This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

Navigating the Lobbying Landscape: Insights From Opinion Dynamics Models

Daniele Giachini, Leonardo Ciambezi, Verdiana Del Rosso, Fabrizio Fornari, *Member, IEEE*, Valentina Pansanella, Lilit Popoyan, and Alina Sîrbu

***Abstract*—While lobbying has been demonstrated to have an important effect on public opinion and policy making, existing models of opinion formation do not specifically include its effect. In this work, we introduce a new model of lobbying-driven opinion influence within opinion dynamics, where lobbyists can implement complex strategies and are characterized by a finite budget. Individuals update their opinions through a learning process resembling Bayes-rule updating but using signals generated by the other agents (a form of social learning), modulated by under-reaction and confirmation bias. We study the model theoretically and numerically, demonstrating rich dynamics both with and without lobbyists. In the presence of lobbying, we observe two regimes: one in which lobbyists can have full influence on the agent network, and another where the peer-effect generates polarization. When lobbyists are symmetric, the lobbyist-influence regime is characterized by prolonged opinion oscillations. If lobbyists temporally differentiate their strategies, frontloading is advantageous in the peer-effect regime, whereas backloading is advantageous in the lobbyist-influence regime. These rich dynamics pave the way for studying real lobbying strategies to validate the model in practice.**

***Index Terms*—Agent-based modeling, Bayesian learning and behavioral bias, climate change, lobbying, opinion dynamics, social networks.**

Received 26 December 2025; revised 6 March 2026; accepted 5 April 2026. This work was supported in part by the European Union—Next-GenerationEU—National Recovery and Resilience Plan (NRRP)—MISSION 4 COMPONENT 2, INVESTMENT N. 1.1, CALL PRIN 2022 PNRR D.D. 1409 14-09-2022—ALMONDO Project, CUP N. J53D23015400001. *(Corresponding author: Daniele Giachini.)*

Daniele Giachini and Leonardo Ciambezi are with the Institute of Economics and Department L'EMbeDS, Sant'Anna School of Advanced Studies, 56127 Pisa, Italy (e-mail: Daniele.Giachini@santannapisa.it; leonardo.ciambezi@santannapisa.it).

Verdiana Del Rosso and Fabrizio Fornari are with the Computer Science Division, School of Science and Technology, University of Camerino, 62032 Camerino, Italy (e-mail: verdiana.delrosso@unicam.it; fabrizio.fornari@unicam.it).

Valentina Pansanella is with the Italian National Research Council (CNR), Institute of Information Science and Technologies "A. Faedo" (ISTI), 56124 Pisa, Italy (e-mail: valentina.pansanella@isti.cnr.it).

Lilit Popoyan is with the School of Business and Management, Queen Mary University of London, E14NS London, U.K. (e-mail: l.popoyan@qmul.ac.uk).

Alina Sîrbu is with the Department of Computer Science and Engineering, University of Bologna, 40126 Bologna, Italy (e-mail: alina.sirbu@unibo.it).

Digital Object Identifier 10.1109/TCSS.2026.3683772

I. Introduction

LOBBYING is a pervasive feature of modern policymaking, shaping decisions across virtually all domains of public life, from environmental regulation and healthcare to financial oversight and digital governance. Beyond direct engagement with policymakers, lobbying strategies increasingly aim to influence public opinion, leveraging social networks and mass communication to create political momentum, reduce resistance to controversial policies, or shift the Overton window on key issues. Lobbyists, whether representing corporate interests, civil society, or ideological coalitions, act strategically to influence not only policymakers but also public perception and discourse. A large body of empirical literature has demonstrated the effectiveness of lobbying in shaping policy outcomes [1], [2], with recent work underscoring the growing sophistication of lobbying strategies in networked political and media environments [3], [4]. This broader landscape of influence highlights the need to understand lobbying not only as a transactional activity but also as a dynamic process of belief formation and diffusion in complex social systems.

Among the various mechanisms through which lobbying exerts influence, opinion formation within the public sphere plays a critical role. Lobbyists increasingly target public attitudes to create political pressure or legitimacy for specific policy choices [5], [6]. The advent of digital platforms has magnified these efforts, enabling tailored messaging and decentralized opinion shaping. This makes the diffusion of opinions across social networks—where individuals interact, update beliefs, and potentially become advocates themselves—a crucial site for understanding lobbying efficacy.

In this article, we focus on the opinion-formation layer of lobbying. In particular, we examine how strategic actors shape public beliefs over time through advocacy, agenda setting, and framing, and how grassroots/"outside" mobilization complements "inside" access strategies. These channels are well documented in the political-economy and advocacy literatures and often operate precisely by shifting public attitudes to create political momentum or legitimacy for policy change. Our model abstracts from the policy decision stage and instead endogenizes lobbying-driven opinion influence as costly, time-staged communication under resource constraints. This scope is consistent with empirical evidence that frames and advocacy can materially affect public support and elite incentives [1], [2], [3], [5], [6].

This article contributes to this growing interdisciplinary literature by introducing a novel model of lobbying influence over opinion dynamics in a boundedly rational social network. In our framework, a population of agents interacts repeatedly to update their subjective beliefs about the likelihood of an uncertain future event. Each agent weighs two competing probabilistic models (one "optimistic" and one "pessimistic") and updates beliefs through a process resembling Bayes-rule updating but using signals generated by other agents (an extreme form of social learning [7], [8], [9], [10]), modulated by under-reaction [11], [12], [13], and confirmation bias (a form of directional motivated reasoning) [14]. Agents are embedded in a directed network and interact through sequential signal exchanges that propagate beliefs over time.

We then endogenize lobbying by introducing external agents (lobbyists) who strategically send costly signals to influence the population's belief distribution by the end of a finite time horizon. Lobbyists differ in the model they support and are subject to a resource constraint (budget). Their objective is to minimize the distance between the final average belief and their preferred model. Lobbyists select randomized strategies from a constrained set of signaling paths, and their interventions shape both the speed and direction of belief evolution in the network.

Our theoretical and numerical results reveal rich dynamics. First, we theoretically prove that in the absence of lobbying and considering ex-ante identical agents, the model naturally shows path-dependence, with asymptotic consensus either on the pessimist or on the optimist model. The numerical experiments in such a framework extend such a result, showing that, adding initial heterogeneity in opinions, the network tends to converge toward consensus or, under high confirmation bias, become polarized. Combining theoretical and numerical results, we find that a single lobbyist with sufficient resources can dominate belief formation and steer the network toward its preferred model, especially when agents are relatively open to new information. However, when two lobbyists with opposing goals act simultaneously, the system often fails to converge within the active time horizon, resulting in persistent instability and oscillation, especially in regions of the parameter space where both lobbyists are highly effective. A similar conclusion emerges from the theoretical analysis of a special case, where the unique Nash equilibrium arising from the strategic interaction of two lobbyists supporting opposing models involves uniform mixed strategies. We examine the role of timing in lobbying strategies, showing that frontloading or backloading (i.e., concentrating signals at the beginning or end of the interaction window) can be advantageous depending on the population's biases. Finally, we investigate the role of different network configurations and show that, in sparser networks, a lobbyist has greater influence on the emerging opinion dynamics. Overall, our results indicate a complex interplay between cognitive constraints and strategic influence.

The remainder of the article is structured as follows. Section II reviews related work. Section III introduces the baseline opinion dynamics model and the modeling of lobbying agents. Section IV presents the theoretical and simulation results across different scenarios and associated sensitivity analyses. Section V concludes and outlines future research directions.

II. Related Work

Opinion dynamics models offer mathematical frameworks for understanding how beliefs spread and evolve across populations [15], [16]. Rooted in theories of social influence and information diffusion, and inspired by models from statistical physics, they have been applied in sociology [17], political science [18], marketing [19], and economics [20]. These models help elucidate the emergence of collective behaviors, consensus formation, and information propagation patterns in networks [21], [22], [23]. Recent extensions study technology-mediated settings, such as bounded confidence models with algorithmic bias [24], [25], which demonstrate how biased digital interactions can foster fragmentation and polarization.

In economic modeling and political science, opinion dynamics has increasingly intersected with agent-based modeling approaches [26], [27], [28], enabling the incorporation of bounded rationality, motivated reasoning, and strategic behavior. Related efforts examine media influence, typically modeled as an external field affecting agents with some probability [21], [22], [23], [25], [29]. Depending on the structure of the population and model parameters, media exposure can accelerate consensus or foster polarization, especially when signals are extreme. Other works include zealots or stubborn agents [23], [30], whose persistent behavior can shape long-run opinion distributions.

Despite these advances, the integration of lobbying into opinion dynamics models remains limited. Existing approaches rarely account for lobbying as a strategic, resource-constrained, time-staged intervention.

Furthermore, these models often rely on simplified opinion-updating rules that do not explicitly incorporate cognitive biases such as confirmation bias or under-reaction. Recent work has shown that confirmation bias alone can generate persistent polarization even under exact Bayesian updating and may also induce oscillatory opinion dynamics under specific informational conditions [31]. As a result, the endogenous modeling of lobbying strategies aimed at belief manipulation within dynamic opinion formation is still underdeveloped.

Our work is also connected to research on opinion dynamics with stubborn or exogenous influencers, such as the framework of [32], which analyzes convergence and fluctuations under linear averaging with fixed-opinion agents. However, our setting differs in several important respects. In contrast to models with constant linear weights, individual responsiveness in our framework is both signal- and state-dependent, shaped by under-reaction and confirmation bias. Moreover, external influence is not exogenously imposed but generated by budget-constrained lobbyists who send costly, time-staged signals that interact with peer-to-peer communication. These elements combine to produce regime patterns—such as persistent oscillations under symmetric lobbying—that do not arise in linear stubborn-agent models.

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

A related strand models persistent external influence via zealots or stubborn agents [33], [34], sometimes motivated by media or lobbying contexts and, in some variants, incorporating constraints on influence intensity. We see these contributions as important reduced-form benchmarks. Our setting differs, however, in three structural respects that are central of the dynamics studied here. First, influence is generated by budget-constrained lobbyists that allocate costly signals over time and across targeted agents, rather than by permanently fixed-opinion nodes. Second, lobbying signals enter the same nonlinear, state- and signal-dependent updating channel governing peer communication (via under-reaction and directional motivated reasoning), whereas many zealot frameworks rely on linear averaging with constant weights or state-independent adoption rules. Third, by explicitly modeling time-staged interventions (finite horizon and timing profiles), our framework allows a systematic analysis of strategy timing (e.g., frontloading and backloading) and of regime behavior (lobbyist-dominance versus peer-effect regions), including oscillatory dynamics under symmetric competing lobbyists, which are not naturally addressed in standard zealot formulations.

While existing research has illuminated the role of social networks and public discourse in shaping climate attitudes and policy preferences [35], the development of formal mathematical frameworks to analyze lobbying dynamics remains a promising frontier. Our work contributes to this direction by combining cognitive biases, Bayesian-like social learning, and strategic lobbying behavior within a unified modeling framework.

III. Model and Methods

In this section, we present the modeling framework underlying our analysis. We begin by describing the *baseline opinion dynamics model*, which characterizes how agents update their beliefs through social interaction. We then introduce *lobbyists* as strategic, budget-constrained external actors whose time staged signals influence the evolution of opinions in the network. Finally, we outline the *simulation setup*, including the network structure, parameter choices, scenario definitions, and the performance metrics used in our numerical experiments.

A. The Baseline Model

We consider a social network of individuals who repeatedly interact communicating on the occurrence of an uncertain harmful event. For instance, they can do that by sharing posts, sending messages, and creating content. We refer to any such item generically as a signal sent by an individual to their neighbors. For simplicity, signals are binary: they either assert that the event will occur or that it will not. Guided by evidence of selective behavior in social networks [36], individuals choose which type of signal to send based on their opinions. For example, an individual who assigns a high probability to the occurrence of the event will be more likely to share a signal affirming that it will occur. Consistent with evidence on social influence and contagion [17], [37], [38], [39], received signals affect the recipients' opinions. Note that whether signals contain a kernel of truth is not essential for generating opinion change that, in turn, increases the likelihood of sharing similar signals. This is consistent with findings that false news diffuses more than true news on social media [40]. We model opinion formation and updating as a modification of Bayes' rule: individuals behave as if signals from others are informative about the true data-generating process. At the same time, we allow the updating mechanism to deviate from exact Bayes to capture common processing biases, such as under-reaction [11], [12], [13], [41] and confirmation bias [14], [42]. This setup is also analogous to selective perception [43], a well-documented feature of social-media environments [36], [44]. Hence, our model represents an extreme form of social learning [7], [8], [9], [10], in which individuals do not receive private signals from the true data-generating process but try to learn from what other agents share.

In formal terms, N individuals interact for τ time steps about the occurrence of an uncertain damaging event. Individuals do not know the true probabilities that determine the occurrence of the event, but can rely on two probabilistic models. Those are the probability distributions $(\pi_o, 1-\pi_o)$ and $(\pi_p, 1-\pi_p)$, where π_k, with $k \in \{o, p\}$, indicates the probability that the damaging event occurs. The model $(\pi_o, 1-\pi_o)$ is *optimistic*, while the model $(\pi_p, 1-\pi_p)$ is *pessimistic*: $0 < \pi_o < \pi_p < 1$ with π_o close to zero and π_p close to one.

The social network is represented by a *directed* graph, where a link from i to j means that agent i communicates to agent j but not vice-versa (for example, followers in online social networks). We represent this by the $N \times N$ adjacency matrix $A = (a_{i,j})$. Communication occurs in interaction rounds. At each time step $t \in \{1, \ldots, \tau\}$, an individual ι_t is independently and uniformly drawn from the population and it sends a signal $s_t \in \{0, 1\}$ to the set of individuals to whom it communicates, that is $J(\iota_t) = \{j \in \{1, 2, \ldots, N\} | a_{\iota_t, j} = 1\}$. The signal $s_t = 1$ indicates that the individual communicates that the event will occur, while $s_t = 0$ indicates that the individual communicates that the event will not occur.

Each individual i has a subjective probability distribution $\mathbf{p}_{i,t} = (p_{i,t}, 1 - p_{i,t})$ on the realization of the event, with $p_{i,t}$ the probability attached at time step t to the case in which the event occurs. Such a distribution evolves according to the signals the individual receives from its network of connections, and it is used to draw the signals it sends to the individuals to whom it communicates. Concerning the evolution of subjective probabilities, for any $i \in \{1, 2, \ldots, N\}$ and $t \in \{1, \ldots, \tau\}$, it is as follows:

$$p_{i,t} = w_{i,t}\pi_o + (1 - w_{i,t})\pi_p \tag{1}$$

with $w_{i,t} \in (0, 1)$. That is, any individual builds its subjective probabilities as the convex combination of the probabilistic predictions of the two models. This definition of $p_{i,t}$ means that $p_{i,t}$ will always stay in the interval $[\pi_o, \pi_p]$, even if $w_{i,t}$ falls outside this interval. The weights $w_{i,t}$ used to build subjective probabilities change depending on the received signals. In particular, for any individual i and for any time step t, the weight

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

assigned to the optimistic model evolves according to

$$w_{i,t} = \begin{cases} w_{i,t-1}, & \text{if } i \notin J(\iota_t) \\ \lambda_{i,t} w_{i,t-1} + \dfrac{s_t(1-\lambda_{i,t})\pi_o w_{i,t-1}}{p_{i,t-1}} & \\ +\dfrac{(1-s_t)(1-\lambda_{i,t})(1-\pi_o)w_{i,t-1}}{1-p_{i,t-1}}, & \text{if } i \in J(\iota_t) \end{cases} \tag{2}$$

where $\lambda_{i,t} \in [0,1]$ modulates the effects of cognitive biases of agent i at time step t and $w_{i,0} \in (0,1)$ is the initial prior. As pointed out before, our agents build their beliefs in a way that resembles Bayesian updating but is applied in an extreme social-learning setting. They behave as if the observed social signals were i.i.d. drawn from the true process, although these signals are noisy observations of other agents' contemporaneous beliefs. In the absence of new evidence, and under common knowledge of identical 0.5 priors, Bayesians would keep posteriors at 0.5. By contrast, our agents treat social signals as informative about the state, departing from proper Bayesian-like learning. We nonetheless interpret $\lambda_{i,t}$ as capturing specific cognitive biases on top of this process. In particular, although proper Bayesian learning cannot be recovered in this social-learning environment, setting $\lambda_{i,t} = 0$ makes agent i update by Bayes' rule, and this acts as a benchmark within the proposed setting. If $\lambda_{i,t} = 1$, then i keeps the initial combination fixed no matter the sequence of signals it receives. Intermediate values of $\lambda_{i,t}$ indicate that, as i receives a signal, it updates the weights in the right direction but in lower magnitude than what Bayes' rule would prescribe. Thus, individuals can under-react to signals [11], [12], [13], producing the well-documented pattern in belief updating often referred to as conservatism [41]. We enrich such a behavioral interpretation imposing that $\lambda_{i,t}$ depends upon the type of signal received and the prior. In this way, it captures a form of directional motivated reasoning that, in turn, gives rise to confirmation bias (the tendency of individuals to discard information that contradicts their priors), another common feature in the formation of beliefs [14], [42]. In particular, we assume

$$\lambda_{i,t} = \phi_i |1 - s_t - w_{i,t-1}| + (1-\phi_i)\lambda_i \tag{3}$$

with $\phi_i, \lambda_i \in [0,1]$. Note that our updating rule differs from linear averaging rules commonly studied in stubborn-agent models [32], because the effective step size is endogenous to the signal's (in)congruence with priors (3). This feature is central to the regime behavior we document later.

The overall degree of under-reaction is the convex combination of two components. On the one hand, the parameter λ_i captures the baseline level of under-reaction of agent i. On the other hand, the function $|1 - s_t - w_{i,t-1}|$ introduces directional motivated reasoning or, equivalently, for our analysis and interpretation, confirmation bias. The parameter ϕ_i regulates the strength of the directional motivated reasoning. Suppose, for instance, that directional motivated reasoning is strong ($\phi_i \simeq 1$), if an individual has a strong prior in the optimistic model ($w_{i,t-1} \simeq 1$) and receives a signal that favors it with respect to the pessimistic model ($s_t = 0$) then its overall degree of under-reaction is low: information has a sensible impact on its beliefs. If, instead, the information contradicts its priors ($s_t = 1$), then the degree of under-reaction is high: information has a negligible effect on its beliefs. A symmetric argument holds if the agent's priors favor the pessimistic model. If, instead, directional motivated reasoning is weak ($\phi_i \simeq 0$), then the individual under-reacts with respect to information because of the effect of λ_i, but the nature of the signal and its relation to the prior have a negligible effect on probability updating. Hence, λ_i can be considered as a parameter capturing under-reaction generated by sources different from directional motivated reasoning. Moreover, ϕ_i allows us to accommodate a wide range of real-life situations, from those in which directional motivated reasoning plays a stronger role (e.g., emotionally involving and divisive topics) to the ones in which it may be negligible (e.g., more technical topics). As mentioned above, this setup is also akin to a form of selective perception.

The functional form in (2) and (3) is not intended to be the unique representation of biased belief updating. Rather, it is constructed to satisfy three design principles that are central to our analysis. First, it nests Bayesian updating as a benchmark case when the bias parameters vanish, thereby providing a normative anchor consistent with the non-Bayesian social learning literature [7], [10]. Second, it allows for proportional under-reaction through a tractable attenuation parameter λ_i, in line with experimental evidence on conservatism in belief revision [11], [12]. Third, and crucially, responsiveness to signals is state-dependent: the term $|1 - s_t - w_{i,t-1}|$ ensures that the effective step size varies smoothly with the degree of congruence between prior beliefs and incoming information. This generates a continuous form of directional motivated reasoning, consistent with empirical findings on confirmation bias [14], without imposing discrete thresholds as in bounded-confidence models. Alternative nonlinear specifications could incorporate similar behavioral features. However, the regime patterns documented in Section IV rely on the presence of endogenous, signal-dependent responsiveness rather than on this specific algebraic form. The chosen formulation provides a parsimonious way to embed these behavioral mechanisms while preserving analytical transparency and bounded belief support.

Selective behavior in the generation of signals is modeled assuming that the signal at time step t is drawn from the distribution $\mathbf{p}_{\iota_t,t}$ (i.e., Prob$\{s_t = 1\} = p_{\iota_t,t}$). Thus, the timeline of events at each time step t is as follows: 1) ι_t is independently and uniformly drawn from the population; 2) the set of receivers $J(\iota_t)$ is created; 3) the signal s_t is drawn from the distribution $\mathbf{p}_{\iota_t,t}$; 4) individual weights are updated according to (2).

B. Introducing Lobbyists

Lobbyists can be considered external agents with respect to the social network that, in each turn, can pay a cost to send signals to individuals. Their objective is to influence the time τ distribution of beliefs so that it becomes concentrated on one of the two models. We emphasize that this module captures the opinion-formation channel of lobbying; the policy-adoption stage is not modeled and is treated as exogenous to the opinion dynamics studied here.

Assume that there are L active lobbyists and define $S^\ell_{i,t} \in \{0,1\}$ as a variable indicating whether lobbyist $\ell \in \{1,2,\dots,L\}$ sends a signal to individual i at the beginning of time step t ($S^\ell_{i,t}=1$) or not ($S^\ell_{i,t}=0$). It follows that a pure strategy for lobbyist ℓ can be indicated as a matrix $S^\ell \in \mathcal{S} = \{0,1\}^{N\times\tau}$. Each lobbyist supports one model between the optimist and the pessimist one, and the signals it sends are in favor of the model it supports. For instance, a lobbyist supporting the optimistic model will signal 0, that is, the event is not occurring. Define the function that assigns to each lobbyist the model it supports as $m:\{1,2,\dots,L\}\to\{o,p\}, \ell\mapsto m(\ell)$. We assume that signals are costly. In particular, the cost of a signal is fixed to a given level $c>0$, and each lobbyist ℓ has a budget B^ℓ. Without loss of generality, we shall set $c=1$ in what follows. That is, B^ℓ shall be understood as the maximum number of signals lobbyist ℓ can send. Notice that this assumption constrains the set of strategies that each lobbyist can actually play, that is, a feasible strategy is such that $\sum_{t=1}^{\tau}\sum_{i=1}^{N} S^\ell_{i,t} \le B^\ell$. Hence, we define the set of feasible strategies for lobbyist ℓ as $\mathcal{S}_\ell = \{S^\ell \in \mathcal{S} \mid \sum_{t=1}^{\tau}\sum_{i=1}^{N} S^\ell_{i,t} \le B^\ell\}$. We allow lobbyists to randomize their choices. Call $\Delta(\mathcal{S}_\ell)$ the mixed extension of $\mathcal{S}_\ell$ and $\sigma^\ell(S)$ the probability that lobbyist ℓ attaches to a feasible strategy S, such that $\sigma^\ell \in \Delta(\mathcal{S}_\ell)$ is a feasible mixed strategy for lobbyist ℓ. Note that this strategic, budget-constrained sending of costly, time-staged signals differs from exogenous stubborn-node formulations. Indeed, each individual reacts to the signal received in the same way as they do to peer signals. Therefore, now, individual probabilities also evolve as a consequence of the strategies of the lobbyists. In particular, call $S^{-\ell} = (S^1,\dots,S^{\ell-1},S^{\ell+1},\dots,S^L) \in \mathcal{S}_1\times\dots\times\mathcal{S}_{\ell-1}\times\mathcal{S}_{\ell+1}\times\dots\times\mathcal{S}_L$ a profile of feasible pure strategies for the lobbyists different from ℓ, such that $p_{i,\tau}(S^\ell,S^{-\ell})$ indicates the final probability of individual i as a function of the strategies played by all lobbyists. Define the vector of lobbyists' signals received by agent i at time t as $\mathcal{L}^{i,t} \in \{0,1,2\}^L$, where $\mathcal{L}^{i,t}_\ell = 0$ if $S^\ell_{i,t}=0$, $\mathcal{L}^{i,t}_\ell = 1$ if $S^\ell_{i,t}=1$ and $m(\ell)=p$, $\mathcal{L}^{i,t}_\ell = 2$ if $S^\ell_{i,t}=1$ and $m(\ell)=o$. Thus, assume that lobbyists send signals at the beginning of each time step t and that the order in which their signals are received by individuals is $Z_t \in \mathrm{Perm}\{1,\dots,L\}$, i.e., an element of the set of all permutations of the indexes of lobbyists (in what follows, we assume that the order according to which lobbyists send their signals is randomly chosen). The weight assigned to the optimistic model deriving from previous interaction needs to be updated for any $z\in\{1,\dots,L\}$ according to

$$w^z_{i,t-1} = \begin{cases} w^{z-1}_{i,t-1}, & \text{if } \mathcal{L}^{i,t}_{\ell_z}=0 \\ \lambda^{1,z}_{i,t} w^{z-1}_{i,t-1} + \dfrac{(1-\lambda^{1,z}_{i,t})\pi_o w^{z-1}_{i,t-1}}{\pi_o w^{z-1}_{i,t-1} + \pi_p(1-w^{z-1}_{i,t-1})}, & \text{if } \mathcal{L}^{i,t}_{\ell_z}=1 \\ \lambda^{2,z}_{i,t} w^{z-1}_{i,t-1} + \dfrac{(1-\lambda^{2,z}_{i,t})(1-\pi_o) w^{z-1}_{i,t-1}}{1-\pi_o w^{z-1}_{i,t-1} - \pi_p(1-w^{z-1}_{i,t-1})}, & \text{if } \mathcal{L}^{i,t}_{\ell_z}=2 \end{cases} \quad (4)$$

where ℓ_z indicates the lobbyist index number appearing in zth position of the Z_t vector, $w^0_{i,t-1} = w_{i,t-1}$

$$\lambda^{1,z}_{i,t} = \phi_i w^{z-1}_{i,t-1} + (1-\phi_i)\lambda_i$$

and

$$\lambda^{2,z}_{i,t} = \phi_i (1-w^{z-1}_{i,t-1}) + (1-\phi_i)\lambda_i.$$

In this way, the cognitive biases discussed above influence belief updating also when signals are sent by lobbyists. The weight updating rule for the signals received from other individuals needs to be adjusted, indeed, (2) now becomes

$$w_{i,t} = \begin{cases} w^L_{i,t-1}, & \text{if } i \notin J(\iota_t) \\ \lambda_{i,t} w^L_{i,t-1} + \dfrac{s_t(1-\lambda_{i,t})\pi_o w^L_{i,t-1}}{\pi_o w^L_{i,t-1} + \pi_p(1-w^L_{i,t-1})} + \dfrac{(1-s_t)(1-\lambda_{i,t})(1-\pi_o) w^L_{i,t-1}}{1-\pi_o w^L_{i,t-1} - \pi_p(1-w^L_{i,t-1})}, & \text{if } i \in J(\iota_t) \end{cases} \quad (5)$$

with $\lambda_{i,t} = \phi_i|1-s_t-w^L_{i,t-1}| + (1-\phi_i)\lambda_i$ as above. Subjective probabilities, instead, remain as in (1) since they are computed at the end of the interaction round.

Finally, we assume that the payoff of lobbyist ℓ using the feasible mixed strategy σ^ℓ when the other lobbyists use the feasible mixed strategies $\sigma^1,\dots,\sigma^{\ell-1},\sigma^{\ell+1},\dots,\sigma^L$ is as follows:

$$U_\ell = -\mathrm{E}\left[\sum_{S^1\in\mathcal{S}_1}\cdots\sum_{S^L\in\mathcal{S}_L}\prod_{l=1}^{L}\sigma^l(S^l) \left|\sum_{i=1}^{N}\frac{p_{i,T}(S^\ell,S^{-\ell})}{N} - \pi_{m(\ell)}\right|\right]$$

where the expectation E is computed with respect to the random variables deciding the order in which lobbyists send their signals, the selection of communicating individuals, and the signaling choices operated by individuals. The underlying intuition is that a lobbyist benefits as the expected distance between the average final belief and the model it supports decreases.

When lobbyists are active, the timeline of events at each time step t becomes: 1) each lobbyist sends the signals to the individuals it selected; 2) individual weights are updated according to (4); 3) ι_t is independently and uniformly drawn from the population; 4) the set of receivers $J(\iota_t)$ is created; 5) the signal s_t is drawn from the distribution $\mathbf{p}_{\iota_t,t}$; and 6) individual weights are updated according to (5).

C. Simulation Setup

In all of our simulations, we consider a complete network of 500 agents. The adjacency matrix A is therefore of size 500×500, with diagonal entries $a_{i,i}=0$, and off-diagonal entries $a_{i,j\neq i}=1$. Each agent is characterized by an individual-specific initial weight $w_{i,0}$ which is randomly drawn from a uniform distribution with support $(0,1)$. The probabilities characterizing the two models (optimistic and pessimistic) are $\pi_o=0.01$ and $\pi_p=0.99$. We consider three main scenarios.

TABLE I
SUMMARY OF MODEL'S PARAMETERS, THEIR DESCRIPTIONS, AND THEIR BASELINE NUMERICAL VALUES

Parameter	Description	Baseline Value
N	Number of individuals in the network	500
T	Number of time steps of lobbyist activity	100
π_o	Probability of damaging event under optimistic model	0.01
π_p	Probability of damaging event under pessimistic model	0.99
ϕ	Strength of directional motivated reasoning/confirmation bias	[0, 1]
λ	Degree of baseline under-reaction	[0, 1]
B	Budget of each lobbyist	10 000
c	Cost per signal sent by each lobbyist	1

In the "baseline" scenario, the network is left to interact with itself, with no outside interference, i.e., no lobbyist. In the "one-lobbyist" scenario, we introduce a single active lobbyist ($L = 1$) which supports the pessimistic model ($m(1) = p$). In the "two-lobbyists scenario," we introduce two active lobbyists that support competing models ($m(1) = p$, and $m(2) = o$). In each scenario contemplating lobbyists, each one of them is endowed with a budget $B = 10\,000$. For the purpose of our numerical analysis, we introduce the parameter $T \leq \tau$, representing the lobbyists' time horizon, that is, the time span over which it allocates its budget B^ℓ. Accordingly, we set $S^\ell_{i,t} = 0\, \forall t \in (T, \tau]$. This constraint allows us to manage the potentially vast space of feasible lobbying strategies and enables analysis in scenarios with large τ. In particular, neglecting the zeros appearing after T, we assume that lobbyists dispose of a set of 100 strategy matrices of size ($T \times 500$) in which the variable indicators for signals are distributed randomly over nodes and time steps, and choose one of those matrices at random at the beginning of each independent simulation run. The use of randomized uniform strategies (as well as being supported by the theoretical exercise in Section IV-C) serves as a controlled benchmark rather than a claim about optimal lobbying behavior. Our objective is to isolate the structural interaction between cognitive bias, endogenous credibility, and resource constraints without introducing additional layers of dynamic targeting or network-aware optimization. Allowing strategies to condition on node centrality or intermediate belief realizations would transform the framework into a dynamic strategic control problem in which optimization coevolves with diffusion dynamics. The present specification therefore provides a minimal baseline that permits identification of the core cognitive–diffusion mechanisms generating the regime results documented as follows.

For most of our experiments, we set $T = 100$. This means that, over 100 periods, each lobbyist is able to send 100 signals per period, thus reaching out to 20% of the nodes at each time step. Moreover, this strategy selection procedure resembles an adaptation of the uniform-probability mixed strategy that emerges as the Nash equilibrium in Example 2. Throughout, we interpret results as effects on lobbying-driven opinion influence; any policy-adoption mapping is intentionally left outside the model.

As shown in (3), each $\lambda_{i,t}$ is subject to two different types of bias, captured by the parameters λ_i and ϕ_i. While λ_i captures the degree of under-reaction to *any* new information, ϕ_i represents the strength of directional motivated reasoning (or confirmation bias), which is signal-dependent. We explore the properties of the model assuming homogeneous biases ($\lambda_i = \lambda$ and $\phi_i = \phi$ for all individuals i) and considering different configurations of the network with respect to these two independent dimensions, i.e., a parameter space generated by $(\lambda, \phi) \in [0, 1] \times [0, 1]$. For reference, Table I summarizes the model parameters, their descriptions, and the baseline values.

Given the model specification, the population can achieve a steady state when each individual fully supports either one model or the other, p_i is either pi_o or pi_p. This happens when w_i is either 1 or 0, respectively. It is easy to see from the equations that an individual reaching one of these two opinions can never change opinion, regardless of the signals received. However, given the multiplicative nature of the belief updating process, these two fixed points can be reached only asymptotically. Still, in the presence of high directional motivated reasoning, two opinion clusters are possible in the population if some individuals end up close to one model and the rest end up close to the other model. This is because the degree of overall under-reaction increases as w_i approaches its boundaries and the process slows down. The path-dependent nature of signal diffusion among peers contributes to fostering and preserving the polarization. Numerically, to avoid precision errors, when weights w approach 0 or 1 with a precision of 10^{-12}, they are set to 0 or 1, respectively. Regarding simulation time, we fix it a priori by setting a maximum number of iterations, but making sure that for all simulations the number of iterations is enough to reach a steady state. The simulation stops when opinions do not change more than the same precision parameter above.

In all simulations, we monitor various criteria. An important aspect is the number of clusters that form in the population: consensus means one cluster, more clusters mean fragmentation and polarization. We compute opinion clusters by using a simple threshold-based binning procedure. First, individual opinions are sorted and grouped together such that values lying within a predefined distance threshold ($\epsilon = 0.0001$) are treated as belonging to the same group. Rather than reporting the raw number of clusters, we compute the *effective* number of clusters defined as

$$C = \frac{N^2}{\sum_i^k N_i^2} \tag{6}$$

where k is the number of clusters and N_i is the size of each cluster. This is a concentration-adjusted measure accounting for both the number and relative size of belief groups. C is equal

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

to k when we have k clusters of equal size, and it decreases toward one as opinions become increasingly concentrated in a dominant group, providing a continuous measure of opinion fragmentation. For example, if 1/10 of the population is in one cluster and 9/10 in a second one, then the effective number of clusters is about 1.22, indicating that one of the two clusters is much larger. A second criterion of interest is the average opinion of agents at the end of the simulation, which we report in terms of average of the $p_{i,\tau}$ values as defined in (1). In addition to indicating how beliefs have collectively shifted, this criterion is also informative of lobbyists' payoff. Furthermore, we look at the number of time steps required for convergence, indicating the speed of opinion cluster formation. For all parameter configurations, we provide mean values of these criteria over 150 independent runs, each one of which is simulated until the model reaches convergence (i.e., until all the individual weights assigned to the competing models stop updating). We have also experimented with a larger number of runs, up to 1000; however, results did not change significantly while running times did, therefore we opted to keep only 150 runs per configuration.

IV. Results and Analysis

In this section, we provide a theoretical analysis of the properties of the model and present the results of our numerical experiments aimed at analyzing the behavior of the model across different configurations. In the theoretical analysis, we first study the asymptotic dynamics of the individual process of belief updating, providing general results and characterizing the behavior of the model without lobbyists. Then, we investigate the lobbyists' equilibrium strategies in two simple cases. In the numerical experiments, we begin by examining the baseline dynamics of the system in the absence of lobbyists, highlighting the emergence of consensus or polarization under different cognitive-bias parameters. We then study the effects of a single lobbyist, followed by the more complex scenario involving two competing lobbyists. After characterizing the main regimes of behavior, we explore how lobbying performance varies with key parameters through a set of sensitivity analyses and investigate the role of timing by comparing frontloading, backloading, and uniform strategies. Together, these results provide a comprehensive picture of how cognitive biases, peer influence, and strategic lobbying jointly shape opinion evolution in the network.

A. Asymptotic Dynamics

Here, we study the asymptotic dynamics of the belief-updating process of an individual agent in the network. As we shall see, such an exercise can provide us with relevant information on the evolution of the whole network of opinions. To do so, assume $\tau \to \infty$, focus on an agent, and let us simplify the notation by letting w_n denote the probability that such an agent assigns to the optimist model being correct after the nth round of interaction, regardless of whether the interaction is with a peer or a lobbyist. Hence, in accordance with our setting, such a quantity evolves according to

$$w_{n+1} = \begin{cases} w_n, & \text{with Prob. } 1-\rho_n^1-\rho_n^2 \\ \lambda^1(w_n)w_n + \dfrac{(1-\lambda^1(w_n))w_n\pi_o}{\pi_o w_n + \pi_p(1-w_n)}, & \text{with Prob. } \rho_n^1 \\ \lambda^2(w_n)w_n + \dfrac{(1-\lambda^2(w_n))w_n(1-\pi_o)}{1-\pi_o w_n - \pi_p(1-w_n)}, & \text{with Prob. } \rho_n^2 \end{cases} \tag{7}$$

where $\lambda^1(w) = \phi w + (1-\phi)\lambda$, $\lambda^2(w) = \phi(1-w) + (1-\phi)\lambda$, ρ_n^1 indicates the probability of receiving signal 1 (the event will occur) in the nth interaction, and ρ_n^2 indicates the probability of receiving signal 0 (the event will not occur) in the nth interaction. Notice that, as long as the agent has at least one incoming link and the lobbyists assign positive probability to sending a signal to each agent in each of the T rounds during which they are active, there exists a number $\varepsilon > 0$ such that $\rho_n^1, \rho_n^2 \geq \varepsilon$. Moreover, in this section we shall assume $w_0, \phi, \lambda \in (0,1)$.

In this context, the analysis of $\{w_n\}$ mirrors those pursued in different settings by [45], [46], and [47] and builds on the mathematical results of [48]. In particular, define the process

$$x_n = \log \frac{w_n}{1-w_n}$$

and notice that $x_n \to +\infty$ if and only if $w_n \to 1$ and $x_n \to -\infty$ if and only if $w_n \to 0$. Then, the following lemma holds.

Lemma IV.1: There exists a number $Y > 0$ such that for all n it is $\text{Prob}(|x_{n+1} - x_n| < Y) = 1$. Moreover, there exists a number δ such that, almost surely, $\text{Prob}(x_{n+1} > x_n + \delta) > \delta$ and $\text{Prob}(x_{n+1} < x_n - \delta) > \delta$.

Proof: Define $f_1(w_n) = x_{n+1} - x_n$ in the case in which signal 1 is observed and $f_2(w_n) = x_{n+1} - x_n$ in the case in which signal 0 is observed. Then,

$$f_1(w) = \log \frac{\pi_o + \lambda^1(w)(1-w)(\pi_p - \pi_o)}{\pi_p - \lambda^1(w)w(\pi_p - \pi_o)}$$

and

$$f_2(w) = \log \frac{1 - \pi_o - \lambda^2(w)(1-w)(\pi_p - \pi_o)}{1 - \pi_p + \lambda^2(w)w(\pi_p - \pi_o)}.$$

Noticing that $1 - e^{f_1(w)} \geq (\pi_p - \pi_o)(1-\phi)(1-\lambda)/\pi_p$ and $e^{f_2(w)} - 1 \geq (\pi_p - \pi_o)(1-\phi)(1-\lambda)/(1-\pi_o)$, one immediately has $\forall w \in [0,1]$

$$\begin{aligned} -\infty < \log \frac{\pi_o}{\pi_p} &\leq f_1(w) \\ &\leq \log\left(1 - \frac{(\pi_p - \pi_o)(1-\phi)(1-\lambda)}{\pi_p}\right) < 0 \end{aligned}$$

and

$$\begin{aligned} 0 &< \log\left(1 + \frac{(\pi_p - \pi_o)(1-\phi)(1-\lambda)}{1-\pi_o}\right) \\ &\leq f_2(w) \leq \log \frac{1-\pi_o}{1-\pi_p} < +\infty. \end{aligned}$$

Choosing $Y = 2\max\{\log((1-\pi_o)/(1-\pi_p)), \log(\pi_p/\pi_o)\}$, the first statement directly follows. For the second one, notice that $\rho_n^1\,\rho_n^2 \geq \varepsilon > 0\;\forall n$, hence, choose

$$\delta = \frac{1}{2}\min\left\{\varepsilon, \log\left(1+\frac{(\pi_p-\pi_o)(1-\phi)(1-\lambda)}{1-\pi_o}\right),\right.$$
$$\left. -\log\left(1-\frac{(\pi_p-\pi_o)(1-\phi)(1-\lambda)}{\pi_p}\right)\right\}$$

and the statement follows. □

Lemma IV.1 shows that the process x_n has bounded increments, with finite positive and finite negative increments in the sense of [48]. This allows us to characterize the asymptotic behavior of x_n (and thus of w_n) by studying its drift, namely the function

$$\begin{aligned}\mu_n(w) &= \mathbb{E}_n\left[x_{n+1}-x_n \,\middle|\, x_n = \log\frac{w}{1-w}\right]\\ &= \rho_n^1 \log\frac{\pi_o+\lambda^1(w)(1-w)(\pi_p-\pi_o)}{\pi_p-\lambda^1(w)w(\pi_p-\pi_o)}\\ &\quad + \rho_n^2\log\frac{1-\pi_o-\lambda^2(w)(1-w)(\pi_p-\pi_o)}{1-\pi_p+\lambda^2(w)w(\pi_p-\pi_o)}.\end{aligned}$$

In particular, the following holds.

Proposition IV.1: If definitely in n it is

$$\frac{\rho_n^2}{\rho_n^1} > \frac{1-(1-\phi)\lambda}{(1-\phi)(1-\lambda)}\,\frac{\log\pi_p-\log\pi_o}{\log(1-\pi_o)-\log(1-\pi_p)} \tag{8}$$

then

$$\text{Prob}\left(\limsup_{n\to\infty} w_n > 0\right) = 1.$$

If definitely in n it is

$$\frac{\rho_n^1}{\rho_n^2} > \frac{1-(1-\phi)\lambda}{(1-\phi)(1-\lambda)}\,\frac{\log(1-\pi_o)-\log(1-\pi_p)}{\log\pi_p-\log\pi_o} \tag{9}$$

then

$$\text{Prob}\left(\liminf_{n\to\infty} w_n < 1\right) = 1.$$

Proof: The two statements follow from applying Theorem 3.1 and Corollary 3.2 of [48] to x_n, which has bounded increments by Lemma IV.1. For the first one, notice that, from Jensen's inequality, one has the following:

$$\begin{aligned}&\mu_n(0)\\ &\geq \rho_n^1(1-(1-\phi)\lambda)\log\frac{\pi_o}{\pi_p} + \rho_n^2(1-\phi)(1-\lambda)\log\frac{1-\pi_o}{1-\pi_p}\end{aligned}$$

and (8) implies $\mu_n(0) > 0$ definitely in n. Thus, there exists a $M > 0$ and large enough such that $\mu_n(w) > 0$ definitely in n for all $x = \log(w/(1-w)) < -M$. Theorem 3.1 of [48] implies $\text{Prob}\,(\limsup_{n\to\infty} x_n > -M) = 1$ and the statement follows. For the second one, notice that, from Jensen's inequality, one has the following:

$$\begin{aligned}&\mu_n(1)\\ &\leq \rho_n^1(1-\phi)(1-\lambda)\log\frac{\pi_o}{\pi_p} + \rho_n^2(1-(1-\phi)\lambda)\log\frac{1-\pi_o}{1-\pi_p}\end{aligned}$$

and (9) implies $\mu_n(1) < 0$ definitely in n. Thus, there exists a $M > 0$ and large enough such that $\mu_n(w) < 0$ definitely in n for all $x = \log(w/(1-w)) > M$. Corollary 3.2 of [48] implies $\text{Prob}\,(\liminf_{n\to\infty} x_n < M) = 1$ and the statement follows. □

Since $1-(1-\phi)\lambda \geq (1-\phi)(1-\lambda)$, Proposition IV.1 implies that, under symmetric models ($\pi_o = 1-\pi_p$), a sufficient condition to prevent the agent from becoming asymptotically convinced of a given model with positive probability, is that the lobbyist supporting the opponent model must act so that the probability assigned to the signal compatible with the supported model increases well beyond the complementary probability mass.

We can further characterize the asymptotic behavior of the system in the case in which the network is fully connected (including self-loops, i.e., $a_{i,j} = 1$ for all i, j), all individuals share the same initial priors, and no lobbyist is active. Indeed, in this case the probabilities of receiving the signals at time $n+1$ coincide with the probabilities that the agent assigned to the two signals at step n, namely

$$\rho_n^1 = \pi_o w_n + \pi_p(1-w_n), \qquad \rho_n^2 = 1-\rho_n^1. \tag{10}$$

Indeed, the following proposition holds.

Proposition IV.2: If the probabilities are as in (10), then, with probability 1 either $\lim_{n\to\infty} w_n = 0$ or $\lim_{n\to\infty} w_n = 1$.

Proof: The statement follows from Theorem 4.2 of [48] applied to $\{x_n\}$, which has bounded increments and finite positive and finite negative increments by Lemma IV.1. More specifically, notice that, by Jensen's inequality, it is as follows:

$$\begin{aligned}\mu_n(0) &\leq \log((1-(1-\phi)\lambda)\pi_o + (1-\phi)\lambda\pi_p\\ &\quad + (1-\phi)(1-\lambda)(1-\pi_o) + (\phi+(1-\phi)\lambda)(1-\pi_p))\\ &= \log(1-\phi(\pi_p-\pi_o)) < 0\end{aligned}$$

and

$$\begin{aligned}\mu_n(1) &\geq -\log((\phi+(1-\phi)\lambda)\pi_o + (1-\phi)(1-\lambda)\pi_p\\ &\quad + (1-\phi)\lambda(1-\pi_o) + (1-(1-\phi)\lambda)(1-\pi_p))\\ &= -\log(1-\phi(\pi_p-\pi_o)) > 0.\end{aligned}$$

Since the conditional second moment of the increment is bounded, the previous conditions, together with Theorem 4.2 of [48], ensure that either $\lim_{n\to\infty} x_n = -\infty$ or $\lim_{n\to\infty} x_n = +\infty$ almost surely. The claim then follows. □

Proposition IV.2 establishes that, under the stated probability assumptions, the network of individuals may exhibit path dependence. In particular, beliefs converge asymptotically to one of the two models. Hence, because all agents are identical, network opinions concentrate on either the optimistic or the pessimistic model. As a consequence, the system converges to consensus, but the limiting consensus (optimistic versus pessimistic) can be path dependent.

B. One Individual and One Rational Lobbyist

Consider the special case in which $N = L = 1$. Assume that the lobbyist supports the pessimistic model and that $w_{1,0}, \lambda_1, \phi_1 \in (0,1)$. By direct inspection of (4), one notices that the weight assigned to the optimistic model decreases if

$S^1_{1,t} = 1$ (for any t). This implies that the payoff of the lobbyist increases for each signal sent. Thus, a lobbyist that wants to maximize its payoff will use all of its budget to send signals. Since the specific timing of the signals does not matter, the optimal strategy of the lobbyist is as follows:

$$S^{1,*} = \begin{cases} (1,\ldots,1), & \text{if } \tau \le B^1 \\ S^1 \in \{(S^1_{1,1},\ldots,S^1_{1,\tau}) | \sum_{t=1}^{\tau} S^1_{1,t} = \lfloor B^1 \rfloor\}, & \text{if } \tau > B^1. \end{cases}$$

C. Two Unbiased Individuals, One Period, and Two Rational Lobbyists With Unitary Budget Supporting Opposite Models

Consider now the case in which $N = L = 2$, $\tau = 1$, $B^1 = B^2 = 1$, and $\lambda_i = \phi_i = 0\ \forall i \in \{1, 2\}$. Without loss of generality, we assume that lobbyist 1 supports the pessimistic model while lobbyist 2 supports the optimistic model. Since both time and budget are unitary (and not sending signals is strictly dominated), each lobbyist has two feasible pure strategies: sending a signal to individual 1 or to individual 2, that is, $S^\ell \in \mathcal{S}_\ell = \{(0, 1), (1, 0)\}$. Under the assumptions that the model structure is common knowledge and the lobbyists are rational, we can compute the Nash equilibrium [49] resulting from the strategic interaction of the two lobbyists.

Notice that with $\lambda_i = \phi_i = 0\ \forall i \in \{1, 2\}$ individuals follow Bayes-rule updating under the extreme social learning framework that characterizes our model. In this one-period, signal-symmetric case, the order of receipt is immaterial for the final posteriors. Thus, to compute final probabilities, we only need to consider the following set of possible received signals:

$$\Omega = \{\emptyset, \{0\}, \{1\}, \{0, 0\}, \{0, 1\}, \{1, 1\}, \{0, 0, 1\}, \{0, 1, 1\}\}.$$

That is, since the lobbyists support opposite models and one of the individuals sends one signal to the other, the most extreme cases are the one in which an individual does not receive any signal, and the one in which it receives two (different) signals from the lobbyists and a signal from the other agent. All the other cases in between are possible. Assuming a uniform prior distribution for both individuals, the two agents are *ex-ante* identical, meaning that the differences in their final probabilities are uniquely generated by the signals they received. Hence, we can simplify the notation by calling $w(\omega)$ the weight assigned to the optimistic model by an agent that has received the signals $\omega \in \Omega$. As a direct consequence, we define the probability such an agent assigns to the realization of the event as $p(\omega) = w(\omega)\pi_o + (1 - w(\omega))\pi_p$. Then, calling t_s, with $s \in \{0, 1\}$, the number of times s occurs in ω, one has the following:

$$w(\omega) = \frac{\pi_o^{t_1}(1-\pi_o)^{t_0}}{\pi_o^{t_1}(1-\pi_o)^{t_0} + \pi_p^{t_1}(1-\pi_p)^{t_0}}.$$

Given the previous assumptions, the symmetry they entail, and the fact that $p(\omega) \in (\pi_o, \pi_p)\ \forall\omega$, the payoff of lobbyist $\ell \in \{1, 2\}$ conditional upon both lobbyists playing the same pure strategy ($S^1 = S^2$, lobbyists send their signals to the same individual) is as follows:

$$u_\ell(S^1 = S^2) = -\Big|\frac{1}{4}\,(p(0,1)(1+p(1)) + (1-p(0,1))p(0) + p(\emptyset)(1+p(0,1,1)) + (1-p(\emptyset))p(0,0,1)) - \pi_{m(\ell)}\Big|$$

while the payoff of ℓ conditional upon the two lobbyists choosing different pure strategies ($S^1 \neq S^2$, lobbyists send their signals to different individuals) is as follows:

$$u_\ell(S^1 \neq S^2) = -\Big|\frac{1}{4}\,(p(1)(1+p(0,1)) + (1-p(1))p(0,0) + p(0)(1+p(1,1)) + (1-p(0))p(0,1)) - \pi_{m(\ell)}\Big|.$$

Thus, for a generic mixed strategy profile, the payoff of lobbyist ℓ reads

$$U_\ell = (\sigma^1(1,0)\sigma^2(1,0) + \sigma^1(0,1)\sigma^2(0,1))u_\ell(S^1 = S^2) + (\sigma^1(0,1)\sigma^2(1,0) + \sigma^1(1,0)\sigma^2(0,1))u_\ell(S^1 \neq S^2).$$

However, notice that $u_1(S^1 = S^2) - u_1(S^1 \neq S^2) = u_2(S^1 \neq S^2) - u_2(S^1 = S^2)$, hence, neglecting the trivial case in which the differences are zero, there generically exists a unique Nash equilibrium characterized by $\sigma^\ell(1,0) = \sigma^\ell(0,1) = 1/2\ \forall \ell \in \{1, 2\}$.

D. Numerical Experiments: Baseline Scenario

In the first setting, we investigated the properties of the model in the absence of any active lobbyist, when the nodes are left free to interact with their peers without outside input. Fig. 1 displays the effective number of clusters and the average subjective probabilities across the parameter space when the stable state is reached. We observe that, for most of the parameter region under consideration, the distribution of subjective probabilities tends to converge to a general consensus, with all of the nodes eventually supporting a single model (either o or p, depending on initial conditions). This dynamic is relatively quick and is also illustrated by the evolution plot in Fig. 2(a).

Two notable exceptions to this behavior are worth mentioning: the first, and more trivial, occurs in the special case where $\lambda = 1$ and $\phi = 0$. In this case, as it is immediately clear from (3), nodes are “deaf” to whatever signal they receive from their peers, and any update of subjective probabilities is impossible. In this case, there is no dynamics, and each agent maintains its initial condition. The second exception emerges at the bottom and to the right of the parametric region, where the confirmation bias parameter ϕ is sufficiently large to prevent the network from reaching a consensus. In this case, in many runs the model converges to a polarized equilibrium, in which a variable proportion of agents supports one model, while the rest of the network supports the competing one [see Fig. 2(b) for an example]. This indicates that confirmation bias has an important role in our model, facilitating polarization of opinions, since agents are influenced little by peers with conflicting opinions. The effect is larger as lambda increases, i.e., as under-reaction is growing, confirmation bias becomes more effective in generating polarization. This seems to indicate that a society that

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

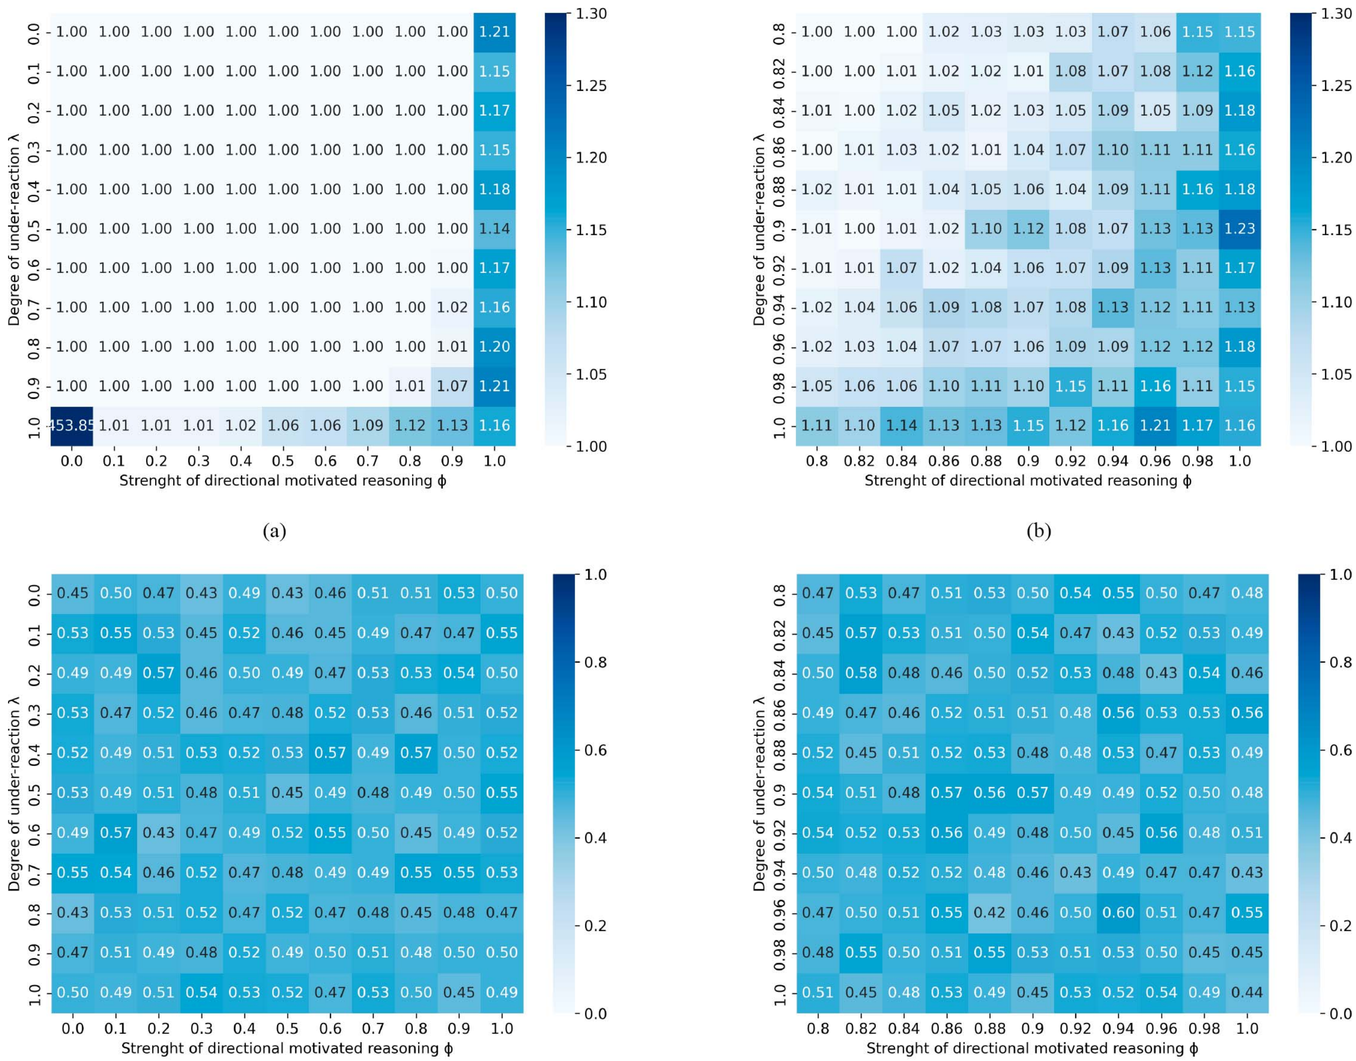

Fig. 1. Baseline (no-lobbyist) scenario. The average effective number of clusters and the final average subjective probabilities are shown as a function of the strength of directional motivated reasoning ϕ and the degree of under-reaction λ. In the case $\lambda = 1.0$, $\phi = 0.0$, the cluster metric is not applicable because agents retain their initial opinions drawn from a uniform distribution over $[0, 1]$. (a) Clusters: full space. (b) Clusters: zoom. (c) Avg. opinions: full space. (d) Avg. opinions: zoom.

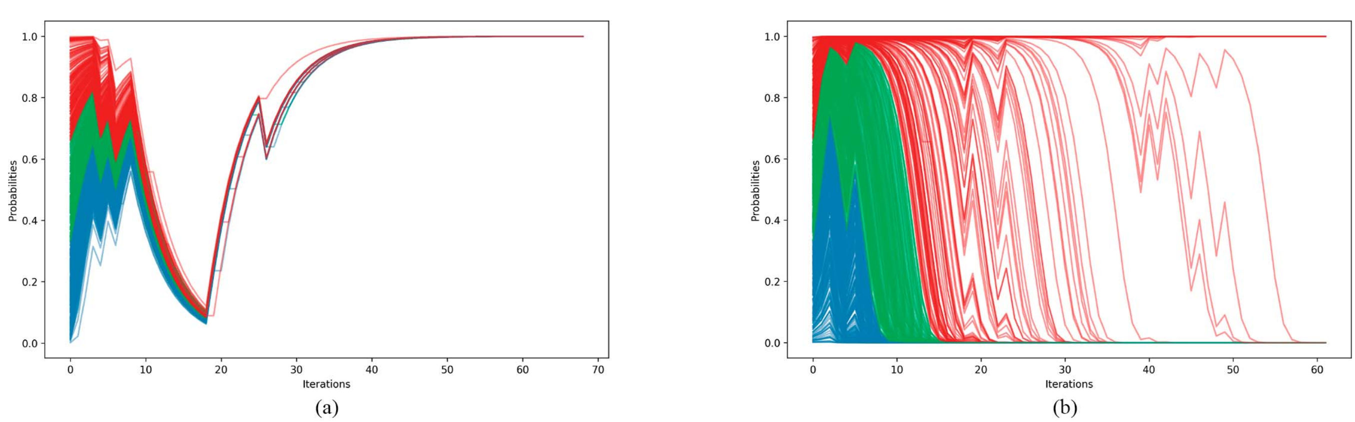

Fig. 2. Opinion evolution. Examples of the evolution of the agents' subjective probabilities in the baseline scenario with and without final consensus. (a) Convergence to one cluster. Each agent has a degree of under-reaction $\lambda = 0.8$ and a strength of directional reasoning $\phi = 0.0$. Initial opinions are uniformly distributed and color coded: red and blue are extreme opinions, and green is moderate. After about 60 time steps, all agents of the network have a final subjective probability $p_i = 0.99$. (b) Convergence to two clusters. Each agent has a degree of under-reaction $\lambda = 1.0$ and a strength of directional reasoning $\phi = 0.9$. Initial opinions are uniformly distributed and color coded: red and blue are extreme opinions, and green is moderate. After about 60 time steps, a cluster of agents of the network ($\approx 16\%$) reaches $p = 0.99$, while the rest (84%) settles at $p = 0.01$.

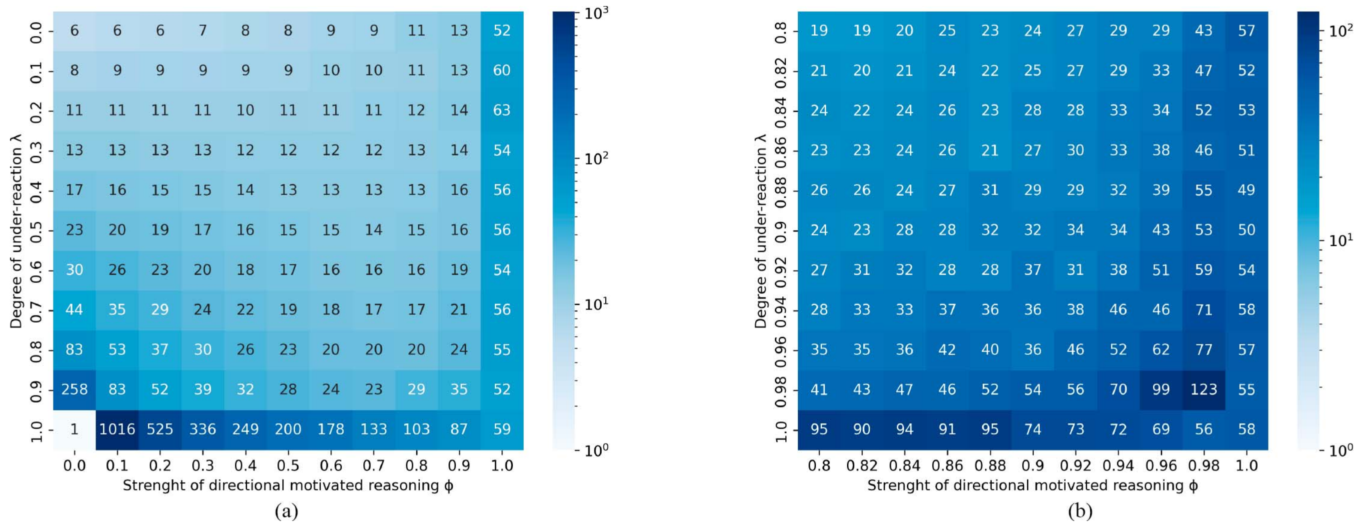

Fig. 3. Average number of time steps required to reach equilibrium in the baseline (no-lobbyist) scenario. The number of time steps is shown as a function of the strength of directional motivated reasoning ϕ and the degree of under-reaction λ. The colorbar is in logarithmic scale. (a) Whole parameter space. (b) Zoom on bottom-right corner.

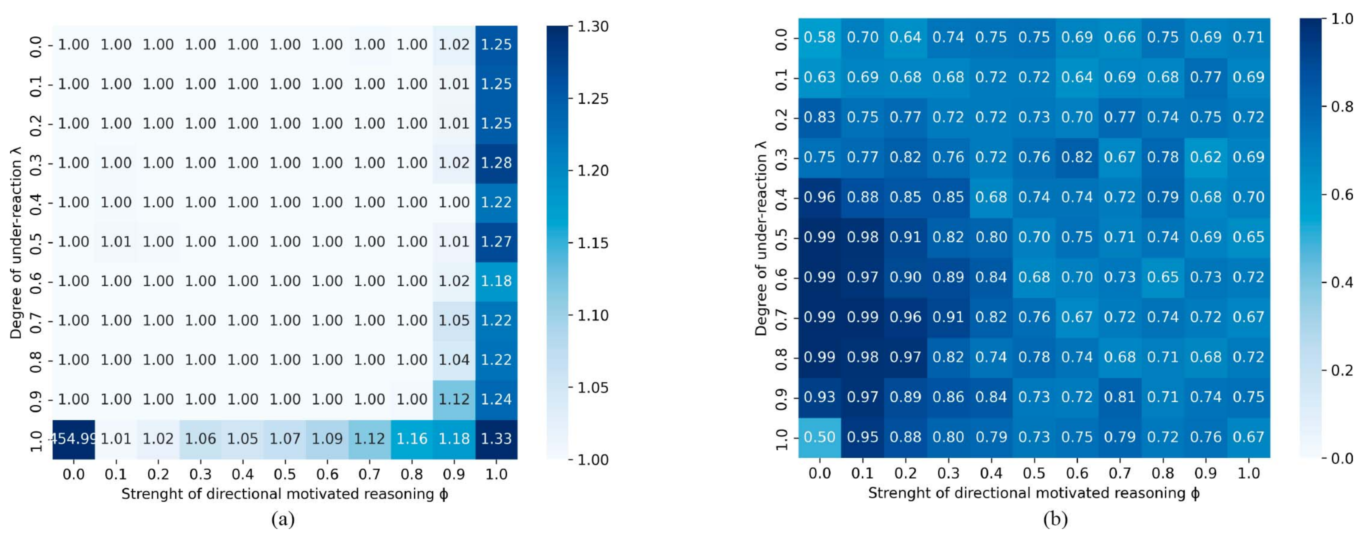

Fig. 4. One-lobbyist scenario. The effective number of clusters and the final average subjective probability are shown as functions of the strength of directional motivated reasoning ϕ and the degree of under-reaction λ in the presence of a lobbyist. The lobbyist supports the pessimistic model, has a budget $B = 10\,000$ for sending signals, and remains active for a time horizon of $T = 100$. (a) Effective number of clusters. (b) Average subjective probabilities.

changes opinion slowly is more susceptible to polarization due to confirmation bias.

From the point of view of average opinions, overall values stay around the middle of the opinion interval (0.5), indicating that even if consensus is on one of the two models (optimistic or pessimistic), the exact model changes from one simulation to another, due to the symmetry of the model definition. That means that the belief distribution turns out to be characterized by path-dependence. This also applies in the case of two opinion clusters, when the opinion of the larger group changes from one simulation to another.

Another basic property of the model is that, while the model reaches consensus generally very fast, the time required to reach the stable state is increasing in λ (as shown in Fig. 3). In particular, as λ approaches 1, nodes take a larger number of time steps to reach a stable configuration, because each of them is less sensitive to new information and under-reacts. The relation between the convergence time and ϕ is, however, not so linear. For low λ values, large ϕ seems to slow down convergence, i.e., if the individuals do not under-react, then confirmation bias slows down convergence as society polarizes. However, when λ is large, the opposite effect is visible: since individuals under-react, polarization is faster as confirmation bias increases: they are not attracted by peers from the opposite opinion groups and so the simulations end faster.

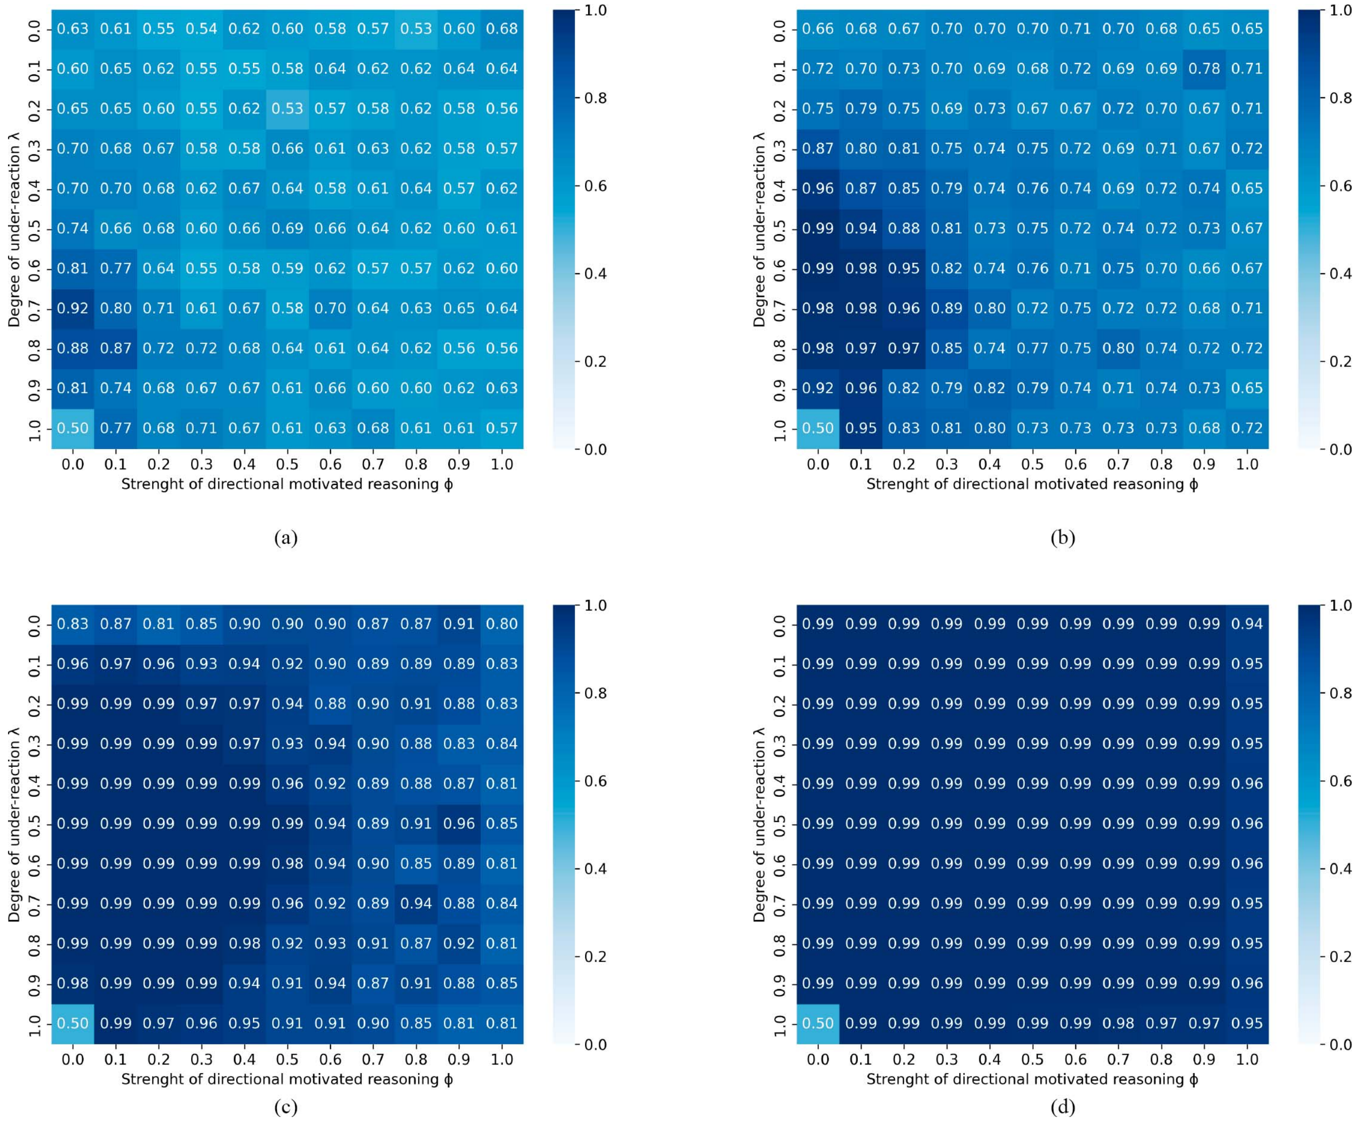

Fig. 5. Budget sensitivity analysis in the one-lobbyist scenario. The final average subjective probability is shown as a function of directional motivated reasoning ϕ and under-reaction λ for different values of the lobbyist's budget B. The lobbyist supports the pessimistic model and remains active for a time horizon of $T = 100$. (a) $B = 5000$. (b) $B = 10\,000$. (c) $B = 20\,000$. (d) $B = 40\,000$.

E. Numerical Experiments: One-Lobbyist Scenario

In the second setting, we introduce a first lobbyist agent that randomly sends signals to the network, with the aim of shifting the distribution of subjective probabilities toward the model it is supporting (namely, toward the pessimistic model).

In Fig. 4, we show the number of clusters and average subjective probabilities at the steady state. Compared to the baseline scenario, where outcomes were balanced between the two competing models, average probabilities are significantly higher over the whole parameter space, indicating that generally the lobbyist is able to steer part of the network toward supporting its own preferred model. To formally benchmark the lobbyist scenarios against the control, we test for each parametrization of the λ and ϕ grid whether the average opinion differs from the theoretical baseline value of 0.5, and we apply the Benjamini–Yekutieli procedure to control the false discovery rate across multiple tests; the results confirm that the null hypothesis cannot be rejected in the no-lobbyist case, while it is rejected for all relevant parametrizations in the one-lobbyist scenario. Furthermore, the effectiveness of the lobbyist is not globally homogeneous over the parameter space: in fact, the average probabilities are around 0.99 in the bottom-left corner region (indicating that, with these configurations, the lobbyist is successful in influencing the network almost always) but tend to be lower as we go farther off said region. Moreover, we see that the lobbyist, while managing to successfully increase the average subjective probability distributions, is unable to eliminate or reduce clustering in the bottom-right corner of the map (Fig. 4). If anything, in this region of the map an increase in polarization emerges, compared to the baseline, as the effective

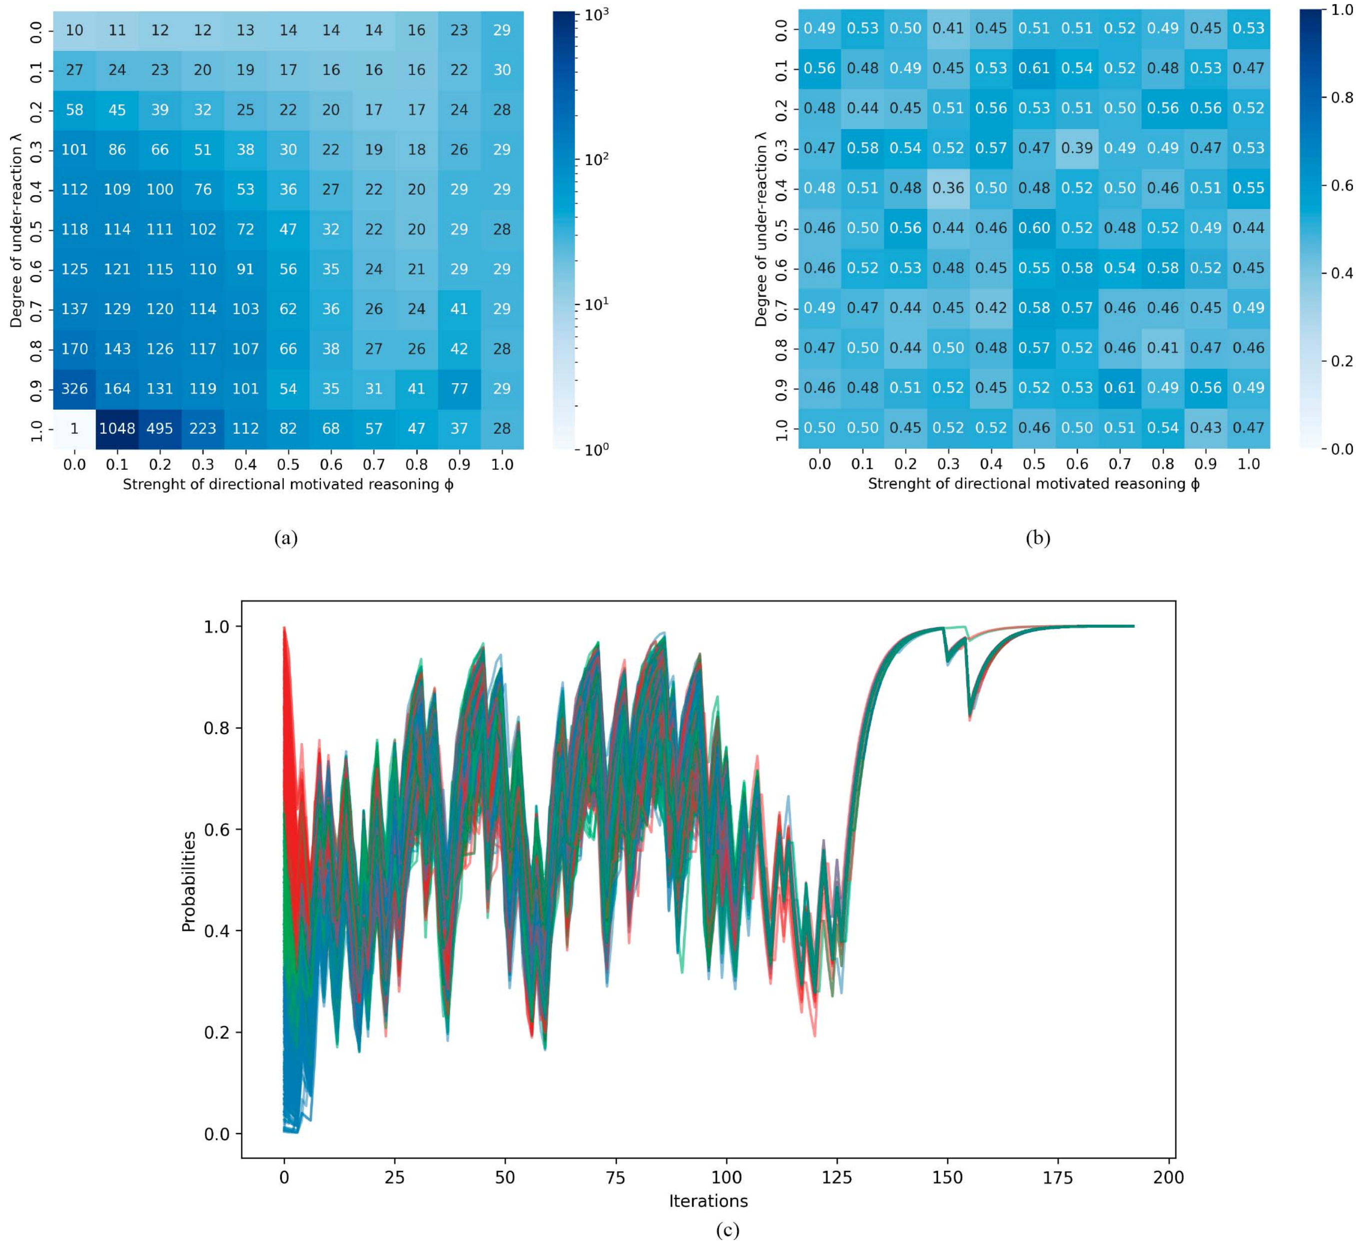

Fig. 6. Two-lobbyist scenario. The average number of time steps and the final average subjective probabilities are shown as functions of the strength of directional motivated reasoning ϕ and the degree of under-reaction λ. The two lobbyists support opposing models (one optimistic, one pessimistic), each with a budget of $B = 10\,000$ for influencing agents and an active time horizon of $T = 100$. (a) Avg. number of time steps. (b) Avg. subjective probabilities. (c) Example evolution plot.

number of clusters increases, i.e., the size of the second cluster increases.

These results indicate that a society with medium-high levels of under-reaction and low confirmation bias [dark blue area in Fig. 4(b)] is easily influenced by a lobbyist, while when confirmation bias is strong or the population is more dynamic (low under-reaction), the peer-effect can contrast the lobbyist. This can generate consensus on the opposite opinion in a minority of simulations, when confirmation bias is low, or increase polarization when confirmation bias is high. We denominate these two regimes the *lobbyist-effect* regime and the *peer-effect* regime.

To investigate further the two regimes, we performed *budget sensitivity analysis*, i.e., changed the budget of the lobbyist from 10 000 to 5000, 20 000, and 40 000, allowing them to reach 10%, 40%, and 80% of the population, respectively. In Fig. 5, we show the resulting average opinions. It is very clear that the region in which the lobbyist dominates the network (where it can impose its model almost certainly) expands linearly with the amount of resources available. With $B > 40\,000$, the lobbyist influence is complete and global over the parameter space. Even outside said region, however, average subjective probabilities are higher on average the higher the budget allocated to the lobbyist, indicating that the lobbyist still has a nonnegligible

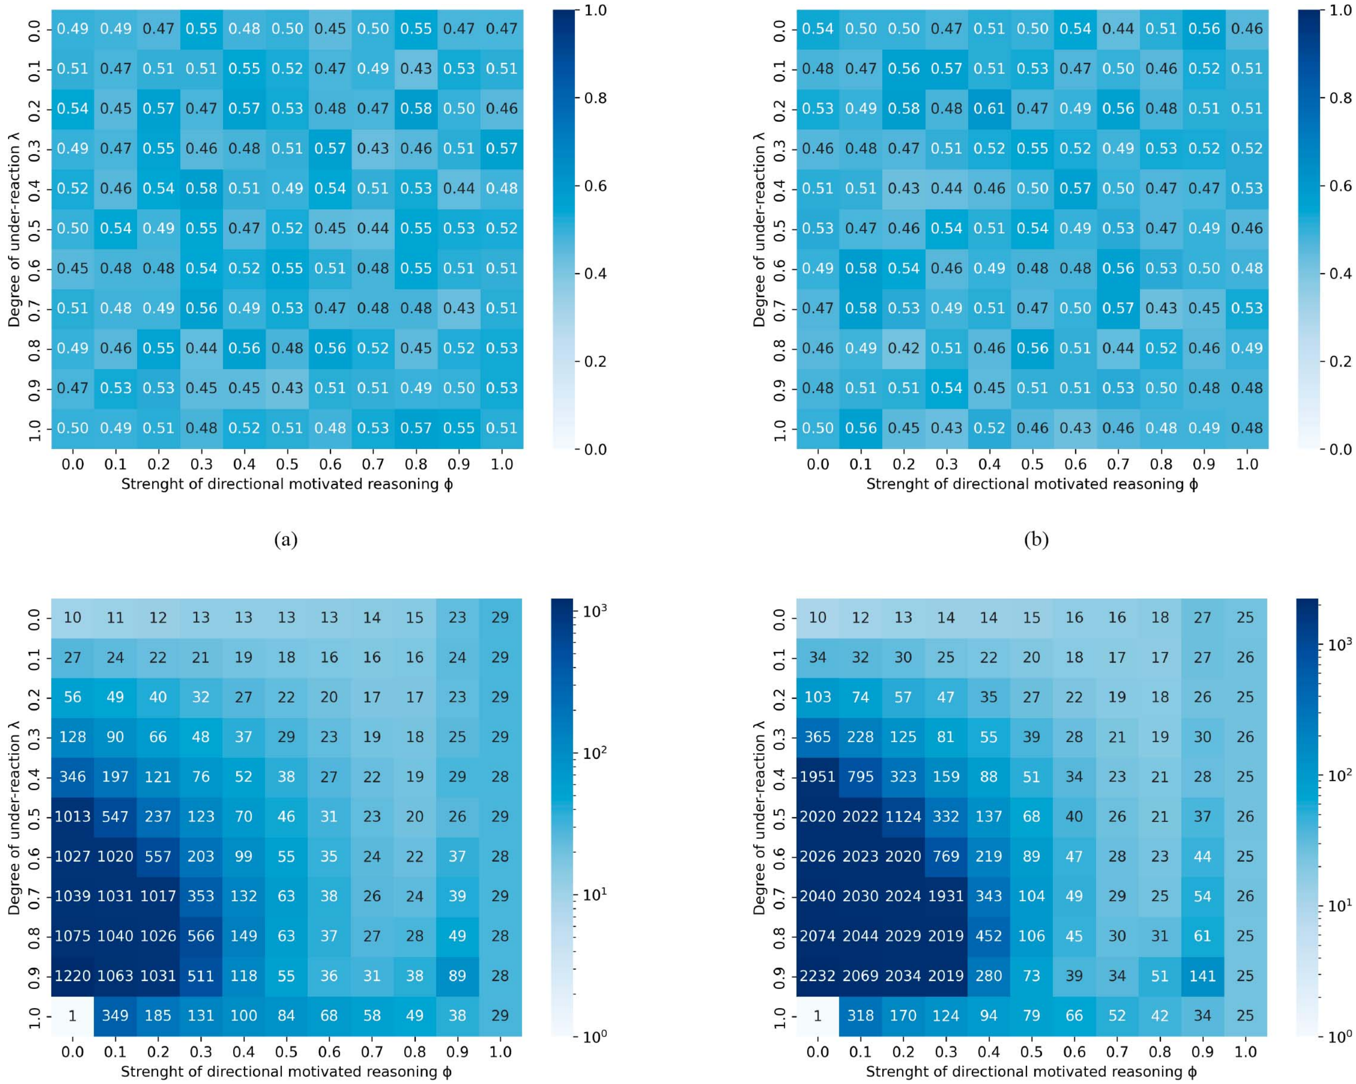

(a) (b) (c) (d)

Fig. 7. Long lobbyist time horizons. The average subjective probability of the final opinion distribution and the average number of time steps are plotted as functions of directional motivated reasoning ϕ and under-reaction λ for different values of T, the maximum number of steps during which lobbyists remain active. The two lobbyists support opposing models and interact with 20% of the agents at every simulation step. (a) Avg. subjective probabilities ($T = 1000$). (b) Avg. subjective probabilities ($T = 2000$). (c) Avg. number of time steps ($T = 1000$). (d) Avg. number of time steps ($T = 2000$).

impact on the network behavior regardless of the specific parameter configuration.

F. Numerical Experiments: Two-Lobbyists Scenario

In the third scenario under consideration, we bring into the picture a second lobbyist which has the same endowments and the same kind of strategies as the first one but that supports the competing model (namely, the "optimist" model).

In this configuration, average subjective probabilities across simulations seem to balance out, becoming comparable to the baseline scenario, as shown in Fig. 6(b). In fact, as the two lobbyists are identical except for the specific model supported, their competing influence seems to cancel out, and the average subjective probabilities gravitate around $\bar{p} = 0.5$, regardless of the parametrization.

The main difference, with respect to the baseline scenario, emerging in this case is that competition between lobbyists over influence on the network results in an increase in the number of time steps needed to reach the stable state, as shown in Fig. 6(a). In the lobbyist-effect area of the parameter space, where in the previous scenario the lobbyist was more efficient, the prolonged simulation time is due to the fact that the joint activity of the two lobbyists prevents the network from reaching an equilibrium as long as the lobbyists send signals. This confirms that, for that subset of configurations, it is easier for the external agent to project its influence on the network. Fig. 6(c) represents the evolution of a typical run of the model in one of those configurations. We note that, differently from what happens in the baseline scenario (Fig. 2), where the network converges to a stable state relatively quickly, the system undergoes persistent fluctuations for as long as the two lobbyists

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

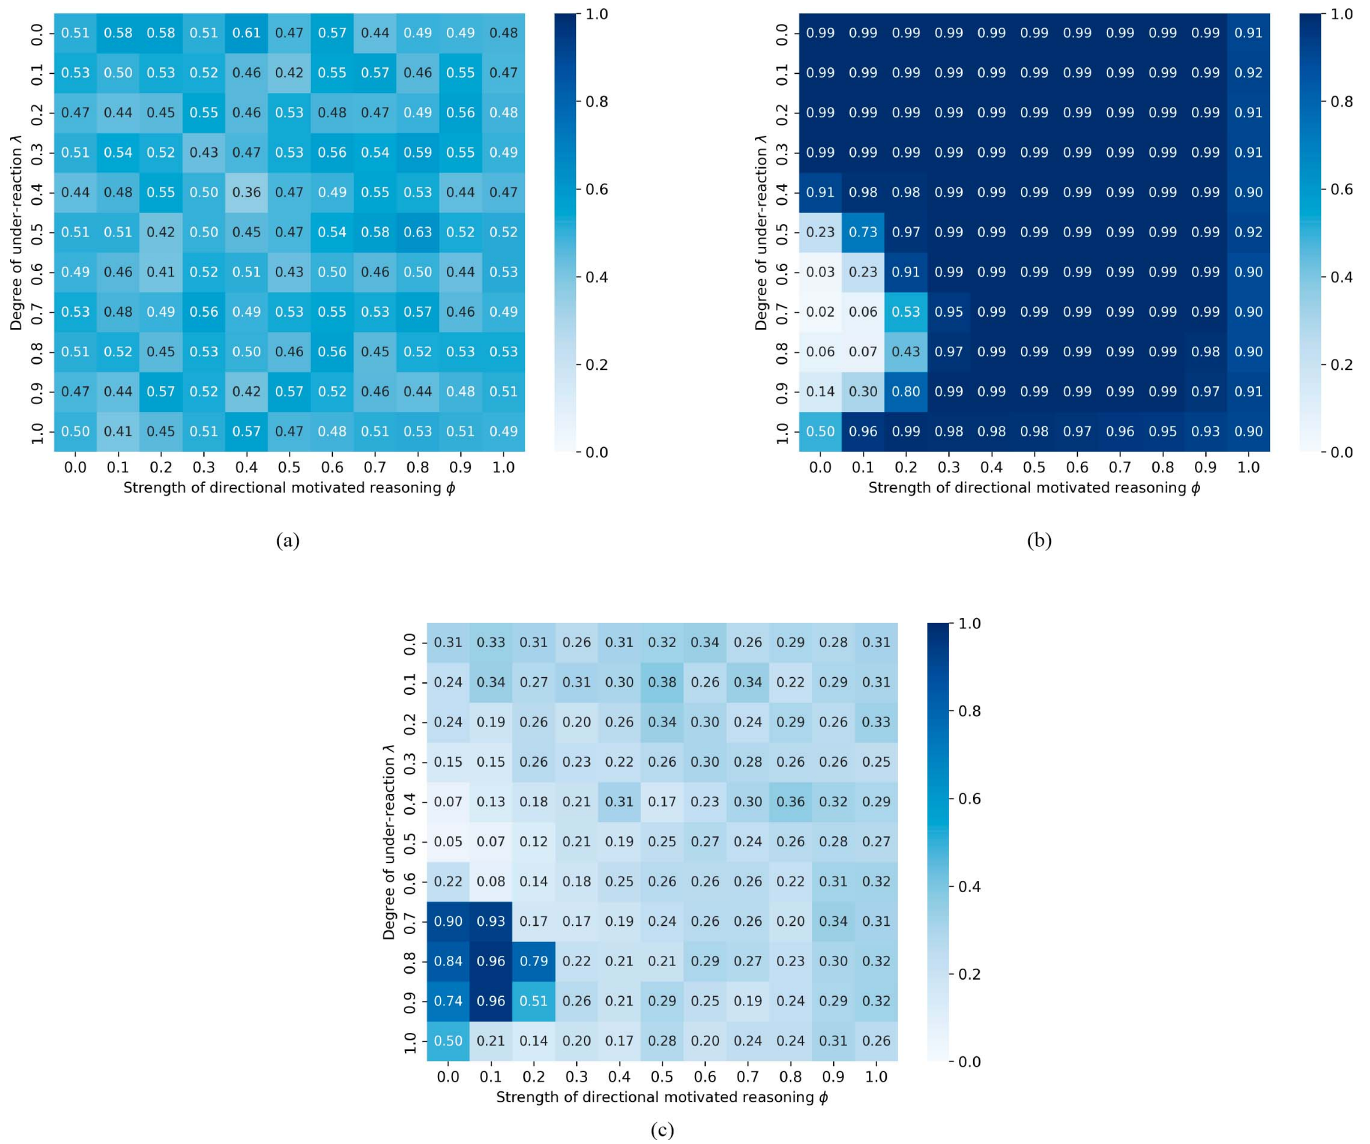

Fig. 8. Average steady-state subjective probabilities under different lobbying timing strategies. Panel (a) shows the outcome when both lobbyists adopt a frontloading strategy. Panel (b) corresponds to the case where the pessimist frontloads, and the optimist follows a uniform strategy. Panel (c) displays the opposite configuration, where the pessimist backloads. Values are averaged over 100 independent runs for each setting. Simulations are performed on a fully connected network with $N = 500$ agents, with each lobbyist equipped with a budget of $B = 10\,000$. (a) Scenario 1. (b) Scenario 2. (c) Scenario 3.

are active on the network. Shortly after they stop emitting signals (beyond 100 time steps), however, the network quickly converges to a consensus, in a similar dynamics to what we have seen in the baseline scenario. Outside the lobbyist-effect area, where the peer-effect prevailed in the baseline scenario, we see a less pronounced increase in the number of time steps required for convergence. The presence of the two opposing lobbyists makes consensus harder, because the lobbyists continue to attract individuals toward their preferred opinion; however, the population still converges because the peer-effect is stronger.

In this last scenario, we try to go a step further, and we perform a sensitivity analysis on the lobbyists' strategic horizon T with the aim of investigating how this parameter affects the model's properties. Specifically, we run several different batches of simulations by increasing the parameter T to 1000 and 2000. At the same time, we adjust the lobbyists' budget B accordingly, in order to maintain the same rate of signals per time-period of 20%, thus keeping the results comparable across specifications. As shown in Fig. 7, with a longer time horizon generally average opinion $\bar{p}$ remains about 0.5 in most configurations, suggesting that a longer horizon by the lobbyists does not radically change the outcome of the simulations, compared to shorter horizons or to the baseline scenario. The average number of time steps, instead, tends to increase linearly with T in the bottom-left corner of the parameter space, the region where the lobbyists are most effective, while it does not seem to change much in the other regions, which converge very fast due to peer effects. This confirms our conclusions for $T = 100$ above. It is worth noticing that

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

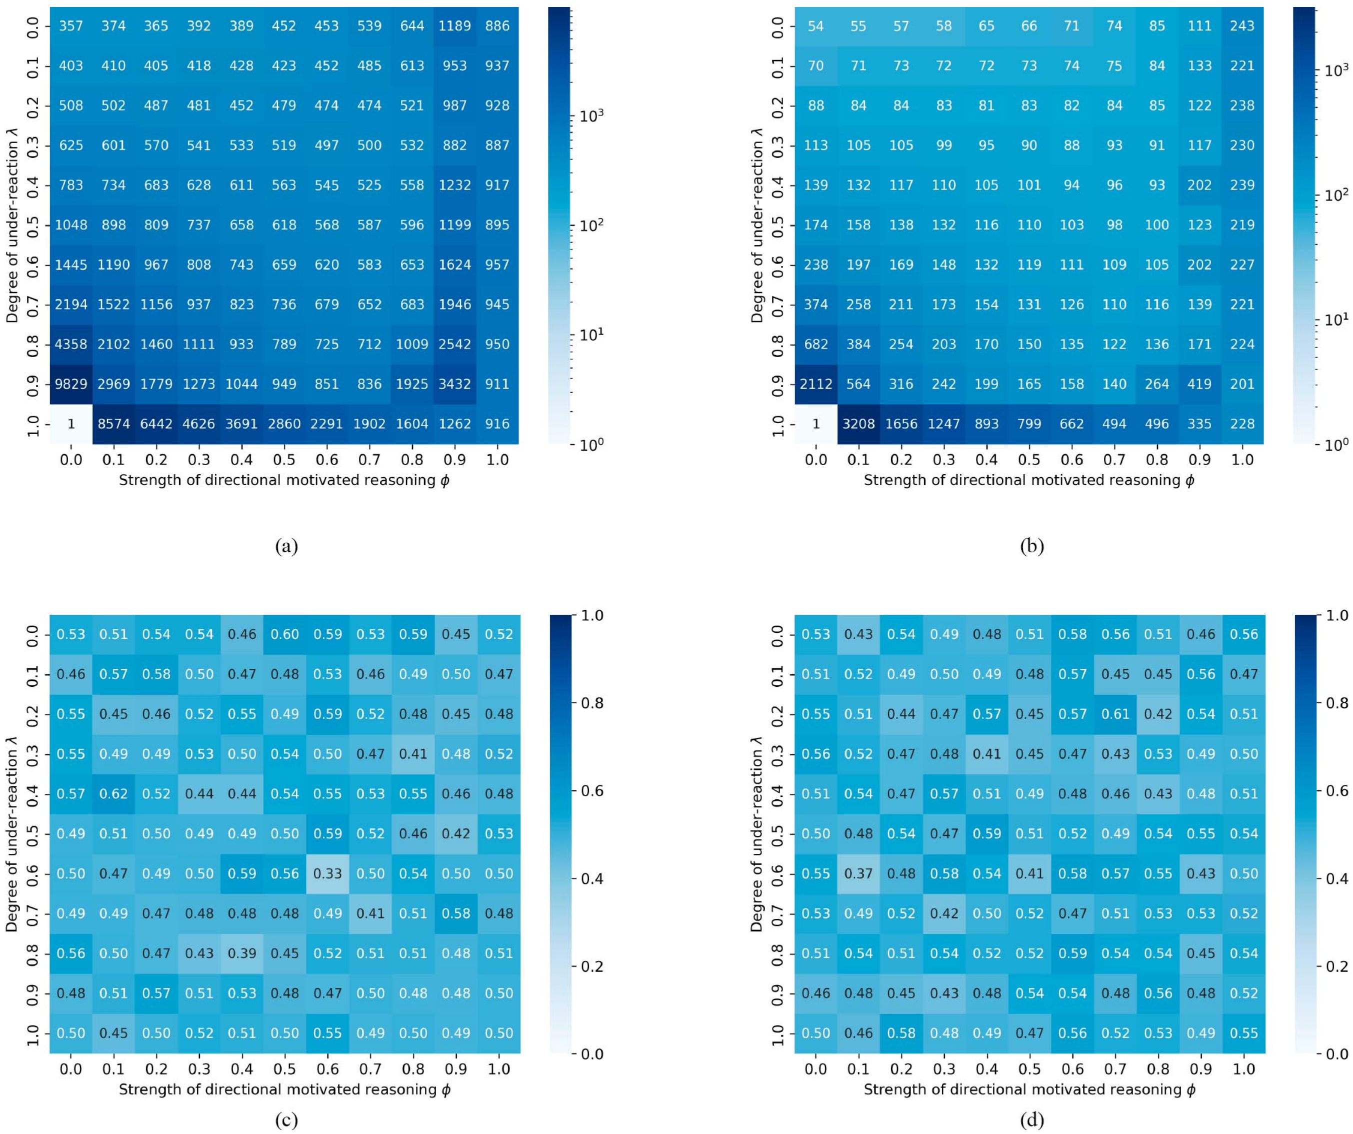

Fig. 9. Baseline (no-lobbyist) scenario, under scale-free networks. The average subjective probability of the final opinion distribution and the average number of time steps are plotted as functions of directional motivated reasoning ϕ and under-reaction λ for different values of m, the connectivity parameter. (a) Avg. number of time steps, $m = 20$. (b) Avg. number of time steps, $m = 100$. (c) Avg. opinions, $m = 20$. (d) Avg. opinions, $m = 100$.

while oscillations with stubborn influences also appear in linear models, here they emerge from motivated, signal-dependent responsiveness combined with time-staged strategic sending, yielding the regime patterns documented in Figs. 6 and 7.

G. Numerical Experiments: Timing Strategy Comparison

As a final exercise within the two-lobbyists scenario, we set up an experiment with the aim to investigate whether timing in the lobbyist strategies has a relevant impact on the outcome of the simulations. While insofar lobbyists have been assumed to spread out their signals uniformly over time, now we define two additional possible strategies, which we call "frontloading" and "backloading." In lobbying, frontloading means concentrating influence efforts at the earliest stages of the process. In the model, this is operationalized by the adoption, by the lobbyist agent, of a strategy matrix in which variable indicators for signals are concentrated in the first 20 (out of 100) time steps, in which the lobbyist exhausts its budget. Conversely, backloading implies holding back resources until later stages, in order to concentrate the effort closer to the steady state. This means, in our model, that the lobbyist adopts a strategy matrix in which variable indicators for signals are concentrated in the last 20 time steps.

For our experiment, we set up three different scenarios: in scenario 1, both the pessimist and the optimist lobbyists adopt a frontloading strategy; in scenario 2, the pessimist lobbyist adopts a frontloading strategy, while the optimist sticks to a random uniform strategy; in scenario 3, the pessimist lobbyist adopts a backloading strategy, while the optimist keeps using a random uniform strategy. We run 100 Monte Carlo simulations for each parameter configuration and for each scenario to

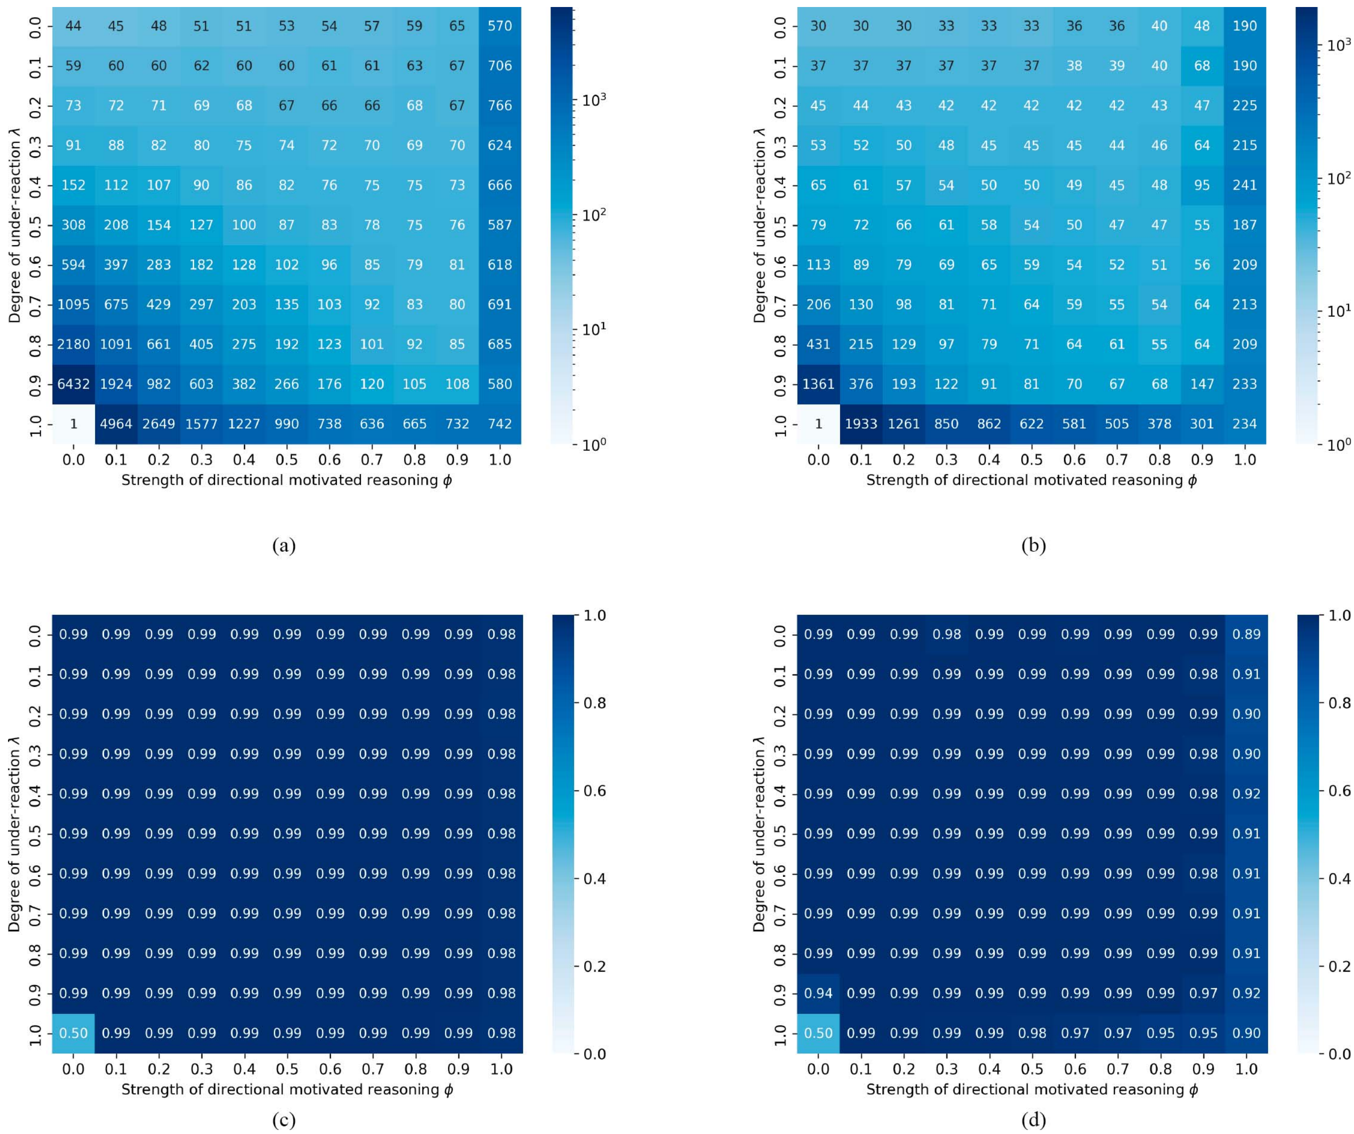

Fig. 10. One-lobbyist scenario, under scale-free networks. The average subjective probability of the final opinion distribution and the average number of time steps are plotted as functions of directional motivated reasoning ϕ and under-reaction λ for different values of m, the connectivity parameter. (a) Avg. number of time steps, $m = 20$. (b) Avg. number of time steps, $m = 100$. (c) Avg. opinions, $m = 20$. (d) Avg. opinions, $m = 100$.

assess the performance of the lobbyists under these different conditions.

Results are presented in Fig. 8. Expectedly, in the completely symmetric scenario 1 [Fig. 8(a)], average opinion outcomes are balanced across the parameter space, with neither lobbyist prevailing in the steady state across Monte Carlo simulations. The picture changes starkly in scenario 2 [Fig. 8(b)]. In this case, as earlier in the model, we see two distinct regions emerge in the parameter space. In the region characterized by lower under-reaction (or high confirmation bias), where beliefs by nodes are more flexible, and the peer effect is stronger, the frontloading (pessimist) lobbyist prevails over his opponent in almost all the simulations. Conversely, the frontloading strategy seems to be defeated by the random uniform optimist lobbyist when under-reaction is higher and motivated reasoning component is lower. In contrast, in scenario 3 [Fig. 8(c)], the situation is reversed: the pessimist lobbyist (now backloading) dominates in the high-λ, low-ϕ configurations but is dominated by the optimist elsewhere.

This behavior can be readily interpreted by considering the core mechanisms driving the evolution of the network over time in our model. Frontloading is particularly effective when agents are highly responsive to new information and early signals can trigger reinforcing dynamics, such as cascades and clustering effects. By influencing opinions early, lobbyists can increase the likelihood that the system evolves toward their preferred equilibrium and can potentially reduce the need for later interventions.

Backloading, by contrast, targets the final opinion distribution directly, which is often the variable of interest for the lobbyist. Backloading is advantageous when agents under-react to signals or when early fluctuations are likely to reverse over time.

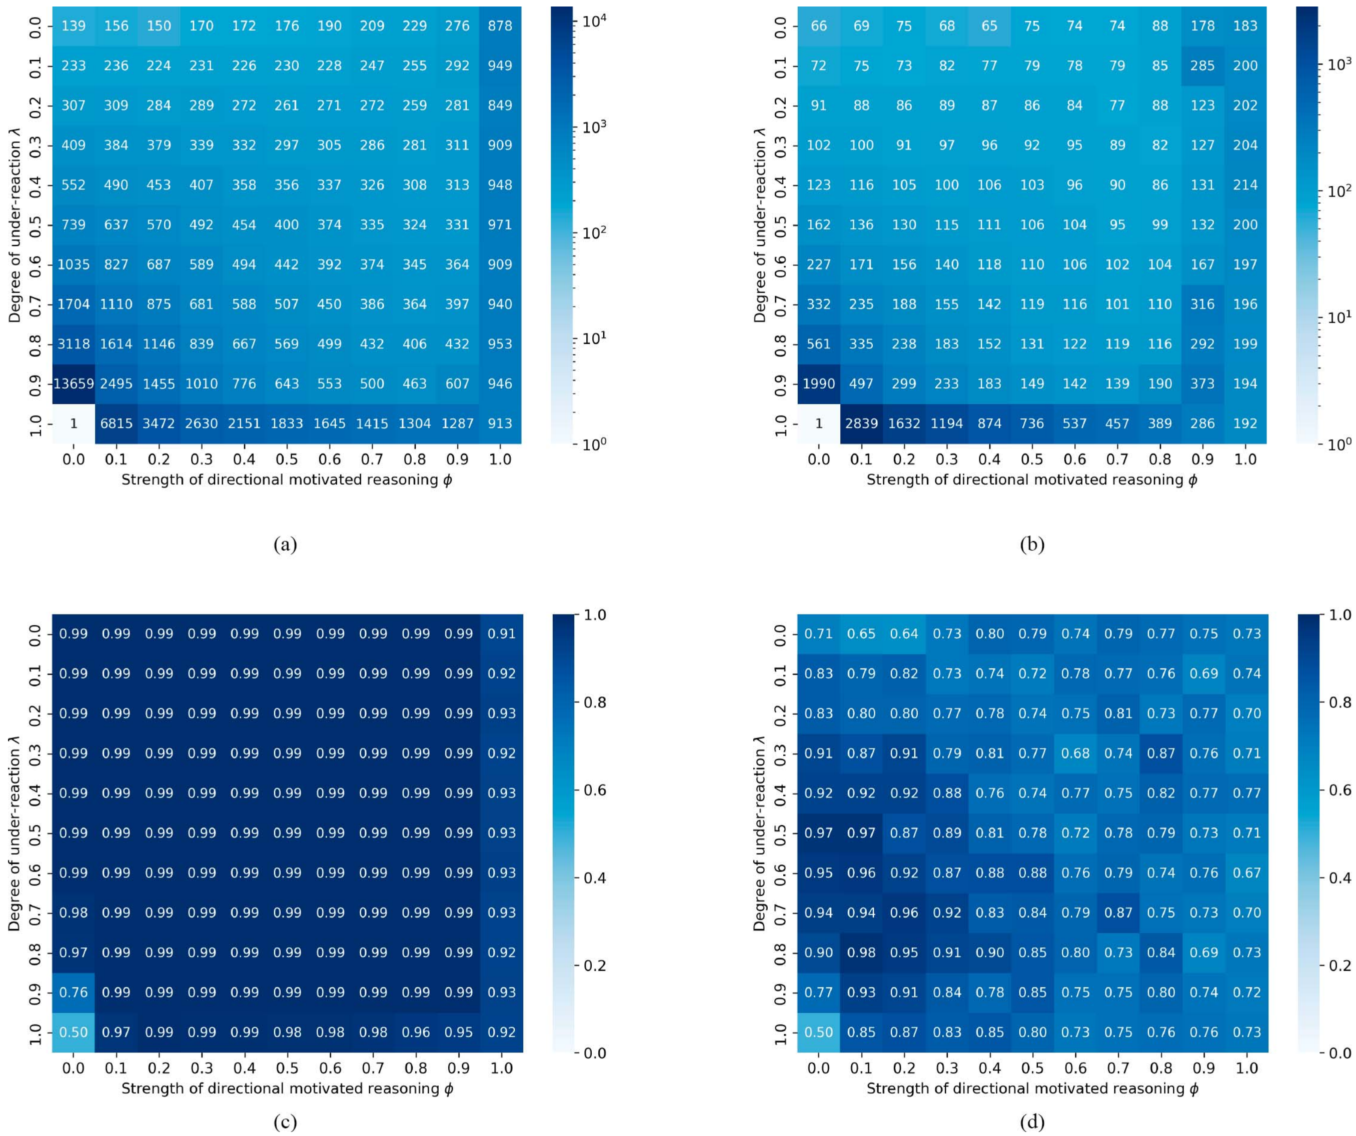

Fig. 11. One-lobbyist scenario, under scale-free networks. The average subjective probability of the final opinion distribution and the average number of time steps are plotted as functions of directional motivated reasoning ϕ and under-reaction λ for different values of m, the connectivity parameter. Lobbyist's budget is set to $B = 3000$. (a) Avg. number of time steps, $m = 20$. (b) Avg. number of time steps, $m = 100$. (c) Avg. opinions, $m = 20$. (d) Avg. opinions, $m = 100$.

By concentrating resources late, lobbyists maximize the impact of each signal on the final state, minimizing wasted influence and counteracting opposing signals. Backloading can be particularly effective in competitive environments where multiple lobbyists attempt to steer opinions, as late interventions can override prior influence and shift the final outcome.

Interestingly, the regions of parameter space where time-dependent strategies outperform the uniform strategy roughly match the regions in the single-lobbyist scenario where lobbyist influence is stronger (backloading prevails) and where peer effects are stronger (frontloading prevails).

H. Numerical Experiments: Network Topology

As a final exercise, we outline another version of our experiment, aimed at investigating the role of network structure in the model. During our exploration so far, the agents were communicating via a complete network; while this provides a useful benchmark, it reduces the empirical relevance of the model, since most real-world and online social networks exhibit scale free topologies [50]. To overcome this limitation, in the current section we run a battery of simulations introducing a Barabási–Albert type of network [51]. These are growing networks generated through preferential attachment, in which new nodes are more likely to connect to already highly connected nodes, producing a scale free degree distribution with emergent hubs. The density of the network is determined by a parameter m which dictates how many edges each new node adds. We run experiments from Sections IV-D (no lobbying) and IV-E (one lobbyist), with the new network topology, under two different parametrizations: a relatively denser network with $m = 100$, and a sparser network with $m = 20$. Results are shown in Figs. 9 and 10.

What emerges clearly from the plots is that, in both cases (but especially in the baseline scenario), simulations last longer on average, with the average number of time steps reaching up to 10 000, depending on the parametrization. This is to be expected, since with a sparser network structure information takes longer to propagate, and thus convergence comes about later on. The sparser the network, the longer the average running time, as it can be easily seen by comparing the two plots in the first row of Fig. 9. The presence of an external lobbyist communicating directly with the nodes, therefore unaffected by the network structure, significantly reduces the time needed for the network to converge, that nevertheless remains higher with respect to its complete-network counterpart. Average outcomes in terms of opinions in the baseline (no-lobbyist) scenario are similar to the complete-network case, where general symmetry is achieved, as average subjective probabilities across independent simulations converge toward 0.5. The case is different in the one-lobbyist scenario, where instead we see a largely increased effectiveness of the lobbyist in influencing the network. Irrespective of the connectivity parameter m, the lobbyist is now able to steer almost any run toward its desired outcome. This result can again be rationalized by the increased sparsity of the network: peer-to-peer communication gets increasingly scattered in a framework where nodes are less well-connected, so that they are no longer able to counterbalance the influence of an outside agent that can reach out to nodes regardless of the network structure. To substantiate this intuition, we run a final batch of simulations in which we reduce the budget allocated to the lobbyist (and therefore the density of its strategy matrix). If the increased effectiveness of the lobbyist is due to higher network sparsity, we should expect a corresponding reduction in the lobbyist budget to recover the two regimes across the parameter space, a "lobbyist-dominated" region in the south-west corner, and a "network-dominated" region elsewhere. Results are displayed in Fig. 11.

Decreasing the lobbyist budget results in an increase in the average duration of the simulation with respect to the results shown in Fig. 10, and in a reduction of the lobbyist's influence that is related to the network's average degree (itself a function of m). In the $m = 100$ scenario, we clearly recover the two-regime partition of the parametric space that we obtained with a complete network, an indication that the introduction of scale-free properties did increase, on average, the lobbyist's influence but did not change the qualitative properties of the model. Conversely, in the $m = 20$ scenario, the network is still not dense enough to effectively counterbalance the lobbyist's activity. In this case, the lower connectivity limits the diffusion of countervailing opinions across the network, allowing the lobbyist's influence to persist more effectively. As a result, the system exhibits lower sensitivity to changes in the behavioral parameters.

V. Conclusion and Future Work

In this work, we introduce a model of opinion formation in which agents update beliefs via a Bayesian-like rule in a social-learning setting, incorporating cognitive biases such as under-reaction and directionally motivated reasoning (confirmation bias). The model includes lobbyists that are characterized by strategies (a set of agents to target at each time step) and are budget constrained. In line with our scope, we study lobbying as an opinion-influence mechanism; the policy stage is exogenous and can be connected via a simple mapping without altering the core dynamics. We studied the model theoretically and numerically and observed the emergence of two polarized clusters when confirmation bias is large, even in the absence of lobbying. Moreover, in such a scenario, we also analytically proved that, under ex-ante homogeneity, a path-dependent consensus is asymptotically observed. A single active lobbyist using a uniform strategy (choice supported by the Nash equilibrium computed for a special case of the model) can attract the entire population to the opinion it promotes if the budget is large enough, if individuals exhibit low confirmation bias and moderate under-reaction, or if the social network is sparse. Otherwise, a polarized minority continues to coexist with the cluster that agrees with the lobbyist. We thus observe two regimes: in a part of the (λ, ϕ) space, the lobbyist can have a strong influence (lobbyist-effect area), while in the rest of the space the lobbyist influence is more reduced (peer-effect area). In the presence of two opposing lobbyists, the final steady states obtained are similar to those without lobbying, except that, in the lobbyist-effect area, convergence happens only after the lobbyists stop influencing the network. As long as they are active, the population continues to oscillate between different opinions without converging. This suggests that continuous lobbying can preclude consensus or convergence. Our oscillations are mechanistically distinct from those in linear stubborn-agent frameworks [32]: they require: 1) nonlinearity in responsiveness (3); and 2) endogenous, budget-constrained sending that injects temporally structured shocks (4)–(5).

Our model shares some similarities with classical models of opinion dynamics; however, exact mechanisms are different. The internal dynamics where an internal continuous weight is updated at discrete time steps is similar to bounded confidence models such as the Deffuant–Weisbuch [52] and Hegselmann–Krause [16] models. However, here, agents do not share the internal state with peers, but an external discrete opinion, more like in the CODA model [53]. Directional motivated reasoning, where agents react more strongly to opinions that are similar to their own, has a rationale similar to bounded confidence in the above-mentioned models; however, in our model, the effect is continuous rather than based on a discrete boundary. Indeed, a large ϕ leads to a lack of consensus, similar to bounded confidence models, but we only obtain two clusters. Moreover, when consensus is reached, it usually stays on one of the two opinions, with all internal weights close to the range limit, not in the middle like bounded confidence models. This is more realistic, and models individuals who actually make a decision in the end rather than remaining undecided and holding a moderate opinion ~0.5. In this regard, our model is more similar to the voter model [54] or to other discrete models. Moreover, this polarizing effect has also been seen in other models where extreme external information can cause minority extremist groups [23], [29]. When two symmetric opposing lobbyists are present, a rich behavior is observed. During the active time horizon of the lobbyists, inside the lobbyist influence area, we observe fluctuations of opinions between the two extreme opinions, with convergence as soon as the lobbyists become silent. Finally, our analysis of alternative timing strategies shows that the distribution of signals is not neutral in general: frontloading can be highly effective when agents are responsive and cascades amplify early interventions, while backloading dominates when agents under-react and late interventions can directly shape the final steady state. This demonstrates that the temporal structure of lobbying is itself an additional strategic dimension of influence, and that its effectiveness depends to some degree on the cognitive bias of the population.

Lobbyists are an important element of novelty in our model. While external influence has been studied through media fields [21], [22], [23], [25], [29] and zealot or stubborn-agent formulations [30], [32], these approaches typically model influence as constant, state-independent, or implemented through fixed opinions and linear averaging rules. In contrast, our framework models lobbying as a budget-constrained and time-staged strategic process, where signals are selectively allocated across agents and periods and enter through the same nonlinear, signal- and state-dependent updating mechanism that governs peer interaction. Because responsiveness depends endogenously on under-reaction and confirmation bias (3), the impact of influence varies across the cognitive parameter space, generating distinct regimes in which either lobbying forces dominate or peer effects prevail. In particular, polarization may persist when confirmation bias is strong and under-reaction limits cross-group influence, whereas lobbying can induce near-consensus when responsiveness allows signals to propagate effectively. The novelty of the approach therefore lies not in introducing external influence per se, but in embedding temporally structured, resource-constrained strategic signaling within a cognitively biased updating environment, which produces regime-dependent dynamics not reducible to standard zealot analogs [32].

An important simplification in our model is the symmetry of the pessimistic and optimistic model probabilities. We plan to explore nonsymmetric models in the future, and investigate polarization patterns in that case. Future work will also explore more complex lobbying strategies extracted from real data [55], [56]. Furthermore, we plan to integrate the model with statistical model-checking environments, such as MultiVeStA [57], [58], to gain further insight into its properties and dynamics. Another relevant extension concerns models in which beliefs gradually revert toward a baseline in the absence of reinforcing signals [59]. In the present framework, beliefs persist unless updated by new signals, which allows early interventions to be amplified by peer dynamics and explains the effectiveness of frontloading in certain regimes. Introducing a mean-reversion or attention-decay mechanism would likely weaken the persistence of early shocks and increase the relative importance of influence exerted closer to the evaluation horizon. Our timing results should therefore be interpreted conditional on belief persistence, and incorporating baseline reversion represents a promising direction for future research.

An additional promising direction for future research concerns relaxing the assumption that lobbyists commit ex-ante to a fixed time-staged strategy without observing intermediate belief realizations. In the current framework, lobbying strategies are predetermined and do not respond to the evolving distribution of opinions in the population. Allowing lobbyists to observe aggregate belief dynamics and adapt their signaling strategy in real time would introduce a feedback mechanism between influence and state variables. Such an extension would transform the problem into a dynamic strategic control setting and would be particularly relevant when lobbyists aim to target intermediate belief levels rather than extreme positions. Exploring adaptive lobbying under informational frictions represents a natural and important next step, building on the present model. At the same time, a complementary line of ongoing research develops adaptive and network-aware lobbying strategies within a distinct theoretical architecture in which strategic lobbying interacts endogenously with evolving belief states and network topology. Because such extensions introduce additional theoretical layers, they are treated separately in order to preserve the analytical transparency and mechanism-identification focus of the present contribution. Finally, while the present article focuses on the theoretical development of the model, ongoing work exploits data from a survey experiment within a structural belief updating framework [60] to estimate behavioral parameters. In future work, these estimates and data could be used to calibrate the model. At the present stage, nonetheless, the results of [60] already confirm that high levels of (overall) under-reaction are actually observed in the population.

References

[1] G. M. Grossman and E. Helpman, *Special Interest Politics.* Cambridge, MA, USA: MIT Press, 2001.

[2] F. R. Baumgartner, J. M. Berry, M. Hojnacki, B. L. Leech, and D. C. Kimball, *Lobbying and Policy Change: Who Wins, Who Loses, and Why.* Chicago, IL, USA: Univ. of Chicago Press, 2009.

[3] H. Klüver, "Informational lobbying in the European union: The effect of organisational characteristics," *West Eur. Politics*, vol. 35, no. 3, pp. 491–510, 2012.

[4] D. C. Lewis, "Advocacy and influence: Lobbying and legislative outcomes in Wisconsin," *Interest Groups Advocacy*, vol. 2, pp. 206–226, 2013.

[5] I. Schlichting, "Strategic framing of climate change by industry actors: A meta-analysis," *Environ. Commun.: A J. Nat. Cult.*, vol. 7, no. 4, pp. 493–511, 2013.

[6] L. C. Stokes and C. Warshaw, "Renewable energy policy design and framing influence public support in the United States," *Nat. Energy*, vol. 2, no. 8, pp. 1–6, 2017.

[7] M. H. DeGroot, "Reaching a consensus," *J. Am. Stat. Assoc.*, vol. 69, no. 345, pp. 118–121, 1974.

[8] A. V. Banerjee, "A simple model of herd behavior," *Quart. J. Econ.*, vol. 107, no. 3, pp. 797–817, 1992.

[9] D. Acemoglu and A. Ozdaglar, "Opinion dynamics and learning in social networks," *Dyn. Games Appl.*, vol. 1, no. 1, pp. 3–49, 2011.

[10] A. Jadbabaie, P. Molavi, A. Sandroni, and A. Tahbaz-Salehi, "Non-Bayesian social learning," *Games Econ. Behav.*, vol. 76, no. 1, pp. 210–225, 2012.

[11] L. G. Epstein, J. Noor, and A. Sandroni, "Non-Bayesian learning," *BE J. Theor. Econ.*, vol. 10, no. 1, pp. 1–20, 2010.

[12] F. Massari, "Under-reaction: Irrational behavior or robust response to model misspecification?" *Econ. Theory*, pp. 1–25, 2025, doi: 10.1007/s00199-025-01677-x.

[13] G. Bottazzi, D. Giachini, and M. Ottaviani, "Market selection and learning under model misspecification," *J. Econ. Dyn. Control*, vol. 156, 2023, Art. no. 104739.

[14] J. N. Druckman and M. C. McGrath, "The evidence for motivated reasoning in climate change preference formation," *Nat. Climate Change*, vol. 9, no. 2, pp. 111–119, 2019.

[15] P. Sobkowicz, "Whither now, opinion modelers?" *Front. Phys.*, vol. 8, 2020, Art. no. 587009.

[16] R. Hegselmann and U. Krause, "Opinion dynamics and bounded confidence models, analysis, and simulation," *J. Artif. Societies Social Simul.*, vol. 5, no. 3, 2002. Available: https://jasss.soc.surrey.ac.uk/5/3/2.html

[17] D. Centola and M. Macy, "Complex contagions and the weakness of long ties," *Amer. J. Sociol.*, vol. 113, no. 3, pp. 702–734, 2007.

[18] G. Deffuant, D. Neau, F. Amblard, and G. Weisbuch, "How can extremism prevail? A study based on the relative agreement interaction model," *J. Artif. Soc. Social Simul.*, vol. 5, no. 4, pp. 1–8, 2002.

[19] D. J. Watts and P. S. Dodds, "Influentials, networks, and public opinion formation," *J. Consum. Res.*, vol. 34, no. 4, pp. 441–458, 2007.

[20] S. Bikhchandani, D. Hirshleifer, and I. Welch, "Learning from the behavior of others: Conformity, fads, and informational cascades," *J. Econ. Perspect.*, vol. 12, no. 3, pp. 151–170, 1998.

[21] A. Das, S. Gollapudi, and K. Munagala, "Modeling opinion dynamics in social networks," in *Proc. 7th ACM Int. Conf. Web Search Data Mining*, 2014, pp. 403–412.

[22] T. Li and H. Zhu, "Effect of the media on the opinion dynamics in online social networks," *Physica A*, vol. 551, 2020, Art. no. 124117.

[23] A. Sîrbu, V. Loreto, V. D. Servedio, and F. Tria, "Opinion dynamics: Models, extensions and external effects," in *Participatory Sensing, Opinions and Collective Awareness*, Cham, Switzerland: Springer, 2017, pp. 363–401.

[24] A. Sîrbu, D. Pedreschi, F. Giannotti, and J. Kertész, "Algorithmic bias amplifies opinion fragmentation and polarization: A bounded confidence model," *PLoS One*, vol. 14, no. 3, 2019, Art. no. e0213246.

[25] V. Pansanella, A. Sîrbu, J. Kertesz, and G. Rossetti, "Mass media impact on opinion evolution in biased digital environments: A bounded confidence model," *Sci. Rep.*, vol. 13, no. 1, 2023, Art. no. 14600.

[26] G. Fagiolo and A. Roventini, "Macroeconomic policy in DSGE and agent-based models redux: New developments and challenges ahead," *J. Artif. Soc. Social Simul.*, vol. 20, no. 1, 2017. Available: https://www.jasss.org/20/1/1.html

[27] A. G. Haldane and A. E. Turrell, "Drawing on different disciplines: Macroeconomic agent-based models," *J. Evol. Econ.*, vol. 29, pp. 39–66, 2019.

[28] M. Bardoscia et al., "The impact of prudential regulation on the UK housing market and economy: Insights from an agent-based model," *J. Econ. Behav. Org.*, vol. 229, 2025, Art. no. 106839.

[29] A. Sîrbu, V. Loreto, V. D. Servedio, and F. Tria, "Opinion dynamics with disagreement and modulated information," *J. Stat. Phys.*, vol. 151, pp. 218–237, 2013.

[30] M. Mobilia, "Commitment versus persuasion in the three-party constrained voter model," *J. Stat. Phys.*, vol. 151, pp. 69–91, 2013.

[31] E. Dong and S. Marzen, "Cognitive biases can move opinion dynamics from consensus to signatures of transient chaos," *Am. J. Phys.*, vol. 92, no. 10, pp. 801–808, 2024.

[32] D. Acemoğlu, G. Como, F. Fagnani, and A. Ozdaglar, "Opinion fluctuations and disagreement in social networks," *Math. Oper. Res.*, vol. 38, no. 1, pp. 1–27, 2013.

[33] G. R. Moreno, E. Manino, L. Tran-Thanh, and M. Brede, "Zealotry and influence maximization in the voter model: When to target zealots?" 2020, *arXiv:2008.10012*.

[34] P. P. Klamser, M. Wiedermann, J. F. Donges, and R. V. Donner, "Zealotry effects on opinion dynamics in the adaptive voter model," *Phys. Rev. E*, vol. 96, no. 5, 2017, Art. no. 052315.

[35] C. Roser-Renouf and E. W. Maibach, "Strategic communication research to illuminate and promote public engagement with climate change," in *Change and Maintaining Change*, Cham, Switzerland: Springer, 2018, pp. 167–218.

[36] C. Xiao, L. Su, J. Bi, Y. Xue, and A. Kuzmanovic, "Selective behavior in online social networks," in *Proc. IEEE/WIC/ACM Int. Conf. Web Intell. Intell. Agent Technol.*, vol. 1, 2012, pp. 206–213.

[37] L. Muchnik, S. Aral, and S. J. Taylor, "Social influence bias: A randomized experiment," *Science*, vol. 341, no. 6146, pp. 647–651, 2013.

[38] R. M. Bond et al., "A 61-million-person experiment in social influence and political mobilization," *Nature*, vol. 489, no. 7415, pp. 295–298, 2012.

[39] S. Aral and D. Walker, "Creating social contagion through viral product design: A randomized trial of peer influence in networks," *Manage. Sci.*, vol. 57, no. 9, pp. 1623–1639, 2011.

[40] S. Vosoughi, D. Roy, and S. Aral, "The spread of true and false news online," *Science*, vol. 359, no. 6380, pp. 1146–1151, 2018. [Online]. Available: https://www.science.org/doi/abs/10.1126/science.aap9559

[41] W. Edwards, *Conservatism in Human Information Processing.* Cambridge, MA, USA: Cambridge Univ. Press, 1982, pp. 359–369.

[42] R. S. Nickerson, "Confirmation bias: A ubiquitous phenomenon in many guises," *Rev. General Psychol.*, vol. 2, no. 2, pp. 175–220, 1998.

[43] C. M. Massad, M. L. Hubbard, and D. Newtson, "Selective perception of events," *J. Exp. Social Psychol.*, vol. 15, no. 6, pp. 513–532, 1979.

[44] S. Gearhart, A. Moe, and B. Zhang, "Hostile media bias on social media: Testing the effect of user comments on perceptions of news bias and credibility," *Hum. Behav. Emerg. Technol.*, vol. 2, no. 2, pp. 140–148, 2020.

[45] G. Bottazzi, P. Dindo, and D. Giachini, "Long-run heterogeneity in an exchange economy with fixed-mix traders," *Econ. Theory*, vol. 66, no. 2, pp. 407–447, 2018.

[46] G. Bottazzi and D. Giachini, "Far from the madding crowd: Collective wisdom in prediction markets," *Quant. Finance*, vol. 19, no. 9, pp. 1461–1471, 2019.

[47] D. Giachini, "Rationality and asset prices under belief heterogeneity," *J. Evol. Econ.*, vol. 31, no. 1, pp. 207–233, 2021.

[48] G. Bottazzi and P. Dindo, "Drift criteria for persistence of discrete stochastic processes on the line," *J. Math. Econ.*, vol. 101, 2022, Art. no. 102696.

[49] J. F. Nash, "Equilibrium points in NF-person games," *Proc. Nat. Acad. Sci.*, vol. 36, no. 1, pp. 48–49, 1950. [Online]. Available: https://www.pnas.org/doi/abs/10.1073/pnas.36.1.48

[50] A. Mislove, M. Marcon, K. P. Gummadi, P. Druschel, and B. Bhattacharjee, "Measurement and analysis of online social networks," in *Proc. 7th ACM SIGCOMM Conf. Internet Meas.*, 2007, pp. 29–42.

[51] A.-L. Barabási and E. Bonabeau, "Scale-free networks," *Sci. Am.*, vol. 288, no. 5, pp. 60–69, 2003.

[52] G. Deffuant, D. Neau, F. Amblard, and G. Weisbuch, "Mixing beliefs among interacting agents," *Adv. Complex Syst.*, vol. 3, nos. 1–4, pp. 87–98, 2000.

[53] A. C. Martins, "Continuous opinions and discrete actions in opinion dynamics problems," *Int. J. Modern Phys. C*, vol. 19, no. 04, pp. 617–624, 2008.

[54] P. Clifford and A. Sudbury, "A model for spatial conflict," *Biometrika*, vol. 60, no. 3, pp. 581–588, 1973.

[55] G. Errichiello, P. M. Falcone, and L. Popoyan, "Navigating climate policy: The influence of lobbying trends and narratives in Europe," *Environ. Sci. Policy*, vol. 163, 2025, Art. no. 103974.

[56] G. Errichiello, L. Popoyan, M. Bonasia, and P. M. Falcone, "Composition and determinants of corporate climate lobbying: Evidence from Italy," *Corporate Social Responsibility Environ. Manage.*, vol. 33, no. 1, pp. 189–216, 2026.

[57] A. Vandin, D. Giachini, F. Lamperti, and F. Chiaromonte, "Automated and distributed statistical analysis of economic agent-based models," *J. Econ. Dyn. Control*, vol. 143, 2022, Art. no. 104458.

[58] S. Sebastio and A. Vandin, "MultiVeStA: Statistical model checking for discrete event simulators," in *Proc. 7th Int. Conf. Perform. Eval. Methodologies Tools, VALUETOOLS*, Torino, Italy, Dec. 10–12, 2013. New York, NY, USA: ACM, 2013, pp. 310–315.

[59] H. L. Van der Maas, J. Dalege, and L. Waldorp, "The polarization within and across individuals: The hierarchical ising opinion model," *J. Complex Netw.*, vol. 8, no. 2, 2020, Art. no. pcnaa010.

[60] D. Giachini, G. Rossello, and L. Ciambezi, "A survey experiment on climate change beliefs, opinions, and actions," Sant'Anna School of Advanced Studies, Pisa, Tech. Rep. 08/2026, 2026.

Daniele Giachini received the Ph.D. degree in economics from Sant'Anna School of Advanced Studies, Pisa, Italy, in 2015.

He is currently an Associate Professor with the Institute of Economics, Sant'Anna School of Advanced Studies, where he also served as a Research Fellow and an Assistant Professor. His research interests include financial economics, mathematical economics, models of financial markets with heterogeneous agents, market selection, statistical learning, applications of information theory, and statistical analysis of simulated agent-based models.

Leonardo Ciambezi received the Ph.D. degree in economics, with a dissertation on inflation as a microfounded emerging phenomenon, from Université Côte d'Azur, Nice, France, in 2025.

He is currently a Research Fellow with Scuola Superiore Sant'Anna, Pisa, Italy. His research interests include computational economics, business cycle theory, and macroeconomic dynamics, both from theoretical and empirical perspectives.

Verdiana Del Rosso received the master's degree in mathematics from the University of Camerino, Camerino, Italy, in 2015.

She is currently a Research Fellow in computer science with the University of Camerino, where she is working on the ALMONDO project to develop a user-friendly simulation tool. She defended her Ph.D. thesis in 2020, focusing on modeling and designing an optimal torque distribution strategy for two electric motors in an all-wheel-drive electric motorcycle in collaboration with Visionarprogetti, Osimo, Italy. Since then, she has been researching mechanical anomaly detection systems for electric motors in household appliances, using vibrational signals and machine learning techniques.

Fabrizio Fornari (Member, IEEE) received the M.Sc. and Ph.D. degrees in computer science from the University of Camerino, Camerino, Italy, in 2013 and 2018, respectively.

Since 2025, he has been an Assistant Professor in computer science with the University of Camerino. Since 2017, he has been a member of the PROcesses and Services Laboratory (PROS Lab, University of Camerino, Camerino, Italy), where he conducts his research activities. He held postdoctoral positions at the University of Camerino. His research interests and teaching activities include formal methods, business process management, software engineering, Internet of Things, digital twins, and applied game design.

Prof. Fornari was the Local Coordinator of the ALMONDO PRIN 2022 PNRR project (2023–2026).

Valentina Pansanella received the Ph.D. degree in data science from Scuola Normale Superiore, Pisa, Italy, in 2024.

She is currently a Researcher with ISTI-CNR, Pisa, on the SoBigData.it project. She held a past position in the ALMONDO project at University of Pisa, Pisa, Italy. Her research interests include the study and development of opinion diffusion models and, more broadly, the societal impact of social and mass media.

Lilit Popoyan received the Ph.D. degree in economics in 2016 from Sant'Anna School of Advanced Studies, Pisa, Italy, where her doctoral research focused on the complex system dynamics of interactions between financial regulation and monetary policy.

She is currently an Associate Professor in business analytics and innovation with the School of Business and Management, Queen Mary University of London, London, U.K. She is also an Associate Researcher with Sant'Anna School of Advanced Studies and one of the Founding Members of the ABM4Policy Research Group. She was an Assistant Professor in economic policy with the University of Naples "Parthenope," Naples, Italy. Her research places strong emphasis on policy-relevant quantitative analysis of the nexus between macroeconomic policy, financial regulation, financial stability, sustainable finance and climate change, production networks, and macroeconomic dynamics.

Alina Sîrbu received the Ph.D. degree in computational biology from Dublin City University, Dublin, Ireland, in 2011.

She is currently an Associate Professor in computer science with the University of Bologna, Bologna, Italy. She held postdoctoral research positions with the Institute for Scientific Interchange, Turin, Italy, and the University of Bologna, and was an Assistant and Associate Professor with the University of Pisa, Pisa, Italy. She was a Visiting Assistant Professor in computer science with New York University Shanghai, Shanghai, People's Republic of China, in 2014, and with New York University Abu Dhabi, Abu Dhabi, United Arab Emirates, in 2017. She has been involved in several interdisciplinary National and European projects, among which she was a National Coordinator of the MEDICA PRIN 2022 project (2023–2025) and a Local Coordinator of the HumMingBird H2020 European project (2019–2023) and the ALMONDO PRIN 2022 PNRR project (2023–2026). Her research interests include complex systems modeling, data science, and machine learning, applied to biomedical, social, and technical systems.

Open Access funding provided by 'Scuola Superiore Sant'Anna' within the CRUI CARE Agreement