The Annals of Applied Statistics
2025, Vol. 19, No. 4, 3416–3437
https://doi.org/10.1214/25-AOAS2069

MEDIA BIAS AND POLARIZATION THROUGH THE LENS OF A MARKOV SWITCHING LATENT SPACE NETWORK MODEL

BY ROBERTO CASARIN[1,2,a], ANTONIO PERUZZI[1,b] AND MARK F. J. STEEL[3,c]

[1]*Department of Economics, Ca' Foscari University of Venice,* [a]*r.casarin@unive.it,* [b]*antonio.peruzzi@unive.it*
[2]*Venice Centre in Economic and Risk Analytics for Public Policies*
[3]*Department of Statistics, University of Warwick,* [c]*m.steel@warwick.ac.uk*

News outlets are now more than ever incentivized to provide their audience with slanted news, while the intrinsic homophilic nature of online social media may exacerbate polarized opinions. Here we propose a new dynamic latent space model for time-varying online audience-duplication networks, which exploits social media content to conduct inference on media bias and polarization of news outlets. We contribute to the literature in several directions: (1) Our model provides a novel measure of media bias that combines information from both network data and text-based indicators; (2) we endow our model with Markov-switching dynamics to capture polarization regimes while maintaining a parsimonious specification; (3) we contribute to the literature on the statistical properties of latent space network models. The proposed model is applied to a set of data on the online activity of national and local news outlets from four European countries in the years 2015 and 2016. We find evidence of a strong positive correlation between our media slant measure and a well-grounded external source of media bias. In addition, we provide insight into the polarization regimes across the four countries considered.

1. Introduction. We propose a new statistical model able to offer meaningful insights into the perceived media bias and regime changes in polarization within online social media. The risk of being unintentionally exposed to biased news and polarized opinions has gained awareness both in the public debate (WEF (2022)) and in the academic sphere (see Puglisi and Snyder (2015), Gentzkow, Shapiro and Stone (2015), Cinelli et al. (2021)) due to the rapid changes in the news consumption landscape (Newman et al. (2017)). Luckily, the current availability of social-media data provides a privileged perspective on phenomena related to people's preferences and homophilous behavior (Zhang et al. (2018), Chen, Okhrin and Wang (2024), Yu et al. (2022)).

Figure 1 provides an illustrative example of both media bias and polarization starting from a preliminary analysis of the dataset described in Section 4. The figure displays a network of Italian news outlets in which the thickness of the edges is proportional to the number of Facebook commenters in common between any two outlets in the years 2015 (top) and 2016 (bottom). While media bias, in terms of political leaning, can be inferred indirectly from the network's structure or directly by analyzing news outlets' content production, an increase in polarization can be detected when the average number of users interacting with various sources decreases (see Figure 2).

Media bias refers to the dissemination of biased news pieces, often with the aim of supporting the interests of certain individuals or groups, such as political parties. The phenomenon is considered detrimental to consumer welfare (Gentzkow, Shapiro and Stone (2015)), as it entails a reduction in the informativeness of news pieces, while some argue that biased news outlets driven by both ideological interests and profits could even affect political outcomes

Received December 2023; revised June 2025.
Key words and phrases. Bayesian Inference, latent variables, political leaning, news outlets.

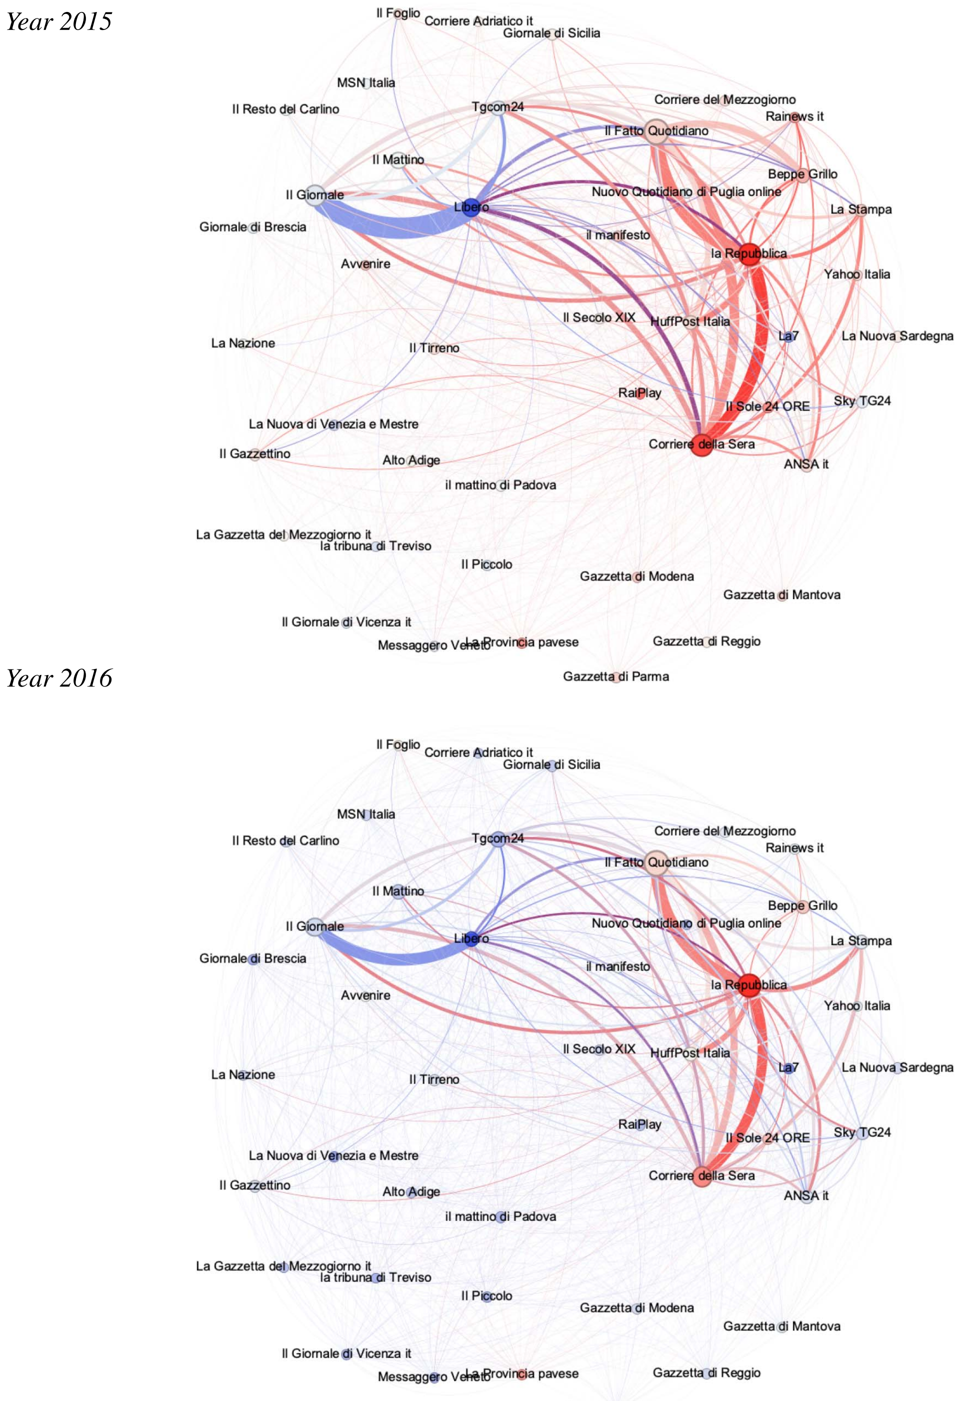

FIG. 1. *Italian audience-duplication networks obtained from the bipartite network of Italian news outlets and their Facebook commenters in* 2015 (*top*) *and* 2016 (*bottom*). *Node size is proportional to the news outlets' engagement in terms of comments. Nodes are labeled with the name of the outlet and colored from red* (*left*) *to blue* (*right*) *according to the text-analysis political-leaning score computed following Gentzkow and Shapiro* (2010) *and Garz, Sörensen and Stone* (2020). *Edge thickness is proportional to the number of commenters in common between any two news outlets.*

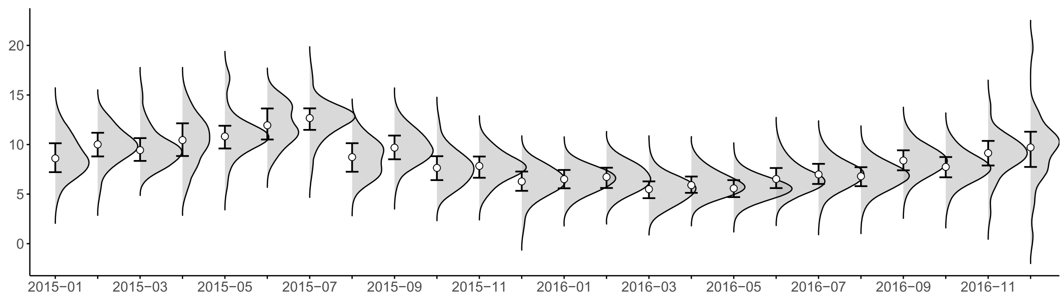

FIG. 2. *Monthly network density of the average number of commenters in common between any two Italian outlets. The gray areas indicate the kernel density estimates, the bars display the interquartile range, and the dots denote the median.*

(Anderson and McLaren (2012)). Recent advancements in measuring media bias include the implementation of text-analysis techniques to account for the similarity between news articles and political content (see Gentzkow and Shapiro (2010), Garz, Sörensen and Stone (2020)) as well as the use of both stochastic block models (SBMs; Lee and Wilkinson (2019)) and latent-space (LS) models (Hoff, Raftery and Handcock (2002), Friel et al. (2016), Sewell and Chen (2016)). SBMs are employed to capture discrete group structures, such as clusters of ideologically aligned actors or political communities (see Peixoto (2019)). LS models, on the other hand, allow for infinitely many levels of political leaning and induce a node ordering by representing social-media relational data in a continuous latent space (Barberá (2015), Ng et al. (2021)).

Polarization refers to the radicalization of people's opinions in the sense that they are further apart from one another. Some fear that this collective change in attitudes may be reflected in more partisan positions of people's representatives, even though there is no obvious evidence of this (Prior (2013)). Others claim that online social media exacerbate polarization by offering incentives for homophilous behavior, that is, the tendency to interact with similar individuals (Dandekar, Goel and Lee (2013)). However, while a predisposition toward homophily has been observed on several social platforms (see Hanusch and Nölleke (2019), Cinelli et al. (2021)), evidence of an exacerbation of polarization in social media environments is mixed (Kubin and von Sikorski (2021)). Several different methodologies have been adopted for measuring polarization (see Esteban and Ray (1994), Yarchi, Baden and Kligler-Vilenchik (2021)), including in the field of network science (see Garimella et al. (2018), Cinelli et al. (2021)). Two common objects of investigation are bipartite networks, which relate social media users to online pages, and audience duplication networks, where nodes represent pages and weighted edges denote the number of users in common between any pair of pages (a one-mode projection of the bipartite network). Figure 3 illustrates the two concepts.

In this paper we present a novel dataset of time-varying media networks. We construct reader-user bipartite networks and audience duplication networks using tick-by-tick information on the Facebook activity of national and local news outlets, as collected by Schmidt et al. (2018). The dataset comprises all posts published by the news outlets from four European countries (France, Germany, Italy, and Spain) along with the corresponding users' interactions for the years 2015 and 2016. Our bipartite network explicitly captures the presence or absence of user–outlet interactions in terms of comments, while the derived audience duplication network is constructed by assigning edge weights between pairs of news outlets proportional to the number of users who comment on posts from both outlets. Following common practice in the literature (Peng and Yang (2022)), we aggregate user comments at the outlet level to obtain a more stable representation of each outlet's overall editorial line,

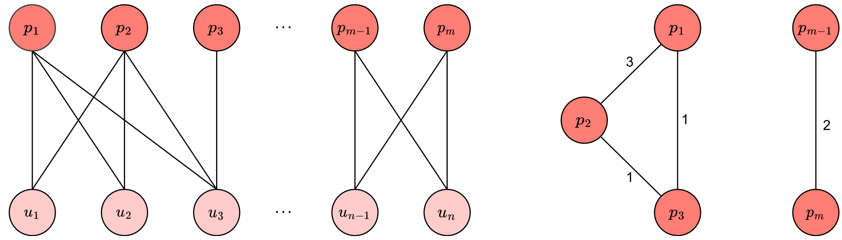

$$B_{n\times m} = \begin{pmatrix} 1 & 1 & 0 & \dots & 0 & 0 \\ 1 & 1 & 0 & \dots & 0 & 0 \\ 1 & 1 & 1 & \dots & 0 & 0 \\ \vdots & & & & & \vdots \\ 0 & 0 & 0 & \dots & 1 & 1 \\ 0 & 0 & 0 & \dots & 1 & 1 \end{pmatrix} \qquad A_{m\times m} = B'B = \begin{pmatrix} 3 & 3 & 1 & \dots & 0 & 0 \\ 3 & 3 & 1 & \dots & 0 & 0 \\ 1 & 1 & 1 & \dots & 0 & 0 \\ \vdots & & & \ddots & & \\ 0 & 0 & 0 & \dots & 2 & 2 \\ 0 & 0 & 0 & \dots & 2 & 2 \end{pmatrix}$$

FIG. 3. *Media networks: An example of a bipartite network of n readers, $u_1, \dots, u_n$ and m news outlets, $p_1, \dots, p_m$ (top left) and the corresponding audience duplication network obtained using one-mode projection (top right). At the bottom, the adjacency-matrix representation of both the bipartite network, $B_{n\times m}$, and the audience-duplication network $A_{m\times m}$. The matrix A is obtained as $A = B'B$.*

avoiding the noise inherent in post-level heterogeneity. The resulting dataset is openly available; see Appendix M in the Supplementary Material (Casarin, Peruzzi and Steel (2025)) and the online repository (Casarin, Peruzzi and Steel (2025)).

Previous studies about media polarization use heuristics to detect communities and informal sequential analysis for time variation, while we propose a formal statistical framework for dynamic polarization analysis, treating polarization as a collective feature of the actors involved. In particular, we introduce a novel dynamic LS network model that leverages both time-varying online audience-duplication network data and textual content to characterize a set of news outlets in terms of both a dynamic latent political leaning dimension and popularity via individual effects.

LS models project the nodes of a network on a lower d-dimensional latent space (Hoff, Raftery and Handcock (2002)). Extensions of the original model include, for example, a dynamic component for the latent coordinates (Friel et al. (2016), Sewell and Chen (2016), Kim et al. (2018)), mixtures of latent coordinates (see Handcock, Raftery and Tantrum (2007)), and common or multiple latent spaces for multilayer networks (see Gollini and Murphy (2016), D'Angelo, Murphy and Alfò (2019), Sosa and Betancourt (2022)). The statistical properties of LS models have been addressed by Rastelli, Friel and Raftery (2016) for binary networks, while Barberá (2015) presents an early application of LS modeling for the estimation of latent ideology on social media and De Nicola, Tuekam Mambou and Kauermann (2023) provides a static assessment of media polarization using LS models.

Our paper adds to the methodological literature in several respects:

- The latent coordinates of our LS model provide a novel measure of media bias which combines information from network data (e.g., Barberá (2015)) and from a state-of-the-art text-based indicator (e.g., Gentzkow, Shapiro and Stone (2015), Garz, Sörensen and Stone (2020)).
- We endow the latent coordinates with Markov-switching (MS) dynamics which allows for capturing polarization regimes over time. A similar approach has been adopted for change-point detection in the context of LS projection models (Park and Sohn (2020)). The choice of modeling regimes rather than continuously varying levels of polarization is coherent with the literature on opinion formation (see Iyengar, Sood and Lelkes (2012), Törnberg et al. (2021)). In addition, the MS model is parsimonious and more easily scalable than other dynamic specifications.

- We extend the results of Rastelli, Friel and Raftery (2016) on statistical properties of LS models to weighted temporal networks with MS dynamics (MS-LS models) by obtaining closed-form expressions for the first and second moment of the strength distribution. Besides the theoretical relevance of the results, we show how such closed-form expressions can be used to assess the adequacy of the proposed model against other competing models (see Section 4.3).

Our model is applied to our novel temporal network dataset to provide estimates of media bias that are coherent with the PEW Research survey (Mitchell et al. (2018)). We also shed light on the in-platform (i.e., within Facebook) polarization regimes across the four countries. We find evidence of cross-country heterogeneity in the shifts from low to high polarization regimes, in line with the sociological argument of Prior (2013).

Section 2 provides an overview of the model within a Bayesian framework and discusses its statistical properties. In Section 3 we discuss posterior inference along with the constraints used in our model and present a simulation exercise. Finally, Section 4 describes the dataset of European news outlets and applies our model in both static and dynamic setups.

2. The longitudinal Markov-switching latent space model.

2.1. *The model.* Let $\mathcal{G} = \{\mathcal{G}_t, t = 1, 2, \dots, T\}$ be an undirected and weighted temporal (in our application audience-duplication) network with $\mathcal{G}_t = (V_t, E_t, \mathbf{Y}_t)$. For each time instance t, the vertex set (collection of news outlets) is constant, that is, $V_t = V \subset \mathbb{N}$, and the edge set $E_t \subset V \times V$ (common commenters between any two outlets) is time-varying. For each edge $(i, j) \in E_t$, we assume the (i, j)th element of the weighted adjacency matrix $\mathbf{Y}_t$, that is, Y_{ijt}, is observed and denotes the number of connections (commenters in common) between news outlets i and j at time t. We adopt a Poisson model for the connections,

$$Y_{ijt}|\lambda_{ijt} \overset{\text{ind}}{\sim} \mathcal{P}oi(\lambda_{ijt}), \tag{1}$$

for $i, j = 1, \dots, N$ $i \neq j$, $t = 1, \dots, T$, where $N = \text{Card}(V)$ is the number of nodes in the network and $\mathcal{P}oi(\lambda)$ denotes a Poisson distribution with intensity parameter $\lambda > 0$. In our LS model, the intensity is driven by static (α_i) and d-dimensional dynamic node-specific latent features ($\mathbf{x}_{it} \in \mathcal{X} \subset \mathbb{R}^d$),

$$\log \lambda_{ijt} = \alpha_i + \alpha_j - \beta \|\mathbf{x}_{it} - \mathbf{x}_{jt}\|^2. \tag{2}$$

The parameters α_i, $i = 1, \dots, N$ have the natural interpretation of individual effects which are news-outlet specific and capture the popularity of the outlet (the engagement of the audience with the newspaper). The latent features $\mathbf{x}_{it}$ are realizations of the latent variables $\mathbf{X}_{it}$, $i = 1, \dots, N$, and enter the log intensity through the squared Euclidean distance, as suggested in Gollini and Murphy (2016) and in D'Angelo, Murphy and Alfò (2019). This accounts for a clearer representation of the proximity of news outlets on a latent manifold and leads to faster computational convergence compared to the standard Euclidean distance. Assuming $\beta > 0$ lends an interpretation of similarity features to the latent variables. The more similar the news outlets (the closer the nodes), the higher the number of commenters they tend to have in common. We also employ an observed political-leaning proxy, L_{it}, to provide additional information on the location of news outlets within the latent space. This political-leaning measure takes values between zero, denoting extreme left, and one, denoting extreme right (e.g., see Gentzkow and Shapiro (2010), Garz, Sörensen and Stone (2020)). Our modeling choice is not to include L_{it} in the log-intensity equation to preserve the tractability of the random graph model properties. Rather, we assume that the political-leaning proxy L_{it} is driven by the same latent variables $\mathbf{X}_{it}$ as appear in the network log-intensity. Since the leaning measure L_{it} is continuous in (0, 1), a Beta-logistic regression model is assumed, which is a convenient and interpretable specification, although other choices can be used (e.g., see Branscum, Johnson and Thurmond (2007), Casarin, Dalla Valle and Leisen (2012)),

$$L_{it}|\mu_{it}, \phi \overset{\text{ind}}{\sim} \mathcal{B}e\big(\mu_{it}\phi, (1-\mu_{it})\phi\big), \tag{3}$$

where $\mathcal{B}e$ denotes the Beta distribution with location and precision parameters $\mu_{it} \in (0, 1)$ and $\phi > 0$, respectively. To endow the latent features $\mathbf{X}_{it}$ with a media-bias interpretation while ensuring Y_{ijt} is independent of L_{it} conditional on $\mathbf{X}_{it} = \mathbf{x}_{it}$, we assume

$$\mu_{it} = \varphi\big(\gamma_0 + \boldsymbol{\gamma}_1'\mathbf{x}_{it}\big), \tag{4}$$

where $\varphi(x) = 1/(1 + \exp(-x))$ is the logistic function. With this modeling choice, we can interpret the latent space as the political spectrum of news outlets.

In the media environment, news outlets are subject to time-varying polarization regimes (e.g., Macy et al. (2021), Leonard et al. (2021) on polarization dynamics). One can think about a state of low polarization, where news outlets are perceived, on average, closer within the political spectrum, and a state of high polarization, in which news outlets are perceived politically further apart. For this reason we specify a Markov-switching (MS) regime process $\{S_t, t = 1, \dots, T\}$ that drives the dynamic latent features. We assume a hidden Markov chain with $K < \infty$ possible states of the world. The MS process allows our latent features $\mathbf{x}_{it}$ to vary jointly through time across the different polarization states. We can reparameterize our dynamic latent features $\mathbf{x}_{it}$ to account for MS dynamics,

$$\mathbf{x}_{it} = \sum_{k=1}^{K} \mathbb{I}(s_t = k)\boldsymbol{\zeta}_{ik}, \tag{5}$$

where $\boldsymbol{\zeta}_{ik}$ denotes the latent position of news outlet i in state k, while $\mathbb{I}(s_t = k)$ is an indicator function which is 1 if the realization of the hidden variable S_t is $s_t = k$ at time t and 0 otherwise. Moreover, we characterize the transition between states through

$$\mathbb{P}(S_t = k|S_{t-1} = l) = q_{lk}, \quad l, k = 1, \dots, K, \tag{6}$$

which can be grouped in the transition probability matrix $\mathbf{Q} = \{\boldsymbol{q}_1, \dots, \boldsymbol{q}_k, \dots, \boldsymbol{q}_K\}$, where $\boldsymbol{q}_k = (q_{k1}, \dots, q_{kK})$ denotes a column vector such that $\boldsymbol{q}_k'\mathbf{1} = 1$ for each state k. Since the number of latent variables increases as $\mathcal{O}(dKN + T)$, our model is more parsimonious than a continuously varying LS model (e.g., see Sewell and Chen (2016)), where the number of latent variables increases as $\mathcal{O}(dTN)$. Thus, our approach is computationally much less demanding for dynamic network applications with large T. Our approach also differs substantially from SBMs. While SBMs would cluster the news outlets into discrete groups, our MS-LS ranks news outlets on a continuum interpretable as a political spectrum. More importantly, the specification in (5) allows identifying polarization regimes (periods in time characterized by common behavior in terms of political leaning and polarization). Appendix K in Casarin, Peruzzi and Steel (2025) provides more detail on alternative dynamic specifications.

We take a Bayesian approach to inference and choose the following prior structure (as explained in Section 3.2 we fix β in our empirical application):

$$\begin{aligned} &\pi\big(\alpha_1, \dots, \alpha_N, \{\boldsymbol{\zeta}_{1k}, \dots, \boldsymbol{\zeta}_{Nk}, \sigma_k^2\}_{k=1,\dots,K}, \gamma_0, \boldsymbol{\gamma}_1, \phi, \mathbf{Q}\big) \\ &\quad = \pi(\gamma_0)\pi(\boldsymbol{\gamma}_1)\pi(\phi)\left(\prod_{i=1}^{N}\pi(\alpha_i)\right)\prod_{k=1}^{K}\pi(\sigma_k^2)\pi(\boldsymbol{q}_k)\prod_{i=1}^{N}\pi\big(\boldsymbol{\zeta}_{ik}|\sigma_k^2\big), \end{aligned} \tag{7}$$

where we find it useful to augment the parameter space with σ_k^2, $k = 1, \dots, K$ and the specific prior distributions assumed are detailed in Appendix C.1 (Casarin, Peruzzi and Steel (2025)). In our implementation of the model, we have little prior information at our disposal, and we opt for the use of relatively vague priors to let the data speak. Our results are robust to substantial changes in these priors. The directed acyclic graph in Figure 4 summarizes our MS-LS model.

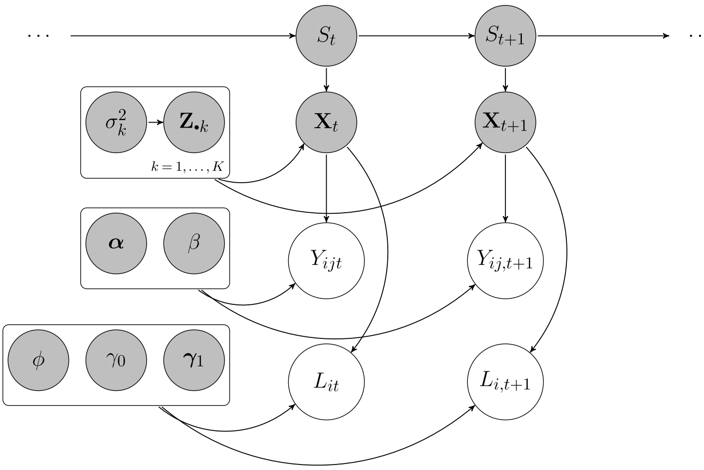

FIG. 4. *Directed acyclic graph of the Markov-switching latent-space model. The graph exhibits the conditional independence structure of the observation model for Y_{ijt} and L_{it} within white circles with parameters, $\boldsymbol{\alpha} = \{\alpha_1, \ldots, \alpha_N\}$, β, γ_0, $\boldsymbol{\gamma}_1$, ϕ, latent coordinates $\mathbf{X}_t = \{\mathbf{X}_{1t}, \ldots, \mathbf{X}_{Nt}\}$, their state-dependent counterparts $\mathbf{Z}_{\bullet k} = \{\mathbf{Z}_{1k}, \ldots, \mathbf{Z}_{Nk}\}$ with variance σ_k^2, and latent states, S_t, within gray circles.*

2.2. *Model properties.* We now present some of the properties of the MS-LS Model in (1), (2) and (5). With Assumption 2.1 we extend the scope of the latent variable model provided in Rastelli, Friel and Raftery (2016) to weighted temporal networks. In this section we do not consider the political-leaning equation (3) since this is quite specific to our application and we aim to present the properties of a general network model.

ASSUMPTION 2.1. Given an undirected temporal network, $\mathcal{G}_t = (V, E_t, \mathbf{Y}_t)$, for $t = 1, 2, \ldots$ having vertex set $V \subset \mathbb{N}$ and weighted edge sets $E_t \subset V \times V$ with characteristic weight Y_{ijt}, we assume a sequence of latent coordinates $\{\boldsymbol{X}_{1t}, \ldots, \boldsymbol{X}_{Nt}\}$ for $t = 1, 2, \ldots$ with $\boldsymbol{X}_{it} \in \mathcal{X} \subset \mathbb{R}^d$ for each node $i \in V$ and time index $t \in \mathbb{N}$.

With Assumption 2.2 we introduce the Markov-switching dynamics and the independence between latent variables across states.

ASSUMPTION 2.2. Given a K-state latent Markov-chain process $S_t \in \{1, 2, \ldots, K\}$ for $K < \infty$ and $t = 1, 2, \ldots$ with transition probabilities $q_{lk} = \mathbb{P}(S_t = k | S_{t-1} = l)$, we assume the latent variables $\boldsymbol{X}_{it} = \sum_{k=1}^{K} \mathbb{I}(S_t = k)\mathbf{Z}_{ik}$. We also define the set $\boldsymbol{\zeta}_{\bullet k} = \{\boldsymbol{\zeta}_{1k}, \ldots, \boldsymbol{\zeta}_{Nk}\}$ consisting of the i.i.d. realizations of the latent random variables $\{\mathbf{Z}_{1k}, \ldots, \mathbf{Z}_{Nk}\}$ with $k \in \{1, \ldots, K\}$, where each $\mathbf{Z}_{ik}$ is distributed according to $\pi_k(\cdot)$, a given probability measure.

Assumption 2.3 introduces the conditional independence between any two edges, given the latent variables and the current state of the world.

ASSUMPTION 2.3. We assume conditional independence between any two edges, given the latent variables applicable to the current state S_t. Given the intensity parameter λ_{ijt}, $\forall (j, i) \in E_t$, $Y_{ijt} | \boldsymbol{X}_{it} = \mathbf{x}_{it}, \boldsymbol{X}_{jt} = \mathbf{x}_{jt} \overset{\text{ind}}{\sim} \mathcal{P}oi(\lambda_{ijt})$ is a Poisson random variable.

Moreover, we assume that our latent variables are normally distributed and conditionally independent given their variance.

ASSUMPTION 2.4. The latent variables are normally distributed: $\mathbf{Z}_{ik} | \sigma_k^2 \sim \mathcal{N}(\mathbf{0}, \sigma_k^2 I_d)$ and take values in $\mathbb{R}^d$, for a fixed d.

Finally, we specify the form of the intensity parameter λ_{ijt}.

ASSUMPTION 2.5. Given the individual effects α_i and α_j and the latent variables, we assume the Poisson rate parameter,

$$\lambda_{ijt} = \exp\left\{\alpha_i + \alpha_j - \sum_{k=1}^{K} \mathbb{I}(s_t = k)\beta \|\boldsymbol{\zeta}_{ik} - \boldsymbol{\zeta}_{jk}\|^2\right\}.$$

Under Assumptions 2.1–2.5, our model is a time-varying MS-LS model with Poisson weights and normally distributed latent variables.

The nodal strength, defined as $Y_{it} = \sum_{j \neq i} Y_{ijt}$, is a quantity of particular interest when dealing with weighted networks, as it provides information on how strongly connected a node is with its neighbors. We can derive the following properties of the probability generating function (pgf) of the nodal strength, where in the sequel we will focus on the strength of a random node (and thus, omit the index i):

PROPOSITION 2.1. *The mth derivative of the conditional pgf G_l of the nodal strength evaluated in $x = 1$, given $S_{t-1} = l$ for the MS-LS model can be written as*

$$\frac{\partial^m G_l(x)}{\partial x^m}\bigg|_{x=1} = \sum_{k=1}^{K} q_{lk} \frac{\partial^m \widetilde{G}_k(x)}{\partial x^m}\bigg|_{x=1} = \sum_{k=1}^{K} \sum_{\underline{h}_i \in \mathcal{H}_i} \binom{m}{\underline{h}_i} e^{\sum_{j \neq i}(\alpha_i + \alpha_j)h_j} (\sigma_k^2)^{-\frac{d}{2}} b_{\underline{h},k} q_{lk},$$

where $\widetilde{G}_k(x)$ is the conditional pgf, given $S_t = k$ and $S_{t-1} = l$, and

$$b_{\underline{h},k} = \left(\frac{1}{\sigma_k^2} + \sum_{j \in \mathcal{J}_i} \frac{2\beta h_j}{2\beta h_j \sigma_k^2 + 1}\right)^{-\frac{d}{2}} \prod_{j \in \mathcal{J}_i} (2\beta h_j \sigma_k^2 + 1)^{\frac{d}{2}}$$

with multi-index $\underline{h}_i = \{h_1, \dots, h_{i-1}, h_{i+1}, \dots, h_N\}$ and index set $\mathcal{H}_i = \{h_j \in \{0, \dots, m\}, j \neq i \mid \sum_{j \neq i} h_j = m\}$, $\mathcal{J}_i = \{j \mid j \neq i, h_j > 0\}$ and $\beta > 0$.

The first derivative of the pgf returns the conditional expectation of the strength for a random node, $\mathbb{E}(Y_t | S_{t-1} = l)$.

COROLLARY 2.1. *Defining $\alpha = \alpha_i + \alpha_j$ for each i and each j, the expected nodal strength of the underlying network $\mathcal{G}_t$ can be expressed as*

$$\mathbb{E}(Y_t | S_{t-1} = l) = G_l'(x)|_{x=1} = \sum_{k=1}^{K} q_{lk} \widetilde{G}_k'(x)\bigg|_{x=1} = (N-1)e^{\alpha} \sum_{k=1}^{K} q_{lk} (4\sigma_k^2 \beta + 1)^{-\frac{d}{2}}.$$

Note that $\mathbb{E}(Y_t | S_{t-1} = l)$ turns out to be a weighted sum of the expected nodal strength obtained by conditioning on each possible state of the world. The result in Rastelli, Friel and Raftery (2016) can be obtained as a special case imposing all but one of the conditional probabilities q_{lk} for $k \in \{1, \dots, K\}$ to be zero. This is due to the fact that the expected conditional strength in their unweighted network is the same as in our setup, only with the restriction that α has to be negative. The expected strength for each regime increases linearly with the number of nodes, N, and exponentially with the intercept parameter α. The lower $\sigma_k^2 \beta$, the larger the similarity between the nodes in that state and, in turn, the higher the expected strength.

COROLLARY 2.2. *The analytical expression of the variance of the strength distribution uses the first and the second factorial moment of the pgf, resulting in*

$$\mathbb{V}ar(Y_t | S_{t-1} = l) = \sum_{k=1}^{K} q_{lk} \mathbb{V}ar(Y_t | S_t = k) + \sum_{k=1}^{K} q_{lk} \left(\widetilde{G}_k'(x)|_{x=1} - G_l'(x)|_{x=1}\right)^2,$$

where $\mathbb{V}ar(Y_t | S_t = k) = \widetilde{G}_k''(x)|_{x=1} + \widetilde{G}_k'(x)|_{x=1} - \widetilde{G}_k'^2(x)|_{x=1}$.

Similarly, an analytical expression for the dispersion index can be obtained,

$$\mathfrak{D}(Y_t|S_{t-1}=l)=\sum_{k=1}^{K}q_{lk}\mathfrak{D}(Y_t|S_t=k)+v,$$

where

$$\mathfrak{D}(Y_t|S_t=k)=1+v_k-\widetilde{G}'_k(x)|_{x=1},\quad v=\frac{\sum_{k=1}^{K}q_{lk}\widetilde{G}''_k(x)|_{x=1}}{\sum_{k=1}^{K}q_{lk}\widetilde{G}'_k(x)|_{x=1}}-\sum_{k=1}^{K}q_{lk}v_k,$$

with $v_k=\widetilde{G}''_k(x)/\widetilde{G}'_k(x)|_{x=1}$. *The expressions of the derivatives* $\widetilde{G}'_k(x)$ *and* $\widetilde{G}''_k(x)$ *are given in* (A.1) *and* (A.1) *in Appendix A* (*Casarin, Peruzzi and Steel* (2025)).

The result for the second factorial moment differs from the results in Rastelli, Friel and Raftery (2016), as our derivation reflects the existing heterogeneity in weighted edges. Moreover, our MS-LS model allows modeling overdispersion in the strength distribution; in particular, we show in Appendix A.3.5 (Casarin, Peruzzi and Steel (2025)) that $\mathfrak{D}(Y_t|S_{t-1}=l)>1$.

We derive the above results and present a sensitivity analysis of the moments of the strength distribution to different combinations of the model parameters in Appendix A (Casarin, Peruzzi and Steel (2025)). To help clarify these results, Figure A.1 presents a sensitivity analysis of these network metrics to the parameters of the generative network model: it displays contour plots of moments of the strength distribution of an MS-LS model with $d=1$ and $K=2$. Labels "L" and "H" denote low- and high-polarization states. The mean, variance and dispersion index of the strength distribution increase with α and N (see also Figure A.2, in the Appendix, Casarin, Peruzzi and Steel (2025)), and the mean decreases with $\sigma^2_{\mathrm{L}}\beta$. For $q_{\mathrm{LL}}=1$, when the network stays in low polarization, both the variance and the dispersion initially increase with $\sigma^2_{\mathrm{L}}\beta$ up to a maximum and then decrease. When a state of high polarization kicks in with $\sigma^2_{\mathrm{H}}\beta>\sigma^2_{\mathrm{L}}\beta$, *that is*, $q_{\mathrm{LL}}<1$, both variance and dispersion decrease with $\sigma^2_{\mathrm{L}}\beta$. The smaller q_{LL}, the sharper the variance and especially the dispersion index decrease with $\sigma^2_{\mathrm{L}}\beta$.

3. Inference.

3.1. *Posterior sampling algorithm.* In this section we will go back to the setup in (1)–(8), so including (3). Let $\mathbf{Y}=(\mathbf{Y}_1,\dots,\mathbf{Y}_T)$ be the collection of observed network weights with characteristic element of $\mathbf{Y}_t$ given by Y_{ijt}, $\mathbf{L}=(\mathbf{L}_1,\dots,\mathbf{L}_T)$ are the observed political-leaning proxies with characteristic element L_{it} and $\mathbf{S}=(S_1,\dots,S_T)$. Consider the parameters $\boldsymbol{\theta}=(\boldsymbol{\alpha},\mathbf{Z},\boldsymbol{\sigma}^2,\gamma_0,\boldsymbol{\gamma}_1,\phi,\boldsymbol{Q})$, while β is fixed in our empirical implementation (see Section 3.2). Here $\mathbf{Z}=(\mathbf{Z}_{\bullet 1},\dots,\mathbf{Z}_{\bullet K})$ denotes the latent coordinate parameters, where $\mathbf{Z}_{\bullet k}$ is a $d\times N$ matrix, while $\boldsymbol{\sigma}^2=(\sigma^2_1,\dots,\sigma^2_K)$ is a vector of state-specific variables. The joint posterior $p(\boldsymbol{\theta}|\mathbf{Y},\mathbf{L})\propto f(\mathbf{Y},\mathbf{L}|\boldsymbol{\theta})\pi(\boldsymbol{\theta})$ is not tractable. Thus, we follow a data augmentation approach and apply a Gibbs sampler for posterior inference (see Appendix C, Casarin, Peruzzi and Steel (2025)). Let us denote with $\boldsymbol{\xi}=(\boldsymbol{\xi}_1,\dots,\boldsymbol{\xi}_T)$ the collection of state indicator variables, where $\boldsymbol{\xi}_t=(\xi_{1t},\dots,\xi_{Kt})'$ and $\xi_{kt}=\mathbb{I}(S_t=k)$. The complete-data likelihood function is the product of the following:

$$f(\mathbf{Y},\mathbf{L},\boldsymbol{\xi}|\boldsymbol{\theta})=\prod_{t=1}^{T}\prod_{i=1}^{N}\left(f_B(L_{it}|S_t,\boldsymbol{\theta})\prod_{j=i+1}^{N}f_P(Y_{ijt}|S_t,\boldsymbol{\theta})\right)\prod_{l=1}^{K}\prod_{k=1}^{K}q_{lk}^{\xi_{lt-1}\xi_{kt}},\tag{8}$$

where $f_P(Y_{ijt}|S_t,\boldsymbol{\theta})$ is the probability mass function of the Poisson in (1) with dynamic intensity given in (2), $f_B(L_{it}|S_t,\boldsymbol{\theta})$ the probability density function of the Beta given in (3).

With this notation, $\mathbf{x}_{it}$ can be written as $\boldsymbol{\zeta}_{i\bullet}\boldsymbol{\xi}_t$, in (2) and (4) where $\boldsymbol{\zeta}_{i\bullet} = (\boldsymbol{\zeta}_{i1}, \dots, \boldsymbol{\zeta}_{iK})$ is a $d \times K$ matrix.

We approximate the joint posterior distribution by Markov-chain Monte Carlo (MCMC) sampling. Our Gibbs sampling algorithm iterates the following steps:

1. Draw α_i from $p(\alpha|\dots)$, $i = 1, \dots, N$ via Adaptive Metropolis–Hastings (MH).
2. Draw ϕ from $p(\phi|\dots)$ via MH with truncated normal proposal.
3. Draw γ_0 and $\boldsymbol{\gamma}_1$ from $p(\gamma_0, \boldsymbol{\gamma}_1|\dots)$ via MH.
4. Draw $\boldsymbol{\zeta}_{ik}$ from $p(\boldsymbol{\zeta}_{ik}|\dots)$, $i = 1, \dots, N$ and for $k = 1, \dots, K$ via Adaptive MH.
5. Draw σ_k^2 from $p(\sigma_k^2|\boldsymbol{\zeta}_{\bullet k})$ for $k = 1, \dots, K$.
6. Draw $\boldsymbol{q}_l$ from $p(\boldsymbol{q}_l|\boldsymbol{\xi})$ for $l = 1, \dots, K$.
7. Draw $\mathbf{s}$ via the forward-filtering and backward-sampling algorithm (see Frühwirth-Schnatter (2006)).

Further details on the algorithmic design can be found in Appendix C (Casarin, Peruzzi and Steel (2025)).

3.2. *Identifying restrictions.* The model presents several identification challenges. The first issue is related to the multiplication of the squared Euclidean distance $\|\boldsymbol{\zeta}_{i\bullet}\boldsymbol{\xi}_t - \boldsymbol{\zeta}_{j\bullet}\boldsymbol{\xi}_t\|^2$ by the parameter β. As there is a clear scale indeterminacy between β and the variance of the latent variables in terms of σ_k^2, we choose to set $\beta = 1$. In addition, latent coordinates enter the parameter λ_{ijt} only through the squared distance. This makes—in principle—positions that differ just by means of reflection, translation, and rotation equally likely (see Hoff, Raftery and Handcock (2002) and Friel et al. (2016)). Nonetheless, the introduction of (3) helps prevent the emergence of many equivalent latent-space representations. To overcome translation issues, we center the latent coordinates to the origin of the axes at each Gibbs-sampling iteration. For $d = 1$, a reflection of the latent coordinate around the origin is possible. To counteract reflection, we assume that the position of a single outlet is known in terms of left and right political leaning (e.g., $\zeta_{i^*k} < 0$ for a left-leaning outlet i^* and for each state k), and we apply a reflection transformation to the latent leaning coordinates every time the latent leaning of i^* is in the wrong orthant, similarly to Barberá (2015). For $d > 1$, not only reflection remains possible but also rotation. This causes indeterminacy of the parameter $\boldsymbol{\gamma}_1$. To prevent rotation of the latent coordinate with media bias interpretation, we impose the following restrictions to (3): $\boldsymbol{\gamma}_1 = (\gamma_{11}, 0, \dots, 0)$ and $\gamma_{11} = 1$. Such a strategy is similar in spirit to the loading restrictions in factor analysis (see Frühwirth-Schnatter, Hosszejni and Lopes (2025)). For the other coordinates, we apply Procrustes transformation to solve the identification issue, as commonly done for standard LS models (see Hoff, Raftery and Handcock (2002)).

Finally, as pointed out in Frühwirth-Schnatter (2006), the joint posterior in Markov-switching models is invariant with respect to a relabeling of the hidden states. We tackle this issue by imposing an ordering restriction on the latent regimes across states. In particular, we label latent regimes in increasing order of median distance, $\tilde{D}_k = med_{j>i}(\|\zeta_{ik} - \zeta_{jk}\|)$. Alternative ordering restrictions may be considered, but we leave these for future research.

3.3. *Analysis of simulated data.* We assess the performance of our inference method by running the algorithm on simulated data from an MS-LS model with $d = 1$ and $K = 2$. Our simulation consists of 20 fictitious news outlets observed for 100 periods. Each period may belong to one of two polarization states $k \in \{\text{L}, \text{H}\}$, where State L is characterized by a lower average distance in the political-leaning dimension across news outlets, and State H has a higher average distance. So news outlets jointly undergo periods of high polarization and low polarization. The simulation setting is inspired by previous studies (Törnberg et al. (2021)) and aligns with our empirical findings. In periods of high polarization, news outlets exhibit a stronger and more radicalized identity, corresponding to a higher average distance in the latent space. This manifests itself with news content that is more aligned with the language of the reference political parties and with an audience that is more homogeneous in terms of news diet, partially reflected by well-separated clusters along the political-leaning direction of the latent space. Vice versa, we expect a less pronounced and more heterogeneous identity of news outlets in periods of low polarization with lower outlet separation in the latent space.

The latent leaning positions in State H, $Z_{i\mathrm{H}}$, are drawn from a normal distribution centered at -0.75 for $i \in \{1, \dots, 10\}$ and at 0.75 for $i \in \{11, \dots, 20\}$ with a standard deviation equal to 0.15, while the latent leaning positions in State L, $Z_{i\mathrm{L}}$, are drawn from a normal centered around the positions -0.25 for $i \in \{1, \dots, 10\}$ and 0.25 for $i \in \{11, \dots, 20\}$ also with a standard deviation of 0.15; the individual effect parameters α_i are randomly drawn from a normal distribution with $\mu_\alpha = 0$ and $\sigma_\alpha = 2$; the transition probability matrix $\boldsymbol{Q}$ has 0.95 on the diagonal and 0.05 as off-diagonal elements, and the sequence of states is randomly drawn from the Markov chain initialized at State L; finally, we set $\phi = 200$, $\gamma_0 = -0.1$, and $\gamma_1 = 0.5$. We sample Y_{ijt} and L_{it} from the data generating process (1)–(5). Our simulation represents a situation in which news outlets diverge in magnitude—via α_i—and in terms of political leaning—via Z_{ik}.

We run our MCMC algorithm for 50,000 iterations, and we discard the first 30,000 iterations as burn-in and thin by a factor of 10 to reduce autocorrelation in the draws. To correctly identify left- and right-leaning, we consider news outlet 3 as known to be left.

Figure 5 reports a summary of the simulation results. From a comparison between the true values and the marginal posterior distributions, the model performs very well in terms of identification of the individual effects and latent variables (Panel A), the latent states (Panel B) as well as the other parameters in the simulation (Panel C). Credible regions for the pairs (α_i, Z_{ik}) estimated with our model (solid ellipses in Panel A) are narrower than those obtained disregarding the political-leaning proxy L_{it}, that is, dropping (3) from the model

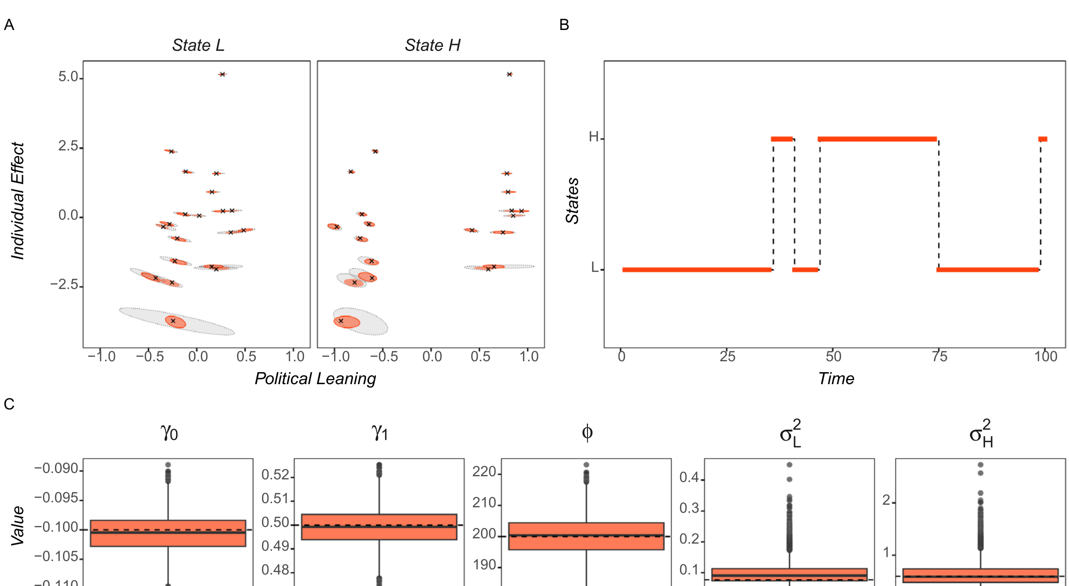

FIG. 5. *Simulated data: Panel A: Estimated Coordinates in the Latent Leaning—Individual Effect Plane with 99% credible ellipses for the mode with (3) (solid inner ellipses), without (3) (dashed outer ellipses) and true values (crosses). Panel B: Estimated latent states (squared dots) and true states (dashed step line). Panel C: Boxplots of the marginal posteriors for the parameters* γ_0, γ_1, ϕ, σ_L^2, *and* σ_H^2, *with dashed lines indicating the true values.*

(dashed ellipses). This suggests that the information-borrowing strategy is effective in improving the estimation accuracy of the latent variables. As expected, the detection of the regimes deteriorates when the regimes are not well separated in the latent space (see Appendix E, Casarin, Peruzzi and Steel (2025)). Properties of the MCMC chains are reported in the Supplementary Material (Appendix D, Casarin, Peruzzi and Steel (2025)). Our MCMC algorithm is implemented in *R* and *C*++, and we have made the scripts freely available (see Appendix M.2, Casarin, Peruzzi and Steel (2025)).

4. Political leaning and polarization of news outlets. We apply our model to a dataset of daily Facebook activities related to 225 national and local news outlets in France, Germany, Italy, and Spain. We provide both static and dynamic analyses to assess the media slant in these outlets as well as the polarization regimes across countries.

4.1. *Dataset description and construction.* Our novel media temporal network dataset, *network dataset* in what follows, comprises the news outlets in the list reported in the Reuters Digital News Report (2017) (Newman et al. (2017)). The tick-by-tick data in the *source dataset* of Schmidt et al. (2018) contains all posts published, along with their associated metadata, as well as all data on anonymized user interactions with these posts, in the form of comments. Table 1 reports a summary description of the source dataset. We aggregate comment interactions to daily data at the outlet level by constructing the set of daily bipartite networks between news outlets and *commenters*—those Facebook users commenting on news outlets' posts. For each country, that is, *France*, *Germany*, *Italy*, and *Spain*, at time t, we obtain the set of audience-duplication networks $\mathcal{G}_t$ presented in Figure 6 by performing the one-mode projection on the side of news outlets (as in Figure 3).

We complement our network dataset with data from Crowdtangle (CrowdTangle Team (2022)) and Chapel Hill Expert Survey (CHES) data (Polk et al. (2017)). Crowdtangle allows for retrieving Facebook posts for public pages and provides additional metadata for each post. In particular, the fields *Link Text* and *Description* contain information on the text of linked pages, such as the texts of news articles published on the Facebook walls of the news outlets. There is not a perfect match between all the pages available in the source dataset of Schmidt et al. (2018) and those available in Crowdtangle. Some news outlets may have changed account or ceased to exist. In this case, information about these news outlets may no longer be available on Facebook at the time of writing. The CHES questions political scientists on different aspects related to politics and European integration. The *CHES dataset* contains all the information at an aggregate level about scientists' opinions on the ideological position of political parties in Europe. Here we will make use of the *lrgen* variable, which provides the ideological stance of a political party from 0 (extreme left) to 10 (extreme right). The information retrieved from Crowdtangle and CHES allows us to construct a text-analysis proxy for media slant. In particular, we obtain our observed proxy for daily media slant L_{it} by

TABLE 1
Source dataset description: Description of the Facebook dataset on national and local news outlets of the four European countries (France, Germany, Italy, and Spain) gathered by Schmidt et al. (2018). The dataset entirely covers the years 2015 *and* 2016. *The news outlet list is reported in the Reuters Digital News Report* (2017)

Country	Pages	Posts	Comments	Commenters
France	65	1,008,018	47,225,675	5,755,268
Germany	49	749,805	31,881,407	5,338,195
Italy	54	1,554,817	51,515,121	4,086,351
Spain	57	1,372,805	34,336,356	6,494,725

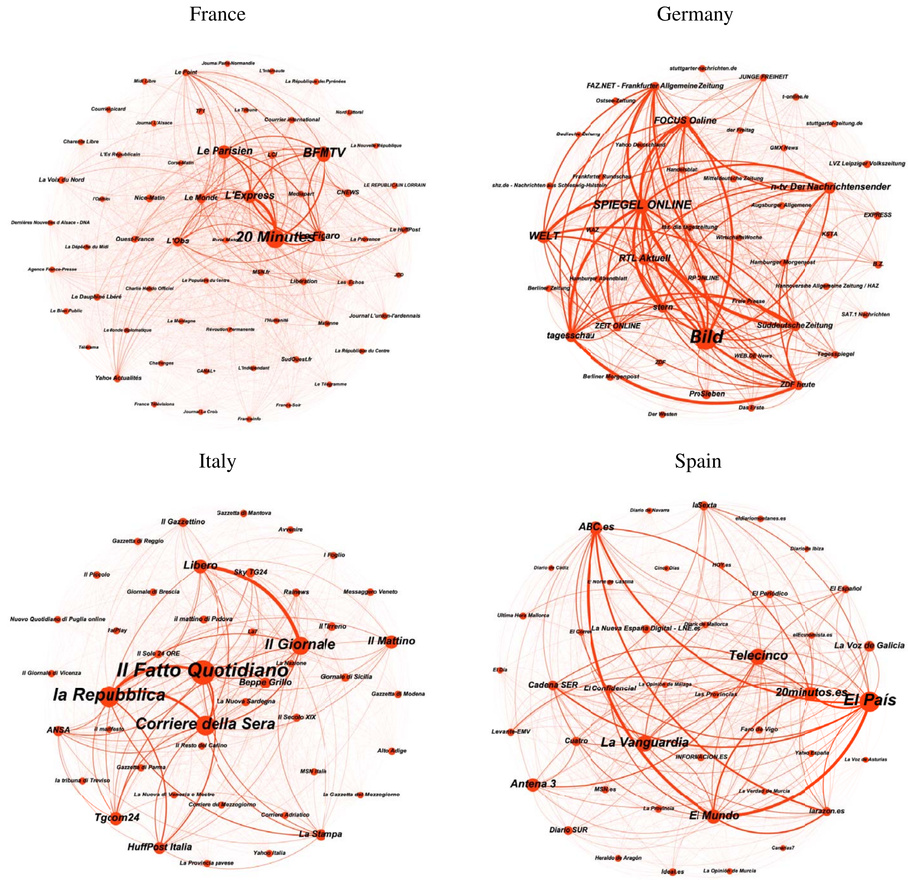

FIG. 6. *Audience-Duplication Networks*: *Cumulative networks for France* (62 *outlets*), *Germany* (47 *outlets*), *Italy* (45 *outlets*), *and Spain* (43 *outlets*) *from January* 1, 2015 *to December* 31, 2016. *Node size is proportional to the cumulative number of comments received in the time interval. Edge thickness is proportional to the number of common commenters between each pair of outlets. Quantities are normalized for each country.*

computing the index proposed by Gentzkow and Shapiro (2010) and adapted to online media outlets by Garz, Sörensen and Stone (2020). Such a media slant index relies on text analysis techniques to assess the similarity between pieces by news outlets and texts published by politicians. We then associate a political leaning to each news outlet as a function of this similarity and the parties' political leaning. Further information on the adopted methodology can be found in Appendix G (Casarin, Peruzzi and Steel (2025)).

The network dataset and the media slant index are publicly available, as described in Appendix M.1 (Casarin, Peruzzi and Steel (2025)).

4.2. *Results from a static analysis*. First, we implement a static version of our model with $d = 1$ on the whole two-year time period, without the MS dynamic component. For this we use an overall audience duplication network $\tilde{\mathcal{G}}_t$ for each country, where the weighted edge for each pair of outlets is $\tilde{Y}_{ij} = \sum_{t=1}^{T} Y_{ijt}$, and the overall observed leaning-feature is constructed as $\tilde{L}_i = T^{-1} \sum_{t=1}^{T} L_{it}$. Panels A, B, C, and D in Figure 7 report the estimated (posterior mean) latent coordinates in the latent leaning-individual effect space. We notice how the individual effect parameter α_i associated with each news outlet and country may be interpreted in terms of news outlet's engagement, as major national news outlets are concentrated at the top in each graph.

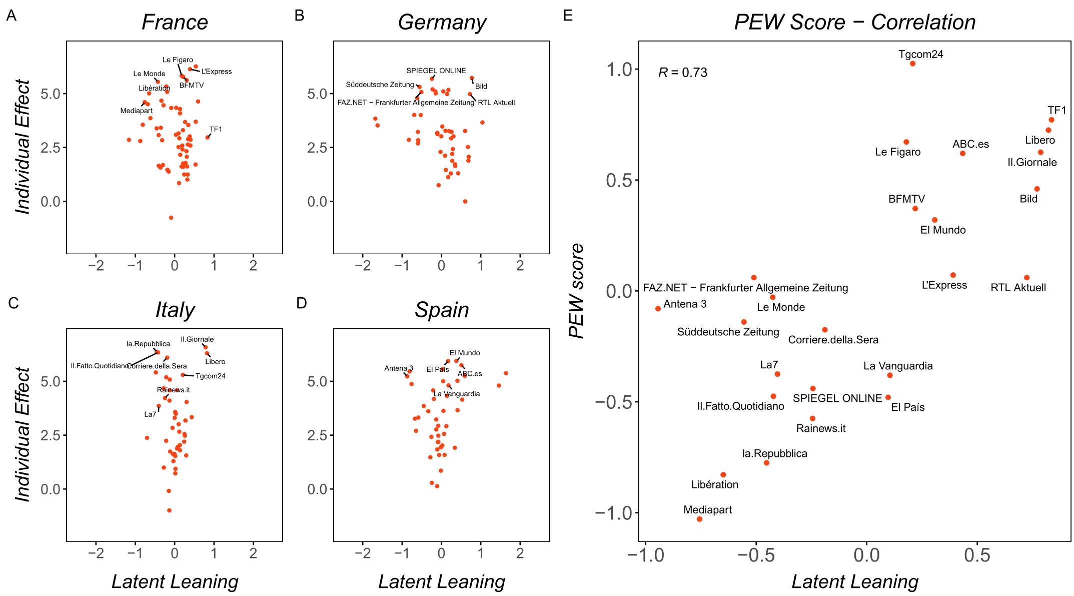

FIG. 7. *PEW Index—latent leaning comparison for the static model*: *Panels A*, *B*, *C*, *and D present the posterior mean of the latent coordinates of our news outlets*, *while Panel E displays the scatter plot comparing the PEW survey results*, *available for* 25 *major national news outlets*, *with our estimated latent leaning variable* (*country-specific means have been subtracted from the PEW scores to improve readability*).

We correlate our posterior mean media slant with the results obtained by the PEW Research Survey (Mitchell et al. (2018)). In this survey, participants were asked to assess the left-right leaning of major national news outlets (25 of these also appear in our dataset) on a 0–6 scale with 0 indicating far left and 6 indicating far right. To the best of our knowledge, a survey assessing the left-right leaning of all the outlets in our dataset is not available. We will refer to the left-right ranking obtained by PEW Research as the PEW Research index. We find that the PEW Research index has a 0.73 correlation with our estimated latent leaning; see Panel E in Figure 7. Moreover, we notice the presence of both a left-leaning cluster (bottom-left) and a right-leaning cluster (top-right). As a further validation, Figure L.1 in Appendix L (Casarin, Peruzzi and Steel (2025)) provides a comparison between the estimated leaning and the partition obtained via traditional cluster analysis, using the Fast Greedy algorithm (Clauset, Newman and Moore (2004)). This algorithm provides an optimal number of clusters larger than two and in itself is not able to discriminate between left and right-leaning outlets or rank news outlets on the political spectrum. Nevertheless, a graphical inspection shows that both the PEW score and our leaning measure can be used to separate the clusters (different symbols) into two large groups: left-wing outlets (bottom-left quadrant) and right-wing outlets (top-right quadrant), coherently with the cluster analysis partition.

Figure L.2 presents the marginal posterior distribution of the parameters γ_0, γ_1, and ϕ. The parameter γ_1 conveys information on the relationship between the latent variable x_i and the observed leaning proxy L_i. Latent leaning appears to be a strong driver for the observed proxy only in the case of Italy, as the posterior mass of γ_1 is located far away from zero, while it seems a weak driver for France and mostly irrelevant for both Germany and Spain. Nonetheless, the strong correlation with the PEW Research index suggests that having information on online users' interactions with news outlets may still be sufficient to provide an effective classification on the political spectrum.

4.3. *Results from a dynamic analysis*. We now estimate the model described in Section 2.1 in its dynamic specification with $d = 2$ and $K = 2$ (called $\mathcal{M}_4$ in the next subsec-

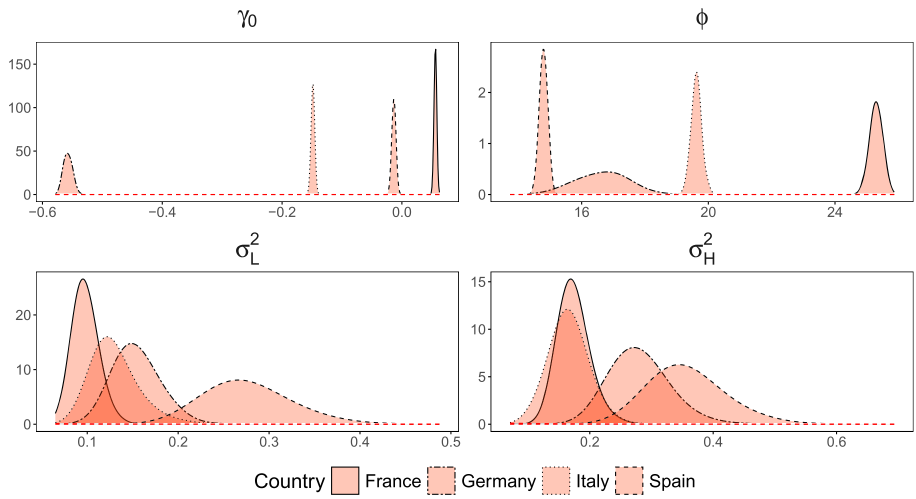

FIG. 8. *Marginal posteriors for the dynamic model $\mathcal{M}_4$: Kernel posterior density estimates for the parameters γ_0, ϕ, σ^2_L and σ^2_H for France, Germany, Italy, and Spain. The corresponding prior distributions (dashed lines) are flat and nearly indistinguishable from the horizontal axis.*

tion) and with $d = 2$ and $K = 5$ ($\mathcal{M}_6$). The former specification is instrumental for a clear representation of the main characteristics of the latent space, while the latter is the preferred model (see the next section), which also provides an improved fit of the dynamic network features. Model selection is discussed in some detail in the next section. The choice of a different model for each country is also legitimate. The dynamic analysis uses daily data, and we deleted from our dataset those outlets that remained inactive—that is, did not receive any comment—for more than 15 consecutive days. Overall, we removed 13 news outlets (DE: four outlets, FR: five, IT: two, SP: two; see Appendix H, Casarin, Peruzzi and Steel (2025)), which displayed unusual behavior. Posterior results for the parameters of the MS-LS model with $d = 2$ and $K = 2$ are presented in Figure 8. As expected, values of σ^2_{H} tend to be larger than those for σ^2_{L}. In Figure 9 we report the posterior means of the latent positions for the four countries in states L and H. The individual-effect values are coherent with the engagement interpretation in both states: well-known national newspapers display larger individual effects (larger point size), while local newspapers display smaller individual effects (smaller point size). Moreover, our latent variable dimension with media bias interpretation shows a positive correlation with the PEW Research Survey Index in this setting. Figure 10 illustrates the correlation of 0.66 in the lower polarisation state and 0.62 in the state of higher polarisation. Finally, the second latent coordinate captures news outlets' similarities in other distinctive characteristics, such as geographic vicinity or editorial ownership.

Latent states signal the presence of lower or higher in-platform polarization regimes. In state L the average distance between outlets is lower in terms of political leaning than in state H, making Facebook users more prone to interact with news outlets with different political tendencies. Figure 11 reports the estimated polarization regimes (bottom panels) through time for the MS-LS model, with $d = 2$ and $K = 5$ estimated for the four countries. State 1 denotes the lowest polarization state, while State 5 indicates the highest polarization state. We notice a tendency to move from a state of high in-platform polarization to a lower polarization for France, Germany, and Spain. Italy and, to some extent, Spain follow an oscillating pattern, with Spain exhibiting frequent switches between very-high and very-low polarization. Overall, our findings contradict the hypothesis of a global shift toward a high-polarization regime on social media in this time frame. In line with the opinion dynamics literature in media contexts (e.g., see Hu, Chen and Hu (2024), Brooks and Porter (2020)), polarization

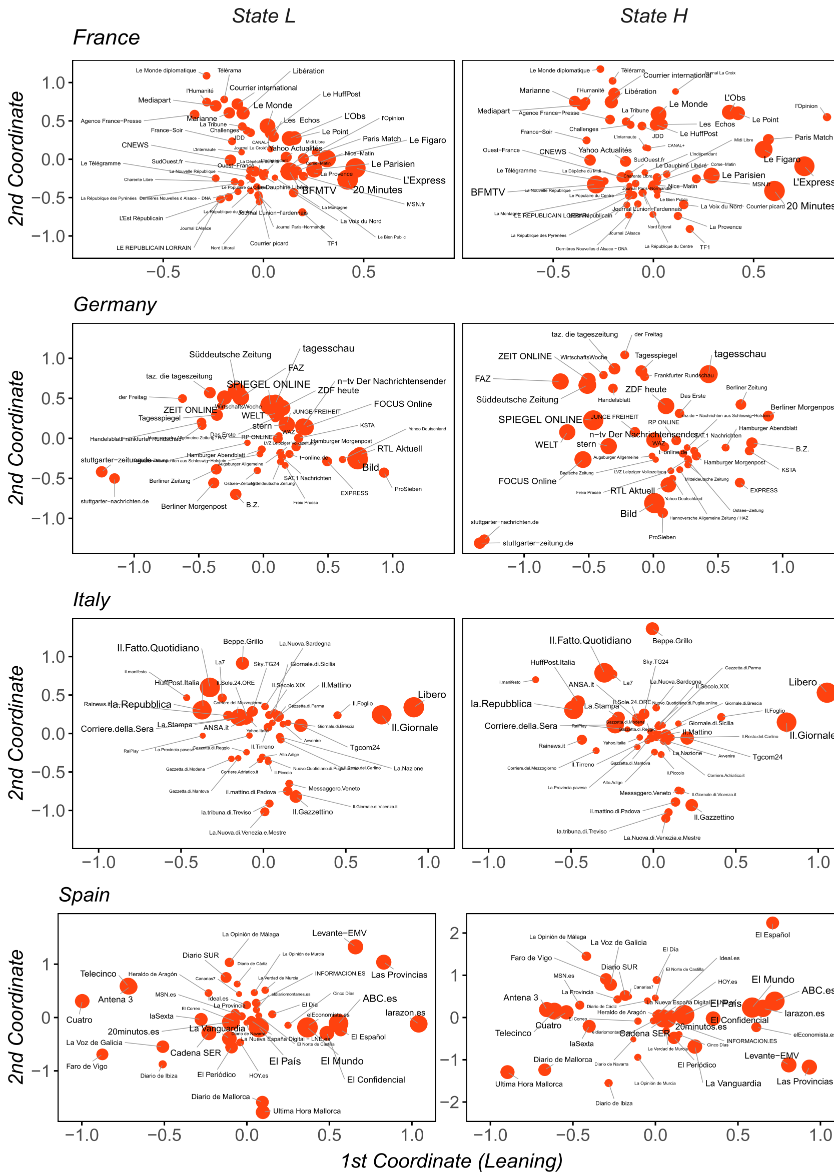

FIG. 9. *Latent positions for the dynamic model $\mathcal{M}_4$: Posterior means of latent coordinates of the news outlets for France, Germany, Italy, and Spain in State L and in State H. The node size is proportional to the posterior mean of the individual effects.*

within our model is strongly interconnected with the degree of interactions within a network (expected strength). In the opinion dynamics literature, weak network connectivity and fragmented structures (e.g., echo chambers) can favor the emergence of polarization driven by ideological differences. Overall, our model agrees with these media theories, since periods of higher polarization (high values in the bottom panel of Figure 11) correspond to periods of lower expected strength in the observed networks (low values in the top panel). Despite not being the primary purpose of our analysis, polarization changes can be used to effectively track the medium-term dynamics of the expected strength (boxplots in the top panels) by borrowing information across periods classified within the same polarization regime. We stress that our main aim is not to replicate the short-term features of the daily observations, but

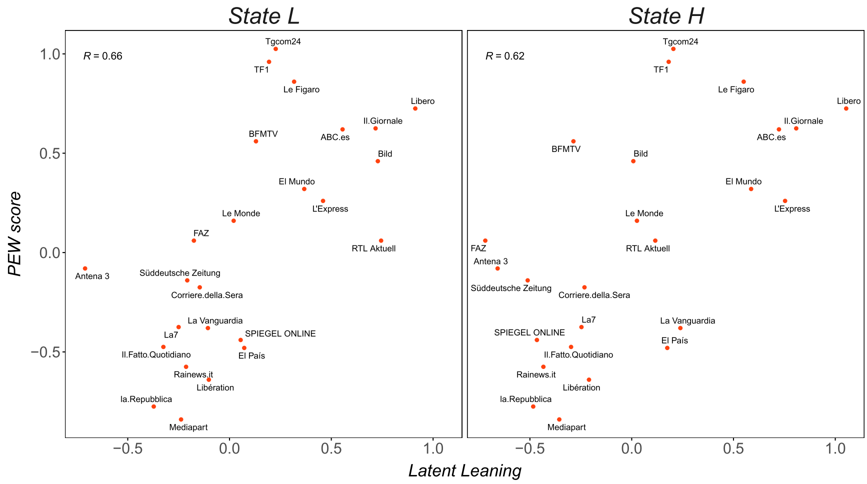

FIG. 10. *PEW Index—latent leaning comparison for dynamic model $\mathcal{M}_4$: Scatter plot comparing the PEW survey results with the estimated latent leaning variable in State L and H. The country-specific mean has been subtracted from the PEW Score to improve readability.*

rather to conduct inference on polarization and media bias regimes; see Appendix K (Casarin, Peruzzi and Steel (2025)) for a discussion on alternative MS-LS dynamic specifications.

4.4. *Model selection.* We perform model selection considering eight alternative models. $\mathcal{M}_1$ is an MS-LS model with $d = 1$ and $K = 2$ omitting the text-analysis interpretation in (3) (i.e., imposing $\gamma_1 = 0$), $\mathcal{M}_2$ is an LS model with $d = 1$ and $K = 1$ omitting the Markov-switching dynamics described in (5), $\mathcal{M}_3$ is an unrestricted MS-LS model with $d = 1$ and $K = 2$, $\mathcal{M}_4$ is an unrestricted MS-LS model with $d = 2$ and $K = 2$, $\mathcal{M}_5$ is an unrestricted MS-LS model with $d = 2$ and $K = 3$, and $\mathcal{M}_6$ is an unrestricted MS-LS model with $d = 2$ and $K = 5$. In addition, we consider the standard Poisson random graph model RG_1 with

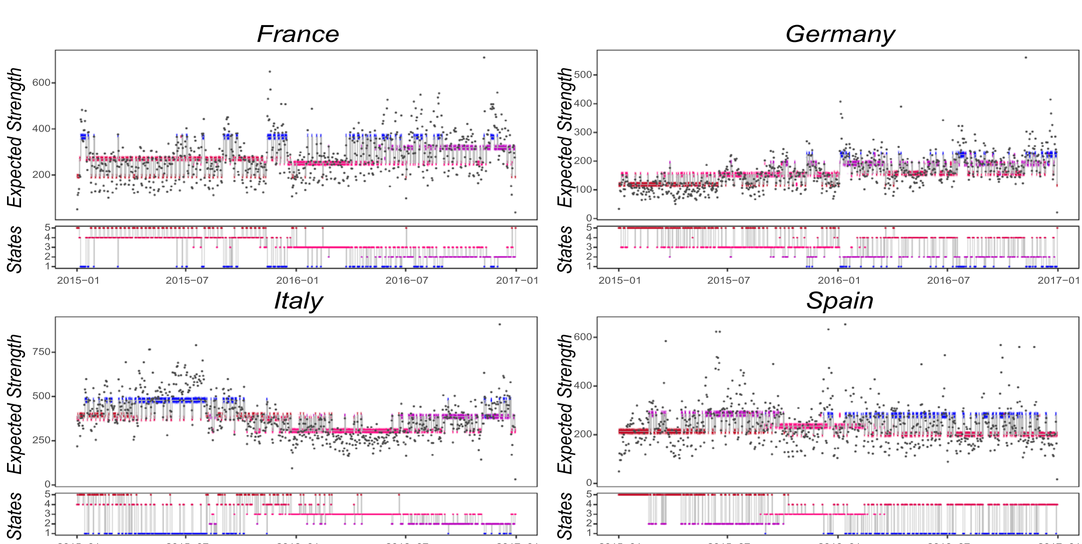

FIG. 11. *Latent states for dynamic MS-LS model $\mathcal{M}_6$: Top panels show the posterior distribution boxplots for the average expected strength in each regime, with the underlying daily observations for information in the background (gray dots); Bottom panels show the polarization states through time, where 1 denotes the lowest polarization state and 5 the highest polarization state. In the online version, colors in both panels relate to the polarization level, with red corresponding to high, purple to intermediate, and blue to low polarization.*

TABLE 2

Network model selection: (*a*) *DIC scores and* (*b*) *Log pointwise predictive densities* (*lppd*), *as in* Gelman, Hwang and Vehtari (2014), *for the models introduced here and two random graph models* RG_1 *and* RG_2 *with intensity* $\lambda_{ijt} = \exp\{\alpha\}$ *and* $\lambda_{ijt} = \exp\{\alpha_i + \alpha_j - \beta\|L_{it} - L_{jt}\|^2\}$, *respectively. Best performance in bold*

	(a) DIC $\times 10^{-6}$				(b) lppd $\times 10^{-6}$			
Model	France	Germany	Italy	Spain	France	Germany	Italy	Spain
$\mathcal{M}_1$	4.4698	2.3669	3.3066	4.6390	−2.2784	−1.2190	−1.6771	−2.3471
$\mathcal{M}_2$	4.6434	2.4766	3.4825	4.9049	−2.3597	−1.2702	−1.7657	−2.4511
$\mathcal{M}_3$	4.4696	2.3669	3.3049	4.6139	−2.2784	−1.2191	−1.6776	−2.3347
$\mathcal{M}_4$	4.2654	2.2582	2.9797	4.2796	−2.1697	−1.1500	−1.5171	−2.1576
$\mathcal{M}_5$	4.2533	2.2143	2.9766	4.1819	−2.1644	−1.1290	−1.5127	−2.1075
$\mathcal{M}_6$	**4.0949**	**2.1570**	**2.6588**	**3.9590**	**−2.0827**	**−1.1107**	**−1.4641**	**−2.0633**
RG_1	24.9489	9.6074	26.1551	15.1828	−12.5047	−4.8341	−13.1011	−7.6049
RG_2	5.1913	2.8212	4.4554	5.4713	−2.6153	−1.4121	−2.2424	−2.7350

intensity $\lambda_{ijt} = \exp\{\alpha\}$ and the Poisson random graph model with individual effects and observed leaning distances RG_2 with intensity $\lambda_{ijt} = \exp\{\alpha_i + \alpha_j - \beta\|L_{it} - L_{jt}\|^2\}$.

Model selection is carried out via two popular predictive measures, the Deviance Information Criterion (DIC) (Spiegelhalter et al. (2002)) and the log pointwise predictive density (lppd) of Gelman, Hwang and Vehtari (2014). Table 2 reports both criteria (see Appendix F, Casarin, Peruzzi and Steel (2025), for further details). A first comparison conducted using the DIC and lppd of model $\mathcal{M}_1$, $\mathcal{M}_2$, and $\mathcal{M}_3$ highlights the contributions of the observed leaning *a-là* Gentzkow, Shapiro and Stone (2015) and the dynamic component. The model without the dynamic component ($\mathcal{M}_2$) is dominated by the other two specifications for each country. Except for Spain and to some extent Italy, very similar scores are obtained for $\mathcal{M}_1$ and $\mathcal{M}_3$. This is to be expected as $\mathcal{M}_3$ aims to offer more interpretable latent coordinates rather than an improved fit for the network. Table 2 also includes results for the random graph models RG_1 and RG_2, which are clearly performing worse than our models. A comparison between models $\mathcal{M}_3$, $\mathcal{M}_4$, $\mathcal{M}_5$, and $\mathcal{M}_6$ highlights the effect of increasing the number of dimensions. The best-performing model on the basis of DIC and lppd is $\mathcal{M}_6$ ($d = 2$ and $K = 5$). Graphical inspection reveals that the latent space results for $\mathcal{M}_6$ are in line with those of the other models (see Appendix I, Casarin, Peruzzi and Steel (2025), for additional results).

In addition, we compare the empirical and model-implied network metrics averaged over the entire sample. Table I.1 reports a posterior predictive check (see Gelman, Hwang and Vehtari (2014)) in which the posterior predictive expected nodal strength, nodal strength's standard deviation and dispersion index—derived in Section 2.2—are compared with the empirical values, for the random graph models RG_1 and RG_2, and models $\mathcal{M}_3$ to $\mathcal{M}_6$. All models are able to mimic the first moment of the strength distribution, but, as expected, the simple Poisson random graph model can not accommodate the large amount of overdispersion in the data. The MS-LS models are able to capture the observed dispersion in the strength distribution, leading to a rather similar fit. Appendix I in the Supplementary Material (Casarin, Peruzzi and Steel (2025)) presents time-varying diagnostic figures for these metrics. Appendix L.2 (Casarin, Peruzzi and Steel (2025)) provides a comparison with SBMs, which generally offer a good fit and provide an alternative modeling solution for applications where media clustering is more relevant than ranking media on the political spectrum.

5. Conclusion. We propose a dynamic Markov-switching latent space model to extract insightful information concerning media ideology and in-platform polarization. The model projects the audience duplication network of news outlets on a d-dimensional Euclidean space where the latent positions can be interpreted in terms of political leaning through a suitable proxy. Inference is carried out within a Bayesian framework, allowing reliable results with relatively standard MCMC methods. We derive the theoretical model properties and assess the effectiveness of the proposed methodology on simulated data. Our model is applied to a Facebook dataset of news outlets in four European countries. We find that the inferred latent leaning strongly correlates with the independent PEW Research Survey Index and correctly ranks news outlets in terms of left- and right-leaning. Moreover, inference on the latent states does not support the hypothesis of a unidirectional shift toward high polarization on Facebook. Finally, model selection suggests that the dynamic specification should be preferred over a static one and that a text-analysis index may aid the leaning identification. Our novel modeling technique is not free from limitations. First, few longitudinal studies have focused on polarization in general and more specifically on media polarization in Europe. This makes it difficult to find a ground truth for validating polarization changes. Moreover, latent factor and state identification is not trivial in this context. Finally, the current model does not benefit from cross-country information-sharing which could be obtained, *for example*, with a suitable hierarchical prior choice. Future research may be directed toward the testing of more advanced text-analysis indicators for media bias, a more fine-grained analysis, *for example*, at post level, the adoption of alternative latent space embedding specifications, such as circular and hyperbolic, and the exploration of different identifying restrictions.

Acknowledgments. The authors are grateful to the Editors and the anonymous Referee for the insightful comments, and to A.L. Schmidt, F. Zollo, A. Scala, and W. Quattrociocchi for providing access to the Facebook *source* dataset used in this work. We are also grateful to the conference participants for helpful discussions at the ISBA World Meeting (June 2022, Montreal, Canada), Bayes Comp (March 2023, Levi, Finland) and the 10th Italian Congress of Econometrics and Empirical Economics (May 2023, Cagliari, Italy).

This research used the HPC system of the “Venice Centre in Economic and Risk Analytics for Public Policies” (VERA).

MFJS gratefully acknowledges the hospitality of Ca’ Foscari University of Venice, where he was a visiting scholar during the latter stages of this research.

The views and opinions expressed are solely those of the authors and do not necessarily reflect those of the EU.

Funding. RC and AP acknowledge support from the Italian Ministry of University—PRIN project “Discrete random structures for Bayesian learning and prediction” under g.a. n. 2022CLTYP4; the European Union (EU)—NextGenerationEU, Growing Resilient, Inclusive and Sustainable project (GRINS PE00000018—CUP H73C22000930001), National Recovery and Resilience Plan (NRRP)—PE9.

SUPPLEMENTARY MATERIAL

Supplement to “Media bias and polarization through the lens of an MS-LS model” (DOI: 10.1214/25-AOAS2069SUPPA; .pdf). The online Supplementary Material include derivations of the results and further analyses.

Repository for “Media bias and polarization through the lens of an MS-LS model” (DOI: 10.1214/25-AOAS2069SUPPB; .zip). This online repository contains the data and code used to replicate the results in the paper. The online repository is also available at https://github.com/BayesianEcon/Dyn-MS-LS-Media and at https://codeocean.com/capsule/9380600/tree/v1.

REFERENCES

ANDERSON, S. P. and MCLAREN, J. (2012). Media mergers and media bias with rational consumers. *J. Eur. Econ. Assoc.* **10** 831–859. https://doi.org/10.1111/j.1542-4774.2012.01069.x

BARBERÁ, P. (2015). Birds of the same feather tweet together: Bayesian ideal point estimation using Twitter data. *Polit. Anal.* **23** 76–91. https://doi.org/10.1093/pan/mpu011

BRANSCUM, A. J., JOHNSON, W. O. and THURMOND, M. C. (2007). Bayesian beta regression: Applications to household expenditure data and genetic distance between foot-and-mouth disease viruses. *Aust. N. Z. J. Stat.* **49** 287–301. MR2405396 https://doi.org/10.1111/j.1467-842X.2007.00481.x

BROOKS, H. Z. and PORTER, M. A. (2020). A model for the influence of media on the ideology of content in online social networks. *Phys. Rev. Res.* **2** 023041. https://doi.org/10.1103/PhysRevResearch.2.023041

CASARIN, R., DALLA VALLE, L. and LEISEN, F. (2012). Bayesian model selection for beta autoregressive processes. *Bayesian Anal.* **7** 385–409. MR2934956 https://doi.org/10.1214/12-BA713

CASARIN, R., PERUZZI, A. and STEEL, M. F (2025). Supplement to "Media bias and polarization through the lens of a Markov switching latent space network model." https://doi.org/10.1214/25-AOAS2069SUPPA, https://doi.org/10.1214/25-AOAS2069SUPPB

CHEN, C. Y.-H., OKHRIN, Y. and WANG, T. (2024). Monitoring network changes in social media. *J. Bus. Econom. Statist.* **42** 391–406. MR4728979 https://doi.org/10.1080/07350015.2021.2016425

CINELLI, M., MORALES, G. D. F., GALEAZZI, A., QUATTROCIOCCHI, W. and STARNINI, M. (2021). The echo chamber effect on social media. *Proc. Natl. Acad. Sci. USA* **118** e2023301118. https://doi.org/10.1073/pnas.2023301118

CLAUSET, A., NEWMAN, M. E. and MOORE, C. (2004). Finding community structure in very large networks. *Phys. Rev. E* **70** 066111. https://doi.org/10.1103/PhysRevE.70.066111

CROWDTANGLE TEAM (2022). *CrowdTangle*. Meta, Menlo Park, CA.

D'ANGELO, S., MURPHY, T. B. and ALFÒ, M. (2019). Latent space modelling of multidimensional networks with application to the exchange of votes in Eurovision Song Contest. *Ann. Appl. Stat.* **13** 900–930. MR3963557 https://doi.org/10.1214/18-AOAS1221

DANDEKAR, P., GOEL, A. and LEE, D. T. (2013). Biased assimilation, homophily, and the dynamics of polarization. *Proc. Natl. Acad. Sci. USA* **110** 5791–5796. MR3065807 https://doi.org/10.1073/pnas.1217220110

DE NICOLA, G., TUEKAM MAMBOU, V. H. and KAUERMANN, G. (2023). COVID-19 and social media: Beyond polarization. *PNAS Nexus* **2** pgad246. https://doi.org/10.1093/pnasnexus/pgad246

ESTEBAN, J.-M. and RAY, D. (1994). On the measurement of polarization. *Econometrica* **62** 819–851. https://doi.org/10.2307/2951734

FRIEL, N., RASTELLI, R., WYSE, J. and RAFTERY, A. E. (2016). Interlocking directorates in Irish companies using a latent space model for bipartite networks. *Proc. Natl. Acad. Sci. USA* **113** 6629–6634. https://doi.org/10.1073/pnas.1606295113

FRÜHWIRTH-SCHNATTER, S. (2006). *Finite Mixture and Markov Switching Models*. *Springer Series in Statistics*. Springer, New York. MR2265601

FRÜHWIRTH-SCHNATTER, S., HOSSZEJNI, D. and LOPES, H. F. (2025). Sparse Bayesian factor analysis when the number of factors is unknown (with discussion). *Bayesian Anal.* **20** 213–344. MR4888145 https://doi.org/10.1214/24-ba1423

GARIMELLA, K., MORALES, G. D. F., GIONIS, A. and MATHIOUDAKIS, M. (2018). Quantifying controversy on social media. *ACM Trans. Soc. Comput.* **1** 1–27. https://doi.org/10.1145/3140565

GARZ, M., SÖRENSEN, J. and STONE, D. F. (2020). Partisan selective engagement: Evidence from Facebook. *J. Econ. Behav. Organ.* **177** 91–108. https://doi.org/10.1016/j.jebo.2020.06.016

GELMAN, A., HWANG, J. and VEHTARI, A. (2014). Understanding predictive information criteria for Bayesian models. *Stat. Comput.* **24** 997–1016. MR3253850 https://doi.org/10.1007/s11222-013-9416-2

GENTZKOW, M. and SHAPIRO, J. M. (2010). What drives media slant? Evidence from U.S. daily newspapers. *Econometrica* **78** 35–71. MR2642856 https://doi.org/10.3982/ECTA7195

GENTZKOW, M., SHAPIRO, J. M. and STONE, D. F. (2015). Media bias in the marketplace: Theory. In *Handbook of Media Economics* **1** 623–645. Elsevier, Amsterdam. https://doi.org/10.1016/B978-0-444-63685-0.00014-0

GOLLINI, I. and MURPHY, T. B. (2016). Joint modeling of multiple network views. *J. Comput. Graph. Statist.* **25** 246–265. MR3474046 https://doi.org/10.1080/10618600.2014.978006

HANDCOCK, M. S., RAFTERY, A. E. and TANTRUM, J. M. (2007). Model-based clustering for social networks. *J. Roy. Statist. Soc. Ser. A* **170** 301–354. MR2364300 https://doi.org/10.1111/j.1467-985X.2007.00471.x

HANUSCH, F. and NÖLLEKE, D. (2019). Journalistic homophily on social media: Exploring journalists' interactions with each other on Twitter. *Dig. Journal.* **7** 22–44. https://doi.org/10.1080/21670811.2018.1436977

HOFF, P. D., RAFTERY, A. E. and HANDCOCK, M. S. (2002). Latent space approaches to social network analysis. *J. Amer. Statist. Assoc.* **97** 1090–1098. MR1951262 https://doi.org/10.1198/016214502388618906

HU, H., CHEN, W. and HU, Y. (2024). Opinion dynamics in social networks under the influence of mass media. *Appl. Math. Comput.* **482** Paper No. 128976, 17. MR4778178 https://doi.org/10.1016/j.amc.2024.128976

IYENGAR, S., SOOD, G. and LELKES, Y. (2012). Affect, not ideology: A social identity perspective on polarization. *Public Opin. Q.* **76** 405–431. https://doi.org/10.1093/poq/nfs038

KIM, B., LEE, K. H., XUE, L. and NIU, X. (2018). A review of dynamic network models with latent variables. *Stat. Surv.* **12** 105–135. MR3850294 https://doi.org/10.1214/18-SS121

KUBIN, E. and VON SIKORSKI, C. (2021). The role of (social) media in political polarization: A systematic review. *Ann. Int. Commun. Assoc.* **45** 188–206. https://doi.org/10.1080/23808985.2021.1976070

LEE, C. and WILKINSON, D. J. (2019). A review of stochastic block models and extensions for graph clustering. *Appl. Netw. Sci.* **4** 1–50. https://doi.org/10.1007/s41109-019-0232-2

LEONARD, N. E., LIPSITZ, K., BIZYAEVA, A., FRANCI, A. and LELKES, Y. (2021). The nonlinear feedback dynamics of asymmetric political polarization. *Proc. Natl. Acad. Sci. USA* **118** Paper No. e2102149118, 11. MR4401236 https://doi.org/10.1073/pnas.2102149118

MACY, M. W., MA, M., TABIN, D. R., GAO, J. and SZYMANSKI, B. K. (2021). Polarization and tipping points. *Proc. Natl. Acad. Sci. USA* **118** e2102144118. https://doi.org/10.1073/pnas.2102144118

MITCHELL, A., SIMMONS, K., MATSA, K. E., SILVER, L., SHEARER, E., JOHNSON, C., WALKER, M. and TAYLOR, K. (2018) In Western Europe, Public Attitudes Toward News Media More Divided by Populist Views than Left-Right Ideology. PEW Research Center.

NEWMAN, N., FLETCHER, R., LEVY, D. and NIELSEN, R. K. (2017). Reuters Institute Digital News Report 2017.

NG, T. L. J., MURPHY, T. B., WESTLING, T., MCCORMICK, T. H. and FOSDICK, B. (2021). Modeling the social media relationships of Irish politicians using a generalized latent space stochastic blockmodel. *Ann. Appl. Stat.* **15** 1923–1944. MR4355082 https://doi.org/10.1214/21-aoas1483

PARK, J. H. and SOHN, Y. (2020). Detecting structural changes in longitudinal network data. *Bayesian Anal.* **15** 133–157. MR4050880 https://doi.org/10.1214/19-BA1147

PEIXOTO, T. P. (2019). Bayesian stochastic blockmodeling. In *Advances in Network Clustering and Blockmodeling* **11** 289–332. Wiley, New York. https://doi.org/10.1002/9781119483298.ch11

PENG, Y. and YANG, T. (2022). Anatomy of audience duplication networks: How individual characteristics differentially contribute to fragmentation in news consumption and trust. *New Media Soc.* **24** 2270–2290. https://doi.org/10.1177/1461444821991559

POLK, J., ROVNY, J., BAKKER, R., EDWARDS, E., HOOGHE, L., JOLLY, S., KOEDAM, J., KOSTELKA, F., MARKS, G. et al. (2017). Explaining the salience of anti-elitism and reducing political corruption for political parties in Europe with the 2014 Chapel Hill expert survey data. *Res. Polit.* **4** 2053168016686915. https://doi.org/10.1177/2053168016686915

PRIOR, M. (2013). Media and political polarization. *Annu. Rev. Pol. Sci.* **16** 101–127. https://doi.org/10.1146/annurev-polisci-100711-135242

PUGLISI, R. and SNYDER, J. M. JR (2015). Empirical studies of media bias. In *Handbook of Media Economics* **1** 647–667. Elsevier, Amsterdam. https://doi.org/10.1016/B978-0-444-63685-0.00015-2

RASTELLI, R., FRIEL, N. and RAFTERY, A. E. (2016). Properties of latent variable network models. *Netw. Sci.* **4** 407–432. https://doi.org/10.1017/nws.2016.23

SCHMIDT, A. L., ZOLLO, F., SCALA, A. and QUATTROCIOCCHI, W. (2018). Polarization rank: A study on European news consumption on Facebook. Preprint. Available at arXiv:1805.08030. https://doi.org/10.48550/arXiv.1805.08030

SEWELL, D. K. and CHEN, Y. (2016). Latent space models for dynamic networks with weighted edges. *Soc. Netw.* **44** 105–116. https://doi.org/10.1016/j.socnet.2015.07.005

SOSA, J. and BETANCOURT, B. (2022). A latent space model for multilayer network data. *Comput. Statist. Data Anal.* **169** Paper No. 107432, 21. MR4369145 https://doi.org/10.1016/j.csda.2022.107432

SPIEGELHALTER, D. J., BEST, N. G., CARLIN, B. P. and VAN DER LINDE, A. (2002). Bayesian measures of model complexity and fit. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **64** 583–639. MR1979380 https://doi.org/10.1111/1467-9868.00353

TÖRNBERG, P., ANDERSSON, C., LINDGREN, K. and BANISCH, S. (2021). Modeling the emergence of affective polarization in the social media society. *PLoS ONE* **16** e0258259. https://doi.org/10.1371/journal.pone.0258259

WEF (2022). The Global Risks Report 2022 17th Edition Technical Report, World Economic Forum.

YARCHI, M., BADEN, C. and KLIGLER-VILENCHIK, N. (2021). Political polarization on the digital sphere: A cross-platform, over-time analysis of interactional, positional, and affective polarization on social media. *Polit. Commun.* **38** 98–139. https://doi.org/10.1080/10584609.2020.1785067

YU, X., LI, T., YING, N. and JING, B.-Y. (2022). Collaborative filtering with awareness of social networks. *J. Bus. Econom. Statist.* **40** 1629–1641. MR4492058 https://doi.org/10.1080/07350015.2021.1954527

ZHANG, Y., POUX-BERTHE, M., WELLS, C., KOC-MICHALSKA, K. and ROHE, K. (2018). Discovering political topics in Facebook discussion threads with graph contextualization. *Ann. Appl. Stat.* **12** 1096–1123. MR3834296 https://doi.org/10.1214/18-AOAS1191