are suppressed in the end, the result suggests that the jamming messages could have deterred the learning process in earlier stages, leading to a slowing down of the learning process. Specifically, the result under the polarized initial belief setting is interesting to note as it tends to take longer than its comparison sets in other scenarios. There are two interesting observations in Figure 29. First, when citizens rely on the δ comparison rule, they tend to reach the steady-state slowly compared to the z-statistics comparison rule. As previously explained, jamming messages work better with δ comparison. Furthermore, the left three columns of the figure's upper row show a gradual increase in total steps, implying that the jamming —even though the effect was negligibly minimal —slows down the social learning process, at least in these subsets of scenarios.

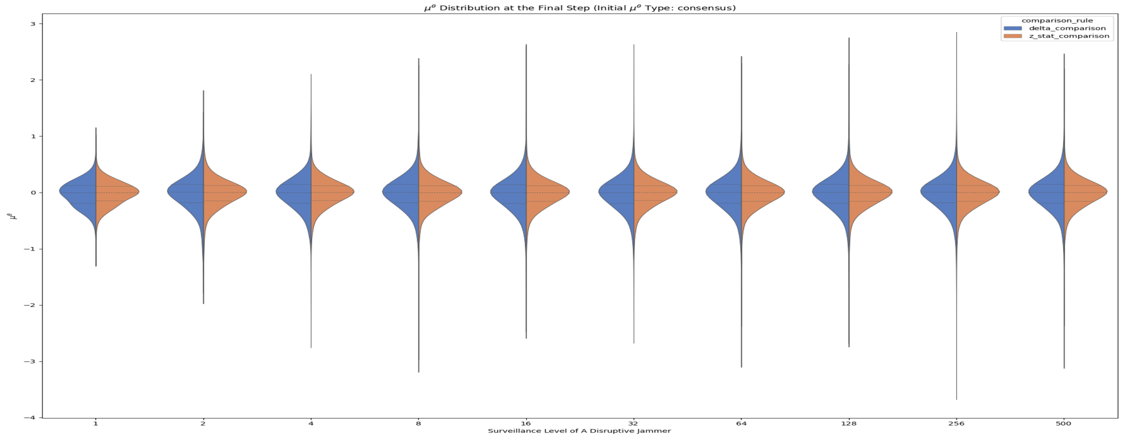

(a) Consensual Initial Beliefs

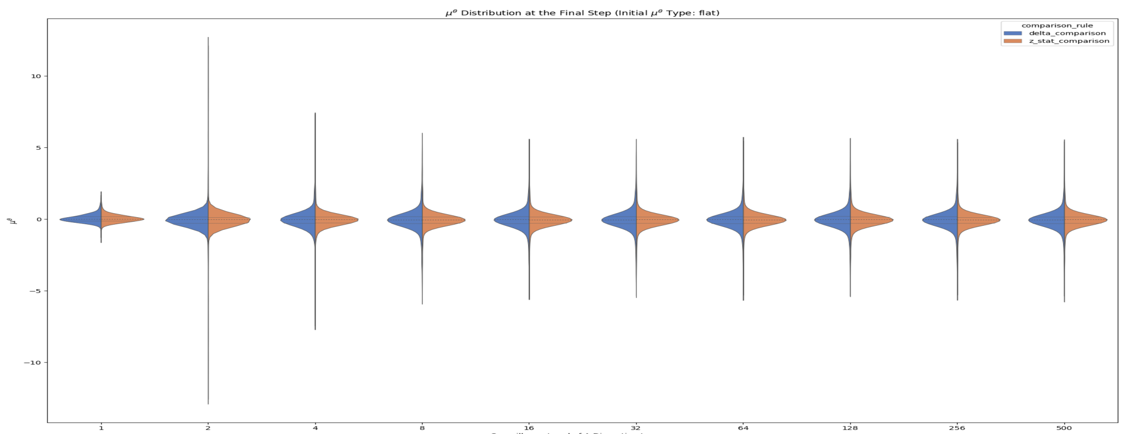

(b) Flat Initial Beliefs

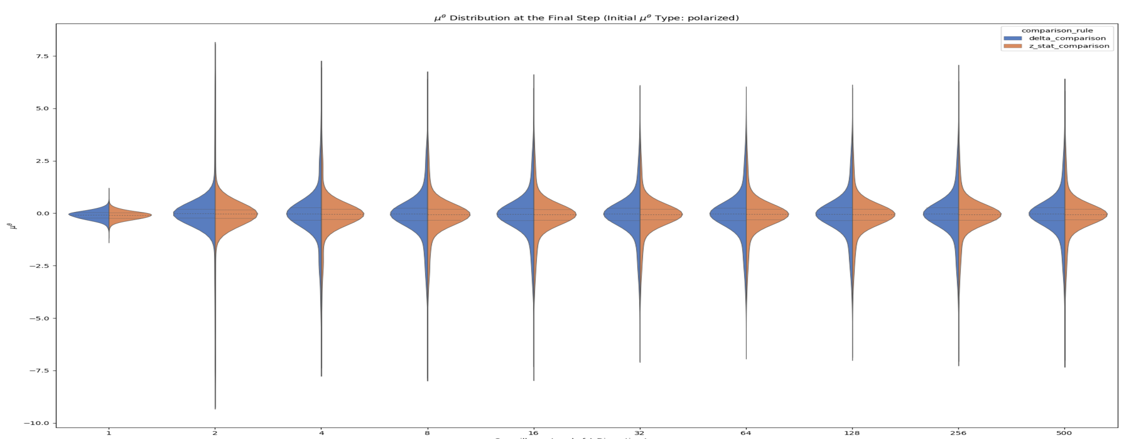

(c) Polarized Initial Beliefs

Figure 27: μ^θ at Final Steps: Extended Network

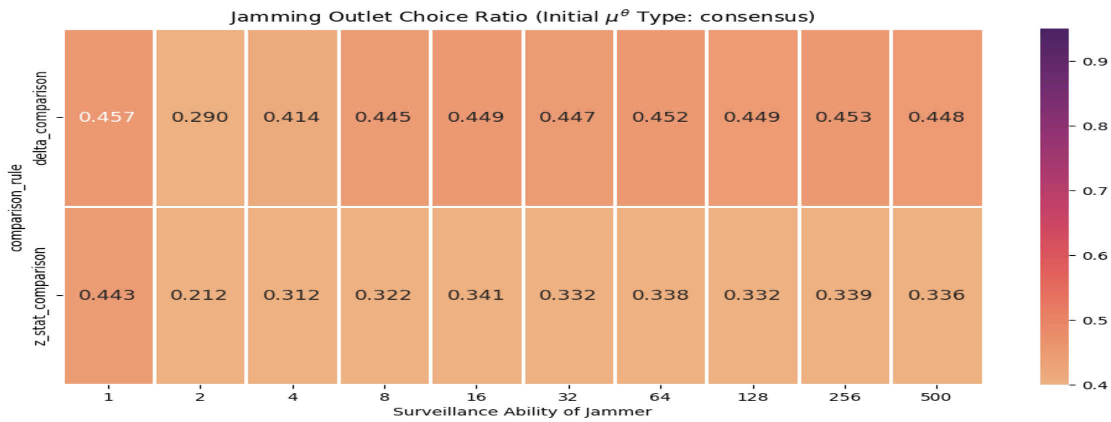

(a) Consensual Initial Beliefs

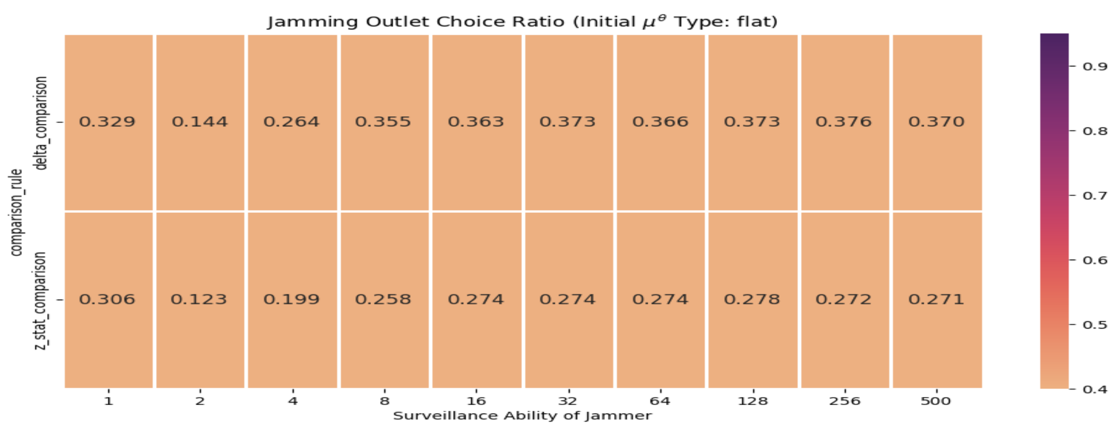

(b) Flat Initial Beliefs

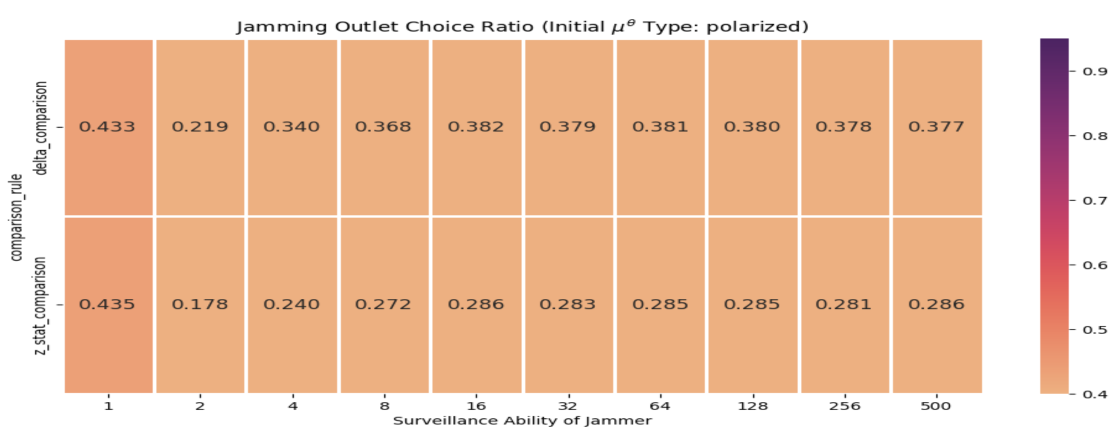

(c) Polarized Initial Beliefs

Figure 28: Jamming Outlet Choice Ratio: Extended Network

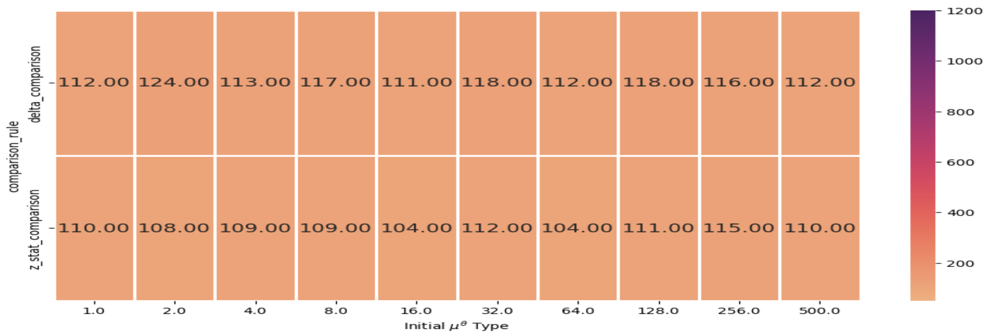

(a) Consensual Initial Beliefs

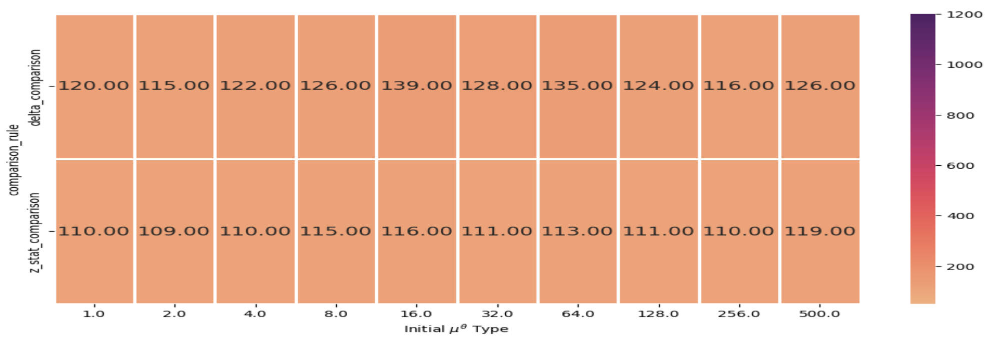

(b) Flat Initial Beliefs

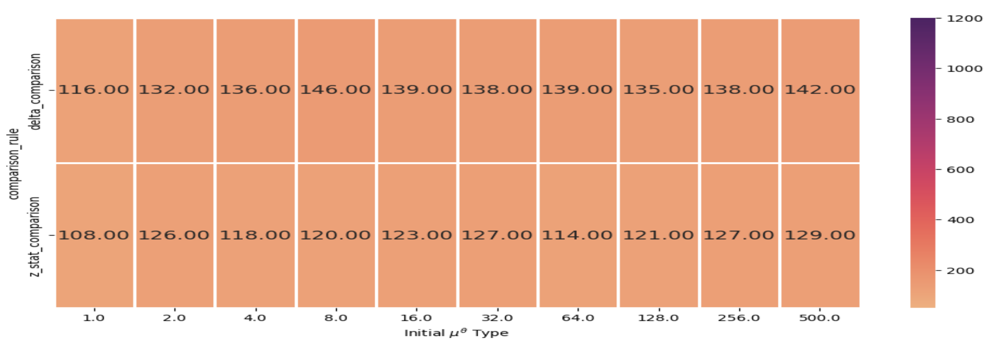

(c) Polarized Initial Beliefs

Figure 29: Total Steps Run: Extended Network

5.4 Discussion and Conclusion

This study delved into the dynamics of disruptive messaging within the realm of social belief formation. Through extensive simulations, the impact of a disruptive jammer on citizens' beliefs was explored across various network structures and initial belief distributions. The findings from the agent-based simulations shed light on the intricate interplay between disruptive messaging, social networks, and belief consolidation processes.

Primarily, the simulation results underscore the resilience of social consensus-building mechanisms against disruptive messaging. Despite the presence of an adversarial troll delivering disinformation to deter social learning, social networks ultimately facilitate the convergence of citizens' beliefs towards the true state of the world. This resilience underscores the pivotal role of social interaction in mitigating the influence of disinformation campaigns and preserving the integrity of collective beliefs, as evidenced in earlier chapters.

Moreover, I found that the effectiveness of disruptive messaging is contingent upon several factors, including the credibility assessment strategy employed by citizens and the degree of belief polarization in the initial belief distribution. The simulations revealed that citizens relying on the δ comparison strategy are more susceptible to disruptive messaging, as it concentrates solely on the perceived similarity of messages to pre-existing beliefs. Additionally, greater belief polarization provides fertile ground for the jammer to exploit, leading to more pronounced disruptions in political belief formation.

The structure of social networks also emerges as a crucial determinant in shaping the impact of disruptive messaging. While homophilous networks may exacerbate belief polarization, they concurrently facilitate social consensus-building through interactions with diverse opinion holders. Conversely, networks with more random connections afford citizens access to a broader range of perspectives, bolstering resilience against disruptive messaging and promoting collective rationality.

In conclusion, this study underscores the imperative of comprehending the interplay between disruptive messaging, social networks, and belief formation processes. By elucidating the mechanisms underlying the spread of disinformation and the resilience of social consensus-building, it furnishes valuable insights for strategizing against the proliferation of misinformation. Particularly, the insights derived from the findings, especially evident in Figure 26, suggest that disinformation providers may exploit the fractionalized citizens' beliefs to prolong their influence. Thus, this study provides critical insights into the conditions under which disinformation can persist and magnify its impact.

6.0 Conclusion

Appendix Appendix to Social Network Model

Extending from the social network structures illustrated in the main text, this section explores how the results changes by varying the degree of homophily among agents. While agents are connected with same group members with the probability of 90%, here, I lower the probability to 70%. This serves as the midpoint between the Random 2 Matching and the Group ID Matching illustrated in the Section 4.3.2.

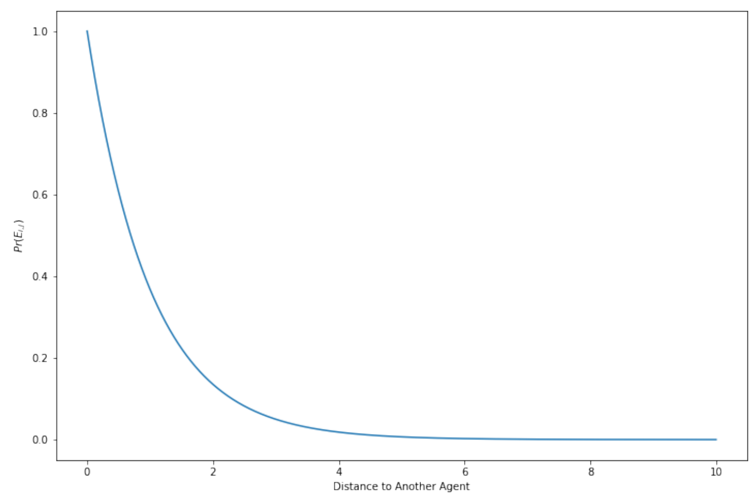

Figure 30: Probability Distribution over the Distance between Agents

Another variation made here is matching agents based on the distance between agents. Suppose Citizen i tries to connect with another Citizen j. To determine if Citizen i samples messages from j or not, I first calculated the euclidean distance between two agents ($D_{i,j}$). Then I assigned the probability of building an out-edge from i to j with following exponential function: $Pr(E_{i,j}) = e^{(D_{i,j})}$, where $Pr(E_{i,j})$ is the probability of having an outedge from i to j and $e^{(D_{i,j})}$ is the natural exponential with the power of the distance between i and j. The probability distribution is illustrated in Figure 30. With the probability obtained from the exponential function, a random draw from a binomial distribution was executed: if the outcome is 1, the outedge from i to j is built; otherwise, find another agent k and repeat the same process until it has 2 outedges in total.

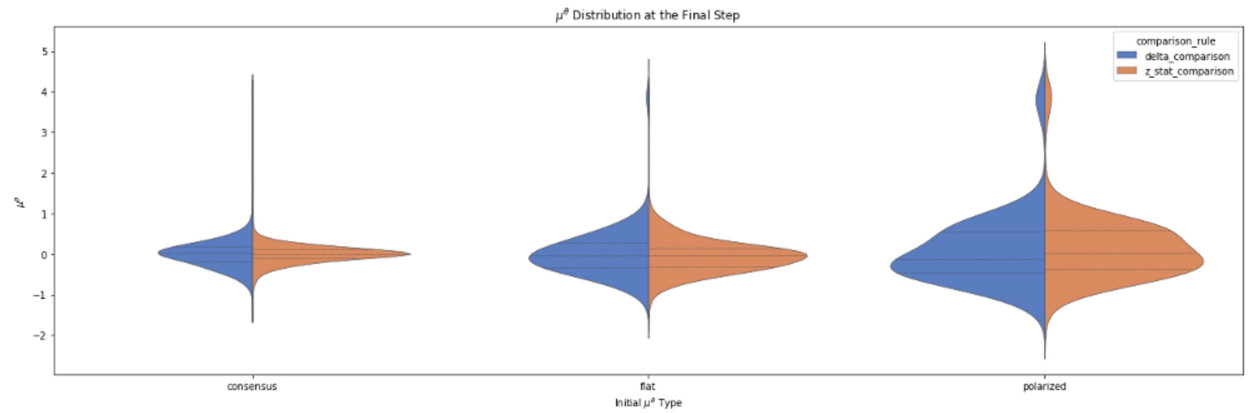

(a) Group ID 70% Matching

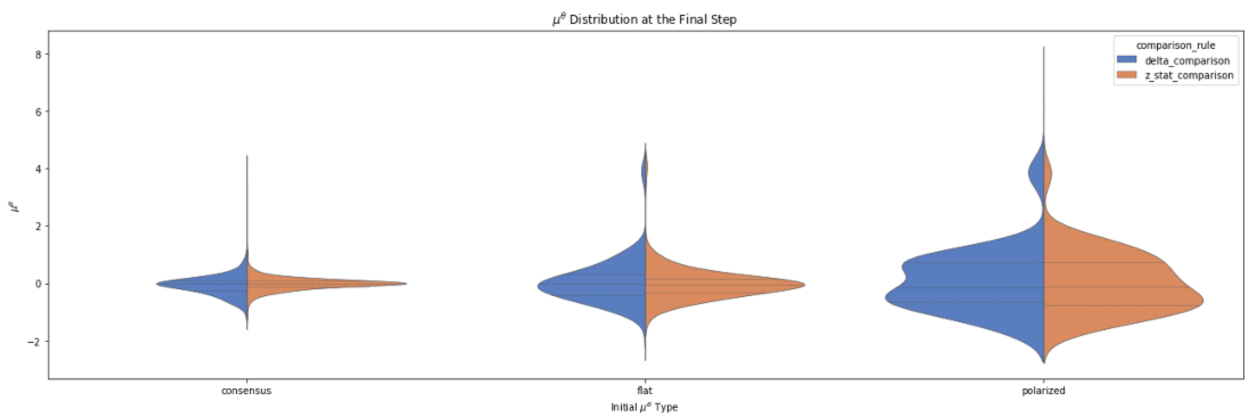

(b) Belief Distance Based Matching

Figure 31: μ^θ at Final Steps: Scenario A

Figure 31 and 32 shows the results from two network structures. The findings are consistent with ones from the Section 4.3.2: The social interaction among citizens does decrease the belief polarization (Figure 31). The swiftness to reach the consensus is slightly slowed down with comparison to the baseline model, but they are equally as fast as the ones from the Section 4.3.2.

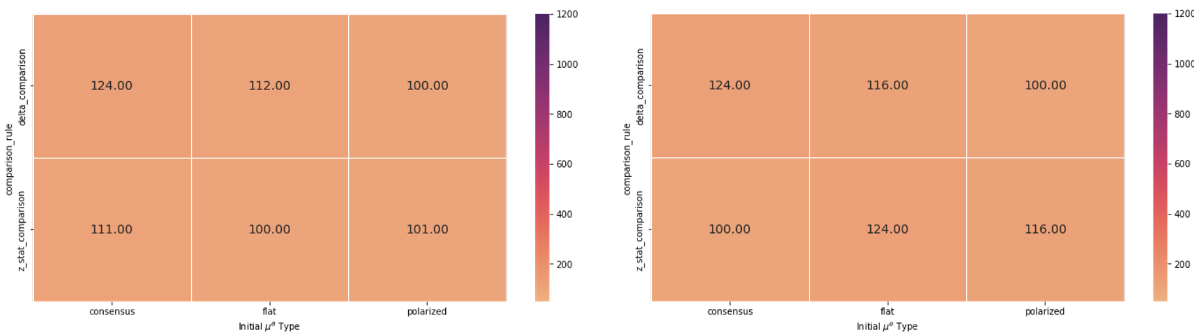

(a) Group ID 70% Matching

(b) Belief Distance Based Matching

Figure 32: The Average Steps Run Until The Equilibrium: Scenario A

Bibliography

[1] Daniel Acuna and Paul Schrater. Bayesian modeling of human sequential decision-making on the multi-armed bandit problem. In *Proceedings of the 30th annual conference of the cognitive science society*, volume 100, pages 200–300. Washington, DC: Cognitive Science Society, 2008.

[2] Johnathan Adams, Gentry White, and Robyn Araujo. The role of mistrust in the modelling of opinion adoption. *Journal of Artificial Societies and Social Simulation*, 24(4), 2021.

[3] Mohammad Afshar and Masoud Asadpour. Opinion formation by informed agents. *Journal of Artificial Societies and Social Simulation*, 13(4):5, 2010.

[4] Ghada Alaa. Derivation of factors facilitating organizational emergence based on complex adaptive systems and social autopoiesis theories. *Emergence: Complexity and Organization*, 11(1):19, 2009.

[5] Arthur E Albert. The sequential design of experiments for infinitely many states of nature. *The Annals of Mathematical Statistics*, pages 774–799, 1961.

[6] Christopher M Anderson. Ambiguity aversion in multi-armed bandit problems. *Theory and decision*, 72(1):15–33, 2012.

[7] James Andreoni and John H Miller. Auctions with artificial adaptive agents. *Games and economic behavior*, 10(1):39–64, 1995.

[8] Jean-Yves Audibert, Rémi Munos, and Csaba Szepesvári. Exploration–exploitation tradeoff using variance estimates in multi-armed bandits. *Theoretical Computer Science*, 410(19):1876–1902, 2009.

[9] Chris Bail. *Breaking the social media prism: How to make our platforms less polarizing*. Princeton University Press, 2022.

[10] Christopher A Bail, Lisa P Argyle, Taylor W Brown, John P Bumpus, Haohan Chen, MB Fallin Hunzaker, Jaemin Lee, Marcus Mann, Friedolin Merhout, and Alexander Volfovsky. Exposure to opposing views on social media can increase political polarization. *Proceedings of the National Academy of Sciences*, 115(37):9216–9221, 2018.

[11] Venkatesh Bala and Sanjeev Goyal. Learning from neighbours. *The review of economic studies*, 65(3):595–621, 1998.

[12] Sven Banisch, Ricardo Lima, and Tanya Araújo. Agent based models and opinion dynamics as markov chains. *Social Networks*, 34(4):549–561, 2012.

[13] Matthew Barnidge, Albert C. Gunther, Jinha Kim, Yangsun Hong, Mallory Perryman, Swee Kiat Tay, and Sandra Knisely. Politically motivated selective exposure and perceived media bias. *Communication Research*, available online first, 2017.

[14] Larry M Bartels. Beyond the running tally: Partisan bias in political perceptions. *Political behavior*, 24(2):117–150, 2002.

[15] Michael Barthel. 6 key takeaways about the state of the news media in 2020. *Pew Research Center*, 2021.

[16] Adam J Berinsky. Rumors and health care reform: experiments in political misinformation. *British Journal of Political Science*, 47(2):241–262, 2017.

[17] Donald A Berry and Bert Fristedt. *Bandit problems: sequential allocation of experiments*. Springer, 1985.

[18] Kelly S Bouas and Samual S Komorita. Group discussion and cooperation in social dilemmas. *Personality and Social Psychology Bulletin*, 22(11):1144–1150, 1996.

[19] Cheryl Boudreau. Closing the gap: When do cues eliminate differences between sophisticated and unsophisticated citizens? *The Journal of Politics*, 71(3):964–976, 2009.

[20] Cheryl Boudreau and Scott A MacKenzie. Informing the electorate? how party cues and policy information affect public opinion about initiatives. *American Journal of Political Science*, 58(1):48–62, 2014.

[21] Andrei Boutyline and Robb Willer. The social structure of political echo chambers: Variation in ideological homophily in online networks. *Political Psychology*, 38(3):551–569, 2017.

[22] Russell N Bradt, SM Johnson, and Samuel Karlin. On sequential designs for maximizing the sum of n observations. *The Annals of Mathematical Statistics*, 27(4):1060–1074, 1956.

[23] Sébastien Bubeck, Rémi Munos, and Gilles Stoltz. Pure exploration in finitely-armed and continuous-armed bandits. *Theoretical Computer Science*, 412(19):1832–1852, 2011.

[24] John G Bullock. Partisan bias and the bayesian ideal in the study of public opinion. *The Journal of Politics*, 71(3):1109–1124, 2009.

[25] Daniel M Butler and Emily Schofield. Were newspapers more interested in pro-obama letters to the editor in 2008? evidence from a field experiment. *American Politics Research*, 38(2):356–371, 2010.

[26] Braz Camargo. Learning in society. *Games and Economic Behavior*, 87:381–396, 2014.

[27] Angus Campbell, Philip E Converse, Warren E Miller, and Donald E Stokes. *The American Voter*. University of Chicago Press, 1960.

[28] David E Campbell. Social networks and political participation. *Annual Review of Political Science*, 16:33–48, 2013.

[29] Taylor N Carlson. Through the grapevine: Informational consequences of interpersonal political communication. *American Political Science Review*, 113(2):325–339, 2019.

[30] Michael X Delli Carpini, Fay Lomax Cook, and Lawrence R Jacobs. Public deliberation, discursive participation, and citizen engagement: A review of the empirical literature. *Annu. Rev. Polit. Sci.*, 7:315–344, 2004.

[31] Shelly Chaiken and Yaacov Trope. *Dual-process theories in social psychology*. Guilford Press, 1999.

[32] Arthur Charpentier, Romuald Elie, and Carl Remlinger. Reinforcement learning in economics and finance. *Computational Economics*, pages 1–38, 2021.

[33] Herman Chernoff. Sequential design of experiments. *The Annals of Mathematical Statistics*, 30(3):755–770, 1959.

[34] Nicholas A. Christakis and James H. Fowler. Friendship and natural selection. *Proceedings of the National Academy of Sciences*, 111(supplement_3):10796–10801, 2014.

[35] Michael D Cobb and James H Kuklinski. Changing minds: Political arguments and political persuasion. *American Journal of Political Science*, pages 88–121, 1997.

[36] Michael Colaresi, Lara Putnam, Bree Bang-Jensen, and Minsu Jang. Revealing the growing blindspots in current computational social media telescopes and how to fill them: Using a locally-focused system to study disinformation in context. 2024.

[37] Elanor Colleoni, Alessandro Rozza, and Adam Arvidsson. Echo chamber or public sphere? predicting political orientation and measuring political homophily in twitter using big data. *Journal of communication*, 64(2):317–332, 2014.

[38] Alexander Coppock. *Persuasion in parallel: How information changes minds about politics*. University of Chicago Press, 2023.

[39] Vincent P Crawford and Joel Sobel. Strategic information transmission. *Econometrica: Journal of the Econometric Society*, pages 1431–1451, 1982.

[40] Russell J. Dalton, Paul A. Beck, and Robert Huckfeldt. Partisan cues and the media: Information flows in the 1992 presidential election. *The American Political Science Review*, 92(1):111–126, 1998.

[41] Scott De Marchi and Scott E Page. Agent-based models. *Annual Review of political science*, 17:1–20, 2014.

[42] Jeffrey S Dean, George J Gumerman, Joshua M Epstein, Robert L Axtell, Alan C Swedlund, Miles T Parker, and Stephen McCarroll. Understanding anasazi culture change through agent-based modeling. *Dynamics in human and primate societies: Agent-based modeling of social and spatial processes*, pages 179–205, 2000.

[43] Guillaume Deffuant, David Neau, Frederic Amblard, and Gérard Weisbuch. Mixing beliefs among interacting agents. *Advances in Complex Systems*, 3(01n04):87–98, 2000.

[44] Morris H DeGroot. Reaching a consensus. *Journal of the American Statistical Association*, 69(345):118–121, 1974.

[45] Michela Del Vicario, Alessandro Bessi, Fabiana Zollo, Fabio Petroni, Antonio Scala, Guido Caldarelli, H Eugene Stanley, and Walter Quattrociocchi. The spreading of misinformation online. *Proceedings of the national academy of Sciences*, 113(3):554–559, 2016.

[46] Carlos Andres Devia and Giulia Giordano. Classification-based opinion formation model embedding agents' psychological traits. *Journal of Artificial Societies and Social Simulation*, 26(3), 2023.

[47] Susanna Dilliplane. Activation, conversion, or reinforcement? the impact of partisan news exposure on vote choice. *American Journal of Political Science*, 58(1):79–94, 2014.

[48] Peter H Ditto and David F Lopez. Motivated skepticism: Use of differential decision criteria for preferred and nonpreferred conclusions. *Journal of personality and social psychology*, 63(4):568, 1992.

[49] Anthony Downs. *An Economic Theory of Democracy*. Harper Collins Publisher, New York, 1957.

[50] James N Druckman. A framework for the study of persuasion. *Annual Review of Political Science*, 25, 2021.

[51] James N. Druckman and Lawrence R. Jacobs. *Who governs? : presidents, public opinion, and manipulation*. The University of Chicago Press, Chicago, IL, 2015.

[52] James N Druckman, Matthew S Levendusky, and Audrey McLain. No need to watch: How the effects of partisan media can spread via interpersonal discussions. *American Journal of Political Science*, 62(1):99–112, 2018.

[53] James N Druckman and Arthur Lupia. Preference formation. *Annual Review of Political Science*, 3(1):1–24, 2000.

[54] James N Druckman and Mary C McGrath. The evidence for motivated reasoning in climate change preference formation. *Nature Climate Change*, 9(2):111–119, 2019.

[55] James N Druckman and Kjersten R Nelson. Framing and deliberation: How citizens' conversations limit elite influence. *American journal of political science*, 47(4):729–745, 2003.

[56] James N Druckman and Michael Parkin. The impact of media bias: How editorial slant affects voters. *The Journal of Politics*, 67(4):1030–1049, 2005.

[57] James N Druckman, Erik Peterson, and Rune Slothuus. How elite partisan polarization affects public opinion formation. *American Political Science Review*, 107(1):57–79, 2013.

[58] Alice H. Eagly, Patrick Kulesa, Laura A. Brannon, Kelly Shaw, and Sarah Hutson-Comeaux. Why counterattitudinal messages are as memorable as proattitudinal messages: The importance of active defense against attack. *Personality and Social Psychology Bulletin*, 26(11):1392–1408, 2000.

[59] Joshua M Epstein. *Generative Social Science: Studies in Agent-Based Computational Modeling*. Princeton University Press, 2006.

[60] William P Eveland Jr and Dhavan V Shah. The impact of individual and interpersonal factors on perceived news media bias. *Political psychology*, 24(1):101–117, 2003.

[61] James D. Fearon. *Deliberation as Discussion*, page 44–68. Cambridge Studies in the Theory of Democracy. Cambridge University Press, 1998.

[62] Dorian Feldman et al. Contributions to the" two-armed bandit" problem. *The Annals of Mathematical Statistics*, 33(3):847–856, 1962.

[63] Lauren Feldman, Teresa A Myers, Jay D Hmielowski, and Anthony Leiserowitz. The mutual reinforcement of media selectivity and effects: Testing the reinforcing spirals framework in the context of global warming. *Journal of Communication*, 64(4):590–611, 2014.

[64] Leon Festinger. *Conflict, decision, and dissonance.* Stanford University Press, 1964.

[65] Seth Flaxman, Sharad Goel, and Justin M Rao. Filter bubbles, echo chambers, and online news consumption. *Public opinion quarterly*, 80(S1):298–320, 2016.

[66] Carl Folke, Stephen R Carpenter, Brian Walker, Marten Scheffer, Terry Chapin, and Johan Rockström. Resilience thinking: integrating resilience, adaptability and transformability. *Ecology and society*, 15(4), 2010.

[67] Jan-Philipp Fränken and Toby Pilditch. Cascades across networks are sufficient for the formation of echo chambers: An agent-based model. *Journal of Artificial Societies and Social Simulation*, 24(3), 2021.

[68] R Kelly Garrett, Dustin Carnahan, and Emily K Lynch. A turn toward avoidance? selective exposure to online political information, 2004–2008. *Political Behavior*, 35(1):113–134, 2013.

[69] Matthew Gentzkow and Jesse M Shapiro. Media bias and reputation. *Journal of political Economy*, 114(2):280–316, 2006.

[70] Matthew Gentzkow and Jesse M Shapiro. What drives media slant? evidence from us daily newspapers. *Econometrica*, 78(1):35–71, 2010.

[71] Alan Gerber and Donald Green. Misperceptions about perceptual bias. *Annual review of political science*, 2(1):189–210, 1999.

[72] Erik J Girvan, Jason Weaver, and Mark Snyder. Elevating norm over substance: Self-monitoring as a predictor of decision criteria and decision time among independent voters. *Analyses of Social Issues and Public Policy*, 10(1):321–336, 2010.

[73] John Gittins, Kevin Glazebrook, and Richard Weber. *Multi-armed bandit allocation indices*. John Wiley & Sons, 2011.

[74] John Gittins and David M. Jones. A dynamic allocation index for the sequential design of experiments. *Progress in statistics*, pages 241–266, 1972.

[75] John C Gittins. Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society: Series B (Methodological)*, 41(2):148–164, 1979.

[76] Sanjeev Goyal. Learning in networks. In Jess Benhabib, Alberto Bisin, and Matthew O Jackson, editors, *Handbook of Social Economics*, pages 679–728. Elsevier, 2010.

[77] Tim Groeling. Media bias by the numbers: Challenges and opportunities in the empirical study of partisan news. *Annual Review of Political Science*, 16:129–151, 2013.

[78] Eric Groenendyk. *Competing motives in the partisan mind: How loyalty and responsiveness shape party identification and democracy*. Oxford University Press, 2013.

[79] Andrew Guess and Alexander Coppock. The exception, not the rule? the rarely polarizing effect of challenging information. *Working Paper*, 2016.

[80] Jurgen Habermas. *The structural transformation of the public sphere: An inquiry into a category of bourgeois society*. MIT press, 1991.

[81] Peter K Hatemi and Rose McDermott. Give me attitudes. *Annual Review of Political Science*, 19:331–350, 2016.

[82] Rainer Hegselmann, Ulrich Krause, et al. Opinion dynamics and bounded confidence models, analysis, and simulation. *Journal of artificial societies and social simulation*, 5(3), 2002.

[83] Eitan D Hersh and Brian F Schaffner. Targeted campaign appeals and the value of ambiguity. *The Journal of Politics*, 75(2):520–534, 2013.

[84] Seth J Hill. Learning together slowly: Bayesian learning about political facts. *The Journal of Politics*, 79(4):1403–1418, 2017.

[85] D Sunshine Hillygus and Todd G Shields. *The persuadable voter: Wedge issues in presidential campaigns*. Princeton University Press, 2008.

[86] Daniel E. Ho and Kevin M. Quinn. Measuring explicit political positions of media. *Quarterly Journal of Political Science*, 3(4):353–377, 2008.

[87] Lauren C Howe and Jon A Krosnick. Attitude strength. *Annual review of psychology*, 68:327–351, 2017.

[88] Robert Huckfeldt. The social communication of political expertise. *American Journal of Political Science*, pages 425–438, 2001.

[89] Shanto Iyengar and Kyu S Hahn. Red media, blue media: Evidence of ideological selectivity in media use. *Journal of communication*, 59(1):19–39, 2009.

[90] Kathleen Hall Jamieson and Joseph N Cappella. *Echo chamber: Rush Limbaugh and the conservative media establishment.* Oxford University Press, 2008.

[91] Elihu Katz and Paul F. Lazarsfeld. *Foundations of communications research.* Free Press of Glencoe, New York, N.Y, 1964.

[92] Donald R Kinder. Communication and opinion. *Annual review of political science*, 1(1):167–197, 1998.

[93] Joseph T Klapper. *The effects of mass communication.* Free Press, Glencoe, IL, 1960.

[94] Casey Klofstad. *Civic talk: Peers, politics, and the future of democracy.* Temple University Press, 2011.

[95] Casey A Klofstad. Talk leads to recruitment: How discussions about politics and current events increase civic participation. *Political Research Quarterly*, 60(2):180–191, 2007.

[96] Casey A Klofstad. Civic talk and civic participation: The moderating effect of individual predispositions. *American Politics Research*, 37(5):856–878, 2009.

[97] Casey A Klofstad. The lasting effect of civic talk on civic participation: Evidence from a panel study. *Social Forces*, 88(5):2353–2375, 2010.

[98] Silvia Knobloch-Westerwick. Selective exposure and reinforcement of attitudes and partisanship before a presidential election. *Journal of Communication*, 62(4):628–642, 2012.

[99] Silvia Knobloch-Westerwick and Jingbo Meng. Looking the other way: Selective exposure to attitude-consistent and counterattitudinal political information. *Communication Research*, 36(3):426–448, 2009.

[100] Silvia Knobloch-Westerwick, Cornelia Mothes, Benjamin K. Johnson, Axel Westerwick, and Wolfgang Donsbach. Political online information searching in germany and the united states: Confirmation bias, source credibility, and attitude impacts. *Journal of Communication*, 65(3):489–511, 2015.

[101] Ziva Kunda. The case for motivated reasoning. *Psychological bulletin*, 108(3):480, 1990.

[102] Howard G Lavine, Christopher D Johnston, and Marco R Steenbergen. *The ambivalent partisan: How critical loyalty promotes democracy*. Oxford University Press, 2012.

[103] Paul F. Lazarsfeld, Bernard Berelson, and Hazel Gaudet. *The People's Choice: How the Voter Makes Up His Mind in a Presidential Campaign*. Columbia Univ. Press, New York, 2nd edition, 1948.

[104] David Lazer. The co-evolution of individual and network. *Journal of Mathematical Sociology*, 25(1):69–108, 2001.

[105] Eun Lee, Fariba Karimi, Claudia Wagner, Hang-Hyun Jo, Markus Strohmaier, and Mirta Galesic. Homophily and minority-group size explain perception biases in social networks. *Nature human behaviour*, 3(10):1078–1087, 2019.

[106] Matthew S. Levendusky. Why do partisan media polarize viewers? *American Journal of Political Science*, 57(3):611–623, 2013.

[107] Lindsey C Levitan and Brad Verhulst. Conformity in groups: The effects of others' views on expressed attitudes and attitude change. *Political Behavior*, 38:277–315, 2016.

[108] Andrew T Little. The distortion of related beliefs. *American Journal of Political Science*, 2019.

[109] Irina Lock. Conserving complexity: A complex systems paradigm and framework to study public relations' contribution to grand challenges. *Public Relations Review*, 49(2):102310, 2023.

[110] Shanhong Luo and Eva C Klohnen. Assortative mating and marital quality in newlyweds: a couple-centered approach. *Journal of personality and social psychology*, 88(2):304, 2005.

[111] Arthur Lupia and Mathew D. McCubbins. *The democratic dilemma: Can citizens learn what they need to know?* Cambridge University Press, 1998.

[112] Jens Koed Madsen, Richard M Bailey, and Toby D Pilditch. Large networks of rational agents form persistent echo chambers. *Scientific reports*, 8(1):12391, 2018.

[113] Jens Koed Madsen and Toby D. Pilditch. A method for evaluating cognitively informed micro-targeted campaign strategies: An agent-based model proof of principle. *PLOS ONE*, 13(4):1–14, 04 2018.

[114] Shie Mannor and John N Tsitsiklis. The sample complexity of exploration in the multi-armed bandit problem. *Journal of Machine Learning Research*, 5(Jun):623–648, 2004.

[115] George E. Marcus, W. Russell Neuman, and Michael MacKuen. *Affective Intelligence and Political Judgment*. Affective Intelligence and Political Judgment. University of Chicago Press, 2000.

[116] Peter V Marsden. Core discussion networks of americans. *American sociological review*, pages 122–131, 1987.

[117] André CR Martins. Bayesian updating rules in continuous opinion dynamics models. *Journal of Statistical Mechanics: Theory and Experiment*, 2009(02):P02017, 2009.

[118] Loretta Mastroeni, Pierluigi Vellucci, and Maurizio Naldi. Agent-based models for opinion formation: A bibliographic survey. *IEEE Access*, 7:58836–58848, 2019.

[119] David R. Mayhew. *Congress : the electoral connection*. Yale studies in political science ; 26. Yale University Press, New Haven, 1974.

[120] Scott D McClurg. The electoral relevance of political talk: Examining disagreement and expertise effects in social networks on political participation. *American Journal of Political Science*, 50(3):737–754, 2006.

[121] Kathleen M McGraw and Clark Hubbard. Some of the people some of the time: Individual differences in acceptance of political accounts. *Political persuasion and attitude change*, pages 145–170, 1996.

[122] Michael F Meffert, Sungeun Chung, Amber J Joiner, Leah Waks, and Jennifer Garst. The effects of negativity and motivated information processing during a political campaign. *Journal of Communication*, 56(1):27–51, 2006.

[123] Tali Mendelberg. The deliberative citizen: Theory and evidence. *Political decision making, deliberation and participation*, 6(1):151–193, 2002.

[124] John Stuart Mill. *On Liberty*. Longmans, Green, Reader, and Dyer, 1869.

[125] William Minozzi. A jamming theory of politics. *The Journal of Politics*, 73(2):301–315, 2011.

[126] William Minozzi, Michael A Neblo, Kevin M Esterling, and David MJ Lazer. Field experiment evidence of substantive, attributional, and behavioral persuasion by members of congress in online town halls. *Proceedings of the National Academy of Sciences*, 112(13):3937–3942, 2015.

[127] William Minozzi, Hyunjin Song, David MJ Lazer, Michael A Neblo, and Katherine Ognyanova. The incidental pundit: Who talks politics with whom, and why? *American Journal of Political Science*, 64(1):135–151, 2020.

[128] Melanie Mitchell. *Complexity: A guided tour*. Oxford university press, 2009.

[129] Jonathan S Morris. Slanted objectivity? perceived media bias, cable news exposure, and political attitudes. *Social science quarterly*, 88(3):707–728, 2007.

[130] Diana C Mutz. Cross-cutting social networks: Testing democratic theory in practice. *American Political Science Review*, 96(1):111–126, 2002.

[131] Diana C Mutz. *Hearing the other side: Deliberative versus participatory democracy*. Cambridge University Press, 2006.

[132] Diana C Mutz and Jeffery J Mondak. The workplace as a context for cross-cutting political discourse. *The Journal of Politics*, 68(1):140–155, 2006.

[133] David W Nickerson and Todd Rogers. Campaigns influence election outcomes less than you think. *Science*, 369(6508):1181–1182, 2020.

[134] Cailin O'Connor and James Owen Weatherall. *The misinformation age: How false beliefs spread*. Yale University Press, 2019.

[135] Benjamin I. Page. *Who deliberates?: Mass media in modern democracy*. University of Chicago Press, 1996.

[136] Scott E Page. *Diversity and complexity*. Princeton University Press, 2010.

[137] Richard G Palmer, W Brian Arthur, John H Holland, Blake LeBaron, and Paul Tayler. Artificial economic life: a simple model of a stockmarket. *Physica D: Nonlinear Phenomena*, 75(1-3):264–274, 1994.

[138] Gordon Pennycook and David G Rand. Fighting misinformation on social media using crowdsourced judgments of news source quality. *Proceedings of the National Academy of Sciences*, 116(7):2521–2526, 2019.

[139] Toby Pilditch and Jens Koed Madsen. Targeting lt;igt;yourlt;/igt; preferences: Modelling micro-targeting for an increasingly diverse electorate. *Journal of Artificial Societies and Social Simulation*, 24(1):5, 2021.

[140] Carlo Proietti and Davide Chiarella. The role of argument strength and informational biases in polarization and bi-polarization effects. *Journal of Artificial Societies and Social Simulation*, 26(2):5, 2023.

[141] Riccardo Puglisi and James M Snyder Jr. Newspaper coverage of political scandals. *The journal of politics*, 73(3):931–950, 2011.

[142] Eric Pulick, Patrick Korth, Patrick Grim, and Jiin Jung. Modeling interaction effects in polarization: Individual media influence and the impact of town meetings. *Journal of Artificial Societies and Social Simulation*, 19(2):1, 2016.

[143] Philip J Salem. *The complexity of human communication*. Hampton Press (NJ), 2009.

[144] John L Sherry. The complexity paradigm for studying human communication: A summary and integration of two fields. *Review of Communication Research*, 3:22–54, 2015.

[145] Orowa Sikder, Robert E Smith, Pierpaolo Vivo, and Giacomo Livan. A minimalistic model of bias, polarization and misinformation in social networks. *Scientific reports*, 10(1):5493, 2020.

[146] Betsy Sinclair. *The social citizen: Peer networks and political behavior*. University of Chicago Press, 2012.

[147] Mario Luis Small. Weak ties and the core discussion network: Why people regularly discuss important matters with unimportant alters. *Social networks*, 35(3):470–483, 2013.

[148] Paul M Sniderman and Sean M Theriault. The structure of political argument and the logic of issue framing. *Studies in public opinion: Attitudes, nonattitudes, measurement error, and change*, pages 133–65, 2004.

[149] Bradly C Stadie, Sergey Levine, and Pieter Abbeel. Incentivizing exploration in reinforcement learning with deep predictive models. *arXiv preprint arXiv:1507.00814*, 2015.

[150] Michele Starnini, Mattia Frasca, and Andrea Baronchelli. Emergence of metapopulations and echo chambers in mobile agents. *Scientific reports*, 6(1):31834, 2016.

[151] Dietrich Stauffer, Adriano Sousa, and Christian Schulz. Discretized opinion dynamics of the deffuant model on scale-free networks. *Journal of Artificial Societies and Social Simulation*, 7(3), 2004.

[152] Natalie J. Stroud. Media use and political predispositions: Revisiting the concept of selective exposure. *Political Behavior*, 30(3):341–366, 2008.

[153] Natalie J. Stroud. Polarization and partisan selective exposure. *Journal of Communication*, 60(3):556–576, 2010.

[154] Natalie Jomini Stroud. *Niche news: The politics of news choice*. Oxford University Press on Demand, 2011.

[155] C.R. Sunstein. *Republic.com 2.0*. Princeton University Press, 2009.

[156] Richard S Sutton and Andrew G Barto. Toward a modern theory of adaptive networks: expectation and prediction. *Psychological review*, 88(2):135, 1981.

[157] Katarzyna Sznajd-Weron and Jozef Sznajd. Opinion evolution in closed community. *International Journal of Modern Physics C*, 11(06):1157–1165, 2000.

[158] Charles S Taber, Damon Cann, and Simona Kucsova. The motivated processing of political arguments. *Political Behavior*, 31(2):137–155, 2009.

[159] Charles S. Taber and Milton Lodge. Motivated skepticism in the evaluation of political beliefs. *American Journal of Political Science*, 50(3):755–769, 2006.

[160] Ben M Tappin, Chloe Wittenberg, Luke B Hewitt, Adam J Berinsky, and David G Rand. Quantifying the potential persuasive returns to political microtargeting. *Proceedings of the National Academy of Sciences*, 120(25):e2216261120, 2023.

[161] William R Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294, 1933.

[162] Michel Tokic and Günther Palm. Value-difference based exploration: adaptive control between epsilon-greedy and softmax. In *Annual conference on artificial intelligence*, pages 335–346. Springer, 2011.

[163] Nicholas A Valentino, Antoine J Banks, Vincent L Hutchings, and Anne K Davis. Selective exposure in the internet age: The interaction between anxiety and information utility. *Political Psychology*, 30(4):591–613, 2009.

[164] Robert P Vallone, Lee Ross, and Mark R Lepper. The hostile media phenomenon: biased perception and perceptions of media bias in coverage of the beirut massacre. *Journal of personality and social psychology*, 49(3):577, 1985.

[165] Joannes Vermorel and Mehryar Mohri. Multi-armed bandit algorithms and empirical evaluation. In *European conference on machine learning*, pages 437–448. Springer, 2005.

[166] Walter Vogel. A sequential design for the two armed bandit. *The Annals of Mathematical Statistics*, 31(2):430–443, 1960.

[167] Eva Vriens and Rense Corten. Are bridging ties really advantageous? an experimental test of their advantage in a competitive social learning context. *Social Networks*, 54:91–100, 2018.

[168] Mark D Watts, David Domke, Dhavan V Shah, and David P Fan. Elite cues and media bias in presidential campaigns: Explaining public perceptions of a liberal press. *Communication Research*, 26(2):144–175, 1999.

[169] Gérard Weisbuch, Guillaume Deffuant, Frédéric Amblard, and Jean-Pierre Nadal. Meet, discuss, and segregate! *Complexity*, 7(3):55–63, 2002.

[170] Peter Whittle. *Optimization Over Time: Dynamic Programming and Optimal Control, volume I.* John Wiley and Sons, Ltd., New York,, 1982.

[171] Jonathan Woon. Political lie detection. *Unpublished manuscript, University of Pittsburgh. Presented at the 2017 Toronto Political Behavior Workshop*, 2017.

[172] John R. Zaller. *The nature and origins of mass opinion.* Cambridge university press, 1992.

[173] Qing Zhao. Multi-armed bandits: Theory and applications to online learning in networks. *Synthesis Lectures on Communication Networks*, 12(1):1–165, 2019.