Simulating Pairwise Communication to Study Opinion Dynamics in Networked Communities

Grygoriy Zholtkevych[1,2,†], Yurii Lytvynenko[2,*,†]

[1]Ivan Franko National University of Lviv, 1 Universyteska Str., Lviv, 79007, Ukraine
[2]V.N. Karazin Kharkiv National University, 4 Svobody sq., Kharkiv, 61022, Ukraine

Abstract

This paper introduces an extended model of opinion dynamics based on pairwise dialogues within networked communities. Building upon an existing binary dialogue model, we propose a more realistic framework that allows agents to adopt, reject, or choose alternative opinions when consensus is not reached. By treating dialogues as the fundamental units of social interaction, our probabilistic approach incorporates resistance and persuasiveness as independent agent attributes influencing opinion evolution. We present a rigorous formalization and simulation analysis that demonstrates how varying these attributes leads to diverse outcomes, such as opinion deadlocks, rapid convergence toward alternative opinions, and gradual shifts that reflect real-world opinion formation.

Beyond its methodological contributions, the extended dialogue model addresses challenges central to contemporary psychological operations (psyops), where information is weaponized to manipulate narratives, polarize societies, and destabilize decision-making. The ongoing war in Ukraine illustrates these dynamics vividly and motivates the need for models that make the mechanics of influence observable at scale. Our framework provides a computational lens for analyzing, detecting, and anticipating manipulation strategies in networked communities, complementing empirical and policy work on countering disinformation.

Keywords

Opinion Dynamics, Communication Simulation, Networked Communities, Stochastic Model, Persuasion, Resistance, Multi-Agent Simulation, Computational Social Science, Psyops, Information Warfare,

1. Introduction

Psychological operations (psyops) have become a defining feature of modern hybrid warfare, where information is weaponized alongside kinetic force to shape perceptions, erode trust, and steer collective behavior. The ongoing war in Ukraine illustrates these dynamics vividly: coordinated disinformation campaigns, troll networks, and propaganda outlets target domestic and international audiences to fracture consensus and destabilize decision-making [1, 2]. Social media platforms such as X (formerly Twitter) and Facebook provide fertile ground for such operations, enabling adversarial actors to exploit social influence mechanisms and algorithmic amplification at scale [3, 4].

In this environment, there is an urgent need for computational models that make the mechanics of influence observable: how opinions evolve under strategic manipulation, which parameter regimes lead to deadlock or rapid shifts, and when asymmetric attributes (e.g., high persuadability vs. high resistance) create path-dependent outcomes. Opinion dynamics offers the right formal lens. Classical models such as DeGroot's averaging process explain consensus formation under ideal conditions [5]; the Friedkin-Johnsen (FJ) model introduces resistance (stubbornness), allowing persistent disagreement [6]; and subsequent work shows how stubborn agents shape polarization and misinformation propagation [7]. Probabilistic and multi-topic extensions — e.g., voter-like updates and varying susceptibility [8, 9, 10], and interdependent-topic or antagonistic ties [11, 12] — increase realism, but many still collapse dialogue outcomes to binary adoption or rejection.

SMICS'25: Workshop on Cryptology and Data Security, October 16–18, 2025, Lviv, Ukraine
*Corresponding author.
†These authors contributed equally.
grygoriy.zholtkevych@lnu.edu.ua (G. Zholtkevych); inbox@yury-lytvynenko.com (Y. Lytvynenko)
https://yury-lytvynenko.com (Y. Lytvynenko)
0000-0002-7515-2143 (G. Zholtkevych); 0009-0004-0732-0602 (Y. Lytvynenko)

Such binarization is limiting in psyops contexts, where targets may reject both dominant and counter-narratives and pivot to alternative or adversarial viewpoints. It also obscures how persuasion and resistance act as independent drivers of change — an asymmetry routinely exploited in coordinated influence operations in and around Ukraine [1]. To analyze these mechanisms, models must explicitly separate persuasion from resistance and admit outcomes beyond simple adoption.

This paper answers that need. We extend a dialogue-based framework by (i) modeling persuasion and resistance as independent agent attributes; (ii) introducing an explicit *alternative-opinion* outcome when consensus is not reached; and (iii) defining a probabilistic transition process over pairwise dialogues that yields rich behaviors (deadlock, rapid convergence to alternatives, gradual shifts). Through simulations, we show how parameter regimes reproduce phenomena salient to psyops and information warfare—rapid narrative pivots, echo-chamber stability, and asymmetric influence effects—thus offering a computational lens for analyzing, detecting, and anticipating manipulation strategies in networked communities [2, 4].

We next recall the baseline dialogue model that our extension builds upon and then formalize the proposed transition dynamics.

The model proposed in [13] represents opinion evolution as a sequence of dialogue between network agents. Each dialogue allows an agent to either:

1. Retain their current opinion (resistance to persuasion).
2. Adopt the opinion of their interlocutor (persuasion).

Opinion exchange occurs in a networked setting, where individuals interact according to predefined communication probabilities. Over time, these interactions shape the distribution of collective opinions within the community.

While this framework captures basic opinion exchange mechanisms, it assumes that an individual must either accept or reject an opinion of their interlocutor without any third option or consideration of an independent evaluation. As a result, the original model has two key limitations that restrict its applicability to real-world decision-making:

- *No opinion change in high-resistance cases:* When both agents are highly resistant to persuasion, dialogues become ineffective — no opinion shifts occur, leading to a static system even after many interactions.
- *Lack of an alternative option:* The model assumes that agents must either keep their opinion or fully adopt their interlocutor's opinion. However, individuals often reject both options and seek an independent alternative.

To address these limitations, we introduce an extended model that allows agents to select an alternative opinion by:

- Incorporating a mechanism of choosing an alternative, allowing agents to reject both their current opinion and their interlocutor's.
- Introducing *persuasion* and *resistance* as independent attributes, enabling a more nuanced understanding of how opinions evolve.
- Defining a probabilistic transition model that determines the likelihood of different opinion shifts based on these factors.
- Simulating the extended model to analyze its behavior under different conditions and compare it to the original framework.

These enhancements specifically address critical gaps in the initial model [13], which only permitted binary opinion outcomes (accept or reject) without explicitly modeling alternative options or independently considering resistance and persuasiveness. By explicitly incorporating these dimensions, our extended model enables more realistic and varied dynamics, aligning closely with real-world opinion formation processes.

1.1. Relevance to Social Science

Persuasion and resistance have long been fundamental concepts in social psychology, highlighted in classical frameworks such as the Social Judgment Theory proposed by Sherif and Hovland [14]. According to this theory, opinion formation and change are influenced by existing attitudes, the perceived credibility of communicators, and the latitude of acceptance or rejection toward incoming messages. Additionally, resistance to persuasion — explored through cognitive dissonance theory — explains how and why individuals maintain their attitudes even when faced with persuasive counterarguments [15].

Furthermore, real-world opinion dynamics often involve exploring alternatives rather than binary acceptance or rejection, a behavior increasingly significant in contemporary sociopolitical environments. Research by Mutz [16] and Sunstein [17] confirms that individuals frequently navigate beyond polarized positions, selecting or constructing alternative perspectives to mitigate cognitive dissonance and social pressure.

Thus, extending computational models with these social-scientific dimensions facilitates more realistic simulations and deeper insights into collective decision-making processes, enriching both the theoretical understanding and practical interventions in social networks and digital communication systems.

2. Model for Pairwise Communication Simulation

To address the limitations of the original model, we introduce an extended opinion dynamics framework that allows agents to not only retain or adopt an interlocutor's opinion but also select an alternative option when consensus is not reached. This modification makes the model more flexible and realistic, reflecting real-world decision-making processes.

This section formalizes the dialogue process, defines the state space, and introduces a probabilistic transition function that governs opinion evolution.

2.1. Assumptions

We begin by introducing assumptions to formalize opinion exchange:

Assumption 1. *If two agents hold the same opinion, they will retain it after a dialogue. It reflects stability in consensus – if both individuals agree, external influence is impossible.*

Assumption 2. *If two agents hold different opinions, three outcomes are possible:*

1. *An agent retains their opinion.*
2. *An agent adopts the interlocutor's opinion.*
3. *An agent chooses an alternative opinion.*

It expands the original model by allowing agents to reject both their own and the interlocutor's opinions, introducing a third option.

Assumption 3. *The exact alternative chosen is not important.*

- *In real-world discussions, when people reject both available options, they often consider multiple alternatives without an immediate preference for one.*
- *Instead of modeling specific alternatives, we aggregate all alternatives into one effective option.*
- *This simplification reduces computational complexity while preserving the model's core dynamics.*

Assumption 4. *Initial preferences are independent.*

- *The initial state distribution of opinions is assumed to be statistically independent between agents.*
- *This simplifies analytical computations while allowing the model to focus on how interactions influence opinion shifts.*

2.2. Base Definitions

Let V denote the set of possible opinions, where $|V| > 2$. The special case when $|V| = 2$ was already covered in detail by the original model introduced in [13]. In this paper, we extend the model to handle more general scenarios where agents can choose among more than two possible opinions, reflecting richer and more realistic dialogue dynamics.

A **preference density function** $w(v)$ of an agent over a set of possible opinions V is a probability distribution defined on V, assigning each opinion $v \in V$ a probability $w(v)$ representing how strongly the agent initially favors that opinion. Formally, it satisfies:

$$\sum_{v \in V} w(v) = 1, \quad w(v) \geq 0, \ \forall v \in V$$

For two interacting agents, Alice (A) and Bob (B), their individual preferences are denoted as:

$$w_A(v), w_B(v) \quad \forall v \in V$$

The **state of the dialogue** between two agents, Alice and Bob, is formally represented by an ordered pair of opinions:

$$S = (v_A, v_B), \quad v_A, v_B \in V$$

where v_A denotes Alice's opinion and v_B denotes Bob's opinion immediately before their interaction. Consequently, the complete dialogue state space is the Cartesian product: $S \in V \times V$

Given Assumption 4, which asserts that the initial opinions of agents are independent, we define the joint preference density of the dialogue state as the product of the individual preference densities:

$$w_{AB}(v_A, v_B) = w_A(v_A) \cdot w_B(v_B)$$

Here, $w_A(v_A)$ and $w_B(v_B)$ represent the probabilities that Alice and Bob independently hold opinions v_A and v_B, respectively, before their interaction.

$w_{AB}(v_A, v_B)$ is also a preference density function defined on S that satisfies conditions

$$\sum_{(v_A, v_B) \in S} w_{AB}(v_A, v_B) = 1, \quad w_{AB}(v_A, v_B) \geq 0, \ \forall (v_A, v_B) \in S$$

2.3. Reducing Alternative Options

Since agents can choose an alternative when they disagree, we formally define the space of other options.

If $v_A \neq v_B$ then

- for $|V| = 3$ the alternative is uniquely given by the single element $v_{alt(A,B)}$ in $V \setminus \{v_A, v_B\}$;
- for $|V| > 3$, agents, however, have more than one alternative $v \in alt(A, B) = V \setminus \{v_A, v_B\}$ to choose from. To maintain computational simplicity while preserving model robustness, under the Assumption 3, we assume:
  - The alternative is selected uniformly at random from the set $V \setminus \{v_A, v_B\}$.
  - It means the effective alternative is represented by a single aggregated option, v_{alt}.

  Importantly, this aggregated alternative should not be interpreted as a singular or uniform stance adopted equally by all agents. Instead, it reflects a generalized category of alternatives, potentially unique to each individual, signifying dissatisfaction with currently available options and openness toward exploring other viewpoints. In practical terms, when agents transition into this residual category, they individually enter exploratory or uncertain states, each potentially distinct from others' interpretations and motivations. This modeling approach aligns with common probabilistic and decision-theoretic methodologies, where varied but infrequently occurring alternatives are grouped into a general "other" category for analytical convenience. The model, therefore, captures the essential dynamics of opinion shifts without the intractable complexity of enumerating all possible alternatives explicitly. Future research could further disaggregate this generalized alternative category to explore detailed, actor-specific opinion dynamics, enriching the model's descriptive power and practical applicability.

Thus, the opinion set is reduced to:

$$V = \{v_A, v_B, v_{alt}\}, \quad S = V \times V, \quad |S| = 9.$$

It allows the transition function to be tractable and interpretable while capturing the impact of alternative selection.

2.4. Transition Function and State Evolution

Given a state (v_A, v_B), the transition function $f : S \to \mathcal{P}(S)$ defines the set of possible opinion shifts after a dialogue:

$$f(v_A, v_B) = \begin{cases} \{(v, v)\} & \text{if } v_A = v_B = v \\ S & \text{if } v_A \neq v_B \end{cases} \tag{1}$$

It means that:

- If $v_A = v_B$, the state is absorbing (no change occurs).
- If $v_A \neq v_B$, the dialogue may lead to any of the nine possible states where agents can retain, adopt, or switch to an alternative.

Each possible outcome can be represented as:

Outcome	**Interpretation**
1. $((v_A, v_B)) \to (v_A, v_B)$	Both keep their opinions
2. $((v_A, v_B)) \to (v_A, \boldsymbol{v_A})$	Bob switches to Alice's opinion
3. $((v_A, v_B)) \to (\boldsymbol{v_B}, v_B)$	Alice switches to Bob's opinion
4. $((v_A, v_B)) \to (\boldsymbol{v_B}, \boldsymbol{v_A})$	Both switch (swap opinions)
5. $((v_A, v_B)) \to (v_A, \boldsymbol{v_{alt}})$	Alice keeps, Bob chooses alternative
6. $((v_A, v_B)) \to (\boldsymbol{v_{alt}}, v_B)$	Alice chooses alternative, Bob keeps
7. $((v_A, v_B)) \to (\boldsymbol{v_{alt}}, \boldsymbol{v_{alt}})$	Both choose alternative
8. $((v_A, v_B)) \to (\boldsymbol{v_B}, \boldsymbol{v_{alt}})$	Alice switches to Bob's, Bob chooses alternative
9. $((v_A, v_B)) \to (\boldsymbol{v_{alt}}, \boldsymbol{v_A})$	Alice chooses alternative, Bob switches to Alice's

This formalism provides a probabilistic representation of opinion shifts in a network.

2.5. Defining Probabilities

To model how agents update their opinions, we introduce two key attributes:

- **Resistance** $\rho \in [0, 1]$: The likelihood that an agent will retain their opinion rather than adopt another.
- **Persuasiveness** $\pi \in [0, 1]$: The likelihood that an agent will convince their interlocutor to adopt their opinion.

In sociopsychological terms, "resistance" aligns closely with constructs such as attitude strength (the firmness or certainty of one's beliefs) and ego involvement (the extent to which an opinion is tied to personal identity). Individuals with high resistance may be less receptive to contrary arguments, reflecting firm conviction or attachment to their existing viewpoint. On the other hand, "persuasiveness" resonates with frameworks like communicator credibility in social psychology, which highlights how a speaker's perceived expertise and trustworthiness shape their influence over others. Agents endowed with high persuasiveness may shift others' opinions more readily, mirroring how credible or charismatic communicators often succeed in swaying audiences.

For Alice (A) and Bob (B), these parameters are denoted as:

$$\rho_A, \pi_A, \quad \rho_B, \pi_B$$

Each agent has three possible actions after a dialogue:

1. Retain their opinion.
2. Adopt the interlocutor's opinion.
3. Choose an alternative opinion.

The probabilities of these actions are determined by resistance and persuasiveness. For Alice, they are defined as:

$$p_1^A = \frac{\rho_A(1-\pi_B)}{Z_A} \quad \text{(Probability of retaining opinion)}$$

$$p_2^A = \frac{(1-\rho_A)\pi_B}{Z_A} \quad \text{(Probability of adopting Bob's opinion)}$$

$$p_3^A = \frac{\rho_A \pi_B}{Z_A} \quad \text{(Probability of choosing an alternative)}$$

where Z_A is a normalization constant ensuring that the probabilities sum to 1:

$$Z_A = \rho_A(1-\pi_B) + (1-\rho_A)\pi_B + \rho_A\pi_B$$

Similarly, for Bob:

$$p_1^B = \frac{\rho_B(1-\pi_A)}{Z_B}, \quad p_2^B = \frac{(1-\rho_B)\pi_A}{Z_B}, \quad p_3^B = \frac{\rho_B\pi_A}{Z_B}$$

with: $Z_B = \rho_B(1-\pi_A) + (1-\rho_B)\pi_A + \rho_B\pi_A$

Interpretation

- If resistance is high ($\rho_A \approx 1$), Alice is more likely to retain her opinion.
- If Bob's persuasiveness is high ($\pi_B \approx 1$), Alice is likelier to adopt Bob's opinion.
- If both resistance and persuasiveness are high, Alice may reject both opinions and select an alternative.

These probabilities define how agents transition between states in the dialogue process.

2.6. Joint Transition Probabilities and Transition Matrix

Since each agent acts independently, the probability of transitioning from state (v_A, v_B) to a new state (v'_A, v'_B) is the product of individual transition probabilities:

$$T((v_A, v_B), (v'_A, v'_B)) = P(v'_A|v_A) \cdot P(v'_B|v_B)$$

where:

$$P(v'_A|v_A) = p_1^A[v'_A = v_A] + p_2^A[v'_A = v_B] + p_3^A[v'_A = v_{alt}]$$
$$P(v'_B|v_B) = p_1^B[v'_B = v_B] + p_2^B[v'_B = v_A] + p_3^B[v'_B = v_{alt}]$$

These probabilities define how the system evolves, forming a Markov process. For each $v_A \neq v_B$, the nine possible outcomes and their associated probabilities are:

Outcome	**Interpretation**		$T\big((v_A, v_B), (\cdot, \cdot)\big)$
1. (v_A, v_B)	Both keep their opinions	:	$p_A^1\, p_B^1$
2. (v_A, v_A)	Bob switches to Alice's opinion	:	$p_A^1\, p_B^2$
3. (v_B, v_B)	Alice switches to Bob's opinion	:	$p_A^2\, p_B^1$
4. (v_B, v_A)	Both switch (swap opinions)	:	$p_A^2\, p_B^2$
5. (v_A, v_{alt})	Alice keeps, Bob chooses alternative	:	$p_A^1\, p_B^3$
6. (v_{alt}, v_B)	Alice chooses alternative, Bob keeps	:	$p_A^3\, p_B^1$
7. (v_{alt}, v_{alt})	Both choose alternative	:	$p_A^3\, p_B^3$
8. (v_B, v_{alt})	Alice chooses Bob's, Bob – alternative	:	$p_A^2\, p_B^3$
9. (v_{alt}, v_A)	Alice chooses alternative, Bob – Alice's	:	$p_A^3\, p_B^2$

2.6.1. Matrix Representation

For convenience, we enumerate the state space S, reducing the joint preference density w_{AB} into a row vector, and express transitions as a 9×9 matrix $T : S \times S \to [0, 1]$:

$$T_{i,j} = T(s_i, s_j)$$

represents the probability of moving from state s_i to state s_j. This matrix satisfies the following key properties:

1. Non-negativity: $T(s_i, s_j) \geq 0, \quad \forall s_i, s_j \in S$
2. Valid Transition Probabilities: $T(s_i, s_j) = 0, \quad \forall s_j \notin f(s_i)$ If the model dynamics do not allow a transition, its probability is zero.
3. Normalization (Stochastic Matrix Property): $\sum_{s_j \in f(s_i)} T(s_i, s_j) = 1, \quad \forall s_i$ Each row of T sums to 1, ensuring that the agent must transition to some valid state for any given state.

3. Simulation Scenarios and Model Dynamics

To illustrate the behavior of the extended model, we conduct simulations where two agents, Alice and Bob, engage in a series of dialogues. These simulations examine how different values of resistance (ρ) and persuasiveness (π) influence opinion evolution. We aim to analyze whether agents reach consensus, polarization, or alternatives under different conditions.

We begin by demonstrating the limitations of the original model, followed by the impact of introducing persuasiveness and alternative options. Finally, we discuss practical applications of these dynamics in real-world decision-making and computational modeling.

The following simulation results and illustrations were generated using publicly accessible software, available online for review and replication at [18].

- Fig. 1 illustrates the limitation of the original dialogue model [14], specifically demonstrating opinion deadlock due to maximum resistance. Alice and Bob each hold different initial opinions and have maximum resistance ($\rho = 1$) with zero persuasiveness. The figure shows how, despite repeated dialogues, neither agent changes their opinion. This visually underscores a key limitation that our extended model seeks to address.
- Fig. 2 demonstrates the immediate convergence to the alternative opinion (v_{alt}) when both Alice and Bob exhibit maximum resistance ($\rho = 1$) but also maximum persuasiveness ($\pi = 1$). Here, the figure emphasizes how high persuasiveness combined with strong resistance leads to rapid exploration of alternative opinions, highlighting a core innovation of our model.
- Fig. 3 captures an asymmetric scenario where Alice is highly persuasive and resistant, while Bob is fully resistant but not persuasive. It illustrates how Alice's persuasiveness influences Bob's transition towards alternative opinions without herself adopting Bob's viewpoint. This figure demonstrates how asymmetric attributes affect opinion dynamics uniquely.

- Fig. 4 illustrates the effect of partial persuasiveness, showing Bob's limited persuasiveness beginning to influence Alice's openness toward alternative options. This captures realistic interactions where subtle changes in persuasiveness can incrementally shift agents' opinions, highlighting the model's nuanced behavior.
- Fig. 5, realistic opinion evolution where Alice and Bob have moderate levels of both resistance and persuasiveness ($\rho = \pi = 0.8$). Unlike previous extreme scenarios, this figure depicts gradual shifts over multiple dialogues, representing incremental opinion convergence toward mixed states, closely reflecting realistic social influence dynamics.

3.1. Original Model Limitation: No Opinion Change Under High Resistance

As discussed, the original model in [13] does not allow opinion change when both agents have maximum resistance to persuasion ($\rho_A = \rho_B = 1$), which leads to a deadlock, where neither agent alters their opinion, regardless of the number of dialogues.

In Fig.1, Alice and Bob initially have a strong preference for their own options and high resistance, ignoring the alternative completely. After repeated dialogues, her opinion remains unchanged, as resistance prevents any shifts.

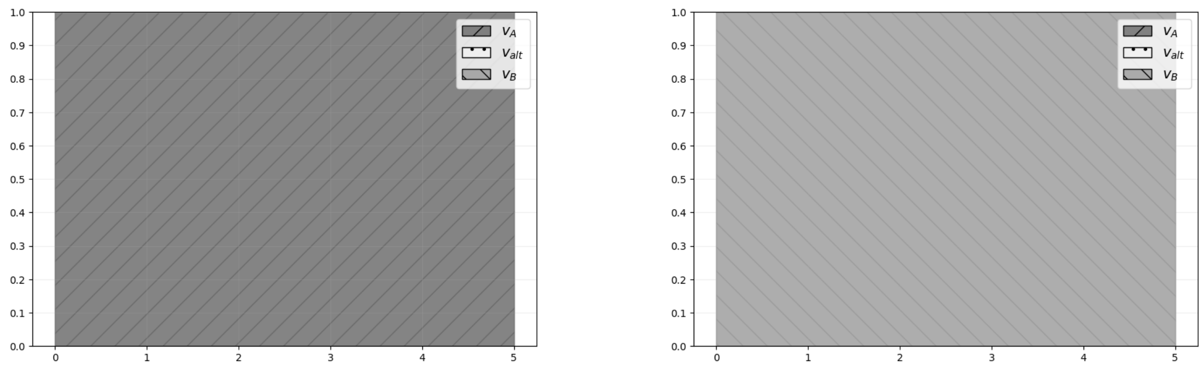

(a) Alice

(b) Bob

Figure 1: High Resistance in the Original Model

3.2. Introducing Persuasion: Moving Beyond Deadlock

The extended model introduces persuasiveness (π), representing an agent's ability to influence their interlocutor. It allows persuasion to counteract resistance, leading to potential opinion shifts even when agents start from opposing views.

If Alice and Bob have zero persuasiveness ($\pi_A = \pi_B = 0$), the model behaves identically to the original model in [13]. In this case, if their resistance is high ($\rho_A = \rho_B = 1$), no opinion change occurs, confirming that persuasion is necessary for breaking deadlocks.

However, when both Alice and Bob are fully resistant ($\rho_A = \rho_B = 1$) but also entirely persuasive ($\pi_A = \pi_B = 1$), the system quickly reaches consensus on an alternative option (Fig.2).

It is important to emphasize that, although Alice and Bob converge rapidly to the alternative opinion (v_{alt}) under conditions of maximum resistance and persuasiveness, this state does not imply they necessarily adopt the same viewpoint. Instead, as explained in Section 2.3, v_{alt} represents a generalized category of diverse opinions, individually interpreted by each agent. Hence, real-world scenarios might require further differentiation among alternative opinions (e.g., v_{altA}, v_{altB}) or the inclusion of stochastic elements to prevent unrealistic deadlocks and more accurately capture nuanced opinion dynamics.

- Initially, Alice and Bob hold distinct opinions.

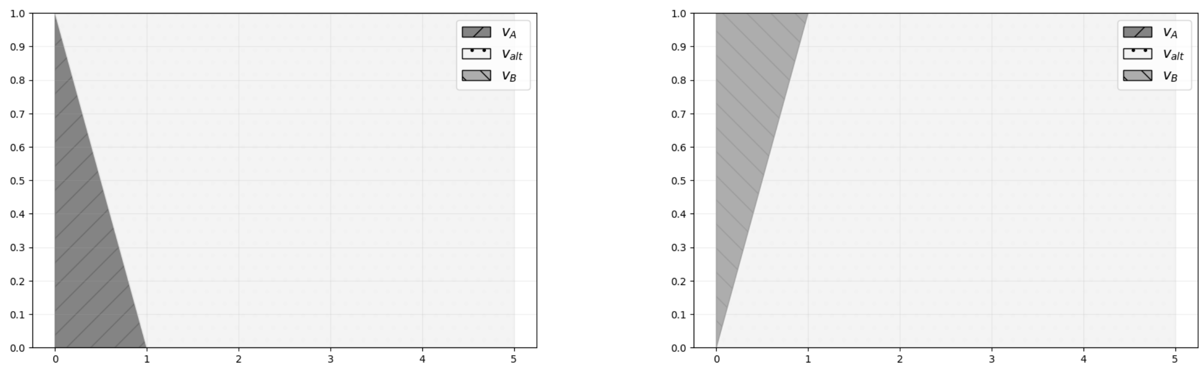

(a) Alice (b) Bob

Figure 2: Alice and Bob are persuasive and resistant

- Since they are both highly persuasive yet unwilling to adopt the other's viewpoint, they converge on an alternative option (v_{alt}) instead.
- This shift happens immediately after the first dialogue iteration, unlike the no-change scenario observed when $\pi_A = \pi_B = 0$.

If Bob, however, is resistant but not persuasive ($\rho_B = 1, \pi_B = 0$), he rejects Alice's option but still moves toward the alternative, as shown in Fig.3. This case demonstrates that a persuasive agent (Alice) can influence the dynamics while a non-persuasive agent (Bob) remains inert.

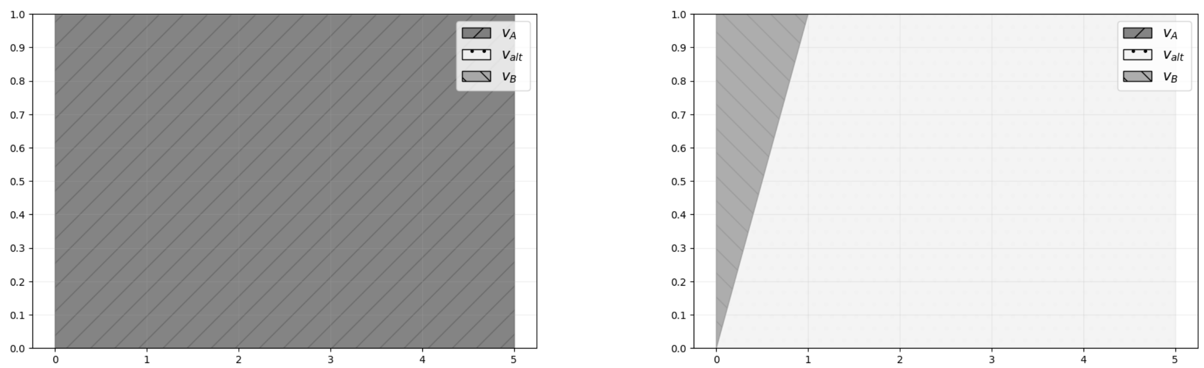

(a) Alice (b) Bob

Figure 3: Alice is resistant and persuasive, Bob is not persuasive but resistant

Tuning up Bob's persuasion power does not affect his own choices but influences Alice, who, after the first iteration, opens up for an alternative option (Fig.4)

3.3. The Role of Partial Persuasion and Resistance

In most real-world scenarios, agents are neither fully resistant nor fully persuadable. To explore this, we simulate cases where Alice and Bob have moderate levels of resistance and persuasion ($\rho_A = \rho_B = 0.8, \pi_A = \pi_B = 0.8$).

As Fig.5 illustrates:

- Gradual shifts occur over multiple dialogue rounds.
- Both agents move toward a mixed state, partially adopting each other's perspectives while considering alternative options.

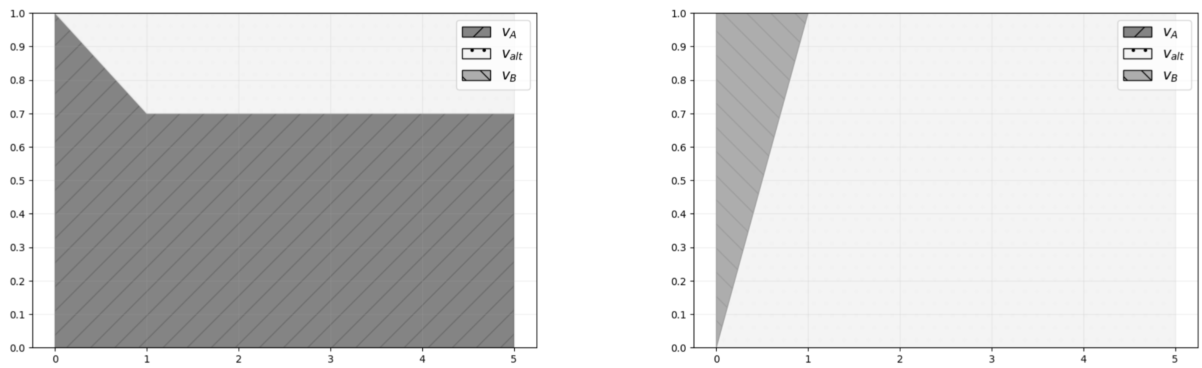

(a) Alice

(b) Bob

Figure 4: Bob is somewhat persuasive

- Unlike the extreme cases in previous sections, where no change or instant consensus occurred, this scenario represents a more realistic, gradual evolution of opinions.

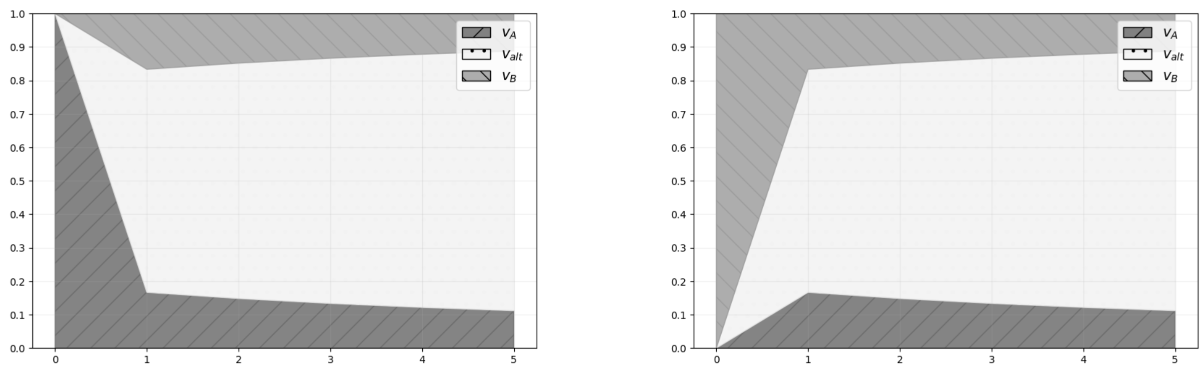

(a) Alice

(b) Bob

Figure 5: Moderate levels of resistance and persuasion

4. Conclusions

In this paper, we enhanced the original dialogue-based model of opinion dynamics by allowing agents to select alternative options beyond mere acceptance or rejection of opposing viewpoints. Our simulations illustrate how varying resistance and persuasiveness parameters significantly shape dialogue outcomes, yielding three behaviors:

- **Opinion Deadlock** (Fig. 1): High resistance and low persuasiveness lead to persistent disagreement without convergence.
- **Rapid Consensus on an Alternative** (Fig. 2, Fig. 3): High resistance combined with strong persuasive abilities quickly directs agents toward alternative opinions, avoiding mutual acceptance.
- **Gradual Opinion Shifts** (Fig. 5): Moderate persuasion and resistance levels result in incremental opinion changes, closely resembling real-world gradual consensus-building.

These results have practical implications for studying and countering psyops and contemporary information warfare. By simulating adversarial influence strategies and community-level responses, the model provides a computational tool for (i) identifying parameter regimes that make populations resilient or vulnerable to manipulation, (ii) testing detection heuristics for coordinated influence, and (iii) stress-testing mitigation strategies before deployment [1].

The novelty of the proposed framework lies in explicitly incorporating an *alternative-opinion* outcome and separating *persuasion* from *resistance* as independent drivers. Unlike traditional binary approaches, this structure captures dynamics observed in influence campaigns—rapid narrative pivots, echo-chamber stability, and asymmetric effects across subpopulations: phenomena documented in the context of the war in Ukraine [2, 4].

4.0.1. Limitations and Implications

Despite the strengths of the proposed model, several limitations should be recognized:

- **Aggregated Alternatives**: The simplification of grouping diverse alternative opinions into a single aggregated category neglects the complex landscape of real-world opinion diversity. This approach, although computationally efficient, reduces the granularity of possible outcomes, potentially obscuring more nuanced opinion dynamics. This simplification may also mask individual-specific variations and nuances in how different actors conceptualize alternatives.
- **Static Resistance and Persuasiveness Parameters**: The current model treats resistance and persuasiveness as static attributes, whereas in real-world scenarios these characteristics may dynamically evolve due to interaction history, social context, or external events.
- **Pairwise Interactions Only**: By focusing exclusively on pairwise dialogues, the model omits group dynamics, concurrent influences from multiple agents, and network-wide effects. This limitation could underestimate the complexity and speed of opinion propagation in highly connected networks, such as social media platforms.

These limitations suggest that while the model offers valuable insights into opinion evolution mechanisms, caution is warranted when directly extrapolating the findings to real-world scenarios without additional context-specific calibration.

4.0.2. Future Research Directions

- Extending the model for larger, more complex network interactions.
- Investigating the dynamic evolution of agent preferences based on interaction histories and network changes.
- Exploring the disaggregation of the aggregated alternative category into multiple distinct alternatives to capture richer, individual-specific opinion dynamics.
- Assigning empirically grounded values to resistance and persuasion parameters by
  - conducting empirical studies and controlled experiments to measure resistance and persuasiveness within defined populations or contexts;
  - analyzing historical datasets or opinion surveys to infer realistic parameter ranges and distributions, enhancing the model's calibration;
  - incorporating machine learning techniques to estimate these parameters dynamically from observed interactions in online communities or social media platforms'
- Introducing individualized or context-dependent persuasiveness and resistance parameters.
- Empirically validating predictions through observational studies on social media interactions and controlled experiments with structured dialogues.
- Applying and calibrating the model to real-world datasets related to psyops (e.g., platform takedown datasets, messaging telemetry, and annotated narrative corpora) to evaluate detection/mitigation strategies in situ.

Declaration on Generative AI

During the preparation of this work, the authors used ChatGPT and Grammarly to: Grammar and spelling check, Paraphrase and reword. After using this tool/service, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

References

[1] J. Kalenský, S. Rönnholm, K. Nyman-Metcalf, S. McInerney, et al., How Ukraine Fights Russian Disinformation: Beehive vs Mammoth, Technical Report Research Report 11, European Centre of Excellence for Countering Hybrid Threats (Hybrid CoE), 2024. URL: https://www.hybridcoe.fi/wp-content/uploads/2024/01/20240124-Hybrid-CoE-Research-Report-11-How-UKR-fights-RUS-disinfo-WEB.pdf.

[2] G. Căzănaru, M.-N. Stancu, Psychological Operations During the Russian War of Aggression in Ukraine, Review of the Air Force Academy (RAFT) XXII (2024) 49–58. URL: https://sciendo.com/article/10.2478/raft-2024-0037. doi:10.2478/raft-2024-0037.

[3] Z. Abrams, The role of psychological warfare in the battle for Ukraine, Monitor on Psychology (2022). URL: https://www.apa.org/monitor/2022/06/news-psychological-warfare.

[4] S. D. Bachmann, Hybrid Warfare and Disinformation: A Ukraine War Perspective, Global Policy 14 (2023) 63–76. URL: https://onlinelibrary.wiley.com/doi/10.1111/1758-5899.13257. doi:10.1111/1758-5899.13257.

[5] M. H. Degroot, Reaching a Consensus, Journal of the American Statistical Association 69 (1974) 118–121. URL: https://www.tandfonline.com/doi/abs/10.1080/01621459.1974.10480137. doi:10.1080/01621459.1974.10480137. arXiv:https://www.tandfonline.com/doi/pdf/10.1080/01621459.1974.10480137.

[6] N. E. Friedkin, E. C. Johnsen, Social influence and opinions, The Journal of Mathematical Sociology 15 (1990) 193–206. URL: https://doi.org/10.1080/0022250X.1990.9990069. doi:10.1080/0022250X.1990.9990069. arXiv:https://doi.org/10.1080/0022250X.1990.9990069.

[7] D. Acemoglu, A. Ozdaglar, Opinion Dynamics and Learning in Social Networks, Dynamic Games and Applications 1 (2011) 3–49. URL: https://doi.org/10.1007/s13235-010-0004-1. doi:10.1007/s13235-010-0004-1.

[8] A. Das, S. Gollapudi, K. Munagala, Modeling opinion dynamics in social networks, in: Proceedings of the 7th ACM International Conference on Web Search and Data Mining, WSDM '14, Association for Computing Machinery, New York, NY, USA, 2014, p. 403–412. URL: https://doi.org/10.1145/2556195.2559896. doi:10.1145/2556195.2559896.

[9] R. Abebe, J. Kleinberg, D. Parkes, C. E. Tsourakakis, Opinion Dynamics with Varying Susceptibility to Persuasion, 2018. URL: https://arxiv.org/abs/1801.07863. doi:10.48550/ARXIV.1801.07863.

[10] J. Semonsen, C. Griffin, A. Squicciarini, S. Rajtmajer, Opinion Dynamics in the Presence of Increasing Agreement Pressure, IEEE Transactions on Cybernetics 49 (2019) 1270–1278. URL: http://dx.doi.org/10.1109/TCYB.2018.2799858. doi:10.1109/tcyb.2018.2799858.

[11] S. E. Parsegov, A. V. Proskurnikov, R. Tempo, N. E. Friedkin, Novel Multidimensional Models of Opinion Dynamics in Social Networks, IEEE Transactions on Automatic Control 62 (2017) 2270–2285. URL: http://ieeexplore.ieee.org/document/7577815/. doi:10.1109/TAC.2016.2613905.

[12] A. De, S. Bhattacharya, P. Bhattacharya, N. Ganguly, S. Chakrabarti, Learning Linear Influence Models in Social Networks from Transient Opinion Dynamics, ACM Trans. Web 13 (2019). URL: https://doi.org/10.1145/3343483. doi:10.1145/3343483.

[13] G. Zholtkevych, O. Muradyan, K. Ohulchanskyi, S. Shelest, About One Approach to Modelling Dynamics of Network Community Opinion, in: Information and Communication Technologies in Education, Research, and Industrial Applications, Springer International Publishing, 2020, pp. 327–347. doi:10.1007/978-3-030-39459-2_15.

[14] M. Sherif, C. Hovland, Social Judgment: Assimilation and Contrast Effects in Communication and Attitude Change, Bloomsbury Academic, 1981.
[15] L. Festinger, A Theory of Cognitive Dissonance, Stanford University Press, 1957.
[16] D. C. Mutz, Hearing the Other Side: Deliberative versus Participatory Democracy, Cambridge University Press, 2006.
[17] C. Sunstein, #Republic: Divided Democracy in the Age of Social Media, Business book summary, Princeton University Press, 2018.
[18] Y. Lytvynenko, Extended Dialog Model Source Code, https://github.com/yurylyt/extended-dialog-model/tree/apr-2025, 2025. Accessed: 2025-04-20.

A. Online Resources

The source code for the model and figures is available via https://github.com/yurylyt/extended-dialog-model/tree/apr-2025 GitHub.