Final Report
Project CODE: Coupling Opinion Dynamics with Epidemics

February 2026

Contents

Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA

1 Coordination, Management and Dissemination

1.1 Organization and Scientific Coordination

The CODE project involved three research units:

- Consiglio Nazionale delle Ricerche (CNR)
- Politecnico di Milano (PoliMi)
- Museo Storico della Fisica e Centro Ricerche "Enrico Fermi" (CREF)

The collaboration integrated complementary expertise in network science, epidemic modeling, computational social science and data analysis. The project has been coordinated by the Principal Investigator, who ensured consistency between the different research lines and integration of the methodological developments.

The organizational structure relied on regular meetings at different levels. Weekly meetings were held within each research unit to monitor the progress of the activities and coordinate software development and data analysis. Monthly online meetings involving all partners ensured continuous alignment of the different research directions and allowed us to refine the modeling framework as new results emerged.

This structure proved particularly effective in a project where theoretical modeling, empirical analysis and software development were strongly interconnected.

1.2 Dissemination Activities

To support dissemination and ensure public accessibility of the project's activities and outputs, we created the project website:

https://www.code-project.it/

The scientific results of the project were presented at several major international conferences in the fields of network science, computational social science and complex systems.

Intermediate results on heterogeneous epidemic models were presented at the **Complex Networks 2024** conference. Work on behavioral epidemic models and opinion dynamics was presented at the **CS2Italy 2025 Conference**, the first Italian conference on Computational Social Science.

Further results on network modeling and data-driven epidemic simulations were presented at **NetSci 2025**, **StatPhys29** and at the **Conference on Complex Systems 2025 (CCS 2025)**.

These dissemination activities allowed us to obtain feedback from the international research community and contributed to refining the modeling framework developed within the project.

1.3 Scientific Events Organized by the Project

A major dissemination activity of the project was the organization of the satellite workshop

BeSAFE: Behavioral and Social Aspects in Fighting Epidemics

held at the Conference on Complex Systems 2025.

The workshop focused on the interplay between epidemic processes, human behavior and information dynamics and brought together leading researchers working at the interface between epidemiology and computational social science.

Invited speakers included:

- Eugenio Valdano
- Chiara Poletto
- Matteo Cinelli
- Kyriaki Kalimeri

The workshop provided an opportunity to discuss recent advances in the modeling of behavioral responses to epidemics and contributed to positioning the CODE project within the international research landscape.

At the end of the project we organized a final workshop hosted by CREF in February 2026. The workshop presented the main results of the project and included invited talks by several external researchers working on related topics.

Invited speakers included:

- Giulia Andrighetto (CNR-ISTC)
- Claudio Castellano (CNR-ISC)
- Riccardo Gallotti (FBK)
- Michele Tizzani (DTU)
- Mattia Mazzoli (ISI Foundation)
- Lorenzo Lucchini (Bocconi University and FBK)
- Elisabetta Colosi (Bocconi University)

The final workshop provided a comprehensive overview of the project results and stimulated discussions on future research directions.

1.4 Publications

The project produced several scientific publications. Among them, the following are most directly related to the project objectives.

Discursive communities. The paper

Leveraging content producer networks and user perception to detect online discursive communities [1],

accepted for publication on Scientific Reports, introduces a method for identifying discussion communities in online social networks. The method allows the identification of a small number of interpretable communities that accurately capture the underlying political divisions among users.

Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA

Behavioral heterogeneity in epidemic spreading. The paper

The Impact of Heterogeneity on Epidemics: Insights from a Modified SIR Model

published in the proceedings of Complex Networks 2024, introduces the HeSIR model, a modified SIR framework incorporating heterogeneous susceptibility and infectivity across behavioral groups. The study shows that the combined presence of these heterogeneities can generate a resurgence regime in which an initial decline in infections is followed by a delayed outbreak. The analysis also highlights the role of homophily and degree heterogeneity in amplifying this effect, showing how structured populations may conceal epidemic potential and delay the recognition of epidemic risk.

The above paper has been extended for the journal publication

Impact of behavioral heterogeneity on epidemic outcome and its mapping into effective network topologies

published in Physical Review E. This paper analyzes the HeSIR model more in details, deriving analytical expressions for the epidemic threshold and shows how behavioral heterogeneity can be mapped into an equivalent effective network topology. The results clarify how differences in individual behavior and homophilic interaction patterns jointly shape epidemic dynamics and may generate non-trivial outbreak regimes even in relatively simple network structures.

Network modeling. The paper

Random hyperbolic graphs with arbitrary mesoscale structures [2],

published in Physical Review E introduces the Random Hyperbolic Block Model within a maximum entropy framework. The model combines the geometric structure of latent space network models with explicit control over community mixing patterns.

Data-driven epidemic modeling. The paper

A data-driven analysis of the impact of non-compliant individuals on epidemic diffusion in urban settings [3],

published in Proceedings of the Royal Society A, shows how even small fractions of non-compliant individuals can strongly affect epidemic dynamics in realistic urban networks.

Vaccination strategies. The paper

When to boost: How dose timing determines the epidemic threshold [4],

published in Physical Review Research, studies SIS-type epidemic models with multi-dose vaccination and fading immunity and identifies optimal vaccination strategies across different epidemiological regimes.

Randomization of bipartite networks. The paper

Pattern detection in bipartite networks: A review of terminology, applications, and methods [5],

published in PLOS Complex Systems, presents a multidisciplinary review of null models and randomization techniques for binary bipartite matrices. It outlines main randomization strategies and evaluates their performance and limitations, and emphasizes the role of structural constraints—especially preservation of row and/or column sums—in designing meaningful null ensembles. Bipartite networks are used to highlight discursive communities, crf. [1].

Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA

Echo chambers detection. The paper

Entropy-based detection of Twitter echo chambers [6],

published in PNAS Nexus, proposes a statistically grounded method to detect echo chambers by identifying significant overlap between groups of like-minded users and groups consuming the same news. Applied to the Italian Twitter debate on Covid-19 vaccination, the approach shows that echo-chamber users are a small minority but disproportionately drive the conversation and often spread misinformation.

News Publishers Trustworthiness. The paper

Unveiling News Publishers Trustworthiness Through Social Interactions [7],

published as a proceeding paper of the ACM Web Science Conference 2024, introduces a scalable framework to assess the trustworthiness of online news publishers by leveraging user interactions on social media. It aims both to identify verifiable publishers and to provide an automated first-pass trustworthiness estimate for previously unclassified outlets.

Online disinformation. The paper

Online disinformation in the 2020 U.S. election: swing vs. safe states [8]

was published in EPJ Data Science. Analyzing Twitter activity around the 2020 U.S. pre-election debate, this study links the winner-take-all electoral system to a disproportionate focus on swing states. Tweets associated with swing states show a significantly higher prevalence of links to unreliable sources, largely amplified by automated accounts. The conversation is dominated by two main communities, with most disinformation originating from the predominantly Republican-affiliated one.

Hypergraph randomization. The paper

Entropy-based models to randomise real- world hypergraphs [9],

published in Communications Physics, introduces Exponential Random Hypergraphs (ERHs), an entropy-based null-model framework that extends exponential random graph ideas to higher-order (polyadic) interactions via hypergraph incidence matrices. Comparing expected and empirical values of metrics defined for hypergraphs yields a scalable, analytically tractable method to detect non-random structural patterns in real-world hypergraphs beyond simple constraints. Higher-order interactions are believed to play a crucial role in describing social systems in which many individuals interact among themselves simultaneously.

Signed bipartite network randomization. The paper

Statistically validated projection of bipartite signed networks [10],

published in npj Complexity, proposes an unsupervised method to build statistically validated one-mode projections of bipartite signed networks when intra-layer links are unobserved. Node pairs are connected by positive/negative edges if they share a statistically significant excess of concordant/discordant relationships, yielding link-specific p-values and a validated projection via multiple-hypothesis testing. Tests on synthetic and real data show the method uncovers non-trivial mesoscopic structures not explained by benchmark constraints, indicating genuine self-organization.

1.5 Data and Software

All data and software produced during the project are made available through the project website or the public repository:

https://github.com/steguar/code-toolbox/

The project produced a collection of software tools for network generation, epidemic simulation and opinion dynamics modeling. The models can be explored separately and provide a basis for future integrated socio-epidemic simulations.

The public repository also contains additional research code and datasets produced during the project that are not directly integrated into the interactive toolbox. These include data-driven network generation components, tools for the identification of discursive communities in online platforms, datasets related to territorial non-compliance analysis, and collections of social media data used in the empirical analyses described in the following sections.

Part of the software is accessible through a web-based dashboard that allows non-technical users to explore model scenarios without the need for programming skills.

The dashboard and the associated backend code are released as open-source software.

2 Modeling Framework

2.1 General Framework

The main objective of the CODE project has been to investigate the interplay between epidemic spreading and the dynamics of information, opinions and behaviors in structured populations.

A central idea guiding the project is that the notion of community provides a natural first-order approximation for describing how social and behavioral heterogeneity emerges in real populations. Individuals tend to organize into groups characterized by similar ideological orientations, sources of information and reference influencers. Members of the same group are therefore likely to be exposed to similar narratives and interpretations of events and, as a consequence, tend to develop similar opinions and behavioral attitudes.

In the context of epidemic containment, this implies that different ideological or informational communities may exhibit systematically different levels of compliance with mitigation measures such as vaccination, distancing or mobility reduction. Rather than treating behavioral variability as an individual-level random effect, the project adopts the perspective that behavioral heterogeneity is largely structured at the community level.

This observation naturally leads to the question of how epidemic and social dynamics depend on the way such communities are connected to one another. In particular, different patterns of interaction between communities may produce qualitatively different macroscopic outcomes. For instance, strong modularity in the network of physical contacts with respect to ideological or behavioral attributes may allow high-risk subpopulations to sustain transmission even when the global epidemic appears under control. Conversely, highly mixed interaction patterns may lead to rapid synchronization of epidemic dynamics across groups.

A similar reasoning applies to information and opinion dynamics, where community structure determines the pathways through which information propagates and influences individual attitudes.

Based on these considerations, the project adopted a modeling strategy centered on the explicit representation of communities and their mixing patterns.

The general research line has been structured around two main elements.

First, we developed models for generating synthetic networks that reproduce the main structural properties of real interaction systems and allow explicit control of community structure and mixing patterns. These networks represent both physical contacts relevant for epidemic transmission and online interactions relevant for information diffusion.

Second, epidemic processes and opinion dynamics have been studied on top of these networks in order to understand how community structure and mixing patterns shape the resulting collective behavior.

This approach provides a unified framework in which epidemic diffusion, information spreading and behavioral responses can be analyzed within a common structural setting.

The resulting modeling framework allows both the construction of realistic data-driven scenarios and the study of simplified models aimed at clarifying the qualitative behavior of the system.

2.2 Network Models and Community Structure

A central component of the CODE project has been the development of network models capable of representing populations structured into communities with heterogeneous interaction patterns.

As discussed in Section 2.1, the notion of community provides a natural representation of ideological and behavioral heterogeneity. Individuals belonging to the same community tend to share similar sources of information and reference influencers and therefore often exhibit correlated behavioral responses during epidemic crises.

For this reason the modeling framework developed in the project relies on network models that explicitly incorporate community structure and allow quantitative control of mixing patterns between communities.

Two complementary classes of models have been developed. The first is a general class of maximum-entropy network models designed to provide flexible synthetic benchmarks with controlled structural properties. The second is a data-driven framework for constructing realistic urban-scale contact networks integrating demographic and geographic information.

In addition, we developed methods for extracting ideological community structure directly from online social interaction data. These methods provide empirical guidance for defining realistic mixing patterns in both physical and online interaction networks.

2.2.1 Random Hyperbolic Block Model

The Random Hyperbolic Block Model (RHBM) is a general network model combining latent geometry with explicit control over mesoscale structure.

The model has been introduced as a maximum-entropy ensemble allowing the simultaneous control of:

- degree heterogeneity,
- clustering,
- community structure,
- mixing patterns between communities.

Nodes are embedded in the hyperbolic disk and each node i is assigned two coordinates:

- a radial coordinate r_i controlling expected degree,
- an angular coordinate θ_i representing similarity.

Nodes are also assigned to a community label

$$I_i \in \{1, \ldots, B\},$$

and the expected number of links between communities is specified by a mixing matrix

$$K = \{K_{IJ}\}.$$

The model is defined as the maximum-entropy ensemble satisfying the constraints:

$$\langle k_i \rangle = k_i^*, \tag{1}$$
$$\langle L_{IJ} \rangle = K_{IJ}, \tag{2}$$
$$\langle E \rangle = E^*, \tag{3}$$

where k_i^* denotes the expected degree of node i, L_{IJ} is the number of links between communities I and J, and E represents the total geometric energy of the configuration.

Under these constraints the probability of a link between nodes i and j takes the form

$$p_{ij} = \frac{1}{1 + \left(\frac{\Delta(\theta_i, \theta_j) \mathcal{I}_\beta}{\pi K_{I_i I_j} f_i f_j} \right)^\beta},$$

where $\Delta(\theta_i, \theta_j)$ is the angular distance between nodes and β is an inverse temperature parameter controlling clustering.

Low temperatures correspond to highly clustered networks where edges are strongly concentrated between geometrically close nodes, while high temperatures produce networks approaching block-model behavior.

The RHBM model extends classical degree-corrected stochastic block models by incorporating a latent geometric structure and extends random hyperbolic graphs by allowing arbitrary mesoscale organization.

This flexibility makes the model particularly suitable for studying the role of community mixing in epidemic and social dynamics.

Within the CODE project the RHBM model has been used both as a standalone synthetic network generator and as a building block of more complex data-driven network models.

2.2.2 Urban Social Network Model

While the RHBM provides a general modeling framework, realistic epidemic simulations require networks that reproduce the demographic and spatial structure of real populations.

For this purpose we adopted and extended the Urban Social Network (USN) model, a data-driven framework for generating synthetic populations and contact networks at the urban scale.

The USN model had been introduced in previous work and was further developed during the CODE project in order to integrate ideological and behavioral heterogeneity into the synthetic populations.

The model consists of two main components: the generation of a synthetic population and the construction of the interaction network.

The implementation of this data-driven framework is included in the project repository for reproducibility purposes, although it is not currently integrated into the public simulation dashboard.

Synthetic population. Individuals are generated using publicly available statistical data including:

- population density,
- age distributions,
- household statistics,
- geographic information.

Each individual is assigned a set of attributes including:

- spatial location,
- age group,
- household membership,
- sociability parameters.

During the project the model has been extended to include data-informed ideological attributes.

Electoral data aggregated at the voting precinct level are used to assign political preferences to individuals. Spatial units used in the synthetic population are mapped to electoral precincts and individuals are assigned ideological attributes according to the observed distributions.

This extension allows the construction of populations where behavioral heterogeneity is correlated with geographic location.

Network construction. The interaction network is constructed as a multilayer structure including household contacts and non-household interactions.

Household members are connected through fully connected subgraphs, representing strong and persistent contacts.

Non-household interactions are generated probabilistically using block models where blocks correspond to combinations of demographic and geographic attributes.

Within the CODE project the acquaintance layer of the USN model has been extended using the Random Hyperbolic Block Model.

This extension allows the simultaneous reproduction of:

- heterogeneous degree distributions,
- realistic clustering,
- controlled community mixing.

The integration between the USN framework and the RHBM model provides a flexible approach for constructing realistic urban contact networks while preserving analytical control of their structural properties.

2.2.3 Data-driven Identification of Ideological Communities

In addition to generative models, the project developed methods for identifying ideological communities directly from social interaction data.

Online social networks provide a direct observation of how individuals organize into discursive communities characterized by shared narratives and information sources.

We developed network-based methods for extracting such communities from large-scale Twitter/X interaction datasets.

The approach is based on the observation that political discourse on social platforms is typically structured around a relatively small number of influential content producers whose messages are amplified by their audiences.

Two complementary methods have been introduced to identify discursive communities.

The MonoDC method constructs a network of content producers based on retweet interactions and identifies communities through the analysis of statistically validated connections.

The BiDC method instead builds similarity networks between content producers based on the overlap of their audiences.

Both approaches rely on maximum-entropy null models in order to filter statistical noise and retain only significant interaction patterns.

The resulting communities correspond closely to known political divisions and provide interpretable ideological partitions of the user population.

These methods provide an empirical basis for defining realistic community structures and mixing matrices.

They can be used to inform synthetic network models describing physical contacts as well as models of online information diffusion.

In particular, the ability to extract interpretable ideological communities provides a direct link between empirical data and the network models used in the CODE project.

3 Modeling Toolbox

A central objective of the CODE project has been the development of an open-source computational toolbox for simulating epidemic and social dynamics on structured populations.

While the network models described in the previous section provide the structural backbone of the framework, the modeling of the associated dynamical processes required identifying an appropriate compromise between realism and simplicity.

An extensive analysis of the literature on epidemic modeling and opinion dynamics was carried out with the goal of identifying the main mechanisms that should be included in a general-purpose simulation framework while preserving interpretability and usability (cfr. the project's Deliverables D1.3, D2.3 and D3.3).

Rather than aiming at the most detailed possible description of real systems, the models implemented in the toolbox were designed to capture the essential mechanisms influencing epidemic and social dynamics, while remaining sufficiently simple to allow systematic exploration of alternative scenarios.

This approach led to the definition of two main dynamical modeling frameworks implemented in the toolbox:

- a multi-community epidemic model with behavioral modulation,
- an opinion dynamics model coupled to epidemic risk perception.

Both models are implemented as open-source software and can be used either programmatically or through an interactive dashboard.

In addition, the toolbox includes a stochastic network simulator implementing the S3I2 epidemic model with multi-dose vaccination and waning immunity.

Because of the large number of parameters and the strong nonlinear couplings involved, analytical investigation of the full models is in general difficult. For this reason the primary goal of this part of the project has been the identification of the key mechanisms to include and the development of a flexible computational environment allowing systematic projection and comparison of alternative scenarios.

The toolbox therefore provides a framework primarily intended for simulation-based exploration of model behavior. In the following sections we describe the two models in detail and present a first empirical analysis aimed at identifying some of their key features.

3.1 A Community-Structured Epidemic Model with Behavioral Modulation and Media Effects

Epidemic spreading processes in real populations are strongly influenced by behavioral responses, information dynamics, and social heterogeneity. Traditional compartmental models typically assume homogeneous mixing and fixed transmission parameters, thereby neglecting the feedback mechanisms through which media exposure, perceived local prevalence, and ideological predispositions shape individual behavior. However, empirical evidence shows that disease transmission rates evolve dynamically as individuals adjust their contacts and protective behaviors in response to both global information and local epidemic conditions.

The implemented model addresses these limitations by integrating a community-structured epidemic framework with dynamic behavioral modulation mechanisms. The approach combines classical compartmental dynamics with a time-dependent transmission rate influenced by media intensity, local prevalence awareness, and community-specific ideological factors. This formulation enables the study of how information processes and social heterogeneity jointly affect epidemic evolution.

The framework supports both deterministic mean-field simulations and stochastic network-based realizations, allowing the exploration of epidemic dynamics across multiple modeling scales.

3.1.1 Model Architecture and Mathematical Framework

Population Structure and Community Organization The population is organized into a set of distinct communities representing groups of individuals with shared social, ideological, or behavioral characteristics. Each community is associated with a subset of nodes extracted from an underlying interaction network. These communities are assumed to be static over time and serve as the primary units through which heterogeneity in transmission dynamics is represented.

Within each community i, individuals are classified according to epidemiological compartments corresponding to susceptible, exposed, infected, and removed states, depending on the specific compartmental formulation adopted ($S_i(t)$, $E_i(t)$, $I_i(t)$, $R_i(t)$). The model supports multiple epidemic configurations, including SIR, SEIR, SIS and SEIS variants.

The community structure introduces heterogeneity in contact patterns by distinguishing between interactions occurring within communities and those occurring between different groups.

Effective Transmission Rate A central feature of the model is the introduction of a time-dependent effective transmission rate that reflects behavioral adaptation. Unlike classical epidemic models in which

transmission parameters remain constant, the implemented framework assumes that transmission evolves dynamically as a function of information exposure and local epidemic conditions.

The effective transmission rate for each community is determined by three multiplicative components.

$$\beta_i(t) = \beta_0 \, f_i^{(media)}(t) \, f_i^{(local)}(t) \, g(\theta_i) \tag{4}$$

The first component captures the influence of global media intensity $f_i^{(media)}(t)$, representing behavioral changes induced by widespread information about the epidemic. The second component accounts for local prevalence awareness $f_i^{(local)}(t)$, reflecting behavioral responses to infection levels observed within social surroundings. The third component represents a community-specific ideological factor $g(\theta_i)$ that modulates baseline attitudes toward risk and preventive behavior.

This factorized structure allows the model to separately quantify the effects of global information, local experience, and intrinsic behavioral predispositions on disease transmission.

Media Intensity Dynamics The media influence mechanism is modeled through a dynamic variable representing the global intensity of epidemic-related information. Rather than responding instantaneously to current prevalence, media intensity evolves as a filtered and delayed function of the overall infection level in the population.

$$f_i^{(media)}(t) = 1 - \phi_i M(t) \, \varepsilon_{media} \tag{5}$$

Where ϕ_i represents the sensitivity of community i to the media effect, $M(t)$ accounts for the global intensity of the media and ε_{media} sets its intensity.

The dynamics of $M(t)$ incorporates a characteristic relaxation timescale, ensuring that media responses exhibit temporal inertia consistent with real information dissemination processes. In addition, the model includes an amplification mechanism that depends on the recent growth rate of the epidemic. When infection levels are increasing, media attention intensifies, whereas declining epidemic trends reduce media amplification. This asymmetric response captures the tendency of information systems to react more strongly to worsening conditions than to improvements.

$$\frac{dM}{dt} = \frac{A(t) \, I_{tot}(t - \Delta) - M(t)}{\tau} \tag{6}$$

with $I_{tot-\Delta}$ accounting for the total incidence in the system delayed by Δ and $A(t)$ is the amplification mechanism that allows for the asymmetric response depending on the epidemic growth $g(t)$:

$$g(t) = I_{tot}(t) - I_{tot}(t-1) \tag{7}$$

this way $A(t)$ takes the form:

$$A(t) = \begin{cases} 1 - a & \text{if } g(t) < g_{thr} \\ 1 + a & \text{otherwise} \end{cases} \tag{8}$$

defining g_{thr} as the threshold for the change in the response.

Through this mechanism, media intensity acts as a global driver of behavioral change that simultaneously affects all communities.

Local Prevalence Awareness In addition to global media influence, the model incorporates behavioral responses based on local epidemic conditions. Individuals adjust their behavior according to the prevalence of infection within their immediate social environment, which is represented either through community-level infection fractions in mean-field simulations or through local neighbor infection ratios in network-based simulations.

This local awareness mechanism introduces spatial heterogeneity into the epidemic dynamics, allowing different communities to experience distinct behavioral responses depending on their local epidemic trajectories. Mathematically $f_i^{(local)}(t)$ for community i is defined as:

$$f_i^{(local)}(t) = 1 - \psi_i \, P_i(t) \, \varepsilon_{local} \tag{9}$$

where ψ_i is the community dependent sensitivity to the local conditions, $P_i(t)$ represents the disease prevalence that, depending on the simulation type can be: the overall prevalence for the mean-field dynamics or the community prevalence for the network scheme. Finally, ε_{local} is a free parameter adjusting the intensity of the local awareness.

Community ideological factor The last factor that determines the infection rate accounts for how different communities align with respect to the outbreak:

$$g(\theta_i) = 1 + \alpha_{risk}(2\theta_i - 1) \tag{10}$$

with θ_i defining the ideological alignment of each community and α_{risk} modulate the intensity of the response.

3.1.2 Force of Infection and Transmission Dynamics

The force of infection in each community depends on both within-community interactions and cross-community contacts. A homophily parameter regulates the relative contribution of internal versus external interactions, thereby controlling the degree of social mixing between communities.

In the mean-field formulation, the force of infection is computed using a mixing matrix W_{ij} derived from the underlying interaction network. This matrix encodes the probability of contact between individuals belonging to different communities, ensuring consistency between the network structure and the aggregated compartmental dynamics.

The resulting transmission process reflects the combined influence of social structure, behavioral adaptation, and information dynamics:

$$\lambda_i(t) = \beta_i(t) \left[h_i I_i(t) + (1 - h_i) \sum_{j \neq i} W_{ij} I_j(t) \right] \tag{11}$$

3.1.3 Epidemic Dynamics

Compartmental Evolution The epidemic evolution is described by standard compartmental equations governing transitions between susceptible, exposed, infected, and removed states. The specific formulation depends on the chosen epidemic variant. In models that include an exposed compartment, individuals progress from susceptibility to infection through an intermediate latent stage, while in simpler formulations the transition occurs directly.

Recovery processes are governed by constant rates γ that determine the duration of infectiousness. In SIS variants, recovered individuals return to the susceptible state, whereas in SIR formulations they remain permanently removed.

These compartmental dynamics operate within each community while interacting through cross-community transmission terms.

Network-Based Stochastic Dynamics In addition to deterministic mean-field dynamics, the model includes a stochastic network-based simulation mode. In this formulation, individuals occupy discrete epidemiological states and transitions occur probabilistically according to infection and recovery rates.

The infection probability for each individual depends on the proportion of infected neighbors, capturing local interaction effects that are not represented in aggregated models. This stochastic formulation allows the investigation of fluctuations, extinction events, and finite-size effects that arise in realistic population networks.

3.1.4 Observables and Model Outputs

The simulation framework produces a range of observables describing epidemic and behavioral dynamics. These include time series of infection prevalence at both global and community levels, the evolution of compartmental fractions, and the trajectory of media intensity. The model also provides information on the effective transmission rates, allowing direct analysis of behavioral modulation effects.

These outputs enable comprehensive characterization of epidemic evolution under varying behavioral and informational conditions.

3.1.5 Model Capabilities and Applications

The implemented framework enables systematic investigation of how behavioral adaptation and information processes influence epidemic spreading in heterogeneous populations. In particular, it allows the study of the interplay between media dynamics and local awareness, the impact of ideological heterogeneity on transmission, and the role of community structure in shaping epidemic trajectories.

By supporting both mean-field and network-based simulations, the model provides a versatile platform for exploring epidemic processes across different levels of abstraction.

3.1.6 Model analysis

In this section we present the results obtained from simulations of the community-structured epidemic model described above. The analysis focuses on how the interaction between behavioral adaptation mechanisms, media dynamics, and community heterogeneity shapes the temporal evolution of the epidemic.

The simulations explore different configurations of the behavioral modulation terms governing the effective transmission rate. The figures illustrate representative epidemic trajectories and highlight the role played by the different components of the model.

Baseline epidemic dynamics without behavioral or media modulation Figure 1 illustrates the baseline epidemic dynamics obtained when all behavioral modulation mechanisms are deactivated. In this configuration the parameters controlling media influence, local prevalence awareness, and ideological risk perception are set to zero ($\varepsilon_{local} = 0$, $\varepsilon_{media} = 0$, and $\alpha_{risk} = 0$). As a consequence, the effective transmission rate remains constant over time and identical across communities, reducing the system to a standard compartmental epidemic model embedded within a community-structured network.

The figure compares deterministic mean-field simulations with stochastic network-based realizations for both SIR and SIS epidemic formulations. The system consists of five communities with heterogeneous levels of homophily ranging from 0.5 to 0.9. Community-specific parameters are assigned independently of the homophily level, ensuring that no systematic correlation between structural connectivity and behavioral traits is present in this baseline scenario.

Panels (a) and (b) display the temporal evolution of the infected fraction for the SIR model. In the deterministic mean-field formulation (panel a), all communities exhibit, as expected, identical epidemic trajectories, and their curves perfectly overlap with the global infection dynamics. The stochastic simulations on a network of $N = 1000$ agents shown in panel (b) exhibits a qualitatively similar epidemic profile but with visible fluctuations across communities. These variations arise from finite-size stochastic effects and the discrete nature of network interactions. Individual communities experience slightly different epidemic peaks and timing of infection growth, yet the global infection curve remains consistent with the deterministic mean-field prediction.

Panels (c) and (d) present the corresponding results for the SIS epidemic model. In the mean-field formulation (panel c), the system converges toward a stationary endemic equilibrium characterized by a stable fraction of infected individuals. As in the SIR case, all community-level trajectories coincide with the global curve due to the absence of heterogeneous transmission mechanisms. The stochastic network realization shown in panel (d) again introduces fluctuations around the deterministic solution. After an initial growth phase, the system approaches an endemic regime in which the infection level oscillates around a quasi-stationary value. Community-level infection fractions exhibit persistent variability due to stochastic transmission events and differences in local network connectivity, while the global infection level remains close to the deterministic equilibrium predicted by the mean-field model.

Overall, Figure 1 establishes the reference dynamics of the model in the absence of behavioral feedback mechanisms.

Isolated effects of media influence, local awareness, and community ideology Figure 2 illustrates the impact of the three behavioral modulation mechanisms introduced in the model when each of them is activated independently. The simulations are performed using the mean-field formulation of the SIS epidemic dynamics including the exposed compartment. In all cases the behavioral parameters associated with the other mechanisms are set to zero, allowing the individual contribution of each factor to be isolated.

Panel (a) shows the epidemic evolution when only the media effect is active ($\varepsilon_{media} = 1$). In this configuration the effective transmission rate is modulated by the global media intensity $M(t)$ and the community sensitivity ϕ_i. Since the media influence acts as a global signal affecting all communities simultaneously, the small differences observed in the epidemic trajectories of the different communities are largely due to ϕ that in the chosen setup ranges between 0.1 and 0.6.

Panel (b), instead, presents the case in which only the local prevalence awareness mechanism is activated ($\varepsilon_{local} = 1$). Unlike the media signal, which depends on global infection levels, the local awareness factor is directly linked to the infection prevalence within each community. As a consequence, the effective transmission rate becomes community-dependent, reflecting the different local epidemic conditions experienced by each group. This produces a visible separation between the infection trajectories of the five communities. Groups experiencing higher local prevalence exhibit stronger behavioral responses and therefore slightly reduced effective transmission rates, while communities with lower infection levels show weaker behavioral adjustments. The resulting dynamics lead to moderate heterogeneity in the endemic infection levels across communities.

Panel (c) illustrates the effects of activating only the ideological modulation term ($\alpha_{risk} = 1$). In this case the effective transmission rate differs across communities due to their intrinsic ideological alignment

Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA

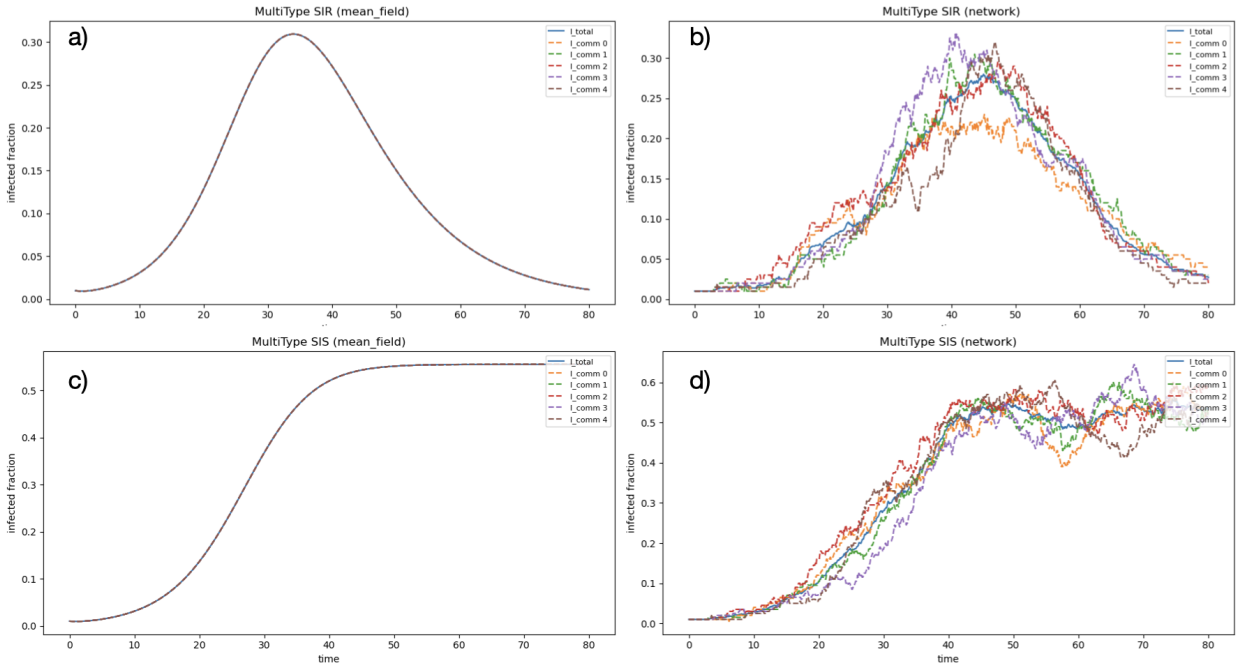

Figure 1: Baseline epidemic dynamics without behavioral modulation. Panels (a,b) show SIR simulations, while panels (c,d) show SIS dynamics. Left panels correspond to deterministic mean-field simulations and right panels to stochastic network realizations on a network of $N = 1000$ nodes obtained with a stochastic block model with 5 groups and values of the homophily ranging from $0.5 \leq h \leq 0.9$. The epidemic parameters are: $\beta_0 = 0.45$ and $\gamma = 0.1$.

with respect to risk perception θ. Communities characterized by stronger precautionary attitudes experience lower effective transmission rates, leading to reduced endemic infection levels, whereas communities with weaker risk sensitivity exhibit higher infection prevalence. This mechanism generates the most pronounced separation between community-level epidemic trajectories among the three scenarios considered. The resulting stationary infection levels differ substantially across groups, reflecting persistent structural heterogeneity introduced by ideological predispositions.

Overall, Figure 2 highlights the distinct roles played by the three behavioral mechanisms. Media influence acts as a global regulator that preserves synchronization across communities, local awareness introduces moderate heterogeneity through community-specific feedback from infection prevalence, and ideological alignment is only related to structural differences in transmission dynamics across groups.

Combined behavioral effects in SIR and SEIR dynamics Figure 3 presents the epidemic dynamics obtained when all behavioral mechanisms are simultaneously active. In this configuration the parameters controlling the three modulation mechanisms are set to $\varepsilon_{media} = 1$, $\varepsilon_{local} = 1$, and $\alpha_{risk} = 1$. Communities are assumed to be uncorrelated, meaning that their structural properties and behavioral parameters are assigned independently. The figure compares two epidemic formulations within the deterministic mean-field framework: a standard SIR model (panel a) and an SEIR model including an exposed compartment (panel

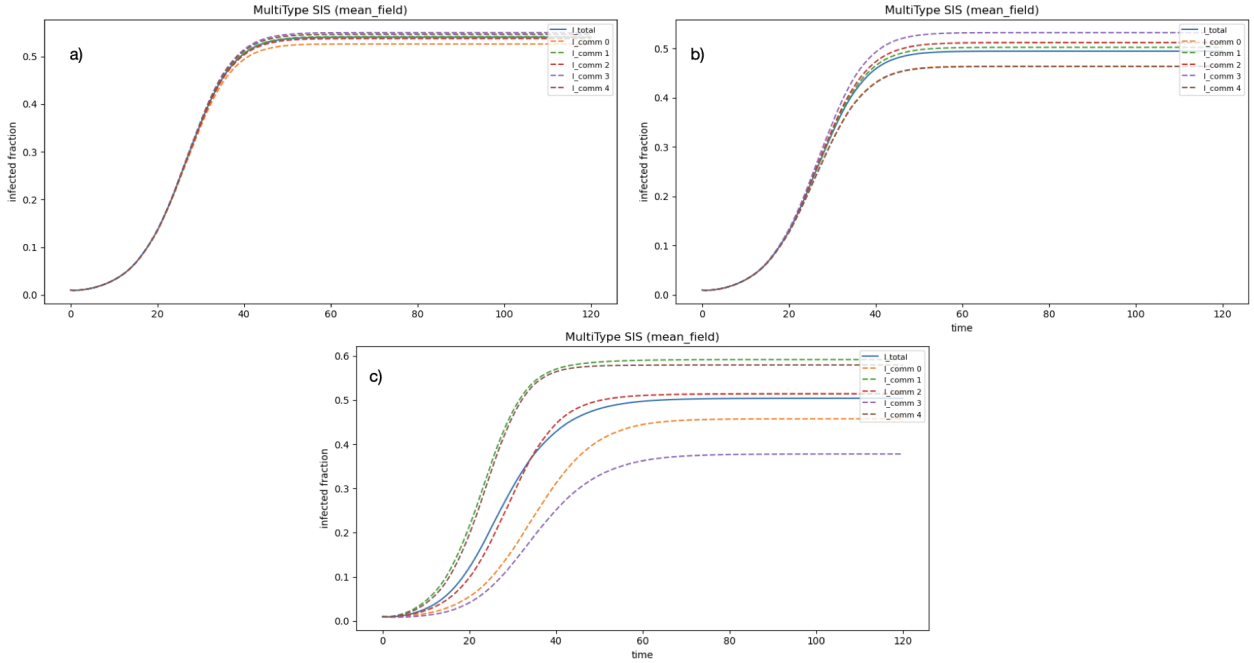

Figure 2: Effects of the behavioral modulation mechanisms when activated independently in the mean-field SIS model with exposed compartment. Panel (a) shows the media influence effect ($\varepsilon_{media} = 1$), panel (b) the local prevalence awareness effect ($\varepsilon_{local} = 1$), and panel (c) the community ideological modulation ($\alpha_{risk} = 1$). The epidemic parameters are: $\beta_0 = 0.45$ and $\gamma = 0.1$.

b).

When all behavioral components are active, the effective transmission rate becomes dynamically heterogeneous across communities due to the combined influence of global media signals, local prevalence feedback, and intrinsic ideological predispositions. As a consequence, the epidemic trajectories differ substantially across communities, producing a clear separation between infection curves. These differences emerge because each community experiences distinct behavioral adjustments driven by both its internal infection dynamics and its ideological alignment with respect to risk perception.

Panel (a) shows the results for the SIR formulation. In this case the epidemic follows the characteristic outbreak profile with a single infection peak followed by a gradual decay toward extinction. The peak infection levels differ significantly between communities, reflecting the heterogeneous effective transmission rates generated by the ideological modulation and the community-specific behavioral responses. Communities characterized by higher effective transmission experience earlier and larger infection peaks, while those with stronger behavioral responses display lower peak prevalence and slower epidemic growth.

Panel (b) illustrates the corresponding results for the SEIR formulation, where the presence of an exposed compartment introduces a latent period between infection and infectiousness. This additional stage modifies the temporal structure of the epidemic by slowing the transmission process and delaying the onset of the epidemic peak. As a consequence, the infection curves shift toward later times compared to the SIR dynamics. Moreover, the exposed compartment amplifies the differences between communities by allowing

small variations in transmission dynamics to accumulate over the latent phase, resulting in more pronounced divergence in the timing and magnitude of the epidemic peaks.

Overall, the comparison between panels (a) and (b) highlights how epidemiological structure interacts with behavioral heterogeneity. While the combined behavioral mechanisms already generate significant community-level variability in the SIR formulation, the inclusion of the exposed compartment further enhances these differences by introducing an additional dynamical timescale that magnifies the effects of heterogeneous transmission modulation.

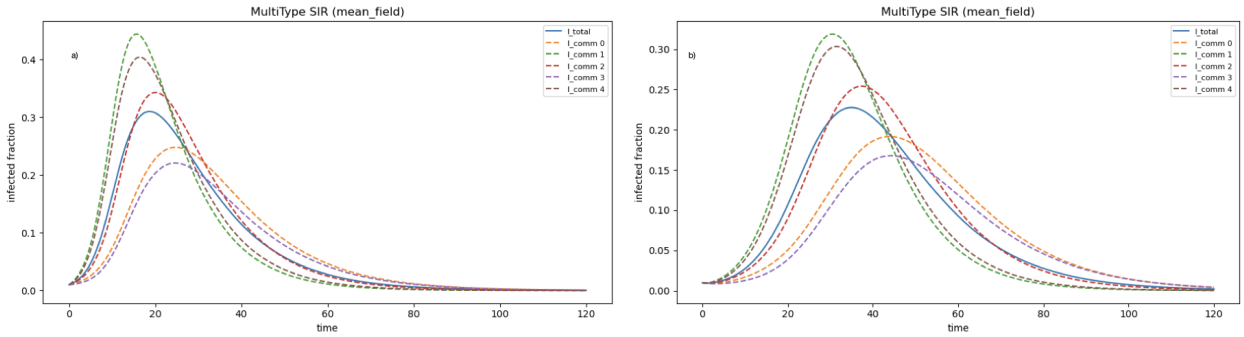

Figure 3: Epidemic dynamics with all behavioral mechanisms active ($\varepsilon_{media} = 1$, $\varepsilon_{local} = 1$, $\alpha_{risk} = 1$) for uncorrelated communities. Panel (a) shows the SIR mean-field dynamics, while panel (b) presents the SEIR formulation with an exposed compartment. The inclusion of the exposed state delays the epidemic dynamics and enhances the differences between community-level infection trajectories. The epidemic parameters are: $\beta_0 = 0.45$ and $\gamma = 0.1$.

Impact of correlations between community structure and behavioral parameters Figure 4 examines the role of correlations between structural and behavioral characteristics of communities. In previous simulations the community-specific parameters controlling local awareness, media sensitivity, and ideological alignment (ψ_i, ϕ_i, and θ_i) were assigned independently of the structural properties of the communities, particularly their homophily level. In this experiment, instead, these quantities are explicitly correlated with the homophily parameter h_i, allowing communities that are more internally connected to also exhibit stronger behavioral responses.

The left panel reproduces the baseline configuration in which the community parameters remain uncorrelated with homophily. In this scenario the epidemic trajectories differ across communities due to behavioral modulation mechanisms, but the structural organization of the network does not systematically reinforce these differences.

The right panel shows the results obtained when correlations are introduced between homophily and the behavioral parameters. Five communities are considered. Community 0 is characterized by strong internal cohesion, with homophily $h = 0.85$ and correspondingly high values of the behavioral parameters ψ, ϕ, and θ. Community 1 exhibits intermediate values for both structural and behavioral parameters ($h = 0.7$). Community 2 represents a neutral reference group with homophily $h = 0.5$ and vanishing behavioral parameters, effectively removing both ideological and behavioral modulation. Finally, communities 3 and 4 are characterized by weak internal cohesion, with homophily values $h = 0.2$ and $h = 0.1$, respectively.

The introduction of these correlations significantly modifies the epidemic dynamics. Communities with

high homophily and strong behavioral parameters exhibit a marked amplification of their epidemic trajectories compared with the uncorrelated scenario. The high internal connectivity reinforces within-community transmission, while the aligned behavioral parameters produce coherent responses that affect the effective transmission rate in a systematic way. As a result, these groups experience more pronounced infection peaks relative to the baseline case.

Conversely, communities characterized by low homophily display a weaker epidemic response. Because their internal connectivity is limited, transmission is more strongly influenced by interactions with other groups, which reduces the impact of their own behavioral parameters. This leads to smaller infection peaks and a reduced sensitivity to the behavioral modulation mechanisms.

The neutral community (community 2) behaves qualitatively differently from the others. Since its behavioral parameters vanish and its homophily remains moderate, the effective transmission rate is not significantly modulated by either ideological or behavioral effects. As a consequence, the infection level remains extremely low and the community experiences little to no epidemic outbreak.

Overall, these results highlight the importance of correlations between social structure and behavioral characteristics. When communities with strong internal connectivity also exhibit coherent behavioral responses, structural and behavioral effects reinforce each other, producing substantially larger differences in epidemic dynamics across the population.

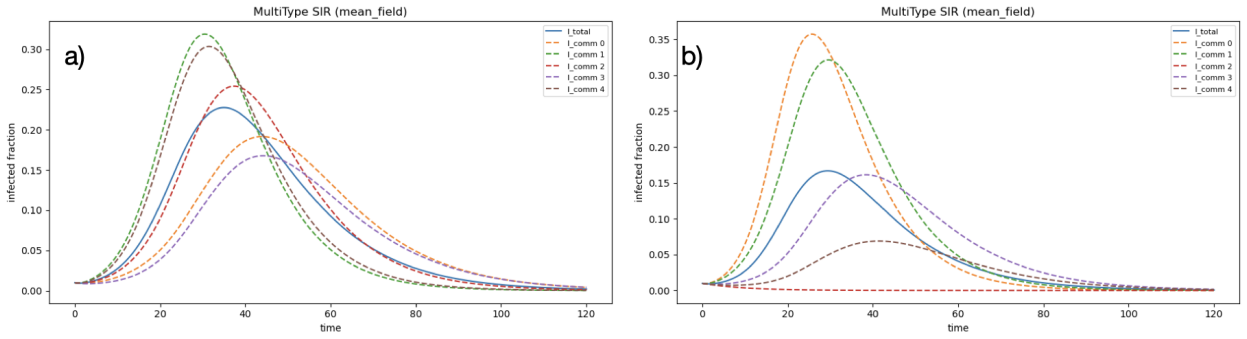

Figure 4: Effect of correlations between community structural properties and behavioral parameters. The left panel shows the epidemic dynamics for the configuration in which the community parameters (ψ_i, ϕ_i, θ_i) are assigned independently of the homophily level h_i. The right panel shows the correlated scenario, where the behavioral parameters are aligned with community homophily. In this case community 0 is characterized by high homophily ($h = 0.85$) and high behavioral parameters, community 1 has intermediate values ($h = 0.7$), community 2 represents a neutral reference group ($h = 0.5$) with vanishing behavioral parameters, while communities 3 and 4 have low homophily ($h = 0.2$ and $h = 0.1$ respectively). The correlation between structural cohesion and behavioral response leads to significantly different epidemic outcomes across communities, amplifying the outbreak in highly cohesive groups while reducing its impact in weakly connected ones.

3.1.7 Conclusion

The community-structured epidemic model developed in this report integrates classical compartmental dynamics with behavioral modulation mechanisms driven by media influence, local prevalence awareness, and ideological heterogeneity. By allowing the effective transmission rate to evolve dynamically in response to

both global information signals and local epidemic conditions, the framework captures key feedback processes that are typically neglected in traditional epidemic models. In particular, the model explicitly links epidemic dynamics to social structure and information flow, providing a flexible platform for investigating how behavioral adaptation modifies disease propagation in heterogeneous populations.

Overall, the presented results demonstrate that epidemic dynamics cannot be fully understood without accounting for the interaction between information processes, behavioral adaptation, and community structure. Global information signals, local behavioral responses, and ideological heterogeneity act on different spatial and temporal scales, and their combined effects can generate complex patterns of epidemic heterogeneity across social groups.

The modular structure of the model makes it suitable for a wide range of future extensions. Possible developments include the introduction of adaptive contact networks, feedback mechanisms between epidemic dynamics and opinion formation, or more detailed representations of information diffusion processes. Such extensions would further enhance the ability of the framework to capture the coupled evolution of social behavior and epidemic spreading in complex societies.

3.2 An Integrated Model for Opinion Dynamics with Delayed Risk Perception

Understanding how individual opinions evolve in response to external crises requires models capable of integrating social interaction mechanisms with dynamic perceptions of risk. In many real-world contexts, particularly during public health emergencies, individuals do not react directly to objective epidemiological indicators. Instead, their behavioral responses are mediated through subjective risk perception processes that incorporate temporal delays, media amplification, and social influence structures.

The model implemented in this section addresses this challenge by combining three coupled dynamical components: a mean-field epidemic process, a network-based opinion dynamics system, and a delayed risk perception mechanism. These components interact through a unidirectional coupling structure in which epidemic prevalence drives risk perception, which in turn influences opinion evolution. This approach allows the study of how external shocks propagate through perceptual and social layers before manifesting as collective opinion changes.

The framework is designed to capture key phenomena observed in crisis contexts, including risk-induced polarization, delayed behavioral responses, and the role of ideological community structure in shaping social influence.

3.2.1 Model Architecture and Mathematical Framework

The implemented model consists of three interacting subsystems. The first subsystem describes epidemic spreading through a classical mean-field susceptible–infected–removed formulation. The second subsystem models opinion evolution on a social network characterized by static ideological communities. The third subsystem represents the dynamics of perceived risk, which acts as an intermediary between epidemiological conditions and social responses.

The coupling between these subsystems follows a hierarchical structure. Epidemic prevalence determines the evolution of perceived risk through a delayed filtering process. The resulting perception signal then enters the opinion dynamics as an external forcing term. In the implemented version of the model, no feedback from opinions to epidemic dynamics is considered, ensuring a one-directional flow of influence from epidemiological conditions to social attitudes.

Social Network and Community Structure The social contacts are represented by an undirected network whose nodes correspond to individuals and whose edges encode interaction pathways. The network

is divided in a fixed number of communities representing different ideologies and the homophily between individuals. This way, the community structure influences opinion evolution by introducing a tendency toward intra-group consensus. The presence of communities allows the model to capture the coexistence of local social influence and group-level ideological cohesion, two mechanisms that jointly shape collective opinion dynamics.

Epidemic Dynamics Once described the structure of the interactions, the epidemic dynamics is described by a normalized SIR model in which the population is divided into susceptible, infected, and removed fractions. The dynamics are governed by the transmission β and recovery rates μ, and the total population is conserved at all times.

$$\frac{dS}{dt} = -\beta SI \tag{12}$$

$$\frac{dI}{dt} = \beta SI - \mu I \tag{13}$$

$$\frac{dR}{dt} = \mu I \tag{14}$$

Where $S(t) + I(t) + R(t) = 1$.

The model equations are integrated numerically using an explicit time-stepping scheme, which allows efficient coupling with the slower perception and opinion dynamics.

The prevalence of the infection at each time, represents the real risk of being infected and serves as the input signal driving risk perception.

Opinion Dynamics Each individual is characterized by a continuous opinion variable defined within a bounded interval representing the spectrum of attitudes toward a crisis-related issue $o_i(t) \in [-1, 1]$. The evolution of opinions is driven by three distinct mechanisms.

The first mechanism corresponds to local social influence, whereby individuals adjust their opinions toward the average opinion of their neighbors in the network. The second mechanism represents community-level attraction, which pulls individual opinions toward the mean stance of their ideological group. The third mechanism is an external forcing term derived from perceived risk. The combination of these processes produces a dynamic balance between conformity pressures at different social scales and external influences originating from the crisis environment and can be described by:

$$\frac{do_i}{dt} = \alpha_s \left(\langle o \rangle_{\mathcal{N}_i} - o_i \right) + \alpha_c \left(\langle o \rangle_{c_i} - o_i \right) + \text{External}_i(t) \tag{15}$$

where α_s and α_c represents the local social influence and the community cohesion respectively, while the average opinion of the neighbors and the communities are $\langle o \rangle_{\mathcal{N}_i}$ and $\langle o \rangle_{c_i}$.

The perceived risk enters the opinion dynamics as the external driving force that can be modulated in different ways. For now, we consider two alternative formulations for this influence. In the first one, risk perception induces a uniform directional shift in opinions, representing collective movement toward precautionary attitudes:

$$\text{External}_i(t) = \alpha_R R_p(t) \tag{16}$$

where $R_p(t)$ comes from the risk perception dynamics while α_R modulates its intensity.

A second option second formulation, the external influence amplifies existing opinion positions, causing individuals with strong prior views to become more extreme. This nonlinear mechanism captures the emergence of polarization driven by crisis-induced emotional and cognitive responses. Mathematically it can be expressed as:

$$\text{External}_i(t) = \alpha_R R_p(t)\,\text{sign}(o_i)|o_i|^p \tag{17}$$

Combined, these alternative modes enable the exploration of different sociobehavioral scenarios, ranging from consensus formation to radicalization dynamics.

Risk Perception Dynamics Risk perception $R_p(t)$ is modeled as a dynamic variable representing the collective subjective assessment of epidemic danger. Rather than responding instantaneously to current prevalence, perception evolves according to a delayed filtering process. This delay accounts for the time required for information to spread through media channels and for individuals to interpret and internalize new information. The perception dynamics are governed by a relaxation equation in which perceived risk gradually approaches a delayed version of epidemic prevalence. The rate of adjustment is controlled by a characteristic response timescale.

$$\frac{dR_p}{dt} = \frac{A(t)I(t-\Delta) - R_p(t)}{\tau} \tag{18}$$

where $A(t)$ quantifies the effect of media amplification while Δ stands for a fixed delay due to time needed to spread the information and τ the relaxation time of the information dynamics.

The transformation from epidemic prevalence to perceived risk includes a media amplification factor that depends on the recent growth rate of the epidemic. When infection levels are increasing rapidly, the amplification factor strengthens the perception signal, reflecting heightened media attention and public concern. Conversely, when the epidemic is declining, the amplification factor reduces perceived urgency. To model this situation, the growth rate of the epidemic $g(t) = I(t) - I(t-1)$ is used to define media amplification $A(t)$ as:

$$A(t) = \begin{cases} 1-a & \text{if } g(t) < g_{thr} \\ 1+a & \text{otherwise} \end{cases} \tag{19}$$

with g_{thr} being the threshold after which amplification starts and a modulates its intensity.

3.2.2 Model analysis

Once implemented the model, we studied how the effects of an external crisis propagate through perceptual and social layers to produce collective opinion outcomes. In particular, we focused on the conditions that allow polarization to emerge and how it can be avoided. To do that, five different communities are considered, each with a distinct initial opinion distribution covering the entire spectrum $1 < o_i < 1$. In particular, we set two communities to extreme values $\langle o \rangle = -0.4$ and $\langle o \rangle = 0.7$, other two with opposite moderate initial opinion $\langle o \rangle = \pm 0.2$ and a neutral one $\langle o \rangle = 0.0$.

To quantify system behavior, several macroscopic observables are considered. The global average opinion $\bar{o}(t)$ measures overall alignment, while the variance $\sigma_o^2(t)$ captures polarization intensity. Community-level averages allow the assessment of inter-group divergence.

The comparison between $I(t)$ and $R_p(t)$ also reveals the magnitude of perceptual lag and media amplification.

Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA

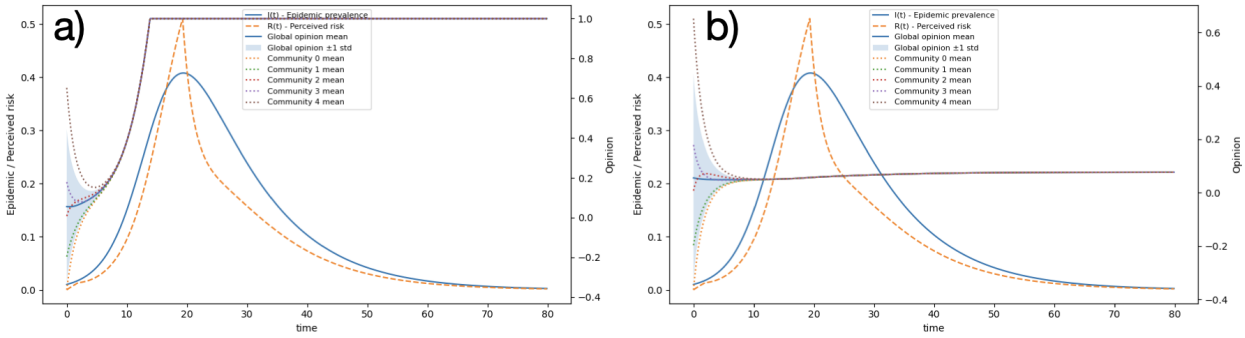

Figure 5: Temporal evolution of epidemic prevalence $I(t)$, perceived risk $R_p(t)$, global mean opinion, opinion standard deviation, and community-level mean opinions under two alternative formulations of the external forcing term. **Panel (a):** Precautionary external influence (type = 0), corresponding to the linear formulation in Eq. (16). Perceived risk induces a uniform directional shift in opinions, leading to transient global alignment across communities. **Panel (b):** Polarization-inducing external influence (type = 1), corresponding to the nonlinear formulation in Eq.(17). The crisis signal amplifies pre-existing opinion magnitudes, reinforcing inter-community divergence. Common parameters: $\alpha_s = 1.0$, $\alpha_c = 1.0$, $\alpha_R = 1.0$, media amplification intensity $a = 0.5$, growth threshold $g_{thr} = 5 \times 10^{-5}$, delay $\Delta = 20$, and relaxation time $\tau = 2.0$. Epidemic dynamics follow the SIR model defined in Eq.(12).

Consensus versus Polarization Regimes We start by analyzing the effect of the structure of the external forcing term as, by construction, it determines whether the system converges toward consensus or amplifies pre-existing polarization.

In the prudential formulation, given by Eq. (16), perceived risk induces a uniform directional shift across individuals. When α_R dominates over α_s and α_c, we expect opinions to align globally during high-risk phases. If community cohesion α_c is strong but symmetric across groups, the system may exhibit temporary convergence followed by re-fragmentation once risk perception declines.

In the nonlinear amplification formulation of Eq. (17), instead, the dynamics differ qualitatively. Because the external field scales with $|o_i|^p$, individuals with stronger prior opinions experience larger reinforcement. This mechanism increases opinion variance and can drive the system toward bimodal or polarized distributions.

Figure 5 illustrates the temporal evolution of epidemic prevalence $I(t)$, perceived risk $R_p(t)$, the global mean opinion, and community-level opinion averages under two alternative formulations of the external

forcing term. All parameters are identical across panels except for the definition of the external influence in Eq. (15) with with the social, community and risk perception effects having the same importance. $\alpha_s = \alpha_c = \alpha_R = 1.0$ (see caption for the remaining parameters). In both panels, epidemic prevalence exhibits the typical unimodal SIR profile, reaching a peak around $t \approx 20$ and then decaying monotonically. Due to the imposed delay $\Delta = 20$, perceived risk reacts with a significant temporal lag. The peak of $R_p(t)$ occurs after the epidemic peak, and its decay is slower than that of $I(t)$, producing a clear hysteresis effect.

The amplification mechanism defined in Eq. (19) enhances perceived risk during phases of positive epidemic growth. Consequently, $R_p(t)$ rises sharply once the growth threshold is exceeded and remains elevated even as prevalence begins to decline. This delayed and amplified perception constitutes the effective driver of opinion dynamics. In the left panel, the external term follows Eq. (16), producing a uniform directional shift proportional to perceived risk –i.e. directing all individuals toward a more prudent behavior. Since the forcing term does not depend on the sign or magnitude of o_i, all individuals experience the same additive pressure. As a result, the global mean opinion increases synchronously with the rise in perceived risk, reaching a maximum shortly after the perception peak. The standard deviation remains relatively contained, indicating that the shock primarily induces collective alignment rather than dispersion. Community-level means exhibit transient deviations during the early phase, reflecting heterogeneous initial conditions and intra-group cohesion. However, once the perception-driven forcing dominates, all communities converge toward a common trajectory. After perceived risk declines, opinions gradually relax toward a shared steady state determined by the balance between social averaging and community attraction. This regime corresponds to a crisis-induced consensus scenario: the external shock temporarily overrides structural ideological differences and promotes alignment.

In the right panel, instead, the external influence follows Eq. (17), introducing a nonlinear amplification proportional to $\text{sign}(o_i)|o_i|^p$. In this configuration, individuals with stronger prior opinions experience larger reinforcement. Although epidemic and perception dynamics are identical to panel (a), opinion evolution differs qualitatively. Instead of uniform alignment, the external shock amplifies pre-existing ideological separation. Community averages initially move in the direction determined by their internal average, and the nonlinear forcing strengthens these tendencies.

The global mean opinion still exhibits a peak associated with the perception maximum; however, convergence across communities is weaker. The dispersion between group averages persists longer, indicating partial reinforcement of polarization during the crisis phase.vImportantly, the presence of strong community cohesion ($\alpha_c = 1.0$) stabilizes intra-group agreement, while the nonlinear external term prevents full cross-group alignment. The system therefore displays crisis-induced polarization amplification rather than global consensus.

Because α_s, α_c, and α_R are equal, neither social averaging nor community cohesion dominates the dynamics. Therefore, the qualitative difference observed across panels is entirely attributable to the functional form of the external field. This confirms that polarization in this framework is not solely a consequence of network modularity, but emerges from the interaction between ideological structure and nonlinear crisis-driven reinforcement.

Overall, Figure 5 demonstrates that delayed and amplified risk perception can either homogenize or polarize social attitudes depending on how the crisis signal couples to individual opinions.

Role of Local Social Influence α_s in the Polarization Regime Figure 6 shows the effect of increasing the strength of local social averaging α_s under nonlinear external forcing of Eq. 17. All simulations share identical parameters except for α_s which we increase from $\alpha_s = 0$ to $\alpha_s = 0.125$. Epidemic prevalence and perceived risk are identical across panels and the same as in Figure 5.

When local social averaging is absent $\alpha_s = 0.0$, opinion evolution is governed only by community cohe-

Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA

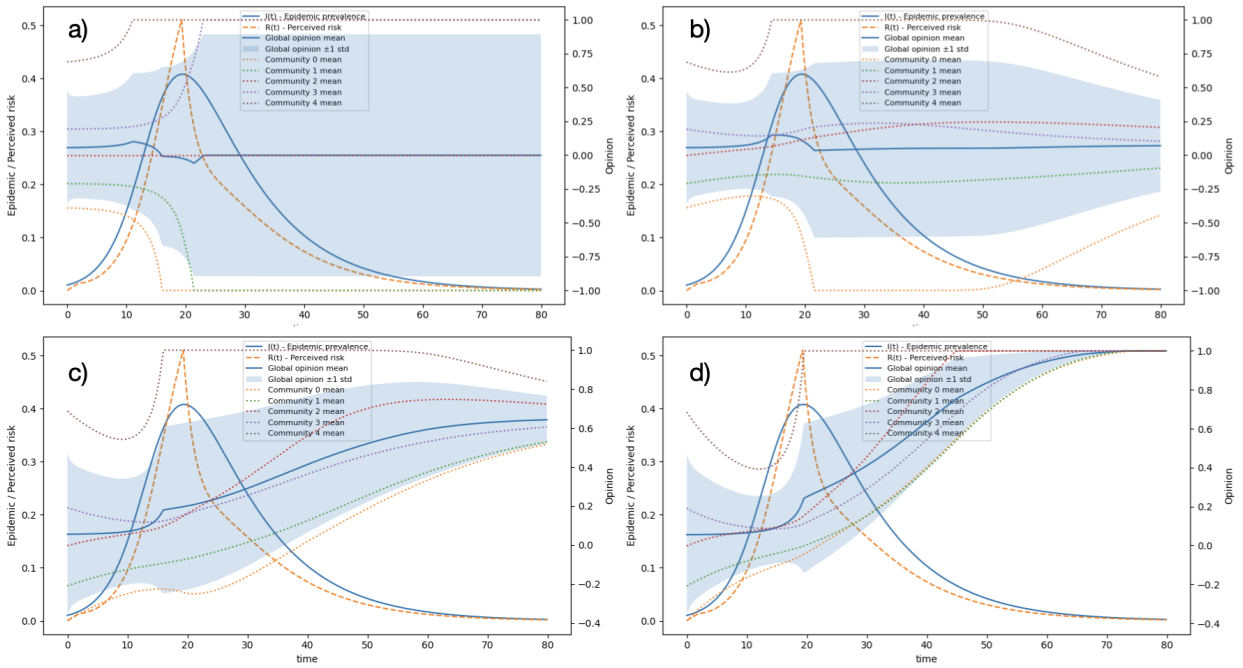

Figure 6: Effect of increasing local social influence α_s under nonlinear external forcing (type = 1). Panel (a) $\alpha_s = 0.0$, panel (b) $\alpha_s = 0.05$, panel (c) $\alpha_s = 0.075$, panel (d) $\alpha_s = 0.125$. Epidemic and perceived risk dynamics are identical across panels. As α_s increases, local averaging progressively counteracts nonlinear crisis-induced amplification, driving the system from persistent polarization (panel a) toward consensus formation (panel d). Other parameters: $\alpha_c = 1.0$, $\alpha_R = 1.0$, $a = 0.5$, $g_{thr} = 5 \times 10^{-5}$, $\Delta = 20$, $\tau = 2.0$.

sion and nonlinear external reinforcement. During the perception peak, the amplification term in Eq. (17) strengthens pre-existing opinion signs. Because there is no cross-neighborhood averaging, ideological communities evolve largely independently. This produces strong divergence between group means and a broad dispersion in global opinion, as indicated by the wide variance band. The crisis shock therefore maximizes polarization when $\alpha_s = 0$.

On the other side, introducing weak social averaging ($\alpha_s = 0.05$) begins to couple neighboring individuals across community boundaries. Although nonlinear reinforcement remains active, local averaging partially counteracts extremization. Community means still diverge during the high-perception phase, but the magnitude of separation is reduced compared to panel (a) and, in the long-term, opinions converge toward a middle ground. This regime illustrates a transition from persistent polarization to gradual consensus restoration.

For intermediate values of $\alpha_s = 0.075$, the competition between nonlinear amplification and social averaging becomes balanced. During the epidemic growth phase, polarization still increases; however, after the perception peak, averaging progressively dominates. Community averages start converging toward a shared trajectory, and the global variance decreases steadily.

Finally, when local social influence becomes sufficiently strong $\alpha_s = 0.125$, it dominates over nonlinear external reinforcement. Even though perceived risk reaches the same peak as in the other panels, the amplification mechanism cannot sustain inter-community divergence. Instead, opinions across all communities

progressively align toward a common positive state.

Overall, Figure 6 demonstrates that local social averaging acts as a stabilizing mechanism against crisis-induced polarization. While nonlinear risk amplification tends to increase ideological separation, sufficiently strong α_s restores cross-group coupling and promotes convergence. The transition observed across panels reflects a dynamical competition between three forces: community anchoring (α_c), nonlinear external reinforcement (α_R), and neighborhood averaging (α_s). Increasing α_s shifts the system from a polarization-dominated regime to a consensus-dominated regime, even under identical epidemic and perception dynamics.

This result highlights that polarization in the model is not solely determined by crisis intensity or media amplification, but critically depends on the strength of local cross-cutting social interactions.

Role of Community Structure Community cohesion plays a central role in determining long-term outcomes. When α_c is small, local network averaging dominates and the system behaves similarly to a classical bounded confidence or averaging model with external forcing. In this regime, external shocks may produce temporary global shifts without persistent polarization.

However, when α_c is large, communities act as attractors in opinion space. Even under strong external forcing, individuals remain anchored to their group mean. In the nonlinear amplification regime, this anchoring effect can generate divergence between communities, as each group amplifies its internally coherent stance.

Therefore, ideological communities introduce structural stability to polarized configurations, particularly when combined with nonlinear risk amplification.

Figure 7 illustrates the effect of increasing community-level attraction α_c under nonlinear external forcing. To show the effect of community structured alone we eliminated local social averaging ($\alpha_s = 0$), while other parameters follow the ones of Figure 5.

Because $\alpha_s = 0$, the only internal stabilizing mechanism is ideological community cohesion. Consequently, the competition shaping opinion evolution is exclusively between nonlinear external reinforcement and group-level attraction.

In the absence of both local averaging and community cohesion (panel (a), $\alpha_c = 0.0$) opinion dynamics are governed entirely by nonlinear external forcing. During the perception peak, Eq. (17) amplifies individual opinions proportionally to their magnitude and sign. In this case, Community averages lose structural meaning because groups no longer exert internal attraction, and the global variance remains high over time. Surprisingly, even an extremely weak community attraction (panel b, $\alpha_c = 0.05$) begins to stabilize intra-group alignment, effectively reducing global opinions variance. While each community still follows a separate opinion and amplification still reinforces initial differences, individuals within the same group begin to move coherently. For $\alpha_c = 0.2$ (panel c), instead, communities act as effective attractors in opinion space. During the crisis phase, the nonlinear external field pushes groups in directions consistent with their internal mean, but intra-group dispersion decreases due to stronger anchoring. As a consequence, polarization becomes structured and persistent. The system stabilizes around two extremely polarized groups. Finally, when community attraction is strong (panel d,$\alpha_c = 1.0$), the system behave exactly in the same way as for moderate

Taken together, the results in Figure 7 highlights the critical role of ideological cohesion in sustaining crisis-induced polarization when cross-cutting social influence is absent. Increasing α_c does not reduce polarization; instead, it organizes and stabilizes it. In contrast to the α_s analysis (Figure 2), where stronger local averaging suppressed divergence, here stronger community attraction enhances structural persistence of polarized states. The model therefore distinguishes clearly between two mechanisms: local averaging (α_s) promotes cross-group mixing and consensus and community cohesion (α_c) reinforces within-group alignment and stabilizes ideological fragmentation.

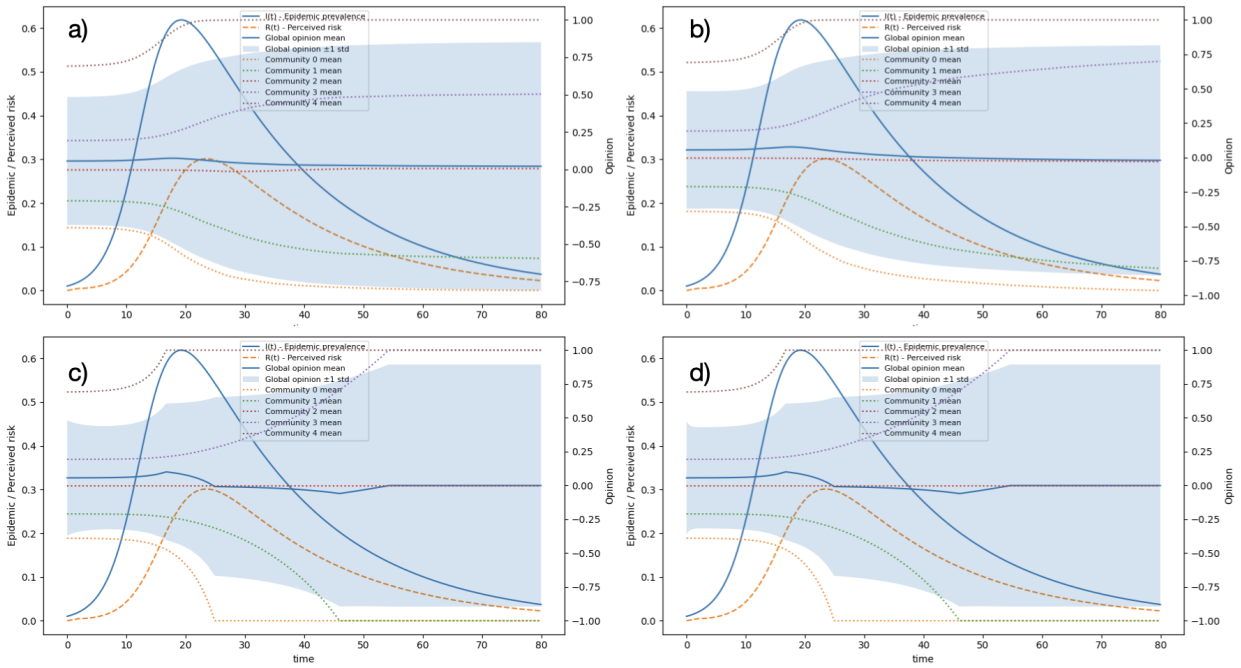

Figure 7: Effect of increasing community cohesion α_c under nonlinear external forcing and in the absence of local social averaging ($\alpha_s = 0$). Panel (a) $\alpha_c = 0.0$, (b) $\alpha_c = 0.01$, (c) $\alpha_c = 0.2$, (d) $\alpha_c = 1.0$. Epidemic and perceived risk dynamics are identical across panels. Increasing α_c strengthens intra-group alignment and stabilizes inter-community divergence, leading to structurally persistent polarization. Other parameters: $\alpha_R = 0.5$, $a = 0.5$, $g_{thr} = 5 \times 10^{-5}$, $\Delta = 20$, $\tau = 2.0$.

The interplay between nonlinear crisis amplification and community anchoring determines whether polarization remains transient or becomes structurally embedded.

3.2.3 Conclusions

We introduced and analyzed an integrated framework coupling crisis dynamics, delayed risk perception, and network-based opinion evolution. The model was designed to investigate how external crises propagate through perceptual filters and social interaction layers before manifesting as macroscopic opinion patterns.

The results demonstrate that collective outcomes are not determined solely by epidemic intensity, but critically depend on how perceived risk couples to individual attitudes and on the internal structure of social influence. The delayed perception mechanism introduces a temporal mismatch between objective epidemiological conditions and subjective assessment, generating effects that prolong the social impact of the crisis beyond the epidemic peak. Media amplification further enhances this effect by increasing perceived risk during phases of rapid epidemic growth.

Two qualitatively distinct behavioral regimes emerge depending on the structure of the external forcing. When perceived risk acts as a uniform additive field, crises promote temporary global alignment, leading to consensus even in the presence of ideological communities. In contrast, when the external signal amplifies

pre-existing opinions nonlinearly, polarization can be reinforced and structurally stabilized.

The numerical exploration of parameter space highlights a clear functional distinction between local social averaging and community cohesion. Increasing local cross-cutting interactions (α_s) systematically counteracts nonlinear amplification and drives the system toward consensus. Conversely, increasing ideological cohesion (α_c) strengthens intra-group alignment and stabilizes inter-group divergence, particularly when cross-community mixing is weak. Polarization therefore emerges not merely from external shocks, but from the interaction between nonlinear reinforcement and structurally segregated social organization.

Overall, the model shows that crisis-induced polarization is neither inevitable nor purely externally imposed. It results from the interplay of three competing forces: external risk amplification, local social mixing, and ideological anchoring. The framework provides a flexible tool for studying how delayed perception and social structure jointly determine whether societies converge toward shared responses or fragment into persistent polarized states during large-scale crises.

Future extensions could incorporate feedback from opinion to epidemic dynamics, adaptive network rewiring, or heterogeneous perception delays. Such developments would allow the exploration of co-evolutionary effects between behavior and contagion, further clarifying the mechanisms underlying resilience or fragmentation in complex social systems.

4 Scientific Results

While the modeling framework developed within the CODE project allows the simulation of integrated epidemic and social dynamics, the main scientific effort focused on the study of simplified epidemic models that retain the essential mechanisms included in the full framework.

The goal of this line of research was twofold. On the one hand, simplified models allow analytical investigation of the mechanisms driving epidemic dynamics. On the other hand, they provide reference benchmarks that help interpret the behavior of the more complex models implemented in the toolbox.

The scientific results obtained during the project can be grouped into three main directions:

- the study of epidemic spreading in the presence of behavioral heterogeneity,
- the validation of these mechanisms in data-informed urban settings,
- the analysis of vaccination strategies with fading immunity.

In parallel, we developed methods and datasets for the study of ideological communities and their evolution during the COVID-19 pandemic.

4.1 Multitype Epidemic Models and Behavioral Heterogeneity

A major part of the scientific activity of the project has been devoted to the study of epidemic models with heterogeneous susceptibility and infectivity. These models provide a minimal representation of behavioral and social heterogeneity and constitute a natural benchmark for the more complex frameworks described in the previous sections.

Differences in compliance, exposure and risk perception lead to systematic heterogeneity in transmission parameters across the population. At the same time, social structure and homophily constrain interaction patterns, producing structured mixing between groups. Understanding how these mechanisms jointly affect epidemic thresholds and outbreak dynamics is therefore a key problem in network epidemiology.

To address this problem we introduced a class of multitype epidemic models in which individuals belong to a finite set of types

$$\tau \in \{1, \dots, r\},$$

each characterized by susceptibility and infectivity modifiers

$$(\alpha_S^\tau, \alpha_I^\tau).$$

The framework is agnostic with respect to the origin of heterogeneity, and types may represent behavioral classes, socio-demographic categories, or groups defined by ideological or geographic attributes. In its simplest formulation the model considers two types, corresponding for instance to low-risk and high-risk individuals.

While multitype epidemic models have a long tradition in mathematical epidemiology, the formulation developed in this project explicitly combines type-dependent susceptibility and infectivity with network topology and assortative mixing, allowing a systematic study of the interplay between individual heterogeneity and mesoscale structure.

Multitype SIR dynamics We first considered a multitype extension of the classical SIR model in which the transmission probability between two individuals depends on their types.

Under a heterogeneous mean-field and percolation-based formulation we derived an explicit expression for the epidemic threshold on networks with arbitrary degree distribution and tunable homophily.

Denoting by p the fraction of high-risk individuals and by h a parameter controlling assortative mixing between groups, the critical transmission-to-recovery ratio takes the form

$$\left(\frac{\beta}{\gamma}\right)_c = \frac{\langle k \rangle}{\langle k^2 \rangle - \langle k \rangle} \frac{1}{(1-h)(1-p+\alpha_I \alpha_S p) + h(\alpha_I \alpha_S + 1)}.$$

This expression shows explicitly how epidemic thresholds depend on three distinct factors:

- degree heterogeneity, through the factor $\langle k^2 \rangle / \langle k \rangle$,
- behavioral heterogeneity, through the parameters α_S and α_I,
- group mixing patterns, through the homophily parameter h.

Analytical predictions are in good agreement with numerical simulations on synthetic networks with both homogeneous and heterogeneous degree distributions.

Beyond the determination of epidemic thresholds, the model reveals a non-trivial transient regime occurring close to the threshold. In this regime epidemic trajectories may initially decline before eventually developing into a large-scale outbreak.

This resurgence phenomenon originates from the different effective reproduction numbers associated with different behavioral groups. Infections initially spread mainly among low-risk individuals, where transmission may be close to criticality, while transmission within the high-risk group develops on longer timescales and eventually dominates the dynamics.

The effect is amplified by assortative mixing, which allows high-risk subpopulations to sustain transmission even when global prevalence appears to be decreasing. This mechanism provides a simple explanation of how structured populations may conceal epidemic potential and complicate early assessment of epidemic risk.

Multitype SIS dynamics We extended the same framework to susceptible–infected–susceptible (SIS) dynamics, allowing the study of epidemic persistence and endemic states.

Within the SIS setting epidemic spreading can be described at the individual level by a probabilistic individual-based dynamics whose linear stability is governed by an effective transmission matrix

$$B = A \circ Q,$$

where A is the adjacency matrix of the contact network and Q is a matrix encoding type-dependent transmission weights.

The epidemic threshold is determined by the leading eigenvalue of this operator,

$$\left(\frac{\beta}{\gamma}\right)_c = \frac{1}{\rho(B)},$$

providing a quenched estimate that explicitly incorporates both network structure and type-dependent transmission parameters.

In order to obtain a threshold expression directly in terms of model parameters we also considered an annealed approximation based on a degree-corrected stochastic block model.

In this formulation the adjacency matrix is replaced by its ensemble average, leading to an effective matrix $\tilde{B}$ that incorporates degree heterogeneity and group mixing patterns.

For a population partitioned into r subgroups the matrix $\tilde{B}$ has rank r, implying that the epidemic threshold is governed by a low-dimensional operator describing interactions at the group level.

This reduction provides an analytically tractable framework for structured populations with arbitrary mixing matrices and multiple behavioral classes.

Numerical validation The theoretical predictions obtained within both the SIR and SIS formulations were validated through extensive numerical simulations on synthetic networks.

Simulations confirm the accuracy of both quenched and annealed estimates of the epidemic threshold and reproduce the transient regimes predicted by the analytical theory.

Overall, the multitype epidemic framework developed in this work provides a coherent description of how susceptibility heterogeneity, infectivity heterogeneity and group mixing jointly determine epidemic thresholds and outbreak dynamics.

These models constitute a natural theoretical benchmark for the data-driven simulations described in the following sections.

4.2 Data-driven Epidemic Simulations in Urban Networks

The multitype epidemic framework described in the previous subsection provides a theoretical characterization of the role of susceptibility heterogeneity, infectivity heterogeneity and group mixing. A natural next step is to investigate how the same mechanisms operate in realistic urban environments.

To this end we performed large-scale data-driven epidemic simulations based on the synthetic populations and contact networks described in Section 2.2.2. These simulations allow a quantitative assessment of the impact of behavioral heterogeneity in realistic settings while maintaining a direct connection with the multitype epidemic framework.

The analysis focused on three major Italian cities, Torino, Milano and Palermo, with populations between approximately 6×10^5 and 1.2×10^6 individuals. For each city we constructed a synthetic contact network combining household interactions and social contacts generated through the Urban Social Network model.

Individuals are assigned spatial locations according to high-resolution population density data and belong to one of four age groups. Households are generated using demographic statistics, producing clique-like subgraphs representing strong and persistent contacts. Non-household interactions are generated through a block model in which blocks correspond to combinations of spatial tiles and age groups, with interaction probabilities depending on age-specific contact matrices and geographic distance.

The resulting networks reproduce realistic spatial structure and contact patterns and have an average degree of approximately eleven contacts per individual.

Behavioral heterogeneity Epidemic spreading is simulated using a heterogeneous SIR model consistent with the multitype framework introduced in the previous subsection.

Individuals belong to one of two behavioral classes:

- ordinary individuals (O),
- non-compliant individuals (M).

Non-compliant individuals are characterized by increased infectivity and susceptibility. Denoting by β the baseline transmission rate, transmission probabilities are modified by the susceptibility and infectivity modifiers α_S and α_I:

$$\beta_{OO} = \beta, \qquad \beta_{MO} = \alpha_I \beta, \qquad \beta_{OM} = \alpha_S \beta, \qquad \beta_{MM} = \alpha_I \alpha_S \beta.$$

Recovery occurs at a constant rate γ independent of the behavioral class.

When $\alpha_S = \alpha_I = 1$ the model reduces to the standard SIR process.

Simulations are initialized with a small number of randomly selected infected individuals and repeated multiple times in order to average over stochastic variability.

Uniform behavioral distributions As a baseline scenario we considered a uniform distribution of non-compliant individuals in which each individual belongs to the non-compliant class with probability p_M.

Simulations show that even relatively small fractions of non-compliant individuals substantially modify epidemic outcomes.

The presence of non-compliant individuals lowers the epidemic threshold and increases the attack rate across a wide range of transmission parameters. For moderate transmission rates the heterogeneous model produces substantial increases in final epidemic size compared with the homogeneous SIR model.

The impact of non-compliance is strongest for intermediate values of the transmission rate. Near the epidemic threshold the introduction of non-compliant individuals may transform a subcritical regime into a large outbreak, while at very high transmission rates the marginal effect becomes smaller because transmission is already widespread.

Behavioral heterogeneity also affects epidemic dynamics. Increasing either the fraction of non-compliant individuals or the parameters α_S and α_I produces earlier and higher epidemic peaks, indicating faster epidemic growth and increased stress on healthcare systems.

The roles of susceptibility and infectivity heterogeneity are not identical. While increases in either parameter reduce epidemic thresholds, increased infectivity tends to produce larger epidemic sizes at higher transmission rates, whereas increased susceptibility has a stronger effect close to the threshold.

Data-driven behavioral heterogeneity Beyond uniform behavioral distributions we constructed spatially heterogeneous distributions of non-compliant individuals using electoral data combined with empirical estimates of behavioral attitudes.

For each spatial tile i we defined a non-compliance score r_i derived from electoral results and rescaled to the interval $[0, 1]$. Individuals located in tile i are assigned to the non-compliant class with probability $p_{M,i}$ proportional to r_i, under the constraint

$$\sum_i p_{M,i} N_i = p_M N,$$

where N_i is the population of tile i and N is the total population.

This construction produces spatially heterogeneous behavioral landscapes with localized clusters of non-compliance.

At the city-wide scale the difference between uniform and data-driven behavioral distributions is relatively limited.

However, substantial differences emerge at smaller spatial scales. Areas with higher concentrations of non-compliant individuals may experience significantly larger local attack rates than predicted under uniform assumptions.

The dependence of local epidemic severity on the fraction of non-compliant individuals is strongest for intermediate transmission rates and weakens for highly transmissible diseases.

Role of network structure To isolate the role of realistic contact patterns we compared the results obtained on the original urban networks with those obtained on degree-preserving randomized networks.

Randomization preserves the degree distribution but destroys spatial and demographic correlations.

At the aggregate level epidemic outcomes are only moderately affected by randomization.

At the local level, however, spatial heterogeneity is strongly reduced. In particular, the dependence of local epidemic outcomes on the fraction of non-compliant individuals becomes substantially weaker.

This result shows that realistic urban contact structure plays a crucial role in translating behavioral heterogeneity into spatially heterogeneous epidemic outcomes.

Overall, these results provide a data-driven validation of the multitype epidemic framework and show that even small fractions of non-compliant individuals can substantially affect epidemic dynamics in realistic urban environments.

4.3 Vaccination with Waning Immunity and Optimal Boosting Strategies

The multitype epidemic models described in the previous subsections focus on behavioral heterogeneity in contact patterns and risk management, which affects susceptibility and infectivity through parameters such as α_S and α_I. Another major behavioral and ideological dimension influencing epidemic outcomes is participation in vaccination campaigns.

Vaccination behavior is strongly correlated with political and social attitudes and often exhibits pronounced heterogeneity across subpopulations. Understanding how vaccination strategies affect epidemic outcomes is therefore a necessary step toward modeling realistic scenarios in which different groups exhibit different levels of vaccine uptake.

To provide a mathematically grounded reference framework for such studies, we investigated a class of epidemic models describing multi-dose vaccination with imperfect protection and waning immunity. The goal of this line of research is to establish baseline results on optimal vaccination schedules that can later be extended to simulation-based studies including heterogeneous vaccine acceptance.

Model formulation We consider a generalized SIS-type model with vaccination and waning immunity, referred to as the S3I2 model. Individuals can occupy five possible epidemiological states:

- S_0 : fully susceptible individuals with no prior immunization,
- S_1 : partially immunized individuals after one immunization event,
- S_2 : maximally immunized individuals after two immunization events,
- I_1 : infected individuals with one prior immunization event,
- I_2 : infected individuals with two prior immunization events.

The immunization counter increases either through vaccination or through natural infection and is reset when immunity wanes.

Primary immunization events occur either through infection or through administration of a first vaccine dose. Booster doses correspond to the second immunization event and provide stronger but still imperfect protection.

Transmission rates depend on the immunization status. Denoting by β the baseline transmission rate and by k the average number of contacts, the effective transmission parameter is

$$\lambda = \beta k.$$

Partial and maximal immunity reduce susceptibility by factors σ_1 and σ_2 satisfying

$$0 < \sigma_2 < \sigma_1 < 1.$$

Recovery occurs at rate γ independently of immunization status.
Immunity wanes at rates η_1 and η_2, with

$$\eta_2 < \eta_1,$$

so that maximal immunity lasts longer than partial immunity.
Vaccination occurs at rate c. A priority parameter

$$\alpha \in [0, 1]$$

controls the allocation of doses between individuals waiting for the first dose and individuals waiting for the booster dose. The effective interdose interval is

$$\tau = \frac{1}{c(1-\alpha)}.$$

The population dynamics is described by the system

$$\dot{s}_0 = -\lambda\rho s_0 - c\alpha s_0 + \eta_1 s_1 + \eta_2 s_2, \tag{20}$$
$$\dot{s}_1 = -\sigma_1\lambda\rho s_1 + c\alpha s_0 - c(1-\alpha)s_1 - \eta_1 s_1 + \gamma\rho_1, \tag{21}$$
$$\dot{s}_2 = -\sigma_2\lambda\rho s_2 + c(1-\alpha)s_1 - \eta_2 s_2, \tag{22}$$
$$\dot{\rho}_1 = \lambda\rho s_0 - \gamma\rho_1, \tag{23}$$
$$\dot{\rho}_2 = \lambda\rho(\sigma_1 s_1 + \sigma_2 s_2) - \gamma\rho_2, \tag{24}$$

where s_i and ρ_i denote the population fractions in each state and

$$\rho = \rho_1 + \rho_2, \qquad s_0 + s_1 + s_2 + \rho_1 + \rho_2 = 1.$$

This formulation provides a minimal description of vaccination with multi-dose schedules and waning immunity consistent with diseases such as COVID-19.

Epidemic threshold At stationarity the epidemic threshold is determined by the condition

$$\rho\,[\lambda(s_0 + \sigma_1 s_1 + \sigma_2 s_2) - \gamma] = 0,$$

which separates a disease-free phase from an endemic phase.

Linear stability analysis yields an explicit expression for the critical transmission parameter

$$\lambda_c(\alpha) = \frac{\gamma}{s_0^{(0)} + \sigma_1 s_1^{(0)} + \sigma_2 s_2^{(0)}},$$

where $(s_0^{(0)}, s_1^{(0)}, s_2^{(0)})$ denotes the stationary distribution of susceptible states in the absence of infection.

An explicit expression can be obtained as

$$\lambda_c(\alpha) = \gamma \frac{\eta_2(c+\eta_1) + c^2\alpha(1-\alpha)}{\eta_2(c+\eta_1) - c\eta_2(1-\sigma_1)\alpha + \sigma_2 c^2\alpha(1-\alpha)}.$$

The epidemic threshold therefore depends strongly on the vaccination schedule through the parameter α.

Adjusting the prioritization between first and second doses can shift the epidemic threshold and determine whether the disease persists or dies out. :contentReference[oaicite:0]index=0

Optimal vaccination strategies The optimal vaccination strategy corresponds to the value

$$\alpha^* = \arg\max_{\alpha\in[0,1]} \lambda_c(\alpha),$$

which maximizes the epidemic threshold and therefore minimizes the region of parameter space where endemic transmission occurs.

The analysis reveals the existence of two distinct regimes separated by a critical vaccination rate c^*.

For low vaccination rates

$$c < c^*$$

the optimal strategy corresponds to

$$\alpha^* = 1,$$

meaning that first doses should be prioritized in order to maximize the fraction of partially protected individuals.

For higher vaccination rates

$$c > c^*$$

a nontrivial prioritization scheme becomes optimal, corresponding to a finite interdose interval.

This transition identifies the minimal vaccination capacity required for booster scheduling to become advantageous.

Numerical validation The theoretical predictions were validated through stochastic simulations on several classes of networks, including Erdős–Rényi graphs, geometric networks and empirical contact networks.

Simulation results confirm that the epidemic threshold reaches a maximum at the theoretically predicted value α^* and decreases for both smaller and larger values of the priority parameter. :contentReference[oaicite:1]index=1

These results demonstrate that vaccination timing can determine whether an epidemic remains endemic or is suppressed.

Role within the project The S3I2 vaccination model provides a theoretical baseline for the study of vaccination behavior in socially structured populations.

While the model assumes homogeneous mixing, its analytical tractability allows precise characterization of optimal vaccination schedules in the presence of waning immunity and imperfect protection.

Within the CODE project this framework serves as a reference model for future simulation-based studies in which heterogeneous populations, generated through the synthetic population and network models described above, exhibit different attitudes toward vaccination.

In particular, the model provides a natural starting point for investigating scenarios in which ideological or behavioral communities differ in vaccine acceptance, allowing a direct extension of the multitype epidemic framework introduced in the previous subsections.

4.4 Discursive Communities and Social Media Analysis

The modeling framework developed in this project relies on the assumption that populations can be partitioned into a relatively small number of communities characterized by similar sources of information, ideological orientations and behavioral attitudes. While this assumption is supported by sociological evidence, it is important to establish empirically whether such community structures can be reliably identified from real interaction data.

Within the CODE project we therefore developed a data-driven methodology to identify *discursive communities* in online social networks and to characterize their structure. These communities represent groups of users that participate in the same public debates while relying on similar information sources and interacting preferentially with the same set of content creators.

The methodology provides an empirical basis for defining realistic community structures and mixing patterns in the network models described in the previous sections, and allows the study of the temporal evolution of opinion structures during epidemic events.

The datasets used for these analyses, together with the scripts employed for their processing, are available in the public project repository for reproducibility purposes.

4.4.1 Identification of discursive communities

The proposed methodology is described in detail in *Leveraging content producer networks and user perception to detect online discursive communities*, [1]. The central assumption is that online debates are typically organized around a relatively small number of influential content creators whose messages are redistributed by a much larger audience [11].

This hierarchical organization implies that the ideological structure of the system can be inferred by focusing on content producers and on their audiences rather than by applying generic community detection algorithms to the full interaction network.

Starting from a dataset of posts and retweets, we construct a bipartite network between users and messages. From this representation two complementary procedures are used to identify discursive communities.

MonoDC. In the MonoDC approach we consider the directed retweet network between content creators. Let w_{ij} denote the number of times user j retweets user i. The resulting weighted adjacency matrix defines a directed network of content producers.

Since raw retweet networks contain a large amount of statistical noise, we filter this network using a maximum-entropy null model that preserves the activity of each user while randomizing higher-order correlations. In particular, if k_i^{out} and k_j^{in} denote the number of messages posted by user i and the number of retweets performed by user j, the null model preserves the expected values

$$\langle k_i^{\text{out}} \rangle \quad \text{and} \quad \langle k_j^{\text{in}} \rangle$$

while maximizing entropy. Observed retweet counts are compared with the null model expectations and only statistically significant links are retained.

Community detection is then performed on the validated network of content creators and community labels are propagated to the remaining users.

BiDC. The BiDC approach identifies communities from audience similarity rather than direct interactions.

We construct a bipartite undirected network between content creators and standard users, where a link indicates that a user has retweeted a given content creator at least once. Projecting this structure onto the set of content creators produces a similarity network in which two users are connected when they share a significant fraction of their audience.

If U_i and U_j denote the sets of users who retweeted content creators i and j, similarity can be measured through the overlap

$$S_{ij} = |U_i \cap U_j|.$$

As in the MonoDC approach, a maximum-entropy null model preserving the degree of each node in the bipartite network is used to identify statistically significant similarities and remove spurious co-occurrences.

Communities are detected on the resulting validated network and labels are propagated to the rest of the user population.

Community structure Extensive validation tests show that the resulting partitions closely reproduce known political alignments.

MonoDC typically identifies relatively fine-grained communities corresponding to individual political parties or ideological subgroups. BiDC instead tends to reveal larger-scale structures corresponding to coalitions or broader ideological alignments.

Both methods produce partitions that are substantially more stable and interpretable than those obtained by applying standard modularity-based community detection algorithms directly to retweet networks.

The detected communities also exhibit a characteristic bow-tie structure. A relatively small set of influential users produces most of the content, while the majority of users act primarily as receivers. This structure provides a natural representation of information flow in online debates and suggests a simple framework for constructing synthetic directed interaction networks.

Within the CODE project these results provide empirical support for the community-based modeling approach adopted in the previous sections and offer a method to calibrate realistic mixing matrices between groups.

4.4.2 Social media datasets and temporal analyses

The discursive community detection methods described above were applied to a dataset of interactions from X/Twitter covering several time windows during the first two years of the COVID–19 pandemic. The dataset includes posts, retweets, and user activity data, and is organized into successive temporal intervals to allow the study of the evolution of discursive communities over time.

The collection of these datasets required substantial methodological adaptation due to the changes in the access policies of the platform following the change of ownership in 2022. These changes slowed down the data acquisition process and required a reorganization of the download procedures.

Thus, the dataset was built by anchoring collection on verified accounts that, during the first week of March 2020, posted tweets that were retweeted at least 10 times (following [12] and related work). For these verified users, we downloaded original tweets (excluding retweets) in 1-month windows sampled every three months from November 2019 to November 2021, and for each tweet we collected the full list of retweeters. The resulting data consists of a set of tweets written by verified users and the corresponding retweeter lists per snapshot. The dataset includes 388,578 original messages (i.e., no retweets), posted by 895 different verified users, in 9 one-month-length windows. The dataset also includes information about the 17,746,555 retweets to the collected posts, by 2,129,884 different users.

The analyses are currently ongoing and will continue beyond the official duration of the project. More in detail, the research intends to assess whether (and how) the mesoscale structure of the Italian Twitter/X debate changed when the COVID-19 pandemic hit Italy. We focus on three potential channels of change: (i) reorganization of discursive communities among debate leaders (verified users), (ii) changes in how their audiences (standard users) are organized and polarized, and (iii) topic-level changes captured through hashtag-based semantic networks. The results reported here should be interpreted as an intermediate checkpoint: overall patterns appear fairly stable, yet several additional tests are required to rule out subtler forms of reorganization.

We adopted the Discursive Communities (DiCo) framework described above, using both (i) *bipartite* DiCo (BiDC), and (ii) *monopartite* DiCo (MonoDC), based on retweet ties among verified users. In addition, to probe topic-level structure, we built a semantic bipartite network between verified users and hashtags, and applied the same validation-and-projection strategy as in [13].

Across snapshots, verified-user BiDC communities appear fairly stable, with two notable transitions: between November 2019 and February 2020 (increasing separation between M5S and Lega) and between February and May 2021 (consistent with the formation of the Draghi government). At the audience level, flows of active standard users across communities are limited: most users remain within the same macro-community, and snapshot-to-snapshot persistence suggests that the "zealot" core may be relatively small in some political communities, while the majority of the audience is composed of only moderately active users, who interact from time to time with the content created by verified users. Polarization among standard users is generally stable over time; a statistically significant decreasing trend is observed for M5S (Mann–Kendall test), while other trends are weaker and do not robustly survive multiple-testing correction. Hashtag semantic networks provide limited evidence of strong, validated structure on the hashtag layer; the projection on the verified-user layer yields small communities mixing different political alignments, making large-scale interpretations challenging.

To reduce snapshot noise, we also detected BiDC communities on the dataset aggregated across all snapshots (see a pictorical representation of the validated network in Fig. 8); this yields cleaner and more politically interpretable clusters, but does not qualitatively alter the stability of polarization patterns. A

notable feature is that the right-wing community shows a fat-tailed polarization-index distribution, whereas progressive and M5S communities are closer to Gaussian-like shapes. To isolate COVID-related dynamics from unrelated high-frequency events (e.g., entertainment, sports), we built a thematic COVID subset via a manually curated keyword list (outperforming a sentence-transformer baseline on a small labeled sample). On this thematic subset, the most visible change concerns verified-user cross-community retweeting: non-political communities show reduced polarization and increased retweeting toward institutional/political communities (notably PD and M5S) during pandemic phases, plausibly reflecting amplification of official communications. This effect is largely masked in the full dataset, where non-COVID topics dominate.

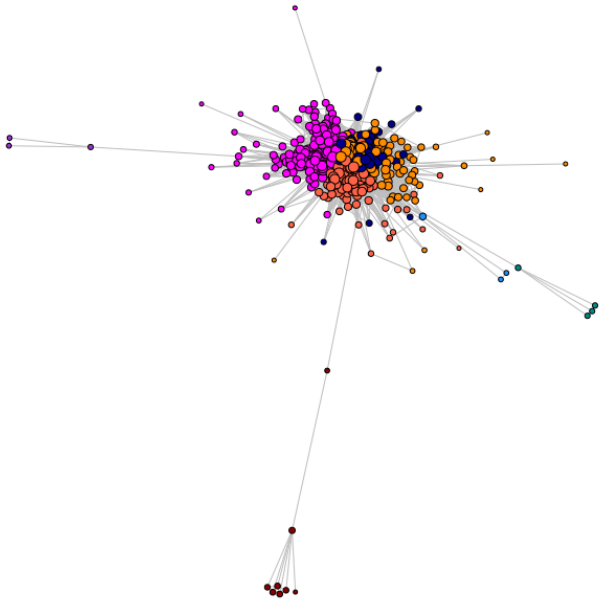

Figure 8: Validated network of verified users, computed on the full aggregated dataset, using BiDC. The red community is associated with progressives, orange with M5S, and navy with the right-wing coalition.

While the analyses performed so far suggest limited evidence for abrupt mesoscale reorganization (especially in leader communities and audience polarization), additional work is required to test for more nuanced changes. Key next steps include: (i) quantifying the consistency between aggregated communities and snapshot communities (e.g., how often verified users assigned to an aggregated community are split across snapshot partitions, and whether the direction of these deviations changes over time), (ii) extending the analysis of audience dynamics beyond consistently super-active users, (iii) integrating topic-sensitive measures beyond validated hashtag projections (e.g., alternative semantic representations or supervised thematic labeling at scale), and finally (iv) studying the URLs from news outlets shared in the debate and how the endorsement to the various news outlets chaged during the pandemic. These extensions are necessary to robustly assess whether COVID induced subtle but systematic changes in debate structure that are not captured by the current observables.

Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA

5 Conclusions

The CODE project investigated the interplay between epidemic spreading, information dynamics and human behavior, with particular emphasis on the role of community structure in shaping both social and epidemic processes.

A central idea underlying the project is that populations can be naturally partitioned into groups characterized by similar ideological orientations, sources of information and behavioral attitudes. This observation provides a common framework for modeling epidemic spreading, opinion dynamics and online interactions. The existence of such community structures motivates the use of network models with explicit mesoscale organization and allows the integration of data-driven information into epidemic simulations.

The project followed a coherent research strategy combining three complementary directions.

First, we developed a class of network models capable of generating synthetic populations and interaction networks with controlled community structure. The Random Hyperbolic Block Model provides a general theoretical framework in which clustering, degree heterogeneity and community mixing can be tuned independently, while the Urban Social Network model allows the construction of realistic contact networks based on demographic and geographic data. Together, these models provide a flexible foundation for the simulation-based study of socially structured populations.

Second, we defined a set of epidemic and opinion-dynamics models that capture the main mechanisms identified in the literature while remaining sufficiently simple to allow systematic exploration of scenarios. These models have been implemented in an open-source computational toolbox that allows both researchers and non-specialist users to explore data-informed epidemic scenarios. The availability of a unified software platform represents one of the main deliverables of the project and provides a basis for future work on coupled epidemic and social dynamics.

Third, we carried out a detailed scientific analysis of simplified models that retain the essential mechanisms included in the full framework. Multitype epidemic models clarified how susceptibility heterogeneity, infectivity heterogeneity and group mixing jointly determine epidemic thresholds and transient dynamics. Data-driven simulations on realistic urban networks demonstrated that these mechanisms remain relevant in realistic settings and may produce strong spatial heterogeneities even when global indicators appear stable. The vaccination model with waning immunity provided a mathematically well-founded reference framework for studying optimal vaccination strategies and establishes a baseline for future work on heterogeneous vaccine acceptance.

In parallel, we developed data-driven methods for identifying ideological communities from online interaction data. These methods provide empirical support for the community-based modeling approach adopted throughout the project and offer a practical way to calibrate realistic mixing patterns between groups. The collection and release of social media datasets covering the main phases of the COVID–19 pandemic will enable further investigations on the temporal evolution of ideological communities and polarization.

Some research directions initiated within the project will continue beyond its formal duration. In particular, the temporal analysis of discursive communities during the pandemic is still in progress and will benefit from the datasets and analysis tools developed in this work.

Overall, the results of the CODE project show that realistic epidemic modeling requires a coupled treatment of social structure and epidemic dynamics. Community structure emerges as a key element linking opinion formation, behavioral heterogeneity and epidemic spreading. The models and tools developed in this project provide a unified framework for investigating these mechanisms and offer a foundation for future research on data-informed epidemic scenarios.

Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA

References

[1] Stefano Guarino, Ayoub Mounim, Guido Caldarelli, and Fabio Saracco. Leveraging content producer networks and user perception to detect online discursive communities, 2026.

[2] Stefano Guarino, Enrico Mastrostefano, and Davide Torre. Random hyperbolic graphs with arbitrary mesoscale structures. *Phys. Rev. E*, 112:054310, Nov 2025.

[3] Fabio Mazza, Marco Brambilla, Carlo Piccardi, and Francesco Pierri. A data-driven analysis of the impact of non-compliant individuals on epidemic diffusion in urban settings. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 481(2324):20250511, 10 2025.

[4] Alessandro Celestini, Francesca Colaiori, Stefano Guarino, Enrico Mastrostefano, Francesca Pelusi, and Lena Rebecca Zastrow. When to boost: How dose timing determines the epidemic threshold. *Phys. Rev. Res.*, 7:033125, Aug 2025.

[5] Zachary P. Neal, Annabell Cadieux, Diego Garlaschelli, Nicholas J. Gotelli, Fabio Saracco, Tiziano Squartini, Shade T. Shutters, Werner Ulrich, Guanyang Wang, and Giovanni Strona. Pattern detection in bipartite networks: A review of terminology, applications, and methods. *PLOS Complex Systems*, 1(2):1–34, 10 2024.

[6] Manuel Pratelli, Fabio Saracco, and Marinella Petrocchi. Entropy-based detection of twitter echo chambers. *PNAS Nexus*, 3(5), April 2024.

[7] Manuel Pratelli, Fabio Saracco, and Marinella Petrocchi. Unveiling news publishers trustworthiness through social interactions. In *ACM Web Science Conference*, Websci '24, page 139–148. ACM, May 2024.

[8] Manuel Pratelli, Marinella Petrocchi, Fabio Saracco, and Rocco De Nicola. Online disinformation in the 2020 u.s. election: swing vs. safe states. *EPJ Data Science*, 13(1), March 2024.

[9] Fabio Saracco, Giovanni Petri, Renaud Lambiotte, and Tiziano Squartini. Entropy-based models to randomise real-world hypergraphs. *Communications Physics*, 8(1), July 2025.

[10] Anna Gallo, Fabio Saracco, and Tiziano Squartini. Statistically validated projection of bipartite signed networks. *npj Complexity*, 2(1), July 2025.

[11] Carolina Becatti, Guido Caldarelli, Renaud Lambiotte, and Fabio Saracco. Extracting significant signal of news consumption from social networks: the case of twitter in italian political elections. *Palgrave Communications*, 5:1–16, 12 2019.

[12] Emanuele Brugnoli, Pietro Gravino, Donald Ruggiero Lo Sardo, Vittorio Loreto, and Giulio Prevedello. Fine-grained clustering of social media: How moral triggers drive preferences and consensus. In *Proceedings of the 16th International Conference on Agents and Artificial Intelligence - Volume 3: AWAI*, pages 1405–1412. INSTICC, SciTePress, 2024.

[13] Tommaso Radicioni, Fabio Saracco, Elena Pavan, and Tiziano Squartini. Analysing twitter semantic networks: the case of 2018 italian elections. *Scientific Reports 2021 11:1*, 11:1–22, 6 2021.