Frontiers in Computing and Intelligent Systems
ISSN: 2832-6024 | Vol. 17, No. 2, 2026

Centrality-driven Sparse Optimal Control of Belief Formation in Social Networks

Bo Wang

College of Intelligent Robotics and Advanced Manufacturing, Fudan University, Shanghai, China

Abstract: This paper studies the problem of steering collective beliefs in social networks when only a small fraction of nodes can be directly influenced. We propose a sparse optimal control framework built on the Network Drift-Diffusion Model (NDDM). Two intervention mechanisms are considered: direct control and latent (indirect) control. To select which nodes to actuate, we compare six centrality measures---Degree, Betweenness, Eigenvector, Closeness, PageRank, and K-shell---and keep only the top 5\%--30\% as control inputs. The optimal feedback law follows from the Hamilton-Jacobi-Bellman (HJB) equation, which reduces to solving Riccati-type differential equations. We test our approach on Erdős-Rényi (ER), Barabási-Albert (BA), and Watts-Strogatz (WS) networks. The results show that the best centrality choice depends strongly on the network topology, and the system undergoes phase transitions as control parameters vary.

Keywords: Complex Networks; Dynamical Systems; Optimal Control; Sparse Actuation; Structural Centrality; Numerical Simulation.

1. Introduction

Social networks often amplify collective beliefs into rapid consensus or polarization, especially during public health emergencies, financial crises, or political disputes. These processes arise from local interactions and individual observations. Although non-Bayesian learning models offer a useful baseline[1, 2], real networks are constantly shaped by external inputs such as algorithmic recommendations or official campaigns. A natural question then is: how should one design these interventions mathematically? Earlier studies have used stubborn agents, bounded-confidence rules, or linear-quadratic games to model such problems[3, 4]. Yet nearly all of these works assume that every node can be controlled or observed. This assumption fails in large-scale online platforms, where budget, privacy, and technical constraints limit intervention to a tiny subset of users[5–7]. In reality, an external controller can inject signals into only an ultra-sparse subset of nodes, raising a fundamental question: under extreme resource scarcity ($p \to 0$), which nodes should be prioritized to maximize global regulatory benefit?

To answer this, we introduce a structural-centrality-driven sparse control framework. Centrality metrics—degree, betweenness, eigenvector, closeness, PageRank, and K-shell[8–11]—offer distinct mathematical perspectives on a node's network importance. Our contributions are threefold: (i) we synthesize a sparse actuation mechanism using these metrics within the NDDM; (ii) we derive the closed-loop optimal control law via HJB theory; (iii) we discover phase transitions and efficiency inversion laws through comprehensive simulations on canonical topologies.

2. Controlled Social Learning Dynamics and Sparse Intervention Mechanisms

Consider a network of n agents with row-stochastic adjacency $\mathbf{A} \in \mathbb{R}^{n\times n}$. The log-belief ratio vector $\mathbf{x}(t) \in \mathbb{R}^n$ evolves under two intervention paradigms, which is conceptually illustrated in Fig. 1:

1) Direct intervention:

$$\mathrm{d}\mathbf{x}(t) = (\mathbf{A} - \mathbf{I}_n)\mathbf{x}(t)\,\mathrm{d}t + \mathbf{S}\mathbf{u}(t)\,\mathrm{d}t, \tag{1}$$

where $\mathbf{u}(t) \in \mathbb{R}^m$ is the control signal and $\mathbf{S} \in \mathbb{R}^{n\times m}$ is the sparse selection matrix.

2) Latent intervention:

$$\mathrm{d}\mathbf{x}(t) = ((1-\gamma)\mathbf{A} - \mathbf{I}_n)\mathbf{x}(t)\,\mathrm{d}t + \gamma\mathbf{S}\mathbf{u}(t)\,\mathrm{d}t, \tag{2}$$

with $\gamma \in (0,1)$ controlling the coupling strength.

Given a budget proportion p, we select $m = \lfloor pn \rfloor$ nodes based on a centrality score $S_{\mathcal{J}}(i)$ for $\mathcal{J} \in \{\text{Degree}, \text{Betweenness}, \text{Eigenvector}, \text{Closeness}, \text{PageRank}, \text{K}-\text{shell}\}$. The columns of $\mathbf{S}$ are the corresponding standard basis vectors. Definitions are standard: degree counts direct neighbors; betweenness counts shortest-path fractions; eigenvector is the leading eigenvector of $\mathbf{A}$; closeness is reciprocal average shortest-path distance; PageRank uses damping factor 0.85; K-shell assigns coreness via recursive peeling.

We minimize the finite-horizon cost

$$J = \|\mathbf{x}(T) - \mathbf{x}^*\|^2 + a\int_0^T \|\mathbf{u}(t)\|^2\,\mathrm{d}t, \tag{3}$$

where $a > 0$ is the energy penalty and $\mathbf{x}^*$ is the target consensus. To streamline the analytical synthesis and achieve a unified numerical solution framework for both the direct and latent intervention mechanisms, we cast the distinct controlled state-space equations into the following generalized high-dimensional linear dynamical system form:

$$\mathrm{d}\mathbf{x}(t) = (\mathbf{B}_x\mathbf{x}(t) + \mathbf{B}_u\mathbf{u}(t))\mathrm{d}t, \tag{4}$$

where $\mathbf{B}_x \in \mathbb{R}^{n\times n}$ acts as the system drift matrix, and $\mathbf{B}_u \in \mathbb{R}^{n\times m}$ is the generalized control configuration input operator. The exact mappings of these parameterized matrices under different actuation modalities are explicitly dictated by:

1) Direct Intervention Mode: $\mathbf{B}_x = \mathbf{A} - \mathbf{I}_n$, $\mathbf{B}_u = \mathbf{S}$,

2) Latent Intervention Mode: $\mathbf{B}_x = (1-\gamma)\mathbf{A} - \mathbf{I}_n$, $\mathbf{B}_u = \gamma\mathbf{S}$,

where $\mathbf{S}$ is the topology-centric sparse allocation matrix determined via structural centralities in the previous subsection. We now derive the optimal feedback law. Apply the Hamilton-Jacobi-Bellman (HJB) principle to the linear-quadratic problem defined above. The resulting closed-loop control takes the form,

$$\mathbf{u}^*(t) = -\frac{1}{a}\mathbf{B}_u^\top(\mathbf{P}(t)\mathbf{x}(t) - \mathbf{v}(t)), \tag{5}$$

where $\mathbf{P}(t)$ and $\mathbf{v}(t)$ solve the coupled differential Riccati and adjoint equations:

$$\dot{\mathbf{P}} + \mathbf{P}\mathbf{B}_x + \mathbf{B}_x^\top\mathbf{P} - \frac{1}{a}\mathbf{P}\mathbf{B}_u\mathbf{B}_u^\top\mathbf{P} = \mathbf{0}, \quad \mathbf{P}(T) = \mathbf{I}_n, \tag{6}$$

$$\dot{\mathbf{v}} + \mathbf{B}_x^\top\mathbf{v} - \frac{1}{a}\mathbf{P}\mathbf{B}_u\mathbf{B}_u^\top\mathbf{v} = \mathbf{0}, \quad \mathbf{v}(T) = \mathbf{x}^*. \tag{7}$$

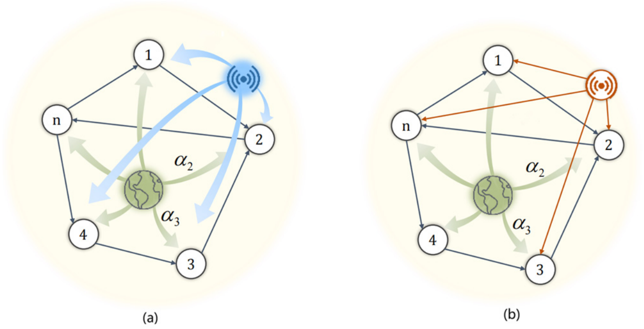

Fig 1. Diagram of direct intervention control and implicit intervention control

3. Numerical Simulations and Analysis of Core Topological Dynamics

We performed ensemble simulations on Erdős–Rényi (ER)[12], Barabási–Albert (BA)[13], and Watts–Strogatz (WS)[14] networks (30 realizations each, $n \in [50,80]$). Default parameters: $T = 10$, $a = 1.0$, $\gamma = 0.8$, $\mathbf{x}(0) \sim \mathcal{U}[-1,1]^n$, $\mathbf{x}^* = \mathbf{1}_n$, and p varied from 5% to 90%. Performance is measured by terminal error e_T, control energy E, and composite cost $J^* = e_T + aE$. Consequently, the absolute tracking error attenuation enabled by a given sparse strategy is expressed as $\Delta e_T = e_T^{noctrl} - e_T$. To quantify the control performance relative to energy investment, we define the relative energy efficiency metric η as follows:

$$\eta = \frac{\Delta e_T}{e_T^{noctrl} - e_T^{full}} \cdot \frac{E_{full}}{E}, \tag{8}$$

which parameterizes the exact scale of macroscopic tracking error contraction achieved per unit of microscopic continuous control energy consumption. To rigorously ascertain the specific structural advantage harvested by a centrality-driven node prioritization over a blind actuation policy, we introduce the normalized energy efficiency ratio η_{ratio}:

$$\eta_{ratio} = \frac{\eta}{\eta_{random}}, \tag{9}$$

where η_{random} represents the localized ensemble mean efficiency of the random selection baseline evaluated under identical topological and budget boundaries. A critical mathematical threshold of $\eta_{ratio} > 1$ firmly indicates that the underlying structural centrality heuristic successfully exploits the graph topological spectrum to deliver enhanced closed-loop control gain over the unoptimized spatial baseline.

3.1. Topological Adaptation

In network control, the backward integration of the differential Riccati equation (Eq. 5) depends directly on the spectral properties of the control input operator $\mathbf{B}_u\mathbf{B}_u^\top$. Selecting nodes via different centralities alters the rank-deficient structure of $\mathbf{B}_u\mathbf{B}_u^\top$, which in turn shapes the transient trajectories of the node beliefs according to the underlying graph topology.

For ER networks, direct control is highly effective when guided by global centrality measures like Eigenvector or Closeness centrality, which identify the structural core of the graph. However, under latent control, the performance of Betweenness centrality decreases. This suggests that indirect belief synchronization is more sensitive to global stationary distributions than localized bridge nodes. Conversely, in WS networks characterized by high local clustering, the K-shell decomposition identifies the highly connected core, making it the most energy-efficient strategy for direct intervention.

For Erdős-Rényi (ER) random networks, direct control is highly effective when guided by global centralities, such as Eigenvector, Closeness, and Betweenness, which naturally identify the structural centers of the graph. Under the latent control paradigm, however, the performance of Betweenness centrality decreases. Conversely, Eigenvector, Closeness, and PageRank centralities still yield satisfactory results, suggesting that indirect belief synchronization relies more on steady-state distributions than on specific path bottlenecks. We use these ER results as a baseline to study the phase transitions in subsequent sections.

In Barabási-Albert (BA) scale-free networks, the performance of different centrality metrics is almost indistinguishable, with no single strategy showing a clear statistical advantage (Fig. 2a,b). This is because the densely connected hub structure balances the topological influence across different node levels. When switching to latent intervention, the performance ranking of these centralities remains consistent, though the overall control cost increases. This upward shift in cost suggests that the latent filtering process in scale-free topologies leads to greater energy dissipation, while the relative efficacy of the control indices remains unaffected.

In Watts-Strogatz (WS) small-world networks, which feature high clustering and short path lengths, K-shell centrality achieves the minimum control cost under direct intervention (Fig. 2c,d). This strategy effectively targets the dense core of the network while ignoring peripheral nodes. Under latent intervention, this structural advantage becomes even more pronounced. Because the coupling parameter \gamma scales down the influence of long-range shortcuts, the core nodes must rely on their tightly clustered local neighborhoods to propagate beliefs, making K-shell control significantly more efficient than global metrics in the presence of suppressed long-range paths.

3.2. Robust Boundaries of Sparse Optimal Control under Extreme Resource Constraints

This subsection examines the framework's behavior under extreme sparsity $p \to 0$. At $p = 0.05$, the Random strategy fails to break spectral cut-set constraints, yielding very low efficiency η.

In BA networks, global metrics (Betweenness, Eigenvector) outperform Degree: Betweenness targets geodesic bottlenecks for cross-community reach, while Eigenvector aligns with the dominant Laplacian mode, achieving high modal traction with minimal energy. Thus, "global strategic location” matters more than local degree.

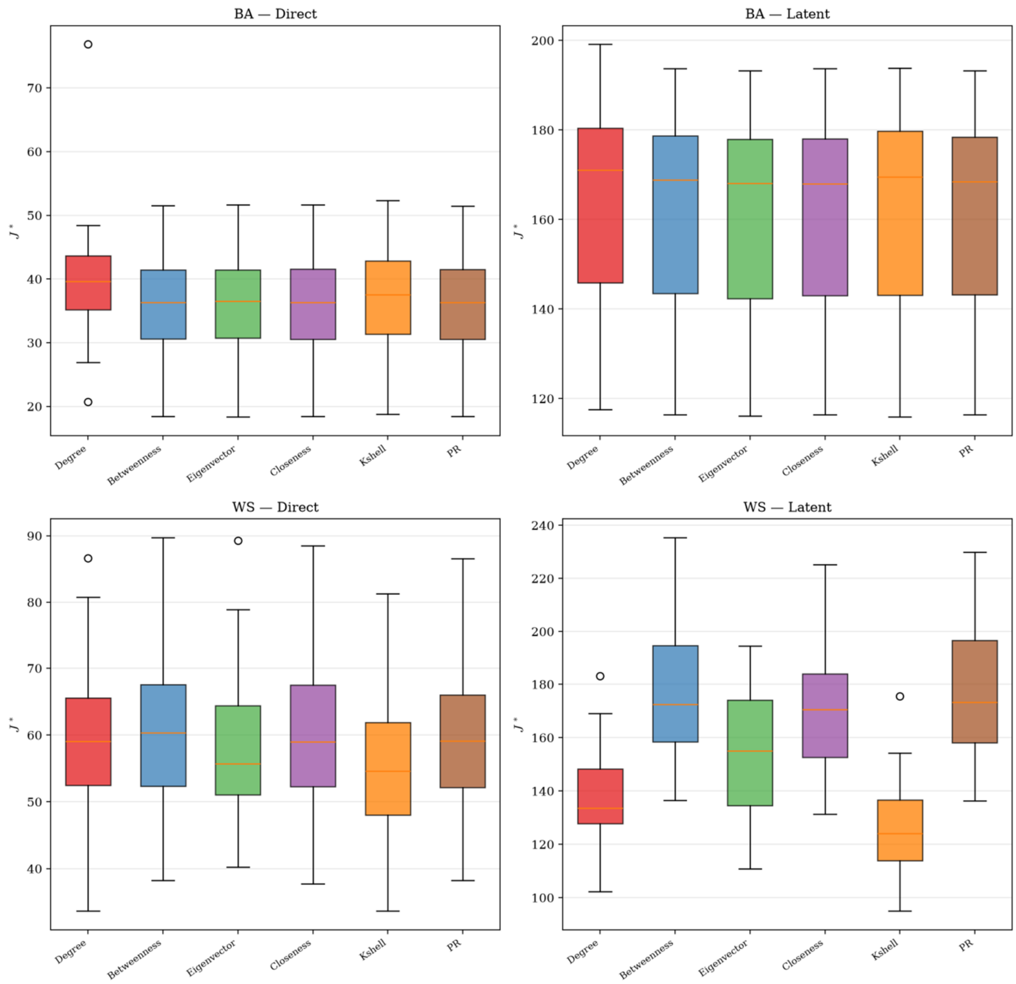

Fig 2. Comparison of optimal control cost across centrality strategies under different network topologies and intervention modes

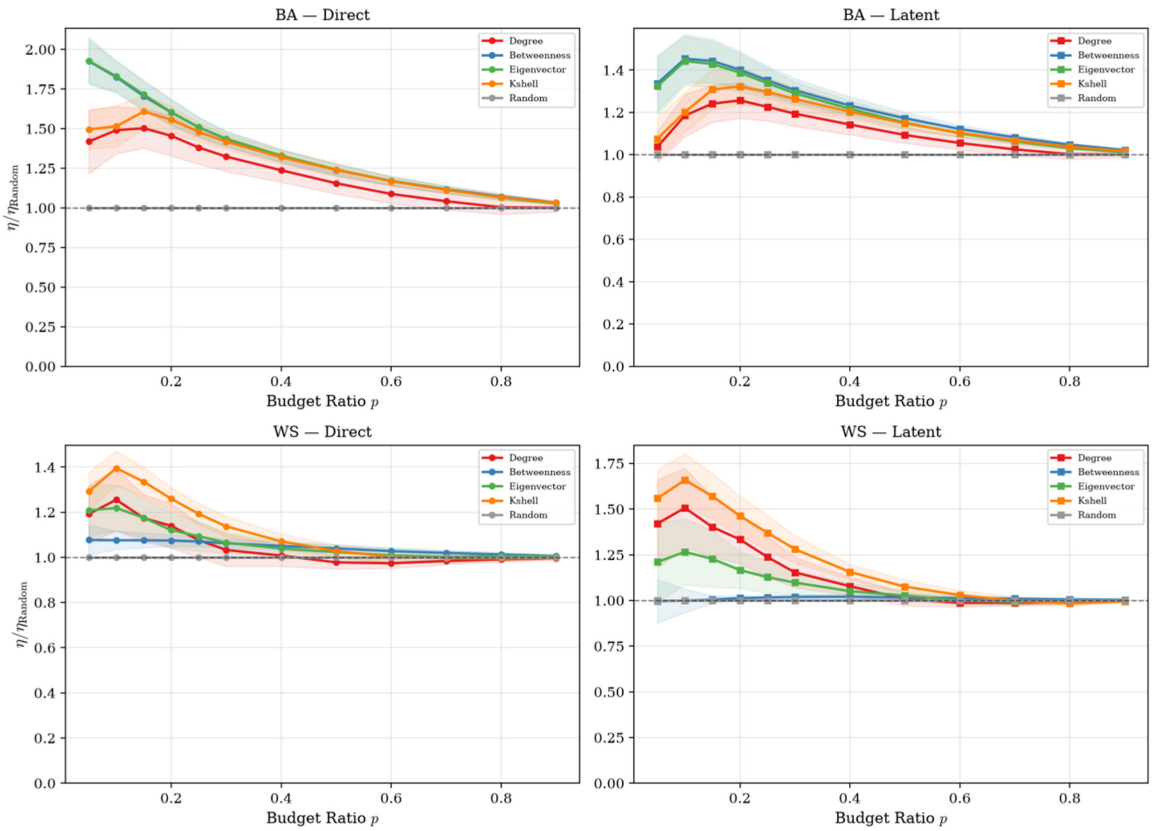

Fig 3. Changes in η_{ratio} of each strategy under different controlled ratios

In WS networks, K-shell dominates by isolating the dense core, which exploits high clustering to rapidly synchronize neighborhoods and broadcast via shortcuts—even though core nodes may lack the highest degrees.

As p increases, all curves converge toward $\eta_{ratio} = 1.0$. Fig. 3 validates these trends: for BA, Betweenness and Eigenvector peak at $p = 0.05$ and maintain a lead up to $p \leq 0.20$; for WS, K-shell remains highest throughout $p \leq 0.3$. The optimal centrality choice thus shifts deterministically with topology: global metrics for scale-free, core-peeling for small-world. Random selection always degrades, confirming that structure-aware node selection is irreplaceable under resource scarcity.

3.3. Phase Transition with Latent Coupling γ

Within the mathematical architecture of the latent intervention modality, the infiltration coupling parameter $\gamma \in (0,1)$ exerts a sophisticated dual-regulatory role on the system's global operator spectrum. On one hand, it defines the intensive scaling of external optimized control energy penetrating the network coordinates; on the other hand, it acts as a scaling multiplier that systematically damps the autonomous, unperturbed belief interactions within the social system. When the coupling strength resides in the low-to-moderate regime ($\gamma \in [0.1,0.7]$), the autonomous drift operator $\mathbf{B}_x = (1-\gamma)\mathbf{A} - \mathbf{I}_n$ retains a high degree of topological conductivity. Under these operational boundaries, high-performance global centrality metrics (e.g., PageRank, Betweenness) successfully propagate regulatory signals across multi-level geodesic pathways, maintaining their normalized energy efficiencies securely above the spatial benchmark ($\eta_{ratio} > 1$).

Conversely, when the system crosses into the extreme coupling limit ($\gamma \to 0.9$), the autonomous network drift term is severely compressed to merely $0.1\mathbf{A}$. Under this condition, the network topology undergoes a near-complete physical extinction of its long-range path conductivity, causing the dynamical system to behaviorally decouple into localized, non-homogeneous dissipation fields. In this highly clustered state under restrictive sparse allocation budgets ($p \leq 0.30$), blindly deploying traditional metrics that rely on global spectral coordinates (e.g., Eigenvector centrality) causes the continuous control inputs to become trapped in spatial localized modes, triggering extreme control energy overloads and local convergence annihilation. Consequently, the relative cost-effectiveness of these sophisticated structural metrics falls precipitously below that of a spatially uniform random allocation policy ($\eta_{ratio} < 1$). It is only when the sparse budget allocation $p > 0.30$ sufficiently expands the spatial actuation footprint that these global structural metrics revive their capacity to capture macroscopic topological configurations. To empirically characterize this complex phase-transition modulation, Fig.4 systematically tracks the continuous variations of η_{ratio} across the sparse budget spectrum $p \in [0.05,0.90]$ under latent intervention within the baseline ER random network topology, utilizing three distinct subfigures corresponding to $\gamma = 0.3$, 0.7, and 0.9:

1) Weak-to-Moderate Coupling Regime (Fig.4(a) and Fig.4(b)): At a low coupling intensity of $\gamma = 0.3$, all centrality metrics firmly maintain an efficiency profile of $\eta_{ratio} > 1$, showing a steady monotonic expansion as p scales upward—marking a stable structural optimization zone. As the parameter scales up to $\gamma = 0.7$, local structural metrics like Degree centrality begin to collapse toward the unoptimized spatial baseline, demonstrating that enhanced latent coupling forces an initial erosion of distance-dependent long-range strategies.

2) Extreme Coupling and Efficiency Inversion Phase Transition (Fig.4(c)): At the extreme limit of $\gamma = 0.9$, a profound, nonlinear ``efficiency inversion'' phase transition is exposed. Within the restrictive budget interval of $p \in [0.10,0.40]$, the efficiency ratios η_{ratio} of both Eigenvector and Betweenness centralities undergo a sharp inversion drop below the critical unit threshold ($\eta_{ratio} < 1$), mapping into a clear sub-optimal efficiency regime where the blind Random strategy demonstrates counter-intuitive superiority. It is only when the node coverage scales beyond the threshold boundary of $p > 0.50$ that these global spectral metrics gradually recover their large-scale coordination benefits.

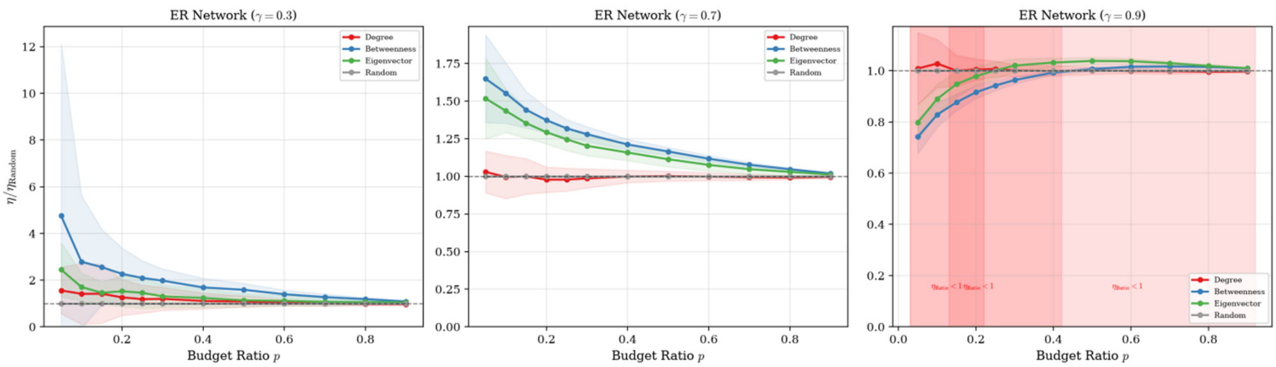

Fig 4. Changes of η_{ratio} with p for different γ values in the ER network under implicit intervention

3.4. Topological Preference Phase Transition of Sparse Control Law Driven by Cost Penalty Coefficient a

This subsection systematically investigates the profound topological preference phase transitions and the resultant strategy efficiency inversion laws driven by the continuous scaling of the isotropic control energy consumption penalty coefficient a.

1) Global Coordination Regime (Small a Bound): When external regulatory energy is extremely inexpensive ($a \to 0$), the optimization priority of the high-dimensional linear-quadratic system is overwhelmingly dominated by the minimization of macroscopic tracking errors. Under this operational regime, the controller possesses an abundant budget to actively force state convergence. Consequently, global structure-aware metrics such as Betweenness, Closeness, and Eigenvector centralities successfully construct a globally coordinated feedback control flow. These metrics leverage cheap, unconstrained signal propagation to permeate the entire topological state space, thereby delivering superior optimized composite cost J^* performance. In sharp contrast, the localized Degree centrality blindly pours massive, unoptimized control energy into immediate local hub nodes. This structural configuration triggers severe local modal saturation and energy dissipation waste along the chosen backbone channel, causing its corresponding cost trajectory to degrade, even falling below the unoptimized Random baseline.

2) Local High-Yield Regime (Large a Bound): Conversely, when the injected regulatory energy becomes prohibitively expensive ($a \to \infty$), the time-varying optimal feedback gain operator $-\frac{1}{a}\mathbf{B}_u^\top \mathbf{P}(t)$ dictated by the HJB theory asymptotically vanishes. As a result, the multi-agent system spontaneously enters an extreme, energy-saving conservative mode. Under this restrictive boundary, the long-range cross-community propagation signals relied upon by global spectral strategies are so severely attenuated that they undergo complete structural damping before navigating through complex network paths to reach peripheral agents (inducing systemic structural failure). Consequently, the J^* performance curves of global metrics collapse over a wide spectral range. In this energy-scarce regime, the Degree strategy exhibits unmatched robustness. Because its priority-selected hub individuals need only jump a singular topological step (1-hop) to directly actuate and force the maximum number of immediate neighbor states, this localized policy minimizes non-homogeneous topological attenuation losses. Therefore, under extreme high-cost constraints, Degree centrality becomes the singular robust topological heuristic capable of systematically outperforming the random spatial allocation baseline, precisely routing each costly unit of external energy into the highest-yield topological nodes.

To empirically validate the existence of this topology-preference phase transition, Fig.5 illustrates the continuous trajectories of the optimized composite cost J^* mapped against a wide logarithmic spectrum of the cost penalty $a \in [1,500]$ within the benchmark ER random network topology under a fixed sparse actuation budget of $p = 0.20$:

1) Accuracy-Priority Phase ($a < a_c$): When the cost penalty is low, the system occupies the Global Coordination Regime (graphically demarcated by the light blue shaded region). Here, the Eigenvector and Betweenness trajectories converge at the bottom of the cost spectrum, significantly outperforming the Degree strategy and the Random baseline.

2) Energy-Saving-Priority Phase ($a > a_c$): As the cost parameter intercepts and crosses a critical mathematical threshold $a_c \approx 20.0$ (explicitly highlighted by the vertical black dashed line), a distinct structural efficiency inversion occurs. The Degree strategy successfully overtakes all global spectral metrics, demonstrating compact, low-cost scaling and maintaining an absolute leader state across the expanding Local Efficiency Regime (graphically demarcated by the light red shaded region).

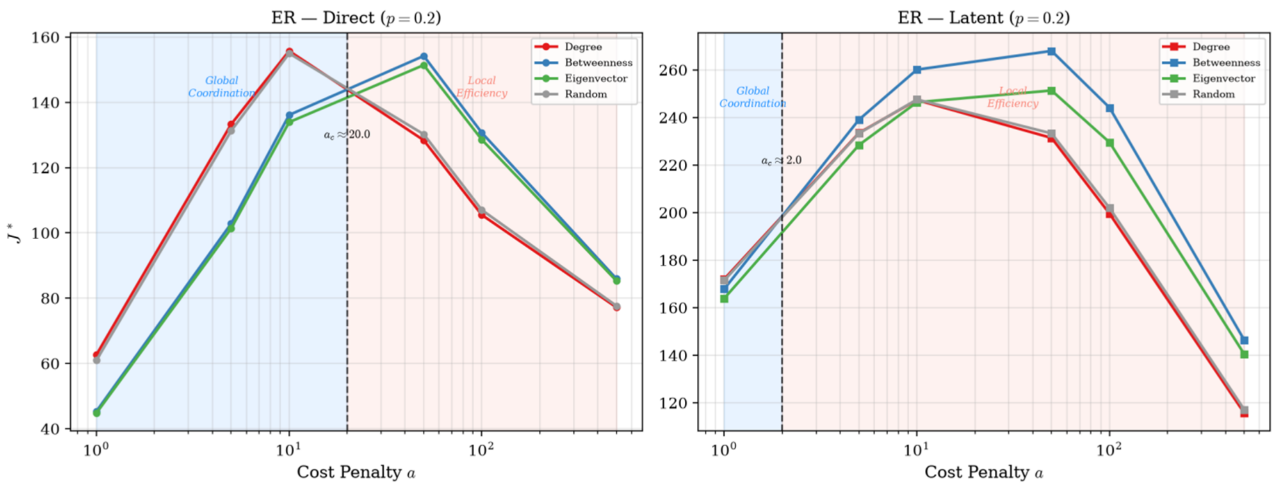

Fig 5. Changes in each strategy J^* under different cost penalties a

4. Conclusion and Discussion

This paper presents a sparse optimal control framework driven by network centralities to regulate collective beliefs under limited control resources. Integrating both direct and latent intervention mechanisms within the Network Drift-Diffusion Model (NDDM), we formulate the belief dynamics in a unified state-space representation. We then derive an optimal feedback control law based on the Hamilton-Jacobi-Bellman (HJB) equation, where a sparse selection matrix restricts direct actuation to only a small fraction of influential nodes. This approach successfully achieves global belief tracking and consensus coordination while minimizing the number of active control channels.

Numerical simulations on Erdős-Rényi (ER), Barabási-Albert (BA), and Watts-Strogatz (WS) networks demonstrate that our framework ensures stable error convergence and high energy efficiency. The results show that no single centrality metric is universally optimal. Instead, the best node selection strategy depends heavily on the interaction between network topology (such as degree distribution, clustering, and core-periphery structures) and control parameters (including the coupling coefficient γ and the control penalty a). Particularly under strict budget constraints ($p \leq 0.3$), centrality-based node selection offers distinct performance gains over random allocation baselines, providing a systematic approach for targeted network intervention.

Despite these insights, several limitations remain. First, our analytical model assumes a static network topology, thereby neglecting real-world dynamic factors such as time-varying edge rewiring, transmission delays, and non-linear opinion updates. Second, our node selection relies on predefined, heuristic centrality metrics, which may not capture the mathematically optimal control configurations in massive, complex online social systems.

Future research will extend this sparse control framework to time-varying graphs and multilayer networks. We plan to explore data-driven, model-free methods, such as Graph Neural Networks (GNNs) and Deep Reinforcement Learning (DRL), to adaptively identify optimal control nodes in dynamic environments. Additionally, validating these models on empirical social media datasets will be crucial to bridging the gap between theoretical network control and practical applications.

References

[1] Sui, D., Liu, Y., & Leng, S. (2026). Socially inspired adaptive framework for distributed online inference. *IEEE Transactions on Network Science and Engineering*, *13*, 5924–5942. https://doi.org/10.1109/TNSE.2026.3653651.

[2] Jadbabaie, A., Molavi, P., Sandroni, A., & Tahbaz-Salehi, A. (2012). Non-Bayesian social learning. *Games and Economic Behavior*, *76*(1), 210–225.

[3] Shao, C., Ciampaglia, G. L., Varol, O., Yang, K. C., Flammini, A., & Menczer, F. (2018). The spread of low-credibility content by social bots. *Nature Communications*, *9*(1), 4787.

[4] Johnson, N. F., Velásquez, N., Restrepo, N. J., Leahy, R., Gabriel, N., El Oud, S., Zheng, M., Manrique, P., Wuchty, S., & Lupu, Y. (2020). The online competition between pro- and anti-vaccination views.*Nature*, *582*(7811), 230–233.

[5] Jiang, H., Mazalov, V. V., Gao, H., & Wang, C. (2023). Opinion dynamics control in a social network with a communication structure. *Dynamic Games and Applications*, *13*(1), 412–434.

[6] Etesami, S. R. (2021). Open-loop equilibrium strategies for dynamic influence maximization game over social networks. *IEEE Control Systems Letters*, *6*, 1496–1500.

[7] Wang, C., Mazalov, V. V., & Gao, H. (2021). Opinion dynamics control and consensus in a social network. *Automation and Remote Control*, *82*(6), 1107–1117.

[8] Zhang, J., & Luo, Y. (2017). Degree centrality, betweenness centrality, and closeness centrality in social network. In *2017 2nd International Conference on Modelling, Simulation and Applied Mathematics (MSAM2017)* (pp. 315–318). Atlantis Press. https://doi.org/10.2991/msam-17.2017.68.

[9] Ruhnau, B. (2000). Eigenvector-centrality — A node-centrality? *Social Networks*, *22*(4), 357–365. https://doi.org/ 10.1016/ S0378-8733(00)00031-9.

[10] Zhang, P., Wang, T., & Yan, J. (2022). PageRank centrality and algorithms for weighted, directed networks. *Physica A: Statistical Mechanics and its Applications*, *586*, 126438. https: // doi.org/10.1016/j.physa.2021.126438.

[11] Freeman, L. C. (1978). Centrality in social networks: Conceptual clarification. *Social Networks*, *1*(3), 215–239. https://doi.org/10.1016/0378-8733(78)90021-7.

[12] Erdős, L., Knowles, A., Yau, H. T., & Yin, J. (2013). Spectral statistics of Erdős–Rényi graphs I: Local semicircle law.*The Annals of Probability*, *41*(3B), 2279–2345. https://doi.org/10. 1214/ 11-AOP734.

[13] Zadorozhnyi, V. N., & Yudin, E. B. (2012). Structural properties of the scale-free Barabasi-Albert graph. *Automation and Remote Control*, *73*(4), 702–716. https://doi. org/10. 1134/ S0005117912040091.

[14] Chen, Y. W., Zhang, L. F., & Huang, J. P. (2007). The Watts–Strogatz network model developed by including degree distribution: Theory and computer simulation. *Journal of Physics A: Mathematical and Theoretical*, *40*(29), 8237–8246. https://doi.org/10.1088/1751-8113/40/29/003.