

Innovation diffusion dynamics toward long-term behavioral shifts

L. Piccinin, V. Breschi *Member, IEEE*, C. Ravazzi, F. Dabbene, and M. Tanelli, *Senior Member, IEEE*

Abstract—Sustainable technologies and services can play a pivotal role in the transition to “greener” habits. Their widespread adoption is thus crucial, and understanding how to foster this phenomenon in a systematic way could have a major impact on our future. With this in mind, in this work we propose an extension of the Friedkin-Johnsen opinion dynamics model toward characterizing the long-term impact of (structural) fostering policies. We then propose alternative nudging strategies that target a trade-off between widespread adoption and investments under budget constraints, showing the impact of our modeling and design choices on inclination shifts over a set of numerical tests.

Index Terms—Emerging control applications, constraint control, optimal control.

I. INTRODUCTION

PROMOTING the adoption of new green technologies and services has become increasingly important, especially as society faces the pressing challenges of climate change [1]. To achieve widespread adoption, policymakers and stakeholders can strategically rely on *on-off interventions* (e.g., first-time user discounts for sharing mobility services), promoting first-hand experiences of the benefits of a service or a technology that can, nevertheless, have only a short-term impact on individual habits [2], and *systemic policies* (e.g., building dedicated parking spaces for sharing vehicles), progressively shaping available infrastructures or regulatory system to facilitate the acceptance of new (green) solutions in the long run [3]. However, the effectiveness of these strategies is often hampered (and, thus, the widespread adoption of these green solutions slowed) by changes that the use of new technologies can induce in established habits due to the natural tendency of people to resist changes in their long-standing routines [4], as well as social dictates [5]. In this context, opinion dynamics and control theory can be crucial in systematically characterizing and harnessing the interplay between individual needs, social dictates, and interventions on

personal choices [5], making control strategies key tools in designing human-centered policies to shape a more sustainable future [6]. Several mathematical models have been proposed in the literature to describe opinion formation in social networks, from DeGroot’s weighted-average-based model [7] to the Friedkin-Johnsen (FJ) model [8], which introduces agents with persistent biases. These models have then been extended and applied to different contexts. For instance, the FJ model is at the core of the deterministic, dynamic game model proposed in [9], which further studies how the stubbornness of agents not influencing others (there denoted as passive agents) affects opinion dynamics over a network. However, these models do not consider how external interventions shape opinion dynamics, a characterization that is not necessary for their conventional applications. Therefore, traditional opinion dynamics models cannot be used for the evaluation and design of interventions. Only a few recent studies (see [10]–[12]) have explored models that embed the effects of external interventions into opinion dynamics to design policies to nudge changes in individual inclinations systematically. Nonetheless, despite considering the effect of external interventions, the opinion dynamics model adopted in [11] was designed with a focus on on-off incentives. Indeed, this model tends to asymptotically “forget” changes in inclinations induced by nudging policies, with opinions asymptotically converging to their open-loop values. Therefore, it does not allow the modeling of the impact of systemic policies, making it unsuitable for assessing the potential long-term impact of such policies. In this work, we explicitly model these effects by extending the FJ model with controlled inputs having saturated integral dynamics. Meanwhile, in [12] opinions evolve according to a continuous-time weighted averaging with neighbors (i.e., the DeGroot model) between marketing campaigns, while they are driven by discrete-time impulsive updates during marketing interventions. While such interventions control the trade-off between social imitation and personal inclination to reach a target opinion, our work rather aims at shifting individual inclinations without altering that balance. Similarly to [11], but assuming access only to individuals’ average inclinations over time, we then use the proposed model to devise different strategies for the design of (centralized) fostering policies aimed at achieving an “optimal” *trade-off* between the average adoption of a new technology/service (referred to as *social benefit*) and the investment made (i.e., the policy’s *cost*), under budget constraints. This choice leads to a constrained allocation problem, the solution of which allows us to (empirically) analyze the impact of budget-constrained inputs in closed loop.

This work is funded by the European Union – Next Generation EU, Mission 4, Component 1, under the PRIN project TECHIE: “A control and network-based approach for fostering the adoption of new technologies in the ecological transition” Cod. 2022KPHA24, CUP Master: D53D23001320006, CUP: B53D23002760006.

Lisa Piccinin and Mara Tanelli are with Politecnico di Milano, 20133, Milan, Italy (e-mail: name.surname@polimi.it).

Valentina Breschi is with the Eindhoven University of Technology, 5600 MB, Eindhoven, The Netherlands (e-mail: v.breschi@tue.nl).

Chiara Ravazzi and Fabrizio Dabbene are with the Institute of Electronics, Computer and Telecommunication Engineering, National Research Council of Italy (CNR-IEIIT), 10129, Turin, Italy (e-mail: name.surname@cnr.it).

Outline: Section II introduces the proposed opinion dynamics model, whose properties are then analyzed in Section III. This model is used in Section IV to introduce budget-constrained policy design strategies, whose impact on individual opinions is analyzed through a numerical example in Section V. The paper ends with some final remarks and directions for future work.

Notation: We denote with \mathbb{N} , \mathbb{N}_0 and \mathbb{R}_+ , the set of natural numbers, the set of natural numbers including zero, and the set of positive real numbers, respectively. Given any vector $x \in \mathbb{R}^n$ and matrix $A \in \mathbb{R}^{m \times n}$, their transposes are denoted as x^\top and A^\top , respectively, while the inverse of $B \in \mathbb{R}^{n \times n}$ is given by B^{-1} and its spectral radius is denoted as $\rho(B)$. The positive (non-negative) definite matrix A is indicated as $A \succ 0$ ($A \succeq 0$). Meanwhile, $\|x\|_2$ and $\|A\|_2$ denote their 2-norms, $\|x\|_B^2 = x^\top Bx$, $x_i \in \mathbb{R}$ denotes the i -th component of x and $A_{ij} \in \mathbb{R}$ indicates the element of A in position (i, j) . We denote by $\mathbf{1}$ and $\mathbf{0}$ vectors of ones and zeros (of appropriate dimensions) while we indicate identity matrices with I . Given a random vector $x \in \mathbb{R}^n$, $\mathbb{E}[x]$ denotes its expected value. The logical operator *or* is indicated as \vee .

II. MODELING LONG-TERM ATTITUDE SHIFTS

Consider an influence network with $N \in \mathbb{N}$ agents that we formally characterize with a directed weighted graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, P)$. Here, \mathcal{V} represents the set of N agents, \mathcal{E} indicates the existence of bonds between them, i.e., agent $w \in \mathcal{V}$ influences agent $v \in \mathcal{V}$ if $(v, w) \in \mathcal{E}$, while the strength of mutual influences is dictated by the non-negative, row-stochastic matrix $P \in \mathbb{R}^{N \times N}$ satisfying

$$P_{vw} > 0, \forall (v, w) \in \mathcal{E}, \quad \sum_{w \in \mathcal{V}} P_{vw} = 1, \forall v \in \mathcal{V}. \quad (1)$$

Meanwhile, the agents' susceptibility to social influences is captured by a diagonal matrix $\Lambda \in [0, 1]^{N \times N}$, where each diagonal entry $\lambda_i \in [0, 1]$ represents¹ the susceptibility of the i -th agent to the opinions of others, for $i = 1, \dots, N$. As common in opinion dynamics [13], we make the following technical assumption about these susceptibilities.

Assumption 2.1: For every node $v \in \mathcal{V}$, there exists a path from $v \in \mathcal{V}$ to a node $w \in \mathcal{V}$ such that $\lambda_w < 1$.

Agents are set apart by their *inherent biases* $u^o \in [0, 1]^N$ to the new technology/services, which we assume are the features that nudging policies can modify. In particular, the closer $u_v^o \in [0, 1]$ is to 1, the more inherently well-disposed the v -th agent is to the new technology/service. Along with their biases, agents' are also characterized by their latent inclination to adoption at each time instant $t \in \mathbb{N}_0$, here assumed to be collected in a state $x(t) \in [0, 1]^N$, with $x_v(t) \in [0, 1]$ closer to 1 indicating a positive attitude of the v -th agent toward the technology/service of interest. The evolution of this variable over time is here characterized by

$$x(t+1) = \Lambda P x(t) + (I - \Lambda) \max\{\mathbf{0}, \min\{u(t) + u^{nc}(t), \mathbf{1}\}\}, \quad (2a)$$

where $u^{nc}(t)$ characterizes the effect of stochastic (short-term) fluctuations in individual inclination to adoption due to

¹ $\lambda_i = 1$ indicates that the opinion of the i -th agent is driven by peer influence, while $\lambda_i = 0$ means the agent's opinion is unaffected by that of its peers.

external, uncontrollable factors (e.g., public transport strikes or bad weather in the context of bike sharing) on changes in individual inclinations at time $t \in \mathbb{N}_0$ (as in [11]). To reflect this modeling assumption, the uncontrollable inputs thus satisfy the following.

Assumption 2.2: The uncontrollable inputs in $\{u^{nc}(t)\}_{t \in \mathbb{N}_0}$ are zero-mean, i.i.d. random vectors bounded within the interval $[-\delta \mathbf{1}, +\delta \mathbf{1}]$, with $0 \leq \delta \leq u_v^o$ for all $v \in \mathcal{V}$.

Note that, due to the condition on δ , we ultimately assume that $u^{nc}(t)$ causes only (very) slight short-term changes in individual opinions at all $t \in \mathbb{N}_0$, since $|u_v^{nc}(t)| \leq \delta$ for all $t \in \mathbb{N}_0$ and $v \in \mathcal{V}$. Meanwhile, $u(t)$ encodes the impact of the initial individual bias and the cumulative ones of policy actions enacted by a policymaker or a stakeholder until $t \in \mathbb{N}_0$. Specifically, we describe the evolution of $u(t)$ with the following (simplistic) cumulative dynamics:

$$u(t+1) = u(t) + u^c(t), \quad \forall t \in \mathbb{N}, \quad \text{with } u(0) = u^o, \quad (2b)$$

where $u^c(t)$ is a non-negative controlled input modeling actions that the policymaker or stakeholders can undertake (and adjust) over time to nudge a shift in individual preferences toward a new technology/service.

Remark 2.3 (The validity of (2)): Including the *max/min* in (2a) ensures that $x(t)$ is well defined $\forall t \in \mathbb{N}_0$, i.e., that $x(t) \in [0, 1]^N$ at all time instants. \square

It is worth pointing out that, when no controlled policy is deployed to nudge the acceptance of a new technology/service, the latent state's expected value asymptotically coincides with that of the standard FJ model, as formalized next.

Lemma 2.4 (Control-free mean asymptotic opinions): Let Assumption 2.1 hold, $x(0) \in [0, 1]^N$ and $u^c(t) = 0$ for all $t \in \mathbb{N}_0$. Then, the latent state's expected value satisfies

$$\mu_\infty = \lim_{t \rightarrow \infty} \mu(t) = \lim_{t \rightarrow \infty} \mathbb{E}[x(t)] = (I - \Lambda P)^{-1} (I - \Lambda) u^o, \quad (3)$$

with $\mu_\infty \in [0, 1]^N$.

Proof: Since $u^c(t) = 0$ for all $t \in \mathbb{N}_0$, according to (2b) then $u(t+1) = u(t) = u^o \in [0, 1]^N$ for all $t \in \mathbb{N}_0$. In turn, it thus straightforwardly follows from (2) that

$$x(t+1) = \Lambda P x(t) + (I - \Lambda)(u^o + u^{nc}(t)), \quad \forall t \in \mathbb{N}_0,$$

and, accordingly, that

$$\mu(t+1) := \mathbb{E}[x(t+1)] = \Lambda P \mu(t) + (I - \Lambda) u^o, \quad \forall t \in \mathbb{N}_0. \quad (4)$$

The steady-state result in (3) straightforwardly results from the same reasoning in [13], concluding the proof. \blacksquare

Therefore, under Assumption 2.1, the expected latent opinions converge to a profile that is a convex combination of the initial inclinations of the agents driven by the strength of their mutual bonds and their susceptibility.

Remark 2.5: The result in Lemma 2.4 indicates that persistent changes in opinions can be achieved by acting on individual biases, supporting our choice of designing fostering policies directly targeting a change in u^o and indirectly leveraging social imitation, while not changing the set of initial adopters nor changing the features of social bonds.

III. A CLOSER LOOK AT THE MODEL'S PROPERTIES

We now analyze the properties of the opinion dynamics model in (2) for two classes of interventions, namely static

and feedback policies. In both cases, according with Assumption 2.2, the policies $\{u^c(t)\}_{t \in \mathbb{N}_0}$ satisfy the following constraint by design.

Design constraints 1: Nudging policies are designed to be *non-negative* and satisfy

$$\sum_{\tau=0}^{t-1} u^c(\tau) \leq 1 - \delta \mathbb{1} - u^o \leq \mathbb{1}, \quad (5)$$

to ensure that $u(t) + u^{nc}(t) \in [0, 1]^N$ for all possible realizations of the stochastic input $u^{nc}(t)$ at all $t \in \mathbb{N}_0$.

Note that, if imposed by design, the constraint in (5) guarantees $\mu(t) \in [0, 1]$ for all possible noise realization thanks to Assumption 2.2. Accordingly, the system dynamics (2) can be equivalently rewritten as

$$x(t+1) = \Lambda P x(t) + (I - \Lambda)(u(t) + u^{nc}(t)). \quad (6)$$

Apart from assuming that (5) holds, we further assume that the designed policies are deployed under budget constraints, as formalized in the following assumption.

Assumption 3.1 (Limited resources): Nudging policies are enacted under a fixed and finite *budget* $\beta \in \mathbb{R}_+$, with $\beta \ll \infty$, depleted over time. Therefore, the resources available at time $t \in \mathbb{N}_0$ to nudge individuals are dictated by

$$U(t) = \max \left\{ 0, \beta - \sum_{k=0}^{t-1} \sum_{v \in \mathcal{V}} u_v^c(t-k) \right\}. \quad (7)$$

A. Expected inclinations with constant policies

A strategy that can be used in practice to nudge the adoption of new technologies is to give potential adopters the same incentive over time (e.g., the same tax discount over the years) until budget depletion. This kind of strategy can be formalized as follows:

$$u^c(t) = \begin{cases} \nu, & \text{if } u(t) + \nu \leq 1 - \delta \mathbb{1} \vee U(t) \geq \sum_{v \in \mathcal{V}} \nu_v, \\ \nu^r, & \text{if } u(t) + \nu^r \leq 1 - \delta \mathbb{1} \vee U(t) \in (0, \sum_{v \in \mathcal{V}} \nu_v), \\ 0, & \text{otherwise.} \end{cases} \quad (8)$$

where $\nu_v > 0$ is the baseline magnitude of the intervention the policymaker has tailored to the v -th agent, with $v \in \mathcal{V}$, while ν^r is the action to be enacted to deplete eventual remaining resources in one time step within the feasibility limits, and $u^o \geq \delta \mathbb{1}$ by design. Let us assume that the following holds.

Assumption 3.2: The static policy $u^c(t) = \nu$ is enacted until a finite instant $T \in \mathbb{N}$, i.e., $u^c(T+1) = \nu^r$ and $u^c(t) = \mathbb{0}$ for all $t > T+1$.

Under this Assumption, we define $\nu^r = \nu U(T) / (\sum_{v \in \mathcal{V}} \nu_v)$ so that the baseline input ν is proportionally scaled to consume remaining resources². Note that, if it is not possible to exploit the whole budget without exceeding the feasibility limits, some resources can still remain unused. Meanwhile, when ν is proportional to the budget β , then $\nu^r = \mathbb{0}$.

By relying on Assumption 3.2, we can formalize an asymptotic result on the expected latent inclinations.

Proposition 3.3 (Asymptotic opinions under static policies): Let Assumptions 2.1-3.2 and (5) hold. Then the asymptotic expected inclinations under the policy in (8) satisfy

$$\mu_\infty = (I - P\Lambda)^{-1}(I - \Lambda)\bar{u}, \quad (9a)$$

with $\mu_\infty \in [0, 1]^N$,

$$\bar{u} = u^o + T\nu + \nu^r, \quad \text{with } \bar{u} \leq \mathbb{1} - \delta \mathbb{1}. \quad (9b)$$

Proof: Thanks to (5), the latent inclination evolves according to (6) and, thus, it is straightforward to prove that

$$\mu(t+1) := \mathbb{E}[x(t+1)] = \Lambda P \underbrace{\mathbb{E}[x(t)]}_{:=\mu(t)} + (I - \Lambda)\mathbb{E}[u(t)], \quad (10)$$

due to Assumption 2.2. Meanwhile, because of Assumption 3.2,

$$\mathbb{E}[u(t)] = \bar{u} = u^o + T\nu + \nu^r, \quad \forall t > T+1, \quad (11)$$

where $\bar{u} \leq \mathbb{1} - \delta \mathbb{1}$ is the maximum achievable value³ for $\mathbb{E}[u(t)]$ under the enacted constant (yet saturated) policy in (8), with its finiteness being a consequence of (5) and Assumption 3.1. Accordingly, it further holds that

$$\lim_{t \rightarrow \infty} \mathbb{E}[u(t)] = \bar{u}, \quad (12)$$

from which (9a) follows thanks to Assumption 2.1, thus concluding the proof. ■

Therefore, expected inclinations toward a new technology/service asymptotically converge to a finite value dictated by the characteristics of the static intervention in (8), the features of interpersonal bonds, and the agents' initial biases.

B. Comparison with [11] under static policies

Toward showing the suitability of the proposed model to characterize long-term shifts in individual attitudes, we now compare the asymptotic expected inclinations reported in (9a) with those attained by using the model proposed in [11] under the assumption that the policymaker enacts (8). To this end, let $\mu^{\text{st}}(t)$ be the mean inclinations at time $t \in \mathbb{N}_0$ dictated by the model in [11], which evolves according to

$$\mu^{\text{st}}(t+1) = \Lambda P \mu^{\text{st}}(t) + (I - \Lambda)(u^o + u^c(t)). \quad (13)$$

Then, the following asymptotic result holds.

Proposition 3.4: Let Assumptions 2.1-3.2 and (5) hold. Then, the mean inclinations in (13) under the policy in (8) satisfy

$$\mu_\infty^{\text{st}} = \lim_{t \rightarrow \infty} \mu^{\text{st}}(t) = (I - P\Lambda)^{-1}(I - \Lambda)u^o, \quad (14)$$

with $\mu_\infty^{\text{st}} \in [0, 1]^N$.

Proof: According to (8), it easily follows that

$$\lim_{t \rightarrow \infty} u^o + u^c(t) = u^o,$$

and, consequently, that (14) holds, thus ending the proof. ■

As (14) coincides with the asymptotic latent inclinations in the absence of external, controlled inputs, this result highlights that the model proposed in [11] implicitly relies on the assumption that any policy enacted with a limited budget cannot lead to an irreversible shift in one's expected inclination. Note that, while this might reflect reality when on-off policies are

³This bound is guaranteed by (8), thus not requiring to constraint elements (e.g., the initial inclination u^o) that are not under the policymaker control.

²Alternative definitions of ν^r would not change our formal results.

undertaken, such an (implicit) assumption is instead likely falsified when systemic actions are performed.

Based on Proposition 3.4, we can then compare the asymptotic inclinations resulting from our modeling choices and those made in [11], as subsequently formalized.

Proposition 3.5: Let Assumptions 2.1-3.2 and (5) hold. Let $\mu(t)$ and $\mu^{\text{st}}(t)$ evolve as in (10) and (13), respectively. Then, by enacting the policy in (8), it asymptotically holds that

$$\mu^{\text{st}}(\infty) < \mu(\infty). \quad (15)$$

Proof: The proof straightforwardly follows from the definition of \bar{u} in (9b) and it is thus omitted. ■

C. Feedback policies with constraints and budget limitations

Let us now consider a static feedback policy of the error between full acceptance and the average individual inclinations, i.e.,

$$u^c(t) = K(\mathbb{1} - \mu(t)), \quad (16a)$$

with $K \in \mathbb{R}^{N \times N}$ designed such that the spectral radius of the closed-loop state transition matrix is strictly contained within the unit circle, i.e.,

$$\rho(\Lambda P - (I - \Lambda)K) < 1. \quad (16b)$$

It is worth remarking that (16) implies that $u^c(t) = \mathbb{0}$ whenever $\mu(t) = \mathbb{1}$ and, instead, $u^c(t) = K\mathbb{1}$ if $\mu(t) = \mathbb{0}$. Let us further assume the following.

Assumption 3.6 (Bound on the feedback policy): The feedback policy $u^c(t)$ in (16) is bounded by design in an interval $[u_{\min}^c, u_{\max}^c]$, such that that $\mu(t) \in [0, 1]^N$ for all $t \in \mathbb{N}_0$. Hence, the mean input $\mathbb{E}[u(t)]$ satisfies

$$u^o + \mathcal{T}_{\min} u_{\min}^c \leq \mathbb{E}[u(t)] \leq u^o + \mathcal{T}_{\max} u_{\max}^c, \quad \forall t \in \mathbb{N}_0, \quad (17)$$

where \mathcal{T}_{\min} and \mathcal{T}_{\max} are diagonal matrices containing the time instants after which the controlled input is set to zero due to saturation of the states or consumption of the budget (similarly to (9b)). Accordingly, the following asymptotic result holds.

Proposition 3.7 (Asymptotic opinions and feedback): Let Assumptions 2.1-3.1 be satisfied and let the enacted policy be defined as in (16) while satisfying Assumption 3.6. Then, the expected inclination achieved by closing the loop is asymptotically limited, i.e.,

$$\mu_{\infty} := \lim_{t \rightarrow \infty} \mathbb{E}[x(t)] \leq 1. \quad (18)$$

Proof: Along the same line of the proof of Proposition 3.3, the result in (18) straightforwardly follows from (17). Therefore, the proof is omitted. ■

IV. TOWARD OPTIMAL NUDGING POLICY DESIGN

By relying on the model introduced in Section II, we now propose two strategies that policymakers/stakeholders can adopt to design interventions that trade-off encouraging the widespread adoption of a new technology/service (i.e., maximizing social benefit) and avoiding waste of resources. To this end, we make the following assumption.

Assumption 4.1: The intrinsic predispositions u^o of the agents or their estimates are available for policy design. While having direct access to u^o might be unrealistic, the agents' intrinsic predisposition can be inferred from real-world

data (see, e.g., [14]). We postpone a formal analysis on the impact of eventual errors in the estimation of the agents' intrinsic predisposition to future works.

A. Optimized Constant Control Policy (CCP)

As a first alternative to design a policy that targets the aforementioned goal in one shot, policymakers/stakeholders can take advantage of the asymptotic properties of the proposed model (discussed in Section III-A). Specifically, considering a prefixed time horizon $T \in \mathbb{N}$ for the policy's deployment, a constant policy can be designed by solving the following problem

$$\underset{\mu_{\infty}, u_{\infty}^c}{\text{minimize}} \quad J^{\text{CCP}}(\mu_{\infty}, u_{\infty}^c) \quad (19a)$$

$$\text{s.t.} \quad \mu_{\infty} = (I - \Lambda P)^{-1}(I - \Lambda)(u^o + T u_{\infty}^c), \quad (19b)$$

$$T \sum_{v \in \mathcal{V}} u_{\infty}^c, v \leq \beta, \quad (19c)$$

$$u_{\infty}^c, v \geq 0, \quad \forall v \in \mathcal{V}, \quad (19d)$$

$$u_v^o + T u_{\infty}^c, v \leq 1 - \delta, \quad \forall v \in \mathcal{V}, \quad (19e)$$

where the last constraint guarantees that (5) is satisfied, and the loss is defined as

$$J^{\text{CCP}}(\mu_{\infty}, u_{\infty}^c) = \|\mathbb{1} - \mu_{\infty}\|_2^2 + \|T u_{\infty}^c\|_R^2 + \left\| \beta - T \sum_{v \in \mathcal{V}} u_{\infty}^c, v \right\|_S^2, \quad (19f)$$

with $R \succ 0$ and $S \succeq 0$ being penalties chosen by the policymaker, and the last term in the cost aims at minimizing the unused resources. Note that, since $\delta < u_v^o$ for all $v \in \mathcal{V}$ by Assumption 2.2, the lower-bound in (5) is guaranteed by construction. As a result, policymakers can enact

$$u^c(t) = u_{\infty}, \quad \forall t \in \{0, 1, \dots, T-1\}. \quad (20)$$

B. Model Predictive Control (MPC) Fostering Policy

Instead of looking at asymptotic behaviors, a policymaker/stakeholder can instead decide to optimize its strategies in a receding horizon fashion⁴ by monitoring average opinions and accordingly adjusting their strategies over time. In this case, the control problem⁵ can be formulated as:

$$\underset{\mathcal{M}_{|t}, \mathcal{U}_{|t}^c}{\text{minimize}} \quad J^{\text{MPC}}(\mathcal{M}_{|t}, \mathcal{U}_{|t}^c) \quad (21a)$$

$$\text{s.t.} \quad \mu_{|t}(k+1) = \Lambda P \mu_{|t}(k) + (I - \Lambda) u_{|t}(k), \quad (21b)$$

$$u_{|t}(k+1) = u_{|t}(k) + u_{|t}^c(k), \quad k \in [0, L-1], \quad (21c)$$

$$u_{|t}^{\Sigma}(k) = \sum_{\tau=0}^k \sum_{v \in \mathcal{V}} u_{|t}^c, v(\tau), \quad k \in [0, L-1], \quad (21d)$$

$$u_{|t}^c, v(k) \geq 0, \quad \forall v \in \mathcal{V}, \quad k \in [0, L-1], \quad (21e)$$

$$u_{|t}^{\Sigma}(k) \leq U_{|t}(0) - u_{|t}^{\Sigma}(k-1), \quad k \in [1, L-1], \quad (21f)$$

$$u_{|t}, v(k) \leq 1 - \delta, \quad \forall v \in \mathcal{V}, \quad k \in [0, L-1], \quad (21g)$$

$$u_{|t}(0) = u(t), \mu_{|t}(0) = \mu(t), U_{|t}(0) = U(t), \quad (21h)$$

⁴This choice allows us to mitigate the (strong) requirement of an infinite horizon control policy, which necessitates policymakers to unrealistically foreseen individual average attitudes over an infinite time span.

⁵The subscript $|t$ denotes a prediction made starting from time t and, therefore, based on the information available at that time instant.




Fig. 1. The social network considered in our example, featuring 20 agents and 7 clusters of agents.

with $\mathcal{M}_{|t} = \{\mu_{|t}(k)\}_{k=0}^L$ and $\mathcal{U}_{|t}^c = \{u_{|t}^c(k)\}_{k=0}^{L-1}$, $U(t)$ being defined as in (7), and

$$J^{\text{MPC}}(\mathcal{M}_{|t}, \mathcal{U}_{|t}^c) = \sum_{k=0}^{L-1} \|\mathbb{1} - \mu_{|t}(k)\|_2^2 + \|u_{|t}^c(k)\|_R^2 + \|\mathbb{1} - \mu_{|t}(L-1)\|_Q^2, \quad (21i)$$

where $L \geq 1$ is the prediction horizon decided by the policymaker/stakeholder, while $R \succ 0$ controls the trade-off between adoption boosting/cost containment and the terminal penalty can be weighed according to

$$(\Lambda P)^\top Q \Lambda P - Q = -I, \quad (21j)$$

since the dynamics of expected opinions is asymptotically stable by Assumption 2.1. Note that, while (21f) allows us to explicitly account for budget consumption in policy design (see Assumptions 3.1), (21g) guarantees that (5) is satisfied as the lower-bound is already verified by construction (see the initial condition in (2b)). It is worth remarking that the cost in (21i) comprises two terms that penalize the average distance of the agents' opinions to the acceptance of the targeted technology/service and the second term, which weights the distance of the policy action to be designed from its (equilibrium) value at an average full adoption⁶.

Remark 4.2 (Practical issues with policy implementation): Designing a policy as in (21) requires continuative monitoring of the average agents' inclination, which is likely unfeasible in practice. We postpone tackling this practical issue along the same lines of [11] in future works.

V. NUMERICAL EXAMPLE

We now analyze the impact of the strategies introduced in Section IV considering the (randomly generated) modular social network depicted in Fig. 1 and assuming that individual opinions evolve according to the long-term shifts model proposed in Section II.

The considered social network comprises $N = 20$ agents and 7 clusters, generated by setting a link probability of 0.2 and the probability of connection between agents of different clusters at 0.7. We impose the elements in Λ (see (2)) to be equal, i.e., $\Lambda = \lambda I$, yet considering two scenarios where external interventions (i.e., $\lambda = 0.25$) and social influences (namely, $\lambda = 0.75$) are the main drivers of adoption, respectively. Meanwhile, we consider three setups for the initial bias u^o by splitting the agents into two groups, namely:

⁶It is straightforward to prove that the controlled input at full adoption is $\bar{u}^c = \mathbb{0}$ and, thus, such proof is omitted.




Fig. 2. Negatively biased scenario: evolution of latent inclinations under different budget constraints.

- 1) *mixed* biases: $u_v^o = 0.2$ for half agents and 0.8 for the remaining ones, thus having a population grouped into distinct factions with respect to the new technology/service;
- 2) *negative* biases: $u_v^o = 0.2$ for half agents and 0.3 for the remaining ones, so that the considered population has a negative polarization toward the technology/service;
- 3) *positive* biases: $u_v^o = 0.6$ for half agents and 0.8 for the remaining ones, considering a population that is instead positively polarized toward such a technology/service.

The performance achieved through the designed policies is evaluated by looking at their *final distance to acceptance*, here defined as

$$\Gamma_{\text{sim}} = \|\mathbb{1} - x(T_{\text{sim}})\|_2^2, \quad (22)$$

where $T_{\text{sim}} = 30$ is the considered simulation horizon, as well as their cumulative cost and usage of the available budget, i.e.,

$$u_{\text{sim}}^\Sigma = \sum_{t=0}^{T_{\text{sim}}-1} \sum_{v \in \mathcal{V}} u_v^c(t), \quad B_\% = 100 \frac{u_{\text{sim}}^\Sigma}{\beta} [\%], \quad (23)$$

respectively. In all our tests, uncontrollable factors are described as realizations of a uniformly distributed white noise, with $\delta = 0.025$ (see Assumption 2.2).

A. Analyzing the impact of different budgets

By considering the same penalty for the input effort introduced before, namely $R = 10I$, we focus on the performance of the approach proposed in this paper for different budgets. In particular, we consider a scenario with a high budget $\beta_{\text{high}} = 25$, so that all available resources do not need to be depleted to achieve the widespread diffusion of the new technology/service, a setting with moderate budget $\beta_{\text{mod}} = 8$ (fully exhausted only for some of the scenarios we consider for the individual biases), and, lastly, a low budget $\beta_{\text{low}} = 5$ case, where all resources are depleted irrespective of individual biases.

As shown in Table I, the lower the available budget, the more individuals will be resistant to embrace the new technology on average. Moreover, the highest final distance to acceptance is achieved in the second scenario, as the population is negatively biased toward the technology/service. It can be observed that, when the budget is constrained, a lower final distance to acceptance is attained with $\lambda = 0.75$. In this case, since the budget is limited, also social influence drives the promotion of the new technology/service, highlighting

TABLE I

FINAL DISTANCE TO ACCEPTANCE Γ_{sim} AND BUDGET CONSUMPTION $B\%$ FOR DIFFERENT BUDGETS β .

| λ | β_{high} | | | | β_{mod} | | | | β_{low} | | | |
|---------------|-----------------------|------|-------|-------|-----------------------|------|--------|--------|-----------------------|------|--------|--------|
| | Γ_{sim} | | $B\%$ | | Γ_{sim} | | $B\%$ | | Γ_{sim} | | $B\%$ | |
| | 0.25 | 0.75 | 0.25 | 0.75 | 0.25 | 0.75 | 0.25 | 0.75 | 0.25 | 0.75 | 0.25 | 0.75 |
| Mixed bias | 0.01 | 0.01 | 38.00 | 38.00 | 0.19 | 0.17 | 100.00 | 100.00 | 1.25 | 1.11 | 100.00 | 100.00 |
| Negative bias | 0.01 | 0.01 | 58.00 | 58.00 | 2.38 | 1.97 | 100.00 | 100.00 | 4.88 | 4.23 | 100.00 | 100.00 |
| Positive bias | 0.01 | 0.01 | 22.00 | 22.00 | 0.01 | 0.01 | 68.75 | 68.75 | 0.05 | 0.04 | 100.00 | 100.00 |




Fig. 3. Γ_{sim} vs u_{sim}^{Σ} : constant vs receding horizon policy.

the importance of the interplay between social contagion and external policies in fostering innovation diffusion. Focusing on the negatively biased population with $\lambda = 0.25$, Fig. 2 showcases the impact of budget restrictions on a realization of closed-loop inclinations. As expected, all individual inclinations approach 1 when the budget is high (i.e., we have a near-universal acceptance of the technology), which is not the case when the available budget is moderate. Note that, in this scenario, the latent opinions at the end of the considered simulation horizon are nonetheless higher than the value achieved in open-loop (i.e., when no policy is enacted).

B. Static vs receding horizon policy

Finally, we analyze the possible advantages of the receding horizon policy in (21) compared to the constant policy (19) focusing on the mixed biases scenario, under a budget $\beta = 10$ and imposing $R = 15I$. Since the receding horizon policy exhausts the budget after 20 time steps, we further set T in (19) to 20, while imposing $S = 10$ to avoid waste of resources with the constant policy. As summarized in Fig. 3, within this setting, the MPC strategy uses more resources at the beginning than the CCP strategy, resulting in a lower final distance to acceptance in less time. The receding horizon strategy slightly outperforms the CCP one in terms of final distance to acceptance even at the end of the horizon, despite the cost of the two policies becoming aligned (and close to zero, as the whole budget is consumed).

VI. CONCLUSIONS

In this work, we propose an opinion dynamics model that describes long-term shifts in opinion induced by external policies by introducing an artificial accumulation state, directly impacting individual opinion dynamics. We then rely on the proposed model to introduce two strategies for policy design aimed at trading off a widespread adoption of a

new (sustainable) technology/service and costs under budget constraints, whose impact on opinion dynamics is evaluated through numerical simulations.

Future works will be devoted to blending models describing only short-term shifts in inclination with the proposed one, considering stochastic control approaches for policy design to cope with uncontrollable external factors, and analyzing the realism of such models and the validity of the proposed policy design strategies on real data.

REFERENCES

- [1] P. A. Nylund, A. Brem, and N. Agarwal, "Enabling technologies mitigating climate change: The role of dominant designs in environmental innovation ecosystems," *Technovation*, vol. 117, p. 102271, 2022.
- [2] G. Cantelmo, R. E. Amini, M. M. Monteiro, A. Frenkel, O. Lerner, S. S. Tavory, A. Galtzur, M. Kamargianni, Y. Shiftan, C. Behrischi *et al.*, "Aligning users' and stakeholders' needs: How incentives can reshape the carsharing market," *Transport Policy*, vol. 126, pp. 306–326, 2022.
- [3] A. Pamidimukkala, S. Kermanshachi, J. M. Rosenberger, and G. Hladik, "Barriers and motivators to the adoption of electric vehicles: A global review," *Green Energy and Intelligent Transportation*, vol. 3, no. 2, p. 100153, 2024.
- [4] J. Markard, F. W. Geels, and R. Raven, "Challenges in the acceleration of sustainability transitions," *Environmental Research Letters*, vol. 15, no. 8, p. 081001, 2020.
- [5] V. Breschi, C. Ravazzi, S. Strada, F. Dabbene, and M. Tanelli, "Driving electric vehicles' mass adoption: An architecture for the design of human-centric policies to meet climate and societal goals," *Transportation Research Part A: Policy and Practice*, vol. 171, p. 103651, 2023.
- [6] A. M. Annaswamy, K. H. Johansson, and G. Pappas, "Control for societal-scale challenges: Road map 2030," *IEEE Control Systems Magazine*, vol. 44, no. 3, pp. 30–32, 2024.
- [7] M. H. DeGroot, "Reaching a consensus," *Journal of the American Statistical Association*, vol. 69, no. 345, pp. 118–121, 1974.
- [8] N. E. Friedkin and E. C. Johnsen, "Social influence and opinions," *Journal of mathematical sociology*, vol. 15, no. 3-4, pp. 193–206, 1990.
- [9] Y. Kareeva, A. Sedakov, and M. Zhen, "Influence in social networks with stubborn agents: From competition to bargaining," *Applied Mathematics and Computation*, vol. 444, p. 127790, 2023.
- [10] B. Sprenger, G. De Pasquale, R. Soloperto, J. Lygeros, and F. Dörfler, "Control strategies for recommendation systems in social networks," *IEEE Control Systems Letters*, 2024.
- [11] C. Ravazzi, V. Breschi, P. Frasca, F. Dabbene, and M. Tanelli, "Optimal policy design for repeated decision-making under social influence," *arXiv preprint arXiv:2503.03657*, 2025.
- [12] I.-C. Morărescu, V. S. Varma, L. Buşoniu, and S. Lasaulce, "Space-time budget allocation for marketing over social networks," *IFAC-PapersOnLine*, vol. 51, no. 16, pp. 211–216, 2018.
- [13] P. Frasca, C. Ravazzi, R. Tempo, and H. Ishii, "Gossips and prejudices: Ergodic randomized dynamics in social networks," *IFAC Proceedings Volumes*, vol. 46, no. 27, pp. 212–219, 2013.
- [14] E. Villa, V. Breschi, C. Ravazzi, M. Tanelli, and F. Dabbene, "Can control aid in attaining sustainable goals? an improved data-informed framework to promote shared mobility," *Control Engineering Practice*, vol. 153, p. 106106, 2024.