

BAYESIAN STRUCTURED ADDITIVE DISTRIBUTIONAL REGRESSION WITH AN APPLICATION TO REGIONAL INCOME INEQUALITY IN GERMANY

BY NADJA KLEIN^{*,1}, THOMAS KNEIB^{*,1},
STEFAN LANG^{†,2} AND ALEXANDER SOHN^{*}

Georg-August-University Göttingen and University of Innsbruck†*

We propose a generic Bayesian framework for inference in distributional regression models in which each parameter of a potentially complex response distribution and not only the mean is related to a structured additive predictor. The latter is composed additively of a variety of different functional effect types such as nonlinear effects, spatial effects, random coefficients, interaction surfaces or other (possibly nonstandard) basis function representations. To enforce specific properties of the functional effects such as smoothness, informative multivariate Gaussian priors are assigned to the basis function coefficients. Inference can then be based on computationally efficient Markov chain Monte Carlo simulation techniques where a generic procedure makes use of distribution-specific iteratively weighted least squares approximations to the full conditionals. The framework of distributional regression encompasses many special cases relevant for treating non-standard response structures such as highly skewed nonnegative responses, overdispersed and zero-inflated counts or shares including the possibility for zero- and one-inflation. We discuss distributional regression along a study on determinants of labour incomes for full-time working males in Germany with a particular focus on regional differences after the German reunification. Controlling for age, education, work experience and local disparities, we estimate full conditional income distributions allowing us to study various distributional quantities such as moments, quantiles or inequality measures in a consistent manner in one joint model. Detailed guidance on practical aspects of model choice including the selection of several competing distributions for labour incomes and the consideration of different covariate effects on the income distribution complete the distributional regression analysis. We find that next to a lower expected income, full-time working men in East Germany also face a more unequal income distribution than men in the West, *ceteris paribus*.

Received September 2013; revised March 2015.

¹Support by the German Research Foundation via the research training group 1644 and the projects KN 922/4-1/2 is gratefully acknowledged.

²Supported by the Oesterreichische Nationalbank, Anniversary Fund, project number: 15309.

Key words and phrases. Generalised additive models for location, scale and shape, income distribution, iteratively weighted least squares proposal, Markov chain Monte Carlo, semiparametric regression, wage inequality.

1. Introduction. The analysis of determinants of labour incomes has a long tradition in economics, dating back at least to [Mincer \(1974\)](#). His classical wage equation includes potential labour market experience as well as years of education as the most important determinants of human capital which then translates into expected income [[Lemieux \(2006\)](#)]. Additional possible determinants include age, actually realised labour market experience, gender, regional information concerning the residence of employees, or area of employment. One considerable restriction of most analyses conducted so far is their sole focus on the expected income given covariates, that is, the conditional mean. In some cases, distributions are required, for example, for inequality decomposition or to account for incomplete information due to truncation or censoring. Then, the (log-)normal distribution [[Greene \(2008\)](#), Chapter 19, [Morduch and Sicular \(2002\)](#)] is often implicitly considered (again with regression effects only on the mean) or one reverts to local analyses by means of quantile regression [[Autor, Katz and Kearney \(2008\)](#), [Galvao, Lamarche and Lima \(2013\)](#)]. More flexible types of distributions have so far mostly been used to describe income distributions on a highly aggregated level, normally the national level [[Kleiber \(1996\)](#)].

We utilise detailed, longitudinal information on incomes available from the German socio-economic panel (SOEP) to derive a flexible, structured additive distributional regression model for labour incomes of full-time male workers. We consider several candidate distributions for describing the nonnegative conditional income distributions, including the log-normal distribution, the gamma distribution, the inverse Gaussian distribution and the Dagum distribution. To obtain flexible models, we allow for regression effects on potentially all parameters of the income distribution, thereby overcoming the previous concentration on expected incomes. As an illustration, consider the income distributions visualised in [Figure 1](#) corresponding to an “average,” full-time male worker with/without higher education in East and West Germany. Here we find that the income distributions differ considerably not only in terms of their expectation but also with respect to other aspects of the distribution, like the variance (see [Section 4](#) for more details on the analysis).

Some earlier attempts to define distributional regression models comprise [Biewen and Jenkins \(2005\)](#) or [Donald, Green and Paarsch \(2000\)](#). [Biewen and Jenkins \(2005\)](#) suggest to decompose the population into a coarse set of subgroups for which parametric income distributions are estimated such that the distributional form varies over the subgroups. [Donald, Green and Paarsch \(2000\)](#) propose to vary location and scale parameters with respect to covariates while the general shape of the distribution remains fixed over the covariate set. Building on their work, we propose to combine these approaches in the sense that conditional income distributions are modelled parametrically as suggested by [Biewen and Jenkins \(2005\)](#), while allowing for variation in the whole distribution (not just location and scale) with respect to covariates as specified by [Donald, Green and Paarsch \(2000\)](#).

Differences between East and West Germany have received considerable attention in the economic literature [[Biewen \(2000\)](#), [Fuchs-Schündeln](#), [Krueger and](#)

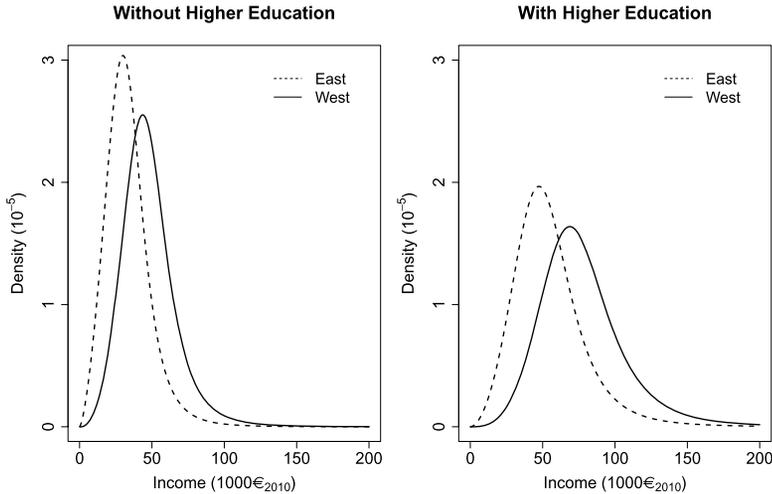


FIG. 1. SOEP data. Four conditional income distributions for 42-year-old males with 19 years of working experience without higher education (left) or with higher education (right) and living in the East (dashed lines) or West (solid lines). Densities shown are posterior means of the densities in our best model DA_M1; compare Section 3 for details on the model specification.

Sommer (2010), Kohn and Antonczyk (2011)] and also consistently played a major role in the domestic political debate. Instead of solely taking a macroeconomic perspective to look at income inequality in the East and West at a highly aggregated level, we build a microeconomic foundation to the analysis of income inequality. Thereby, we consider the effect of various covariates on the conditional individual income distribution underlying the aggregate income distribution. It is our hypothesis that there are not only significant differences between East and West in the conditional mean income but also in the conditional income inequality aggravating the economic divide more than two decades after the reunification.

As a conceptual framework for our analyses, we extend the Bayesian structured additive distributional regression models recently proposed in Klein, Kneib and Lang (2015) for zero-inflated and overdispersed count data regression to general types of univariate distributions. In this class of regression models, all parameters of a potentially complex response distribution are related to additive regression predictors in the spirit of generalised additive models (GAMs). While the latter assume responses to follow a distribution from the exponential family and focus exclusively on relating the mean of a response variable to covariates [see, e.g., Fahrmeir, Kneib and Lang (2004), Fahrmeir et al. (2013), Ruppert, Wand and Carroll (2003), Wood (2004, 2008)], distributional regression enables the consideration of basically any response distribution and allows to specify regression predictors for all parameters of this distribution. The main advantage of distributional regression is that it provides a broad and generic framework for regression models encompassing continuous, discrete and mixed discrete-continuous response distri-

butions and therefore considerably expands the common exponential family framework.

Distributional regression is closely related to generalised additive models for location, scale and shape (GAMLSS) as suggested by [Rigby and Stasinopoulos \(2005\)](#). We prefer the notion of distributional regression for our approach since in most cases, the parameters of the response distribution are in fact not directly related to location, scale and shape but are general parameters of the response distribution and only indirectly determine location, scale and shape. For example, in case of the Dagum distribution, there are three distributional parameters, but none of them is directly related to a measure of location which is jointly determined by all three parameters.

In GAMLSS, inference is commonly based on penalised maximum likelihood estimation achieved via backfitting loops over the additive predictor components. In this paper, we consider a generic Bayesian treatment of distributional regression relying on Markov chain Monte Carlo simulation algorithms. To construct suitable proposal densities, we follow the idea of iteratively weighted least squares proposals [[Brezger and Lang \(2006\)](#), [Gamerman \(1997\)](#)] and use local quadratic approximations to the full conditionals in order to avoid manual tuning. Utilising explicit derivations of the score function and expected Fisher information in these approximations considerably enhances numerical stability as compared to using numerical derivatives and the observed Fisher information (which are frequently used in the R add-on package `gam1ss` implementing penalised likelihood inference). The Bayesian approach also has the advantage to provide credible intervals without relying on asymptotic arguments. The full potential of distributional regression is only exploited when the regression predictor is broadened beyond the scope of simple linear or additive specifications. We will consider structured additive predictors [[Brezger and Lang \(2006\)](#), [Fahrmeir et al. \(2013\)](#)] where each predictor is determined as an additive combination of various types of functional effects, including nonlinear effects of continuous covariates, spatial effects, random effects or varying coefficient terms.

Alternatives to distributional regression are provided by quantile and expectile regression which also allow us to go beyond studying the mean by focusing on local features of the response distribution, indexed by a prespecified asymmetry parameter (the quantile or expectile level); see [Koenker and Bassett \(1978\)](#), [Newey and Powell \(1987\)](#) for the original references and [Koenker \(2005\)](#), [Schnabel and Eilers \(2009\)](#), [Sobotka and Kneib \(2012\)](#), [Yu and Moyeed \(2001\)](#) for more recent overviews. Single quantiles or expectiles are elicitable [[Gneiting \(2011a\)](#), [Osband and Reichelstein \(1985\)](#)] by considering asymmetrically weighted loss functions and consistent estimates can be obtained under rather mild conditions on the conditional distribution of the responses (basically reducing to independence and the correct specification of the quantity of interest). However, when interest focuses on the complete conditional distribution or if distributional quantities such as the Gini coefficient for inequality that are not elicitable by specifying a corresponding

loss function are desired, the direct specification of distributional regression turns out to be advantageous.

Many of the aspects discussed in the remainder of this paper (such as choice of a suitable response distribution and adequate predictor specifications, Bayesian inference, interpretation of estimation results) are relevant beyond our application. We therefore provide an analysis on the proportion of farm outputs achieved by cereals in the application Supplement A [Klein et al. (2015b), Section A.2], to this paper as a second example on distributional regression.

The remainder of the paper is structured as follows: Section 2 provides a detailed introduction to distributional regression and Bayesian inference along our case study on labour incomes. Model choice concerning the type of the response distribution and the specification of the regression predictors is treated in Section 3. Given the selected models, Section 4 provides empirical results on the regional disparities of conditional incomes in East and West Germany. Additional material on the application is provided in the application supplement Section A.1. Section 5 provides a summary and comments on directions for future research. Finally, we summarise general aspects of distributional regression with other types of responses in the methodological Supplement B [Klein et al. (2015c)] which also comprises details on Bayesian inference, derivations of required quantities for the iteratively weighted least squares proposals and simulation studies.

2. Distributional regression. As a conceptual framework for our analysis of labour incomes and their regional disparities, we consider distributional regression models where, conditional on all available covariate information summarised in the vector \mathbf{v}_i , the response variables y_1, \dots, y_n are assumed to be independently distributed with K -parametric densities $p(y_i | \vartheta_{i1}, \dots, \vartheta_{iK}) \equiv p_i$. The conditional distribution p_i of observation y_i given \mathbf{v}_i is indexed by the (in general covariate-dependent) distributional parameters $\vartheta_{i1}, \dots, \vartheta_{iK}$. Each parameter ϑ_{ik} , $k = 1, \dots, K$ is then related to a semiparametric, additive predictor $\eta_i^{\vartheta_{ik}}$ defined in terms of (potentially different) subvectors of the covariate vector \mathbf{v}_i . Similarly, as in generalised linear models, a suitable (one-to-one) response function is utilised to map the predictor to the parameter of interest, that is, $\vartheta_{ik} = h^{\vartheta_{ik}}(\eta_i^{\vartheta_{ik}})$, where the superscript ϑ_{ik} in the predictors and response functions indicates that we are dealing with K predictors specific to the different distributional parameters instead of only one single predictor as in mean regression. The response function is chosen to ensure appropriate restrictions on the parameter space such as the exponential function $\vartheta_{ik} = \exp(\eta_i^{\vartheta_{ik}})$ to ensure positivity. We discuss specific choices for distributional regression of labour incomes after having introduced our data in more detail.

2.1. German labour income data. For studying conditional income distributions in Germany, we utilise information from the German Socio-Economic Panel

[Wagner, Frick and Schupp (2007)]. More specifically, we consider real gross annual personal labour income in Germany as defined in Bach, Corneo and Steiner (2009) for the years 2001 to 2010. We deflate the incomes by the consumer price index [Statistisches Bundesamt (2012)], setting 2010 as our base year. Thus, all incomes are expressed in real-valued 2010 Euros from here on.

Following the standard literature, we only look at the income of males in full-time employment [see, among others, Card, Heining and Kline (2013), Dustmann, Ludsteck and Schönberg (2009)] in the age range 20–60. This yielded 7216 individuals for whom we considered the income trajectories from the ten year period. For each individual, we used every observation for which all required dependent and independent variables were available, yielding a total of $n = 40,965$ observations. Naturally, this implies that for some individuals we do not have full longitudinal coverage over the whole ten year period.

As covariates, we consider educational level measured as a binary indicator for completed higher education (according to the UNESCO International Standard Classification of Education 1997 provided in the SOEP) in effect coding (*educ*), age in years (*age*), previous labour market experience in years (*lmexp*), the calendar time (t), information on the geographical district (*Raumordnungsregion*) representing the area of residence (s) and a binary indicator in effect coding for districts belonging to the eastern part of Germany (*east*). A description of the data set is given in Table A1; details on the specifications for the different effect types will be provided in Section 2.3.

A common assumption in economic analyses of income is that incomes y_i are log-normally distributed with covariate-dependent location parameter η_i (corresponding to the mean of the log-transformed incomes) and a constant scale parameter σ^2 . For an observation i collected at time point t_i , a suitable semiparametric predictor (dropping the dependence on the parameter ϑ_k) could then be specified as

$$(1) \quad \eta_i = \beta_0 + educ_i \beta_1 + f_1(age_i) + educ_i f_2(age_i) + f_3(lmexp_i) + f_{\text{spat}}(s_i) + f_{\text{time}}(t_i),$$

where β_0 represents the overall intercept, β_1 captures the effect of higher education, $f_1(age)$ and $f_2(age)$ are nonlinear effects of age capturing also the interaction with the educational status, $f_3(lmexp)$ is the nonlinear effect of previous labour market experience, $f_{\text{spat}}(s)$ is a spatial effect capturing heterogeneity at the level of the districts s , and $f_{\text{time}}(t)$ is an effect specific for the calendar year t . In a second step, the spatial effect can further be decomposed into

$$(2) \quad f_{\text{spat}}(s) = east_s \gamma_1 + g_{\text{str}}(s) + g_{\text{unstr}}(s),$$

where γ_1 captures the difference between the eastern and western part of Germany and $g_{\text{str}}(s)$ and $g_{\text{unstr}}(s)$ represent spatially structured (smooth) or unstructured

TABLE 1
Selected candidate distributions; for a more comprehensive list see Table B1

Name	Density	Parameters	Response functions
Log-normal	$p(y \mu, \sigma^2) = \frac{1}{\sqrt{2\pi\sigma^2}y} \exp(-\frac{(\log(y)-\mu)^2}{2\sigma^2})$	$\mu \in \mathbb{R}, \sigma^2 > 0$	$h^\mu(\eta) = \eta, h^{\sigma^2}(\eta) = \exp(\eta)$
Inverse Gaussian	$p(y \mu, \sigma^2) = \frac{1}{\sqrt{2\pi\sigma^2}y^{3/2}} \exp(-\frac{(y-\mu)^2}{2y\mu^2\sigma^2})$	$\mu, \sigma^2 > 0$	$h^\mu(\eta) = h^{\sigma^2}(\eta) = \exp(\eta)$
Gamma	$p(y \mu, \sigma) = (\frac{\sigma}{\mu})^\sigma \frac{y^{\sigma-1}}{\Gamma(\sigma)} \exp(-\frac{\sigma}{\mu}y)$	$\mu, \sigma > 0$	$h^\mu(\eta) = h^\sigma(\eta) = \exp(\eta)$
Dagum	$p(y a, b, c) = \frac{acy^{ac-1}}{bac(1+(y/b)^a)^{c+1}}$	$a, b, c > 0$	$h^a(\eta) = h^b(\eta) = h^c(\eta) = \exp(\eta)$

(unsmooth) district-specific effects. Note that, in addition to the East–West indicator *east*, more district-specific information could be included if desired. While this decomposition could simply be plugged into (1) to obtain a reduced-form specification, it can also be interpreted as a hierarchical multilevel specification where we differentiate between an individual-specific level in (1) and a region-specific level in (2). Further details on the predictors and associated priors will be discussed in Section 2.3.

2.2. *Potential response distributions.* One of the great advantages of structured additive distributional regression is the wide range of distribution types that can be modelled. Since labour income is by definition positive, we will restrict ourselves to four nonnegative distributions summarised in Table 1. For a more comprehensive list of distributions supported by the distributional regression framework, see Section B.1.1.

As noted, the standard conditional distribution type in econometric income analyses is the log-normal distribution. Next to its theoretical appeal from an economic perspective [see Arnold (2008), page 122], it has the advantage that it makes the vast statistical inference machinery built around Gaussian regression available to researchers. However, Atkinson (1975) and others have noted that, at least for the aggregate income distribution, the log-normal distribution fit is problematic at times, especially for the upper tail of the distribution.

Partly as a consequence, various other distribution types have thus been suggested for the modelling of income distributions. Salem and Mount (1974) proposed the gamma distribution as a suitable alternative to the log-normal distribution. One of its advantages is that its estimation is possible within the framework of generalised linear models as the distribution belongs to the exponential family (as long as covariate effects are restricted to the mean).

The third distribution we consider also belongs to the exponential family (if the second parameter is assumed to be independent of covariates). The inverse Gaussian distribution has to our knowledge not been used in the context of modelling income distributions yet. But for other nonnegative distributions with a similar economic rationale, like the distribution of claim sizes arising in car insurance [Heller,

Stasinopoulos and Rigby (2006), Klein et al. (2014)], it has shown to perform well due to its flexibility in modelling extreme right skewness. As it is conceivable that some conditional income distributions also portray such extreme skewness, we decided to also consider this distribution type.

The last distribution we consider is the Dagum distribution [Dagum (1977)] which belongs to the beta-type size distributions that have seen considerable attention in the literature on modelling (aggregate) income distributions [see Kleiber and Kotz (2003)]. One of its appealing properties is that towards the upper end of the distribution its shape mirrors the one of the Pareto distribution which is generally assumed to provide a good approximation for the income distribution for the top percentiles of the (aggregate) income distribution [Piketty and Saez (2007)].

2.3. Structured additive predictors and associated priors.

Generic representation. While considering a specific instance of a structured additive predictor for the analysis of income, a generic structured additive predictor for parameter ϑ_{ik} is given by

$$(3) \quad \eta_i^{\vartheta_k} = \beta_0^{\vartheta_k} + f_1^{\vartheta_k}(\mathbf{v}_i) + \dots + f_{J_k}^{\vartheta_k}(\mathbf{v}_i),$$

where β_0 represents the overall level of the predictor and the functions $f_j^{\vartheta_k}(\mathbf{v}_i)$, $j = 1, \dots, J_k$, relate to different covariate effects defined in terms of the complete covariate vector \mathbf{v}_i . Note that each distribution parameter may depend on different covariates and a different number of effects J_k , but we suppress this possibility (as well as the parameter index) in the following.

In structured additive regression, each function f_j is approximated by a linear combination of D_j appropriate basis functions, that is,

$$f_j(\mathbf{v}_i) = \sum_{d_j=1}^{D_j} \beta_{j,d_j} B_{j,d_j}(\mathbf{v}_i)$$

such that in matrix notation we can write $\mathbf{f}_j = (f_j(\mathbf{v}_1), \dots, f_j(\mathbf{v}_n))' = \mathbf{Z}_j \boldsymbol{\beta}_j$, where $\mathbf{Z}_j[i, d_j] = B_{j,d_j}(\mathbf{v}_i)$ is a design matrix and $\boldsymbol{\beta}_j$ is the vector of coefficients to be estimated. To ensure identifiability specific constraints representing for example centring of the functional effects are added, see Section B.2.2 for further details. The basis function representation then leads to the following matrix representation of the generic predictor (3):

$$(4) \quad \boldsymbol{\eta} = \beta_0 \mathbf{1} + \mathbf{Z}_1 \boldsymbol{\beta}_1 + \dots + \mathbf{Z}_J \boldsymbol{\beta}_J.$$

For each of the parameter vectors $\boldsymbol{\beta}_j$ we can then either assume a hierarchical specification, where $\boldsymbol{\beta}_j$ is related to another structured additive predictor (as in the

case of the spatial effect in our example), or we directly assume the multivariate normal prior

$$(5) \quad p(\boldsymbol{\beta}_j | \tau_j^2) \propto \left(\frac{1}{\tau_j^2}\right)^{(\text{rk}(\mathbf{K}_j))/2} \exp\left(-\frac{1}{2\tau_j^2} \boldsymbol{\beta}'_j \mathbf{K}_j \boldsymbol{\beta}_j\right)$$

with (potentially rank-deficient) precision matrix \mathbf{K}_j and prior smoothing variance τ_j^2 . The latter is assigned an inverse gamma hyperprior $\tau_j^2 \sim \text{IG}(a_j, b_j)$ (with $a_j = b_j = 0.001$ as a default option) in order to obtain a data-driven amount of smoothness.

A detailed discussion of terms that fit into the generic predictor framework (in the context of mean regression) is provided in Fahrmeir, Kneib and Lang (2004) and Fahrmeir et al. (2013), Chapters 8 and 9.

In the following, we will discuss suitable specifications and prior assumptions for the hierarchical predictor defined in (1) and (2). Note that we drop the dependence on the distributional parameter indicated by the superscript ϑ_k , the observation index i and the function index j to simplify notation. Hierarchical extensions are treated in detail in Lang et al. (2014).

Linear effects. For all parametric, linear effects, we assume a flat, noninformative prior. This may be considered the limiting case of a multivariate Gaussian prior with high dispersion which can also be used to achieve regularisation in the case of high-dimensional parameter vectors. In our analyses, we assume linear effects for the intercept and the educational indicator, as well as for the East–West indicator.

Continuous covariates. For the effects of age and previous work experience, assuming a linear effect is probably too restrictive. We therefore consider P(enalised)-splines [Eilers and Marx (1996)] as a flexible device for including potentially nonlinear effects $f(x)$ of a continuous covariate x . In a first step, $f(x)$ is approximated by a linear combination of D B-spline basis functions $B_d(x)$ that are constructed from piecewise polynomials of a certain degree l upon an equidistant grid of knots, $f(x) = \sum_{d=1}^D \beta_d B_d(x)$. To avoid the requirement of choosing an optimal number of knots together with optimal knot positions, Eilers and Marx (1996) regularise the function estimate by augmenting a difference penalty to the fit criterion. In our Bayesian framework, the stochastic analogue is to assume a first or second order random walk

$$\begin{aligned} \beta_d &= \beta_{d-1} + \varepsilon_d, & d &= 2, \dots, D, \\ \beta_d &= 2\beta_{d-1} - \beta_{d-2} + \varepsilon_d, & d &= 3, \dots, D \end{aligned}$$

with Gaussian errors $\varepsilon_d \sim \text{N}(0, \tau^2)$ and noninformative priors for β_1 or β_1 and β_2 [Lang and Brezger (2004)]. The joint prior of all basis coefficients $\boldsymbol{\beta} =$

$(\beta_1, \dots, \beta_D)'$ can then be shown to be a (partially improper) multivariate Gaussian distribution with zero mean and precision matrix $\mathbf{K} = \mathbf{D}'\mathbf{D}$, where \mathbf{D} is a difference matrix of appropriate order. In our analysis, we use twenty inner knots, a cubic spline basis and a second order random walk prior as the default specification for penalised splines.

In the case of the age effect, we allow for separate functions for individuals with high and low levels of education. This is achieved by the inclusion of the varying coefficient term [Hastie and Tibshirani (1993)] $f_2(\text{age})$ such that the age effect is given by $f_1(\text{age}) - f_2(\text{age})$ for individuals with low educational level and $f_1(\text{age}) + f_2(\text{age})$ for individuals with high educational level. In this case, a penalised spline can be assumed for function $f_2(\text{age})$ as well.

Random effects. Penalised splines can in principle also be considered to represent the temporal effect $f_{\text{time}}(t)$ in (1). However, since in economic research temporal effects such as ours are generally considered by year-specific effects, we do not impose the smoothness assumption implied by penalised splines. We therefore consider a random effects specification where separate regression effects $\beta_t = f_{\text{time}}(t)$ are assumed for the distinct time points. An i.i.d. Gaussian prior with random effects variance τ^2 is then placed on the coefficients $\beta = (\beta_1, \dots, \beta_T)'$. Similarly, random effects priors can be used for any other grouping variable with levels $\{1, \dots, G\}$ present in the data.

Note that we have not included individual-specific random effects. The reason for this is that we are specifically interested in the unobserved heterogeneity among individuals with similar covariate sets which finds expression in income inequality among them. In some sense our analysis is thus systematically different from standard regression techniques which pursue to eradicate the stochastic component or at least reduce it to a minimum. The inclusion of individual-specific effects goes a long way towards seemingly achieving this aim, as the share of the variance left to the error term is drastically reduced. However, the inferential gain obtained thereby could be expressed as follows: including individual-specific effects, we have found that incomes are largely different because individuals are different. While there are some analyses where such eradication of variance is useful, it sheds little insights on the nature of inequality at the disaggregated level since we are unable to disentangle the differences between individuals in a meaningful way.

Spatial effects. For the spatial effect $f_{\text{spat}}(s)$ defined upon the discrete, spatial variable $s \in \{1, \dots, S\}$ which denotes the different regions in the data set, we assume a hierarchical predictor specification following Lang et al. (2014). In fact, equation (2) merely defines a second structured additive predictor where now the distinct spatial regions define the unit of observation. As a consequence, any type of regression effect that is specific for the region can be included on this level. In our case, the East–West indicator is one such example that is assigned a parametric effect with flat prior.

In addition, we consider the spatially structured and spatially unstructured effects $g_{\text{str}}(s)$ and $g_{\text{unstr}}(s)$, respectively. In both cases, separate regression effects $\beta_{\text{str},s} = g_{\text{str}}(s)$ and $\beta_{\text{unstr},s} = g_{\text{unstr}}(s)$ are assumed for each of the regions, but the effects differ in terms of their prior assumptions. For the structured spatial effect, we assume spatial correlations defined implicitly by assuming a Gaussian Markov random field prior [Rue and Held (2005)] for a suitable neighbourhood structure derived from the spatial orientation of the data. The most common case would be to treat two regions as neighbours if they share a common boundary. If ∂_s denotes the set of all neighbours of region s , the Markov random field prior then assumes

$$(6) \quad \beta_{\text{str},s} | \beta_{\text{str},r}, r \neq s, \tau^2 \sim N\left(\sum_{r \in \partial_s} \frac{1}{N_s} \beta_{\text{str},r}, \frac{\tau^2}{N_s}\right)$$

with number of neighbours of region s denoted as N_s . Consequently, the conditional mean of $\beta_{\text{str},s}$ given all other coefficients is the average of the neighbouring regions. It can be shown that the conditional normal distributions specified in (6) correspond to a multivariate, partially improper normal distribution with zero mean and precision matrix given by the adjacency matrix induced by the neighbourhood structure.

For the unstructured spatial effect, we consider an i.i.d. Gaussian prior, that is, we assume a random effects prior specification. The rationale for considering both a structured and unstructured part of the spatial effect is that they are surrogates for unobserved spatial heterogeneity which may either be spatially structured (i.e., spatially smooth) or unstructured.

2.4. Bayesian inference. To perform Bayesian inference, we consider Markov chain Monte Carlo (MCMC) simulation techniques and develop suitable proposal densities based on iteratively weighted least squares (IWLS) approximations to the full conditionals. The derivation of the approximations and the complete algorithm are documented in Section B.2. Here, we only sketch the essential parts.

IWLS proposals for regression coefficients. The regression coefficients β_j are proposed from $N(\boldsymbol{\mu}_j, \mathbf{P}_j^{-1})$ with expectation and precision matrix

$$\boldsymbol{\mu}_j = \mathbf{P}_j^{-1} \mathbf{Z}'_j \mathbf{W}(\mathbf{z} - \boldsymbol{\eta}_{-j}), \quad \mathbf{P}_j = \mathbf{Z}'_j \mathbf{W} \mathbf{Z}_j + \frac{1}{\tau_j^2} \mathbf{K}_j,$$

where \mathbf{W} is a diagonal matrix of working weights $w_i = \mathbb{E}(-\partial^2 l / \partial \eta_i^2)$, $\mathbf{z} = \boldsymbol{\eta} + (\mathbf{W})^{-1} \mathbf{v}$ is a working response depending on the score vector $\mathbf{v} = \partial l / \partial \boldsymbol{\eta}$ and $\boldsymbol{\eta}_{-j} = \boldsymbol{\eta} - \mathbf{Z}_j \boldsymbol{\beta}_j$ is the predictor without the j th component. The working weights and the score vector are specific for the chosen response distribution and induce an automatic adaptation to the form of the full conditional without requiring manual tuning.

Updates for the smoothing variances. The smoothing variances τ_j^2 can be sampled in a Gibbs update where $\tau_j^2 | \cdot \sim \text{IG}(a'_j, b'_j)$, with updated parameters $a'_j = \frac{\text{rk}(\mathbf{K}_j)}{2} + a_j$, $b'_j = \frac{1}{2} \boldsymbol{\beta}'_j \mathbf{K}_j \boldsymbol{\beta}_j + b_j$.

Working weights. The specification of the working weights \mathbf{W} involves the expectations of the negative second derivatives of the log-likelihood which improved both mixing and acceptance rate in comparison with the (seemingly simpler) approach of using the negative second derivative without deriving the expectation. Furthermore, invertibility of the precision matrix \mathbf{P}_j is ensured for many distributions when using the expectation since the working weights are then nonnegative. Explicit derivations for both the distributions utilised for analysing labour incomes and the additional distributions summarised in Table B1 can be found in Section B.2.3.

Propriety of the posterior. Propriety of the posterior in distributional regression can be ensured when combining the assumptions considered in Klein, Kneib and Lang (2015) for count data regression with appropriate restrictions on the densities. These need to be bounded or integrable with respect to the predictors, whereby at least one observation fulfilling the latter assumption is required. Note that integrability of the densities can be assured by the assumption that none of the distributional parameters is on the boundary of the parameter space (an assumption that would also have to be made to apply standard maximum likelihood asymptotics).

Software. Our Bayesian approach to distributional regression is implemented in the free, open source software BayesX [Belitz et al. (2015)]. As described in Lang et al. (2014), the implementation makes use of efficient storing mechanisms for large data sets and sparse matrix algorithms for sampling from multivariate Gaussian distributions. An R interface to BayesX is provided in the R add-on package `bamlss` [Umlauf et al. (2014)].

Empirical evaluation. We compared the empirical performance of the proposed Bayesian approach to the frequentist GAMLSS framework in two simulation scenarios and also investigated the performance of the deviance information criterion [DIC, Spiegelhalter et al. (2002)] for choosing between competing models. The studies and their outcomes are documented in more detail in Section B.3. A summary on the ability of the DIC for model choice is given in Section 3 and for the comparison with the frequentist approach (denoted as ML) in the following:

(1) *Comparison with ML in additive models.* In purely additive models, the point estimates and corresponding posterior means, as well as their mean squared errors (MSEs), are very similar. However, coverage rates based on asymptotic

maximum likelihood theory for ML are far too narrow in several distribution parameters. In particular, for the Dagum distribution, rates for all three parameters are far from the desired coverage level, while the credible intervals of the Bayesian approach are still reliable (albeit being usually slightly too conservative); compare, for example, Figure B2.

(2) *Comparison with ML in geoadditive models.* 10% of the estimation runs of ML failed before convergence. MSEs of the spatial effect (based on a Markov random field) are slightly smaller for the Bayesian approach compared to ML. While the MSEs of the other effects do not deteriorate for our proposed method, we observe partly increasing MSEs for ML.

3. Model choice. In any application of distributional regression, one faces important model choice decisions: choosing the most appropriate out of a set of potential response distributions and selecting adequate predictor specifications for each parameter of these distributions. For our application on conditional income distributions, we consider the inverse Gaussian (IG), log-normal (LN), gamma (GA) and Dagum (DA) distribution as candidate distributions. A general predictor that could now be utilised for any of the parameters of these distributions was already introduced in equations (1) and (2). Instead of performing a complete step-wise model selection for each distribution, we study the following model specifications:

(M1) All distributional parameters are related to a predictor of type (1). For the spatial effect, we only include the unstructured effect since it turned out in exploratory analyses that the smooth component has only negligible impact.

(M2) Instead of modelling all parameters in terms of covariates, the model structure of M1 is only applied to the parameters μ in the case of LN, IG and GA, and b in case of DA. The parameters a, c, σ, σ^2 are considered to be equal across all individuals. This corresponds to a usual GAM specification with focus on conditional means.

(M3) All parameters are modelled in analogy to M1 except that the random effect for calendar time and the complete spatial effect (including the East–West indicator) are not included in the parameters a, c, σ, σ^2 .

In total, we therefore end up with 12 models to compare. In the following, we will discuss different options for conducting this comparison and will also comment on their wider applicability in the context of model choice for distributional regression.

3.1. *Deviance information criterion.* The deviance information criterion (DIC) is a commonly used criterion for model choice in Bayesian inference that has become quite popular due to the fact that it can easily be computed from the MCMC output. If $\theta^{[1]}, \dots, \theta^{[T]}$ is a MCMC sample from the posterior for the

complete parameter vector θ , the DIC is given by $\overline{D(\theta)} + pd = 2\overline{D(\theta)} - D(\overline{\theta}) = \frac{2}{T} \sum D(\theta^{[t]}) - D(\frac{1}{T} \sum \theta^{[t]})$, where $D(\theta) = -2 \log(f(\mathbf{y}|\theta))$ is the model deviance and $pd = \overline{D(\theta)} - D(\overline{\theta})$ is an effective parameter count.

The DIC can be used to discriminate between types of response distributions as well as different predictor specifications for a fixed distribution. The latter can also be implemented in a stepwise model choice strategy. However, since the DIC is sample-based, small differences of DIC values for competing models may induce a region of indecisiveness. If in such a situation sparser models are desired, the DIC-based selection of covariate effects can be assisted by only including significant effects, that is, effects for which the credible interval of a certain level does not contain the zero (parametric effects) or the zero line (nonparametric effects); compare also Section B.3.3.2.

For distributions and models considered in our applications, we conducted simulations on the performance of the DIC which are documented in detail in Section B.3.3. The basic outcome is that the DIC can discriminate between competing response distributions although differences can be rather small depending on what distributions are compared. Concerning the identification of relevant covariates, we focused on spatial effects and found that the DIC usually is in clear favour of the true model if a relevant effect is omitted. In the reverse situation, that is, irrelevant information is included, the DICs of the true models are only slightly smaller, but then the irrelevant covariate mainly yields an insignificant effect (i.e., the 95% credible interval of each region contains zero) and would thus be excluded under the aim of a sparser model. For count data distributional regression models, the performance of the DIC was also positively evaluated by Klein, Kneib and Lang (2015) who compare several misspecified models to the true model in terms of the DIC.

The DIC values for the 12 income regression models under consideration are documented in Table 2 and indicate a clear preference for the model DA_M1. In general, it is noticeable that the DIC favours our flexible model specifications (M1) compared to the simplified versions (M2, M3).

3.2. Quantile residuals. For continuous random variables, it is a well-known result that the cumulative distribution function $F(\cdot)$ evaluated at the random variable y_i yields a uniform distribution on $[0, 1]$. As a consequence, quantile residuals defined as $\hat{r}_i = \Phi^{-1}(F(y_i|\hat{\vartheta}_i))$, with the inverse cumulative distribution function (c.d.f.) of a standard normal distribution Φ^{-1} and $F(\cdot|\hat{\vartheta}_i)$ denoting c.d.f. with estimated parameters $\hat{\vartheta}_i = (\hat{\vartheta}_{i1}, \dots, \hat{\vartheta}_{iK})'$ plugged in, should at least approximately be standard normally distributed if the correct model has been specified [Dunn and Smyth (1996)]. In practice, the residuals can be assessed graphically in terms of quantile–quantile-plots: the closer the residuals are to the bisecting line, the better the fit to the data. We suggest to use quantile residuals as an effective tool for deciding between different distributional options where strong deviations from the bisecting line allow us to sort out distributions that do not fit the data well.

TABLE 2

Comparison of DIC values (calculated based on the complete data set) and average scores obtained from ten-fold cross-validation

Distribution	DIC	Quadratic score	Logarithmic score	Spherical score	CRPS
LN_M1	179,090	0.130	-2.436	0.362	-2.158
LN_M2	180,533	0.126	-2.460	0.357	-2.141
LN_M3	179,451	0.130	-2.435	0.362	-2.163
IG_M1	184,614	0.146	-2.274	0.378	-1.620
IG_M2	189,702	0.138	-2.314	0.366	-1.677
IG_M3	186,494	0.144	-2.282	0.374	-1.642
GA_M1	177,453	0.161	-2.172	0.396	-1.274
GA_M2	178,736	0.156	-2.181	0.392	-1.279
GA_M3	177,971	0.160	-2.174	0.395	-1.277
DA_M1	172,421	0.168	-2.103	0.405	-1.266
DA_M2	173,791	0.164	-2.120	0.402	-1.274
DA_M3	172,790	0.167	-2.108	0.404	-1.270

Quantile residuals are closely related to the probability integral transform (PIT) which considers $u_i = F(y_i | \hat{\theta}_i)$ without applying the inverse standard normal c.d.f. If the estimated model is a good approximation to the true data generating process, the u_i will then approximately follow a uniform distribution on $[0, 1]$. As a graphical device, histograms of the u_i are then typically considered.

Quantile residual plots for the models of type M1 are shown in Figure 2. Similar outcomes for model types M2/M3 and PITs for the models M1 can be found in Figures A1 and A2, respectively. We prefer quantile residuals in the quantile–quantile-plot representation since they avoid the requirement to define breakpoints in the construction of the histogram.

While none of the distributions provides a perfect fit for the data, the Dagum distribution turns out to be most appropriate for residuals in the range between -2 and 2 but deviates from the diagonal line for extreme residuals. In contrast, the log-normal and inverse Gaussian distribution seem to have problems in capturing the overall shape of the income distribution, resulting in sigmoidal deviations from the diagonal. Residuals of the gamma model are reasonable in the range between -2 and 2 (similar to the Dagum distribution) but deviate more strongly from the diagonal for extreme residuals.

3.3. Proper scoring rules. Gneiting and Raftery (2007) propose proper scoring rules as summary measures for the evaluation of probabilistic forecasts, that is, to evaluate the predictive ability of a statistical model. We consider three common scores, namely, the Brier or quadratic score (QS), the logarithmic score (LS) and the spherical score (SPS). For continuous response distributions with density

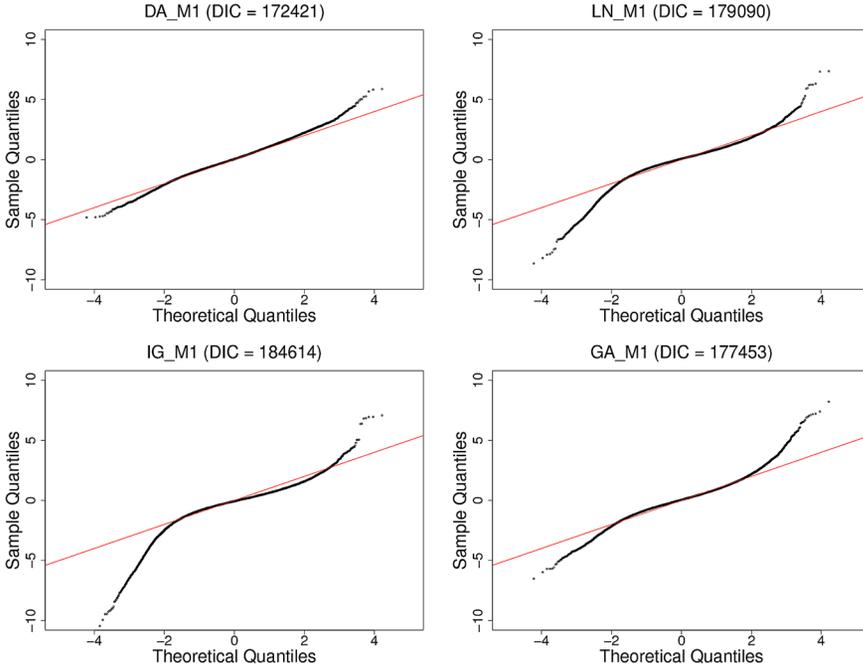


FIG. 2. Comparison of quantile residuals for the full models DA_M1 (topleft), LN_M1 (topright), IG_M1 (bottomleft), GA_M1 (bottomright).

$p_r(y) = p(y|\vartheta_{r1}, \dots, \vartheta_{rK})$ and a given new realisation y_{new} , these are defined as

$$\begin{aligned} \text{LS}(p_r, y_{\text{new}}) &= \log(p_r(y_{\text{new}})), \\ \text{SPS}(p_r, y_{\text{new}}) &= \frac{p_r(y_{\text{new}})}{(\int |p_r(y)|^2 dy)^{1/2}}, \\ \text{QS}(p_r, y_{\text{new}}) &= 2p_r(y_{\text{new}}) - \int |p_r(y)|^2 dy. \end{aligned}$$

Appropriate definitions for discrete as well as mixed discrete continuous responses are provided in Section B.1.2. As a fourth alternative, we consider the continuous ranked probability score (CRPS)

$$\text{CRPS}(p_r, y_{\text{new}}) = - \int_{-\infty}^{\infty} (F_r(y) - \mathbb{1}_{\{y \geq y_{\text{new}}\}})^2 dy,$$

where F_r is the cumulative distribution function corresponding to the density p_r [Gneiting and Ranjan (2011)]. Laio and Tamea (2007) showed that the CRPS score can also be written as

$$\text{CRPS}(p_r, y_{\text{new}}) = -2 \int_0^1 (\mathbb{1}_{\{y_{\text{new}} \leq F_r^{-1}(\alpha)\}} - \alpha)(F_r^{-1}(\alpha) - y_{\text{new}}) d\alpha,$$

where $F_r^{-1}(\alpha)$ is the quantile function of p_r evaluated at the quantile level $\alpha \in (0, 1)$. This formulation allows not only to look at the sum of all score contributions (i.e., the whole integral) but also to perform a quantile decomposition and to plot the mean quantile scores versus α in order to compare fits of specific quantiles [Gneiting and Ranjan (2011)]. This decomposition is especially helpful in situations where the quantile score can be interpreted as an economically relevant loss function [Gneiting (2011b)].

In practice, we obtain the probabilistic forecasts in terms of predictive distributions p_r for observations y_r by cross-validation, that is, the data set is divided into subsets of approximately equal size and predictions for one of the subsets are obtained from estimates based on all the remaining subsets. Let y_1, \dots, y_R be data in a hold-out sample and p_r the predictive distributions with predicted parameter vectors $\hat{\vartheta}_r = (\hat{\vartheta}_{r1}, \dots, \hat{\vartheta}_{rK})'$, $r = 1, \dots, R$. Competing forecasts are then ranked by averaged scores $S = \frac{1}{R} \sum_{r=1}^R S(p_r, y_r)$ such that higher scores deliver better probabilistic forecasts when comparing different models.

In our application, we conducted ten-fold cross-validation; observations are assigned randomly to the different folds. The scores discussed above are documented in Table 2 where the values are averages of the ten folds (and scores within the folds are themselves averages over the individual score contributions). In line with the DIC and the residual plots, the scores of the DA_M1 model are the highest and thus deliver the best forecast among the 12 models under consideration. Also similar to the DIC, models of type M2 (the simplest versions) show lower scores compared to the ones of type M3 and they themselves are inferior compared to the most flexible models of type M1.

In addition to the averages over the ten folds, the proper scoring rules can also be used to assess the predictive distributions in more detail. We illustrate this along a decomposition of the CRPS over quantile levels (Figure 3) and a decomposition of the scores over the cross-validation folds; compare the supplement Section A.1. The quantile level decomposition of the CRPS again indicates a comparable performance of the Dagum distribution and the gamma distribution as compared to the inverse Gaussian distribution which performs somewhat worse and the log-normal distribution which shows a considerably deteriorated behaviour. This ordering holds true over the complete range of quantiles. The fact that the log-normal distribution fails to provide a competing predictive ability is most probably related to the strong impact of the extreme observations. These are hard to capture by the log-normal distribution in general. However, since extreme observations are typically also influential observations, they seem to impact estimates in the log-normal model to such an extent that even predictions for the central part of the distribution are affected negatively.

4. Regional disparities of the distribution of labour income in Germany.

As discussed in the Introduction, our main focus is on investigating differences in

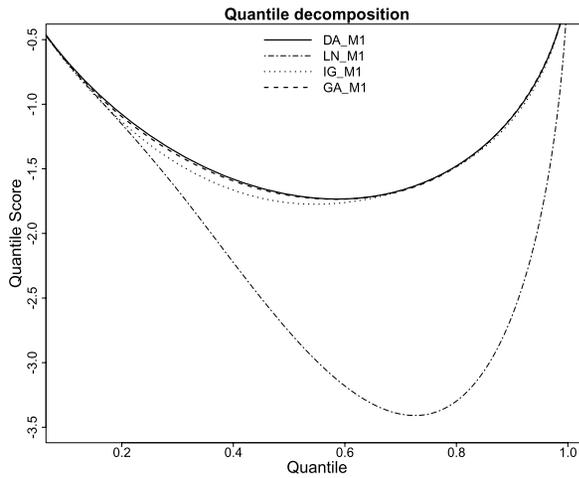


FIG. 3. *Quantile decomposition of CRPS in the full models DA_M1, LN_M1, IG_M1, GA_M1.*

conditional income distributions between former East and West Germany in the first decade of the new millennium. More specifically, we focus on differences in the inequality of the conditional income distribution as measured by the Gini coefficient [Silber (1999)] next to significant differences in the first two moments of conditional income distributions. Based on our model choice, we illustrate the estimation results along the Dagum model DA_M1.

In their seminal paper, DiNardo, Fortin and Lemieux (1996), stress the need to look at differences between the whole conditional income distributions rather than just the conditional mean income, or certain indices. Using our proposed estimation procedure, this is feasible. Figure 1 displays an exemplary contrast of four conditional income distributions in a ceteris paribus type analysis. The four distributions have all but two covariates fixed at their average value. For age (42 years) and labour market experience (19 years) we use the arithmetic mean of the observations in our sample, while we fixed the random effects at their prior expectation, that is, at zero. Keeping these covariates fixed, we can observe the nature of the change if the regional variable is changed from East to West. For both educational levels, this figure furthermore indicates that there is a noticeable difference not only in the mean value of the distributions but also in other aspects, like variability, skewness, etc. Thus, a simple analysis of means falls short of portraying a comprehensive picture of the differences in income between East and West.

Note that for determining the densities displayed in Figure 1 we consider the posterior mean of the densities obtained in the different MCMC iterations instead of plugging in the posterior mean parameters in the corresponding parametric densities. The availability of such posterior mean estimates is another advantage of the Bayesian inferential approach based on MCMC simulations.

There are various additional aspects of the distribution that can be considered. In principle, it is possible to obtain any distributional measure from the conditional distribution as long as it is defined for the given distribution type and the corresponding parameter set. Here, we consider the mean, the standard deviation and the Gini coefficient of the estimated conditional income distributions. While the mean provides important information on the location of the income distribution, the standard deviation provides information on the scale of the distribution and the Gini coefficient is the most frequently used scalar measure on income inequality [Silber (1999)]. We therefore look at three important aspects of the conditional income distributions and observe how they change over the covariate space.

4.1. *The spatial effect on conditional means and standard deviations.* Assuming a Dagum distribution, the first two moments of the conditional income distributions of y_i can be found in Dagum (2008), respectively. Figure 4 displays the posterior mean estimates for the expected incomes for each of the 96 regions (*Raumordnungsregionen*) and education. As described above, the other covariates are fixed at their mean.

Unsurprisingly, there is a clearly visible divide between East and West Germany, as expected incomes are higher in the former Federal Republic of Germany for both education levels and at the average of the other covariates. Abstracting from the variations at the district level, we get an expected income of 33,600€ if

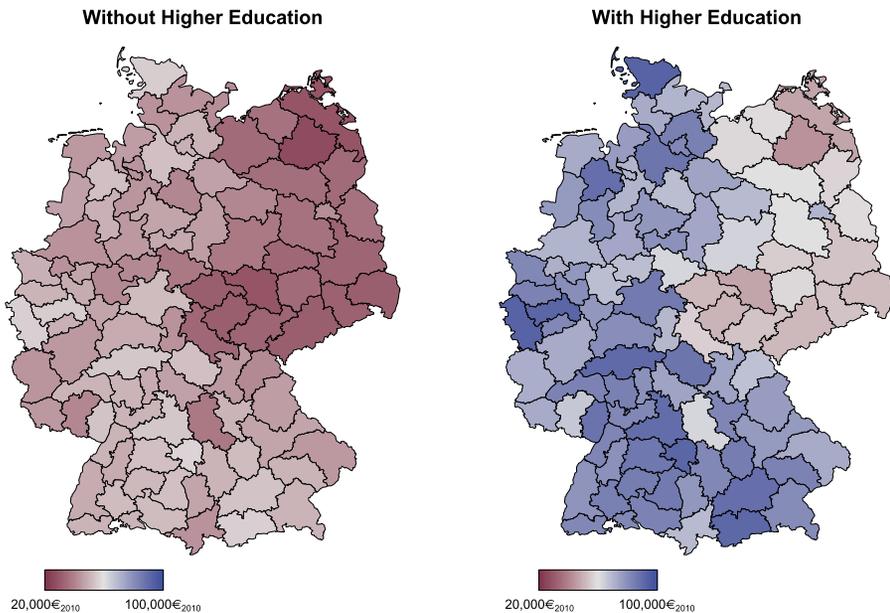


FIG. 4. SOEP data. Posterior means for the expected incomes for 42-year-old males with 19 years of working experience. Left: males without higher education. Right: males with higher education.

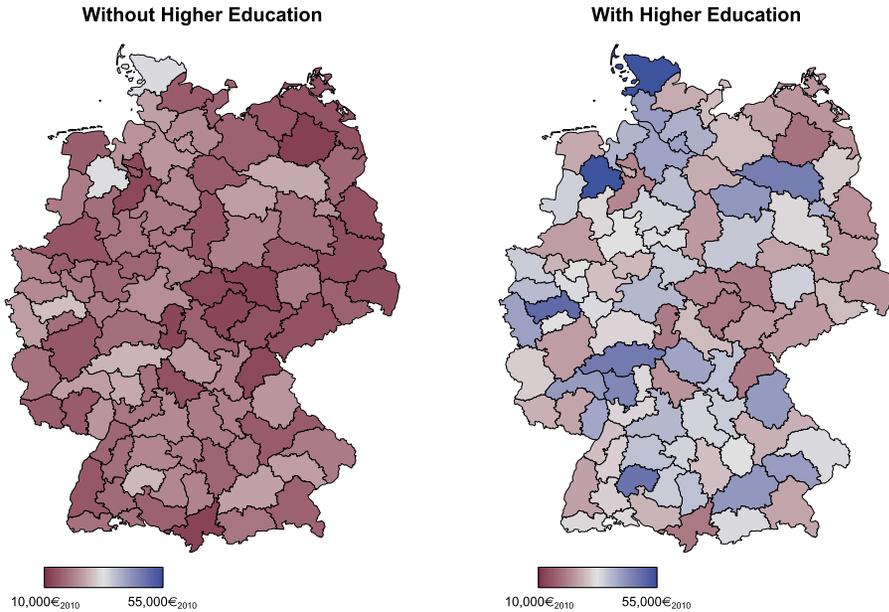


FIG. 5. SOEP data. Posterior means for standard deviations for 42-year-old males with 19 years of working experience. Left: males without higher education. Right: males with higher education.

the average man lives in the East and has no higher education. With higher education the income increases to 55,200€. The corresponding values if a person with the same attributes lives in the West are 48,100€ and 78,300€. The difference between East and West is thus 14,500€ (12,000€; 17,100€) and 23,100€ (19,000€; 27,400€) without and with higher education, respectively, where the numbers in the brackets denote the corresponding 95% credible intervals. In addition to posterior means, we also looked at posterior medians. Overall, differences were negligible, which is in line with the theory suggesting asymptotic normality for the posterior distribution.

The posterior mean estimates for the standard deviations of the conditional income distributions are shown in Figure 5. We prefer presenting the square roots of the second moments, that is, we consider the standard deviations rather than the variances for interpretability reasons.

For standard deviations, the division between East and West is not as distinct as for the means. The main difference in the scale of the conditional distributions is found between the education levels and not along the different regions or former two parts of Germany. Nonetheless, if we set the spatial random effect to zero again and only consider the structural effect, the resultant conditional distribution in the West has a standard deviation of 19,300€, while that of the East has a standard deviation of 16,000€ for those without higher education. For those with higher education the respective numbers are 32,000€ and 26,600€. The difference between the standard deviations is thus 3300€ (1300€; 5200€) in the group

of lower educated males and 5400€ (1700€; 9100€) for the one with higher educated males.

Our results show that evaluated at the mean of other covariates, the first and second moment are significantly different in East and West Germany for both education levels, highlighting the diverse nature of the change of conditional income distributions.

4.2. The spatial effect on the conditional income inequality. The Gini coefficient is an inequality measure based on the Lorenz curve [Sarabia (2008)], which can vary between the value 0 (everybody has the same) and 1 (one person has everything). Note that the Gini coefficient is scale invariant such that in standard mean regression on log-incomes it would be postulated as constant across the covariate space. In analogy to the conditional mean income and standard deviation, the Gini coefficient of the conditional income distribution can easily be obtained from the parameter estimates of the Dagum distribution [Dagum (2008), page 104].

Figure 6 portrays the posterior mean estimates for the Gini coefficients for each region. As we can see, the differences are not as clear cut as for the conditional mean incomes. Nonetheless, the pattern emerging indicates that income inequality among 42-year-old males with 19 years of experience is higher in the East for both education levels. Indeed, if we only consider the impact of the binary East–West

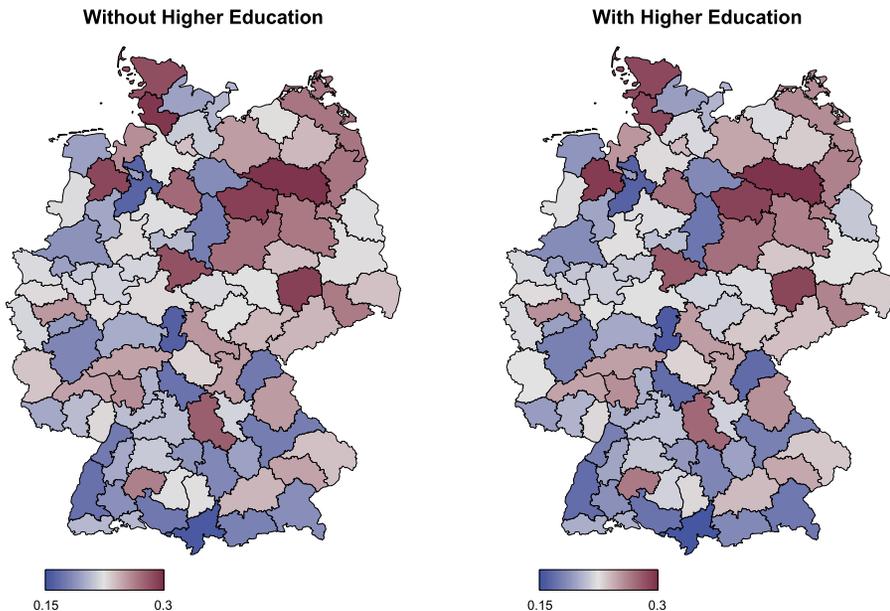


FIG. 6. SOEP data. Posterior means for the Gini coefficients for 42-year-old males with 19 years of working experience. Left: males without higher education. Right: males with higher education.

variable on the Gini coefficient, we obtain a difference of the posterior means of 0.039 and 0.036 for those without higher education and those with higher education, respectively. The corresponding 95% credible intervals are [0.015, 0.067] and [0.013, 0.063], respectively. Thus, we have a significantly larger income inequality for 42-year-old males with 19 years of experience, as measured by the Gini coefficient, in the East than in the West. Putting these differences into perspective, the standard deviations of the Gini coefficients of the regions' conditional income distributions within East and West are 0.030 and 0.031 for those without higher education, and 0.032 and 0.031 for those with higher education. Thus, the differences between East and West are not only significant, they also surpass their variation within East and West.

4.3. Further analysis of the conditional income distribution.

The effect of varying age and experience. Next to spatial effects, the impact of the other covariates can also be of interest. In the following, we focus on the effects of age and experience, while the effect of year is treated in Section A.1.1. For results on additional covariate sets, see Section A.1.2.

In Figure 7, we display the expected conditional mean income and the Gini coefficient with respect to age and experience. In order to keep the dimension of the varying covariate to one, we simply assume that from the age of 21 onwards people gain one year of work experience as they grow older by one year. Here, we thus portray the development of expected incomes and the Gini coefficient for full-time working males who have been working since the age of 21. With regard to the categorical variables region and education, we consider only the West and lower education, respectively. The random effects for the *Raumordnungsregion* and year

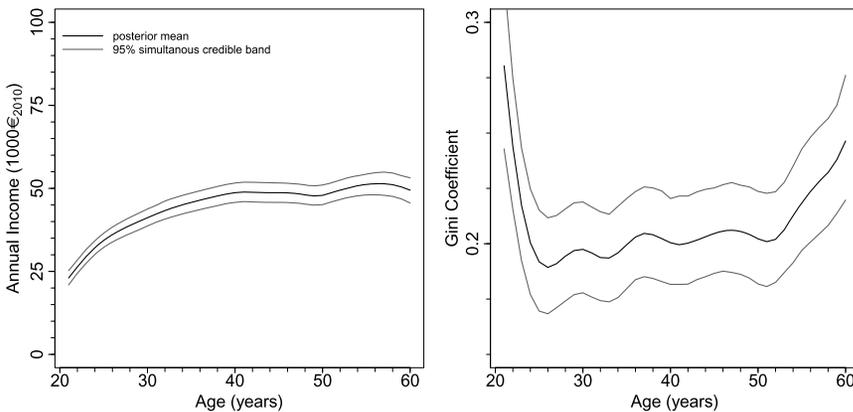


FIG. 7. SOEP data. Posterior means for expected income (left) and Gini coefficients (right) for males who have been working since the age of 21, without higher education and living in the West, together with 95% simultaneous credible bands.

are considered at zero, that is, their prior expectation. The grey lines indicate the 95% simultaneous credible bands. As expected, there is a general upward trend such that expected incomes are rising with increasing age. In addition, we see the concave structure that is generally also found by the literature.

For the Gini coefficient, we observe a U-shaped development over age. This indicates that the conditional income distribution is not simply rescaled over the age range but rather that it changes its shape such that the Gini coefficient rises. Our results are again in line with economic theory. At the very beginning of the career, income inequalities should be rather high, as large parts are still not yet allocated in accordance to their capabilities and, consequently, are employed and paid more or less arbitrarily. These mismatch-induced inequalities quickly fade away. From then on we would expect rising inequality, as following the classical theories on the shape of the unconditional income distribution [Arnold (2008)]; the latter is made up of incomes derived from a varying number of autoregressive permutations. These permutations, which would generally occur over the age range under consideration, would lead to a rising inequality in incomes with rising age.

Other quantities derived from conditional income distributions. Using distributional regression, it is easily possible to obtain estimates for certain quantiles, like the median, which is an alternative to the mean as a location measure. Furthermore, one can calculate interquantile ranges as an alternative measure of inequality. Naturally, such quantiles can be estimated in a more direct manner using quantile regression, although additional efforts may be required to avoid crossing quantile curves, in particular when considering a dense set of quantile levels. Distributional regression automatically avoids the problem of quantile crossing and makes model comparison easier in such situations. We contrast distributional regression against quantile regression in more detail for our case study in Section A.1.3.

Next to measures of inequality like the Gini coefficient or the Theil index, which are easily computable, it is also straightforward to calculate measures of polarisation, which have recently received considerable attention in the literature [for further references and explanations, see, e.g., Duclos, Esteban and Ray (2004), Wolfson (1994)]. Following Gradín (2000), it would be possible to calculate the polarisation between two groups as defined by sets of covariates.

It is also possible to assess density differences at different income levels or probability mass differences for different income ranges. For instance, one could consider the probability mass above a certain income, for example, 48,000€, which according to John Keynes would suffice to turn one's mind away from pecuniary worries [Skidelsky (2010)]. Consequently, it could be highlighted that not only the conditional mean income for the average man without higher education is lower in the East but also that the probability mass of incomes below that threshold is much lower. Such an analysis may be of particular interest for research questions on poverty and vulnerability [Pudney (1999)].

4.4. *Economic consequences.* Our findings show that keeping other variables fixed at their average level, there are significant differences in income inequalities within East and West Germany. Duclos, Esteban and Ray (2004) have noted the importance of within-group inequality for levels of alienation and identification within society. The higher income inequality in the East would thereby induce a weakened in-group identity. Lack of in-group identity in turn is likely to cause feelings of isolation and mistrust [Misztal (2013)], and thus leads to a deterioration of well-being which is beyond that captured by solely considering average incomes, or even distribution-adjusted well-being measures [Klasen (2008)].

While a profound analysis of the effect of different income distributions to well-being must be left for further research, our application shows that structured additive distributional regression offers a methodology to the analysis of income inequality which goes beyond the analysis at a highly aggregated level and thus allows to start the assessment of this important issue at a microeconomic level.

5. Conclusion. Distributional regression and the closely related class of GAMLSS provide a flexible, comprehensive toolbox for solving complex regression problems with potentially nonstandard response types. They are therefore useful to overcome the limitations of common mean regression models and to enable a proper, realistic assessment of regression relationships. In this paper, we provided a Bayesian approach to distributional regression and described solutions for the most important applied problems, including the selection of a suitable predictor specification and the most appropriate response distribution. Based on efficient MCMC simulation techniques, we developed a generic framework for inference in Bayesian structured additive distributional regression relying on distribution-specific iteratively weighted least squares proposals as a core feature of the algorithms.

Concerning the specific application of distributional regression to conditional income distributions, there are significant differences between men with similar age, work experience and education levels between East and West which go beyond the mean income. Taking the Gini coefficient as an indicator for inequality, income inequality among these men is larger in the East than it is in the West, further deepening differences in well-being. While this study highlights the scope of the new methodology to an application of income analysis and beyond, much work remains to be done on the application of distributional regression techniques.

Despite the practical solutions outlined in this paper, model choice and variable selection remain relatively tedious and more automatic procedures would be highly desirable. Suitable approaches may be in the spirit of Belitz and Lang (2008) in a frequentist setting or based on spike and slab priors for Bayesian inference as developed in Scheipl, Fahrmeir and Kneib (2012) for mean regression.

It will also be of interest to extend the distributional regression approach to the multivariate setting. For example, in the case of multivariate Gaussian responses, covariate effects on the correlation parameter may be very interesting in specific

applications. Similarly, multivariate extensions of beta regression lead to Dirichlet distributed responses representing multiple percentages that sum up to one; see Klein et al. (2015a) for a first attempt in this direction.

In the context of economic applications, it should be noted that, analogously to generalised linear models, the additive impact of explanatory variables on the economic measure of interest, like the Gini coefficient, is generally not attained. Consequently, the size, and possibly also the direction of the estimated spatial effect, may well be very different for different points in the covariate space. While it is straightforward to calculate these differences with corresponding credible intervals for any desired combination of other covariates to give a more comprehensive assessment of differences in inequality, further work needs to be done to facilitate the interpretation of results.

In addition, in-depth-testing is required to find adequate parametric forms for conditional income distributions, as the application of structured additive distributional regression crucially rests on the assumption that the parametric distribution fits the data. While for the case of full-time working men the Dagum distribution indeed seems to provide a decent fit, further work must be done to allow for an analysis with a less restricted covariate space and thus a more comprehensive analysis of income distributions in Germany and beyond.

Yet, this paper demonstrates that structured additive distributional regression offers a statistical framework addressing the challenge to assess entire conditional distributions [Fortin, Lemieux and Firpo (2011), page 56] by broadening the class of potential response distributions beyond simple exponential families and thus offers additional scope for applied statistical analyses on the problem of income inequality and beyond.

Acknowledgments. We thank two referees, the Associate Editor and Tilmann Gneiting for their careful review which was very helpful in improving upon the initial submission.

SUPPLEMENTARY MATERIAL

Supplement A: Case studies (DOI: [10.1214/15-AOAS823SUPPA](https://doi.org/10.1214/15-AOAS823SUPPA); .pdf). Additional material on the application to regional income inequality in Germany is provided in Section A.1. A second case study on the proportion of farm outputs achieved by cereals is treated in Section A.2.

Supplement B: Methodology (DOI: [10.1214/15-AOAS823SUPPB](https://doi.org/10.1214/15-AOAS823SUPPB); .pdf). This supplement comprises details on Bayesian inference, derivations of required quantities for the iteratively weighted least squares proposals and simulation studies.

REFERENCES

- ARNOLD, B. C. (2008). Pareto and generalized Pareto distributions. In *Modeling Income Distributions and Lorenz Curves* (D. Chotikapanich, ed.) 119–145. Springer, New York.

- ATKINSON, A. B. (1975). *The Economics of Inequality*. Clarendon Press, Oxford.
- AUTOR, D. H., KATZ, L. F. and KEARNEY, M. S. (2008). Trends in U.S. wage inequality: Revising the revisionists. *Rev. Econ. Stat.* **28** 300–323.
- BACH, S., CORNEO, G. and STEINER, V. (2009). From bottom to top: The entire income distribution in Germany, 1992–2003. *Rev. Income Wealth* **55** 303–330.
- BELITZ, C. and LANG, S. (2008). Simultaneous selection of variables and smoothing parameters in structured additive regression models. *Comput. Statist. Data Anal.* **53** 61–81. [MR2528592](#)
- BELITZ, C., BREZGER, A., KLEIN, N., KNEIB, T., LANG, S. and UMLAUF, N. (2015). BayesX-software for Bayesian inference in structured additive regression models. Version 3.0. Available at <http://www.bayesx.org>.
- BIEWEN, M. (2000). Income inequality in Germany during the 1980s and 1990s. *Rev. Income Wealth* **46** 1–19.
- BIEWEN, M. and JENKINS, S. P. (2005). A framework for the decomposition of poverty differences with an application to poverty differences between countries. *Empir. Econ.* **30** 331–358.
- BREZGER, A. and LANG, S. (2006). Generalized structured additive regression based on Bayesian P-splines. *Comput. Statist. Data Anal.* **50** 967–991. [MR2210741](#)
- CARD, D. E., HEINING, J. and KLINE, P. (2013). Workplace heterogeneity and the rise of German wage inequality. *Q. J. Bus. Econ.* **128** 967–1015.
- DAGUM, C. (1977). A new model of personal income distribution: Specification and estimation. *Economie Appliquée* **30** 413–437.
- DAGUM, C. (2008). A new model of personal income distribution: Specification and estimation. In *Modeling Income Distributions and Lorenz Curves* (D. Chotikapanich, ed.) 3–25. Springer, New York.
- DI NARDO, J., FORTIN, N. M. and LEMIEUX, T. (1996). Labor market institutions and the distribution of wages, 1973–1992: A semiparametric approach. *Econometrica* **64** 1001–1044.
- DONALD, S. G., GREEN, D. A. and PAARSCH, H. J. (2000). Differences in wage distributions between Canada and the United States: An application of a flexible estimator of distribution functions in the presence of covariates. *Rev. Econ. Stud.* **67** 609–633.
- DUCLOS, J.-Y., ESTEBAN, J. and RAY, D. (2004). Polarization: Concepts, measurement, estimation. *Econometrica* **72** 1737–1772. [MR2095531](#)
- DUNN, P. K. and SMYTH, G. K. (1996). Randomized quantile residuals. *J. Comput. Graph. Statist.* **5** 236–245.
- DUSTMANN, C., LUDSTECK, J. and SCHÖNBERG, U. (2009). Revisiting the German wage structure. *Q. J. Econ.* **124** 843–881.
- EILERS, P. H. C. and MARX, B. D. (1996). Flexible smoothing with B-splines and penalties. *Statist. Sci.* **11** 89–121. [MR1435485](#)
- FAHRMEIR, L., KNEIB, T. and LANG, S. (2004). Penalized structured additive regression for space-time data: A Bayesian perspective. *Statist. Sinica* **14** 731–761. [MR2087971](#)
- FAHRMEIR, L., KNEIB, T., LANG, S. and MARX, B. (2013). *Regression: Models, Methods and Applications*. Springer, Heidelberg. [MR3075546](#)
- FORTIN, N. M., LEMIEUX, T. and FIRPO, S. (2011). Decomposition methods in economics. In *Handbook of Labor Economics* (O. Ashenfelter and D. E. Card, eds.) **4A** 1–102. North-Holland, Amsterdam.
- FUCHS-SCHÜNDELN, N., KRUEGER, D. and SOMMER, M. (2010). Inequality trends for Germany in the last two decades: A tale of two countries. *Rev. Econ. Dyn.* **13** 103–132.
- GALVAO, A. F., LAMARCHE, C. and LIMA, L. R. (2013). Estimation of censored quantile regression for panel data with fixed effects. *J. Amer. Statist. Assoc.* **108** 1075–1089. [MR3174685](#)
- GAMERMAN, D. (1997). Sampling from the posterior distribution in generalized linear mixed models. *Stat. Comput.* **7** 57–68.

- GNEITING, T. (2011a). Making and evaluating point forecasts. *J. Amer. Statist. Assoc.* **106** 746–762. [MR2847988](#)
- GNEITING, T. (2011b). Quantiles as optimal point forecasts. *Int. J. Forecast.* **27** 197–207.
- GNEITING, T. and RAFTERY, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *J. Amer. Statist. Assoc.* **102** 359–378. [MR2345548](#)
- GNEITING, T. and RANJAN, R. (2011). Comparing density forecasts using threshold- and quantile-weighted scoring rules. *J. Bus. Econom. Statist.* **29** 411–422. [MR2848512](#)
- GRADÍN, C. (2000). Polarization by sub-populations in Spain, 1973–1991. *Rev. Income Wealth* **46** 457–474.
- GREENE, W. H. (2008). *Econometric Analysis*, 6th ed. Pearson Prentice Hall, Upper Saddle River.
- HASTIE, T. and TIBSHIRANI, R. (1993). Varying-coefficient models. *J. Roy. Statist. Soc. Ser. B* **55** 757–796. [MR1229881](#)
- HELLER, G., STASINOPOULOS, D. and RIGBY, R. (2006). The zero-adjusted inverse Gaussian distribution as a model for insurance data. In *Proceedings of the 21th International Workshop on Statistical Modelling* (J. Hinde, J. Einbeck and J. Newell, eds.) 226–233.
- KLASSEN, S. (2008). The efficiency of equity. *Rev. Polit. Econ.* **20** 257–274.
- KLEIBER, C. (1996). Dagum vs. Singh–Maddala income distributions. *Econom. Lett.* **57** 39–44.
- KLEIBER, C. and KOTZ, S. (2003). *Statistical Size Distributions in Economics and Actuarial Sciences*. Wiley, Hoboken, NJ. [MR1994050](#)
- KLEIN, N., DENUIT, M., LANG, S. and KNEIB, T. (2014). Nonlife ratemaking and risk management with Bayesian generalized additive models for location, scale, and shape. *Insurance Math. Econom.* **55** 225–249. [MR3179814](#)
- KLEIN, N., KNEIB, T. and LANG, S. (2015). Bayesian generalized additive models for location, scale and shape for zero-inflated and overdispersed count data. *J. Amer. Statist. Assoc.* **110** 405–419. [MR3338512](#)
- KLEIN, N., KNEIB, T., KLASSEN, S. and LANG, S. (2015a). Bayesian structured additive distributional regression for multivariate responses. *J. R. Stat. Soc. Ser. C. Appl. Stat.* To appear.
- KLEIN, N., KNEIB, T., LANG, S. and SOHN, A. (2015b). Supplement to “Bayesian structured additive distributional regression with an application to regional income inequality in Germany.” DOI:10.1214/15-AOAS823SUPPA.
- KLEIN, N., KNEIB, T., LANG, S. and SOHN, A. (2015c). Supplement to “Bayesian structured additive distributional regression with an application to regional income inequality in Germany.” DOI:10.1214/15-AOAS823SUPPB.
- KOENKER, R. (2005). *Quantile Regression. Econometric Society Monographs* **38**. Cambridge Univ. Press, Cambridge. [MR2268657](#)
- KOENKER, R. and BASSETT, G. JR. (1978). Regression quantiles. *Econometrica* **46** 33–50. [MR0474644](#)
- KOHN, K. and ANTONCZYK, D. (2011). The aftermath of reunification: Sectoral transition, gender, and rising wage inequality in East Germany. IZA Discussion paper series No. 5708. Available at <http://hdl.handle.net/10419/51717>.
- LAIO, F. and TAMEA, S. (2007). Verification tools for probabilistic forecasts of continuous hydrological variables. *Hydrol. Earth Syst. Sci.* **11** 1267–1277.
- LANG, S. and BREZGER, A. (2004). Bayesian P-splines. *J. Comput. Graph. Statist.* **13** 183–212. [MR2044877](#)
- LANG, S., UMLAUF, N., WECHSELBERGER, P., HARTTGEN, K. and KNEIB, T. (2014). Multilevel structured additive regression. *Stat. Comput.* **24** 223–238. [MR3165550](#)
- LEMIEUX, T. (2006). The “Mincer equation.” Thirty years after Schooling, Experience, and Earnings. In *Jacob Mincer: A Pioneer of Modern Labor Economics* (S. Grossbard, ed.) 127–145. Kluwer Academic, Boston.
- MINCER, J. (1974). *Schooling, Experience, and Earnings*. Columbia Univ. Press, New York.

- MISZTAL, B. (2013). *Trust in Modern Societies: The Search for the Bases of Social Order*. Wiley, Hoboken.
- MORDUCH, J. and SICULAR, T. (2002). Rethinking inequality decomposition, with evidence from rural China. *Econ. J.* **112** 93–106.
- NEWKEY, W. K. and POWELL, J. L. (1987). Asymmetric least squares estimation and testing. *Econometrica* **55** 819–847. [MR0906565](#)
- OSBAND, K. and REICHELSTEIN, S. (1985). Information-eliciting compensation schemes. *J. Public Econ.* **27** 107–115.
- PIKETTY, T. and SAEZ, E. (2007). Income and wage inequality in the United States, 1913–2002. In *Top Incomes Over the Twentieth Century* (A. B. Atkinson and T. Piketty, eds.) 141–225. Oxford Univ. Press, Oxford.
- PUDNEY, S. (1999). On some statistical methods for modelling the incidence of poverty. *Oxf. Bull. Econ. Stat.* **61** 385–408.
- RIGBY, R. A. and STASINOPOULOS, D. M. (2005). Generalized additive models for location, scale and shape. *J. Roy. Statist. Soc. Ser. C* **54** 507–554. [MR2137253](#)
- RUE, H. and HELD, L. (2005). *Gaussian Markov Random Fields: Theory and Applications*. *Monographs on Statistics and Applied Probability* **104**. Chapman & Hall/CRC, Boca Raton, FL. [MR2130347](#)
- RUPPERT, D., WAND, M. P. and CARROLL, R. J. (2003). *Semiparametric Regression*. *Cambridge Series in Statistical and Probabilistic Mathematics* **12**. Cambridge Univ. Press, Cambridge. [MR1998720](#)
- SALEM, A. B. Z. and MOUNT, T. D. (1974). A convenient descriptive model of income distribution. *Econometrica* **42** 1115–1127.
- SARABIA, J. M. (2008). Parametric Lorenz curves: Models and applications. In *Modeling Income Distributions and Lorenz Curves* (D. Chotikapanich, ed.) 167–190. Springer, New York.
- SCHEIPL, F., FAHRMEIR, L. and KNEIB, T. (2012). Spike-and-slab priors for function selection in structured additive regression models. *J. Amer. Statist. Assoc.* **107** 1518–1532. [MR3036413](#)
- SCHNABEL, S. K. and EILERS, P. H. C. (2009). Optimal expectile smoothing. *Comput. Statist. Data Anal.* **53** 4168–4177. [MR2744314](#)
- SILBER, J. (1999). Introduction—thirty years of intensive research on income inequality measurement. In *Handbook of Income Inequality Measurement* (J. Silber, ed.) 1–18. Kluwer Academic, Boston.
- SKIDELSKY, R. (2010). *Keynes: The Return of the Master*, 1st ed. Public Affairs, New York.
- SOBOTKA, F. and KNEIB, T. (2012). Geoadditive expectile regression. *Comput. Statist. Data Anal.* **56** 755–767. [MR2888723](#)
- SPIEGELHALTER, D. J., BEST, N. G., CARLIN, B. P. and VAN DER LINDE, A. (2002). Bayesian measures of model complexity and fit. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **64** 583–639. [MR1979380](#)
- STATISTISCHES BUNDESAMT (2012). Verbraucherpreisindizes für Deutschland—Lange Reihen ab 1948, Preise.
- UMLAUF, N., KLEIN, N., LANG, S. and ZEILEIS, A. (2014). *bamlss: Bayesian additive models for location scale and shape (and beyond)*. R package Version 0.1-1. Available at <http://bayesr-forge.r-project.org>.
- WAGNER, G. G., FRICK, J. R. and SCHUPP, J. (2007). The German socio-economic panel study (SOEP)—scope, evolution and enhancements. *Schmollers Jahrbuch* **127** 139–169.
- WOLFSON, M. C. (1994). When inequalities diverge. *Am. Econ. Rev.* **84** 353–358.
- WOOD, S. N. (2004). Stable and efficient multiple smoothing parameter estimation for generalized additive models. *J. Amer. Statist. Assoc.* **99** 673–686. [MR2090902](#)
- WOOD, S. N. (2008). Fast stable direct fitting and smoothness selection for generalized additive models. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **70** 495–518. [MR2420412](#)

YU, K. and MOYEED, R. A. (2001). Bayesian quantile regression. *Statist. Probab. Lett.* **54** 437–447.
MR1861390

N. KLEIN
T. KNEIB
A. SOHN
GEORG-AUGUST-UNIVERSITY GÖTTINGEN
CHAIR OF STATISTICS
PLATZ DER GÖTTINGER SIEBEN 5
37073 GÖTTINGEN
GERMANY
E-MAIL: nklein@uni-goettingen.de

S. LANG
UNIVERSITY OF INNSBRUCK
DEPARTMENT OF STATISTICS
UNIVERSITÄTSSTRASSE 15
6020 INNSBRUCK
AUSTRIA