Limit Cycles in Opinion Dynamic Networks with Competing Stubborn Agents

Cody Moser*
School of Collective Intelligence, Mohammed VI Polytechnic University, Morocco

Paul E. Smaldino†
Department of Cognitive and Information Sciences, University of California, Merced, CA 95343, USA and Santa Fe Institute, Santa Fe, NM 87501, USA

(Dated: November 30, 2025)

The Counter-Optimal Stubborn Agent Placement (COSAP) problem concerns how competing agents attempt to position themselves within a network to maximize their influence over collective opinion dynamics. While prior work has primarily examined optimal placements in static or one-shot contexts, less is known about how opposing agents dynamically adjust their strategies in competitive environments. In this study, we model COSAP systems as a competitive diffusion process where two stubborn agents seek to influence a population through iterated network positioning. Simulations across networks of varying structure reveal that while agents generally race to occupy the most central node, networks with short path lengths create conditions in which neither agent can maintain an uncontested advantage. Instead, agents engage in adaptive cycles, repeatedly shifting positions in response to one other. These limit cycles vary in length and complexity depending on global network properties, with highly connected networks supporting persistent cycling and longer cycles, whereas networks with longer path lengths cause cycles to collapse into fixed points. These results underscore the importance of network structure in shaping influence dynamics and suggest that efforts to counter misinformation or manipulate opinion may grow more difficult in an increasingly connected world.

Keywords: Opinion dynamics, Zealots, Competitive diffusion, Propaganda, Networks

I. Introduction

Much attention has been paid to the role that online social networks have played on levels of increasing polarization in the United States and elsewhere [1, 2], with increasing attention paid to how polarization is shaped by the broader architectures of information flow in society [3–5]. Researchers have examined how the positioning of information sources in a social network influences the effectiveness of those sources to sway public opinion [6–9]. In many populations, opposing opinions compete for public sentiment, such that the best way to influence the public can depend on what the opposition is doing [10]. In this paper, we consider the dynamics of two competing opinion sources, each striving to position itself in the network to maximize its influence, conditional on the position of its opponent.

Opinion dynamic models, in which populations of agents express a discrete or continuous opinion and influence one another through local interactions, have previously been employed to examine how social network structure can influence the distribution of opinions within a population [11, 12]. Previous work has identified critical locations for the placement of stubborn agents in networks (also referred to as zealots or propagandists), operationalized either through the placement of agents which do not change their opinions or by modeling networks as being influenced by external propaganda sources who select specific agents to influence with their own opinions [13–18]. The key question for the analysis of such models is the optimal placement of stubborn agents to maximize influence throughout the network. This has been described as the "optimal stubborn agent placement" (OSAP) problem [19].

Analytic and computational solutions to the OSAP problem exist for systems in which stubborn agents of only one type are placed in a population (representing, for example, systems with a single foreign propaganda agent, about which the rest of the population is unaware) [20, 21]. The simplest solution to the single-agent version of the problem is to place the stubborn agent in nodes with the highest degree centrality, thereby maximizing the likelihood that the agent is sampled during the opinion update process [20]. Stubborn agent placement that maximizes other measures of network centrality appears to perform similarly [22]. In systems with multiple stubborn agents that have differing opinions, the problem is considered NP-hard, eluding simple analytic treatments [23]. Instead, several greedy algorithms have been proposed, including those that identify information bottlenecks based on random walks in the network [19, 21] and others that identify optimal diffusion patterns for capturing the maximal number of adjacent non-stubborn agents [20, 24]. Because the computational complexity

* cody.MOSER@um6p.ma
† psmaldino@ucmerced.edu

of what we term the Counter-Optimal Stubborn Agent Placement (COSAP) problem due to the dynamic nature of feedback between opposing agents, most existing approaches rely on static solutions, offering limited insight into the broader dynamics of counter-propaganda systems.

In this paper, we tackle the COSAP problem dynamically, considering how competing stubborn agents actively choose their network positions in response to one another. We introduce a competitive diffusion framework based on the simple voter model with two opposing stubborn agents. We conceptualize optimal agent placement as a process of reactionary relocation and analyze this dynamical system in terms of strategic agent positioning. We find that periodic cycles of network repositioning often emerge. By examining the length and complexity of these cycles in relation to network connectivity, we uncover how structural features shape the influence of propaganda agents on the spread of polarized states, how competitive dynamics can give rise to persistent and difficult-to-counter cycles of influence, and how these patterns relate to the broader challenge of mitigating misinformation in increasingly connected systems.

II. Previous Work

Wu and Huberman [20] studied a simple social influence model in which agents asynchronously adopted the dominant trait of their neighbors, with the addition of stubborn agents whose traits remained fixed. Their results showed that placing stubborn agents on nodes with the highest degree centrality maximized both the probability and speed of network-wide adoption of their trait, supporting this heuristic across various graph structures. Expanding on this, subsequent work on competitive diffusion in directed graphs [23, 25, 26] demonstrated that identifying optimal placements is NP-hard, and that greedy algorithms (which build solutions step by step by selecting the optimal choice at each stage) outperform strategies that rely solely on the stubborn agent's own centrality.

COSAP systems, where stubborn agents hold opposing opinions, are known to prevent consensus formation and sustain polarization [6, 13, 14, 27]. Even a small number of such agents can generate persistent fluctuations in collective opinion [13], typically centered between the two opposing views and following a Gaussian distribution whose variance decreases with the number of stubborn agents. Models incorporating fixed external "media sources" show that increasing their number or strength accelerates polarization and inhibits consensus [14, 27], while simulations using bounded confidence models suggest that ideological influence is maximized when fixed-opinion agents represent a moderate share of the network and maintain balanced connectivity [6].

Other work has specifically addressed optimal placement strategies for stubborn agents, particularly within the OSAP framework. For instance, analyses of binary opinion models have shown that placing a single stubborn agent at nodes with high degree centrality can maximize influence [19, 20]. However, these heuristics do not generalize to COSAP systems involving competing stubborn agents. To address this, greedy algorithms have been proposed to identify network locations where opposing agents are most likely to block each other's influence [21]. While such approaches reveal the inherent complexity of the problem, they provide limited insight into how different network structures shape the dynamics of polarization and strategic adaptation in COSAP settings, as they typically rely on static placement strategies and overlook the adaptive responses of competing agents over time. As a result, the relationship between network structure and the complexity of COSAP dynamics remains largely unexplored.

To address this gap, we model COSAP dynamics as an adaptive process in which agents reposition themselves in response to one another. By examining how these strategies evolve over time and mapping the range of possible solutions, we characterize the COSAP problem space in relation to network connectivity. Specifically, we investigate how the spread of opinions depends on the relative centrality of competing agents, identify cycles in optimal stubborn agent positions as a form of dynamic equilibrium, and examine how these cycles emerge primarily in highly connected networks with short average path lengths.

III. Model Description

We examine populations with two competing stubborn agents and explore how solutions to the COSAP problem become interdependent, such that the optimal placement for one agent type depends critically on the positioning of its opponent. We examine COSAP dynamics across a variety of network structures and levels of connectivity, analyzing how agent positioning shapes competitive dynamics. In particular, we seek to understand whether agents converge on stable placements or continuously adapt in response to one another, giving rise to cycling behavior without a fixed equilibrium.

Our model simulates opinion dynamics in a competitive COSAP environment using the generalized voter model [28], where agents interact over a structured network. To identify optimal placements for stubborn agents, we assume centrality as the key dimension along which influence operates. For each simulation, we fix the position of one stubborn agent based on its percentile rank in a given centrality measure (degree, betweenness, or closeness) and then search for the position of the opposing agent that maximizes its ability to shift the average opinion in its favor. This process is repeated across simulations: once a candidate placement for one agent is fixed, we simulate various configurations for the opposing agent and determine the position that offers the best

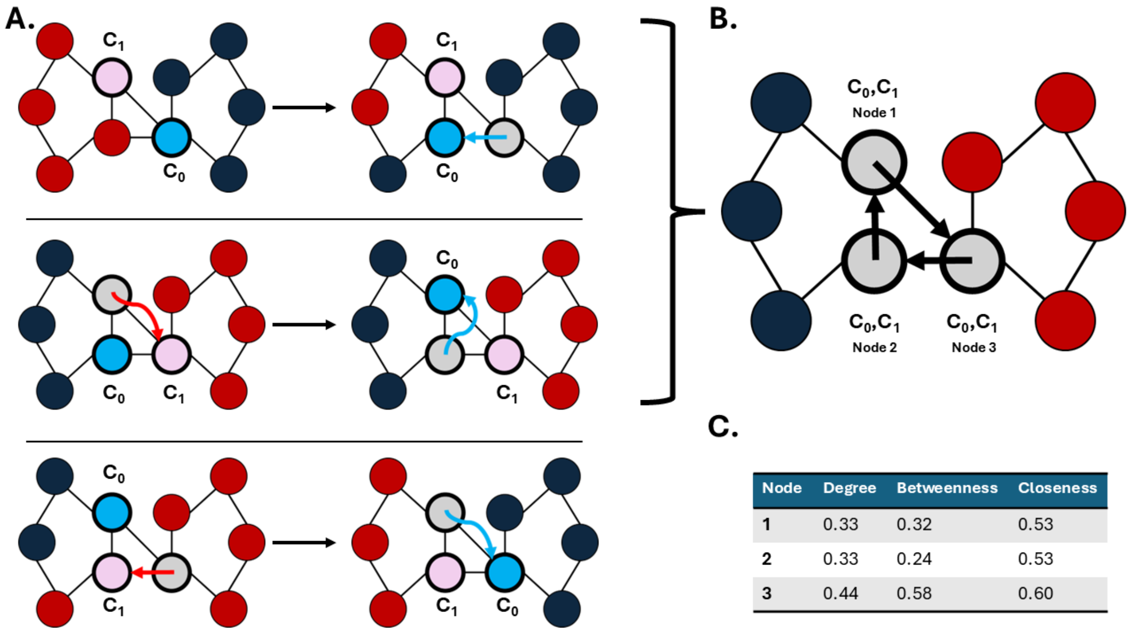

Node	Degree	Betweenness	Closeness
1	0.33	0.32	0.53
2	0.33	0.24	0.53
3	0.44	0.58	0.60

FIG. 1. Schematic of the COSAP model of stubborn agent placement. (A) Agent c_0 first selects the most central node (Node 3) to optimize its influence over a network. A competing agent, c_1 (pink) then selects an optimal position of centrality to disrupt c_0's (light blue) influence. In each iteration, agents re-evaluate their positions and move to new nodes that best disrupt their competitor's influence. This creates a repeating three-node cycle. Gray nodes represent the previous position held by the moving agent. (B) The limit cycle formed between Nodes 1, 2, and 3. COSAP repositioning forces both agents to rotate through these nodes in a stable pattern. (C) Centrality measures for nodes within c_0's and c_1's cycle. Even though Node 3 is the most central node for across degree, betweenness, and closeness centralities, competitive interference yields a dynamic attractor involving nodes of varying centrality.

response. This allows us to identify whether agents settle into a stable configuration or whether no fixed point emerges, in which case we observe cycling behavior where each agent's best response depends on the other's position. Through this approach, we examine how competitive placement strategies and the emergence of dynamic equilibria are shaped by underlying network structure.

We implement three types of network topologies: (1) Erdős–Rényi random networks, varying the edge formation probability incrementally from 0.1 to 1.0; (2) small-world networks, varying the rewiring probability over the same range [29]; and (3) connected caveman networks augmented with long-range rewiring to introduce global connectivity [30, 31]. This range of structures allows us to systematically examine how differences in network connectivity and topology affect the dynamics of agent competition and influence.

A. Initialization

1. **Network Setup:** N agents are placed on a network of a given structure, with each agent randomly assigned a binary-valued opinion. A proportion γ of agents begin with Opinion 0, and the remaining $1-\gamma$ with Opinion 1; for all simulations, we set $\gamma = 0.5$. We compute degree, betweenness, and closeness centrality measures for all nodes and rank each node by its relative centrality percentile within the network. These rankings are used to guide the placement of stubborn agents in subsequent simulations.

2. **Stubborn Agent Selection:** Two agents are designated as stubborn based on their centrality percentiles, measured using either degree, betweenness, or closeness centrality. The stubborn agent holding Opinion 0 is selected according to one centrality percentile (c_0), and the agent holding Opinion 1 according to another (c_1). These values represent relative ranks within the network: $c_x = 1$ corresponds to the most central node, $c_x = 0.5$ to a node with median centrality, and $c_x = 0$ to the least central node.

B. Dynamics

1. **Opinion Updating:** At each time step, one agent is selected uniformly at random. The selected agent randomly selects a neighbor. If the two hold different opinions, the selected agent adopts the neighbor's opinion. Otherwise, no change occurs.

2. **End and Repetition:** The simulation runs for 20,000 time steps, which is sufficient to observe the long-term distribution of opinions in the population, as supported by patterns displayed in Figure 2 (bottom: $M_{steps} = 9,264, sd_{steps} = 7,965$). This process is repeated for 300 iterations across

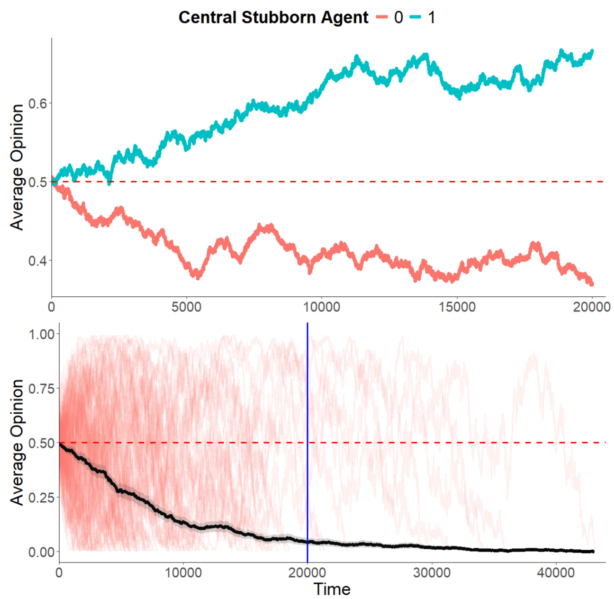

FIG. 2. Opinion trajectories across simulation runs. (Top): Average opinion over time in COSAP simulations, where one stubborn agent is fixed at closeness centrality $c_x = 1$ and the other at $c_x = 0$. Lines represents the average across runs. Colors indicate which agent holds $c_x = 1$. (Bottom): Average opinion over time in simulations with a single stubborn agent fixed at centrality $c_0 = 1$. Thin lines represent 50 individual runs, the bold line shows the average opinion of all simulations at each timestep with light gray showing the standard error. The vertical blue line marks the stopping point used in our COSAP analyses.

all combinations of (c_0, c_1) values, where each value ranges from the 0th to the 100th percentile in increments of 10. This allows us to evaluate how each agent's placement performs relative to the other's and to identify best-response dynamics and potential cycles in agent positioning.

C. Data Collected

For each simulation, we collect the average opinion across the entire population at the end of each run. Because opinions are binary, this value reflects the proportion of agents adopting Opinion 1. Given the stochastic nature of the simulations, where final average opinions tend to jitter around 50% regardless of agent placement, we mitigate noise by calculating the average opinion over the final 100 timesteps of each run. For each network configuration, we also assess the average path length of the network, the shortest path distance between the two stubborn agents, and their centrality values. In simulations where agent repositioning yields limit cycles, we additionally record the *length* of the cycle, defined as the number of distinct agent configurations, and the *distance* of the cycle, measured as the average change in centrality percentile between successive positions.

We define the optimal position for a stubborn agent as the centrality percentile at which it captures the highest proportion of the population's opinions, given the fixed position of the opposing stubborn agent. Because opinions do not converge in systems with multiple stubborn agents [14], we define influence operationally as the long-run deviation from 50% opinion share in a network Figure 2 (top). For each network configuration, we assess the optimal placement of agent (c_0) relative to a fixed position of agent (c_1). We then reverse the procedure, using the identified optimal position of (c_0)as a baseline to determine improved placements for (c_1) in identical network structures. This iterative approach allows us to explore the interdependence of agent positions in competitive COSAP dynamics.

IV. Results

A. Centrality and Influence

Optimal agent positions are assessed based on their ability to influence the average opinion of the network. With two opposed stubborn agents, both opinions typically persist in nearly equal numbers. Because individual simulation runs are highly stochastic, we evaluate average agent influence by comparing outcomes across repeated runs. Figure 2 (top) shows the temporal evolution of average opinion across simulations in which either c_0 or c_1 occupies the most central position in the network. While individual trajectories remain noisy throughout, a small but consistent drift emerges over time in favor of the more centrally positioned stubborn agent.

To evaluate the optimality of agent positioning more systematically, we analyze three centrality metrics: degree, betweenness, and closeness centrality. We compare the influence of c_0 across a range of centrality percentiles, given fixed centrality values for c_1, by assessing the average opinion of the population over the final 100 timesteps of each simulation. While our analysis spans networks of varying types and connectivity, Figure 3 presents results from a representative case of random networks with 100 agents and a wiring probability of 0.5.

Across all three centrality measures and all values of c_1, we observe that the influence of c_0 increases with its relative centrality. This pattern supports findings that more central positions confer a better COSAP advantage. However, the strength of this effect varies modestly across centrality types and depends on the network structure and the positioning of the competing agent. Notably, when the opposing agent is placed at a highly central node (i.e., when c_1 is high), the ability of the stubborn agent assigned to c_0 to sway the majority opinion is substantially reduced, regardless of its own centrality percentile. While the presence of two COSAP agents prevents convergence to any one opinion settles in a generally polarized state, we find that the degree of polarization, or the dominance of one opinion over another, is shaped by the relative positions of the two COSAP agents.

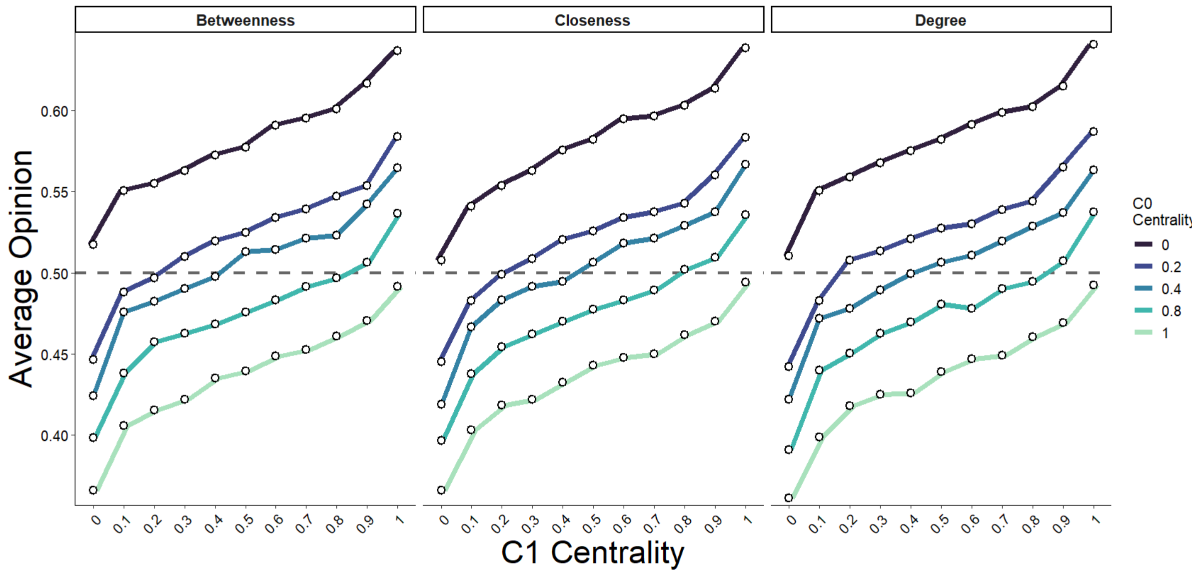

FIG. 3. Average opinion as a function of the centrality of a stubborn agent (c_0), with Opinion 0, under varying centralities of the other stubborn agent c_1, with Opinion 1, in random networks with 100 agents and a wiring probability of 50%. Each panel corresponds to a different centrality metric (Betweenness, Closeness, Degree) used to define stubborn agent positioning, grouped by the centrality of c_0. The x-axis represents the centrality of c_1, and colored lines denote different levels of c_0 centrality (from low to high). White points show individual simulation outcomes, and the black outlines represent mean opinion values. Points show individual simulation outcomes; white points and black lines represent mean opinions for all simulations. The dashed horizontal line at 0.5 indicates the baseline average opinion.

B. Identifying Limit Cycles

COSAP dynamics are commonly approached as a static optimization problem [20, 21]. However, this approach overlooks the dynamic nature of competition between agents: when one agent adjusts its position to shift network influence in its favor, the opposing agent's optimal response also changes. As such, static analyses like those in the previous section may fail to capture the adaptive interplay between competing agents attempting to influence the average opinion within the network. To address this, we examine agent positions dynamically, focusing on how each agent's optimal location changes in response to the other's placement. Specifically, we identify optimal positions for one agent (c_0) while holding the other (c_1) fixed, then use that result to fix c_0's new position and identify the optimal response by c_1. In doing so, we observe that in many networks, this process converges quickly on a stable solution, typically a "race to the center" dynamic in which agents compete to occupy the most central node. This can result in one agent effectively locking out the influence of its opponent.

However, in a subset of networks we observe the emergence of limit cycles in agent positioning. Rather than converging on a fixed optimal node, agents continually adapt to one another by cycling through a sequence of positions in an ongoing attempt to recapture influence. Such behavior is analogous to periodic cycles in strategic games, where no fixed point exists and players continuously adjust strategies in response to one another [32–34]. The characteristics of these cycles vary: some involve only a few repeated positions (short cycles), while others span a wider range of centralities. Figure 4 illustrates examples of this behavior across random networks with varying edge probabilities ($p = 0.7, 0.8, 0.9$). Each panel shows transitions in centrality-based positions for two competing stubborn agents, with arrows representing movement from one simulation to the next and colors indicating which agent was dominant in each round, defined as the one controlling a majority of opinions in the final 100 timesteps. This suggests that simply occupying the most central node does not guarantee dominance, as influence is context-dependent: in some cases, positioning at a slightly less central node can offer a better strategic response to the opposing agent's placement, leading to continued repositioning rather than convergence.

C. Network Properties and Limit Cycle Length

To investigate how network structure influences the emergence and persistence of limit cycles, we conducted simulations across networks varying in path length and connectivity, including random, small world, and connected caveman networks with long-range rewiring. Our analysis reveals a relationship between network path length and the structure of limit cycles: as the average path length of the network increases, both the length of the cycles, defined as the number of distinct agent configurations, and their distance, defined as the average centrality shift between successive positions, tend to diminish (Fig. 5).

In highly connected networks with short path lengths, agents cycle through a wider set of positions (Fig. 5, top),

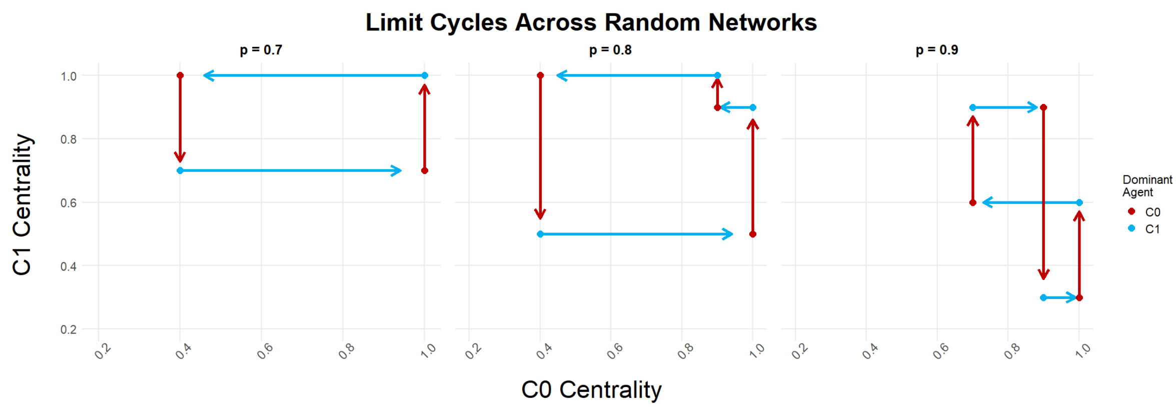

FIG. 4. Each panel shows the recurrent sequence of optimal placements for two competing stubborn agents in Erdős–Rényi random networks with wiring probabilities $p = 0.7$, $p = 0.8$, and $p = 0.9$ Arrows indicate the transitions between successive best-response moves, and colors denote which agent is dominant at each step. The resulting trajectories form stable periodic cycles, whose structure depends on network density. Higher values (denser networks) produce longer cycles, whereas sparser networks produce shorter and more constrained limit cycles.

as no single position remains permanently optimal and each agent's move induces a counter-move from its opponent. In these networks, competition generates cycles rather than settling into a fixed point. By contrast, networks with longer path lengths tend to produce smaller, less variable cycles or stable configurations. The size of these cycles, in terms of the centrality distance between each subsequent position, also increases in highly connected networks, reflecting larger swings in strategic repositioning (Fig. 5, bottom). This is likely because in highly connected networks, centrality is effectively devalued where many nodes have similar access to the rest of the network, and influence tends to overlap extensively between agents. As a result, no single position provides a lasting strategic advantage, enabling agents to continually reposition in response to one another. In contrast, networks with longer path lengths feature more differentiated and spatially distinct central nodes, reducing overlap in influence and allowing agents to settle into stable positions with minimal interference. These findings suggest that network connectivity plays a key role in shaping the landscape of competitive influence, with highly connected networks enabling longer cycles of adaptability and ongoing positional competition.

While the scale of cycles varies with network structure, we find no evidence of multiple distinct cycles emerging within the same network. Instead, agents consistently converge on a single dynamic loop or settle into a fixed position. This suggests that, within the structural constraints of a given network, there is typically one dominant adaptive pathway, limiting the diversity of long-term strategic behaviors. Although the COSAP system is finite and deterministic and we thus observed no evidence of chaotic behavior, cycles of length three or greater are notable because they can mark the onset of qualitatively richer competitive dynamics [35].

D. Basins of Attraction Reveal Strategic Resilience

Some limit cycles exert more influence over the state space than others. A cycle that attracts many different starting configurations occupies a larger portion of the system's dynamics andplays a more dominant role in long-run behavior. To capture this, we examined how many independent pathways of successive COSAP behavior lead into a given cycle. These pathways represent distinct channels of convergence from the broader state space of COSAP dynamics. The set of all such pathways that flow into a cycle constitutes a basin of attraction for the given network configuration. Cycles drawing more independent pathways exhibit larger basins and have greater dynamical reach across the network, as they draw in more of the position space and structure a larger share of the COSAP dynamics. Fig. 6 shows the relationship between basin strength and network structure. For each network realization, we computed the number of independent pathways feeding into its dominant attractor and plotted the proportion of these leading into the cycle against the network's average path length (on a logarithmic scale). The dashed vertical line marks the point at which cycles begin to collapse, giving way to fixed-point attractors. Random networks with close path lengths tend to support attractors with large basins and many independent pathways, while higher-path-length structures especially small-world networks show a pronounced reduction in pathway diversity and basin size. This pattern indicates that denser or more locally clustered networks sustain richer competitive dynamics between stubborn agents, whereas long-path-length networks restrict the available attractor structure.

These findings suggest that networks with shorter path lengths support more complex competitive dynamics. Their attractors are fed by numerous independent pathways, meaning that a wide range of initial conditions can be pulled into dynamical cycles. In contrast, net-

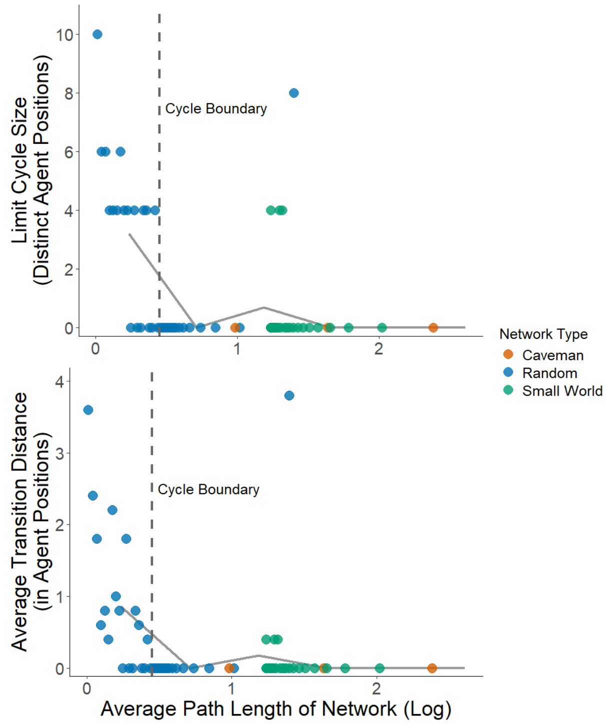

FIG. 5. Limit cycles length as a function of network structure. As the average path length of a random network increases, the complexity of limit cycles diminishes. (Top): The number of distinct agent configurations involved in a limit cycle. (Bottom): The average "centrality distance" between successive configurations within each cycle. The dashed vertical line marks the boundary where cycles emerge in random networks.

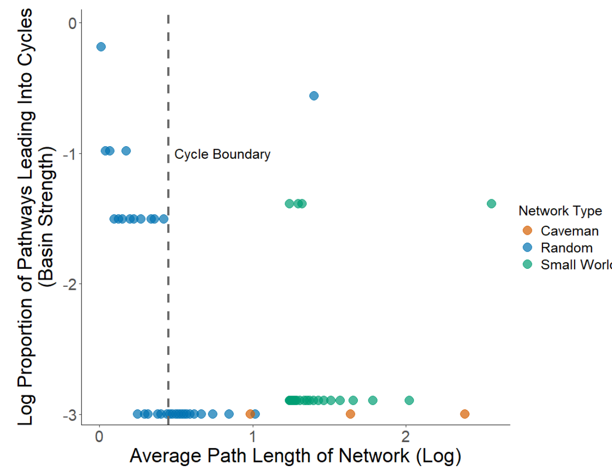

FIG. 6. Basin strength of COSAP attractors as a function of network structure. Each point represents the log proportion of independent pathways external to a limit cycle for all networks. Networks at the highest limit exhibit independent pathways with no cycle. Basin strength corresponds to the number of independent pathways leading into an attractor. The x-axis shows the network's average path length (log scale). The dashed vertical line marks the boundary where cycles emerge in random networks

works with longer path lengths compress space of options, leading to fewer independent pathways and the eventual disappearance of cycles altogether. These findings suggest that networks with larger basins and more independent pathways enable greater opportunities for disruption between competing stubborn agents. When many pathways lead into an attractor, agents can enter and disrupt limit cycles from a wide range of starting positions. In contrast, when only a few pathways exist, agents have fewer potential responses, and the resulting cycles exhibit reduced complexity and limited potential for strategic disruption.

V. Discussion

Previous work on the optimal placement of stubborn agents in diffusion and voter models has generally found that the most central nodes exert the greatest influence [20, 21, 26]. COSAP studies, which examine competitive dynamics between opposing agents, have typically modeled agent positions as fixed, assuming static placements throughout the diffusion process [21]. In this study, we examine the dynamic repositioning of agents and show that COSAP placement strategies based on relative centrality can lead to a range of outcomes, including polarization driven by the sustained dominance of a single agent. The extent of this polarization is directly tied to the centrality of each agent, with more central placements generally conferring greater influence. Furthermore, we identify a form of dynamic polarization that emerges when agents adaptively reposition themselves in response to one another. Unlike static models that assume equilibrium placements, our simulations reveal that stable solutions do not always emerge. Instead, agents can enter limit cycles, repeatedly shifting between multiple positions in an ongoing strategic contest for influence. These cycles reflect a form of persistent competition, in which each agent's optimal response depends on the placement of its opponent. Notably, such dynamics are most pronounced in highly connected networks, where short average path lengths "devalue" the strategic advantage of any single central node. In these settings, influence becomes more diffuse, allowing agents to continually challenge and displace one another. This cycling behavior underscores the difficulty of identifying stable influence in certain networks, where no single placement remains optimal and strategic advantage continually shifts between agents.

These findings suggest that in highly-connected network structures, optimal solutions to the COSAP problem may be inherently unstable or elusive. While it is unlikely that adversarial agents, such as coordinated misinformation campaigns, explicitly strategize based on formal OSAP formulations, our results suggest that the difficulty of countering propaganda may vary substantially across different online ecosystems. In particular, the presence of multiple stable or near-optimal positions for stubborn agents implies that influence can persist or re-emerge even after targeted countermeasures. Furthermore, in cycles with extensive basins, agents can enter cycles from a broad range of positions, amplifying the difficulty of suppressing external influence. This highlights a broader feature of online networks: in systems where multiple influence hubs exist, exerting and maintaining sway over public opinion may be easier than in more centralized environments. Our analysis also highlights why strategic efforts to control network structures directly may be more effective for influencing public opinion than merely adding to the public debate while leaving the network otherwise unchanged [36, 37]. As digital platforms have become increasingly dense and interconnected [38], understanding these dynamics is critical for developing effective strategies to enhance resilience against manipulation, misinformation, and polarization.

The findings presented here highlight the strategic instability of influence in COSAP systems. Even when one agent appears to dominate based on network placement, that advantage can quickly be reversed if the opposing agent identifies a better counter-position. In many cases, there is no universally optimal node; instead, influence unfolds within a shifting landscape where no single strategy remains dominant. Future work may seek to identify the structural conditions that give rise to stable placements and to develop methods for approximating influence outcomes without exhaustive simulation. Our framework also offers a foundation for assessing network vulnerability or susceptibility to influence, providing a potential tool for understanding how control or persuasion might be achieved (or prevented) in networks with differing structural properties, and informing platform design choices aimed at reducing strategic manipulation or polarization.

Acknowledgments

The research was funded by a grant from the Templeton World Charity Foundation (TWCF0715) awarded to PES, a grant from the Institute for Humane Studies (IHS017447) awarded to CM, and the Center for Analytic Political Engagement at UC Merced.

[1] M. P. Fiorina and S. J. Abrams, Political polarization in the American public, Annu. Rev. Polit. Sci. **11**, 563 (2008).

[2] A. K. Dixit and J. W. Weibull, Political polarization, Proceedings of the National Academy of sciences **104**, 7351 (2007).

[3] I. Waller and A. Anderson, Quantifying social organization and political polarization in online platforms, Nature **600**, 264 (2021).

[4] P. E. Smaldino, A. Russell, M. Zefferman, J. Donath, J. G. Foster, D. Guilbeault, M. Hilbert, E. A. Hobson, K. Lerman, H. Miton, C. Moser, J. Lasser, S. Schmer-Galunder, J. N. Shapiro, Q. Zhong, and D. Patt, Information architectures: A framework for understanding socio-technical systems, npj Complexity (Forthcoming).

[5] E. Kubin and C. Von Sikorski, The role of (social) media in political polarization: a systematic review, Annals of the International Communication Association **45**, 188 (2021).

[6] H. Z. Brooks and M. A. Porter, A model for the influence of media on the ideology of content in online social networks, Physical Review Research **2**, 023041 (2020).

[7] M. Cinelli, E. Brugnoli, A. L. Schmidt, F. Zollo, W. Quattrociocchi, and A. Scala, Selective exposure shapes the facebook news diet, PloS one **15**, e0229129 (2020).

[8] P. E. Smaldino, *Modeling Social Behavior: Mathematical and Agent-Based Models of Social Dynamics and Cultural Evolution* (Princeton University Press, Princeton, NJ, 2023).

[9] A. Das, S. Gollapudi, and K. Munagala, Modeling opinion dynamics in social networks, in *Proceedings of the 7th ACM international conference on Web search and data mining* (2014) pp. 403–412.

[10] D. P. Baron, Competing for the public through the news media, Journal of Economics & Management Strategy **14**, 339 (2005).

[11] D. Acemoglu and A. Ozdaglar, Opinion dynamics and learning in social networks, Dynamic Games and Applications **1**, 3 (2011).

[12] K. Eekhoff, Opinion formation dynamics with contrarians and zealots, SIAM J. Undergraduate Research Online **12** (2019).

[13] M. Mobilia, A. Petersen, and S. Redner, On the role of zealotry in the voter model, Journal of Statistical Mechanics: Theory and Experiment **2007**, P08029 (2007).

[14] D. Bhat and S. Redner, Nonuniversal opinion dynamics driven by opposing external influences, Physical Review E **100**, 050301 (2019).

[15] S. M. Shekatkar, Do zealots increase or decrease the polarization of social networks?, Journal of Complex Networks **8**, cnz036 (2020).

[16] T. Tunstall, How social network structure impacts the ability of zealots to promote weak opinions, Physical Review E **111**, 024311 (2025).

[17] A. Sîrbu, V. Loreto, V. D. Servedio, and F. Tria, Opinion dynamics: models, extensions and external effects, Participatory sensing, opinions and collective awareness , 363 (2017).

[18] D. L. Arendt and L. M. Blaha, Opinions, influence, and zealotry: a computational study on stubbornness, Computational and Mathematical Organization Theory **21**, 184 (2015).

[19] E. Yildiz, D. Acemoglu, A. E. Ozdaglar, A. Saberi, and A. Scaglione, Discrete opinion dynamics with stubborn agents, Available at SSRN 1744113 (2011).

[20] F. Wu and B. A. Huberman, Social structure and opinion formation, arXiv preprint cond-mat/0407252 (2004).

[21] E. Yildiz, A. Ozdaglar, D. Acemoglu, A. Saberi, and A. Scaglione, Binary opinion dynamics with stubborn agents, ACM Transactions on Economics and Computation (TEAC) **1**, 1 (2013).

[22] D. Guilbeault and D. Centola, Topological measures for identifying and predicting the spread of complex contagions, Nature Communications **12**, 4430 (2021).

[23] D. Kempe, J. Kleinberg, and É. Tardos, Maximizing the spread of influence through a social network, in *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining* (2003) pp. 137–146.

[24] X. Lin, Q. Jiao, and L. Wang, Competitive diffusion in signed social networks: A game-theoretic perspective, Automatica **112**, 108656 (2020).

[25] P. Domingos and M. Richardson, Mining the network value of customers, in *Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining* (2001) pp. 57–66.

[26] D. Kempe, J. Kleinberg, and É. Tardos, Influential nodes in a diffusion model for social networks, in *international colloquium on automata, languages, and programming* (Springer, 2005) pp. 1127–1138.

[27] D. Bhat and S. Redner, Polarization and consensus by opposing external sources, Journal of Statistical Mechanics: Theory and Experiment **2020**, 013402 (2020).

[28] R. A. Holley and T. M. Liggett, Ergodic theorems for weakly interacting infinite systems and the voter model, The annals of probability , 643 (1975).

[29] D. J. Watts and S. H. Strogatz, Collective dynamics of 'small-world' networks, Nature **393**, 440 (1998).

[30] A. Flache and M. W. Macy, Small worlds and cultural polarization, in *Micro-Macro Links and Microfoundations in Sociology* (Routledge, 2014) pp. 146–176.

[31] M. A. Turner and P. E. Smaldino, Paths to polarization: How extreme views, miscommunication, and random chance drive opinion dynamics, Complexity **2018**, 2740959 (2018).

[32] L. S. Shapley, Some topics in two-person games, in *Advances in Game Theory*, Vol. 52 (Princeton University Press, 1964) pp. 1–28.

[33] Y. Sato and J. P. Crutchfield, Coupled replicator equations for the dynamics of learning in multiagent systems, Physical Review E **67**, 015206 (2003).

[34] P. J.-J. Herings and A. Predtetchinski, Best-response cycles in perfect information games, Mathematics of Operations Research **42**, 427 (2017).

[35] T.-Y. Li and J. A. Yorke, Period three implies chaos, The american mathematical monthly **82**, 985 (1975).

[36] R. DiResta, *Invisible rulers: The people who turn lies into reality* (PublicAffairs, 2024).

[37] A. Marzouk, Death of twitter and the rise of x: The birth of the digital authoritarian context, Dialogues on Digital Society **0** (2025).

[38] J.-G. Liu, Z.-M. Ren, Q. Guo, and D.-B. Chen, Evolution characteristics of the network core in the facebook, PLoS One **9**, e104028 (2014).

[39] F. Baumann, P. Lorenz-Spreen, I. M. Sokolov, and M. Starnini, Modeling echo chambers and polarization dynamics in social networks, Physical Review Letters **124**, 048301 (2020).