Article

Not peer-reviewed version

Mathematical Theory of Social Conformity I: Belief Dynamics, Propaganda Limits, and Learning Times in Networked Societies

Dimitri Volchenkov * and Vakhtang Putkaradze *

Posted Date: 23 April 2025

doi: 10.20944/preprints202504.1928.v1

Keywords: two-timescale theory of consensus; structural limits of propaganda efficiency; bounded vs. divergent learning times; logistic-optimal centrality and autopoietic amplification

 doi:10.20944/preprints202504.1928.v1

Article

Mathematical Theory of Social Conformity I: Belief Dynamics, Propaganda Limits, and Learning Times in Networked Societies

Dimitri Volchenkov [1,*] and Vakhtang Putkaradze [2,*]

[1] Department of Mathematics and Statistics, Texas Tech University, 1108 Memorial Circle, Lubbock, TX 79409, USA;
[2] Department of Mathematical and Statistical Sciences, University of Alberta, Edmonton AB T6G2G1 Canada
* Correspondence: dimitri.volchenkov@ttu.edu; putkarad@ualberta.ca

Abstract: This paper develops a novel probabilistic theory of belief formation in social networks, departing from classical opinion dynamics models in both interpretation and structure. Rather than treating agent states as abstract scalar opinions, we model them as belief adoption probabilities with clear decision-theoretic meaning. Our approach replaces iterative update rules with a fixed-point formulation that reflects rapid local convergence within social neighborhoods, followed by slower global diffusion. We derive a matrix logistic equation describing uncorrelated belief propagation and analyze its solutions in terms of mean learning time (MLT), enabling us to distinguish between fast local consensus and structurally delayed global agreement. In contrast to memory-driven models where convergence is slow and unbounded, uncorrelated influence produces finite, quantifiable belief shifts. Our results yield closed-form theorems on propaganda efficiency, saturation depth in hierarchical trees, and structural limits of ideological manipulation. By combining probabilistic semantics, nonlinear dynamics, and network topology, this framework provides a rigorous and expressive model for understanding belief diffusion, opinion cascades, and the temporal structure of social conformity under modern influence regimes.

Keywords: two-timescale theory of consensus; structural limits of propaganda efficiency; bounded vs. divergent learning times; logistic-optimal centrality and autopoietic amplification

1. Introduction

Information has been used as a tool of power in human society since the dawn of times. Information, as a power tool, goes hand in hand with disinformation. Plato, in the third book of his *Republic*, writes that disinformation, provided sparingly by the enlightened authoritarian ruler, can be beneficial for society much like an unpleasantly-tasting medicine can be beneficial for the patient [1], 389b:

. . . So if anyone is entitled to tell lies, the rulers of the city are. They may do so for the benefit of the city, in response to the actions either of enemies or of citizens. No one else should have anything to do with lying, and for an ordinary citizen to lie to these rulers of ours is . . . a mistake . . .

So, in plain words, Plato believed that the rulers, and only the rulers, are allowed to provide the citizens with medicinal lies. Almost 2500 years later, most people in democratic societies would certainly disagree with that assessment. It is not hard to agree that the possession of accurate information is the foundation for the well-functioning society. However, Plato's description of 'medicinal lies' was not just for ruling the country - it was also the tool for forming the citizens of the right character. Throughout the history, various rulers have tried shape the society they govern and the citizen's thinking and behavior to their will.

In order to reach that desired outcome, for almost two and a half millenia since Plato, the societies provided their citizens with messages that were designed to shape their populace in a certain way.

 doi:10.20944/preprints202504.1928.v1

Some of these messages achieved positive results and improve the society, whereas some of the social experiments leading to disastrous outcomes for the societies and their citizens. Until very recently, the message was centralized, coming from the very top of the government. The more autocratic the government was, the more uniform the centralized messages tend to be. The nature of technology delivering this message changed with the progress, from the literature, statues and even coins in the Roman empire to the radio and then television in the present day [2]. Still, the essence of the method was the same. A centralized message is presented to the citizens and that message, to large extent, is unchanged on the time scale compared to the dynamics of opinion formation in the society.

While centralized message can be effective, it has its limitations that stems from the theory of opinion dynamics. The centralized message can be viewed as a member of the community which has infinite obstinacy, whereas other members of the community can change their opinions accordingly. In the modern models of opinion dynamics utilizing the bounded confidence hypothesis [3–6] the agents interact and exchange opinions only with other agents that have the opinions that are close to their own. If the centralized message deviates from the opinion of the people it tries to influence, that centralized message will have very little to no effect.

That paradigm of centralized message has changed recently with the advent of social media. The presence of such bots on social media platforms such as X (formerly Twitter) and Instagram has been known for a long time [7–9]. While some bots are relatively harmless, many are used for propaganda purposes in order to spread the desired message among the community. The advantage of bots versus a centralized message is their efficiency: these bots align with the local opinion and can influence it much more precisely compared to the centralized message. In terms of opinion dynamics theory, such bots can act as an agent that is tuned to be close to a certain group of people and can influence the opinions of the groups that are unreachable by the centralized message [10].

The efficiency of bots can also be understood in terms of alternative models of consensus formation, such as Friedkin-Johnsen [11] and deGroot models [12]. These models do not pose restriction on the interaction of agents such as the bounded confidence models do, but the interaction between the agents is much slower if the opinion difference is large. Thus, in the agent spreading the centralized message is far from the opinion of a given group, that group is likely to take a long time to align with the centralized message. In contrast, the bot's message will be closer to the group's opinion and thus can influence the opinion of the group faster.

A closely related mechanism in the digital manipulation of public perception is the phenomenon of astroturfing, the artificial simulation of grassroots support or opposition. Astroturfing involves coordinated efforts to create a false impression of widespread public backing or resistance to particular ideas, policies, or products. By leveraging bots, fake accounts, and paid influencers, astroturfing campaigns can distort the perceived popularity of a movement, misleading individuals and policymakers alike. This practice is especially effective in the age of social media, where engagement metrics such as likes, shares, and retweets serve as heuristics for credibility and influence [13,14]. The term astroturfing originated in 1985 when Texas Senator Lloyd Bentsen, referring to a flood of letters sent to him under the guise of public advocacy but actually orchestrated by insurance industry lobbyists, remarked: "*A fellow from Texas can tell the difference between grass roots and 'AstroTurf'… this is generated mail.*" [15]. Bentsen's analogy highlighted the artificial nature of such orchestrated campaigns, comparing them to *AstroTurf*™, a synthetic surface designed to mimic real grass.

The mechanisms underpinning models of social opinion formation and the corresponding efficiency of bots can also relate to the theories of social influence, going beyond bounded confidence models. The "spiral of silence" theory proposed by Noelle-Neumann [16] suggests that individuals, fearing social isolation, are less likely to express dissenting opinions if they perceive their views to be in the minority. In an environment where bot-driven interactions artificially amplify certain viewpoints, real users may misinterpret the distribution of social preferences, leading to a cycle of self-censorship and consensus reinforcement. Similarly, Cialdini's principle of social proof [17] suggests that people rely on cues from their social environment to guide their behavior and opinions—an effect magnified

 doi:10.20944/preprints202504.1928.v1

in online spaces where engagement metrics can be artificially manipulated. Bots artificially inflate the perceived number of people having certain opinions for an individual, and thus force the individual state their alliance to certain beliefs. Whether such allegiance is sincere or not is a complex question, and is likely dependent how far the perceived majority opinion for the individual is from their true beliefs.

The structural dynamics of social influence and consensus formation can be understood through network models. For example, in bounded confidence models, the agents establish connections between them if their opinions are close enough. Clearly, the connections between these individuals are dynamic, *i.e.*, the network is changing from one time step to another. Another way to treat social networks is to have graphs that are static and are independent of the opinion dynamics; these are the interactions due to, for example, family, long-term friendships, and employment. Graph theory provides a powerful framework for analyzing the diffusion of opinions in complex systems, where individuals (nodes) influence one another through their connections (edges), forming clusters of agreement or polarization [18–20]. This general representation allows to examine how external interventions—such as bot-driven campaigns—can shift collective attitudes by altering perceived majorities. Studies on online polarization have demonstrated that artificial amplification reinforces ideological segregation, fostering echo chambers that inhibit cross-group dialogue [21,22]. Furthermore, empirical research on information cascades suggests that once a critical mass of perceived consensus is established, individuals may conform to majority opinion despite personal skepticism [23,24].

The historical use of graphs to represent social and organizational structures dates back to the 19th century. In 1855, Daniel McCallum, the General Superintendent of the New York and Erie Railroad, developed an organizational diagram (Figure 1) that visually depicted hierarchical authority and communication flows [25]. In terms of our discussion, such a network would be static with the interactions are forced by the organizational structure of their work. This early attempt to structure administrative complexity laid the groundwork for modern network analysis, illustrating how visual representations of influence and control can elucidate hidden patterns within large systems.

By analyzing the interplay of algorithmic amplification, artificial engagement, and human psychology, this paper seeks to explore the mechanisms driving opinion shifts in digital environments. Graph-based approaches allow for a systematic examination of social conformity, consensus formation, and influence propagation, providing crucial insights into how disinformation networks can manipulate public perception on a large scale. This phenomenon of disinformation-driven conformity raises important theoretical and practical questions about the relationship between perceived majority opinion and actual individual belief. Why do individuals align so readily with signals of collective support, even when those signals may be artificially engineered? What structural properties of social networks amplify or mitigate these effects? And how does the temporal structure of belief formation, fast local convergence versus slow global consensus, affect the resilience or vulnerability of populations to manipulation?

Our paper develops a comprehensive theoretical framework to explore how structural and temporal factors influence consensus formation and belief propagation in networked societies, emphasizing correlated and uncorrelated social influence. While classical opinion dynamics models typically treat the agent state $x_i(t) \in [0,1]$ as an abstract scalar "opinion" lacking clear semantics, our approach adopts a distinctly probabilistic interpretation. Here, $P_i(t)$ explicitly represents the probability of belief adoption or the agent's readiness for action, embedding a rigorous Bayesian or decision-theoretic meaning into social influence modeling. Although we describe equilibrium beliefs through a static fixed-point equation, the variable t refers to discrete or continuous cycles of external interventions (such as media influence rounds), each sufficiently spaced in time to ensure that local probabilities equilibrate almost instantaneously compared to the scale of these cycles. Thus, our use of t marks successive cycles of global information exposure rather than the short internal equilibration time within local social groups, clearly distinguishing between rapid local consensus formation and slower global belief dynamics.

 doi:10.20944/preprints202504.1928.v1

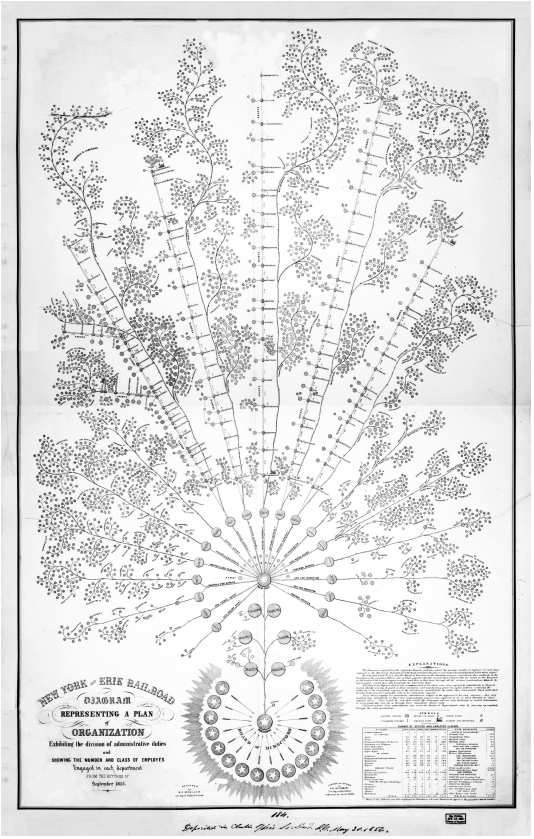

Figure 1. Organizational diagram of the New York and Erie Railroad, created by D. C. McCallum and G. H. Henshaw in 1855. Retrieved from the Library of Congress (https://www.loc.gov/item/2017586274/).

Furthermore, whereas standard frameworks describe iterative processes converging gradually to equilibrium, we formulate our model (4) as a static fixed-point equation. This explicitly captures instantaneous local consensus formation, modeling equilibrium states reached rapidly within social neighborhoods after each cycle of external influence. Consequently, our approach emphasizes the separation of two distinct timescales: fast local adaptation within immediate social circles and slower global diffusion of beliefs across the broader network, a crucial distinction often absent in standard homogeneous models.

We introduce a novel nonlinear continuous-time formulation, referred to as the *matrix logistic* differential equation, derived as the limiting case of frequent but independent interventions. The equation is termed *matrix logistic* because it describes the evolution of a *vector* of belief probabilities $\overline{P}(t) \in [0,1]^N$, governed by a fixed row-stochastic matrix $\mathbf{S}$ that encodes the structure of influence. While $\overline{P}$ can be formally represented as a $1 \times N$ or $N \times 1$ array, we use the term *vector* in the standard linear-algebraic sense, *i.e.*, in the sense of an element of the real vector space $\mathbb{R}^N$, referring simply to an ordered tuple of real-valued belief components, without implying any coordinate transformations or geometric structure as in physical contexts. This allows analytical exploration of belief trajectories, precise estimation of mean learning times (MLT), and optimal convergence trajectories, further enriched through the introduction of a logit-transformation linking our probabilistic model with geometric and information-theoretic interpretations.

Moreover, we propose a novel network centrality metric termed *logistic-optimal centrality*. Unlike traditional centrality measures such as degree or eigenvector centrality, this metric dynamically quantifies a node's temporal responsiveness, indicating how efficiently it transforms weak initial signals into sustained belief states, and explicitly depends on initial belief distributions.

Our analysis introduces several rigorous results: we derive structural thresholds limiting propaganda effectiveness in conformist groups, demonstrating inherent resilience in large heterogeneous populations and vulnerability in smaller ones. Another key finding, formalized as Theorem 2, demon-

 doi:10.20944/preprints202504.1928.v1

strates that correlated belief dynamics lack a finite temporal scale, leading to slow, subexponential convergence and divergent learning times due to cumulative social memory effects. Conversely, we prove that belief propagation under uncorrelated influence exhibits strictly bounded convergence times, enabling rapid consensus even from marginal initial beliefs, a phenomenon termed *autopoietic convergence*.

The manuscript is structured as follows: Sec. 2 revisits statistical foundations underpinning collective intelligence; Sec. 3 presents our probabilistic model of instantaneous local equilibrium formation; Sec. 4 analytically delineates structural constraints on propaganda efficiency; Sec. 5 examines temporal divergence under correlated dynamics; Sec. 6 analyzes rapid logistic shifts and bounded learning in uncorrelated scenarios. Thus, our approach represents an information-theoretic and probabilistic theory of social conformity, explicitly separating temporal scales, focusing on structural influence topology, and deriving rigorous analytical bounds on the speed and robustness of social consensus formation against systematic manipulation.

2. The Power of Majority: From Local Agreement to Global Consensus

We instinctively align with the majority, as prevailing beliefs often shape perceived truth, an example of *argumentum ad populum*. This tendency acts as both a cognitive shortcut and a social reinforcement mechanism, fostering acceptance and reducing social risk [26]. Evolution has favored conformity, as adherence to group norms enhanced survival in cooperative societies [27]. Over time, this ingrained bias has driven cultural traditions and institutional structures that sustain social cohesion.

The statistical rationale behind the superiority of collective judgment over individual estimations is rooted in aggregation theory: a group's prediction is always at least as accurate as the mean individual prediction [28], if the individual predictions have a strong stochastic component. The mathematical foundation of this principle can be derived from the analysis of squared errors. Given a random variable $x \in [0, \infty)$ representing the uncertain outcome of some process and a set of n individuals each making an independent prediction χ_i, the collective estimate is defined as the arithmetic mean $\bar{\chi} = \frac{1}{n}\sum_{i=1}^{n}\chi_i$. The squared error of the collective estimate is given by $(x-\bar{\chi})^2$, while the mean squared individual error is $\frac{1}{n}\sum_{i=1}^{n}(x-\chi_i)^2$. By applying Jensen's inequality for convex functions, it follows that

$$(x-\bar{\chi})^2 \leq \frac{1}{n}\sum_{i=1}^{n}(x-\chi_i)^2, \tag{1}$$

where the equality holds if and only if all individuals produce identical predictions. Thus, the existence of a majority implicitly assumes a diversity of perspectives, ensuring that (1) remains a strict inequality in most practical cases. This phenomenon, often referred to as the *wisdom of the crowd*, has been rigorously analyzed within the framework of error theory, where random individual errors tend to cancel out, resulting in a collective estimate that converges toward the true value [29].

This statistical advantage is further enhanced when individual predictions are uncorrelated, as demonstrated by the variance reduction principle. If individual estimates are characterized by variance σ^2 and correlation coefficient ρ between errors, the variance of the collective prediction [30] is given by

$$\sigma_{\bar{\chi}}^2 = \frac{\sigma^2}{n} + \rho\sigma^2\left(1-\frac{1}{n}\right). \tag{2}$$

As long as $\rho \approx 0$, the aggregation mechanism significantly reduces the uncertainty, explaining why heterogeneous groups with independent perspectives tend to outperform even the most knowledgeable individuals. However, if errors are highly correlated $(\rho \to 1)$, the variance reduction becomes negligible, and the collective judgment offers little to no improvement over individual estimates [31]. In cases of extreme social influence, such as echo chambers and groupthink, the wisdom of the crowd may fail, highlighting the necessity of preserving diversity in decision-making.

 doi:10.20944/preprints202504.1928.v1

The preference for majority alignment is not just statistical optimization but an entropic force driving social cohesion [32]. Evolution has reinforced conformity, as adherence to group norms enhances collective welfare and survival [33], while deviation often leads to exclusion and reduced evolutionary success. However, real-world decision-making rarely operates under ideal conditions of independence and diversity; instead, social influence shapes individual beliefs, leading to varying degrees of conformity.

Consensus formation in the models of opinion formation unfolds in two stages. First, a local consensus emerges rapidly within an individual's immediate social circle, driven by frequent interactions. Over time, these local clusters interact, leading to a slower process of global consensus formation across the broader network[4,11,12]. Yet, true unanimity remains elusive, as isolated clusters of opinion, while established far apart from each other, have very little chance to merge to a single consensus.

This inherent delay between local and global consensus creates a window of opportunity for manipulating public opinion. Temporary dominance of certain narratives within local clusters can fabricate an illusion of widespread agreement, leading individuals to conform to misleading social signals before a genuine consensus emerges. Such artificially induced perceptions reinforce the appearance of majority support even where none truly exists. Individuals that are more focused on the consumption of social media compared to other sources are more vulnerable to such manipulation, as it is relatively easy for the bot campaign to give an impression that there is a global consensus by isolating them from perceiving the true state of variability of opinions. Thus, understanding the structural and temporal dynamics of consensus formation is crucial. While collective intelligence fosters stability of the information flow, its susceptibility to distortion highlights the need for critical awareness in shaping resilient social and informational networks.

3. Fast Timescale of Local Consensus Formation in Networked Belief Dynamics

Understanding how cognitive biases interact with statistical principles in decision-making calls for a refined probabilistic model that accounts for both individual conviction and social influence. This model advances traditional conformity theories by incorporating network dynamics, resistance to persuasion, and iterative belief updating. A key insight is that consensus forms in two stages: a rapid local alignment shaped by frequent interactions, followed by a slower global convergence. The focus here is on the initial phase, where local consensus emerges within one's immediate social sphere, laying the groundwork for broader agreement.

To formalize the mechanisms underlying local consensus formation, interactions among individuals can be represented as a graph $G = (V, E)$, where nodes $V = 1, 2, \ldots, N$ correspond to individuals, and edges E denote influence relationships. The topology of this network – whether centralized, decentralized, or modular – plays a critical role in determining the rate and extent of opinion convergence. To incorporate these structural features into a formal model, we define the influence weight of an edge (i, j) as $\lambda_{ji} > 0$, representing the degree to which individual i affects individual j. The *social conformity* matrix Λ satisfies the row-stochastic constraint, $\sum_{i=1}^{N} \lambda_{ji} = 1$, ensuring that each individual's *posterior* probability of adopting a belief or behavior is derived from a weighted sum of social inputs. This constraint guarantees that belief updating follows a Markovian process, enabling iterative convergence to equilibrium.

Building on the probabilistic model of social conformity developed in [34], we assume that each individual j initially holds a *prior* probability $0 \leq P_{0j} \leq 1$ of adopting a particular course of action. Upon interacting with their social neighborhood, this prior is updated to a *posterior* probability P_j, reflecting the combined influence of personal conviction and external social pressure. To account for resistance to influence, we introduce an *obstinacy* parameter. Let $M_{jj} = \mu_j$ denote the level of obstinacy, where $0 \leq \mu_j \leq 1$, with $\mu_j = 1$ signifying complete independence and $\mu_j = 0$ total conformity. This resistance reflects empirical findings in social psychology, where cognitive complexity correlates with a lower susceptibility to social pressure [35]. Obstinacy thus acts as a stabilizing force, preventing rapid opinion shifts driven by a single interaction.

 doi:10.20944/preprints202504.1928.v1

While our model shares a formal similarity with opinion dynamics models [11,12] in that $P_i(t) \in [0,1]$ represents an evolving state, its interpretation and underlying dynamics are fundamentally different. Unlike those models, where the state is typically treated as an abstract scalar opinion, here $P_i(t)$ is modeled as the probability of adoption, with a clear Bayesian or decision-theoretic interpretation.

For fully independent individuals, the posterior remains equal to the prior, $P_j^{(\mu=1)} = P_{0j}$. For fully conforming individuals, the belief is determined entirely by the collective stance of the group, weighted by the social conformity matrix: $P_j^{(\mu=0)} = \sum_{i=1}^{N} \lambda_{ji} P_{0i}$. A general formulation incorporating both individual obstinacy and social influence is given by:

$$P_j = \mu_j P_{0j} + (1-\mu_j) \sum_{i=1}^{N} \lambda_{ji} P_{0i}. \tag{3}$$

This equation captures the weighted balance between an individual's intrinsic beliefs and the influence exerted by their social neighborhood. It naturally extends to a vector formulation, where the full network is described by $\overline{P} = (P_1, P_2, \ldots, P_N)^T$ and $\overline{P_0} = (P_{01}, P_{02}, \ldots, P_{0N})^T$. The key relationship governing posterior belief updating is:

$$\overline{P} = \mathbf{M}\overline{P_0} + (\mathbf{1} - \mathbf{M})\Lambda\overline{P} \tag{4}$$

Eq. (4) expresses the *fixed-point relation* for posterior belief distribution after integrating personal conviction with social influence. While structurally similar to iterative update models (e.g., [19]), where each step is defined by

$$\overline{P}^{(k+1)} = \mathbf{M}\overline{P}_{\text{init}} + (\mathbf{I} - \mathbf{M})\Lambda\overline{P}^{(k)}, \tag{5}$$

our formulation describes the convergence limit $\overline{P}^* = \lim_{k\to\infty} \overline{P}^{(k)}$, assuming the process reaches *equilibrium*. The presence of $\overline{P}$ on both sides in (4) is not a modeling error but reflects this equilibrium assumption.

Unlike previous models (e.g., [19]), which treat obstinacy as a fixed damping parameter relative to the initial opinion, we allow it to dynamically shape the influence aggregation by embedding it directly in the resolvent operator that governs the effective belief integration. This structure allows us to solve the system (4) explicitly:

$$\overline{P} = \mathbf{S}\overline{P}_0, \quad \mathbf{S} \equiv [\mathbf{I} - (\mathbf{I} - \mathbf{M})\Lambda]^{-1}\mathbf{M} = \left[\sum_{k=0}^{\infty} ((\mathbf{1} - \mathbf{M})\Lambda)^k\right]\mathbf{M}, \tag{6}$$

using the Neumann series expansion for the inverse. This implicit formulation enables the analysis of structural parameters (e.g., obstinacy and conformity topology) shape final belief states *without* simulating the full time evolution. While structurally similar to the standard iterative models of opinion dynamics, our model is inherently probabilistic, not diffusive. It computes the *posterior probability* of belief adoption, rather than an abstract scalar opinion. The underlying logic is rooted in a Bayesian interpretation, where each individual integrates their prior conviction with aggregated social input. The resulting linear system *does not* describe a time-stepped iteration but rather a *fixed-point constraint*, a stationary outcome of fast local learning.

In traditional opinion models, agents update their states gradually over time. In contrast, we assume that each individual's immediate social neighborhood rapidly reaches local equilibrium. From the perspective of our model, this convergence occurs on a fast timescale and is treated as instantaneous relative to the slower dynamics of belief propagation across the global network. Thus, the primary entities in our analysis are not individuals per se, but socially embedded neighborhoods that have already stabilized internally. The long-term evolution of beliefs emerges not from moment-to-moment

 doi:10.20944/preprints202504.1928.v1

updates, but from slower "media cycles" or rounds of global exposure, each of which triggers a new local equilibrium. This distinction justifies the presence of posterior beliefs $\overline{P}$ n both sides of the update equation (4) and explains why the model retains its probabilistic nature even when written in linear algebraic form.

Since each row sum of $(\mathbf{1} - \mathbf{M})\Lambda$ in (6) equals one, the sum over each row in $\mathbf{S}$ is also bounded, ensuring that it is indeed a stochastic matrix. The eigenvalues of $\mathbf{S}$ remain strictly positive for all $\mu_i > 0$, and the matrix reflects how the interplay of personal conviction and social conformity shapes the equilibrium belief state. The expansion in (6) reveals that local consensus forms rapidly through repeated interactions within close-knit groups. Early iterations in the series are dominated by direct social inputs, leading to belief synchronization at a fast timescale. In contrast, global convergence unfolds more slowly, shaped by network topology and communication bottlenecks.

This two-timescale structure reflects a fundamental distinction between fast local adaptation and slower global convergence. Rapid local agreement emerges through frequent interpersonal interactions and the reinforcement of shared beliefs. Empirical research on influence networks shows that individuals tend to quickly align their expressed views with those of their immediate social circles, often converging exponentially toward a local consensus [36,37]. Moreover, *homophily*, the tendency to associate with like-minded peers, further accelerates this process by fostering self-reinforcing opinion clusters that stabilize over short timescales [38]. These findings underscore the critical role of local network structure and interpersonal influence in shaping early belief alignment before any broader consensus takes hold.

4. Structural Limits of Propaganda in Fully Connected Conformist Groups

The probabilistic model of social influence developed in the previous section enables us to assess the structural efficiency of propaganda in tightly connected groups. In such networks, belief formation is shaped by frequent mutual interactions and exposure to persistent external messaging. We focus here on fully connected graphs with uniform influence and obstinacy levels to derive analytical insights into the conditions under which propaganda succeeds or fails.

Consider a fully connected network of N_f individuals, where each agent interacts symmetrically with all others. This setup represents a homogeneous peer group, such as a political community, workplace, or family unit, where individuals share equal exposure to mutual influence. The conformity matrix takes the form:

$$\lambda_{ji} = \frac{1 - \delta_{ji}}{N_f - 1},$$

where δ_{ji} is the Kronecker delta ensuring no individual prioritizes their own opinion. Under these conditions, the general belief-updating equation (3) simplifies to:

$$P_j = P_{0j}\mu_j + (1 - \mu_j)f, \quad f = \frac{\sum_{i=1}^{N_f} P_{0i}\mu_i}{\sum_{i=1}^{N_f} \mu_i}, \tag{7}$$

where f represents the expected fraction of individuals who conform to the dominant opinion, given their initial beliefs P_{0i} and obstinacies μ_i. This result formalizes a key mechanism of social persuasion: the final stance of an individual is a weighted average of their initial belief and the group's dominant opinion. Highly conformist individuals ($\mu_j \approx 0$) quickly align with the majority, whereas highly obstinate individuals ($\mu_j \approx 1$) maintain their initial positions. To incorporate external media influence, following [34], consider a scenario where a family-like group of N_f individuals is exposed to a dominant external source (e.g., television or online media), modeled as a perfectly obstinate $(N_f + 1)$-th agent with belief fixed at $P_0 = 1$ and obstinacy $\mu = 1$. The external source exerts influence but

 doi:10.20944/preprints202504.1928.v1

remains unaffected by the group. Adjusting the normalization to exclude this external source (as it does not update), we have the effective group-level adoption fraction after one exposure cycle:

$$f = \frac{1 + N_f P_0 \mu - \mu}{1 - \mu + N_f \mu}. \tag{8}$$

The terms $-\mu$ and $1 - \mu$ explicitly reflect the exclusion of the self-influence of the external agent.

Persuasion rarely occurs in a single step. Instead, the media landscape may evolve iteratively through repeated exposure, where dominant narratives are reinforced over multiple cycles, gradually eroding independent belief formation. In contrast to standard opinion dynamics models [11,12], where iterative averaging governs the process, our formulation treats the belief state as a probability of adoption. We therefore model the evolving adoption fraction as:

$$f^{(k+1)} = f^{(k)} + (1 - f^{(k)})(1 - \mu), \quad f^{(0)} = P_0, \tag{9}$$

which expresses that each round persuades a $(1 - \mu)$ fraction of individuals who remain unconvinced. Solving this recursion yields:

$$f^{(k)} = 1 - (1 - P_0)\, \mu^k, \tag{10}$$

indicating exponential convergence to unanimity with a convergence rate set by μ.

To determine the conditions under which a target adoption level $f_{\min}$ is reached after k rounds of media exposure, we impose $f^{(k)} = 1 - (1 - P_0)\mu^k \geq f_{\min}$, which yields an upper bound on the admissible level of obstinacy:

$$\mu \leq \left(\frac{1 - f_{\min}}{1 - P_0} \right)^{1/k}. \tag{11}$$

This condition guarantees that the iterative reinforcement process achieves the desired threshold $f_{\min}$ within k exposures.

An alternative constraint on group size N_f follows from the single-cycle equilibrium (8). Requiring that the adoption level after one round satisfies

$$\frac{1 - \mu + N_f \mu P_0}{1 - \mu + N_f \mu} \geq f_{\min}, \tag{12}$$

and rearranging terms, we obtain the inequality: $N_f \mu (f_{\min} - P_0) \leq (1 - \mu)(1 - f_{\min})$. The structural implications of this bound depend on the relationship between the target threshold $f_{\min}$ and the initial support P_0:

If $f_{\min} \leq P_0$, the group already satisfies the target level of belief prior to any exposure. The inequality is trivially satisfied for all N_f, and persuasion is guaranteed regardless of group size.

If $f_{\min} > P_0$, the campaign aims to raise adoption beyond its initial level. The inequality yields an upper bound on group size $N_f \leq (1 - f_{\min})/\mu(f_{\min} - P_0)$, which exhibits the scaling $N_f = O(1/\mu)$, highlighting structural resistance to influence in large conformist populations.

These observations can be summarized as follows:

Theorem 1 (Breakdown of Mass Propaganda in Large Conformist Groups)**.** *Let a homogeneous group of size N_f with uniform obstinacy $\mu \in (0,1)$ and initial belief level $P_0 \in (0,1)$ be exposed to persistent external influence. Suppose the campaign seeks to elevate belief adoption to a target threshold $f_{\min} \in (P_0, 1)$. Then:*

After a single exposure cycle, the goal is achievable only if

$$N_f \leq \frac{1}{\mu} \frac{1 - f_{\min}}{(f_{\min} - P_0)}. \tag{13}$$

 doi:10.20944/preprints202504.1928.v1

After k exposure rounds, a complementary constraint on obstinacy ensures success:

$$\mu \leq \left(\frac{1 - f_{\min}}{1 - P_0} \right)^{1/k}. \tag{14}$$

The analysis given above highlights a key fragility of top-down persuasion: *in sufficiently large or resistant populations, even persistent messaging fails to shift public belief.*

Such dynamics help explain why top-down ideological enforcement often fails in large, homogeneous groups and why authoritarian regimes seek to fragment communities, making centralized narratives easier to impose. Propaganda assumes it can influence broad, diverse populations, yet studies show its effectiveness diminishes in such settings due to varied opinions and social dynamics. This conclusion aligns with the findings of [39], who observed that while terrorist propaganda follows predictable patterns, its impact varies across groups, highlighting the challenge of uniform messaging in diverse populations. Similarly, studies on media exposure in politically varied contexts suggest that greater media diversity weakens propaganda. Mutz and Martin [40] found that exposure to opposing viewpoints fosters broader political knowledge, counteracting propagandistic effects. Recent research [41] further supports this, showing that even market-driven media diversity can challenge political narratives and limit propaganda's reach. These examples underscore propaganda's limitations in large social groups. While small communities may absorb uniform messaging, larger, heterogeneous populations resist centralized narratives due to varied beliefs and complex social dynamics [42].

5. Consensus Without Clocks: Absence of Temporal Scale in Correlated Belief Dynamics

Within the framework of Markovian consensus dynamics,

$$\overline{P}_k = \mathbf{S}^k \overline{P}_0, \tag{15}$$

the belief trajectory converges toward a global ideological equilibrium. If the stochastic matrix $\mathbf{S}$ is aperiodic and irreducible, the process admits a unique stationary distribution $\boldsymbol{\pi} \in \mathbb{R}^N$ satisfying $\mathbf{S}\boldsymbol{\pi} = \boldsymbol{\pi}$, to which all belief vectors $\overline{P}_k$ converge as $k \to \infty$. In the context of equation (15), the term *stochastic matrix* refers not to a randomly generated matrix (as in random matrix theory or Gaussian ensembles), but rather to a deterministic *row-stochastic* matrix: a nonnegative matrix $\mathbf{S}$ in which each row sums to 1. This structure encodes normalized influence weights in a Markovian belief update model. We do not assume any probabilistic distribution over the entries of $\mathbf{S}$; it is fixed and derived from the network's interaction topology. Although such a stochastic transition matrix guarantees eventual convergence, persistent individual obstinacy, heterogeneous influence aggregation strategies, and structural asymmetries within the interaction network may prevent full unanimity. Propaganda exploits these dynamics by injecting targeted narratives, effectively altering transition probabilities and reinforcing selected patterns of consensus.

In a fully correlated learning scenario, individuals continuously draw on prior experiences when assimilating new arguments. Belief adoption thus depends not only on immediate influence but on the entire history of social interactions. Let us consider the probability

$$P_i(t) = 1 - I_i(t), \quad I_i(t) \equiv \int_{[0,1]^N} (1 - p_i)^t \left(\sum_{j=1}^{N} S_{ij} p_j \right) \prod_{j=1}^{N} dp_j \tag{16}$$

accepts a belief precisely at time t. Here, $I_i(t)$ represents the probability that i has remained unconvinced up to time t, with belief acceptance governed by the aggregated influence of peers through the matrix $\mathbf{S}$.

This integral spans *all* possible combinations of individual acceptance probabilities p_j, thereby encoding the entire ensemble of hypothetical opinion trajectories. Initially, $I_i(0) = 1$, indicating total

 doi:10.20944/preprints202504.1928.v1

resistance and thus $P_i(0) = 0$. As time progresses, $I_i(t)$ decreases monotonically, while $P_i(t) \to 1$, reflecting the eventual adoption of belief. To evaluate $I_i(t)$ analytically, observe that the integral decomposes into a sum of N integrals, with each term of the form $S_{ij} \int_{[0,1]^N} (1-p_j)^t p_j \prod_{j=1}^N dp_j$. When $j = i$, only the integration over p_i is nontrivial, yielding

$$S_{ii} \int_0^1 (1-p_i)^t p_i \, dp_i = \frac{S_{ii}}{(t+1)(t+2)}. \tag{17}$$

For $i \neq j$, the integrals factor due to independence, so that each term becomes $S_{ij}/2(t+1)$. Summing over all $i \neq j$, we obtain:

$$\sum_{i \neq j} \frac{S_{ij}}{2(t+1)} = \frac{1 - S_{ii}}{2(t+1)}. \tag{18}$$

Combining both contributions, we find the closed-form solution:

$$P_i(t) = 1 - \frac{S_{ii}}{(t+1)(t+2)} - \frac{(1-S_{ii})}{2(t+1)} \tag{19}$$

where the diagonal entry S_{ii} quantifies the level of autonomy, the extent to which individual i adopts arguments independently of social input.

In the limiting case of complete autonomy ($S_{ii} = 1$), belief adoption accelerates:

$$P_i(t) = 1 - \frac{1}{(t+1)(t+2)} \approx 1 - O\left(\frac{1}{t^2}\right), \tag{20}$$

indicating rapid, quadratic convergence. By contrast, under complete social dependence ($S_{ii} = 0$), adoption slows:

$$P_i(t) = 1 - \frac{1}{2(t+1)} \approx 1 - O\left(\frac{1}{t}\right). \tag{21}$$

revealing a markedly slower, hyperbolic trajectory. Intermediate values $0 < S_{ii} < 1$ yield mixed dynamics: initial progress is dominated by autonomous reasoning, but long-term convergence is shaped increasingly by social influence.

Importantly, in all correlated scenarios, belief adoption ultimately reaches certainty, yet does so without a characteristic timescale. Convergence remains slow and subexponential, in contrast to the fast, bounded dynamics of uncorrelated influence flows (see Sec. 6). This phenomenon is formalized in the following result:

Theorem 2 (Lack of Temporal Scale under Correlated Learning)**.** *Let $P_i(t)$ denote the belief adoption probability under fully correlated learning with stochastic influence matrix $\mathbf{S}$. Then for any $S_{ii} \in [0,1]$,the expected time to convergence diverges:*

$$\int_0^\infty t \, P_i(t) \, dt = \infty. \tag{22}$$

Despite eventual convergence $P_i(t) \to 1$, the mean learning time diverges, indicating the absence of a characteristic temporal scale.

This slow, memory-driven convergence reveals a fundamental limitation of collective belief formation: although consensus is structurally guaranteed, its realization is temporally unbounded. The absence of a finite timescale makes belief adoption arbitrarily slow—constrained not by the lack of persuasive content, but by the inertia of accumulated social influence.

Empirical findings support this theoretical insight. Even modest social exposure can reduce opinion diversity without improving collective accuracy, thereby impeding the efficiency of group-level learning [43]. Similarly, studies of adaptive performance in Kaggle contests reveal that groups often respond more sluggishly than individuals to changing conditions, struggling to incorporate

 doi:10.20944/preprints202504.1928.v1

feedback effectively [44]. The sheer complexity of informational interactions dilutes clarity, further delaying convergence.

Crucially, while unanimity is ultimately assured regardless of autonomy levels, the *rate* of ideological alignment depends on the dominant cognitive mechanism: independent reasoning accelerates belief change, whereas cumulative social memory retards it. This tension reveals a structural trade-off: *autonomous individuals, resistant to external pressure, internalize beliefs rapidly; whereas in tightly coupled social systems, where updates incorporate full historical memory, convergence slows dramatically*.

Under such fully correlated learning regimes, the only path to rapid consensus lies in reducing social cohesion—weakening mutual influence, amplifying individual autonomy, and effectively fragmenting the network to escape the inertia of its own collective memory.

6. Uncorrelated Propaganda Scenarios - Fast Logistic Opinion Shifts

To derive the nonlinear update in the uncorrelated scenario, we model belief formation as a sequence of independent Bernoulli trials, where the probability of conversion at each step depends only on the current state and not on prior exposures. Let $P_i(t)$ denote the probability that individual i has adopted the belief by time t. Importantly, we do not define $\delta P_i(t+1)$ as a finite difference, *i.e.*, $\delta P_i(t+1) \neq P_i(t+1) - P_i(t)$. Instead, we interpret it probabilistically: $\delta P_i(t+1)$ denotes the probability that the agent converts specifically at step $t+1$, given that they have not yet converted by time t. This corresponds to the standard formulation in hazard-rate models and stochastic decision processes. Under this assumption, we write:

$$\delta P_i(t+1) = (1 - P_i(t))\hat{P}_i(t+1), \tag{23}$$

where $\hat{P}_i(t+1)$ is interpreted as the instantaneous probability of conversion given that node i is still unconverted, and it depends solely on the current social input.

Under the assumption of memoryless influence, we model this quantity by taking: $\hat{P}_i(t+1) = \sum_j S_{ij} P_j(t)$, so that the belief adoption process is driven by the current average belief level among neighbors. Substituting this expression back into (23) yields a nonlinear update rule:

$$\delta P_i(t+1) = (1 - P_i(t)) \cdot \sum_j S_{ij} P_j(t). \tag{24}$$

This expression captures the instantaneous probability gain under a memoryless influence process, where the likelihood of belief adoption at each step is proportional both to the residual capacity for change and to the aggregate social pressure exerted by the network. Since the update is derived from a probabilistic transition model, it imposes no algebraic constraint on $\mathbf{S}$ or $\overline{P}$; rather, it defines a distinct nonlinear process with logistic growth characteristics.

As the frequency of interventions increases and the timescales of social influence and media exposure begin to overlap, the model transitions naturally into a continuous-time framework. In this limit, the discrete update rule converges to a matrix-valued logistic differential equation:

$$\frac{d\overline{P}}{dt} = (1 - \overline{P}) \circ (\mathbf{S}\overline{P}), \tag{25}$$

where $\circ$ denotes elementwise (Hadamard) multiplication. Importantly, $\mathbf{S}$ is not time-dependent and does not evolve stochastically in this formulation. It is drawn once (*e.g.*, from a structural model or empirical network) and then held constant, reflecting persistent patterns of social conformity or communication. The system is therefore governed by an ordinary differential equation rather than a stochastic one, and the dynamics are fully deterministic once the initial condition is specified.

The differential formulation (25) captures the continuous adaptation of belief probabilities under persistent uncorrelated influence. This approximation is valid when the characteristic timescale of belief equilibration within social neighborhoods is much shorter than the timescale over which new

 doi:10.20944/preprints202504.1928.v1

influence is introduced. In other words, we assume that local belief states adjust rapidly to each round of influence, allowing the process to be approximated by a smooth trajectory rather than a sequence of discrete steps. In classical models, such separation holds when propaganda efforts and social signals are delivered in distinct, well-separated rounds. However, in contemporary media environments, characterized by high-frequency messaging, algorithmic curation, and real-time information exposure, this assumption is no longer strictly valid. Social input arrives continuously, and belief updates are not confined to discrete events but occur as a fluid response to overlapping streams of stimuli. Despite this, if local equilibration remains fast compared to the rate of external variation, the continuous-time model remains a good approximation of the average dynamics.

This passage to continuous time mirrors the classical transition from discrete-time Markov chains to continuous-time diffusion processes in the limit of vanishing step size. Conceptually, it reflects a shift from punctuated interventions to ongoing adaptation: public opinion is no longer shaped by isolated messages but by a cumulative informational flow, where even rapid media cycles can be integrated into a smooth dynamical framework provided that belief adaptation remains locally fast.

6.1. Bounded Learning Times in Uncorrelated Logistic Dynamics

In continuous-time opinion dynamics governed by matrix logistic propagation, the transition from uncertainty to belief can be quantified by the *Mean Learning Time* (MLT), the expected time at which an individual adopts a belief under persistent social influence.

Definition 1 (Mean Learning Time)**.** *For monotonic, sigmoidal belief trajectories, the MLT for node i is defined as the first moment of the belief activation rate:*

$$\bar{t}_i = \int_0^\infty t \, \frac{dP_i(t)}{dt} \, dt = \int_0^\infty t \, (1 - P_i(t)) \cdot \left(\mathbf{S}\overline{P}(t)\right)_i \, dt, \tag{26}$$

where $\overline{P}(t)$ evolves according to the matrix logistic equation (25).

While this integral generally requires numerical evaluation, an explicit expression is available when $P_i(t)$ closely approximates a logistic sigmoid,

$$P_i^{\log}(t) = \frac{1}{1 + \kappa_i \exp(-\rho_i t)}, \quad \kappa_i \equiv \frac{1 - P_i(0)}{P_i(0)}. \tag{27}$$

where $\rho_i > 0$ is a growth rate. In this case, the MLT becomes

$$\bar{t}_i = \int_0^\infty \frac{t \, \kappa_i \rho_i \, \exp(-\rho_i t)}{(1 + \kappa_i \, \exp(-\rho_i t))^2} dt. \tag{28}$$

To evaluate (28), substitute $u = \kappa_i e^{-\rho_i t}$, yielding

$$\bar{t}_i = \frac{1}{\rho_i} \int_0^{\kappa_i} \frac{\log(\kappa_i / u)}{(1 + u)^2} \, du = \frac{1}{\rho_i} \ln(1 + \kappa_i). \tag{29}$$

The formula implies that belief convergence under uncorrelated logistic dynamics (25) is both structurally guaranteed and temporally bounded. In sharp contrast to correlated influence, which may yield diverging learning times, the uncorrelated case exhibits robust convergence for any $P_i(0) > 0$ and $\rho_i > 0$.

Although the effective instantaneous influence rate $\rho_i(t) = \left(\mathbf{S}\overline{P}(t)\right)_i$ is naturally bounded, $0 < \rho_i(t) \leq 1$, as $\overline{P}(t) \in [0,1]^N$ and $\mathbf{S}$ is row-stochastic, the dynamic character of $\rho_i(t)$ implies that the value of the integral (29) may either exceed or fall below the reference level $\ln(1 + \kappa_i)$, which is determined

 doi:10.20944/preprints202504.1928.v1

by the initial conditions. The only exception is the case of full obstinacy, $\mu = 1$, in which the MLT is dominated entirely by autonomous dynamics, and the inequality becomes sharp:

$$\bar{t}_i^{\text{obst}} = \ln\left(1 + \frac{1 - P_i(0)}{P_i(0)}\right). \tag{30}$$

In the regime of vanishing initial support $P_i(0) \ll 1$, the upper bound (30) simplifies asymptotically: $\bar{t}_i^{\text{obst}} \sim \ln(1/P_i(0))$, as $P_i(0) \to 0$, suggesting that belief formation time scales uniquely with the information-theoretic surprise of the initial belief.

In summary, the boundedness of $\bar{t}_i$ reflects a fundamental temporal regularity: *persistent, uncorrelated social influence drives consensus at a rate that is inversely proportional to effective connectivity and logarithmically sensitive to prior uncertainty.*

6.2. Belief Saturation and Temporal Centrality in Hierarchical Networks with a Teacher

We investigate the structural delays inherent in belief propagation along hierarchical chains, where a marginal node serves as a fully obstinate and initially convinced teacher ($\mu_1 = 1$, $P_1(0) = 1$), while all other nodes share a uniform low initial support ($P(0) = 0.01$) and common obstinacy parameter μ. Influence is defined by an isotropic bidirectional random walk, yielding a symmetric, localized averaging structure.

Figure 2 presents representative belief trajectories under two regimes: Panel (**a**) illustrates the weakly autonomous case $\mu = 0.05$, dominated by neighbor influence, while Panel (**b**)corresponds to a balanced regime $\mu = 0.4$, where internal conviction begins to contribute appreciably. In both settings, the teacher remains fixed at full belief $P_1(t) = 1$, while learners evolve gradually under logistic propagation.

Notably, the propagation dynamics differ qualitatively. In the weakly autonomous regime, early learners near the teacher (*e.g.*, #2 and #3) are, somewhat counterintuitively, delayed in their belief acquisition, whereas deeper bulk nodes activate collectively and earlier than the closest followers, giving rise to a non-monotonic adoption profile. Conversely, in the balanced regime, belief spreads coherently as a sigmoidal front: early learners convert first, followed sequentially by deeper nodes, with the order of adoption aligning with topological depth.

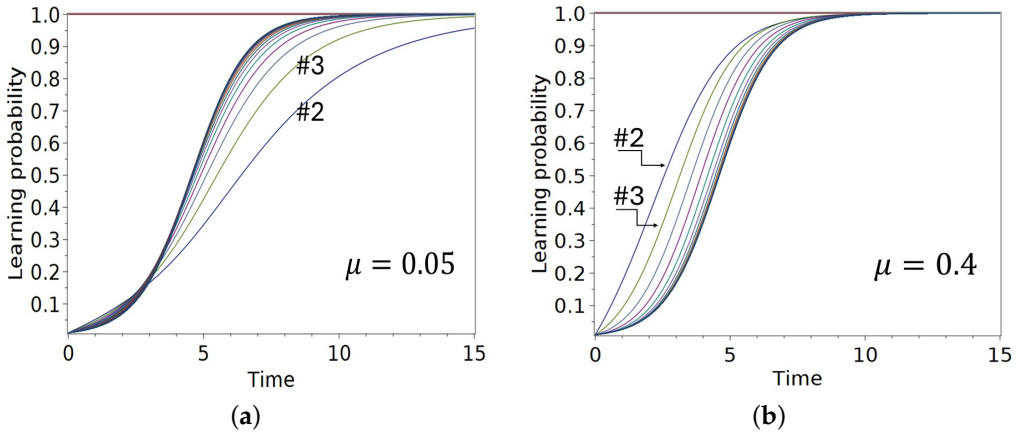

Figure 2. Time evolution of belief adoption probabilities $P_k(t)$, $k > 1$, in a linear chain of learners, with a fully obstinate teacher at node #1 ($P_1(t) = 1$), uniform initial beliefs $P(0) = 0.01$ and fixed influence structure based on an isotropic random walk. In both panels, despite differing dynamics, most nodes in the chain converge within a narrow time window, reflecting the emergence of bulk synchronization. (**a**) At low obstinacy $\mu = 0.05$, deeper nodes activate faster than early learners (*e.g.*, #2 and #3). (**b**) At moderate obstinacy $\mu = 0.4$, belief propagates as a coherent front, with delays increasing with depth.

These distinctions manifest clearly in the mean learning time (MLT) profiles. Figure 3 plots $\bar{t}_k$ as a function of depth k for several values of μ, revealing how structural synchronization emerges with

 doi:10.20944/preprints202504.1928.v1

decreasing obstinacy. At $\mu = 0.05$, MLTs decrease with depth, reflecting early suppression and delayed ignition followed by coherent collective activation. At $\mu = 0.4$, MLTs increase with depth, as early learners are privileged by their proximity to the source. A critical transition occurs around $\mu = 0.11$, where MLTs flatten across the chain, signaling temporal synchronization.

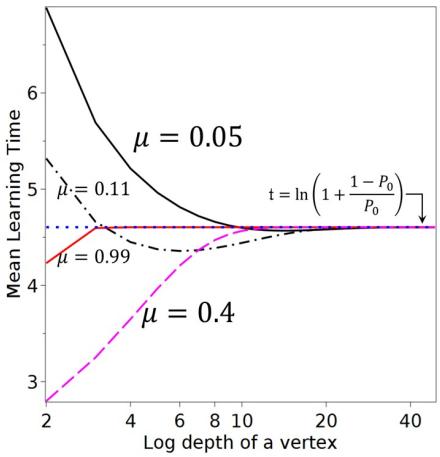

Figure 3. Mean learning time $\bar{t}_k$ versus vertex depth k in a chain with a fully obstinate teacher ($P_1(t) = 1$) and uniform initial beliefs $P(0) = 0.01$. Each curve corresponds to a different obstinacy parameter μ.

In the autonomous limit $\mu \to 1$, all nodes evolve independently as identical sigmoidal curves, saturating the autonomous MLT bound $\bar{t}^{\text{obst}} = \ln(1 + \frac{1-P_0}{P_0})$, marked by the horizontal dotted line. Similar bottleneck and synchronization patterns were observed in tree-structured simulations where the teacher occupied the root node, affirming the generality of the phenomena across hierarchical topologies.

6.2.1. Coherent Learning in the Limit of Vanishing Individual Autonomy

We now turn to the regime of vanishing individual autonomy, in which the obstinacy parameter $\mu \to 0$ and the belief dynamics are entirely governed by mutual influence. In this regime, the system displays emergent coherence for $k \gg 1$, whereby all deep nodes behave synchronously and evolve as if they were a single logistic unit.

Let the network be a one-dimensional chain of N nodes, with influence governed by a normalized isotropic bidirectional random walk:

$$\dot{P}_k(t) = (1 - P_k(t)) \cdot \frac{P_{k-1}(t) + P_{k+1}(t)}{2}, \quad k = 2, \ldots, N-1, \tag{31}$$

with boundary conditions: $P_1(t) \equiv 1$ for a fully obstinate teacher fixed at full belief, while all other nodes (learners) start from the same low initial belief $P_k(0) = \varepsilon \ll 1$, for $k > 1$, and evolve purely through peer averaging.

Guided by the numerical results from Figure 2.**a**), we assume that for sufficiently large k, all trajectories approximate a common profile $P(t)$. We define the deviation from this profile as $\delta_k(t) = P_k(t) - P(t)$, and aim to estimate its decay as a function of depth k. We consider the limiting trajectory $P(t)$ as the solution to the standard logistic equation: $\dot{P}(t) = P(t)(1 - P(t))$, $P(0) = \varepsilon$, with solution

$$P(t) = \frac{1}{1 + Ce^{-t}}, \quad C = \frac{1-\varepsilon}{\varepsilon}. \tag{32}$$

Substituting $P_k(t) = P(t) + \delta_k(t)$ into (31) and linearizing in δ_k, we obtain a linear non-autonomous system:

$$\dot{\delta}_k(t) \approx -P(t) \cdot \delta_k(t) + \frac{1 - P(t)}{2}(\delta_{k-1}(t) + \delta_{k+1}(t)). \tag{33}$$

 doi:10.20944/preprints202504.1928.v1

To analyze the behavior of deviations for large k, we note that as $t \to \infty$, the logistic solution satisfies $P(t) \to 1$ and $1 - P(t) \sim Ce^{-t}$ with exponential accuracy. Substituting this into (33), we obtain the asymptotic form:

$$\dot{\delta_k}(t) \approx -\delta_k(t) + \frac{C}{2} e^{-t} (\delta_{k-1}(t) + \delta_{k+1}(t)). \tag{34}$$

This suggests a depth-dependent solution of the form $\delta_k(t) \sim \eta_k e^{-kt}$, where η_k is independent of time. Neglecting the forward term $\delta_{k+1}(t)$ as subdominant for large k, we obtain a recurrence: $-(k-1)\eta_k \approx C\eta_{k-1}/2$, or $\eta_k \approx C^{k-1}\eta_0/2^{k-1}(k-1)!$. Thus, the deviation decays as

$$\delta_k(t) \approx \frac{C^{k-1}}{2^{k-1}(k-1)!} \eta_0 e^{-kt}, \tag{35}$$

as $t \to \infty$, for all $k \geq 2$. This result confirms exponential coherence in depth under purely diffusive dynamics with a fixed boundary.

Theorem 3 (Coherent Learning in the Limit $\mu \to 0$). *Let $P_k(t)$ evolve under matrix logistic dynamics with $\mu = 0$ on a finite chain, with $P_1(t) \equiv 1$ and all $P_k(0) = \varepsilon \ll 1$ for $k \geq 2$. Then, for all nodes $k \geq k_0 \gg 1$, the deviation $\delta_k(t) = P_k(t) - P(t)$, where $P(t)$ is the logistic solution (32), satisfies the estimate:*

$$|\delta_k(t)| \lesssim \frac{C^{k-1}}{2^{k-1}(k-1)!} e^{-kt}, \tag{36}$$

exhibiting exponential decay in depth and time.

Corollary 1 (Asymptotic Learning Time in the Collective Limit). *In the limit $\mu \to 0$, for all $k \geq k_0 \gg 1$, the mean learning times converge to:*

$$\bar{t}_k \to \int_0^\infty t\, \dot{P}(t)\, dt = \log\left(\frac{1}{\varepsilon}\right). \tag{37}$$

That is, the entire chain behaves as a *single logistic unit* with initial support ε,converging synchronously to full belief in finite expected time.

Empirical studies of innovation diffusion in structured communities suggest that direct followers of a pioneering agent (such as a teacher or opinion leader) may exhibit slower adoption than more distant actors. In particular, [45] shows that highly connected individuals, due to reputational constraints, often delay adoption until reinforced by multiple peers, while downstream actors adopt rapidly once the innovation gains visibility. Similarly, classic diffusion studies by [46] and summarized by [47] observe that early adopters initiate the awareness phase, but the bulk of the population accelerates adoption after the innovation has been socially legitimized. This phenomenon supports our observation that, in low-autonomy regimes, deep learners within a network may convert earlier than those directly adjacent to the source.

6.2.2. Belief Saturation and Bottlenecks in the Regime $\mu \lesssim 1$

In the regime where autonomous dynamics are present but not overwhelming ($\mu \lesssim 1$), belief propagation along the chain exhibits a sharply different pattern. Each node approximately follows an independent logistic curve delayed in time, $P_k(t) \approx [1 + \kappa e^{-\rho(t-\Delta_k)}]^{-1}$, where the delays Δ_k accumulate recursively based on the integration of upstream trajectories. For small initial belief $P(0) \ll 1$, these delays grow sublinearly with depth and eventually saturate beyond a structural threshold.

 doi:10.20944/preprints202504.1928.v1

Remark 1 (Asymptotic Estimate for the Saturation Depth for $\mu \lesssim 1$)**.** *Let belief propagate over a chain under matrix logistic dynamics with fixed obstinacy $\mu \lesssim 1$ and small initial belief $P(0) = \varepsilon \ll 1$. Then the mean learning time $\bar{t}(k)$ at depth k saturates beyond a critical depth approximately given by*

$$k^* \sim \frac{1}{\mu} \ln \ln \left(\frac{1}{\varepsilon} \right). \tag{38}$$

The above estimate is asymptotic and captures the heuristic behavior of saturation observed in simulations (see Figure 3) and supported by the following asymptotic analytical arguments.

First, the early growth of belief follows $P(t) \sim P_0 e^{\mu t}$, yielding the MLT $\bar{t}_0 \sim \mu^{-1} \ln(1/\varepsilon)$.

Second, delays introduced by successive levels decay exponentially, $\Delta t(k) \sim \alpha^k$ for $\alpha < 1$, as belief is inherited with damping due to $\mu < 1$. Aggregating these delays gives a convergent geometric sum $\sum_{s=1}^{k} \alpha^s$, which stabilizes once $\alpha^k \lesssim 1/\ln(1/\varepsilon)$. Inverting yields the estimate (38). While not a rigorous bound, this expression captures the emergence of a structural bottleneck observed in simulations (Figure 3).

The saturation depth k^* defines the minimal number of hierarchical layers required to overcome cognitive inertia and initiate autopoietic propagation. Above this threshold, belief spreads rapidly and near-synchronously. Below it, propagation is significantly slowed by upper-layer resistance. This phenomenon is analogous to nucleation theory, where growth only becomes self-sustaining once a critical cluster size is reached [48].

Empirical evidence supports this structural insight. In hierarchical organizations, early adopters, such as managers or ideological elites, disproportionately influence downstream adoption [49]. Strategic misinformation campaigns exploit these dynamics, prioritizing structural access over volume [50]. Effective belief propagation hinges not on intensity alone, but on timing and topological positioning. Structural depth and temporal centrality jointly govern the diffusion capacity of a network. Understanding their interplay enables more accurate modeling of persuasion dynamics in hierarchical systems.

The MLT offers a natural foundation for defining a dynamic measure of *temporal centrality*, defined by the MLT $\bar{t}_i$: nodes with shorter $\bar{t}_i$ exert greater influence in initiating diffusion cascades.

Definition 2 (Temporal Centrality)**.** *In matrix logistic belief dynamics, the temporal centrality score of node i is given by:*

$$\nu_i \equiv \frac{1}{\bar{t}_i}. \tag{39}$$

Nodes with higher ν_i are temporally more central: they are capable of faster belief adoption and play a disproportionate role in initiating propagation cascades.

Unlike structural centrality measures such as degree or eigenvector centrality, ν_i reflects the network's dynamical geometry and resistance structure, making it a context-sensitive indicator of influence.

6.3. Autopoietic Amplification of Marginal Beliefs

An important insight emerging from matrix logistic dynamics is that even weak initial signals can be amplified through sustained endogenous feedback, leading to rapid convergence without substantial external pressure. This phenomenon, which we refer to as *autopoietic convergence*, captures the capacity of a networked system to self-organize and propagate beliefs, transforming marginal prior support into systemic consensus.

Consider the matrix logistic equation (25) with uniform obstinacy $\mu \in (0, 1)$ and an influence structure encoded by a row-stochastic matrix Λ. The effective influence matrix is then given by $\mathbf{S} = \mu[\mathbf{I} - (1-\mu)\Lambda]^{-1}$, which is well-defined for any irreducible and aperiodic (i.e., ergodic in the Markov sense) matrix Λ. In this context, ergodicity refers to the property that the Markov chain defined

 doi:10.20944/preprints202504.1928.v1

by Λ has a unique stationary distribution and converges to it from any initial state. This ensures that the matrix inverse exists and the influence dynamics are globally well-posed. As shown in earlier sections, the belief update rate satisfies $\rho_i(t) = \left(\mathbf{S}\overline{P}(t)\right)_i \in [0,1]$, and in particular, its lower bound is determined by the spectral radius of $\mathbf{S}$. For many structured graphs, the smallest nonzero eigenvalue of $\mathbf{S}$ is bounded below by $\mu/(2-\mu)$, guaranteeing that belief dynamics never become arbitrarily slow.

To estimate the mean learning time for small initial support $P_i(0) = \epsilon \ll 1$, we approximate the belief trajectory as a logistic sigmoid, so that the MLT is given by

$$\bar{t}_i \approx \frac{1}{\rho_i} \log\left(1 + \frac{1-\epsilon}{\epsilon}\right) \sim \frac{1}{\rho_i} \log\left(\frac{1}{\epsilon}\right), \tag{40}$$

as $\epsilon \to 0$. Using the lower bound $\rho_i \geq \mu/(2-\mu)$ then yields a universal estimate:

$$\bar{t}_i \leq \frac{2-\mu}{\mu} \log\left(\frac{1}{\epsilon}\right). \tag{41}$$

This result formalizes a key structural guarantee: *even marginal beliefs will be amplified and adopted within finite time, provided the network is ergodic and agents retain minimal responsiveness*. The convergence time grows only logarithmically in the inverse prior ϵ, reflecting the information-theoretic surprise of the initial condition.

The mechanism becomes particularly transparent in the case of a complete graph of size N, where the conformity matrix is $\Lambda = (\mathbf{1} - \mathrm{diag}(\mathbf{1}))/(N-1)$, so that each node is equally influenced by all others. The corresponding matrix $\mathbf{S}$ has leading eigenvalue $\sigma_1 = 1$ and remaining spectrum concentrated near μ as $N \to \infty$. In this limit, belief trajectories simplify to $P_i(t) \approx \left(1 + \epsilon^{\mu t}\right)^{-1}$, and the MLT obeys the refined bound:

$$\bar{t}_i \leq \frac{1}{\mu} \log\left(\frac{1}{\epsilon}\right). \tag{42}$$

These analytic results reinforce empirical observations from social psychology. Effects such as group polarization [51,52] and pluralistic ignorance [53] illustrate how even unpopular beliefs can rapidly crystallize into dominant positions under symmetric, decentralized influence structures. The matrix logistic model provides a generative explanation: belief adoption is driven not by initial volume, but by feedback-mediated amplification, a hallmark of *autopoietic convergence*.

6.4. Statistical Isolation of Influence: Techniques for Memory Suppression in Propaganda

Ensuring that each round of propaganda remains statistically independent requires the disruption of memory effects and the prevention of ideological accumulation over time. This necessitates the use of statistical de-correlation techniques that fragment public memory, suppress long-term resistance, and maintain continuous susceptibility to influence. By strategically applying these mechanisms, regimes can ensure that each intervention stands alone, preventing individuals from forming coherent counter-narratives or resisting future campaigns effectively.

One fundamental approach is stochastic resetting, whereby public opinion is periodically reinitialized, either through narrative reversals, purges of ideological figures, or sudden shifts in official policy. This ensures that prior ideological developments do not persist across multiple rounds of influence. Imperial Rome implemented this technique under the name *damnatio memoriae*, where all images of inconvenient historical figures were purged from records [54]. A more familiar recent example is given by Stalinist Russia,, erasing purged officials from records and even photographs, thereby preventing the consolidation of alternative political loyalties [55]. In contemporary contexts, similar methods manifest in sudden shifts in government rhetoric, where public figures previously promoted as authoritative sources are swiftly discredited when their stance no longer aligns with the evolving state narrative.

Closely related is the injection of noise, a technique that floods the information space with contradictory, overwhelming, or misleading narratives to hinder stable belief formation. By generating

 doi:10.20944/preprints202504.1928.v1

a high-frequency, high-volume stream of conflicting messages, regimes can create cognitive fatigue, where the audience disengages from critical analysis and instead follows the dominant message of the moment. The "firehose of falsehood" model exemplifies this strategy, leveraging state-controlled media to generate rapid, multi-channel, repetitive, and often contradictory propaganda [56]. This technique, which prioritizes volume over consistency, ensures that each cycle of influence is self-contained, rendering past narratives obsolete and reducing resistance to new messaging.

To further disrupt the formation of ideological continuity, regimes employ randomization of social influence structures. This is achieved through continuous reshuffling of public discourse influencers, such as political leaders, journalists, or digital personalities, preventing long-term trust from forming between the public and any particular information source. Frequent rotations in government figures and controlled opposition movements serve this purpose, ensuring that loyalty does not accumulate within any specific faction. In the modern era, algorithmic censorship on social media platforms contributes to this fragmentation by selectively boosting or suppressing different narratives based on shifting political priorities [57].

To increase the adoption of the message, target segmentation may be used, where smaller groups receive messages that are fine-tuned to that particular groups. As we mentioned before, bots, exploiting the algorithmic power of modern social media, may become devastatingly effective in delivering their messaging to smaller targeted groups.

7. Discussion and Conclusion

The presented research explores the dynamics of belief formation and opinion consensus in social networks under diverse conditions of social conformity, obstinacy, structural topology, and influence propagation mechanisms. Our analytical and numerical findings provide several crucial insights into how individuals integrate social information, form collective judgments, and converge toward consensus, highlighting both the potential and limits of majority-driven decision-making.

We began our analysis by revisiting the classical result known as the *wisdom of crowds*, demonstrating rigorously how the mean squared error of collective predictions consistently outperforms individual forecasts provided there is diversity and statistical independence among predictions. This advantage arises fundamentally from the cancellation of random individual errors, a phenomenon amplified when individual judgments are minimally correlated. The implications are clear: societal decision-making processes inherently benefit from diversity and decentralization, emphasizing the societal advantage in maintaining heterogeneous sources of information and opinion. Conversely, we found that high correlation among individual errors severely diminishes the accuracy advantage, reflecting situations typical of "echo chambers," groupthink, or propaganda-saturated environments. Thus, preserving independent thinking and diverse informational sources is not merely beneficial but essential for accurate collective judgment.

Our model further decomposed consensus formation into two distinct phases – rapid local consensus and slow global convergence, clarifying empirical observations that immediate social circles rapidly synchronize beliefs, creating stable local clusters that are resistant to external influence. Such localized equilibrium, confirmed through analytical computations, explains why opinions initially solidify quickly within close-knit groups. However, global consensus emerges much more slowly, constrained by communication bottlenecks and structural divisions within society. This fundamental temporal duality offers critical insights into why social influence strategies often first target local clusters, as rapid internal agreement within smaller communities can then propagate more broadly.

Our detailed study of fully connected conformist networks elucidates the structural limits of propaganda. We derived exact formulas that quantify the effectiveness of external influence in uniform social groups, showing that the impact of persuasive efforts is inversely proportional to the size and obstinacy of the targeted population. Specifically, small, conformist groups rapidly succumb to external messaging, while larger, more resistant populations remain resilient. Our recursive model of repeated exposure further highlights how sustained propaganda efforts can erode initial beliefs, eventually

 doi:10.20944/preprints202504.1928.v1

achieving unanimity if the targeted community remains small and sufficiently conformist. Conversely, we established a critical threshold in group size, beyond which propaganda rapidly loses effectiveness due to the prevalence of internal social reinforcement mechanisms.

Examining correlated belief dynamics, we found an unexpected phenomenon: while consensus is guaranteed structurally, the time required to achieve it diverges. When individuals' belief adoption depends on cumulative social memory, their opinions shift at a progressively diminishing rate, ultimately resulting in a convergence that is unbounded in temporal scale. Our rigorous proof reveals a deep structural constraint of collective cognition: strong dependence on historical interactions creates inertia that slows belief adaptation. Empirically, this underscores the social dangers of collective memory saturation, where excessive reliance on past information effectively paralyzes decision-making processes.

In contrast, our analysis of uncorrelated propaganda scenarios demonstrated that belief adoption under memoryless influence rapidly converges to certainty, governed by logistic dynamics that impose strict upper bounds on mean learning times. We provided explicit, closed-form expressions for the fastest possible convergence trajectories, characterizing conditions under which belief shifts are maximally accelerated. This result implies that societies subject to continuous, memoryless streams of external influence, typical of high-frequency social media environments, can experience rapid ideological shifts. Such conditions amplify marginal initial beliefs into systemic consensus through what we termed *autopoietic convergence*, highlighting the powerful potential of continuous external signals to shape societal beliefs even when initial support is minimal.

These findings collectively provide profound insights into the mechanisms of opinion formation and manipulation within societies. They demonstrate how structural parameters, such as group size, connectivity, obstinacy, and influence topology, critically shape societal responsiveness to external narratives. The intrinsic vulnerability of small, homogeneous communities to external influence highlights strategic vulnerabilities that are frequently exploited in political propaganda, marketing, and social engineering contexts. Meanwhile, larger, diverse societies inherently resist top-down ideological control, emphasizing the societal benefits of heterogeneity and robust internal communication structures.

Moreover, the demonstrated fragility of collective intelligence under correlated opinion dynamics cautions against overreliance on consensus-driven decision-making in environments prone to memory saturation and information overload. The identified critical role of independent thinking and diversity in maintaining collective accuracy underscores the importance of institutional safeguards that protect informational plurality and encourage critical engagement.

In conclusion, our study reveals fundamental insights into the structural and dynamic underpinnings of consensus formation and social influence. By rigorously deriving explicit conditions for optimal consensus, exploring the contrasting dynamics of correlated versus uncorrelated influence scenarios, and identifying robust structural limits to external persuasion, we deepen understanding of how societal belief systems evolve and stabilize. Our analysis provides essential tools for diagnosing vulnerabilities in social communication networks and designing resilient informational structures.

Future research should extend these insights empirically, testing model predictions in real-world social networks and systematically exploring how varying degrees of correlation, structural diversity, and influence timing shape belief dynamics across different contexts. Additionally, exploring the ethical implications and policy responses to our findings can further strengthen societal resistance to undue influence and enhance the robustness of collective decision-making processes.

Author Contributions: Conceptualization, D.V., V.P. ; methodology, D.V., V.P. ; software, D.V.; formal analysis, D.V., V.P.; investigation, D.V., V.P.; resources, D.V., V.P.; writing—original draft preparation, D.V.; writing—review and editing, D.V., V.P.; visualization, D.V.; project administration, D.V., V.P. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

 doi:10.20944/preprints202504.1928.v1

Data Availability Statement: Not applicable.

Acknowledgments: The authors are grateful to their institution for the administrative and technical support. The authors also gratefully acknowledge Dr. Ori Swed (https://www.oriswed.com/) for his valuable insights and stimulating discussions, which significantly contributed to the conceptual development of this work.

Conflicts of Interest: The authors declare no conflict of interest.

Abbreviations

The following abbreviations are used in this manuscript:

MLT Mean Learning Time

References

1. Plato (**2000**) "The Republic". Cambridge University Press 2000 (3rd Ed).
2. Taylor, P. M. (**2013**) "*Munitions of the mind: A history of propaganda from the ancient world to the present era*." Manchester University Press.
3. Deffuant, G., Neau, D., Amblard, F., and Weisbuch, G. (**2000**). Mixing beliefs among interacting agents. *Advances in Complex Systems*, 3(01n04), 87-98.
4. Hegselmann, R., and Krause, U. (**2002**). Opinion dynamics and bounded confidence: models, analysis and simulation. *Journal of Artifical Societies and Social Simulation* (JASSS) 5(3).
5. Hegselmann, R. and Krause, U. (**2005**). Opinion dynamics driven by various ways of averaging. *Computational Economics*, 25, 381–405
6. Hegselmann, R. (**2023**). Bounded confidence revisited: What we overlooked, underestimated, and got wrong. *Journal of Artificial Societies and Social Simulation*, 26(4) 11 http://jasss.soc.surrey.ac.uk/26/4/11.html.
7. Mazza, M., Avvenuti, M., Cresci, S. and Tesconi, M. (**2022**). Investigating the difference between trolls, social bots, and humans on Twitter. *Computer Communications*, 196, 23-36.
8. Liu, X. (**2019**). A big data approach to examining social bots on Twitter. *Journal of Services Marketing*, 33(4), 369-379.
9. Tunç, U ., Atalar, E., Gargı, M. S., and Aydın, Z. E. (**2022**). Classification of fake, bot, and real accounts on instagram using machine learning. *Politeknik Dergisi*, 27(2), 479-488.
10. Swed, O., Dassanayaka, S. & Volchenkov, D. Keeping it authentic: the social footprint of the trolls' network. *Soc. Netw. Anal. Min.* 14, 38 (**2024**). https://doi.org/10.1007/s13278-023-01161-1.
11. Friedkin, N. E., and Johnsen, E. C. (**1990**). Social influence and opinions. *Journal of mathematical sociology*, 15(3-4), 193-206.
12. DeGroot, M. H. (**1974**). "Reaching a Consensus." *Journal of the American Statistical Association*, 69(345), 118-121.
13. Zuckerman, E. (**2017**). "Misinformation and manipulation: The rise of astroturfing in digital media." *Journal of Information Technology & Politics*, 14(2), 95-108.
14. Howard, P. N. (**2018**). *Lie Machines: How Fake News and Social Media Manipulate the World*. Yale University Press.
15. Bentsen, L. (**1985**). *U.S. Senate Proceedings on Lobbying and Grassroots Movements*. Congressional Records.
16. Noelle-Neumann, E. (**1974**). "The Spiral of Silence: A Theory of Public Opinion." *Journal of Communication*, 24(2), 43-51.
17. Cialdini, R. B. (**2001**). *Influence: Science and Practice*. Allyn & Bacon.
18. Watts, D. J., & Strogatz, S. H. (**1998**). "Collective dynamics of 'small-world' networks." *Nature*, 393(6684), 440-442.
19. Anderson, B. D., & Ye, M. (**2019**). Recent advances in the modelling and analysis of opinion dynamics on influence networks. International Journal of Automation and Computing, 16(2), 129-149.
20. Centola, D. (**2010**). "The Spread of Behavior in an Online Social Network Experiment." *Science*, 329(5996), 1194-1197.
21. Conover, M. D., Ratkiewicz, J., Francisco, M., Gonçalves, B., Menczer, F., & Flammini, A. (**2011**). "Political polarization on Twitter." *Proceedings of the International AAAI Conference on Web and Social Media*, 5(1), 89-96.
22. Del Vicario, M., Bessi, A., Zollo, F., Petroni, F., Scala, A., Caldarelli, G., Stanley, H. E., & Quattrociocchi, W. (**2016**). "The spreading of misinformation online." *Proceedings of the National Academy of Sciences*, 113(3), 554-559.
23. Sunstein, C. R. (**2017**). *#Republic: Divided Democracy in the Age of Social Media*. Princeton University Press.

 doi:10.20944/preprints202504.1928.v1

24. Barabási, A.-L., & Albert, R. (**1999**). "Emergence of Scaling in Random Networks." *Science*, 286(5439), 509-512.
25. McCallum, D. C., Henshaw, G. H. & New York And Erie Railroad Company, P. (**1855**) New York and Erie Railroad diagram representing a plan of organization: exhibiting the division of academic duties and showing the number and class of employés engaged in each department: from the returns of September. [New York: New York and Erie Railroad Company] [Map] Retrieved from the Library of Congress, https://www.loc.gov/item/2017586274/.
26. Bauman, Z. (**2003**). *Liquid Modernity*. Polity Press.
27. Boyd, R., & Richerson, P. J. **(2005**). *The Origin and Evolution of Cultures*. Oxford University Press.
28. Page, S. E. (**2007**). *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*. Princeton University Press.
29. Surowiecki, J. (**2004**). *The Wisdom of Crowds*. Anchor Books.
30. Krause, J., Ruxton, G. D., & Krause, S. (**2010**). "Swarm Intelligence in Animals and Humans". *Trends in Ecology & Evolution*, 25(1), 28-34.
31. Lorenz, J., Rauhut, H., Schweitzer, F., & Helbing, D. (**2011**). "How social influence can undermine the wisdom of crowd effect". *Proceedings of the National Academy of Sciences*, 108(22), 9020-9025. https://doi.org/10.1073/pnas.1008636108
32. Volchenkov, D. (**2016**). *Survival under Uncertainty: An Introduction to Probability Theory and Its Applications*. In Springer series *Understanding Complex Systems*.
33. Tversky, A., & Kahneman, D. (**1974**). "Judgment Under Uncertainty: Heuristics and Biases." *Science*, 185(4157), 1124-1131.
34. Krasnoshchekov, P.S. (**1998**). *A simplest mathematical model of [collective] behavior. Psychology of conformism. Math. Model*. 10(7), 76–92 (in Russian).
35. Sunstein, C. R. (**2019**). *Conformity: The Power of Social Influences*. New York University Press.
36. Moussaïd, M., Brighton, H., & Gaissmaier, W. (**2013**). The amplification of risk in experimental diffusion chains. *Proceedings of the National Academy of Sciences*, *110*(24), 9354-9359. https://doi.org/10.1073/pnas.1303006110
37. Momennejad, I. (**2019**). Learning structures of social influence. *Nature Human Behaviour*, *3*, 668-678. https://doi.org/10.1038/s41562-019-0621-8
38. McPherson, M., Smith-Lovin, L., & Cook, J. M. (**2001**). Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, *27*, 415-444. https://doi.org/10.1146/annurev.soc.27.1.415
39. Hahn, L., Schibler, K., Lattimer, T.A., Toh, Z., Vuich, A., Velho, R., Kryston, K., O'Leary, J., and Chen, S., "Why We Fight: Investigating the Moral Appeals in Terrorist Propaganda, Their Predictors, and Their Association with Attack Severity," *Journal of Communication*, vol. 74, no. 1, pp. 63–76, Feb. **2024**. https://doi.org/10.1093/joc/jqad029
40. Mutz, D. C., and Martin, P. S. (**2001**). Facilitating communication across lines of political difference: the role of mass media. *Am. Polit. Sci. Rev.*, 95, 97–114.
41. Steinfeld, N., & Lev-On, A. "Exposure to diverse political views in contemporary media environments," *Frontiers in Communication*, vol. 9, Art. 1384706, **2024**. https://doi.org/10.3389/fcomm.2024.1384706
42. Gallagher, P.: Revealed: Putin's army of pro-Kremlin bloggers. *The Independent* March 27, **2015**. [Online]. Available: https://www.independent.co.uk/news/world/europe/revealed-putin-s-army-of-prokremlin-bloggers-10138893.html [Accessed: Feb. 26, 2025]
43. Jan Lorenz, Heiko Rauhut, Frank Schweitzer, & Dirk Helbing (**2011**). "How social influence can undermine the wisdom of crowd effect." *PNAS* 108(22): 9020–9025.
44. Tianyu He, Marco S. Minervini, & Phanish Puranam (**2024**). "How Groups Differ from Individuals in Learning from Experience." *Organization Science* 35(2): 502–521.
45. Centola, D. (**2018**). *How Behavior Spreads: The Science of Complex Contagions*. Princeton University Press.
46. Coleman, J. S., Katz, E., & Menzel, H. (**1957**). The diffusion of an innovation among physicians. *Sociometry*, 20(4), 253–270.
47. Rogers, E. M. (**2003**). *Diffusion of Innovations* (5th ed.). Free Press.
48. K. F. Kelton and A. L. Greer, *Nucleation in Condensed Matter: Applications in Materials and Biology*, Pergamon Materials Series, Vol. 15, Elsevier, **2010**.
49. D. Centola and M. Macy, "Complex contagions and the weakness of long ties," *American Journal of Sociology*, vol. 113, no. 3, pp. 702–734, **2007**.
50. S. Vosoughi, D. Roy, and S. Aral, "The spread of true and false news online," *Science*, vol. 359, no. 6380, pp. 1146–1151, **2018**.

 doi:10.20944/preprints202504.1928.v1

51. Sunstein, C. R., "The law of group polarization". *Journal of Political Philosophy*, 10(2):175–195, **2002**.
52. Daniel J. Isenberg. Group polarization: A critical review and meta-analysis. *Journal of Personality and Social Psychology*, 50(6):1141–1151, **1986**.
53. Prentice, D. A., Miller, D. T., Pluralistic ignorance and alcohol use on campus: Some consequences of misperceiving the social norm. *Journal of Personality and Social Psychology*, 64(2):243–256, **1993**.
54. Varner, E. R. Monumenta Graeca et Romana: Mutilation and transformation: damnatio memoriae and Roman imperial portraiture. 1040 10 Brill, **2004**.
55. Blakemore, E., "How Photos Became a Weapon in Stalin's Great Purge," *History*, A&E Television Networks, 20 April **2018** (updated 11 April 2022), available at https://www.history.com/news/josef-stalin-great-purge-photo-retouching, accessed on 3/12/2025.
56. Paul, Ch., Matthews, M., "The Russian '*Firehose of Falsehood*' Propaganda Model: Why It Might Work and Options to Counter It," RAND Corporation Perspective PE-198, **2016**, available at https://www.rand.org/pubs/perspectives/PE198.html, accessed on 3/12/2025.
57. Tiffert, G., "30 Years After Tiananmen: Memory in the Era of Xi Jinping," *Journal of Democracy* 30, no. 2 (**2019**): 38–49.