Coarse targeting in social networks

Wei Li[1] and Philip Solimine[1]

[1] *The University of British Columbia*

April 7, 2025

Extended abstract

Abstract

We study how a planner can optimally counter misinformation in a social network under *coarse targeting*: she broadcasts the same message to all agents, but chooses their exposure levels. Before messaging begins, the planner chooses a vector of target weights that determine how much each agent is exposed to her message, in order to maximize total discounted utility. Optimal targeting depends jointly on the network structure and the distribution of initial opinions. Counterintuitively, agents with extreme views may sometimes receive *less* exposure. In stylized opinion-leader networks, optimal weights align with authority centrality. But centrality alone is also not sufficient: in symmetric networks, targeting is uniform only when initial opinions are. More generally, optimal weights reflect persistent local opinion dispersion among subsets of agents, which slows opinion convergence. We illustrate the model using U.S. Facebook friendship data and climate change opinions: the planner over-targets Texas and under-targets California, despite their similar centrality—underscoring how local disagreements shape optimal targeting.

JEL: D85, D83, C63.

Keywords: Network-based targeting, Coarse targeting, Misinformation, Content moderation, Social network dynamics.

1 Introduction

Combating misinformation in social networks faces structural and institutional constraints. While individuals form opinions through peer interactions and outside messages, public actors such as health agencies or digital platforms often cannot tailor messages to individuals. Reputational concerns demand consistent messaging across regions (e.g., uniform CDC statements about vaccine efficacy), while privacy regulations like the EU's General Data Protection Regulation (GDPR) restrict the use of personal attributes for targeting. Political sensitivities and fears of overreach further limit interventions (Gorwa, 2019). Even private platforms have begun scaling back fine-grained targeting: Meta and Google have removed political ad targeting by attributes such as race, religion, or political affiliation due to legal and reputational risks.[1] At the same time, personalized advertising has grown more expensive, with rising costs and diminishing marginal returns.[2] These constraints push institutions and platforms toward *coarse targeting*—broadcasting a common message while varying exposure intensity across agents.

We thank Ying Gao, Ben Golub, Li, Hao, Matt Jackson, Jesse Perla, Mike Peters, Amin Rahimian, Marit Rehavi, Euncheol Shin, Kyungchul (Kevin) Song, Wing Suen, Xu Tan, Philip Ushchev, Yiqing Xing and Junjie Zhou for helpful comments.

Email: wei.li@ubc.ca; philip.solimine@ubc.ca.

[1] https://www.facebook.com/government-nonprofits/blog/preparing-for-upcoming-removal-of-certain-ad-targeting-options, Archived May 13, 2025

[2] https://www.axios.com/2021/10/07/higher-prices-weaker-targeting-push-companies-to-rethink-digital-ads, Archived May 13, 2025

We study a planner who operates under this coarse targeting constraint: she chooses the exposure level each agent receives by selecting a vector of target weights before broadcasting uniform messages. This constraint captures many real-world interventions. For example, the CDC allocates more media spend to regions with greater public health risk; platforms like Twitter and YouTube adjust the visibility of fact-checked content; and push notifications are tuned to the user engagement patterns (Pennycook et al., 2021). While message content is fixed, target weight—via algorithmic delivery, regional targeting, or media purchasing—is optimized. The planner maximizes her total discounted payoff and faces a dynamic trade-off: reduce short-run extremism or accelerate long-run convergence to a desired state.

In this model, a forward-looking planner who broadcasts a message per period and chooses a vector of *target weights* $\mathbf{b}^* \in \mathbb{R}^n$ ex ante, where b_i determines agent i's exposure. Agents update opinions by averaging their neighbors' opinions, given by the network adjacency matrix A, along with the planner's broadcast messages weighted by b_i. The planner's stage payoff is decreasing in the quadratic distance between agent opinions and her agenda, and she chooses $\mathbf{b}^*$ once to maximize her total discounted payoff. Theoretically, the planner *jointly designs and controls* the opinion dynamics: $\mathbf{b}^*$ endogenizes how future opinions evolve in response to a fixed message. This contrasts with continuous targeting (e.g. Galeotti et al. (2020), Li and Tan (2024) and discrete targeting such as minimum driver set problems Liu and Barabási (2016). The optimal $\mathbf{b}^*$ solves a fixed-point equation linking itself to all its future effects on agents' opinions. We show that $\mathbf{b}^*$ varies smoothly with initial opinions and, perhaps counterintuitively, may assign *less* weight to agents with more extreme initial opinions. In stylized opinion-leader networks, optimal targeting aligns with *authority centrality*—how much agents are listened to by others. Yet centrality alone is not sufficient: optimal weights also respond to persistent, localized disagreement. For example, when applied to U.S. state-level climate change beliefs, the model predicts more targeting of Texas and less of California, despite their similar centrality.

We begin with a static benchmark where the planner maximizes her next-period payoff. The vector of optimal target weights is proportional to the vector of counterfactual opinions; what agents would believe in the absence of any intervention. Network structure matters less because there is no opportunity for opinions to diffuse.

In the dynamic setting, targeting must account for how opinions interact over time. For any $\mathbf{b}$, the planner's value function takes a quadratic form in initial opinions, with coefficients governed by a Riccati matrix that captures the net future benefit of targeting given optimal messaging. To compute $\mathbf{b}^*$, we exploit the recursive structure of the value function to derive its gradient with respect to target weights, yielding a nonlinear fixed-point equation that characterizes the optimal targeting vector. This structure enables a nested fixed-point algorithm that outperforms projected gradient methods in both speed and scalability in simulations.

Two forces shape optimal targeting: network structure and initial opinions of the agents. In an empty network, where agents are isolated, targeting mirrors initial opinions: more extreme agents receive more weights. In a fully connected network, only the average opinion matters, so uniform targeting suffices. Generally, these forces interact in nontrivial ways.

We first show that the optimal target weights vary continuously with the initial opinions. In symmetric networks with rank greater than one—where all agents have equal centrality—it is optimal to target agents equally *if and only if* their initial opinions are identical. In this case, any initial disagreement disappears in one period, and equal targeting shifts the average opinion toward the planner's agenda as quickly as possible. The resulting payoff matches what the planner would achieve under fully personalized targeting, despite the coarse targeting constraint. Small perturbations to these initial conditions introduce disagreement, which persists and slows convergence. These dynamics are governed by the network's eigenstructure. Agents with extreme views may receive more weight if they are relatively stubborn and their opinions persist locally. But they may receive less weight if their neighbors follow their extreme opinions and propagate disagreement further. While the principal eigenvector determines the most effective direction to steer the average opinion, persistent local disagreement is shaped by the non-principal eigenvectors. The associated eigenvalues determine the rate at which such disagreement dissipates. Thus, optimal targeting is not only about steering the average, but also about reducing persistent local disagreements.

In asymmetric networks, structure plays a more active role in shaping optimal weights. Agents differ in *authority centrality*—how much others listen to them—and *hub centrality*—how much they listen to others (Kleinberg, 1999). There is no closed-form solution in general, but to isolate these forces, we analyze a stylized opinion-leader network in which each agent gives the same descending weights to others' opinions: highest to agent 1, then agent 2, and so on. In this structure, optimal target weights are independent of initial opinions and exactly aligned with agents' authority centrality. Convergence is driven most by agent 1, making him the most valuable target. However, dispersion in other directions still lowers the planner's short-run payoff. As a result, the planner assigns positive weight to lower-authority agents—not to steer the average opinion, but to accelerate the dissipation of disagreement. Even if agent 1 is listened to with probability one, targeting only him is suboptimal: optimal targeting also reflects the structure and persistence of local opinion distortions.

To illustrate our findings, we analyze climate change opinions from the Yale Climate Opinion Maps (Howe et al., 2015), shown in Figure 1, to examine how a planner might counter misinformation in the interstate Facebook friendship network (Bailey et al., 2018). We find that optimal targeting departs from centrality-based benchmarks, as seen in Figure 2. California, despite its high centrality, is targeted less than expected, while Texas receives substantially more weight than its centrality. Similarly, Michigan and South Carolina—states with relatively moderate opinions—are targeted well above what their authority centrality would suggest. In contrast, more extreme states like Utah and Idaho receive less weight. This pattern reflects a core feature of the model: the planner gains more by targeting moderately opinionated agents who are well connected to—and can influence—isolated pockets of disagreement. These agents serve as leverage points for accelerating convergence, even when their centrality is relatively low. A broadly actionable insight is that moderate-opinion clusters should be prioritized over misinformation hot spots if they serve as key "bridges" in the network.

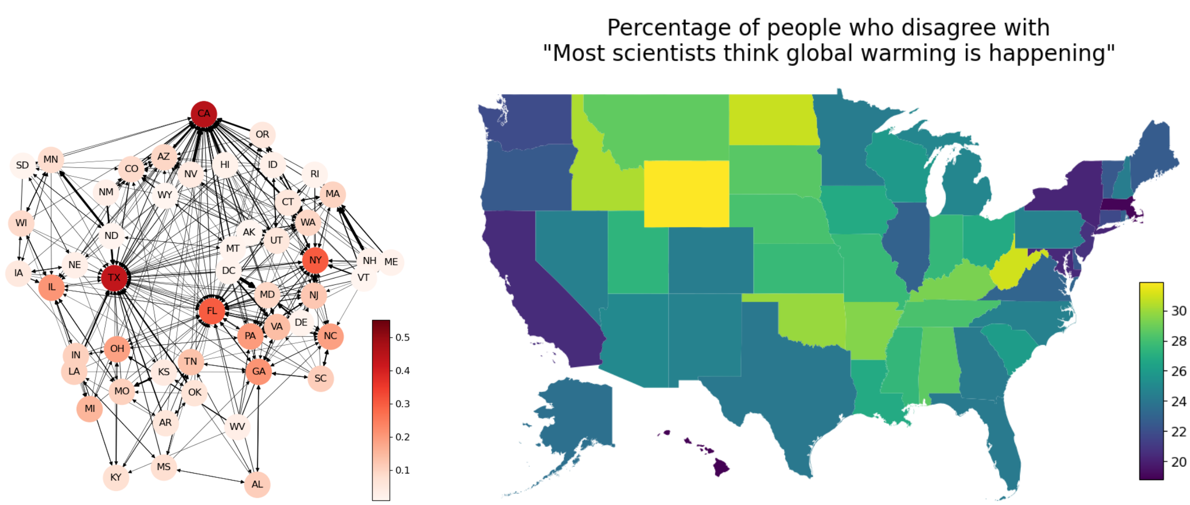

Figure 1: A measure of misinformation across in the Facebook network of US states.

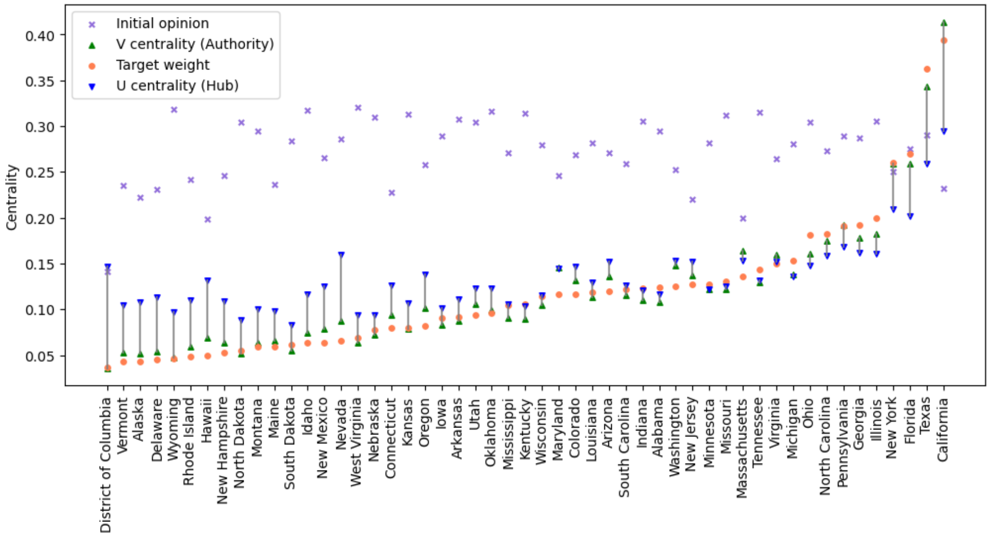

Figure 2: Facebook social network between states, and optimal target weights compared to centrality and opinions.

Related literature

A. Optimal intervention in networks. Galeotti and Goyal (2009) analyze a setting where a single strategic influencer, knowing only the distribution of agents' degrees, intervenes in a network. Galeotti et al. (2020) consider a static game in which a social planner optimally intervenes by changing agents' private returns to investment, which exhibit strategic spillovers in a network. Candogan et al. (2022) use a two-block model to examine platform influence and limiting beliefs, demonstrating how insights from small-network results generalize to larger, dense stochastic block networks. Our model is closely related to Li and Tan (2024) who study dynamic competition among strategic influencers with distinct agendas, but the influencers use continuous targeting. It is equivalent to our model where the optimal target weights are uniform, and she sends potentially different messages to every agent; that is, she can send many signals. In contrast, we focus on the design of an optimal vector of target weights, but the planner can only broadcast messages; that is, send a single signal. Bloch and Shabayek (2023) investigate optimal targeting when the planner lacks information on agents' specific network positions. Unlike these models, our framework assumes that the planner, despite knowing the network structure, cannot target individuals directly due to reputational or legal constraints. Several other papers at the intersection of economics and computer science consider planners with different intervention objectives. For instance, Gaitonde et al. (2020) consider a nefarious agent making a single perturbation to initial opinions that propagates through a Friedkin and Johnsen (1990) model of opinion dynamics to induce as much discord as possible.[3]

B. Control theory. Network-based targeting is often studied through identifying the minimum number of controlled nodes, or the set of *driver nodes*, needed to achieve certain outcomes in a network (Liu and Barabási, 2016; Gao et al., 2020).[4] However, these approaches contrast with platform-wide influence models like ours, where the influence

[3]Among many others, Acemoglu et al. (2024) show how a social media platform owner interested in maximizing engagement tend to design their algorithms to create more homophilic communication patterns ("filter bubbles"). Mostagir et al. (2022) consider a social learning model where agents learn about an underlying state of the world from individual observations as well as from exchanging information with each other, focusing on aggregate measures of vulnerability to misinformation.

[4]In the vast majority of these papers, it is assumed that there is a single specific target state that is exogenously determined; a reasonable assumption for many engineering applications. Notable departures from this include a particularly relevant study by Li et al. (2024), who depart from the assumption of a known target state by presuming that it is fixed but unknown and solving an inverse optimal control problem in a linear quadratic setting.

weights are continuous rather than discrete. A notable departure from driver set models is the work of Solimine and Meyer-Baese (2022), which optimizes interventions at an aggregate level rather than controlling individual nodes. While some studies incorporate continuous targeting as an intermediate step (see Pasqualetti et al. (2014); Summers et al. (2015); Klickstein et al. (2017), Gao et al. (2018) and Klickstein and Sorrentino (2023) among others), they typically aim to drive the network to a specific state under time constraints, a different objective from our planner maximizing her total discounted payoff. More closely related to us is Lindmark and Altafini (2020) who describes two centrality measures that quantify the importance of each node as a potential driver node, balancing a node's potential influence with the cost of control. In our model, asymmetric networks and the agents' different centralities also play an important role. Our planner chooses the optimal target weights, which control all agents with different intensities, with a focus on economic applications.

Our work is closely related to the literature on structural controllability. This concept was originally introduced by Lin (1974), and the idea is to establish a controllability condition that is agnostic to the weights in the input matrix. However, our work separates from this literature by optimizing the weights directly. Wang et al. (2012) discussed perturbations to the network topology itself that maximize measures of structural controllability, but Baggio et al. (2019) acknowledged that the input structure design problem has not received much attention in the literature. Some subsequent work by Baggio and Zampieri (2022); Baggio et al. (2022) has worked to address this problem by minimizing quadratic energy metrics, and Li (2024) who studies the intersection of structural controllability with optimal control.

C. Network learning and opinion dynamics. We assume agents learn naively and update their opinions according to the learning rule proposed by DeGroot (1974).[5] This assumption allows us to focus on the design of optimal influence itself. Due to the high cognitive burden of Bayesian learning for individuals, many recent papers explore quasi-Bayesian learning rules in which agents are boundedly rational.[6] However, even quasi-Bayesian learning is still cognitively and computationally demanding, as shown by Li and Tan (2020, 2021). Furthermore, agents in lab and field experiments often exhibit very limited cognitive ability.[7] Our work can be extended to account for other forms of bounded rational learning such as stubborn agents.

References

Acemoglu, Daron, Asuman Ozdaglar, and James Siderius (2024), "A model of online misinformation." *Review of Economic Studies*, 91, 3117–3150.

Alatas, Vivi, Abhijit Banerjee, Arun G. Chandrasekhar, Rema Hanna, and Benjamin A. Olken (2016), "Network structure and the aggregation of information: Theory and evidence from indonesia." *American Economic Review*, 106, 1663–1704.

Baggio, Giacomo, Fabio Pasqualetti, and Sandro Zampieri (2022), "Energy-aware controllability of complex networks." *Annual Review of Control, Robotics, and Autonomous Systems*, 5, 465–489.

Baggio, Giacomo and Sandro Zampieri (2022), "Reachable volume of large-scale linear network systems: The single-input case." In *2022 IEEE 61st Conference on Decision and Control (CDC)*, 4224–4229, IEEE.

[5]Variants of DeGroot learning were analyzed by Friedkin and Johnsen (1990), DeMarzo et al. (2003), Golub and Jackson (2010), and Ghaderi and Srikant (2014). See Chapter 7 of Jackson (2008).

[6]See Bala and Goyal (1998), Alatas et al. (2016), Molavi et al. (2018), Levy and Razin (2018), Mueller-Frank and Neri (2021), and Della Lena (2024), among others.

[7]See Enke and Zimmermann (2019), Chandrasekhar et al. (2020), and Grimm and Mengel (2020), among others.

Baggio, Giacomo, Sandro Zampieri, and Carsten W Scherer (2019), "Gramian optimization with input-power constraints." In *2019 IEEE 58th Conference on Decision and Control (CDC)*, 5686–5691, IEEE.

Bailey, Michael, Rachel Cao, Theresa Kuchler, Johannes Stroebel, and Arlene Wong (2018), "Social connectedness: Measurement, determinants, and effects." *Journal of Economic Perspectives*, 32, 259–280.

Bala, Venkatesh and Sanjeev Goyal (1998), "Learning from neighbours." *The Review of Economic Studies*, 65, 595–621.

Bloch, Francis and Shaden Shabayek (2023), "Targeting in social networks with anonymized information." *Games and Economic Behavior*, 141, 380–402.

Candogan, Ozan, Nicole Immorlica, Bar Light, and Jerry Anunrojwong (2022), "Social learning under platform influence: Consensus and persistent disagreement." *arXiv preprint arXiv:2202.12453*.

Chandrasekhar, Arun G., Horacio Larreguy, and Juan Pablo Xandri (2020), "Testing models of social learning on networks: Evidence from two experiments." *Econometrica*, 88, 1–32.

DeGroot, Morris H. (1974), "Reaching a consensus." *Journal of the American Statistical Association*, 69, 118–121.

Della Lena, Sebastiano (2024), "The spread of misinformation in networks with individual and social learning." *European Economic Review*, 168, 104804.

DeMarzo, Peter M., Dimitri Vayanos, and Jeffrey Zwiebel (2003), "Persuasion bias, social inffuence, and unidimensional opinions." *Quarterly Journal of Economics*, 118, 909–968.

Enke, Benjamin and Florian Zimmermann (2019), "Correlation neglect in belief formation." *The Review of Economic Studies*, 86, 313–332.

Friedkin, Noah E. and Eugene C. Johnsen (1990), "Social influence and opinions." *Journal of Mathematical Sociology*, 15, 193–206.

Gaitonde, Jason, Jon Kleinberg, and Eva Tardos (2020), "Adversarial perturbations of opinion dynamics in networks." In *Proceedings of the 21st ACM Conference on Economics and Computation*, 471–472.

Galeotti, Andrea, Benjamin Golub, and Sanjeev Goyal (2020), "Targeting interventions in networks." *Econometrica*, 88, 2445–2471.

Galeotti, Andrea and Sanjeev Goyal (2009), "Influencing the influencers: A theory of strategic diffusion." *The Rand Journal of Economics*, 40, 509–532.

Gao, Leitao, Guangshe Zhao, Guoqi Li, Lei Deng, and Fei Zeng (2018), "Towards the minimum-cost control of target nodes in directed networks with linear dynamics." *Journal of the Franklin Institute*, 355, 8141–8157.

Gao, Leitao, Guangshe Zhao, Guoqi Li, Fanghong Guo, and Fei Zeng (2020), "Optimal target control of complex networks with selectable inputs." *IEEE Transactions on Control of Network Systems*, 8, 212–221.

Ghaderi, Javad and Rayadurgam Srikant (2014), "Opinion dynamics in social networks with stubborn agents: Equilibrium and convergence rate." *Automatica*, 50, 3209–3215.

Golub, Benjamin and Matthew O. Jackson (2010), "Naive learning in social networks and the wisdom of crowds." *American Economic Journal: Microeconomics*, 2, 112–49.

Gorwa, Robert (2019), “Regulating them softly.” *Models for platform governance*, 39–43.

Grimm, Veronika and Friederike Mengel (2020), “Experiments on belief formation in networks.” *Journal of the European Economic Association*, 18, 49–82.

Howe, Peter D, Matto Mildenberger, Jennifer R Marlon, and Anthony Leiserowitz (2015), “Geographic variation in opinions on climate change at state and local scales in the usa.” *Nature climate change*, 5, 596–603.

Jackson, Matthew O. (2008), *Social and Economic Networks*. Princeton University Press.

Kleinberg, Jon M. (1999), “Authoritative sources in a hyperlinked environment.” *J. ACM*, 46, 604–632, URL `http://dblp.uni-trier.de/db/journals/jacm/jacm46.html#Kleinberg99`.

Klickstein, Isaac, Afroza Shirin, and Francesco Sorrentino (2017), “Energy scaling of targeted optimal control of complex networks.” *Nature Communications*, 8, 15145.

Klickstein, Isaac and Francesco Sorrentino (2023), “Selecting energy efficient inputs using graph structure.” *International Journal of Control*, 96, 987–999.

Levy, Gilat and Ronny Razin (2018), “Information diffusion in networks with the bayesian peer influence heuristic.” *Games and Economic Behavior*, 109, 262–270.

Li, Wei and Xu Tan (2020), “Locally bayesian learning in networks.” *Theoretical Economics*, 15, 239–278.

Li, Wei and Xu Tan (2021), “Cognitively-constrained learning from neighbors.” *Games and Economic Behavior*, 129, 32–54.

Li, Wei and Xu Tan (2024), “Dynamic network influence: the art of strategic messaging.” UBC working paper.

Li, Xiang (2024), *Structural controllability and optimal control of complex networks*. Ph.D. thesis, Nanyang Technological University.

Li, Yao, Chengpu Yu, Hao Fang, and Jie Chen (2024), “Inverse optimal control for linear quadratic tracking with unknown target states.” *arXiv preprint arXiv:2402.17247*.

Lin, Ching-Tai (1974), “Structural controllability.” *IEEE Transactions on Automatic Control*, 19, 201–208.

Lindmark, Gustav and Claudio Altafini (2020), “Centrality measures and the role of non-normality for network control energy reduction.” *IEEE Control Systems Letters*, 5, 1013–1018.

Liu, Yang-Yu and Albert-László Barabási (2016), “Control principles of complex systems.” *Reviews of Modern Physics*, 88, 035006.

Molavi, Pooya, Alireza Tahbaz-Salehi, and Ali Jadbabaie (2018), “A theory of non-bayesian social learning.” *Econometrica*, 86, 445–490.

Mostagir, Mohamed, Asuman Ozdaglar, and James Siderius (2022), “When is society susceptible to manipulation?” *Management Science*, 68, 7153–7175.

Mueller-Frank, Manuel and Claudia Neri (2021), “A general analysis of boundedly rational learning in social networks.” *Theoretical Economics*, 16, 317–357.

Pasqualetti, Fabio, Sandro Zampieri, and Francesco Bullo (2014), "Controllability metrics, limitations and algorithms for complex networks." *IEEE Transactions on Control of Network Systems*, 1, 40–52.

Pennycook, Gordon, Ziv Epstein, Mohsen Mosleh, Antonio A Arechar, Dean Eckles, and David G Rand (2021), "Shifting attention to accuracy can reduce misinformation online." *Nature*, 592, 590–595.

Solimine, Philip and Anke Meyer-Baese (2022), "Input design for the optimal control of networked moments." In *2022 IEEE 61st Conference on Decision and Control (CDC)*, 5894–5901.

Summers, Tyler H, Fabrizio L Cortesi, and John Lygeros (2015), "On submodularity and controllability in complex dynamical networks." *IEEE Transactions on Control of Network Systems*, 3, 91–101.

Wang, Wen-Xu, Xuan Ni, Ying-Cheng Lai, and Celso Grebogi (2012), "Optimizing controllability of complex networks by minimum structural perturbations." *Physical Review E*, 85, 026115.