

Eibich, Peter; Ziebarth, Nicolas

Conference Paper

Examining the Structure of Spatial Health Effects using Hierarchical Bayes Models

Beiträge zur Jahrestagung des Vereins für Socialpolitik 2013: Wettbewerbspolitik und Regulierung in einer globalen Wirtschaftsordnung - Session: Health and the Economy, No. F07-V1

Provided in Cooperation with:

Verein für Socialpolitik / German Economic Association

Suggested Citation: Eibich, Peter; Ziebarth, Nicolas (2013) : Examining the Structure of Spatial Health Effects using Hierarchical Bayes Models, Beiträge zur Jahrestagung des Vereins für Socialpolitik 2013: Wettbewerbspolitik und Regulierung in einer globalen Wirtschaftsordnung - Session: Health and the Economy, No. F07-V1, ZBW - Deutsche Zentralbibliothek für Wirtschaftswissenschaften, Leibniz-Informationszentrum Wirtschaft, Kiel und Hamburg

This Version is available at:

<https://hdl.handle.net/10419/79844>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Examining the Structure of Spatial Health Effects using Hierarchical Bayes Models

Peter Eibich¹ and Nicolas R. Ziebarth²

This Version: November 11, 2012

Abstract

This paper makes use of Hierarchical Bayes Models to model and estimate spatial health effects. We focus on Germany, combining rich individual-level household panel data with administrative county-level information to estimate spatial county-level health dependencies. As dependent variable, we use the generic, continuous, and quasi-objective SF12 health measure. Our findings reveal strong and highly significant spatial dependencies and clusters. The strong and systematic county-level impact is comparable to an age effect on health of up to 31 years. Even 20 years after the peaceful German reunification, we detect a clear spatial East-West health pattern that equals an age impact on health of up to 9 life years.

Keywords: spatial health effects, Hierarchical Bayes Models, Germany, SOEP, SF12

JEL codes: C21, C11, I12, I14

¹ Corresponding author: DIW Berlin, Mohrenstrasse 58, 10117 Berlin, Germany, e-mail: peibich@diw.de.

² Cornell University, Policy Analysis and Management (PAM), 106 Martha van Rensselaer Hall, Ithaca, NY 14853, DIW Berlin, and IZA Bonn, e-mail: nrz2@cornell.edu, Phone: +1-(607) 255-1180, Fax: +1-(607) 255-4071.

1 Introduction

It is long recognized by applied empirical researchers that regional and neighborhood effects may play a crucial role in the analysis of a wide array of relevant outcomes. For example, in the economics literature, Burchardi and Hassan (2011) exploit the natural experiment of the German reunification to demonstrate that social ties have a long-lasting impact on individual economic prosperity. Studies in other fields, e.g. sociology, analyze the relevance of relative deprivation theories by looking at a respondents' relative income position in the neighborhood (cf. Knies et al., 2008; Durlauf, 2003).

When it comes to regional or neighborhood impacts on individuals' health, the literature is fragmented, both within and across fields. A relatively large set of studies in the epidemiological and public health literature focuses on geographical variation in health or health care (Riva et al., 2007; Vogtländer et al., 2010; Sundmacher et al., 2012). This rich strand of the literature basically emerged from the seminal "Small area variations in health care delivery" article by Wennberg und Gittelsohn (1973). The Small Area Variation Literature is extremely influential, particularly in the US. The DARTMOUTH ATLAS OF HEALTH CARE provides detailed descriptive information on how the health care infrastructure, as well as health care utilization, vary on a disaggregated geographical level, e.g., on the US county level (Dartmouth Atlas of Health Care, 2012). The high policy relevance comes from the fact that utilization and spending measures are not systematically correlated with better health outcomes. Understanding the driving forces of regional differences might, therefore, offer opportunities to limit health care spending while improving health outcomes (Wennberg et al., 2002).

A new subfield in the health economics literature is emerging. This subfield acknowledges the importance of regional factors in the analysis of health and health care. More precisely, this subfield econometrically models spatial interdependencies explicitly. In their synthesis of multilevel studies, Riva et al. (2007) point out that the vast majority of studies account only for within-area correlation and disregard between-area dependency. Not taking spatial patterns in the empirical models into account means that one implicitly assumes that the geographical units are statistically independent. This might be a strong and misleading assumption, since administrative or statistical boundaries might not reflect appropriately underlying ecological, social, and economic processes. Spillover effects are likely to occur.

Within the field of health economics, the statistical and econometric spatial modeling approaches, as well as the field of inquiry, may differ. For example, Baltagi and Moscone (2010) and Moscone and Tosetti (2010) analyze the relationship between income and health expenditures using the Common Correlated Effects (CCE) Approach and impose a spatial autoregressive (SAR) structure on the error terms (cf. Pesaran, 2006). Moscone et al. (2007) use a similar approach, but analyze cross-municipality variation in mental health expenditures. Felder and Tauchmann (2012) combine a nonparametric efficiency analysis with parametric regressions and model spatial dependence in healthcare provision at the county level in Germany.

This paper models and estimates spatial patterns of health in Germany. In terms of our empirical modeling techniques, we apply Hierarchical Bayes Models and advance the literature by combining three methodological approaches: First, hierarchical models are used to account for intragroup correlation and to estimate the effect of predictors from different hierarchical levels on individual level endpoints (cf. Hox, 2002; Bryk and Raudenbush, 1992). We use a three-stage hierarchical model that also accounts for the temporal dimension of the data. Second, we employ methods for spatial data analysis to analyze spatial effects, such as MORAN'S I or Conditional Autoregressive (CAR) models (Arbia, 2006; Cressie, 1993). We are aware of very few studies that use such methods in a health-related context (Costa-Font and Pons-Novell, 2007; Ocaña-Ricola and Mayoral-Cortés, 2010). In addition, spatial models are almost exclusively used with aggregated time-series data. Moreover, most studies investigate regional differences in health care utilization and health care spending, but not in health outcomes (Fuchs et al., 2004). We are not aware of any study in the field of health economics that models individual-level health using spatial econometrics and panel data.³ Third, we carry out our analysis using Bayesian methods (Banerjee et al., 2004; Cressie and Wikle, 2011).

In terms of content, we focus on a continuous and generic health measure, the SF12. The SF12 was developed by public health scientists and has advantages over the conventional categorical subjective health measures (Frick and Ziebarth, 2012). It can be interpreted as a quasi-objective single measure of an individual's health status (RAND, 1995). This paper intentionally focuses on a single continuous measure of individual quasi-objective health rather than measures of the health care infrastructure, which do not vary at the individual level. By taking into account a rich set of socio-demographic individual-level information, we net out important impact factors on health such as age, gender, employment, and marital status. Importantly, we incorporate health behavior measures such as smoking, alcohol consumption, and BMI. In addition, we net out individual variation in health that can be explained by differences in health care utilization. Moreover, we incorporate ten county-level indicators that vary across the 401 German counties, such as the unemployment rate and population density. Hence, in our three level hierarchical models, we strive to unravel the conditional spatial structure of population health after having corrected for the rich set of individual and county-level factors just listed.

Our results show that spatial interdependences and health clusters are of great importance. The systematic county-level impact on individual health is the most important driving force in our variable list above. County of residence may have an impact on individual health that equals the impact of up to 31 life years. Particularly stunning is the clear East-West health pattern that seems to be persistent even 20 years after the fall of the Berlin Wall. It reflects the four decade long impact of Communism versus Capitalism on health and results in an age impact of up to 9 life years.

³ Notable exceptions in the fields of Biology and Sociology include Mueller et al. (2001), who examine spatial patterns of child growth in Papua New Guinea, and Browning et al. (2003), who analyze variation in health status across Chicago neighborhoods during the 1990s.

The rest of the paper is organized as follows: In Section 2 we provide a short description of our underlying dataset, the German Socio-Economic Panel Study (SOEP). Sections 3.1 and 3.2 explain the econometric models and review the basics of Bayesian inference. Section 4 presents and discusses the results. Section 5 concludes.

2 Data

2.1 Dataset

For the empirical analysis we make use of individual health data and further socioeconomic characteristics provided by the German Socio-Economic Panel Study (SOEP). The SOEP is a representative panel study of private households. Starting in 1984, SOEP interviews subjects annually, and expanded to include residents in former East German lands, starting in 1992. All respondents answer one main individual questionnaire covering about 150 questions covering a range of topics, such as the labor market and family situation, attitudes and perceptions, as well as health. Additionally, a household questionnaire is completed by the head of the household. Since 2000 the survey has reached more than 20,000 individuals across 10,000 households. For further details, see Wagner, Frick and Schupp (2007).

The SOEP provides a variety of health measures. Both the standard 5-categorical Self-Assessed Health (SAH) measure and the 11-categorical health satisfaction measure are included every survey wave. In addition, since 2002, in every second year, the continuous quasi-objective SF12 measure and the objective grip strength measure are included in the survey. Furthermore, information on health-related behavior (e.g. alcohol or tobacco consumption) is available since wave W in 2006. Therefore, we restrict our analysis to the three waves: W (2006), Y (2008), and BA (2010). We use only observations without item-non-response. In total we obtain 54,683 person-year observations from 23,410 different individuals.

2.2 Dependent Variable: SF12

Our dependent variable is the generic health measure SF12. A specific algorithm generates the continuous SF12 measure on the basis of 12 different health-related questions. More precisely, the algorithm generates eight subscales and two superordinate dimensions, namely physical (pc) and mental health (mc). We then average over both components to obtain our dependent variable $SF12$. In the standard SOEP version, the SF12 takes on continuous values between 0 and 100, has mean 50, and a standard deviation of 7 (see Table 1).⁴

2.3 Individual Covariates

⁴ A detailed description of the algorithm and an overview over the differences between the original “SF12v2™ Health Survey” and the SOEP version can be found in Andersen et al. (2007).

By controlling for demographic factors, educational characteristics, labor market participation, health-related behavior as well as health care utilization, we net out all differences in health as measured by the SF12 that can be traced back to these factors.

The **demographic factors** are age, age squared, sex, marital status, and the number of children under 14 in the household. Table 1 shows that average age is 49.3 and that about half of our sample is female. The majority of respondents are married and live in the same household as their spouse. On average, 0.36 children live in a SOEP household.

In terms of **education and labor market participation**, we control for whether an individual completed a vocational training or holds a university degree, the working status of the respondent, and the monthly equivalent household income.⁵ Less than a quarter holds a university degree, but about 60 percent completed vocational training. Vocational training is equivalent to a college degree in the US. Slightly less than half of our sample works and the average equivalent gross household income per capita is €1,800 per month.

Concerning **health-related behavior**, we control for alcohol and tobacco consumption and Body Mass Index (BMI). Alcohol consumption is captured by dummy variables. In the SOEP questionnaire, the participants are asked to state how often they drink wine and sparkling wine, beer, spirits and mixed drinks. If an individual states that they do not drink any alcohol at all, *no alcohol consumption* is assigned the value “1”, otherwise it is “0”. In contrast, if an individual states to drink any kind of alcohol on a regular basis, *regular alcohol consumption* takes on the value “1”, otherwise “0” (Ziebarth and Grabka, 2009). 17 percent drink alcohol on a regular basis and 13 percent never drink any alcohol. The dummy variable *Smoking* captures whether the respondent consumes tobacco (cigarettes, cigars or pipes). It takes on the value “1” for about 30 percent of the respondents in our sample. The *BMI* is measured as body weight in kilogram divided by the squared height in meters. The average BMI lies slightly below 26.

Health care utilization is measured by the number of hospital stays during the previous calendar year and the number of doctor visits in the 12 months prior to the interview.⁶

Summary statistics for all variables are given in Table 1.

[Insert Table 1 about here]

2.4 Regional Information

⁵ The monthly equivalent household income uses the OECD-modified scale and assigns households a value of 1 to the first adult, 0.5 to each additional adult and 0.3 to each additional child. Further details are provided in *OECD Project on Income Distribution and Poverty (2009)*.

⁶ It should be noted that the respondents are asked to state the number of doctor visits during the last three months prior to the interview. The answer is then multiplied by four to generate the annual number of doctor visits. This may overestimate the actual annual number because approximately two-thirds of all interviews are carried out in the first quarter of a year and there is a clear seasonal pattern in doctor visits.

The SOEP provides information on the place of residence at multiple levels, ranging from the federal states to specific postal codes (Knies and Spiess, 2007). The SOEP also provides frequency and probability weights to ensure representativeness for most federal states. The state-level is probably too aggregate to detect any significant regional effects – especially if these are caused by small scale phenomena (e.g. structural environmental and economic factors). Therefore in this paper, we focus on counties (“*Landkreise*” and “*Kreisfreie Städte*”). This offers several advantages: First, they are the smallest administrative units in Germany. Furthermore, administrative data exists at the county-level, which can be used as independent covariates to explain variation between counties. Lastly, the SOEP contains observations from almost all counties. Nevertheless there are some disadvantages. For instance, the data is not representative and there are no weights available to account for this. The number of observations per county varies between 2 (Sömmerda) and 2,025 (Berlin) with a mean of 136 observations and a standard deviation of 148. Thus the findings of this analysis should be interpreted carefully. For some counties the number of observations is too low to draw definite conclusions, especially when it comes to the health development of the county over time.

The number of counties, and their respective borders, has evolved since 2006. Therefore, the SOEP Group constantly updates county codes. To ensure consistency of the data with the shapefiles provided by the FEDERAL AGENCY FOR CARTOGRAPHY AND GEODESY (“*BUNDESAMT FÜR KARTOGRAPHIE UND GEODÄSIE*”), we recoded the county codes to reflect the boundaries as of January 1, 2012.

As of January 1, 2012, Germany consisted of 402 counties. Our dataset contains observations from 401 counties with Memmingen, a county in Bavaria, missing. We ignore this county in our analysis and remove it from the neighborhood matrix for the spatial models.

In order to explain systematic county-level differences in health, we incorporate county-level information into our analysis. For this purpose, we exploit information provided by the FEDERAL INSTITUTE FOR RESEARCH ON BUILDING, URBAN AFFAIRS, AND SPATIAL DEVELOPMENT (2012) (“*BUNDESINSTITUT FÜR BAU-, STADT- UND RAUMFORSCHUNG*”) in their INKAR (“*INDICATORS AND MAPS ON SPATIAL DEVELOPMENT*”) database.⁷ As can be inferred from Table 1, we use county-level data on the *area size*, the *population density*, the *proportion of the area used for settlement and transport*, and the *proportion of the area suitable for recreation* in order to ***distinguish urban from rural counties***. Note that there is significant variation in population density across counties, ranging from 40 to 4,257 inhabitants per square kilometer. There is also wide variation in terms of the size of the area that is used for urban purposes, i.e., the degree of urbanization. The values for this variable vary from 5 to 76 percent. The size of the area for recreational purposes varies between 0.2 and 15 percent. Obviously, these indicators capture a great deal of cross-county heterogeneity that is not captured by individual-level indicators and might affect health. For example, in general, population density is highly correlated with the density of health care providers

⁷ As for the individual level, we recoded the county codes to the status as of 2012. To accommodate our dataset to the 2011 county mergers, we calculate the values for the new counties as a weighted average of the values of the merged counties. The variables are weighted either by population number or by area size.

and the health care infrastructure. On the other hand, low urbanization rates or large recreational facilities may represent low air pollution and low noise (e.g. Coneus and Spieß, 2011).

We use average *available monthly income per capita* and the annual *unemployment rate* to characterize the **degree of regional deprivation**. The *population share of immigrants* represents the diversity of the community. Again, we observe huge variation across counties for all these indicators. Average per capita income varies from €1,146 to €2,544. Between 1 and 25 percent of the population of a county are immigrants and unemployment rates vary from 2.5 to 21.30 percent. The *ratio of overnight stays to inhabitants* in the tourism sector may both indicate that a county is located in a landscape that is beneficial to health, e.g. seaside beaches or in the mountains. Alternatively, it may measure the degree of the traffic volume (caused by tourism) in the county. The *car density* may serve as an indicator of air and noise pollution. It varies between 316 and 881 cars per 1,000 inhabitants.

The *physician density* (measured as inhabitants per physician) represents the **health care infrastructure** of the county. In counties with fewer physicians, residents may experience longer travel and waiting times if they want to consult a physician, i.e., a lower degree of access to and utilization of the health care infrastructure and, consequently, health. On the other hand, following the conclusions of Wennberg and Gittelsohn (1973), physician density might be uncorrelated or negatively correlated with individual health.

3 Hierarchical Bayes Models and Methods

3.1 Econometric Models

The main aim of our analysis is to detect and model spatial patterns in the distribution of regional effects on health. With a few notable exceptions (discussed in Section 1), most studies using methods from spatial econometrics or spatial statistics rely on aggregate measures of morbidity (e.g. disease counts) as a dependent variable. In contrast, we model individual health. At the same time, we control for systematic differences between counties as outlined above. In addition, we control for common time shocks across counties. However, the main objective of the paper is to model spatial heterogeneity in health at the individual level. The resulting models are inherently hierarchical with observations nested within individuals nested within regions. The following model serves as our baseline model throughout the analysis:

$$SF12_{ist} \sim N(\mu_{ist}, \sigma^2)$$

$$\mu_{ist} = \theta + b_s + c_i + X_{it}\beta + Z_{st}\gamma + \delta_t \quad (3.1.1)$$

We assume that the *SF12* indicators for individuals $i = 1, \dots, 23,410$ in counties $s = 1, \dots, 401$ and years $t = 2006, 2008, 2010$ are conditionally independent and normally distributed with mean μ_{ist} and variance σ^2 . The mean μ_{ist} is a linear function of an intercept θ , the (1 x 15) vector of individual

regressors X_{it} with a (15×1) parameter vector β , the (1×10) vector of county-level covariates Z_{st} with a (10×1) parameter vector γ , aggregate time effects δ_t , an individual effect c_i and regional effect b_s .

Our interest lies mainly in the regional effects b_s and its spatial distribution. In the following section, we specify four different assumptions on the distribution of these effects and their spatial and temporal dependency, which will lead us to our candidate models

3.1.1 Model 1: Unstructured Regional Effects (URE)

In our first model, we assume that the county-level effects follow a normal distribution with mean zero and variance σ_b^2 :

$$b_s \sim N(0, \sigma_b^2).$$

This implies that, on average, the regional effects are zero, which is reasonable since we include an intercept into our model. This model would be equivalent to a model with dummy variables for each county, i.e., county-fixed effects. We do not specify the type of spatial dependency in this model, which is why we refer to this model as the **Unstructured Model (Model 1)**. In this initial stage of our analysis, we seek to obtain estimates of the regional effects whose spatial distribution we would like to analyze. These estimates are “raw” in the sense that they are not smoothed by assumptions on the spatial dependency; they are used to examine the spatial dependency patterns of the regional effects by maps and test statistics.

Then, in the next stage, we impose a structure on the spatial patterns and use models that smooth the regional effects. This approach is proposed by Banerjee et al. (2004).

3.1.2 Model 2: Intrinsic Conditional Autoregressive Regional Effects (ICAR)

In order to account for spatial dependency between counties, we model regional effects using an **Intrinsic Conditional Autoregressive (ICAR) Model (Model 2)** as proposed by Besag et al. (1991). This model is autoregressive in a sense that the value for each area b_s depends on its spatial lags $b_r, r \neq s$. It is characterized by the conditional distributions:

$$b_s | b_{r \neq s} \sim N(\bar{b}_s, \frac{\sigma_b^2}{m_s}). \quad (3.1.2)$$

The regional effect b_s is normally distributed around the mean, where $\bar{b}_s = m_s^{-1} \sum_{r \in \eta_s} b_r$. m_s yields the number of neighbors of county s . η_s is the set of neighbors of county s (cf. Banerjee et al., 2004). The mean equals the average effect of the neighboring counties of county s . σ_b^2 is a variance parameter. The variance depends on this parameter as well as the number of neighbors. Thus isolated counties with fewer neighbors exhibit greater variance than counties with more neighbors.

The conditional distribution in eq. (3.1.2) specifies that the regional effects b_s are Markovian, i.e., they depend only on their neighbors.⁸ From eq. (3.1.2), the joint density $p(\mathbf{b}_1, \dots, \mathbf{b}_n)$ can be obtained using Brook's Lemma (cf. Besag, 1974) and the Hammersley-Clifford Theorem. However, Besag (1995) notes that the resulting distributions would not exhibit appreciable correlations unless the parameters are close to the boundaries of the parameter space. Therefore, he suggests the pair-wise specification in eq. (3.1.3):

$$p(\mathbf{b}_1, \dots, \mathbf{b}_n) \propto \exp \left\{ -\frac{1}{2\sigma_b^2} \sum_{s \neq r} w_{sr} (b_s - b_r)^2 \right\}, \quad (3.1.3)$$

i.e., the joint distribution of $\mathbf{b}_1, \dots, \mathbf{b}_n$ is proportional to the pairwise difference density, where w_{sr} are the non-diagonal entries of the proximity matrix \mathcal{W} . The non-diagonal entries of the proximity matrix \mathcal{W} specify which areas are considered to be associated (and therefore spatially dependent). The choice of \mathcal{W} will be explained in more detail in section 4. σ_b^2 is the variance parameter.

Equation (3.1.3) is a limiting form⁹ of standard conditional autoregressive normal distributions and it is usually referred to as an Intrinsic Autoregression (IAR) or Intrinsic Conditional Autoregression (ICAR) (Besag, 1995). Eq. (3.1.3) does not yield a proper distribution, since the b_s are not centered and need an additional “sum to zero” constraint (Banerjee et al., 2004). Nevertheless, if used as prior information about the unknown parameter (i.e., as a prior distribution, see section 3.2.1) in Bayesian Analysis, the resulting inference (through the posterior distribution) is proper.

For our candidate Model 2, we assume an ICAR distribution for the regional effects and choose a neighborhood matrix \mathcal{W} based on first-order adjacency (see below).

3.1.3 Models 3 and 4: (Spatiotemporal) Convolution Prior (SCP)

In the candidate model described above, we assume that the regional effects can be described by an ICAR-model, which implies that the regional effect of county s depends only on the values of *all* neighboring counties. This assumption may be too strict, since the effect might be a combination of a spatially dependent part and a randomly distributed part (noise).

Therefore we decompose the effects

$$\mathbf{b}_s = \varphi_s + \omega_s,$$

where

$$\omega_s \sim N(0, \sigma_\omega^2)$$

and

$$\varphi_s | \varphi_{r \neq s} \sim N(\bar{\varphi}_s, \frac{\sigma_\varphi^2}{m_s}), \quad (3.1.4)$$

⁸ The model specified through eq. (3.1.2) is also called the Autonormal or Autogaussian model, i.e. a model for normally distributed data on a Markov Random Field (c.f. Arbia, 2006)

⁹ I.e. the parameter that determines the strength of the spatial association is fixed to the boundaries of the parameter space.

that is, we discriminate between regional heterogeneity captured by ω_s and regional clustering captured by φ_s (Banerjee et al. 2004, p. 283f). We call this a **Convolution Prior** since the distribution of b_s is a convolution of the distributions described above (Besag et al. 1991; Mollié, 1996). Using this structural assumption and the neighborhood matrix based on adjacency, we obtain our **Model 3: the Convolution Prior (CP)**.

Our candidate models 2 and 3 incorporate spatial dependence and aggregate time effects. However, the spatial effects b_s, φ_s and ω_s do not depend on time t , i.e., they are assumed to be time-invariant. This may be a plausible assumption if the main driving factors of the regional differences are time-invariant, e.g., structural environmental and economic factors. Nevertheless we test this assumption by extending the CP-Model to include space-time interaction effects. As such, we obtain our fourth and final model: the **Spatiotemporal Convolution Prior (SCP)**.

Consider the model

$$\mu_{ist} = \theta + X_{it}\beta + Z_{st}\gamma + b_{st} + c_i + \psi t, \quad (3.1.5)$$

where

$$b_{st} = \omega_{st} + \varphi_{st}$$

and

$$\varphi_{st} | \varphi_{rt \neq st} \sim N(\bar{\varphi}_{st}, \sigma_{\varphi t}^2 / m_s),$$

$$\omega_{st} \sim N(0, \sigma_{\omega t}^2).$$

Here, both the spatially dependent part and the random part of the regional effect depend on location s and time t , i.e., the model produces a separate effect for each county and each year (comparable to an interaction between dummy variables for county s and year t). In addition, the model includes a linear time trend with parameter ψ instead of the aggregated time effects. This model has the advantage that we can estimate the spatial pattern for each year. However, it requires observations for each county in each year. Therefore for this model, we restrict the dataset to counties with observations in each year. The adjusted dataset contains 54,528 out of 54,683 observations and 397 out of 401 counties.

3.2 Estimation

In this section we give a brief overview of the basics of Bayesian Methods and then detail the specifications under which the described models above are analyzed.

3.2.1 Bayesian Inference

In Bayesian Analysis, given the data, the main interest lies in learning about the distribution of the unknown parameters. In contrast to frequentist methods, such as Maximum Likelihood Estimation, the

parameters are treated as stochastic quantities that follow a probability distribution.¹⁰ The fundamental principle of Bayesian methods is summarized by Bayes formula:

$$p(\theta|y) = \frac{p(\theta)p(y|\theta)}{p(y)}, \quad (3.2.1)$$

where y is a random variable and θ is the unknown parameter. $p(y|\theta)$ is the likelihood function of the data. $p(\theta)$ is called a **prior distribution**, which expresses prior belief or knowledge (e.g., from similar or older studies) about the distribution of the unknown parameter. $p(y)$ is the unconditional distribution of the data and serves as a normalizing constant to ensure that $p(\theta|y)$ integrates to one. $p(\theta|y)$ is called the **posterior distribution** and summarizes all knowledge of the unknown parameters given the data. Bayes formula may be regarded as a method to update prior knowledge about the parameters θ using the data y (cf. Congdon, 2010).

Bayesian inference is carried out by summarizing the posterior distribution. The point estimates depend on the appropriate “loss function”, e.g., if we want to minimize $E[(\theta - \hat{\theta})^2]$, the estimator is given by the posterior mean. For multimodal or non-normal distributions, the posterior median may be more appropriate (cf. Banerjee et al., 2004).

Bayesian confidence intervals (usually called **credible intervals**) are **Equal Tail (ET) Intervals** or **Highest Posterior Density (HPD) intervals**. For symmetric and unimodal posterior distributions, both intervals are identical. Otherwise the HPD interval will have shorter length, but it is more difficult to compute.

The ET interval is given as the interval $[q_L, q_U]$ where:

$$\int_{-\infty}^{q_L} p(\theta|y)d\theta = \frac{\alpha}{2}$$

and

$$\int_{q_U}^{\infty} p(\theta|y)d\theta = 1 - \frac{\alpha}{2}.$$

As usual, α denotes the confidence level. An attractive feature of the credible interval is that the often erroneous interpretation of confidence intervals in frequentist models does actually hold for them, namely the probability that θ lies in $[q_L, q_U]$ is $(1 - \alpha)$.

A crucial point in Bayesian Analysis is the choice of the prior distribution. The posterior distribution depends on both the prior distribution and the likelihood. Likewise, the parameters of the posterior

¹⁰ It is worth noting that the parameters are not regarded as random quantities—they are fixed but unknown. The incomplete knowledge about these parameters is assumed to be random (cf. Gelman and Robert, 2012).

distribution (e.g. the mean) are a weighted average of the prior values and the likelihood-based values (for an example, see Banerjee et al., 2004).

This offers a way to formally include subjective beliefs or existing knowledge. However, often researchers prefer to not introduce any subjective elements, either since no additional information on the problem is available, or since the researcher prefers to remain as “objective” as possible. In this case, improper or uninformative priors may be chosen (e.g., a flat distribution over the entire real line). This leads to another problem: the determination of the posterior distribution.

3.2.2 The Gibbs sampler

An analytic derivation of the posterior through equation (3.2.1) is only possible for specific, “conjugate” prior distributions, i.e., when prior and posterior belong to the same distributional family. For non-conjugate, and especially for improper priors (i.e. distributions that do not integrate to anything finite), the required integrations - namely for $p(\mathbf{y})$ - are intractable. In this case, Markov Chain Monte Carlo (MCMC) methods may be used to obtain samples from the posterior distribution without knowing its exact analytical form.

The general idea of MCMC methods is to sample from a transition distribution

$$K(\theta^{(t)} | \theta^{(t-1)}, \dots, \theta^{(0)}) = K(\theta^{(t)} | \theta^{(t-1)}),$$

which has the posterior distribution $p(\theta | \mathbf{y})$ as its stationary distribution. $\theta^{(0)}$ is a set of initial values for the parameters $\theta = (\theta_1, \dots, \theta_k)'$. After a “burn-in phase” of B samples, $\theta^{(t)}, t \geq B$ may be seen as a random sample from the posterior $p(\theta | \mathbf{y})$ (cf. Congdon, 2010).

The Gibbs sampler described below generates samples from the joint posterior distribution by sampling from the full set of conditional distributions. Assume we are interested in k parameters $\theta = (\theta_1, \dots, \theta_k)'$. Then, for arbitrary initial values $\theta^{(0)}$, for $t=1, \dots, T$; repeat:

Step 1: Draw $\theta_1^{(t)}$ from $p(\theta_1^{(t)} | \theta_2^{(t-1)}, \theta_3^{(t-1)}, \dots, \theta_k^{(t-1)}, \mathbf{y})$.

Step 2: Draw $\theta_2^{(t)}$ from $p(\theta_2^{(t)} | \theta_1^{(t)}, \theta_3^{(t-1)}, \dots, \theta_k^{(t-1)}, \mathbf{y})$.

...

Step k: Draw $\theta_k^{(t)}$ from $p(\theta_k^{(t)} | \theta_1^{(t)}, \theta_2^{(t)}, \dots, \theta_{k-1}^{(t)}, \mathbf{y})$.

It can be shown that the k draws $\theta^{(t)}$ converge in distribution to a draw from the joint posterior $p(\theta | \mathbf{y})$ (c.f., Banerjee et al., 2004).

Samples from the conditional distributions may be obtained through other methods, such as the Metropolis-Hastings Algorithm. For an extensive treatment of MCMC methods, see Congdon (2010) or Robert and Casella (2005).

3.2.3 Specifications

Bayesian methods offer some important benefits. First, they offer a way to formally incorporate prior knowledge or beliefs into the analysis. Probably most important for our analysis, hierarchical and spatiotemporal models can be adapted quite naturally in a Bayesian framework (Congdon, 2010). Furthermore, Hierarchical Bayes Models “borrow” strength across units through the common “hyperdensity”¹¹ and therefore estimate the regional and individual effects more efficiently. Bayesian Analysis itself is especially suited for small samples – since all estimates are conditional upon the data given, we do not have to rely on asymptotic theory. Thus, we formulate our models as Hierarchical Bayes Models and estimate them using Gibbs sampling.

The prior distributions for our analysis are mostly uninformative. They allocate equal probability mass to all plausible values. In detail, we choose a flat prior $U(-\infty, \infty)$ for the intercept θ and a normal prior distribution with an inflated variance $N(0, 10000)$ for the slopes $\beta_1, \dots, \beta_{15}$ and $\gamma_1, \dots, \gamma_{10}$ and the time fixed effects δ_t .

For our first model, the **Unstructured Regional Effects (URE) Model**, we use a normal prior distribution for the regional effects with a hyperparameter for the variance. For the second model, the **Intrinsic Conditional Autoregressive Regional Effects (ICAR) Model**, we use the ICAR prior as described in section 3.1.2. For the third and fourth model, the **Convolution Prior (CP)** and the **Spatiotemporal Convolution Prior (SCP)** models, we use an inflated normal distribution for ω_s and ω_{st} as well as an ICAR prior for φ_s and φ_{st} . All standard deviation parameters (including hyperparameters) were assigned a $U(0, 100)$ prior distribution, which covers all plausible values (since the SF12 indicators lies in $[0, 100]$) and assigns them equal probabilities.

In order to speed up convergence, we standardized all non-dummy variables, i.e., we subtracted the overall mean and divided them by their standard deviation.

We run three parallel chains with dispersed initial values and monitored the values for the intercept, the slope parameters, the regional effects, and the variance parameters using **Trace Plots** and **Gelman-Rubin statistics** (cf. Gelman and Rubin, 1992; Brooks and Gelman, 1998). The length of the burn-in period was determined for each model separately.

¹¹ It is possible to specify a prior distribution in terms of unknown parameters, so-called “hyperparameters”. The prior distribution assigned to these hyperparameters is called “hyperprior” or “hyperdensity”.

For models 1 and 2, we could not detect any departure from convergence after the 5,000th iteration. For models 3 and 4, the burn-in period took 15,000 iterations, after which the sampler seemed to have stabilized. After the burn-in period, we sampled 15,000 draws from the posterior distribution. In order to decrease autocorrelation and speed up mixing, we thinned the chain by storing only every 10th draw.

4 Results

4.1 Model 1: Unstructured Regional Effects (URE)

Table 2 yields the results for our first model, the *Unstructured Regional Effects (URE) Model* (see Section 3.1.1). Please note that the coefficients are standardized for the non-dummy variables.

Almost all variables are statistically significant at conventional levels. A notable exception is *number of children*.

[Insert Table2 about here]

Looking at the sign of the coefficients, we note that, in general, males are healthier than females. This is in line with the stylized fact that women have greater health care expenditures, not only in childbearing years (cf. Owens, 2008). Also, individuals who abstain from alcohol have a lower health status. This group is likely to contain respondents with current or past serious diseases. On the other hand, regular alcohol consumption is positively correlated with health; this might seem surprising, but is consistent with the literature (cf. Ziebarth and Grabka, 2009). The number of hospital stays as well as the number of doctor visits in the last year, in other words, health care utilization, is negatively correlated with respondents' health.

Of the county-level variables, only the unemployment rate and the car density show a significant effect. The estimates of the spatial models show that the latter effect is not consistently estimated across models.

However, the main focus of our analysis lies on the regional effects b_s and its spatial distribution.

Choropleth Map to Visualize the Unstructured Regional Effects Model

The “posterior means”—i.e., the point estimates for the coefficients— for the URE-Model are plotted in the CHOROPLETH MAP in Figure 1. To produce this map, we order the values of our dependent SF12 health variable and classify them into five categories, with the quintiles serving as cutoff-thresholds. Each category is assigned a shade of green, with lighter shades corresponding to smaller values (lower quintiles) and darker shades corresponding to greater values (higher quintiles). Naturally, the intervals covered by each category differ; the highest and the lowest category cover a wider range of values than the medium

category, since the interval covered by the medium category has higher probability mass¹². However, the choice of categories is appropriate to give a visual impression of the spatial distribution of the random effects.

[Insert Figure 1 about here]

We see the following: First, the county level health effects differ between -2.4 and 2.7.¹³ This means that the overall SF12 mean (the intercept) varies between 46.94 and 52.04, depending on the county.¹⁴ As compared to the age effect, interestingly, the regional effect has an influence on health that equals up to 31 years in both directions.¹⁵

Second, it seems that the counties in East Germany tend to have lower values than the counties in West Germany. However, there are clusters of lower values in Bavaria, Lower Saxony and Rhineland-Palatinate and some rather high values in Brandenburg and Mecklenburg-Western Pomerania. Overall, the map is dominated by clusters, e.g., we find clusters of high values around Hanover (Lower Saxony) and Hamburg, Hof (Bavaria), Cologne (North Rhine-Westphalia) and the Palatinate region. Clusters of smaller values can be found in Lower Bavaria, Thuringia and in the east of Mecklenburg-Western Pomerania and Brandenburg.

Testing for Spatial Dependency using Moran's I

These clusters already suggest that the effects are not spatially randomly distributed and that there exists spatial dependency between neighboring counties. Therefore we calculate MORAN'S I to measure the strength of the spatial association and to formally test the hypothesis of no spatial autocorrelation.¹⁶

MORAN'S I can be seen as a spatial analogue for the lagged autocorrelation coefficient statistic for time series and is defined by:

$$I = \frac{n \sum_i \sum_j w_{ij} (Y_i - \bar{Y})(Y_j - \bar{Y})}{(\sum_{i \neq j} w_{ij}) \sum_i (Y_i - \bar{Y})^2} \quad (4.1.1)$$

Y_i is the value of random variable \mathbf{Y} at location $i = 1, \dots, n$, i.e. in our case the regional effect. w_{ij} are the off-diagonal entries of the neighborhood matrix \mathcal{W} and represent the spatial connection between sites i and j .¹⁷ \mathcal{W} can be defined as a binary matrix; in this case $w_{ij} = 1$ indicates that sites i and j are neighbors

¹² The distribution is approximately normal, i.e. unimodal and symmetric.

¹³ Of course not all of these effects are statistically significant. 92 of the 401 regional effects are significant at a 5 percent level and 156 are significant at a 10 percent level.

¹⁴ Standard Errors and Monte Carlo Error are excluded from the calculation.

¹⁵ Note that this does not refer to a change in life expectancy but to the marginal effects of the age polynomial in our regression models.

¹⁶ Arbia (2006) points out that MORAN'S I is not a proper statistical test since it does not consider an explicit alternative hypothesis. Since we use MORAN'S I as an exploratory tool before specifying our spatial models, this can actually be seen as an advantage (cf. Arbia, 2006).

¹⁷ w_{ii} are usually set to 0, otherwise each site would be a neighbor of itself.

(cf. Arbia, 2006). Another possibility would be to define w_{ij} as weights, e.g. proportional to the inverse distance.

Moran (1950) derives the asymptotic distribution of I , however, the variance expression is very complex. Banerjee et al. (2004) recommend carrying out a Monte Carlo permutation test by drawing a random sample of permutations¹⁸ including the observed one. Then, the observed I can be positioned relative to the other permutations and a pseudo p -value may be derived.

We calculate MORAN'S I using neighborhood matrices based on ***three different definitions of neighborhood*** (cf. Arbia, 2006) and test it using 10,000 permutations (including the observed values). The first definition is based on adjacency, i.e., sites are considered to be neighbors if they share a *common border*. This concept can be generalized to second- and third-order neighbors.¹⁹ Our second definition is based on a *threshold-distance*, i.e., sites i and j are neighbors if the distance between their centroids is smaller than the threshold distance d . We choose d to equal 50 km, 68.2, and 150 km. We include 68.2 because it is the smallest distance such that each county has at least one neighbor. The last definition is based on a *Nearest-Neighbor algorithm* for 3, 5 and 7 Nearest-Neighbors. Note that all matrices are binary and that the adjacency- and threshold-based definitions lead to symmetric matrices.

[Insert Table 3 here]

Table 3 reports the number of non-zero links (i.e. connections between sites), the value for MORAN'S I and the empirical p -value. A comparison of the number of links shows that the Nearest-Neighbor matrices define very small neighborhoods, whereas distance-based matrices result in larger neighborhoods. In general, MORAN'S I is not very high but all values are significant with the exception of the third-order adjacency matrix. We also observe that the value of I decreases with the size of the neighborhood. The highest value is obtained for the 3-Nearest-Neighbors matrix, whereas the distance-based and higher-order adjacency matrices result in considerably smaller values.

Thus we conclude that the spatial dependence between regional health effects occurs on a small, local scale. Although the 3-Nearest-Neighbors matrix delivers the highest value of I , we cannot use it to model spatial association in our models. The ICAR specifications require symmetric neighborhood matrices and with 1201 non-zero links the 3-Nearest-Neighbors matrix is clearly non-symmetric. In contrast, the first-order adjacency matrix is symmetric and delivers a value nearly as high.²⁰ All in all, we reject the hypothesis of independent regional effects and employ a first-order adjacency matrix for the spatial models as described in Section 3.

¹⁸ The observed values of the variable are randomly assigned to the regions.

¹⁹ Second-order neighbors are sites that share a common border with first-order neighbors of site i , but not with site i itself ("neighbors of neighbors").

²⁰ The higher value of the 3-Nearest-Neighbors matrix is possibly due to the much smaller number of links since it is neither reasonable that the number of neighbors is the same for all counties nor is there evidence for directed spatial dependence (as implied by the departure from symmetry).

Local Indicators of Spatial Association (LISAs)

MORAN'S I is a measure of global spatial autocorrelation. In contrast to that, LOCAL INDICATORS OF SPATIAL ASSOCIATION (LISAs) are calculated for each region individually and can be used to test the significance of the spatial association for a specific county. We calculate local P s as proposed by Anselin (1995). Again we use a Monte Carlo permutation test to obtain pseudo p -values. The p -values for each county are depicted in Figure 2. Here, the map corresponds to an adjacency-based neighborhood matrix. The classes correspond to different types of clusters. Counties with insignificant spatial association are depicted in grey. Counties with positive regional effects and significant spatial association are shown in dark red if the spatial correlation is positive—i.e., the neighbors exhibit positive regional effects—and in pink if the spatial correlation is negative—i.e., the neighbors exhibit negative regional effects.

Similarly, counties with negative regional effects and significant spatial association are shown in dark blue if the spatial correlation is positive and in light blue if the spatial correlation is negative (c.f. Anselin, 2005). We see that for most clusters, the spatial association is significant on a local level.

[Insert Figure 2 about here]

4.2 Model 2: Intrinsic Conditional Autoregressive Regional Effects (ICAR)

In the next step, we estimate the ICAR-Model to obtain smoothed estimates of the regional effects (see section 3.1.2). Our aim is to find a model that produces a more distinct spatial pattern, i.e., smoother estimates, but trades off little in terms of model fit as compared to our basic Unstructured Regional Effects (URE) Model.

The coefficient estimates of the individual-level variables are very robust.²¹ Not surprisingly, the county-level variables show some variation across the two models, which is due to the different assumptions upon the spatial structure. Therefore, in models 3 and 4, we improve upon the modeling of the spatial representation before interpreting these results.²²

The map in Figure 3 shows the regional effects for **Model 2** (ICAR model based on adjacency).

[Insert Figure 3 about here]

First, it should be noted that the strength of the regional effects is even greater than in the URE-Model (Model 1). Here, the regional effects vary between -3.3 and 3.4 as compared to -2.4 and 2.7 in the first model.

²¹ These estimates are available upon request.

²² A possible source of this variation could be omitted county-level variables, which introduce spatial correlation. Cressie (1993) points out that modeling spatial autocorrelation can mitigate this problem.

Second, the spatial pattern displayed in figure 3 is very similar to the spatial pattern produced by Model 1, but considerably smoother. Both maps show several large clusters and suggest that there exists a pattern in the spatial distribution:

Thirdly, in West Germany, unobserved but systematic, county-level effects have a strong positive impact on its residents' health. Contrarily, in East Germany, county-level effects have a strong negative impact on residents' health. Remember that we control for a rich set of socio-economic individual background characteristics such as age, gender, marital, and employment status as well as health-behavior and the degree of health care utilization. In addition, ten county-level covariates net out persistent differences across counties due to unemployment, urbanization and population density. Given this modeling approach, it is surprising and staggering that we still find a clear East-West health pattern, 20 years after German reunification.

We use the Deviance Information Criterion (DIC) introduced by Spiegelhalter et al. (2002) to compare Model 1 and Model 2. The DIC is a criterion for model selection similar to the Akaike Information Criterion (AIC) or Bayesian Information Criterion (BIC), i.e., it trades off model fit²³ and complexity²⁴. As with AIC and BIC, the model with the lowest DIC should be preferred. In contrast to AIC and BIC, the DIC is valid for hierarchical models and can be easily computed as a by-product of MCMC methods. The results indicate that the ICAR-Model 2 performs less well, both in terms of model fit and complexity. This is no surprise: we impose additional structural assumptions; hence we obtain more complex models.

4.3. Model 3: Convolution Prior (CP)

Now, we run Model 3 using a convolution prior as described in section 3.1.3. This model has the advantage that it can distinguish between unstructured (i.e., random) regional effects and spatially dependent regional effects. If we plot the overall regional effect, the resulting maps look very similar to the ICAR-Model (Model 2). However, the values of the spatially dependent random effects appear to be significantly smoother.

If we compare the model fit via the DIC (see Appendix, Table 5), we find that the model fit for the Convolution Prior (CP) Model is almost the same as for the ICAR Model (Model 2). According to the rule of thumb proposed by Spiegelhalter et al. (2002), only a DIC difference of 10 should be considered as significant. Furthermore, the CP-Model produces smoother estimate, which is why we prefer it to the ICAR-Model. Figure 4 illustrates the results of Model 3.

[Insert Figure 4 about here]

²³ Model fit is measured through the posterior expected deviance.

²⁴ Complexity refers to the effective number of parameters calculated as the difference between posterior expected deviance and the deviance at the posterior means.

Figure 4 presents a map of the spatially dependent part of the regional effect φ_S . We observe a prominent cluster of high values in the northwest (Lower Saxony, Schleswig-Holstein) and clusters of lower values in the east (Bavaria, Thuringia and Brandenburg). This confirms our findings above and reinforces a clear East-West German spatial health pattern that is likely to be a legacy of the 40 year long division of Germany.

[Insert Figure 5 about here]

Figure 5 shows a LISA cluster map for Model 3. The counties with significant local I are color-coded with respect to the type of spatial correlation. We identify three clusters of high values, mainly in the northwestern part of Germany and three clusters of low values in the eastern part. This again supports the notion of an East-West trend in the spatially dependent component. Deviations from this trend in Eastern Bavaria and Western Mecklenburg-West Pomerania are likely due to spillover effects.

The finding of low regional health patterns in regions with former communist governments during the Cold War is consistent with the existing literature. For example Baltagi et al. (2012) and Bonneux et al. (2010) report significantly lower life expectancy in Eastern Europe, specifically Hungary, Czech Republic, Slovak Republic and Poland, when compared to Western Europe. In addition, Baltagi, Moscone and Tosetti (2012) find significant spillover effects in the healthcare production process across neighboring countries. Treurniet et al. (2004) report higher avoidable mortality rates for Hungary and the Czech Republic as compared to Western Europe between 1980 and 1997. However, we are not aware of any paper that spatially models and compares the health effects within Germany—a country with a common history, language, and culture that had been divided for 40 years.

4.4 Model 4: Spatiotemporal Convolution Prior (SCP)

Finally we estimate the Spatiotemporal Convolution Prior (SCP) Model in order to investigate changes in the spatial pattern over time. The results are in Table 4. First, the parameter estimates of the individual covariates are robust when we compare them to the findings from our first model, the URE-Model in Table 2. Second, most of the county-level covariates are not significantly different from zero and are similar in size to the findings from our first URE-Model in Table 2.

Third, clearly, the county-level unemployment rate shows by far the strongest association with individual-level health. In the SCP-Model, a 1 percentage point higher unemployment rate is associated with a reduction in health by 0.07 points. This equals an age effect of 1 year. This point estimate is highly significant at the 5 percent level. The second most important county-level correlate of individual health is the average household income in the county, which is significant at the 10 percent level. Inhabitants of counties with a €1,000 higher household income have on average 0.8 points more on their SF12 scale. This equals an age effect of 10 years.

[Insert Table 4 about here]

Model 4 also includes a linear time trend instead of aggregated time fixed effects. This trend is highly significant and negative (-0.17), implying that over the five year period from 2006 to 2010 health among SOEP respondents decreased by an average of -0.85 points. Remarkably, this time trend is more than twice as high than the effect of an additional year of age. There are several possible explanations for this trend, e.g. it could be due to a change in the underlying population (i.e., the panel attrition effect) or ageing effects, which are not captured by the age polynomial in the model. Another explanation is that the economic shock caused by the financial crisis led to increased levels of stress and worries. A separate analysis of the physical and mental health component of our dependent variable shows that the time trend is much stronger for mental health than for physical health, which yields some evidence for this explanation.²⁵

A comparison of the DIC of this model with the previous models shows that the SCP-Model offers a much better model fit than the purely spatial models (see Appendix, Table 5). This indicates that there are considerable space-time interactions. The gain in model fit more than compensates for the higher complexity of the model. Surprisingly, the model fit is even better than in the Unstructured Regional Effects Model. This implies that our Model 4 should be preferred among all estimated models.

Figure 6 illustrates the spatial patterns for each year. Note that only the spatially dependent part of the regional effect is displayed. First of all, the size the effects varies across years, but largely overlap: in 2006, the values fall between -1.5 and 2 , in 2008, they fall between -1.2 and 0.8 and in 2010, they fall between $-$

1.4 and 1 .

[Insert Figure 6 about here]

Second, the spatial pattern itself varies slightly across time. However, since the majority of the regional effects are not significant this is probably random variation. In those counties with an especially small sample size, the results are highly sensitive to panel attrition. More important, the clusters with significant spatial dependence, e.g., the regions around Hanover, Rhineland-Palatinate, Thuringia and Bavaria, are stable across the years. Maps of the pseudo p -values of the regional effects (φ_{st}) show that the regional effect in these clusters is significant in all years, whereas most of the effects that change across the years are insignificant.

Finally, to explicitly test for the East-West-German health pattern, we include an East-West dummy variable in Model 4. The dummy is “1” for counties in pre-1989 communist East Germany and “0” for counties in capitalist West Germany. The highly significant coefficient estimate equals -0.69 and represents one of the strongest health predictors. Its impact is comparable to an age impact of more than half a standard deviation or up to 9 life years.

²⁵ The results are available upon request.

5 Conclusion

In this paper we combine representative individual-level household panel data and register county-level data to model and estimate spatial health effects on individual health. Methodologically, we make use of hierarchical models and Bayesian methods.

In a first step, we examine regional effects for spatial associations using adjacency-based, distance-based and Nearest-Neighbor-based definitions of neighborhoods. In all cases, we find a highly significant spatial dependency of individual health. In the next step, we impose structural ICAR assumptions onto our models and re-estimate them, using an ICAR prior based on adjacency and a convolution prior based on adjacency. For both models we find a trade-off between model fit and smoother estimates. Although most of our ten county-level indicators—e.g., population density or unemployment rates—are not individually significant at conventional statistical levels, jointly they clearly improve the fit of our models. Interestingly, we find the model fit for our spatiotemporal model with separate regional effects for each year to be much better than the model fit for our basic Regional Effects Model without structural assumptions.

In terms of content, we find that the regional impact on residents' health is systematic and strong. The general county-level predictor has the strongest impact among all individual- and county-level predictors considered. The regional impact is comparable to an age effect on health of up to 30 years: this means that—after considering a rich set of socio-demographic background information for an average respondent at the age of 50—the regional impact alone may shift a respondent's health either down to the health of a 20 year old or up to the health of a 80 year old.

Interestingly and surprisingly, even 20 years after the German reunification, we still find a clear East-West spatial health pattern. This finding could be interpreted as the long-lasting health effect of the 40 year long division of Germany into a Communist and a Capitalist Part. The long-lasting legacy of Communism on health equals an age effect of up to 9 life years.

Although most outcomes variables of interest reveal a clear spatially-dependent structure, applied researchers are reluctant to model these spatial dependencies explicitly. Using Germany and a quasi-objective and continuous measure of population health as an example, this paper provides a framework for applied researchers on how to use Hierarchical Bayes Models and estimation techniques. A next step would be to establish an identification strategy in order to test whether these regional effects have a causal interpretation or arise due to selection processes or omitted variable bias. Clearly, more research on and applications of modeling techniques for spatial patterns would be fruitful.

6 References

- Andersen, H. H., Mühlbacher, A., Nübling, M., Schupp, J., and Wagner, G. G. (2007): Computation of standard values for physical and mental health scale scores using the SOEP version of sf12v2. *Journal of Applied Social Science Studies (Schmollers Jahrbuch)*, 127, 171-182.
- Anselin, L. (1995): Local indicators of spatial association – LISA. *Geographical Analysis*, 27, 93-115.
- Anselin, L. (2005): *Exploring Spatial Data with GeoDa™: A Workbook*. Available at <http://www2.geog.ucl.ac.uk/~jessletzbichler/GEOG2003/geodaworkbook.pdf>, last accessed on May 1, 2012.
- Arbia, G. (2006): *Spatial Econometrics: Statistical Foundations and Applications to Regional Convergence*. Springer. Berlin Heidelberg.
- Baltagi, B. H., and Moscone, F., (2010): Health care expenditure and income in the OECD reconsidered: Evidence from panel data., *Economic Modelling*, 27(4), 804-811.
- Baltagi, B.H., Moscone, F. and Tosetti, E. (2012): Medical technology and the production of health care. *Journal of Empirical Economics*, 42, 395-411.
- Banerjee, S., Carlin, B. P. and Gelfand, A. E. (2004): *Hierarchical Modeling and Analysis for Spatial Data*. Chapman & Hall/CRC. Boca Raton.
- Besag, J. (1974): Spatial interaction and the statistical analysis of lattice systems (with Discussion). *Journal of the Royal Statistical Society*, B 36, 192-236.
- Besag, J. and Kooperberg, C. (1995): On conditional and intrinsic autoregressions. *Biometrika*, 82(4),733-746.
- Besag, J., York, J.C. and Mollié, A. (1991): Bayesian image restoration with two applications in spatial statistics (with discussion). *Annals of the Institute of Statistical Mathematics*, 43, 1-59.
- Bonneux, L.G., Huisman, C. and de Beer, J.A. (2010): Mortality in 272 European regions, 2002–2004. An update. *European Journal of Epidemiology*, 25, 77-85.
- Brooks, S.P. and Gelman, A. (1998): General Methods for Monitoring Convergence of Iterative Simulations. *Journal of Computational and Graphical Statistics*, 7(4), 434-455.
- Browning, C.R., Cagney, K.A. and Wen, M. (2003): Explaining variation in health status across space and time: implications for racial and ethnic disparities in self-rated health. *Social Science & Medicine*, 57, 1221-1235.
- Bryk, A.S. and Raudenbush, S.W. (1992): *Hierarchical Linear Models: Applications and Data Analysis Methods*. Sage Publications. Newbury Park. London. New Delhi.
- Burchardi, K.B. and Hassan, T.A. (2011): The economic impact of social ties: evidence from German reunification. NBER Working Papers 17186.
- Coneus, K. and Spieß, K.C. (2012): Pollution exposure and child health: Evidence for infants and toddlers in Germany. *Journal of Health Economics*, 31, 180-196.
- Congdon, P.D. (2010): *Applied Bayesian Hierarchical Methods*. Chapman & Hall/CRC. Boca Raton. London. New York.
- Costa-Font, J. and Pons-Novell, J. (2007): Public health expenditures and spatial interactions in a decentralized national health system. *Health Economics*, 16, 291–306.
- Cressie, N. A.C. (1993): *Statistics for Spatial Data*. Wiley. New York.
- Cressie, N. A.C. and Wikle, C. K. (2011): *Statistics for Spatio-Temporal Data*. Wiley, Hoboken.

- Dartmouth Atlas of Health Care (2012), available at www.dartmouthatlas.org, last accessed May 2, 2012.
- Durlauf, S. N. (2003): Neighbourhood effects, in: J. V. Henderson / J.-F. Thisse(eds.), *Handbook of Regional and Urban Economics*, edition 1, volume 4, chapter 50, 2173-2242.
- Federal Institute for Research on Building, Urban Affairs and Spatial Development (2012): *INKAR 2011. Indicators and Maps on Spatial Development*, CD-ROM.
- Felder, S. and Tauchmann, H. (2011): Federal state differentials in the efficiency of health production in Germany: an artifact of spatial dependence? *The European Journal of Health Economics*, forthcoming.
- Frick, J. R. and Ziebarth, N.R. (2012): “Welfare-Related Health Inequality: Does the Choice of Measure Matter?” *The European Journal of Health Economics*, forthcoming.
- Fuchs, V.R., McClellan, M.B. and Skinner, J.S. (2004): Area differences in utilization of medical care and mortality among U.S. elderly. In: Wise, D.A. (ed.): *Perspectives on the economics of aging*, University of Chicago Press, Chicago, 367-413.
- Gelman, A. and Robert, C.P. (2012): “Not only defended but also applied”: The perceived absurdity of Bayesian inference. *The American Statistician*, forthcoming.
- Gelman, A. and Rubin, D.B. (1992): Inference from Iterative Simulation Using Multiple Sequences. *Statistical Science*, 7(4), 457-511.
- Knies, G. and Spieß, C.K. (2007): Regional data in the German Socio-Economic Panel Study (SOEP), Data Documentation 17, DIW Berlin, German Institute for Economic Research, available at <http://ideas.repec.org/e/pkn24.html>, last accessed May 2, 2012.
- Knies, G., Burgess, S. and Propper, C. (2008): Keeping up with the Schmidt’s – an empirical test of relative deprivation theory in the neighbourhood context. *Journal of Applied Social Science Studies*, 128(1), 75-108.
- Hox, J. (2002): *Multilevel Analysis. Techniques and Applications*. Lawrence Erlbaum Associates, Publishers. Mahwah (New Jersey), London.
- Mollié, A. (1996): Bayesian mapping of disease. In A. Lawson, A. Biggeri, D. Bohning, E. Lesaffre, J.F. Viel, and R. Bertollini, (editors): *Markov Chain Monte Carlo in Practice: Interdisciplinary Statistics*. Chapman and Hall. London. 359-379.
- Moran, P.A.P. (1950): Notes on continuous stochastic phenomena. *Biometrika*, 37(1), 17-23.
- Moscone, F., Knapp, M. and Tosetti, E. (2007): Mental health expenditure in England: A spatial panel approach. *Journal of Health Economics*, 26(4), 842-864.
- Moscone, F. and Tosetti, E. (2010): Health expenditure and income in the United States. *Health Economics*, 19(12), 385-1403.
- Mueller, I., Vounatsou, P., Allen, B.J. and Smith, T. (2001): Spatial patterns of child growth in Papua New Guinea and their relation to environment, diet, socio-economic status and subsistence activities. *Annals of Human Biology*, 28, 263-280.
- Ocaña-Riola, R. and Mayoral-Cortés, J.M. (2010): Spatio-temporal trends of mortality in small areas of Southern Spain. *BMC Public Health*, 10:26
- OECD Project on Income Distribution and Poverty (2009): What are equivalence scales?. Available online at <http://www.oecd.org/dataoecd/61/52/35411111.pdf>, last accessed on April 30, 2012.
- Owens, G. M. (2008): Gender differences in health care expenditures, resource utilization, and quality of Care. *Journal of Managed Care Pharmacy*, 14(3), 2-6.

- Pesaran, M.H.(2006): Estimation and inference in large heterogeneous panels with a multifactor error structure. *Econometrica*, 74(4), 967-1012.
- RAND (1995): *RAND. User's Manual for the Medical Outcomes Study (MOS) Core Measures of Health-Related Quality of Life*. (by Hays, R. D., Sherbourne, C. D, Mazel, R.). RAND Monograph Report, Santa Monica, Available at: http://www.rand.org/pubs/monograph_reports/MR162.html, last accessed February 2, 2012
- Riva, M.; Gauvin, L. and Barnett, T.A. (2007): Toward the next generation of research into small area effects on health: a synthesis of multilevel investigations published since July 1998. *Journal of Epidemiology & Community Health*, 61, 853-861.
- Robert, C.P. and Casella, G. (2005). *Monte Carlo Statistical Methods*. 2nd Edition. Springer. New York.
- Spiegelhalter, D.J., Best, N., Carlin, B.P. and van der Linde, A. (2002): Bayesian measures of model complexity and fit (with discussion). *Journal of the Royal Statistical Society, Series B*, 64, 583-639.
- Sundmacher, L., Gaskins, M., Hoffman, K, and Busse R. (2012). Spatial distribution of avoidable cancer mortality. *Journal of Public Health*, 20(3), 279-288.
- Statistical Offices of the Federation and the Federal States of Germany (2012): Regionaldatenbank Deutschland, available at www.regionalstatistik.de, last accessed May 14, 2012.
- Tampubolon,G., Subramanian, S.V. and Kawachi , I.(2011). Neighbourhood social capital and individual self-rated health in Wales. *Health Economics*. Forthcoming.
- Teuniet, H.F., Boshuizen, H.C. and Harteloh, P.P.M. (2004): Avoidable mortality in Europe (1980-1997): a comparison of trends. *Journal of Epidemiology & Community Health*, 58, 290-295.
- Voigtländer, S., Berger, U. and Razum, O. (2010): The impact of regional and neighbourhood deprivation on physical health in Germany: a multilevel study. *BMC Public Health*, 10: 403.
- Wennberg, J. and Gittelsohn, A. (1973): Small area variations in health care delivery, *Science*, 182 (4117), 1102-1108.
- Wennberg, J. E.,Fisher, E. and Skinner, J.S. (2002): Geography and the debate over Medicare reform. *Health Affairs* 21 (2, Suppl. Web Exclusive): W96–114.
- Wagner, G. G., Frick, J. R., and Schupp, J. (2007):The German Socio-Economic Panel study (SOEP) - evolution, scope and enhancements. *Journal of Applied Social Science Studies (Schmollers Jahrbuch)*, 127, 139-169.
- Ziebarth, N.R.and M.M Grabka (2009): In Vino Pecunia? The Association Between Beverage-Specific Drinking Behavior and Wages. *Journal of Labor Research*, 30(3), 219-244.

Tables and Figures

Table 1: Summary Statistics

Variable	Mean	Standard Deviation	Min	Max	N
<u>Dependent Variables</u>					
SF12 - generic health measure	49.73	7.12	22.33	61.62	54,683
<u>A: Individual Characteristics</u>					
<i>Demographics</i>					
Age	49.31	17.58	17	100	54,683
Number of children under 14	0.36	0.75	0	8	54,683
Sex	0.48	0.50	0	1	54,683
Marital status	0.60	0.49	0	1	54,683
<i>Education & Labor Market Participation</i>					
Monthly household income - equivalence scale	1,776	1,259	0	66,666	54,683
Not working	0.43	0.50	0	1	54,683
University degree	0.21	0.41	0	1	54,683
No vocational training	0.23	0.42	0	1	54,683
<i>Health-related Behavior</i>					
No alcohol consumption	0.13	0.34	0	1	54,683
Regular alcohol consumption	0.17	0.38	0	1	54,683
Smoking	0.28	0.45	0	1	54,683
Body Mass Index (BMI)	25.96	4.60	12.59	73.05	54,683
<i>Health Care Utilization</i>					
Number of stays in hospital during the past year	0.16	0.62	0	42	54,683
Number of doctor's visits during the past year	9.80	15.19	0	396	54,683

...continued on next page

...Table 1 continued

B: Regional Characteristics

Area size	890.35	722.00	35.71	5,810.77	1,199
Population density	521.64	676.29	39.94	4,257.16	1,199
Proportion of area used for settlement and transport	20.77	15.48	4.60	76.10	1,199
Proportion of area suitable for recreation	1.98	2.39	0.20	15.00	1,199
Cars per 1000 inhabitants	524.85	57.09	316.10	880.60	1,199
Inhabitants per physician	685.19	183.48	255.91	119.70	1,199
Unemployment quota	9.22	4.13	2.49	21.30	1,199
Percentage of immigrants	7.25	4.51	0.70	25.34	1,199
Average available income per month	1,531.29	202.25	1,145.61	2,544.17	1,199
Number of overnight stays per inhabitant	4.73	5.65	0.43	41.06	1,199

Source: SOEP v27, own calculations.

Table 2: Coefficient estimates for the covariates in the Unstructured Regional Effects model

Variable	Mean	Rescaled	Significance	S.E	2.5% Percentile	97.5% Percentile
<i>Individual Variables</i>						
Age	-1.29151	-0.07346	***	0.22244	-1.72952	-0.85609
Squared age	-0.76136	-0.00043	***	0.21964	-1.18953	-0.33074
Male gender	1.15740	-	***	0.07377	1.01200	1.30000
Marital Status	0.22080	-	***	0.03735	0.14725	0.29495
Number of children	0.03340	0.04454		0.03312	-0.03104	0.09913
Household income	0.40193	0.00032	***	0.03018	0.34199	0.45921
Not working	-0.22076	-	***	0.06465	-0.34846	-0.09050
University Degree	0.92858	-	***	0.09048	0.75165	1.10853
No vocational training	-0.49347	-	***	0.08610	-0.65895	-0.32355
No alcohol consumption	-1.15268	-	***	0.07962	-1.30753	-0.99750
Regular alcohol consumption	0.26792	-	***	0.07164	0.12904	0.40670
Smoking	-0.93175	-	***	0.06899	-1.07200	-0.79669
Number of hospital stays	-0.39349	-0.63467	***	0.02196	-0.43610	-0.35105
Number of doctor's visits	-1.54294	-0.10158	***	0.02426	-1.59000	-1.49500
BMI	-0.58438	-0.12704	***	0.03452	-0.65345	-0.51670
<i>Regional Variables</i>						
Population density	-0.00032	0.00000		0.00038	-0.00108	0.00043
Area size	0.03542	0.00005		0.10527	-0.16683	0.23832
Unemployment quota	-0.14640	-0.03545	*	0.07912	-0.30145	0.00865
Overnight stays per inhabitant	0.01856	0.00328		0.06762	-0.11220	0.15130
Area used for settlement and transport	0.43430	0.02806		0.38702	-0.32388	1.17653
Average household income	-0.03776	-0.00019		0.10221	-0.24136	0.16183
Percentage of migrants	0.17502	0.03881		0.15464	-0.13172	0.49045
Area used for recreation	0.16564	0.06930		0.22082	-0.27477	0.59987
Cars per 1,000 Inhabitants	0.26409	0.00463	*	0.13338	0.00578	0.52916
Inhabitants per physician	-0.05378	-0.00029		0.09983	-0.25253	0.13605

...continued on next page

...Table 2 continued

Time effects	Mean	Rescaled	Significance	S.E	2.5% Percentile	97.5% Percentile
Dummy 2006	0.50230	-	***	0.08577	0.33470	0.66996
Dummy 2008	0.44034	-	***	0.04879	0.34605	0.53496
Constant	49.34192	-		0.32177	48.71000	49.98525
σ_C	4.64511	-		0.03025	4.58600	4.70500
σ_B	1.15781	-		0.06422	1.03647	1.28952
σ	3.87708	-		0.01541	3.84700	3.90700

Source: SOEP v27, own calculations. Notes: Number of Observations N=54,683. Number of draws from the Posterior Distribution D=4,500. Column 1 gives the posterior mean (i.e. the point estimate) for the 15 individual predictors, the ten county-level covariates, the aggregate time effects, the constant and the standard deviations of the dependent variable, σ , the individual effects, σ_C , and the regional effects, σ_B . Keep in mind that the non-dummy variables are standardized and therefore the measurement units are standard deviations. Column 2 gives a rescaled coefficient estimate, i.e. the effect size corresponding to the original measurement unit (left blank for dummy variables). Column 3 gives the posterior probability that the parameter has a different sign than the point estimate: ***=<0.1%, **=<1%, *<5% and .=<10%. This should not be confused with p-values and Null Hypothesis Significance Tests. Column 4 gives the standard errors. Columns 5-6 give the 95%-Equal Tail Credibility Interval around the median.

Table 3: Moran's I for different definitions of neighborhood

Weights	Links	I	p-value
first-order adjacency	2,084	0.14890	0.00020
second-order adjacency	4,880	0.03630	0.03570
third-order adjacency	7,580	0.01840	0.11060
threshold distance - 50 km	3,694	0.11260	0.00010
threshold distance - 150 km	27,092	0.03900	0.00060
threshold distance - 68.2 km	8,410	0.08200	0.00010
3 nearest neighbors	1,201	0.17540	0.00010
5 nearest neighbors	2,000	0.12920	0.00010
7 nearest neighbors	2,798	0.12480	0.00010

Source: SOEP v27, own calculations; Moran's I for nine different weight matrices. Column 1 gives the number of non-zero links, i.e., spatial connections between sites. Column 2 gives the estimate of Moran's I and Column 3 the corresponding empirical p-value, which is derived by a Monte-Carlo permutation test with 10,000 permutations.

Table 4: Coefficient estimates for the Spatiotemporal Convolution Prior Model (Model 4)

Variable	Mean	Rescaled	Significance	S.E.	2.5% Percentile	97.5% Percentile
<i>Individual Variables</i>						
Age	-1.30368	-0.07416	***	0.22380	-1.75053	-0.86091
Squared age	-0.74777	-0.00042	***	0.22231	-1.18600	-0.31034
Male gender	1.15677	-	***	0.07381	1.01000	1.30300
Marital Status	0.22165	-	***	0.03755	0.14700	0.29405
Number of children	0.04375	0.05833		0.03393	-0.02416	0.10975
Household income	0.41487	0.00033	***	0.03053	0.35500	0.47325
Not working	-0.21555	-	***	0.06593	-0.34586	-0.08631
University Degree	0.94687	-	***	0.09183	0.76514	1.12452
No vocational training	-0.48296	-	***	0.08607	-0.65086	-0.31750
No alcohol consumption	-1.13895	-	***	0.07856	-1.29653	-0.98733
Regular alcohol consumption	0.30181	-	***	0.07053	0.16500	0.44119
Smoking	-0.93499	-	***	0.06898	-1.07100	-0.80205
Number of hospital stays	-0.39249	-0.63305	***	0.02193	-0.43665	-0.35035
Number of doctor's visits	-1.53402	-0.10099	***	0.02461	-1.58100	-1.48500
BMI	-0.59272	-0.12885	***	0.03467	-0.66015	-0.52529
<i>Regional Variables</i>						
	Mean	Rescaled	Significance	S.E.	2.5% Percentile	97.5% Percentile
Population density	-0.00018	0.00000		0.00023	-0.00063	0.00027
Area size	0.03523	0.00005		0.07043	-0.10205	0.17160
Unemployment quota	-0.29205	-0.07071	**	0.09392	-0.47055	-0.10545
Overnight stays per inhabitant	0.03732	0.00661		0.04664	-0.05365	0.12990
Area used for settlement and transport	0.40928	0.02644	.	0.26319	-0.11871	0.91966
Average household income	0.16238	0.00080	*	0.08126	0.00475	0.32265
Percentage of migrants	-0.00957	-0.00212		0.11138	-0.23010	0.21161
Area used for recreation	0.06262	0.02620		0.13965	-0.21401	0.34126
Cars per 1,000 Inhabitants	0.08990	0.00157		0.09331	-0.09257	0.27530
Inhabitants per physician	-0.03766	-0.00021		0.07406	-0.18446	0.10655
<i>Time effects</i>						
Linear Trend	-0.17387	-	***	0.02614	-0.22490	-0.12170
Constant	398.63844	-	***	52.47963	294.03749	501.10501
σ_c	4.69938	-	***	0.03036	4.63900	4.75800
σ_B / σ_ω				<i>Results for each year available upon request</i>		
σ_φ				<i>Results for each year available upon request</i>		
σ	3.85295	-	***	0.01557	3.82300	3.88300

Source: SOEP v27, own calculations; number of observations N=54,528. Number of draws from the posterior distribution D=4,500. Column 1 gives the posterior mean (i.e. the point estimate) for the 15 individual predictors, the ten county-level covariates, the aggregate time effects, the constant and the standard deviations of the dependent variable (σ) and the individual effects (σ_B). Keep in mind that the non-dummy variables are standardized and therefore the measurement units are standard deviations. Column 2 gives a rescaled coefficient estimate, i.e., the effect size corresponding to the original measurement unit (left blank for dummy variables). Column 3 gives the posterior probability that the parameter has a different sign than the point estimate: ***=<0.1%, **=<1%, *=<5% and .=<10%. This should not be confused with p-values and null hypothesis significance tests (see 3.2). Column 4 gives the standard errors. Columns 5 and 6 give the 95%-Equal Tail Credibility Interval around the median.

Figure 1: Quantile map for Model 1 (Unstructured Regional Effects)

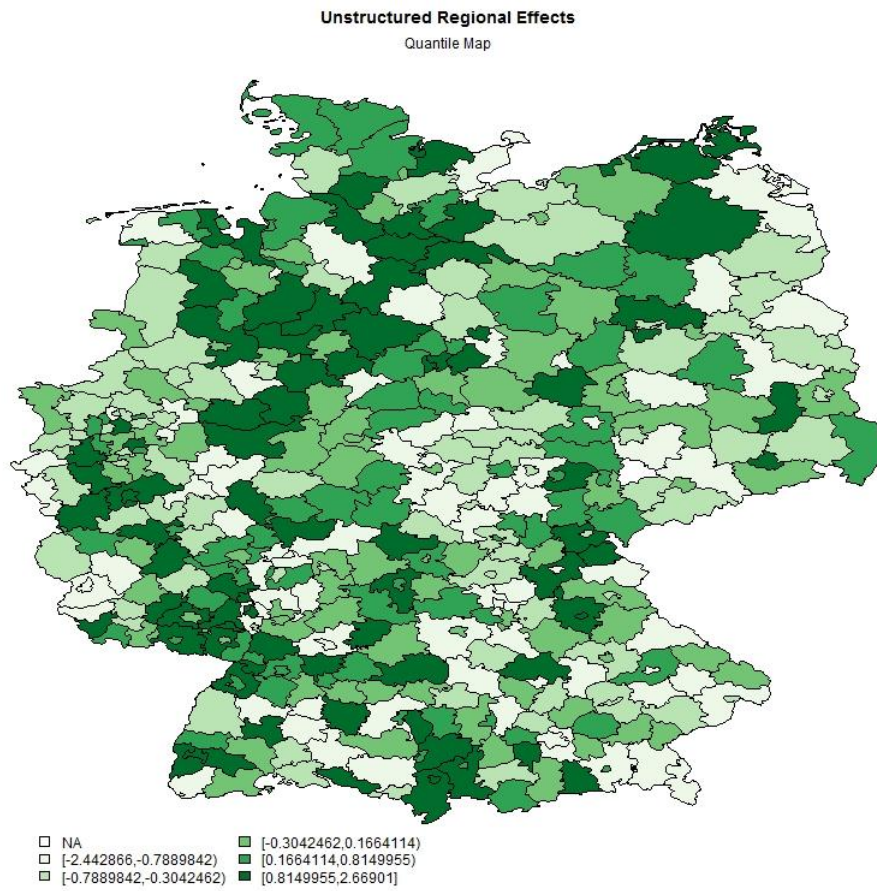


Figure 2: LISA Cluster map for Model 1 (Unstructured Regional Effects)

Unstructured regional effects

LISA Cluster map

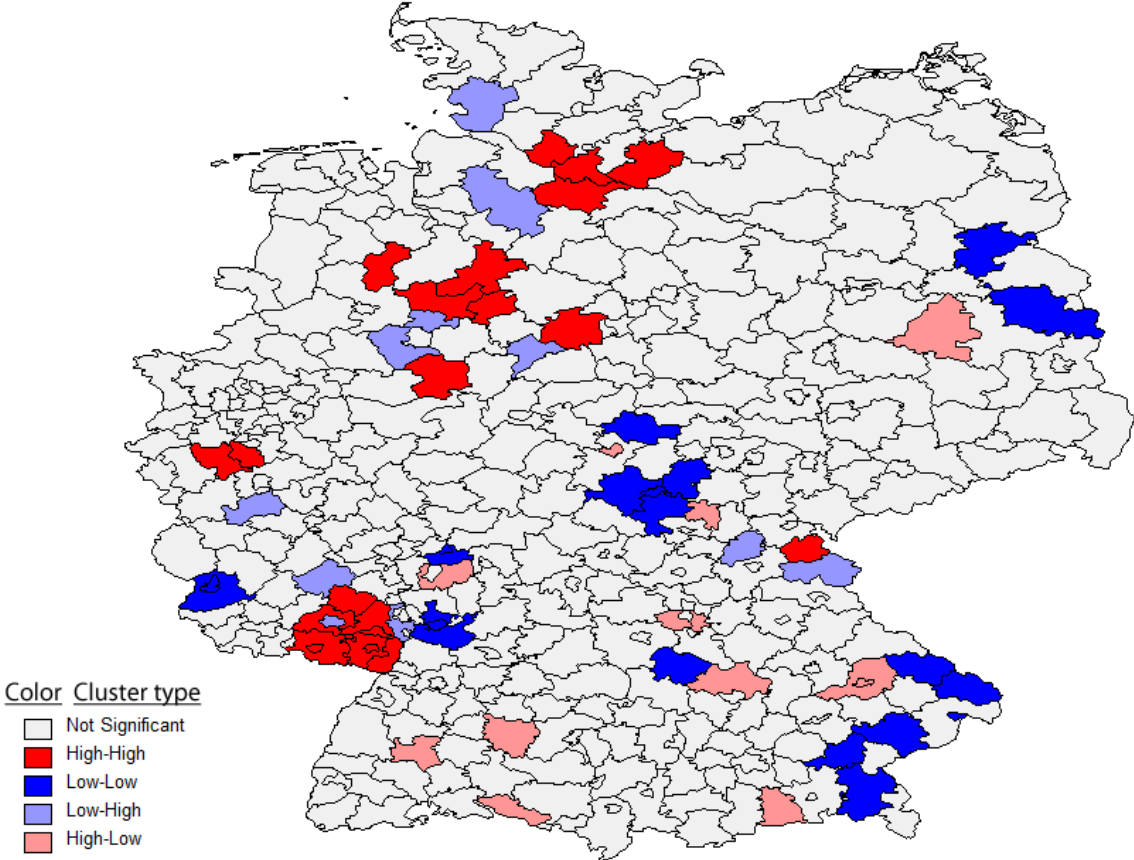


Figure 3: Quantile map for Model 2 (ICAR Regional Effects)

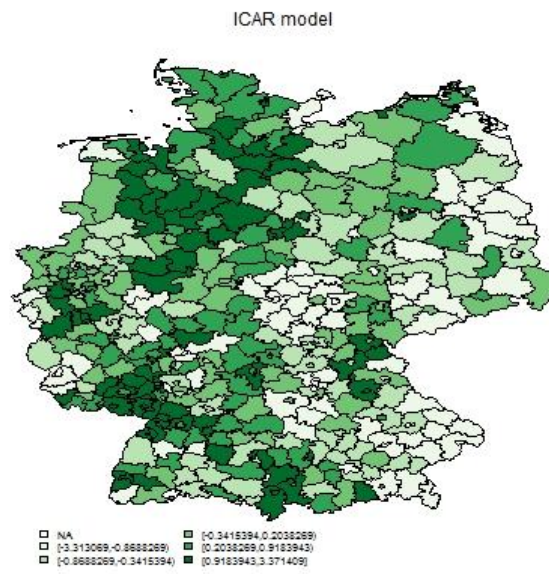


Figure 4: Quantile map for Model 3 (Convolution Prior)

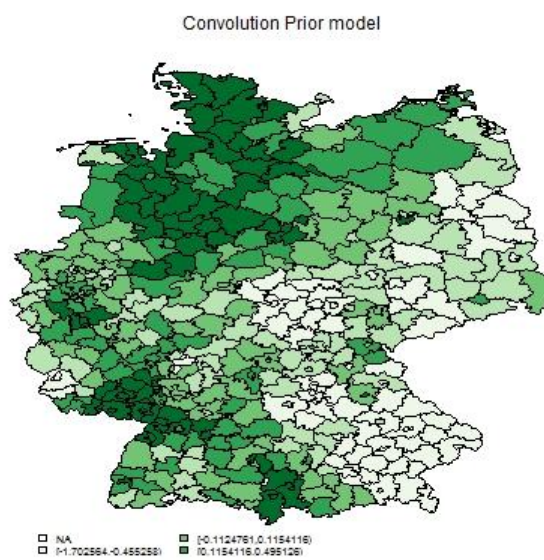


Figure 5: LISA Cluster map for Model 3 (Convolution Prior)

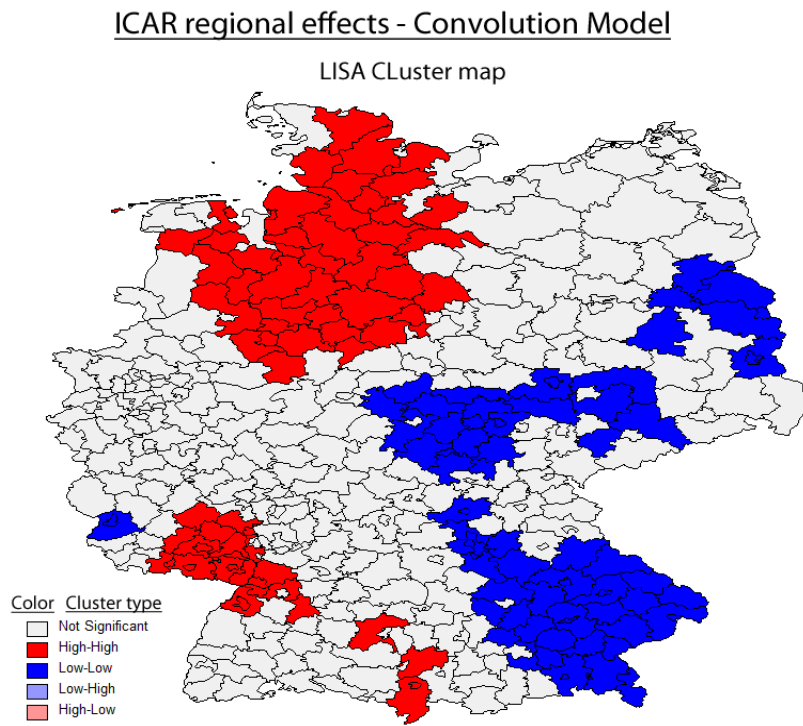
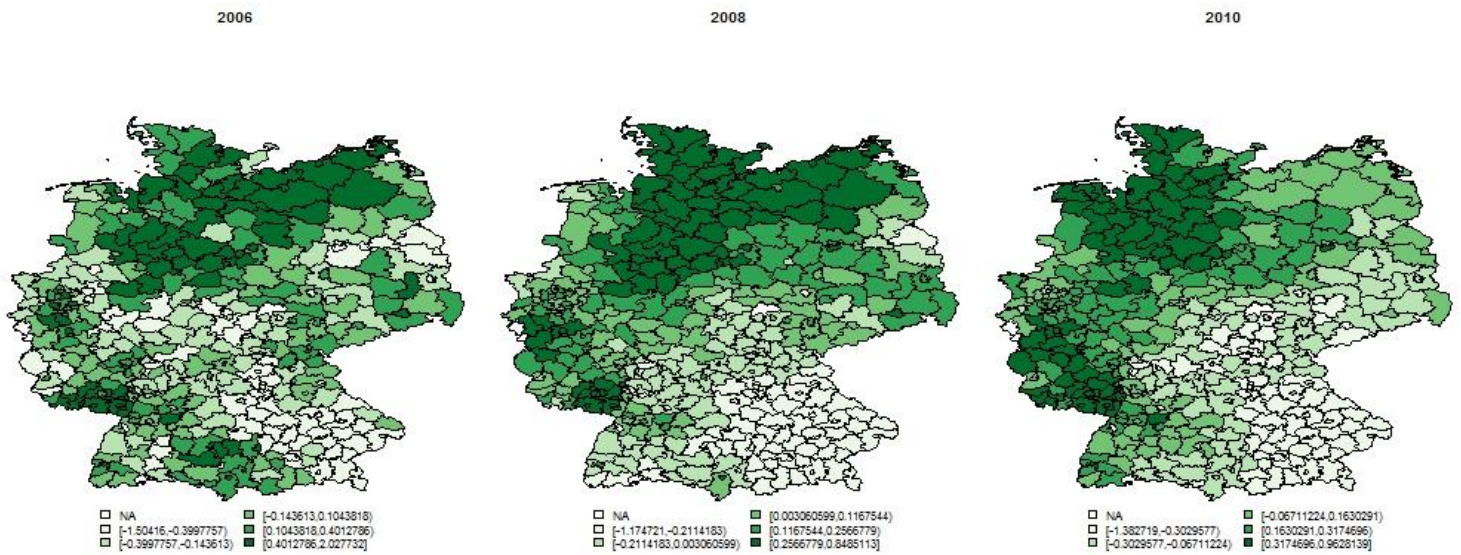


Figure 6: Quantile maps for Model 4 (Spatiotemporal Convolution Prior)



Appendix

Table 5: Deviance Information Criterion (DIC) for 4 candidate models

Model	Dbar	pD	DIC
Unstructured Regional Effects	303,382	17,586.4	320,969
ICAR Regional Effects	304,383	18,563.4	322,946
Convolution Prior model	304,378	18,557.0	322,935
Spatiotemporal Convolution Prior model	302,170	18,355.8	320,526

Source: SOEP v27, own calculations.