



Modeling spatio-temporal dynamics of obesity in Italian regions via bayesian beta regression

Luciano Rota¹ · Raffaele Argiento² · Michela Cameletti² 

Received: 16 April 2026 / Accepted: 16 June 2026
© The Author(s) 2026

Abstract

In this paper, we investigate the spatio-temporal dynamics of obesity across Italian regions from 2010 to 2022, proposing an integrated Bayesian hierarchical framework tailored to the analysis of aggregated regional obesity data. In particular, we implement a Bayesian hierarchical Beta regression model to analyse regional obesity rates, integrating spatial and temporal random effects, alongside gender and several exogenous predictors. The model leverages the Stochastic Search Variable Selection technique to identify significant predictors supported by the data. The analysis reveals both gender and regional heterogeneity in obesity rates over the study period. Structured spatial and temporal random effects, along with gender, emerge as the primary determinants of obesity prevalence across Italian regions, while the role of exogenous covariates is found to be minimal at the regional level, once structured spatial and temporal effects and gender are accounted for. While socioeconomic and lifestyle factors remain fundamental at a micro-level, the findings suggest that the integration of spatial and temporal structures is crucial for capturing macro-level obesity variations.

Keywords Bayesian hierarchical model · Gender-specific intercepts · Intrinsic Conditional Auto-Regressive prior · Obesity rates · Stochastic Search Variable Selection

✉ Michela Cameletti
michela.cameletti@unibg.it

Luciano Rota
l.rota43@campus.unimib.it

Raffaele Argiento
raffaele.argiento@unibg.it

¹ Dipartimento di Economia, Metodi Quantitativi e Strategie di Impresa, Università di Milano-Bicocca, Piazza dell'Ateneo Nuovo, 1, 20126 Milan, MI, Italy

² Dipartimento di Scienze Economiche, Università degli studi di Bergamo, Via dei Caniana, 24127 Bergamo, BG, Italy

1 Introduction

Obesity is recognized as a major public health problem in Europe, where it has reached alarming levels, posing substantial health, social, and economic burdens on European societies. In 2014, one in six adults were considered obese according to the European Health Interview Survey (2016).¹ Moreover, projections suggest an increasing average obesity trend (Janssen et al. 2020).

Even if the problem is common among countries, such a phenomenon is not homogeneous: overweight and obesity vary substantially between and within countries. For instance, compared to other European countries, Italy shows relatively lower obesity rates for women (OECD 2022). According to the Eurostat data (Eurostat 2022), the second lowest proportion of obese people in 2019 was recorded in Italy (slightly above 11%), when nearby countries like Spain, France, Germany, and Austria are all above 14 %, and the European mean is around 17%. Despite this, the situation is still problematic and needs appropriate attention and intervention.

Moreover, national averages hide substantial heterogeneity within the countries, both in space and by gender, through time. Focusing on Italy, between northern and southern regions, there is a notable and time consistent difference in overweight and obesity rates. For example, taking the best and worst regions in the period considered by this study (2010–2022), our data show that the female (male) obesity rate is 7.50% (9.40%) in Trentino-Alto Adige and 12.7% (14.4%) in Molise. While considering overweight for the same regions, the percentages are 23.5% (40.3%) for Trentino-Alto Adige and 31% (46%) for Molise. It is well-known that across Italian regions, there are strong structural differences in general culture, social norms, and economic development (for a general discussion, see, for instance, De Blasio and Nuzzo 2009). This diversity among Italian regions may influence obesity and related phenomena (Gallus et al. 2013). For this reason, proposing an analysis that takes into account heterogeneity is essential (Brunello and Labartino 2014). Flexible models, such as varying coefficient models, have also been proposed in Assaf et al. (2015) to explore spatially and temporally varying effects on public health outcomes, including obesity, using large-scale survey data. Their findings show that this framework can capture dynamic changes over space and time in the effects of key risk factors, providing insight into the evolution of obesity across population subgroups. Their focus, however, differs from ours: while they concentrate on time and subgroup-specific variation in covariate effects, we are mainly interested in studying and interpreting the spatio-temporal pattern of obesity rates after accounting for covariates. Along with space and time, gender aspects should also be considered, given potential physical and genetic differences (Vari 2016) as well as cultural ones (Barra and Barone 2021).

In this paper, we analyse longitudinal areal data consisting of yearly regional obesity rates from 2010 to 2022, which were assembled through a systematic process of data collection, validation, and harmonization across multiple institutional sources. To model these rates, expressed as proportions bounded between 0 and 1, we employ a Beta regression model within a Bayesian hierarchical framework, using a mean–

¹ Eurostat news release, 203/2016, 20 October 2016, <https://ec.europa.eu/eurostat/documents/2995521/7700898/3-20102016-BP-EN.pdf/c26b037b-d5f3-4c05-89c1-00b0b98d646>.

precision parameterization. This approach provides a flexible structure for modelling rate data, enabling the inclusion of covariates as well as unobserved spatial and temporal heterogeneity through structured random effects. The Beta regression framework has been extensively studied over the past two decades. For instance, Simas et al. (2010) extend Beta regression by modelling both the mean parameter and the precision parameter as linear combinations of covariates, linked via appropriate functions. Similarly, da Silva and de Oliveira Lima (2017) incorporate Beta regression into a geographically weighted regression framework, adapting the model to account for spatial variation. Furthermore, Figueroa-Zúñiga et al. (2013) extend Beta regression by embedding it within a Bayesian hierarchical framework, enabling the inclusion of random effects. The Beta regression model we propose builds on the Bayesian hierarchical framework of Figueroa-Zúñiga et al. (2013), while introducing important modifications to the random-effects specification and the variable selection strategy to address the specific challenges of modelling spatio-temporal structures in areal data and selecting from a large set of covariates. Bayesian Beta regression models with spatial dependence have been developed mainly in geostatistical frameworks, where Gaussian spatial random fields enter in the linear predictor to capture distance-based correlation (Lagos-Alvarez et al. 2017; Paradinas et al. 2018). For areal data, spatial dependence is more naturally encoded through neighbourhood structures, as in Cepeda-Cuervo and Núñez-Antón (2013), where a spatial autoregressive Beta regression with contiguity-based lags in both mean and dispersion is used.

To account for spatial and temporal correlation, we incorporate in the linear predictor of the Beta regression model an additive formulation for the random effects. Specifically, we assume an Intrinsic Conditional Auto-Regressive (ICAR) prior for the spatially structured random effect, exploiting its ability to capture structured spatial dependence, which is particularly relevant for regional-level data where geographic proximity strongly influences observed outcomes (Banerjee et al. 2014; Haining and Li 2020). The temporal random effect is modelled as a first-order autoregressive process (AR(1)). Furthermore, in order to capture potential group differences across gender, we include gender-specific intercept terms in the model. In this way, our approach jointly accounts for space, time, and gender influences on regional obesity rates, all of which have been shown to be highly relevant (Grimaccia and Rota 2025). The additive combination of the spatial and temporal random effect offers a parsimonious and interpretable framework, making it a suitable modelling strategy (Banerjee et al. 2014). Although more structured interactions between spatial and temporal components could be incorporated (for a review, see Knorr-Held 2000), we follow the approach of Waller et al. (1997) and Berloco et al. (2023), arguing that prioritizing simplicity and interpretability over added flexible and complexity aligns with the primary objective of this study: to effectively model regional-level patterns in obesity rates while focusing on structured spatio-temporal dynamics.

With the objective of finding a parsimonious representation of the observed data using a statistical model that is also capable of accurate prediction, we adopt the Stochastic Search Variable Selection (SSVS) approach (George and McCulloch 1993, 1997). This method is a Bayesian regularization technique that has proven effective in addressing both prediction and variable selection problems (see Rockova et al. 2012, for a detailed discussion). SSVS is particularly attractive because it is well-

suited for high-dimensional modelling, provides valid standard errors, and enables the simultaneous estimation of regression coefficients and complexity parameters through computationally efficient Markov chain Monte Carlo (MCMC) techniques.

Although the individual components of our proposed approach (Bayesian Beta regression, ICAR spatial effects, AR(1) temporal dynamics, group-specific intercepts and SSVS) are all well established, their integration into a single Bayesian hierarchical framework tailored to regional obesity analysis provides a flexible modeling strategy that, to our knowledge, has not been systematically investigated. In addition, the adoption of Bayesian variable selection within a Beta regression setting addresses an area that has received relatively little attention in the literature.

Our contribution lies in combining these elements into an areal spatio-temporal model that jointly accounts for spatial and temporal heterogeneity, gender-specific differences, and covariate uncertainty, thereby providing a flexible and comprehensive framework for analysing regional obesity dynamics.

Beyond this methodological integration, we construct a harmonized regional panel dataset on obesity and related covariates for Italy through a systematic process of data collection, validation, and reconciliation across multiple institutional sources. This integrated modelling and data effort enables robust inference and improved prediction of obesity prevalence at the regional level, delivering evidence directly relevant for place-based public health policies.

The paper is structured as follows. In Sect. 2, we detail the construction of the dataset, which was assembled from multiple official sources. In Sect. 3 we outline the Bayesian hierarchical spatio-temporal Beta regression model (Sect. 3.1), providing details about prior assumptions (Sect. 3.2) for the structured spatial and temporal random effects, alongside the gender-specific intercepts. Sections 3.3 and 3.4 discuss the SSVS prior specification for the fixed effects and some computation details, respectively. We then present the results of our analysis in Sect. 4 and some final considerations in Sect. 5. Additional material, including descriptive statistics, posterior summaries, predictive comparisons, geographical cross-validation, and robustness diagnostics, is provided in the appendices.

2 Data

The dataset analysed in this study comprises spatio-temporal areal data, which can be conceptualized as a set of time series associated with each spatial unit (i.e., region). The dataset is organized as a longitudinal dataset, collecting information on the 20 Italian regions from 2010 to 2022 (13 years). The Italian regions are divided, according to the NUTS 1 classification² into 5 different geographical zones: North West (NO), North East (NE), Centre (C), South (S) and Islands (SI). In Table 4 (Appendix B) we report the regions and the corresponding geographical zone.

In our analysis, a “weights matrix” W is needed to model spatial dependence through the ICAR prior (see Sect. 3). This matrix encodes the neighborhood relationships between regions, and due to its binary nature, a challenge arises in defining

²NUTS classification: <https://ec.europa.eu/eurostat/web/nuts>.

neighbors for the two main Italian islands, Sicilia and Sardegna. We treat Sicilia as a neighbor of Calabria, reflecting their geographic proximity across the Strait of Messina. In contrast, Sardegna is modeled as isolated, given its substantial spatial separation from all other regions. This specification was selected after comparing the values of Moran's index (see Banerjee et al. 2014) under alternative neighborhood definitions for the two islands (see Grimaccia and Rota 2025)

The dataset contains aggregated data at the regional level (NUTS 2) with variables grouped into four categories: socio-demography, mortality, lifestyle, and general health-related information (see Table 5 in Appendix B). We assembled the dataset by extracting information from several official sources: the Health For All platform (HFA³ the Benessere Equo e Sostenibile (BES⁴ report, Eurostat databases⁵ and the archive of the Ministry of Economy and Finance (MEF)⁶

The dependent variable is the obesity rate measured yearly for each region and considered separately by gender (source: HFA). We display in Fig. 1 the obesity trends by geographic area and gender for the time window we consider. The definition of obesity follows the WHO classification⁷ i.e., having a Body Mass Index (BMI) greater than or equal to 30 defines an obese individual. In the HFA platform, obesity rates are official ISTAT sample-based estimates derived from multipurpose household surveys.

For the mortality rate variables in Group 2 of Table 5, which exhibited high multicollinearity and strong linear relationships, we applied a principal component analysis (PCA). These variables represent cause-specific mortality rates, calculated as the number of deaths due to a given cause per 10,000 individuals within the same age and




Fig. 1 Obesity rate trend from 2010 to 2022 by geographic area and gender

³ HFA: <https://www.istat.it/it/archivio/14562>.

⁴ BES: <https://www.istat.it/en/well-being-and-sustainability/the-measurement-of-well-being/bes-report>.

⁵ Eurostat: <https://ec.europa.eu/eurostat/data/database>.

⁶ MEF archive: www.rgs.mef.gov.it/VERSIONE-I/archivio.

⁷ WHO Obesity definition: <https://www.who.int/news-room/fact-sheets/detail/obesity-and-overweight>.

gender group. Among the nine principal components generated, we retained the first two, which together explained more than 70% of the total variance. The first principal component (PC1) primarily captures diseases related to cardiovascular issues and diabetes, while the second (PC2) reflects factors more associated with stress and mental disorders. Summing up, the final dataset contains $p = 32$ covariates (see Table 5 in Appendix B for more details).

3 Methods

3.1 Beta regression model

Let Y_{its} denote the obesity rate in region i ($i = 1, \dots, 20$), year t ($t = 1, \dots, 13$) and gender s ($s = 1$ for females and $s = 2$ for males). We specify the following Beta regression model (Ferrari and Cribari-Neto 2004):

$$Y_{its} \mid \mu_{its}, \phi \stackrel{\text{ind}}{\sim} \text{Beta}(\mu_{its}\phi, (1 - \mu_{its})\phi)$$

$$\ln \left(\frac{\mu_{its}}{1 - \mu_{its}} \right) = \eta_{its}$$

with the linear predictor η_{its} defined as:

$$\eta_{its} = \beta_0 + \mathbf{x}_{its}^T \boldsymbol{\beta} + \psi_i + \alpha_t + d_s \gamma.$$

With this formulation we have that $\mu_{its} = E(Y_{its}) \in (0, 1)$ and $\text{Var}(Y_{its}) = (\mu_{its}(1 - \mu_{its})) / (1 + \phi)$. It follows that μ_{its} is the mean obesity rate for region i at year t and for gender s , and the parameter ϕ can be interpreted as a precision parameter, i.e., given μ_{its} , the larger the value of ϕ , the smaller the variance of Y_{its} . The structured space random effect is represented by $\boldsymbol{\psi} = (\psi_1, \psi_2, \dots, \psi_{20})^T$, while the structured time random effect by $\boldsymbol{\alpha} = (\alpha_1, \alpha_2, \dots, \alpha_{13})^T$. The global intercept term is β_0 , while the gender-specific intercepts are introduced through the dummy variable d_s , which takes the value 0 for females and 1 for males.

Following Berloco et al. (2023), we adopt an additive specification for the spatial and temporal effects. More details about the associated prior specification are provided in Sect. 3.2. We denote the covariate vector with $\mathbf{x}_{its} = (x_{its1}, x_{its2}, \dots, x_{itsp})^T$ and the corresponding fixed effect coefficients with $\boldsymbol{\beta} = (\beta_1, \beta_2, \dots, \beta_p)^T$.

The Beta parametrization adopted in our model offers distinct advantages. First, by expressing the distribution in terms of its mean μ_{its} and precision ϕ , the model facilitates a more intuitive interpretation. Second, this formulation simplifies the specification of prior distributions within the Bayesian framework. Prior knowledge can be naturally incorporated through μ_{its} via the linear predictor. Simultaneously, the precision parameter ϕ can be assigned a weakly informative prior, allowing the model to flexibly adapt to the data's inherent variability (Ferrari and Cribari-Neto 2004).

3.2 Prior specification

For the spatially structured random effect ψ we assume an ICAR prior, a standard and computationally convenient specification for structured areal effects, as it encourages local smoothing by borrowing strength across neighboring regions (Rue and Held 2005; Banerjee et al. 2014). Specifically, the spatially structured random effect is modelled as follows:

$$\psi \mid \tau_\psi \sim \mathcal{N}(\mathbf{0}, [\tau_\psi(\mathbf{D} - \mathbf{W})]^{-1}),$$

where $\mathbf{W}$ is a binary symmetric 20×20 matrix known as the “weights matrix” or “adjacency matrix” (see, e.g., Banerjee et al. 2014; Haining and Li 2020). The entries of $\mathbf{W}$ are defined as follows: $w_{ij} = 1$ if regions i and j share a common boundary (i.e., are first-order contiguous neighbours) and $w_{ij} = 0$ otherwise. Diagonal elements are set to zero since a region is not considered a neighbour of itself. The diagonal matrix $\mathbf{D}$ is defined as $\mathbf{D} = \text{Diag}(\mathbf{W} \cdot \mathbf{1})$, where Diag represents the diagonal matrix operator, $\mathbf{1}$ denotes a column vector of ones and $\cdot$ denotes the dot product. The diagonal entries D_{ii} represent the number of neighbours for each region i . The scalar τ_ψ is a precision parameter that scales the degree of spatial smoothing in the model: higher values of τ_ψ lead to stronger spatial smoothing, while lower values allow for greater heterogeneity across regions. The full-conditional distribution of the i -th spatial random effects ψ_i is given by:

$$\psi_i \mid \psi_{j \neq i}, \tau_\psi \sim \mathcal{N}\left(\sum_{j \sim i} \frac{\psi_j}{D_{ii}}, \frac{1}{D_{ii}\tau_\psi}\right), \quad (1)$$

where $j \sim i$ indicates that regions i and j are neighbours. As shown in Equation (1), the conditional mean of ψ_i is the average of the values of its neighbours, while its variance decreases as the number of neighbours increases.

For what concerns the structured time random effects, following Knorr-Held (2000) and Berloco et al. (2023), we model the temporal random vector α using a first-order auto-regressive (AR(1)) process with the following structure:

$$\alpha_1 \mid \tau_\alpha \sim \mathcal{N}(0, \tau_\alpha^{-1}),$$

$$\alpha_t \mid \alpha_{t-1}, \rho, \tau_\alpha \sim \mathcal{N}(\rho\alpha_{t-1}, \tau_\alpha^{-1}), \quad t = 2, \dots, 13,$$

where τ_α is the precision parameter and ρ is the auto-regressive coefficient. We assign a $\text{Beta}(a, b)$ prior to ρ , reflecting a prior belief in positive autocorrelation among the temporal random effects while preventing explosive dynamics in the AR(1) process. In our implementation, we set $(a, b) = (1, 1)$, leading to a flat prior avoiding a preference for any particular value of the parameter.

The gender-specific intercepts are modelled in an additive way. First, we recall that the individual ψ_i 's cannot be uniquely determined. This issue is resolved by

imposing a sum to zero constraint $\sum_{i=1}^n \psi_i = 0$ (Thomas et al. 2014) and by modelling the global intercept term β_0 as suggested by Besag and Kooperberg (1995):

$$p(\beta_0) \propto 1, \quad \text{for } \beta_0 \in (-\infty, +\infty),$$

indicating a flat, improper prior for β_0 over the entire real line, a common specification in ICAR models, see Held et al. (2010). A gender-specific shift is introduced through the binary indicator variable d_s . The gender effect coefficient γ allows the definition of gender-specific intercepts as

$$\xi_s = \beta_0 + d_s \gamma,$$

so that $\xi_1 = \beta_0$ for females and $\xi_2 = \beta_0 + \gamma$ for males. We assign to γ a weakly informative normal prior:

$$\gamma \sim \mathcal{N}(0, \sigma_\gamma^2 = 100).$$

For the structured spatial and temporal random effects, we assign a Gamma prior to the respective precision parameters τ_ψ and τ_α . In particular, following the recommendation of Kelsall et al. (1999), for the spatial precision τ_ψ we use a $\text{Gamma}(\alpha_\tau, \beta_\tau)$ prior with $\alpha_\tau = 0.5$ and $\beta_\tau = 0.0005$; this reflects a conservative and weakly informative belief about spatial dependence (Thomas et al. 2014). Although originally proposed for spatial effects, we extended this prior to the temporal precision τ_α , acknowledging the presence of temporal correlation in our setting, favoring shrinkage while preserving flexibility to capture stronger temporal structure when supported by the data.

Finally, we assign a prior to the Beta precision parameter ϕ following the approach proposed by Figueroa-Zúñiga et al. (2013), which extends the prior structure suggested by Gelman (2006) specifically for the Beta distribution, by assuming:

$$\phi = (aB)^2, \quad B \sim \text{Beta}(1 + \epsilon, 1 + \epsilon),$$

where a is a large constant (e.g., 50) and ϵ is a small constant (e.g., 0.1).

To complete the Bayesian hierarchical formulation, we define the prior for the covariate coefficients β using the SSVS approach, as discussed in Sect. 3.3.

3.3 Stochastic Search Variable Selection prior

To model the fixed effects coefficients β_k , $k = 1, \dots, p$, we adopt a SSVS prior specification (George and McCulloch 1993). This approach combines a spike-and-slab mixture prior to probabilistically determine covariate inclusion. In a nutshell, SSVS enables each regressor to be significantly associated with the response (i.e., coefficient different from zero) or shrunk toward zero. This is achieved by assigning a prior that assumes independence between regressors, with marginals defined as a mixture of two Normal distributions, both centered at zero. The spike component corresponds to negligible or irrelevant effects, while the slab component allows for greater vari-

ance and captures substantial or practically meaningful associations. Specifically, we assign the following hierarchical prior structure:

$$\begin{aligned}\beta_k &| \sigma_k^2 \sim \mathcal{N}(0, \sigma_k^2), & k = 1, \dots, p \\ \sigma_k^2 &| \omega_k = (1 - \omega_k)\tau^2 + \omega_k c^2 \tau^2, \\ \omega_k &| \theta_k \sim \text{Bernoulli}(\theta_k), \\ \theta_k &\sim \text{Uniform}(0, 1).\end{aligned}$$

In this formulation, ω_k is a binary latent variable indicating whether the k -th covariate is “practically significant” ($\omega_k = 1$) or not. The prior probability of inclusion of the k -th covariate is denoted by θ_k , to which we assign a Uniform hyper-prior, avoiding imposing strong prior assumptions about the relevance of individual predictors. The parameters τ^2 and $c^2\tau^2$ represent the variances of the spike and slab components, respectively. The intersection points of the spike and slab densities define a threshold ζ for practical significance: $\zeta = \pm\tau\sqrt{2c^2\log(c)/(c^2-1)}$, such that coefficients falling within the interval $[-\zeta, +\zeta]$ are interpreted as practically negligible. Therefore, in accordance with Rockova et al. (2012), we fix τ and c to constant values in order to control the separation between negligible and relevant effects (see Sect. 3.4). Variable selection is then achieved in practice by either identifying the model with the highest posterior probability, including or excluding covariates based on their marginal posterior inclusion probabilities, or evaluating the extent to which the marginal posterior distributions of the regression coefficients are shrunk toward zero (for further details see Rockova et al. 2012).

3.4 Computational details

A summary of all the model parameters is reported in Table 1. To conduct Bayesian inference, we implemented the proposed spatio-temporal hierarchical model using the `Nimble` package (de Valpine et al. 2017, 2024), which enables efficient specification and sampling of complex models. To improve convergence of the MCMC algorithm, we normalized the variables using min-max scaling and standardization, as appropriate.

With regard to the SSVS specification, we set the slab scaling factor to $c = 4000$, and the intersection term to $\zeta = 0.001$, reflecting the threshold of practical signifi-

Table 1 Summary of the model parameters

Parameter	Description
ξ_s	Gender-specific intercept term ($s = 1, 2$)
ψ_i	Structured spatial random effect ($i = 1, \dots, 20$)
τ_ψ	Precision parameter for the spatial random effect
α_t	Structured temporal random effect ($t = 1, \dots, 13$)
τ_α	Precision parameter for the temporal random effect
ρ	AR(1) coefficient for the temporal random effect
ϕ	Beta precision parameter
β_k	Regressor coefficient ($k = 1, \dots, p$)
$\sigma_k, \omega_k, \theta_k$	Parameters from the SSVS prior

cance identified through iterative exploration and consideration of the magnitude of the coefficients (we refer to Appendix E for further details about the choice and sensitivity analysis of SSVS parameters). This resulted in $\tau = 0.00025$ (see Sect. 3.3), a value corresponding to a very small standard deviation for the spike around 0. The slab variance, given by $\tau^2 c^2$, is approximately equal to 1. This configuration aligns the slab variance with the suggested scale of variation for coefficients under the logit link function, as suggested by McElreath (2016).

We ran two parallel chains with overdispersed initial values for a total of 150,000 iterations each. We discarded an initial burn-in period of 50,000 iterations from each chain to ensure convergence to the stationary distribution. To reduce storage requirements and mitigate the effects of residual autocorrelation in the MCMC output, we thinned posterior samples by retaining every 10th iteration. This resulted in a total of 10,000 posterior draws used for inference. The total computation time required to fit the model was approximately 30 min per MCMC chain on a standard laptop.

We assessed convergence and sampling efficiency using multiple diagnostics, including rank-normalized $\hat{R}$ statistics and effective sample size (ESS) measures as proposed by Vehtari et al. (2021), together with the more classic ESS and $\hat{R}$ (Plummer et al. 2006; Gelman and Rubin 1992; Brooks and Gelman 1998). Detailed convergence diagnostics are reported in Appendix F.

4 Results

Table 3 presents the posterior summary statistics for the gender-specific intercepts, the autoregressive parameter ρ and the precision parameter ϕ . Recall that the gender-specific intercepts are defined as $\xi_1 = \beta_0$ for females and $\xi_2 = \beta_0 + \gamma$ for males (see Sect. 3.2). The mean of the intercept for females is -1.09 , indicating a lower baseline effect compared to males, whose corresponding mean is 0.33 . The posterior variability appears to be slightly higher for the male intercept. Figure 2 depicts the posterior density estimates of the gender-specific intercepts. The clear separation between the two distributions highlights the substantial contribution of gender-specific effects in explaining differences in obesity prevalence.

The structured spatial random effects can be visualized through the posterior densities plotted in Fig. 3. The plot shows the posterior densities for all regions, highlighting the contribution of regional random effects to the overall obesity estimates according to the model. Overall, the ICAR prior effectively induces a spatial structure that captures regional dependencies while allowing for evident heterogeneity, as reflected in the varying magnitudes of ψ_i across macro areas and individual regions. To complement the density plots, Fig. 4 provides a map representation of the posterior mean values of the spatial random effects, facilitating the interpretation of their spatial distribution. While spatial proximity generally corresponds to similar values of the spatial random effect, notable exceptions suggest that additional determinants beyond geography influence regional obesity rates. As shown by the posterior summary statistics (Table 6 of Appendix C), regions with negative posterior mean values, such as Piemonte (-0.61), Liguria (-0.62), Trentino-Alto Adige (-0.49), and Sardegna (-0.71 , the lowest), are associated with lower obesity rates. In contrast, with the




Fig. 2 Posterior distributions of the gender-specific intercepts: ξ_1 for females (in pink) and ξ_2 for males (in light blue)

exception of Calabria, all southern regions, particularly Molise (0.89, the highest), as well as Umbria (0.29), exhibit strongly positive spatial effects. The uncertainty in spatial effects varies across regions, as reflected in the standard deviations (SD) of the posterior distributions. As regards the posterior summary statistics for the spatial precision parameter τ_ψ , we refer the reader to Appendix C.

Posterior density estimates of the time-specific random effects are shown in Fig. 5. A clear temporal pattern emerges: during the first half of the observational window (2010–2015), the random effects remain relatively stable, centered around values slightly below zero. In contrast, the second half of the period exhibits a marked upward shift in the magnitude of the effects. Between 2010 and 2015, the posterior probability that the temporal effects exceed zero remains relatively stable, fluctuating around 0.3, with the exception of 2014, which shows a modest increase. Starting in 2016, these probabilities rise significantly: both 2016 and 2017 reach approximately 0.76, and from 2018 onward, the vast majority (if not all) of posterior samples are positive. These trends are summarized in Table 2.

As confirmed by the posterior summary statistics shown in Table 7 of Appendix C, the results suggest a varying influence of the time random effects across differ-




Fig. 3 Posterior density of the space random effects ψ_i , coloured by geographical zone (NO=North West, NE=North East, C=Centre, S=South, SI=Islands, see Table 4)

ent time points. While some periods exhibit negative random effects, indicating a lower-than-expected influence on obesity rates, others show positive effects, suggesting periods of higher-than-expected influence. While the magnitude of time random effects differs across time points, the variability within each time point is relatively homogeneous.

The autoregressive parameter ρ exhibits a negatively skewed posterior distribution, with a posterior mean of 0.85 (Table 3), indicating strong temporal persistence in regional obesity rates. This suggests that obesity levels in a given year are highly correlated with those in previous years, resulting in a smooth temporal evolution rather than abrupt year-to-year changes. Such persistence is coherent with the influence of slowly evolving underlying factors, rather than short-term shocks. Although broader socioeconomic, behavioral, or epidemiological processes may have contributed to this upward trend, the aggregated nature of the data and the lack of information on specific policies or interventions prevent a more detailed analysis. We therefore interpret the estimated value of ρ primarily as capturing long-run temporal dependence in obesity dynamics, rather than effects associated with specific policy




Fig. 4 Posterior mean values of the spatial random effects ψ_i for each Italian region

actions or events. As regards the posterior summary statistics for the temporal precision parameter τ_α , we refer the reader to Table 8 in Appendix C.

The posterior summary statistics of the Beta precision parameter ϕ (Table 3) are not straightforward to interpret, primarily because they depend on the scale of the data. The parameter becomes more interesting when compared between models: a model with a higher posterior precision parameter suggests a better ability to explain the variability of the Beta-distributed response variable, due to the relation that links the variance to the precision ϕ (see Sect. 3). We just highlight that the Bayesian updating process was very pronounced, as the prior distribution, which spans a wide range (between 0 and 2500), is effectively updated to a more concentrated range around the mean of 37.67. The reduction in uncertainty is very high, as reflected by the low coefficient of variation equal to 0.06.

As discussed in Sect. 3.3, we adopt the SSVS approach to perform inference on the fixed-effect coefficients β_k . Consequently, the evaluation of their posterior distribution is intrinsically linked to the variable selection induced by the spike and slab




Fig. 5 Posterior density of the time random effects α_t , coloured by year

Table 2 Posterior probability that the temporal effect α_t exceeds zero (PP0)

t	1	2	3	4	5	6	7
Year	2010	2011	2012	2013	2014	2015	2016
PP0	0.28	0.25	0.25	0.38	0.61	0.32	0.76
t	8	9	10	11	12	13	
Year	2017	2018	2019	2020	2021	2022	
PP0	0.76	0.96	0.98	1.00	1.00	0.99	

Table 3 Posterior summary statistics for the gender-specific intercepts, the temporal dependence parameter ρ , and the Beta precision parameter ϕ

Parameter	Mean	SD	CV	Skewness	Kurtosis
ξ_1 (Females)	-1.09	0.16	-0.15	-0.22	3.48
ξ_2 (Males)	0.33	0.18	0.53	-0.11	3.03
ρ	0.85	0.14	0.16	-1.48	5.62
ϕ	37.67	2.39	0.06	0.14	3.00

prior. The SSVS hierarchical prior specification can make inference on the β_k coefficients challenging, particularly when their posterior distributions are bi-modal, with samples alternating between values near zero and substantially different from zero. This complicates the inferential process. However, we point out that the primary objective of this work is to investigate the spatio-temporal dynamics of obesity rates across Italian regions, rather than exploring or assessing causal relationships between the covariates and the response variable. The inclusion of these covariates is primarily to control for any additional effect that may arise along with the random effects, which are of greater interest in this study. Thus we do not conduct detailed posterior evaluation of the β_k coefficients. Nevertheless, given the valuable insights offered by the SSVS approach, we perform a brief analysis of the inclusion probabilities to assess which covariates are associated with obesity, after accounting for spatial and temporal random effects, as well as gender. Specifically, the ω_k parameter plays a central role in determining whether the k -th covariate should be included in the model, based on the probability parameter θ_k (see Sect. 3.3). By counting the proportion of MCMC iterations in which $\omega_k = 1$, we obtain an empirical inclusion probability for each covariate. Figure 6 presents these posterior inclusion probabilities. The only covariates that are consistently included in the model are “life expectancy”,




Fig. 6 Posterior probability evaluation of the inclusion parameters ω_k for all the covariates (for the description of the covariates see Table 5 in Appendix B)

“life expectancy in good health” and “overweight” (see Table 5). All the remaining covariates exhibit very low posterior inclusion probabilities. This suggests that once the random effects and gender-specific intercepts are incorporated, only three covariates need to be included according to our model specification. The other covariates do not prove to be related to obesity rates and only add unnecessary complexity to the model. We observe that this is consistent with the structure of the model. In particular, the spatial and temporal random effects act as latent proxies for a wide range of unobserved regional characteristics, such as cultural habits, health system features, environmental conditions, socioeconomic context, or persistent lifestyle patterns. As a consequence, much of the cross-regional variability that might otherwise be captured by observable covariates is absorbed by the spatial and temporal terms, reducing the marginal contribution of many individual predictors. This highlights that, at the regional scale, unobserved or hard-to-measure contextual factors play a dominant role relative to the socioeconomic covariates available in the dataset.

To assess the sensitivity of the SSVS procedure to hyperparameter choices, we conducted a sensitivity analysis to determine whether different SSVS settings yield distinct shrinkage patterns. The full analysis is reported in Appendix E. Briefly, we tested four alternative configurations based on different intersections between the spike and slab densities (see Sect. 3.3). While posterior inclusion probabilities vary slightly across specifications, the same three covariates consistently emerge as the only ones not shrunk toward zero. All the remaining covariates have posterior inclusion probabilities below the commonly used 0.5 threshold, with a minor exception in one specification (see Appendix E for further details).

4.1 Predictive performance

In order to assess the performance of the modelling approach we propose, we considered three alternative model specifications alongside the main model presented in Sect. 3. These variants differ in two aspects: (i) the treatment of the gender effect, modelled either as a fixed or a random effect; and (ii) the prior distribution on the auto-regressive coefficient ρ , specified either as a Beta(1, 1) distribution (constraining ρ to the unit interval) or as a standard Normal distribution (allowing unrestricted support). This results in four model configurations, named M1, M2, M3 and M4, as detailed in Table 9 in Appendix D.1. Note that M1 (i.e., gender as fixed effect and Beta prior for ρ) corresponds to our main model described in Sect. 3. To evaluate and compare the predictive performance of the four models, we implemented both temporal and spatial validation strategies. First, we conducted a temporal train–test split, using data from 2022 ($t = 13$) as the test set and data from the preceding years ($t = 1, \dots, 12$) for model training (see Appendix D.1). Second, we performed a spatial cross-validation procedure by leaving out one representative region from each macro-area: Lombardia (North), Toscana (Centre), and Campania (South). Further details are provided in Appendix D.2. Models were assessed using standard predictive metrics. For what concerns the temporal test, the results reported in Table 10 indicate that M1 outperforms the competing models in terms of predictive accuracy. The latter is likely due to the temporal dynamics allowed by the Normal prior, which can lead to explosive behaviours and overfitting. The random gender effect offered

no improvement over the fixed-effect specifications. Overall, M1 offered the best balance between parsimony and predictive accuracy. For what concerns the spatial cross-validation, we obtained smaller performance differences across model specifications than those observed in the temporal validation exercise. This indicates that, once trained on a sufficiently broad set of regions, all four models are able to capture the main spatial patterns in obesity prevalence with comparable accuracy.

5 Concluding remarks

This study proposes a Bayesian hierarchical Beta regression framework to analyse the spatio-temporal dynamics of obesity rates across Italian regions from 2010 to 2022. The results highlight pronounced gender differences, regional heterogeneity and dependence, and clear temporal trends, suggesting the importance of spatial, temporal, and demographic factors in explaining obesity patterns.

The choice of the Beta distribution is motivated not only by its support on the unit interval - making it particularly suitable for modelling rates and proportions - but also by its flexibility, as it accommodates a wide range of distributional shapes, depending on the parameter values (Johnson et al. 1995). We note that, beyond standard Beta regression, several alternative approaches have been proposed for modeling bounded outcomes. Within the Beta regression framework, more flexible specifications have been introduced, such as the flexible Beta distribution of Migliorati et al. (2018). In addition, compositional data methods allow the joint modeling of related categories, such as underweight, overweight, and obesity prevalence, through log-ratio transformations (Mills 2008) or more flexible formulations based on structured mixtures with Dirichlet components (Ascari et al. 2024). Transformation-based approaches, such as logit-normal regression, provide another alternative by modeling a suitable transformation of the outcome on an approximately Gaussian scale (see, e.g., Lesaffre et al. 2007). While these approaches may be particularly useful for capturing complex distributional features or jointly modeling multiple dependent components, our objective is to analyze a single obesity prevalence rate observed over space and time. In this setting, Beta regression provides a natural and readily interpretable modeling framework.

In the present case study, the proposed Bayesian hierarchical Beta regression model incorporates structured spatial and temporal random effects additively, along with gender-specific intercepts. This model proves effective in capturing regional variations and temporal dependencies in obesity rates. While the individual components of the proposed model are well established, their integration into a unified Bayesian hierarchical framework constitutes the main methodological contribution of this work, offering a coherent and flexible approach for analysing regional obesity rates and their spatio-temporal dynamics.

Some methodological limitations warrant consideration. The ICAR prior for the structured spatial random effects assumes complete spatial dependence, which may not fully reflect the underlying spatial processes. Additionally, the additive specification of spatial and temporal random effects, with separate priors, may overlook potential space-time interactions. The exclusion of unstructured spatio-temporal vari-

ability, which could increase flexibility in capturing localized deviations, represents a further limitation. Furthermore, although the regional-level analysis improves measurement reliability and mitigates data gaps, it may obscure finer-scale variations at the provincial level. A further limitation concerns the use of self-reported data, which are known to be affected by systematic reporting bias. In particular, individuals tend to under-report body weight and over-report height, leading to a downward bias in BMI and, consequently, in estimated obesity prevalence. This bias may vary across demographic groups (e.g., by gender, age, or socioeconomic status) and across regions, potentially introducing differential measurement error that is not explicitly accounted for in the model (Gorber et al. 2007). Moreover, the adjacency matrix assigns equal weight only to immediately neighbouring regions and therefore does not account for distance decay or other forms of connectivity, such as economic or cultural links. While this is a standard and interpretable choice for areal data, future work could explore more flexible weighting schemes to better capture complex spatial relationships. Finally, the assumption of a constant precision parameter ϕ is adopted as a parsimonious modelling choice that preserves interpretability and avoids potential over-parameterization. In our setting, the primary inferential interest lies in the fixed and random effects governing the mean structure rather than in modelling the dispersion parameter itself. Moreover, dispersion effects are typically more difficult to specify and estimate reliably than mean effects, as discussed by Cribari-Neto and Souza (2012).

Despite these limitations, the proposed framework remains parsimonious, interpretable, and effective in predictive terms. Its main strength lies in balancing modelling flexibility with methodological clarity, thereby providing a practically useful approach for the spatio-temporal analysis of regional obesity rates in public health applications.

Our findings highlight the substantial role of the gender-specific random effects in estimating obesity outcomes, underscoring gender as a critical determinant of regional obesity rates. Spatial random effects reveal homogeneous patterns across the North-West, with the exception of Emilia-Romagna, which exhibits higher values. In the North-East, Friuli-Venezia Giulia and Veneto show similar values, while Trentino-Alto Adige displays lower magnitudes with a greater range. The Centre geographical zone is characterized by marked heterogeneity, with Lazio showing significantly lower values and Umbria showing higher ones. Southern regions generally exhibit higher spatial random effects, except for Calabria. The Islands show distinct spatial structures, an aspect enforced by the fact that Sicilia is treated as a neighbour to Calabria in the ICAR prior, while Sardegna is considered spatially isolated. Temporal random effects indicate a stable trend in the early years, followed by an increase in later years, suggesting a growing upward trend in obesity rates. Under the SSVS prior specification, most coefficients are shrunk toward zero, with only a few covariates being identified as having a relevant association with obesity, namely “life expectancy”, “life expectancy in good health”, and “overweight”. This result suggests that spatial and temporal random effects explain a substantial portion of the variability in obesity rates. Clearly, additional factors, such as environmental characteristics, food environment indicators, or policy-related variables, may play a relevant role but could not be considered due to the limited availability of regional-level

data over the study period. The integration of richer covariate information, therefore, represents a natural and promising direction for future research as more comprehensive data become accessible.

While this study does not aim to establish causal relationships, it highlights the importance of modelling obesity trends through demographic and spatio-temporal structures rather than relying solely on fixed socioeconomic or lifestyle predictors. The results further support the value of incorporating spatio-temporal dependence in public health research.

Bayesian variable selection in Beta regression models has received limited attention in the literature; a more in-depth methodological investigation in this direction is part of our ongoing research. Future research could also extend this framework by considering finer spatial resolutions, such as provinces or municipalities, incorporating individual-level or environmental information, and allowing for spatio-temporal interactions. Alternative prior specifications could also be explored to increase flexibility in the presence of outliers or asymmetric variation. More broadly, the proposed Bayesian hierarchical Beta regression model may be useful for other areal outcomes bounded between 0 and 1 and characterized by spatial and temporal dependence, including applications in public health and the social sciences.

Appendix organization

We provide additional material in the appendices: Appendix B contains a table reporting the classification of Italy into geographical zones and a table describing the covariates used in the analysis; Appendix C reports additional posterior summary statistics not included in the main text for brevity, Appendix D presents the results of the predictive comparison of some variants of the main model, Appendix E includes sensitivity analyses for different SSVS hyperparameter specifications and, finally, Appendix F provides extensive diagnostics and performance metrics.

Geographical areas and covariates

This appendix provides the description of the Italian geographical zones and of the covariates used in the analysis (see Tables 4 and 5).

Table 4 Italian regions grouped by geographical zone. Source: Geographical classification of Italian regions according to the Italian National Institute of Statistics (ISTAT)

Zone	Regions
North West (NO)	Piemonte, Valle d'Aosta, Lombardia, Liguria, Emilia-Romagna
North East (NE)	Trentino-Alto Adige, Veneto, Friuli-Venezia Giulia
Centre (C)	Toscana, Umbria, Marche, Lazio
South (S)	Abruzzo, Molise, Campania, Puglia, Basilicata, Calabria
Islands (SI)	Sicilia, Sardegna

Table 5 Covariates used in the analysis, grouped into four thematic categories: Socio-Demography, Mortality, Lifestyle, and General Wealth-related Information

Variable	Unit of measure	By gender	Source
<i>Group 1—Socio-Demography</i>			
Population	count	Yes	Eurostat
More than 65 years old	Percentage	Yes	HFA
Foreigners	Percentage	Yes	HFA
Foreigners from EU	Percentage	Yes	HFA
More than 65 y/o living alone	Percentage	Yes	HFA
Mean family components	average	no	HFA
Unemployment	Percentage	Yes	Eurostat
Gross Income	euro (thousands)	no	Eurostat
Wedding rates	Percentage	no	HFA
University (from Bachelor's to PhD)	Percentage	Yes	Eurostat
<i>Group 2—Mortality</i>			
Cancer of digestive system	Percentage	Yes	HFA
Cancer of stomach	Percentage	Yes	HFA
Diabetes	Percentage	Yes	HFA
Mental disorders	Percentage	Yes	HFA
Blood diseases	Percentage	Yes	HFA
Heart diseases	Percentage	Yes	HFA
Digestive diseases	Percentage	Yes	HFA
Liver diseases	Percentage	Yes	HFA
Suicide	Percentage	Yes	HFA
<i>Group 3—Lifestyle</i>			
Overweight	Percentage	Yes	HFA
Overweight of younger	Percentage	Yes	HFA
Cigarette consumption	count	Yes	HFA
Smoking rate	Percentage	Yes	BES
Complete breakfast	Percentage	Yes	HFA
Daily cheese consumption	Percentage	Yes	HFA
Red meat consumption	Percentage	Yes	HFA
Fish consumption	Percentage	Yes	HFA
Daily vegetables consumption	Percentage	Yes	HFA
Dinner as principal meal	Percentage	Yes	HFA
No sport participation	Percentage	Yes	HFA
Adequate nutrition	Percentage	Yes	BES
Alcohol consumption	Percentage	Yes	BES
<i>Group 4—General Wealth-related Information</i>			
Bad wealth	Percentage	Yes	HFA
Life expectancy (LE)	Years	Yes	HFA
LE in good wealth	Years	Yes	HFA
LE without limitations	Years	Yes	HFA
Drug consumption	Percentage	Yes	HFA
Life satisfaction	Percentage	Yes	BES
Health expenditure per capita	Euro/population	Yes	MEF

The table reports the unit of measurement, availability by gender, and data source for each variable. Source: HFA = Health for All database; BES = Italian Equitable and Sustainable Well-being report; MEF = Italian Ministry of Economy and Finance; Eurostat = Statistical Office of the European Union

Posterior summary statistics

This appendix provides the parameter posterior summary statistics not reported in the main text (see Tables 6 and 7).

Table 6 Posterior summary statistics for region-specific parameters ψ_1 to ψ_{20}

Parameter (Region)	Mean	SD	CV	Skewness	Kurtosis
ψ_1 (Piemonte)	-0.61	0.07	-0.12	0.18	3.46
ψ_2 (Valle d' Aosta)	-0.48	0.09	-0.18	0.34	3.96
ψ_3 (Lombardia)	-0.41	0.08	-0.19	0.25	3.41
ψ_4 (Liguria)	-0.62	0.08	-0.12	0.11	3.42
ψ_5 (Emilia-Romagna)	0.25	0.08	0.33	0.32	3.59
ψ_6 (Trentino-Alto Adige)	-0.49	0.12	-0.25	0.02	3.19
ψ_7 (Veneto)	-0.02	0.07	-3.34	0.06	3.13
ψ_8 (Friuli-Venezia Giulia)	0.13	0.07	0.55	0.03	3.08
ψ_9 (Toscana)	-0.15	0.08	-0.56	0.18	3.30
ψ_{10} (Umbria)	0.29	0.07	0.26	0.19	3.29
ψ_{11} (Marche)	0.03	0.07	2.73	0.03	3.09
ψ_{12} (Lazio)	-0.40	0.07	-0.18	0.20	3.55
ψ_{13} (Abruzzo)	0.49	0.06	0.13	-0.03	3.05
ψ_{14} (Molise)	0.89	0.08	0.09	-0.25	3.57
ψ_{15} (Campania)	0.35	0.11	0.32	-0.31	3.65
ψ_{16} (Puglia)	0.65	0.08	0.12	-0.55	4.34
ψ_{17} (Basilicata)	0.70	0.09	0.13	-0.26	3.58
ψ_{18} (Calabria)	-0.07	0.12	-1.69	-0.13	3.12
ψ_{19} (Sicilia)	0.16	0.10	0.60	-0.22	3.60
ψ_{20} (Sardegna)	-0.71	0.09	-0.12	-0.32	3.79

Table 7 Posterior summary statistics for time-specific parameters α_1 to α_{13}

Parameter (Year)	Mean	SD	CV	Skewness	Kurtosis
α_1 (2010)	-0.07	0.12	-1.83	-0.02	3.79
α_2 (2011)	-0.08	0.13	-1.65	0.11	3.64
α_3 (2012)	-0.08	0.14	-1.63	0.29	3.54
α_4 (2013)	-0.04	0.14	-3.51	0.26	3.36
α_5 (2014)	0.04	0.14	3.13	0.37	3.44
α_6 (2015)	-0.06	0.14	-2.24	0.01	3.31
α_7 (2016)	0.10	0.14	1.38	0.36	3.31
α_8 (2017)	0.11	0.14	1.36	0.23	3.09
α_9 (2018)	0.23	0.14	0.62	0.32	3.25
α_{10} (2019)	0.28	0.15	0.52	0.26	3.14
α_{11} (2020)	0.40	0.15	0.37	0.36	3.41
α_{12} (2021)	0.41	0.14	0.35	0.40	3.48
α_{13} (2022)	0.30	0.14	0.48	0.31	3.27

Table 8 Posterior summary statistics for the hyperparameters of space (τ_ψ), and time (τ_α) random effects

Parameter	Mean	SD	CV	Skewness	Kurtosis
τ_ψ (Space precision)	2.37	0.82	0.34	0.80	4.14
τ_α (Time precision)	100.84	65.26	0.65	2.19	12.08

Focusing on the precision parameters for the spatial and temporal random effects (τ_ψ and τ_α , respectively), we observe substantial Bayesian updating, particularly for the temporal component (see Table 8). Beginning with conservative Gamma priors, the posterior distribution of the spatial precision τ_ψ exhibits only a moderate increase. This restrained growth can be partly attributed to the ICAR prior structure, in which the full-conditional precision is scaled by the number of neighboring regions. In contrast, the temporal precision τ_α , whose prior specification does not involve such a multiplicative adjustment, undergoes a more pronounced shift, with a posterior mean markedly higher than that of τ_ψ . However, direct comparison between these two parameters should be approached with caution, as their prior formulations differ substantially, rendering such comparisons potentially misleading. Both posterior distributions are characterized by high skewness and kurtosis, suggesting strong positive asymmetry and heavy tails.

Predictive performance comparison

Predictive test for year 2022

In this appendix we evaluate the predictive performance of four model specifications using a temporal training-test split. Specifically, the training set comprises data from year 1 to 12, while the test set consists of observations from 2022 ($t = 13$), across all regions. As described in Sect. 4, we compare four Bayesian Beta regression models which differ in two aspects: (i) the specification of the gender effect (fixed vs. random effect), and (ii) the prior distribution assigned to the auto-regressive coefficient ρ . The resulting four model variants (named M1, M2, M3 and M4) are summarized in Table 9, with M1 corresponding to our main model described in Sect. 3.

With respect to gender, the random effect specification models the gender-specific intercept as

$$\xi_s = \beta_0 + \gamma_s, \quad s = 1, 2,$$

where β_0 is the global intercept and γ_s represents a random deviation for each gender. The random effects follow a classical hierarchical formulation (e.g., Hoff 2009; Gelman 2006):

$$\gamma_s \mid \mu_s, \sigma_s^2 \sim \mathcal{N}(\mu_s, \sigma_s^2), \quad \mu_s \sim \mathcal{N}(0, 4), \quad \sigma_s^2 \sim U(1, 10).$$

In contrast, the gender fixed effect specification is the one we describe in Sect. 3.2.

Table 9 Summary of model variants

Model	Gender effect	Prior on ρ
M1	Fixed	$Beta(1, 1)$
M2	Fixed	$\mathcal{N}(0, 1)$
M3	Random	$Beta(1, 1)$
M4	Random	$\mathcal{N}(0, 1)$

Table 10 Predictive performance metrics for the four model variants, evaluated on 2022 test data

Model	RMSE	RMSE (F)	RMSE (M)	MAE	MBPV	SLPD
M1	0.010	0.010	0.010	0.008	0.537	43.72
M2	0.013	0.012	0.014	0.010	0.717	33.17
M3	0.010	0.009	0.010	0.008	0.571	44.62
M4	0.014	0.013	0.015	0.012	0.757	27.78

RMSE (F) and RMSE (M) refer to female and male subgroups, respectively

As for the autoregressive coefficient ρ , we consider two prior formulations: a $Beta(1, 1)$ prior (constraining ρ to the interval $(0, 1)$) as specified in Sect. 3.2, and a standard normal prior $\mathcal{N}(0, 1)$, which places no support restrictions and permits values outside the unit interval (Table 9).

To evaluate the predictive performance, we compute standard error-based metrics for the test set: root mean squared error (RMSE), mean absolute error (MAE), and the sum of the log predictive densities (SLPD). Additionally, we include a measure we refer to as the Bayesian p -value (BPV), defined for each observation as the proportion of posterior predictive samples exceeding the observed value. Values of BPV near 0.5 indicate that the observation lies near the centre of the predicted distribution, suggesting a good fit. In contrast, values close to 0 or 1 suggest under or over prediction, respectively. We compute a mean BPV (MBPV) across all the test observations as a summary.

Table 10 reports the predictive performance metrics for each model, evaluated on the 2022 test set. The results suggest that the prior specification for ρ has a greater impact on predictive performance than the treatment of the gender effect. Specifically, models using a Normal prior for ρ (M2, M4) tend to overestimate 2022 obesity rates. This is likely because the unbounded support of the Normal distribution allows ρ to take values greater than one, permitting overly persistent or explosive temporal dynamics and leading to overfitting of upward trends in the training data. In contrast, the $Beta(1, 1)$ prior restricts ρ to the $(0, 1)$ interval, constraining the model to more realistic levels of temporal autocorrelation. Models M1 and M3, which use this prior, yield more stable predictions and improved out-of-sample accuracy. Between M1 and M3, predictive performances are nearly identical. We select M1 for reporting in the main paper, as we think modeling gender as a random effect with only two levels (male and female) is not necessary. A fixed effect specification is more parsimonious in this case.

In Fig. 7a–d we provide for each of the four models a visual representation of the posterior predictive accuracy by representing the predicted obesity rates against the observed ones, together with 95% credible intervals, separately by gender. This helps to assess the uncertainty in the predictions for the four specifications. Overall, Models M2 and M4 display a tendency toward overestimation, particularly for males, as

evidenced by the majority of point predictions lying above the 45-degree reference line, together with upward biased intervals. By contrast, M1 and M3 appear more closely aligned with the observed values. In terms of uncertainty, however, the width of the credible intervals is broadly comparable across models, suggesting that differences in performance are mainly driven by bias rather than by substantial changes in predictive variability.

Spatial cross-validation

To further assess the robustness of the proposed modeling strategies and to compare the four model specifications (M1 to M4), we conducted a spatial cross-validation analysis. Unlike the temporal training-test split described in Sect. D.1, this validation scheme evaluates the ability of the models to extrapolate to previously unseen geographical areas. Specifically, we adopt a leave-one-region-out strategy by selecting three representative Italian regions: Lombardia (North), Toscana (Centre), and Campania (South). For each region in turn, all observations corresponding to that region are excluded from the training set, and models M1 to M4 are estimated using data from the remaining regions. Posterior predictive distributions are then generated for the held-out region across all available years. Predictive accuracy is assessed using the same set of metrics employed in the temporal validation, namely RMSE, MAE, MBPV and SLPD. RMSE is additionally reported separately for female and male subgroups in order to evaluate potential gender-specific discrepancies in predictive performance.




Fig. 7 Predicted versus observed obesity rates for 2022 by gender for the four model specifications reported in Table 9. Points represent posterior means, while vertical bars indicate the 95% credible intervals

Table 11 Predictive performance metrics by region (all years)

Held-out region	Model	RMSE	RMSE (F)	RMSE (M)	MAE	MBPV	SLPD
<i>Lombardia</i>							
	M1	0.0086	0.0098	0.0072	0.0070	0.654	28.8
	M2	0.0086	0.0097	0.0072	0.0068	0.651	28.7
	M3	0.0089	0.0100	0.0076	0.0074	0.660	28.3
	M4	0.0089	0.0101	0.0075	0.0074	0.660	28.3
<i>Toscana</i>							
	M1	0.0065	0.0053	0.0075	0.0054	0.565	32.7
	M2	0.0066	0.0052	0.0077	0.0054	0.571	32.5
	M3	0.0060	0.0051	0.0068	0.0051	0.544	33.3
	M4	0.0063	0.0051	0.0072	0.0053	0.557	33.0
<i>Campania</i>							
	M1	0.0094	0.0102	0.0086	0.0082	0.576	27.1
	M2	0.0096	0.0103	0.0089	0.0082	0.591	27.0
	M3	0.0094	0.0104	0.0083	0.0082	0.577	26.5
	M4	0.0096	0.0103	0.0090	0.0083	0.592	26.7

RMSE (F) and RMSE (M) refer to female and male subgroups, respectively

Table 11 reports the resulting predictive performance metrics for each model-region combination.

Across the three regions, predictive performance is generally higher for Toscana, followed by Lombardia, while Campania exhibits larger prediction errors. This pattern is consistent across model specifications and likely reflects a higher degree of structural heterogeneity or stronger region-specific dynamics in the southern region. In terms of model comparison, differences between specifications remain modest and do not point to a clear overall winner. While minor variations can be observed, with models M1 and M2 performing slightly better for Lombardia, M3 and M4 for Toscana and M1 slightly outperforming others for Campania, these differences are quantitatively small. This suggests that, when evaluating predictive performance across all years, neither the prior specification for the autoregressive coefficient ρ nor the treatment of the gender effect substantially affects out-of-sample accuracy. Consequently, modeling ρ using either a Beta or a Normal prior leads to no systematic predictive advantage in this setting, and treating gender as a random effect does not provide clear benefits given the limited number of gender categories. In light of these findings, we favor Model M1, as it offers the most parsimonious specification without sacrificing predictive performance and remains consistent with the results of the temporal 2022 validation discussed in Sect. D.1.

We present in Fig. 8 the posterior predicted values against the observed data for all held-out regions and models. This visualization makes it possible to assess the uncertainty associated with the predictions across both genders and regions. Overall, the observed values consistently fall within the corresponding credible intervals. For Toscana (middle row of the plot), the observations are closely aligned with the posterior means, particularly for females, while males exhibit slightly greater predictive uncertainty. For Lombardia (upper row), the observed values are near the predictions but display a minor tendency toward overestimation, in agreement with the MBPV analysis. In Campania (lower row), predictions show higher uncertainty and

occasional discrepancies. Notably, for 2020 females across all models, the observed values approach the upper boundary of the 95% credible interval, suggesting slight underestimation; conversely, other years and genders tend to exhibit mild overestimation, again accordingly to MBPV. In general, predictions remain similar and reliable across regions, and differences between the four model specifications are minimal.

Sensitivity analysis for the SSVS prior specification

We conducted a systematic robustness analysis of the SSVS prior. We first fixed the slab variance multiplier at $c^2 = 4000$. We then varied the spike variance τ^2 such that the intersection between the spike and slab densities (i.e., ζ) took four distinct values: 0.0001, 0.001 (the one we used in the main model (see Sect. 3.3), 0.01, and 0.1: smaller intersections correspond to tighter spike. The resulting standard deviations for the two components are reported in Table 12.

For each specification, we re-estimated the model presented in Sect. 3.1 (denoted as M1). For what concerns the SSVS indicators ω_k and their associated inclusion probabilities θ_k , we set all initial values equal to 1, i.e. starting the search from the full model. We let the sampler run for 20,000 iterations, with a burn-in of 5,000 and no thinning. We then computed the posterior inclusion probabilities ω_k for each covariate, which are displayed in Fig. 9a–d.

Across all prior configurations, the same three covariates consistently exhibit the highest posterior inclusion probabilities: life expectancy, life expectancy in good health and overweight prevalence. This concordance across substantially different shrinkage regimes indicates that these covariates capture robust empirical signal rather than reflecting a particular prior choice. Nonetheless, the absolute magnitude of inclusion probabilities varies with the strength of the prior separation. Under the strongest spike concentration ($\zeta = 0.0001$), overall inclusion probabilities increase,




Fig. 8 Predicted versus observed obesity rates across all regions (rows) and models (columns), stratified by gender (columns). Shaded areas represent the 95% credible intervals, the black solid line denotes the posterior mean, and the red points indicate the observed values

and several additional covariates attain intermediate probability levels (still below 0.5), with the only exception of the covariate “dinner as principal meal” reaching a posterior inclusion probability slightly above 0.7. The two intermediate separation settings (ζ equal to 0.001 and 0.01) produce almost identical results, both characterised by substantial shrinkage of most coefficients. The weakest separation ($\zeta = 0.1$) yields the lowest inclusion probabilities across all covariates: the same three covariates identified in the main analysis remain the most influential, while all others are effectively excluded, in line with the findings of Rockova et al. (2012).

These results suggest that although the range of the inclusion probabilities is somehow affected by the strength of prior shrinkage, the variable ranking is highly robust with respect to the SSVS hyperparameter choice.

Model diagnostics

For the model presented in Sect. 3.1, we conducted a comprehensive set of diagnostic checks to assess MCMC convergence, sampling efficiency and model adequacy. First, we evaluated convergence and effective sample size using the diagnostic frameworks proposed by Plummer et al. (2006) and Vehtari et al. (2021). Second, we examined residual patterns to detect potential model misspecification. Finally, we performed posterior predictive checks to assess the model’s ability to reproduce key features of the observed data.

Convergence diagnostics

We assessed MCMC convergence and sampling efficiency using both classical and modern diagnostic tools. Specifically, we report the traditional potential scale reduction factor $\hat{R}$ and the associated effective sample size (ESS) as implemented in the `coda` package (Plummer et al. 2006), based on the original proposals of Gelman and Rubin (1992) and Brooks and Gelman (1998). We also use the rank-normalized split- $\hat{R}$ together with the basic, bulk, and tail ESS introduced by Vehtari et al. (2021) and implemented in the `posterior` package.

We report the resulting convergence diagnostics in Table 13 for the intercept β_0 and the male-specific intercept component γ , the spatial and temporal random effects ψ and α , the autoregressive coefficient ρ , the precision parameter ϕ and the random effect precision hyperparameters τ_ψ and τ_α . SSVS parameters and fixed effect coefficients (i.e., β) are excluded from the analysis because their SSVS hierarchical specification makes more difficult the analysis of the two chains.

Table 12 Spike and slab standard deviations for the four SSVS prior specifications according to the intersection value ζ

Intersection ζ	0.0001	0.001	0.01	0.1
Spike SD	0.00002	0.00024	0.00245	0.02455
Slab SD	0.098	0.982	9.821	98.21




Fig. 9 Posterior probability evaluation of the inclusion parameters ω_k for all the covariates according to the four considered values of the intersection parameter ζ

Overall, effective sample sizes are generally large under both classical and more conservative diagnostics (Vehtari et al. 2021), with only a few spatial and late-period temporal random effects showing moderate reductions. Convergence diagnostics based on $\hat{R}$ (including rank-normalized split- $\hat{R}$) are uniformly close to one, providing no evidence of lack of convergence; taken together, these results indicate satisfactory mixing and support reliable posterior inference.

Residual analysis

We conducted a comprehensive residual analysis to assess the adequacy of the model and to identify potential systematic patterns not captured by the specified model structure. We define residuals as the difference between the observed outcomes and the posterior mean of the predictive distribution (as given by M1).

Table 13 Effective sample size (ESS) and convergence diagnostics for model parameters

Parameter	ESS (basic)	ESS (bulk)	ESS (tail)	rn-split $\hat{R}$	ESS	$\hat{R}$	95% UCB $\hat{R}$
β_0	110	112	189	1.01	122	1.01	1.02
γ	169	169	424	1.01	167	1.01	1.04
ψ_1	121	125	242	1.01	813	1.01	1.03
ψ_2	76	95	49	1.01	286	1.03	1.07
ψ_3	126	126	804	1.00	273	1.01	1.03
ψ_4	284	289	975	1.00	1120	1.00	1.00
ψ_5	151	152	536	1.00	361	1.01	1.01
ψ_6	190	192	330	1.00	1068	1.00	1.00
ψ_7	386	385	1761	1.00	780	1.00	1.01
ψ_8	142	143	528	1.00	280	1.00	1.00
ψ_9	548	549	3185	1.00	1002	1.00	1.00
ψ_{10}	817	819	3765	1.01	1425	1.00	1.02
ψ_{11}	408	408	1336	1.00	553	1.00	1.01
ψ_{12}	342	342	4145	1.01	802	1.00	1.00
ψ_{13}	1007	1014	10,691	1.00	4025	1.00	1.00
ψ_{14}	187	190	383	1.00	721	1.01	1.01
ψ_{15}	51	56	67	1.03	203	1.03	1.09
ψ_{16}	90	93	98	1.01	597	1.01	1.01
ψ_{17}	162	162	508	1.00	595	1.00	1.00
ψ_{18}	129	131	284	1.01	622	1.01	1.01
ψ_{19}	74	77	165	1.01	257	1.01	1.02
ψ_{20}	208	212	487	1.01	1444	1.02	1.10
α_1	115	121	236	1.02	193	1.00	1.00
α_2	130	135	215	1.02	200	1.01	1.01
α_3	101	104	204	1.02	166	1.02	1.07
α_4	81	83	212	1.02	181	1.03	1.11
α_5	82	83	221	1.02	178	1.03	1.13
α_6	139	142	239	1.01	222	1.01	1.02
α_7	74	75	216	1.03	199	1.03	1.14
α_8	73	75	206	1.03	190	1.03	1.14
α_9	64	66	216	1.03	185	1.04	1.15
α_{10}	62	63	226	1.03	188	1.03	1.15
α_{11}	67	66	221	1.03	186	1.03	1.11
α_{12}	72	72	232	1.02	196	1.02	1.10
α_{13}	83	84	213	1.02	181	1.02	1.09
ρ	506	484	2784	1.01	720	1.00	1.01
ϕ	10,338	10,320	11,275	1.00	10,255	1.00	1.00
τ_ψ	9396	9298	11,238	1.00	10,454	1.00	1.00
τ_α	176	187	1079	1.01	617	1.01	1.06

ESS (basic), ESS (bulk), ESS (tail), and rank-normalized split (rn-split) $\hat{R}$ are computed following Vehtari et al. (2021) using the `posterior` package. Classical $\hat{R}$, its 95% upper confidence bound (UCB), and ESS are computed using the `coda` package (Plummer et al. 2006)




Fig. 10 Residuals versus obesity rate predictions (posterior mean). The solid curve represents a smooth fitted function

Figure 10 displays the residuals as a function of the posterior mean predictions. The residual cloud is symmetrically distributed around zero, with no discernible structure or systematic trend across the range of predicted values. The smooth fitted curve remains close to zero, indicating the absence of systematic structure in the residuals.

We additionally summarize residual behavior jointly across space and time using a heatmap representation. Figure 11 displays the mean posterior residual for each region-year combination, obtained by averaging posterior mean residuals within each spatial unit and year. The color scale is centered at zero, such that positive (negative) values indicate systematic underprediction (overprediction) by the model. The heatmap provides a compact diagnostic for detecting localized or persistent model misspecification across regions or over time. Overall, residuals appear small in magnitude and are distributed symmetrically around zero across both dimensions. No region exhibits sustained positive or negative residuals over consecutive years, nor is there evidence of a coherent temporal trend shared across regions. Instead, residual deviations appear sporadic and weakly structured, suggesting that the model adequately captures the dominant spatio-temporal patterns in the data.




Fig. 11 Heatmap of mean posterior residuals by region and year. Each cell represents the average posterior mean residual for a given region-year combination. The color scale is centered at zero, with warmer (cooler) colors indicating negative (positive) residuals

Posterior predictive checks

We further evaluated model adequacy using posterior predictive checks based on the full posterior predictive distribution. For each observation, we summarized the posterior predictive draws by computing the posterior predictive mean and the 2.5%, 50%, and 97.5% predictive quantiles. These summaries were then compared with the corresponding observed outcomes. Figure 12 displays the observed values against the posterior predictive mean predictions, with vertical error bars representing the 95% posterior predictive intervals. The dashed diagonal line indicates perfect agreement between predictions and observations. Overall, the observed data align closely with the posterior predictive means, with points symmetrically scattered around the identity line and no visible systematic deviations across the range of predicted values. The width of the predictive intervals appears relatively stable across observations. Quantitatively, the vast majority of observations fall within their corresponding 95% posterior predictive intervals, indicating good overall predictive calibration. Only a small subset of observations lies outside these intervals: specifically, 18 out of 520 observations (approximately 3.5%) fall below the 2.5% or above the 97.5% posterior predictive quantiles. Overall, the limited number, moderate magnitude, and heterogeneous distribution of these exceedances support the conclusion that the model adequately captures both the central tendency and the uncertainty of the observed outcomes.




Fig. 12 Posterior predictive check comparing observed outcomes to posterior mean predictions. Points represent individual observations, vertical bars denote 95% posterior predictive intervals, and the dashed line indicates the identity line. Red points denote the prediction means falling outside the 95% credible interval

Acknowledgements L.R. gratefully acknowledges the University of Bergamo for funding a Ph.D. scholarship in Statistics at the University of Milano-Bicocca. The authors are grateful for the advices given by two anonymous referees and the associate editor.

Funding Open access funding provided by Università degli Studi di Bergamo within the CRUI-CARE Agreement. Not applicable

Data Availability The assembled dataset and the code used for implementing the analysis and running the models are available at the following GitHub repository: <https://github.com/LucianoRota00/Spatio-Temporal-Dynamics-of-Obesity-in-Italian-Regions.git>.

Declarations

Conflict of interest Not applicable

Ethics approval and consent to participate Not applicable

Consent for publication Not applicable

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this

article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Ascarì R, Di Brisco AM, Migliorati S et al (2024) A multivariate mixture regression model for constrained responses. *Bayes Anal* 19(2):377–405. <https://doi.org/10.1214/22-BA1359>
- Assaf S, Campostrini S, Xu F et al (2015) Analysing behavioural risk factor surveillance data by using spatially and temporally varying coefficient models. *J R Stat Soc Ser A Stat Soc* 179(1):153–175. <https://doi.org/10.1111/rssa.12114>
- Banerjee S, Carlin BP, Gelfand AE (2014) Hierarchical modeling and analysis for spatial data, 2nd edn. Chapman and Hall/CRC. <https://doi.org/10.1201/b17115>
- Barra C, Barone A (2021) Gender differences in weight status and early school leaving in Italy. *Region Sci Policy Pract* 14(3):644–666. <https://doi.org/10.1111/rsp3.12491>
- Berloco C, Argiento R, Montagna S (2023) Forecasting short-term defaults of firms in a commercial network via bayesian spatial and spatio-temporal methods. *Int J Forecast* 39(3):1065–1077. <https://doi.org/10.1016/j.ijforecast.2022.05.003>
- Besag J, Kooperberg C (1995) On conditional and intrinsic autoregressions. *Biometrika* 82(4):733–746. <https://doi.org/10.1093/biomet/82.4.733>
- Brooks SP, Gelman A (1998) General methods for monitoring convergence of iterative simulations. *J Comput Graph Stat* 7(4):434–455
- Brunello G, Labartino G (2014) Regional differences in overweight rates: the case of Italian regions. *Econ Hum Biol* 12:20–29. <https://doi.org/10.1016/j.ehb.2012.10.001>
- Cepeda-Cuervo E, Núñez-Antón V (2013) Spatial double generalized beta regression models: extensions and application to study quality of education in Colombia. *J Educ Behav Stat* 38(6):604–628. <https://doi.org/10.3102/1076998613499779>
- Cribari-Neto F, Souza TC (2012) Testing inference in variable dispersion beta regressions. *J Stat Comput Simul* 82(12):1827–1843. <https://doi.org/10.1080/00949655.2011.599033>
- da Silva AR, de Oliveira Lima A (2017) Geographically weighted beta regression. *Spatial Stat* 21(A):279–303. <https://doi.org/10.1016/j.spa.2017.07.011>
- De Blasio G, Nuzzo G (2009) Historical traditions of civiness and local economic development. *J Reg Sci* 50(4):833–857. <https://doi.org/10.1111/j.1467-9787.2009.00624.x>
- de Valpine P, Turek D, Paciorek CJ et al (2017) Programming with models: writing statistical algorithms for general model structures with nimble. *J Comput Graph Stat* 26(2):403–413. <https://doi.org/10.1080/10618600.2016.1172487>
- de Valpine P, Paciorek C, Turek D et al (2024) NIMBLE: MCMC, particle filtering, and programmable hierarchical modeling. <https://doi.org/10.5281/zenodo.1211190>, <https://cran.r-project.org/package=nimble>, r package version 1.2.1
- Eurostat (2022) European health interview survey (ehis) – body mass index (bmi) by sex, age, and educational attainment level. https://doi.org/10.2908/hlth_ehis_bmi1b
- Ferrari S, Cribari-Neto F (2004) Beta regression for modelling rates and proportions. *J Appl Stat* 31(7):799–815. <https://doi.org/10.1080/0266476042000214501>
- Figueroa-Zúñiga JI, Arellano-Valle RB, Ferrari SL (2013) Mixed beta regression: a Bayesian perspective. *Comput Stat Data Anal* 61:137–147
- Gallus S, Odone A, Lugo A et al (2013) Overweight and obesity prevalence and determinants in Italy: an update to 2010. *Eur J Nutr* 52:677–685. <https://doi.org/10.1007/s00394-012-0372-y>
- Gelman A (2006) Prior distributions for variance parameters in hierarchical models (comment on article by Browne and draper). *Bayes Anal* 1(3):515–534. <https://doi.org/10.1214/06-BA117A>
- Gelman A, Rubin DB (1992) Inference from iterative simulation using multiple sequences. *Stat Sci* 7(4):457–472
- George EI, McCulloch RE (1993) Variable selection via Gibbs sampling. *J Am Stat Assoc* 88(423):881–889. <https://doi.org/10.2307/2290777>

- George EI, McCulloch RE (1997) Approaches for Bayesian variable selection. *Stat Sin* 7(2):339–373 (<https://www.jstor.org/stable/24306083>)
- Gorber SC, Tremblay M, Moher D (2007) A comparison of direct vs. self-report measures for assessing height, weight and body mass index: a systematic review. *Obes Rev* 8(4):307–326. <https://doi.org/10.1111/j.1467-789X.2007.00347.x>
- Grimaccia E, Rota L (2025) Spatiotemporal analysis of obesity: the case of Italian regions. *Obesities*. <https://doi.org/10.3390/obesities5020037>
- Haining R, Li G (2020) Modelling spatial and spatial-temporal data: a Bayesian approach, 1st edn. Chapman and Hall/CRC. <https://doi.org/10.1201/9780429088933>
- Held L, Rue H (2010) Conditional and intrinsic autoregressions. In: Gelfand AE, Diggle PJ, Fuentes M et al (eds) *Handbook of spatial statistics*. CRC Press/Chapman & Hall, Boca Raton, pp 201–216. <http://www.crcpress.com/product/isbn/9781420072877>
- Hoff PD (2009) *A first course in Bayesian statistical methods*, vol 580. Springer, New York. <https://doi.org/10.1007/978-0-387-92407-6>
- Janssen F, Bardoutsos A, Vidra N (2020) Obesity prevalence in the long-term future in 18 European countries and in the USA. *Obes Facts* 13(5):514–527. <https://doi.org/10.1159/000511023>
- Johnson NL, Kotz S, Balakrishnan N (1995) *Continuous univariate distributions*, vol 2, 2nd edn. Wiley, New York
- Kelsall JE, Wakefield JC (1999) Discussion of “Bayesian models for spatially correlated disease and exposure data”. In: Bernardo JM, Berger JO, Dawid AP et al (eds) *Bayesian statistics*, vol 6. Oxford University Press, Oxford, p 151
- Knorr-Held L (2000) Bayesian modelling of inseparable space-time variation in disease risk. *Stat Med* 19(17/18):2555–2570
- Lagos-Alvarez BM, Fustos-Toribio R, Figueroa-Zuniga J et al (2017) Geostatistical mixed beta regression: a Bayesian approach. *Stoch Env Res Risk Assess* 31(2):571–584
- Lesaffre E, Rizopoulos D, Tsonaka R (2007) The logistic transform for bounded outcome scores. *Biostatistics* 8(1):72–85. <https://doi.org/10.1093/biostatistics/kxj034>. (epub 2006 Apr 5)
- McElreath R (2016) *Statistical rethinking: a Bayesian course with examples in R and Stan*, 1st edn. Chapman and Hall/CRC. <https://doi.org/10.1201/9781315372495>
- Migliorati S, Di Brisco AM, Ongaro A (2018) A new regression model for bounded responses. *Bayes Anal* 13(3):845–872. <https://doi.org/10.1214/17-BA1079>
- Mills TC (2008) Forecasting obesity trends in England. *J R Stat Soc Ser A Stat Soc* 172(1):107–117. <https://doi.org/10.1111/j.1467-985X.2008.00550.x>
- OECD (2022) *Health at a Glance: Europe 2022: State of Health in the EU Cycle*. OECD, Paris. <https://doi.org/10.1787/507433b0-en>
- Paradinas I, Pennino MG, López-Quílez A et al (2018) Modelling spatially sampled proportion processes. *REVSTAT Stat J* 16(1):71–86
- Plummer M, Best N, Cowles K et al (2006) Coda: Convergence diagnosis and output analysis for mcmc. *R News* 6(1):7–11 (<https://journal.r-project.org/archive/>)
- Rockova V, Lesaffre E, Luime JJ et al (2012) Hierarchical Bayesian formulations for selecting variables in regression models. *Stat Med* 31(11–12):1221–1237. <https://doi.org/10.1002/sim.4439>
- Rue H, Held L (2005) *Gaussian markov random fields: theory and applications*, 1st edn. Chapman and Hall/CRC. <https://doi.org/10.1201/9780203492024>
- Simas AB, Barreto-Souza W, Rocha AV (2010) Improved estimators for a general class of beta regression models. *Comput Stat Data Anal* 54(2):348–366. <https://doi.org/10.48550/arXiv.0809.1878>
- Thomas A, Best N, Lunn D et al (2014) *GeoBUGS User Manual: Version 3.2.3*. MRC Biostatistics Unit, Cambridge, UK. <https://chjackson.github.io/openbugsdoc/GeoBUGS/Manuals/Manual.html>
- Vari R (2016) Gender-related differences in lifestyle may affect health status. *Ann Istit Super Sanità* 52(2):158–166. https://doi.org/10.4415/ANN_16_02_06
- Vehtari A, Gelman A, Simpson D et al (2021) Rank-normalization, folding, and localization: an improved $\hat{r}$ for assessing convergence of mcmc (with discussion). *Bayes Anal* 16(2):667–71. <https://doi.org/10.1214/20-BA1221>
- Waller LA, Carlin BP, Xia H et al (1997) Hierarchical spatio-temporal mapping of disease rates. *J Am Stat Assoc* 92(438):607–617. <https://doi.org/10.2307/2965708>