Courant Research Centre
‘Poverty, Equity and Growth in Developing and Transition Countries: Statistical Methods and Empirical Analysis’

Georg-August-Universität Göttingen
(founded in 1737)

Discussion Papers

No. 231

Reconsidering the Income-Illness Relationship using Distributional Regression
An Application to Germany

Alexander Silbersdorff, Julia Lynch, Stephan Klasen, Thomas Kneib

June 2017

Platz der Göttinger Sieben 5 · 37073 Goettingen · Germany
Phone: +49-(0)551-3921660 · Fax: +49-(0)551-3914059

Email: crc-peg@uni-goettingen.de Web: http://www.uni-goettingen.de/crc-peg

Reconsidering the Income-Illness Relationship using Distributional Regression
An Application to Germany

Alexander Silbersdorff[1], Julia Lynch[2], Stephan Klasen[1] and Thomas Kneib[1]

[1]Georg August University of Göttingen, Germany.
[2]University of Pennsylvania, USA.

Abstract

In this paper we reconsider the relationship between income on health, taking a distributional perspective rather than one centered on conditional expectation. Using Structured Additive Distributional Regression, we find that the association between income on health is larger than generally estimated because aspects of the conditional health distribution that go beyond the expectation imply worse outcomes for those with lower incomes. Looking at German data from the Socio Economic Panel, we find that the risk of very bad health is roughly halved when doubling the net equivalent income from 15,000 Euro to 30,000 Euro, which is more than tenfold of the magnitude of change found when considering expected health measures. This paper therefore argues that when studying health outcomes, a distributional perspective that considers stochastic variation among observationally equivalent individuals is warranted.

JEL-Classification: I14, C13, C21

1 Introduction

Scores of papers assess the relationship between income and health status in a multivariate framework. Both in epidemiology and health economics, the vast majority of these employ standard regression methods (linear and generalized linear models) to assess the effect of variations of income and other covariates on the expectation of health status. However, in recent years a growing number of papers in the health economics literature have noted the need to look beyond the expected outcome (Duclos and Échevin, 2011; Makdissi and Yazbeck, 2014; Carrieri and Jones, 2016; Heckley et al., 2016; Schiele and Schmitz, 2016).

Building on this young literature, we propose the use of distributional regression models which allow for the estimation of a full distribution for a given set of covariates – both for continuous and discrete health measures. Rather than focussing on health outcomes like the conditional expectation for some health score, or the conditional odds for some binary health measure, we thus propose to consider the full conditional health distribution. Concretely we propose to apply the recently developed technique of structured additive distributional regression (SADR)(Klein et al., 2015) to estimate the relationship between self-reported health status and income, conditional on other standard covariates. Using SADR, we are able to look at both categorical health measures (like the standard ordered 5-response format) and continuous or quasi-continuous variables (like the SF-12). Thus, SADR provides a generic regression framework allowing for a distributional perspective of the relationship between health and other covariates. In contrast to some of the recent literature, which uses standard regression of the recentered influence function of some explanatory variable for the entire population's health distribution (e.g. Carrieri and Jones, 2016; Heckley et al., 2016), SADR allows us to focus on the conditional distribution of health for specific sub-populations defined by income, age, education, or the like.

As we show in the paper, this distributional approach leads to a starkly different assessment of the magnitude of the association between income and health, when controlling for a set of other covariates. For example, we contrast the shift in the risk of being severely ill with the change in the expected health for "average Joe" and "average Jane" that is associated with moving from a net equivalent household income of 15,000€ (the median income of the poorer half of the population) to an income of 30,000€ (the median income of the richer half of the population). Using the distributional approach, we find that the relative change in the risk of severe illness is between 39 percent and 42 percent, while the relative change of expected health is only 3 percent. Distributional regression facilitates a shift in perspective that highlights the substantively large association between income and health at the lower end of the conditional health distribution, which might be missed if only effects on the expectation of the conditional health distribution were examined.

The contribution of the paper to the literature is therefore twofold. On the one hand, we introduce a generic regression framework to the health literature that allows for the distributional analysis of both discrete and continuous health variables. On the other hand, the paper shows how a shift in perspective beyond the classical mean leads to a quite different assessment of magnitude of the association between income and health.

The remainder of this paper is structured as follows: In the next section, we briefly outline the literature on the relationship between income and health, and motivate our analysis of outcomes beyond the expectation. The following section explains how SADR can be used to analyze conditional health distributions. In the subsequent section, we apply the approach to health data from the 2012 wave of the German Socio-Economic Panel (SOEP), modeling the relationship between both a discrete health score (self-rated health) and a quasi-continuous health score (SF-12) on the one hand, and net equivalent household income on the other. In both cases we control for a set of other variables like age, education, etc. We next illustrate the importance of taking a distributional perspective by highlighting several distributional measures that put the emphasis on health impacts at the lower end of the spectrum. In the fifth and final section, we conclude.

2 Literature and Motivation

The association between income and health is one of the most robustly documented findings in the literatures on population health and health economics (Marmot, 2002; Kawachi et al., 2010). Income has been found to be strongly associated with measures of health across a variety of populations, even above a threshold of material deprivation (Backlund et al., 1996; Ettner, 1996; McDonough et al., 1997; Ecob and Davey Smith, 1999; Case, 2001), and recent studies exploiting exogenous variation in income have discussed causal effects of income on health (Frijters et al., 2005; Lindahl, 2005; Case, 2001; Kuehnle, 2014; Cesarini et al., 2016).

Given the intricate nature of the causal mechanisms linking income to health, the magnitude of the association between health and income remains of critical concern for contemporary political decision making, particularly in areas such as the minimum wage, social minimum, or tax treatment of low earnings. In the literatures on health economics, epidemiology and public health, estimates of the relationship between income and health have tended to take one of three forms: bivariate concentration indices summarizing the relationship between income and health in a population (Wagstaff and van Doorslaer, 1994; Lynch and Kaplan, 1997; Gravelle, 1998; Ecob and Davey Smith, 1999; Humphries and van Doorslaer, 2000; Gravelle, 2003; Lindahl, 2005; Wagstaff, 2005, 2011); estimates of the effect of income and other covariates on mean health status (Rogot et al., 1992; Ettner, 1996; Case, 2001; Contoyannis et al., 2004), and likelihood ratios that express the conditional probability of being in a particular health state given a particular level of income and other covariates (Benzeval et al., 2000; Frijters et al., 2005). The first two rely on health measures that are plausibly interpreted as continuous, while the latter technique is often used when the health outcome in question is measured using discrete (often binary) categories. Each of these approaches to measuring the relationship between income and health is useful, but, particularly when applied to the most widely used survey measures of health status, also has limitations.

Concentration indices are the “workhorse [method] in most health economic studies” (Fleurbaey and Schokkaert, 2009, p.73) for quantifying the distribution of health in a population. Their succinct form and resemblance to the Gini coefficient provide an intuitive scalar measure that is well suited for portraying the magnitude of health inequalities related to socio-economic characteristics (Wagstaff and van Doorslaer, 1994; Kakwani et al., 1997). Yet the concentration index by its construction only allows for the relation of two variables or scales. Estimating the relationship between health and income while jointly conditioning on a set of other covariates thought also to be relevant for health is thus not possible using conventional concentration indices. While we reconsider concentration indices later on, we leave this methodology aside for the moment and concentrate on those analyses of the health-income nexus which employ regression methodology to estimate the income-health relation while controlling for the effects of other covariates.

A second workhorse method, particularly prominent in the epidemiological literature, is the analysis of expected health outcomes expressed as odds ratios. In its most prevalent form, a logit model is used to predict outcomes on a binary health measure conditioning on a set of variables. One of the advantages of this method is that risk of having a particular outcome is easily related to the conditional expectation derived from the model. Due to the simplicity and popularity of this method, it is common practice to reduce health variables of higher complexity to a binary form in order to facilitate the construction of odds ratios (Chamberlain, 1980; Benzeval et al., 2000; Frijters et al., 2005). This reduction is problematic if the health outcome of interest does not adhere to the implicit assumptions required by such a reduction, including indifference among health outcomes grouped together in either of the dichotomous response categories. For some health measures, including general health status, a dichotomous representation of healthy/unhealthy is clearly insufficient as important variations would be disregarded.

Possibly for this reason, the construction of more fine-grained scales has become widespread for the analysis of health for various contexts. The most popular approach is to construct (quasi-)continuous health measures. These can subsequently be analyzed using simple mean regression models, like OLS. While these classical regression techniques have the capacity to generate information about the relationship between quasi-continuous health outcomes and other covariates, in practice the reported results generally attend solely to the conditional expectation derived from these models. As is the case with logit modeling described above, this reduction is problematic as potentially important variations beyond the mean are disregarded. For example, looking only at the mean health outcome conditional on covariates ignores research on the utility associated with varying health statuses, some of which suggests that an equal-sized change in health status above or below the mean may in practice generate asymmetric changes in well-being (Finkelstein et al., 2009).

One problem shared by both regression approaches is thus that they use only limited information from the full distribution of health in the sample, either by dichotomizing the outcome from the outset or by considering only the conditional expectation of the outcome. Through such information reduction, these approaches focus attention on one particular aspect of the relationship

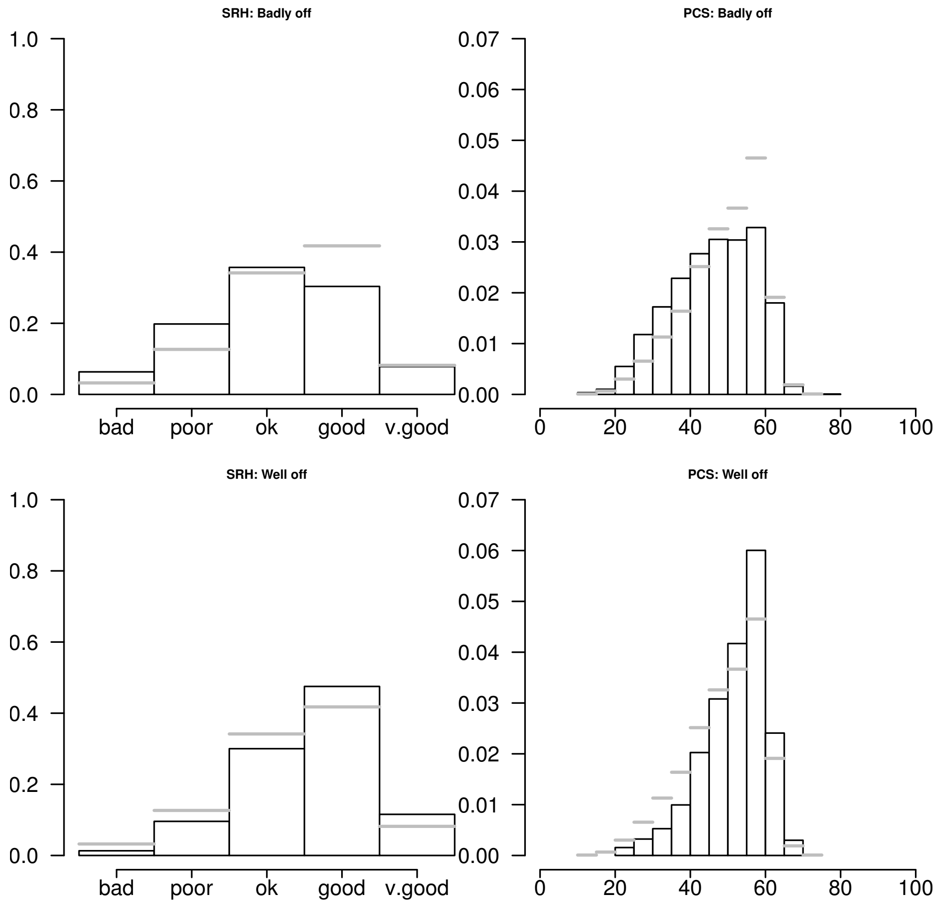

Figure 1: A contrast of coarsely conditioned health distributions. Top: Lowest 20% of net incomes. Bottom: Highest 20% of net incomes. Left: self-rated health (SRH). Right: Physical component score (PCS) of the SF-12. Grey lines indicate reference lines of middle group.

between health and income (and/or other covariates) at the cost of ignoring other potentially important changes in the health distribution that may occur in connection with changes in income (and/or other covariates). While this narrowed perspective is adequate and indeed necessary in many scenarios, properly estimating the effect of income on health requires a broader approach. Using data from Germany in 2012, a simple trisection of the health distribution between those with high, medium and low incomes reveals that the difference in the health outcomes by income goes beyond differences in the mean. Figure 1 shows the distribution of two measures of generalized health – self-rated health (SRH) and a physical health score (PCS) – among those with high (top 20% w.r.t. net equivalent income) and low (bottom 20% w.r.t. net equivalent income) incomes. The variation in health outcomes is substantially more pronounced in the lower part of the income distribution, while those who are economically well off are able to practically eliminate the risk of bad health. While an assessment based on the distributions' means captures the general trend of the health-income relationship, the reduction in information incurred by focussing on the mean leads us to neglect potentially important aspects of the relationship between income and health.

There are also theoretical reasons for studying the distribution, rather than just the mean, of health outcomes. At a societal level, the health-utility relationship is concave rather than linear, a characteristic that can be deduced from the fact that far more health care resources (spending) are dedicated to improvements at the lower end of the health spectrum than to improvements at the higher end (Berk and Monheit, 2001). If the health-utility relationship is concave, mean-based assessment premised on a linear relationship would not be warranted. Equally importantly, if the health-utility relationship is concave, analyses of the association between income (or other covariates) and health ought to give more weight to the lower end of the health distribution. For both multicategorical and continuous health measures, we thus propose the use of a set of risk measures which focus explicitly on the lower end of the conditional health distribution to complement conventional measures focused at the center of the conditional health distribution.

3 Taking a Distributional Perspective

The conventional regression approaches discussed above fall into the category of generalized linear models, where the conditional expectation of a health outcome variable Y given a set of explanatory variables $x_1, \dots, x_K$ is related to a regression predictor η via the response function h, i.e.

$$\mu = \mathbb{E}(Y \mid x_1, \dots, x_K) = h(\eta).$$

The predictor in turn is usually modeled as a linear combination of the covariates[1] entailed in covariate vector $(x_1, \dots, x_K)^T$, i.e.

$$\eta = \beta_0 + \sum_{k=1}^{K} \beta_k x_k.$$

For example, in case of binary outcomes differentiating only between healthy and non-healthy individuals, a logit or probit model is specified, in which the probability of an outcome $\pi = P(Y = 1 \mid x_1, \dots, x_K) = \mathbb{E}(Y \mid x_1, \dots, x_K)$ is related to the predictors via the cumulative distribution function of the logistic and the standard normal distribution, respectively.

The most important feature of generalized linear models for our purposes is that they focus exclusively on modeling the expectation of the response variable. Unlike in the case of binary responses, where the distribution of the health outcome is completely determined by the expectation (i.e. the success probability), when outcomes are more complex the expectation alone generally does

[1] More flexible alternatives have been developed in the context of generalized additive models (see Hastie and Tibshirani, 1990) or structured additive regression models (see Fahrmeir et al., 2004), but we will restrict ourselves to linear predictors in the following.

not represent the complete distribution of the health outcomes well. We will analyze both multicategorical and continuous measures for health outcomes and in these cases the deviations from the expectation are typically at least as important as determinants of expected health. More importantly, these deviations may also be driven by covariates such that more general features of the health outcome distribution such as variance and skewness should also be modeled in terms of regression predictors.

A distributional perspective is needed to allow us to not just consider the conditional expectation of the health variable of interest, $\mathbb{E}(Y \mid x_1, \ldots, x_K)$, but also to relate the complete underlying conditional distribution, $\mathcal{D}(Y \mid x_1, \ldots, x_K)$ to the covariates. The evolution of computation capacity in the past decades has made the estimation of distributional regression models feasible and several approaches have been put forward in the statistical literature. For example, one could use quantile regression, as proposed by Koenker and Bassett (1978), to construct the distribution from the conditional quantiles, or use recentred influence functions (Firpo et al., 2009) that allow for the estimation of unconditional quantiles. Alternatively, conditional transformation models as proposed by Hothorn et al. (2014) or the related distributional regression models proposed by Chernozhukov et al. (2013) could be used.

Here, we will rely on structured additive distributional regression (SADR) models as introduced in Klein et al. (2015), in which a parametric distribution type is assumed for the conditional distribution $\mathcal{D}(Y \mid x_1, \ldots, x_K)$, but all parameters (not only the mean) are then related to regression predictors based on a suitably chosen response function. More specifically, we assume that the conditional distribution $\mathcal{D}(\theta_1(x_1, \ldots, x_K), \theta_2(x_1, \ldots, x_K), \ldots, \theta_L(x_1, \ldots, x_K))$ is characterised by a vector of L parameters $\theta_l(x_1, \ldots, x_K)$, $l = 1, \ldots, L$, and specify

$$g_l(\theta_l) = \eta^{\theta_l} \tag{1}$$

$$\eta^{\theta_l} = \beta_0^{\theta_l} + \sum_{k=1}^{K} \beta_k^{\theta_l} x_k. \tag{2}$$

Consequently, the vector of all regression coefficients $\boldsymbol{\beta}$ entails parameters not only for one predictor but for all L predictors required to specify the response distribution. The main advantage of SADR is that it lends great estimation stability which is critical for the usually available sample sizes.[2] The main disadvantage is of course the need for an adequate parametric response distribution. For further discussion of the advantages and disadvantages of SADR over the alternative methods, we refer the interested reader to section A.2 in the appendix.

[2]Given that distributional assessments are generally much more demanding than simple mean assessments, standard database sizes in the order of ten thousand observations or less suffer from estimation instability already once a set of ten or more covariates is introduced.

4 A Distributional Health Assessment for Germany

To illustrate the difference between a distributional perspective and conventional estimation methodologies, we consider a very simple application using health data from the German Socio-Economic Panel (SOEP, 2014).

4.1 The German Socio-Economic Panel

The German Socio-Economic Panel (SOEP) is a longitudinal household survey repeated annually since 1984 (Wagner et al., 2007). For this study we use only the cross-sectional data from the 2012 survey, which contains information on over 10,000 households (see SOEP, 2014; Rahmann and Schupp, 2013). The SOEP contains a rich array of sociodemographic information about individuals in these households, as well as several measures of health status. In this paper we consider both the standard five-response self-rated health item and the SF-12 physical health scale, as representative ordinal and (quasi-)continuous health measures, respectively. Thus we show that our proposed perspective is feasible for both discrete and continuous variables, both of which are frequently used in the literature. Indeed, as we will show, our proposed perspective which focusses on the poor yields similar outcomes irrespective of whether we use self-rated health or the SF-12 physical health scale. In the following, both health measures are related to a set of sociodemographic variables that are standard in the literature (see below). Using only those adult individuals for whom we have full information on these variables (see below), the 2012 SOEP yields 16,723 observations: 7,820 males and 8,903 females.

4.1.1 Self-rated health

In social epidemiological research, the most commonly used indicator of health status is generalized self-rated health (SRH), captured in a single item with a Likert response scale: “How would you describe your current health¿‘: “Very good”, “good”, “satisfactory”, “poor” or “bad”. Single-item SRH measures have been found in multiple populations to be reliable and responsive to changes in health status, and to predict health expenditure and outcomes (Idler and Benyamini, 1997; DeSalvo et al., 2006). Because well-being is intimately tied to one’s sense of identity, single-item measures can tap respondents’ ability to identify whether or not they are healthy quickly and holistically, and drawing on information that may not (yet) be available to their physicians or to researchers as diagnoses of specific conditions (Benyamini, 2011). DeSalvo et al. (2005, 2009) compare a standard single-item SRH measure to more comprehensive batteries and find that despite its brevity and simplicity, the single-item SRH is equally useful for predicting mortality, health care utilization, and health expenditures.

4.1.2 The SF12

Every two years since 2002, the SOEP has included a battery of health-related questions, the "SF-12v2^{TM} Health Survey" (SF-12 Wagner et al., 2007). The SF-12 is a 12-item subset of Quality Metric's SF-36v2^{TM}, which is used widely in the recent literature (e.g. Marcus, 2013; LaMontagne et al., 2014; Eibich and Ziebarth, 2014) and provides measures of self-rated health in eight domains. The SF-12 comprises 12 items that aim to capture "practical, reliable and valid information about functional health and well-being from the patient's point of view" (OPTUM, 2015). Principal component analysis is used to compute two superordinate scales on physical health (PCS) and mental health (MCS), designed to have a mean of 50 and a standard deviation of 10. See Andersen et al. (2007) for details on the computation.

The SF-12 is an alternative to the longer SF-36 and to single-item measures of general self-rated health (SRH). The SF-12 has been found to be reliable, internally consistent, and to have good convergent and discriminant validity (Gandek et al., 1998; Franks et al., 2003; Bohannon et al., 2004; Cunillera et al., 2010). Across a variety of health outcomes and countries and with different patient populations, the SF-12 predicts physical and mental health outcomes, health related quality of life, and medical expenditure (see Ware Jr. et al., 1996; Fleishman et al., 2006). In cross-sectional and longitudinal tests of validity, the SF-12 generally yielded larger standard errors than the SF-36 (Ware Jr. et al., 1996). Nevertheless, the SF-12 is a practical and widely accepted tool for measuring population health and for predicting health outcomes and expenditure. It has also been found to map reliably onto the EQ-5D scale, a five-dimension health status indicator that is commonly used in generating the preference weightings needed to construct QALYs and other similar measures (Brazier and Roberts, 2004; Lawrence and Fleishman, 2004; Gray et al., 2006). The SF-12 has been found in previous studies to be correlated with income in a general population, even after adjusting for relevant covariates (Burdine et al., 2000; Schnittker, 2004; König et al., 2010). In our analysis we use only the PCS subscale of the SF-12. Differential item functioning by education, age and sex has been observed for the MCS (Fleishman and Lawrence, 2003; Bourion-Bédès et al., 2015), and since the SOEP does not include the institutionalized population, the sample is likely to be non-representative of the population with very low MCS scores.

4.1.3 The Explanatory Variables

We base our choice of explanatory variables on the applied literature on individuals' health production functions (e.g. Fayissa and Gutema, 2005; Lorgelly and Lindley, 2008; Thornton, 2010;

Ravesteijn et al., 2013). In this literature health is modeled as a function of economic, social and environmental factors. Here, we focus on a simple set-up that considers six easily observable covariates.

The main explanatory variable of interest in our model is disposable income, measured as the annual net equivalized household income of an individual, adjusted for household size and composition using the OECD equivalence scale. Following Jones and Wildman (2005), we use the log transformation of income (LOGINC).
In addition to income, we consider the respondent's age as a quadratic polynomial (AGE and AGESQ), to control for differences in health induced by the inevitable biologically-induced deterioration of health over the life course (see Kiuila and Mieszkowsko, 2007).
To adjust for the well known relationship between education (or cultural capital in a broader sense) and health, we control for respondent's educational attainment measured using the ISCED97 education categories provided by the SOEP. Here, we use four education levels. The first level (EDU_1) includes all individuals who have only general elementary education or less (i.e. those whose ISCED is between 0 and 2). The second level (EDU_2) entails all persons with completed secondary education (i.e. ISCED level 3) while the third level (EDU_3) entails all with ISCED levels 4 and 5, i.e. vocational training with *Abitur* or higher vocational training. The highest level (EDU_4) entails all those with completed higher education (i.e. ISCED level 6).
The variable education is complemented by a variable measuring whether the respondent is a German national (GER). This variable is sought to control for some additional cultural differences as well as statistical artefacts like the healthy migrant effect (see Bjornstrom and Kuhl, 2014).
We also account for the marital status of the respondent: Living in a partnership (married or living together) (MAR_1), separated or divorced (MAR_2), single (MAR_3), or widowed (MAR_4).
Lastly, we control for the region of the country in which respondents live. This allows us to control for some environmental factors like the general prosperity of the area of residence and associated aspects like health care infrastructure. We use a hierarchical regional effect which accounts for differences between former West and East Germany and subsequently differentiates between individual federal states of residence, as in Eibich and Ziebarth (2014).

For further information on the variables see Section A.1 in the appendix.

4.2 Model specification

4.2.1 Choice of the Response Distribution

As discussed in Section 3, a distributive regression approach requires that we specify a suitable parametric distribution that is able to approximate the empirically observed conditional health distributions.

Self-rated health outcomes are measured on an ordinal five point scale, which means that their distribution can be characterized by four probability parameters. We use a sequence of logit models to differentiate between the five levels of the self-rated health score rather than to differentiate only between two amalgamations of the levels as is standard in the literature. We first regress the lowest response versus all higher health scores to differentiate low values of the score from all higher scores. In the second step, we consider only individuals that reached at least the second response level of the discrete health measure and contrast the second level it to all higher levels. Continuing this sequence for higher levels provides us with a set of sequential logit models that characterize the multinomial nature of the categorical health outcome while simultaneously acknowledging the ordinal structure in a simple and interpretable fashion.[3]

Scores on continuous health measures, such as the SF-12, generally deviate significantly from a symmetric distribution, such that regression specifications based on the normal distribution often do not provide sufficient flexibility. For the PCS, we find that the conditional health distributions generally feature a negative skewness and are thus in contras to the more common symmetric or positively skewed distributions for which which most parametric formulations are tailored. To be able to employ well-established estimation routines for the standard parametric distributions, we follow Erreygers and van Ourti (2011) and use a linear transformation g_{PCS} of the health score

$$g_{PCS}(H) = H^* = \frac{(H_0 - H)}{H_{scale}}, \tag{3}$$

where H and H^* denote the untransformed and the transformed PCS health score respectively, while H_0 is a constant ensuring that H^* has a positive support if required. Lastly, H_{scale} is another constant rescaling the transformed health score. In the following, we will use $H_0 = 100$ and $H_{scale} = 10$ ensuring that our transformed health score is not only positive but also restricted to the interval $(0, 10)$ which enhances numerical stability. Subsequently, we estimate the conditional distributions of the transformed PCS using the well-known two parameter gamma distribution.[4] Once this conditional distribution is estimated, one can easily obtain the conditional distribution of the original PCS measure by simply applying the inverse transform, g^{-1}_{PCS}. Note that the gamma

[3] Standard cumulative regression models for ordinal responses would be much more limited in their flexibility since they would restrict covariate effects to be the same for the transition between all different stages of the response.

[4] Using a representation of the gamma distribution where μ is the expectation parameter and s the shape parameter, we can write its density as:

$$p(y \mid \mu, s) = \Big(\frac{s}{\mu}\Big)^s \frac{y^{s-1}}{\mathrm{G}(s)} \exp\Big(-\frac{s}{\mu} y\Big), \tag{4}$$

where y denotes the transformed PCS outcome, which is H^* in our case, and where G denotes the Gamma function.

distribution is invariant under scaling such that we effectively model a shifted, reversed, scaled gamma distribution for the health scores.

For both the categorical self-rated health scores and the (quasi) continuous SF-12, we thus specify parametric conditional health distributions which require, respectively, four and two parameters to be estimated with respect to the covariates. With the two distribution types chosen, let us now turn to the specification of the predictors of the distributions' parameters.

4.2.2 Predictor specification

Let us now turn to the specification of predictors For the sake of simplicity, we will specify one generic predictor set-up which is applied to all parameters, i.e.

$$\begin{aligned}\eta_l = \beta_0^{\theta_l} &+ \beta_1^{\theta_l}\mathrm{AGE} + \beta_2^{\theta_l}\mathrm{AGE}SQ + \beta_3^{\theta_l}\mathrm{LOGINC} + \beta_4^{\theta_l}\mathrm{GER} + \beta_5^{\theta_l}\mathrm{EDU}_2 + \beta_6^{\theta_l}\mathrm{EDU}_3 + \beta_7^{\theta_l}\mathrm{EDU}_4 \\ &+ \beta_8^{\theta_l}\mathrm{MAR}_2 + \beta_9^{\theta_l}\mathrm{MAR}_3 + \beta_{10}^{\theta_l}\mathrm{MAR}_4 + \beta_{11}^{\theta_l}\mathrm{EAST} + \gamma_{\mathrm{DISTRICT}}^{\theta_l} \qquad (5)\end{aligned}$$

where η_l is the predictor for for the lth parameter of the response distribution. The explanatory variables (defined as outlined in Section 4.1.3) are all included in a linear fashion, supplemented by two effects representing spatial variation in health outcomes. EAST is an effect-coded binary variable scored one if the federal state is in the east of Germany, thus capturing the structural differences between the former German Democratic Republic (GDR) and the Federal German Republic (FDR). The differences within the former GDR and FDR are captured by random effects, denoted by $\gamma_{\mathrm{DISTRICT}}$. This regularizing approach is chosen over a plain use of fixed effects for all federal states in order to enhance estimation stability (see Klein et al., 2015).

In order to relate the predictors to their corresponding parameters, we specify appropriate response functions. For the categorical responses, these are simply given by logit response functions while the exponential response function is used to ensure positivity of the two parameters for the gamma distribution.

4.3 Parameter Estimates

The estimation is done in the software BayesX (Belitz et al., 2015) which employs Markov Chain Monte Carlo (MCMC) simulation techniques to estimate posterior distributions in a Bayesian framework. See Klein et al. (2015) for details on the estimation procedure. In the following set-up, we use non-informative flat priors for the linear effect. For the spatial effect, we use Gaussian random effects priors centered on zero with inverse gamma distributions (with hyperparameters $a = b = 0.001$) used as hyperpriors for their variance. To obtain the posterior distribution, we draw on one million MCMC realizations which are thinned out at a rate of 800 after a burn-in of 200,000 MCMC realizations. For the posterior distributions we thus obtain 1,000 MCMC realizations for each parameter.

Before we go on to discuss our main findings concerning the impact of income on the two health variables considered, we first portray the effects of all covariates on the predictors of the parameters required to yield the distribution. While some of the parameters are interpretable in their own right (for example μ for the gamma distribution), we focus on evaluating the resultant distribution rather than the single parameters' estimates.

		males		
	$\eta^{\tilde{\pi}_1}$	$\eta^{\tilde{\pi}_2}$	$\eta^{\tilde{\pi}_3}$	$\eta^{\tilde{\pi}_4}$
const.	1.499[1.415; 1.592]	2.842[2.786; 2.902]	2.896[2.832; 2.969]	-0.899[-1.036;-0.742]
AGE	0.084[0.083; 0.085]	0.038[0.037; 0.038]	0.025[0.025; 0.026]	0.068[0.066; 0.070]
AGE^2	-0.001[-0.001;-0.001]	0.000[0.000; 0.000]	0.000[0.000; 0.000]	-0.001[-0.001; 0.000]
LOGINC	-0.182[-0.187;-0.178]	-0.458[-0.461;-0.455]	-0.453[-0.458;-0.450]	-0.298[-0.307;-0.290]
GER	-0.207[-0.211;-0.203]	-0.003[-0.006;-0.001]	0.266[0.262; 0.269]	-0.124[-0.132;-0.117]
EDU$_2$	0.024[0.020; 0.028]	0.069[0.067; 0.071]	-0.011[-0.014;-0.008]	0.148[0.142; 0.154]
EDU$_3$	0.064[0.059; 0.070]	-0.133[-0.136;-0.129]	-0.063[-0.068;-0.058]	-0.172[-0.184;-0.161]
EDU$_4$	-0.439[-0.443;-0.434]	-0.246[-0.248;-0.242]	-0.153[-0.157;-0.148]	-0.117[-0.127;-0.107]
MAR$_2$	-0.103[-0.110;-0.097]	0.215[0.211; 0.218]	-0.030[-0.035;-0.026]	0.153[0.145; 0.161]
MAR$_3$	-0.140[-0.146;-0.134]	0.084[0.081; 0.088]	0.326[0.321; 0.331]	0.056[0.046; 0.066]
MAR$_4$	0.158[0.149; 0.167]	-0.201[-0.205;-0.196]	-0.093[-0.097;-0.088]	-0.152[-0.161;-0.143]
EAST	0.247[0.072; 0.423]	-0.002[-0.099; 0.095]	-0.031[-0.100; 0.036]	0.037[-0.139; 0.207]
		females		
	$\eta^{\tilde{\pi}_1}$	$\eta^{\tilde{\pi}_2}$	$\eta^{\tilde{\pi}_3}$	$\eta^{\tilde{\pi}_4}$
const.	-1.787[-1.863;-1.692]	1.266[1.209; 1.330]	4.003[3.918; 4.089]	-1.052[-1.216;-0.890]
AGE	0.142[0.141; 0.143]	0.092[0.091; 0.093]	0.021[0.020; 0.022]	0.034[0.032; 0.035]
AGE^2	-0.001[-0.001;-0.001]	0.000[0.000; 0.000]	0.000[0.000; 0.000]	0.000[0.000; 0.000]
LOGINC	0.004[0.000; 0.008]	-0.461[-0.464;-0.458]	-0.593[-0.598;-0.589]	-0.160[-0.169;-0.152]
GER	-0.106[-0.110;-0.101]	-0.267[-0.270;-0.264]	0.009[0.004; 0.013]	-0.071[-0.080;-0.063]
EDU$_2$	0.082[0.079; 0.086]	0.040[0.038; 0.043]	0.091[0.088; 0.095]	0.312[0.305; 0.319]
EDU$_3$	-0.064[-0.070;-0.059]	0.027[0.024; 0.031]	-0.071[-0.076;-0.066]	-0.343[-0.354;-0.332]
EDU$_4$	-0.424[-0.428;-0.419]	-0.334[-0.337;-0.331]	-0.007[-0.012;-0.003]	-0.159[-0.169;-0.149]
MAR$_2$	-0.084[-0.093;-0.074]	0.167[0.163; 0.172]	0.006[0.001; 0.011]	0.248[0.238; 0.258]
MAR$_3$	-0.394[-0.403;-0.386]	0.209[0.205; 0.213]	0.051[0.046; 0.057]	-0.071[-0.082;-0.060]
MAR$_4$	0.442[0.424; 0.462]	-0.230[-0.236;-0.223]	0.016[0.009; 0.024]	-0.057[-0.071;-0.044]
EAST	0.092[-0.077; 0.264]	-0.024[-0.122; 0.072]	-0.009[-0.120; 0.100]	-0.089[-0.333; 0.148]

Table 1: Linear effects on $\eta^{\tilde{\pi}_1}$,$\eta^{\tilde{\pi}_2}$,$\eta^{\tilde{\pi}_3}$ and $\eta^{\tilde{\pi}_4}$ for PCS.

Table 1 displays the estimates for the covariate effects on the predictors of the sequential logits for the self-rated health outcomes. Here, we display the medians of the posterior distributions with the 95% (symmetric) credible intervals denoted in the brackets. In order to conserve space, we do not display the estimates for the the random effect estimates for the individual federal states but show them separately in Table 5 in the appendix.

While the parameter $\tilde{\pi}_l$ can be interpreted individually, we will not analyze these effects in detail.

Here, we restrict ourselves to noting that the effects of various variables differ significantly across the range of parameters estimated, both for males and females. Regarding LOGINC in particular, the effects are significantly different at the 5% level for different parameters.

	males		**females**	
	η^{μ}	η^{s}	η^{μ}	η^{s}
const.	1.637[1.562; 1.713]	3.382[2.709; 4.029]	1.777[1.700; 1.852]	3.186[2.554; 3.760]
AGE	0.007[0.006; 0.008]	-0.045[-0.057;-0.034]	0.005[0.004; 0.007]	-0.031[-0.043;-0.021]
AGE^2	0.000[0.000; 0.000]	0.000[0.000; 0.000]	0.000[0.000; 0.000]	0.000[0.000; 0.000]
LOGINC	-0.034[-0.041;-0.027]	0.152[0.091; 0.211]	-0.041[-0.048;-0.034]	0.123[0.066; 0.183]
GER	-0.010[-0.016;-0.004]	0.045[-0.017; 0.103]	0.007[0.001; 0.012]	0.042[-0.011; 0.099]
EDU_2	0.014[0.009; 0.019]	-0.077[-0.125;-0.032]	0.008[0.002; 0.013]	-0.079[-0.123;-0.036]
EDU_3	-0.006[-0.013; 0.001]	0.084[0.016; 0.147]	-0.013[-0.021;-0.005]	0.038[-0.029; 0.101]
EDU_4	-0.038[-0.044;-0.031]	0.062[0.000; 0.126]	-0.026[-0.033;-0.019]	0.017[-0.045; 0.080]
MAR_2	0.008[-0.002; 0.019]	0.023[-0.063; 0.105]	0.002[-0.006; 0.011]	0.000[-0.071; 0.063]
MAR_3	0.006[-0.004; 0.015]	-0.028[-0.107; 0.056]	-0.002[-0.011; 0.006]	-0.059[-0.128; 0.014]
MAR_4	-0.019[-0.036;-0.003]	-0.115[-0.246; 0.007]	-0.003[-0.013; 0.007]	-0.001[-0.081; 0.075]
EAST	0.010[-0.001; 0.021]	0.008[-0.074; 0.094]	0.009[-0.002; 0.019]	0.004[-0.044; 0.049]

Table 2: Linear effects on η^{μ} and η^{s} for PCS.

Table 2 shows the estimates for the predictors η^{μ} and η^{s} analogously to the table above. Again it may be noted that the effects are significantly different for males and females and that both for μ and for s, various covariates are significantly different from zero. For μ, which yields the conditional expectation, it should be noted that due to the linear transformation the effects are reversed, so that for example LOGINC has a negative impact on the predictor but thus a positive impact on the expected health, as one would expect. Concerning s, note that although a direct interpretation of the parameter is not feasible, one can observe that LOGINC as well as other variables have a significant impact which indicates complex changes across the covariate space that go beyond the changes in the conditional mean on which standard regression techniques focus.

4.4 Considering the Distributional Changes

Since we employ non-linear link functions for our predictors, the impact of the variables varies across the covariate space. This is well known from the literature on generalized linear models (Nelder and Wedderburn, 1972). We thus employ effect displays as proposed by Fox (1987). This means that we consider the effect of varying income while the other covariates are fixed at a given value. Here we consider the effects for both males and females who can be considered the "average Joe/average Jane", i.e. who are 52 years of age, are married, live in North-Rhine Westphalia (the most populous state in Germany), have standard secondary education and have German nationality.[5]

[5]See Section A.1 in the appendix for the covariates' distributions underlying this choice. For the continuous variable age we consider the arithmetic mean in our sample, while for the other categorical variables we consider

In assessing the health differentials associated with different income levels, we focus on relative rather than absolute differences.[6] This choice is based on the recommendation to use relative inequality measures when only concerned mainly about assessing health inequality rather than the absolute level of a health risk (see Harper et al., 2010). While the absolute levels of any health measure are clearly of importance in any inter-temporal or international comparison, the comparison we pursue here is between different metrics. A cross-metric comparison can only be based on a relative assessment, as their absolute measures cannot be compared for an assessment of inequality.

Figure 2 makes visible how the distribution of self-rated health changes with income, displaying the change in the probability of falling into one of five self-rated health states as one moves from the bottom to the top of the income distribution. These estimates are derived from the median results displayed in Table 1. We consider the income range from 5,000€ to 100,000€. The former constitutes the lower bound as only 1% of our estimates fall below this sample due to social security levels in Germany; the latter is chosen as the upper bound because it constitutes roughly the threshold to the most well-off 1% of the population. This income range thus encompasses the whole population bar the bottom and the top percent of the income distribution.
From the visualization alone, one can observe that the nature of the change in the health distribution across the income distribution is far from equiproportional. For example, while the share of respondents in "very good" health is nearly constant across the income range, the probability of being in "good" health increases 147% (from 0.20 to 0.49) from the bottom to the top of the income range. About 42% of respondents at both ends of the income range report being in "fair" health, but far fewer wealthier respondents are located at the bottom end of the health distribution: the share of people in "poor" health declines 81% (from 0.26 to 0.04) as one moves from the bottom to the top of the income distribution. For "very poor" health, the decrease is even larger at 88% (from 0.09 to 0.01). This shows that dichotomizing the outcome, e.g. by subsuming the levels 1-2 (not healthy) and 3-5 (healthy), may hide important relative variation within the aggregated categories.

In Figure 3, we focus on the difference in the conditional distributions of health status for men and women with a net equivalent income of 15,000€ (roughly corresponds to the 25^{th} percentile, i.e. the median for the poorer half of the population) versus 30,000€ (roughly corresponds to the

the mode. See Section A.5 in the appendix for other covariate combinations. Note also that it would be possible to consider average marginal effects rather than the marginal effects at the representative values. For the purposes of our paper, the marginal effects at the representative values were deemed more intuitive and are thus considered in the the following.

[6]The distinction between relative and absolute inequality has been discussed extensively in the health inequalities literature (see Mechanic, 2002; Oliver et al., 2002; Harper et al., 2010) and it has been noted that choosing relative over absolute measures of health inequality constitutes "an inherently value-laden enterprise, and judgments about justness, fairness, and social acceptability are inextricably bound to the selection of measures and statistical strategies" (Harper et al., 2010, p.6).

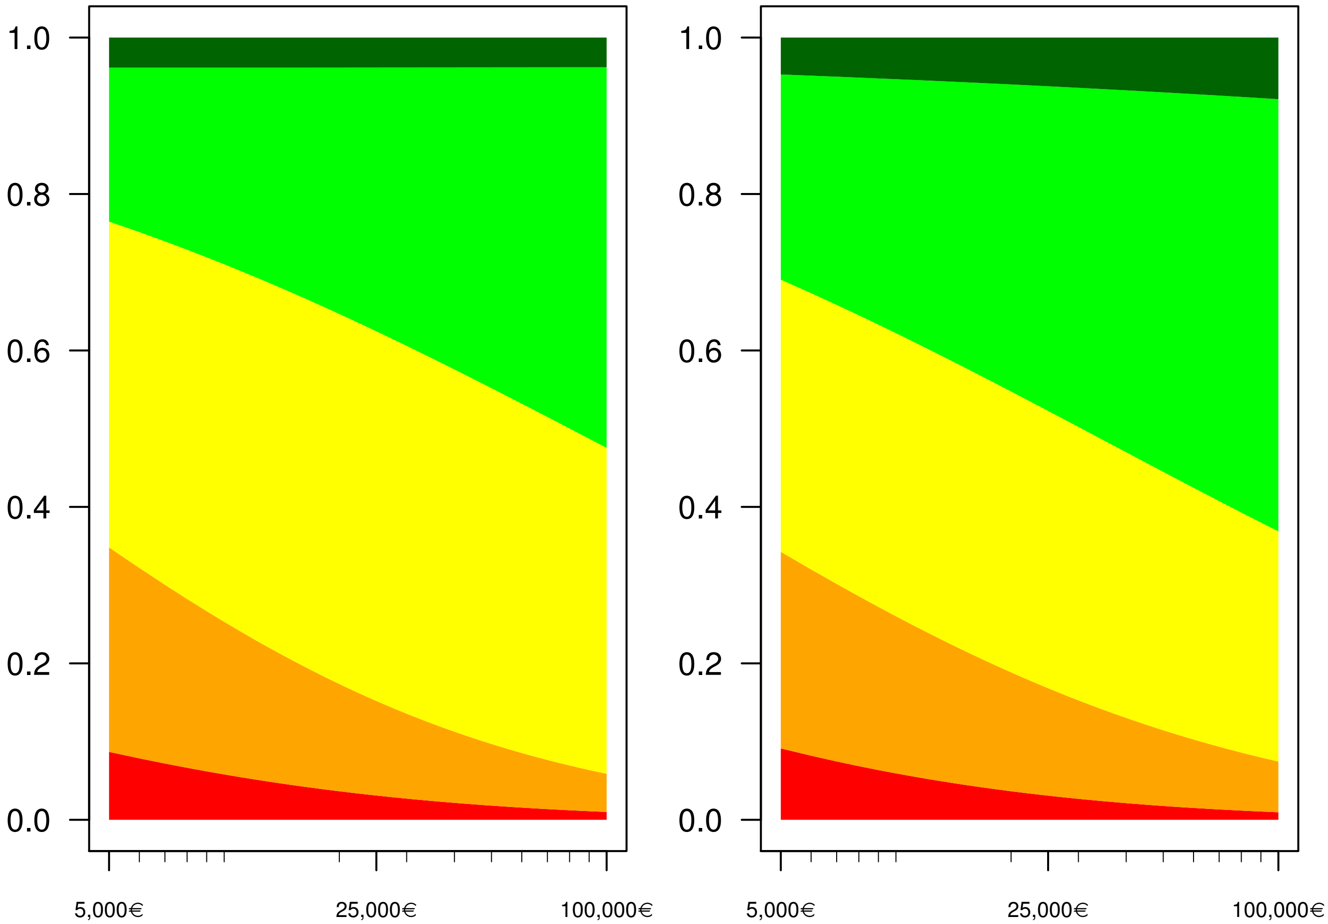

Figure 2: Conditional income effects for self-rated health for average Joe (left) and average Jane (right). From red (bad health) to dark green (very good health).

75$^{\text{th}}$ percentile, i.e. the median for the richer half of the population), with the other covariates fixed at the values to yield "average Joe" and "average Jane". The largest absolute differences occur near the center of the health distribution, i.e. for poor, fair, and good health. Despite the lower absolute levels, there are also noticeable changes at the bottom end of the health scale when moving from the lower to higher income level. Meanwhile, there is little change at the higher end of the distribution. This indicates that (more) money cannot buy (more) good health; but income does seem to contribute significantly to safeguarding against bad health outcomes – especially very bad ones, as we will see.

Let us contemplate the risk of falling in one of the lowest response categories for health across the two distributions (for income of 15,000€versus 30,000€). We can define the following three health

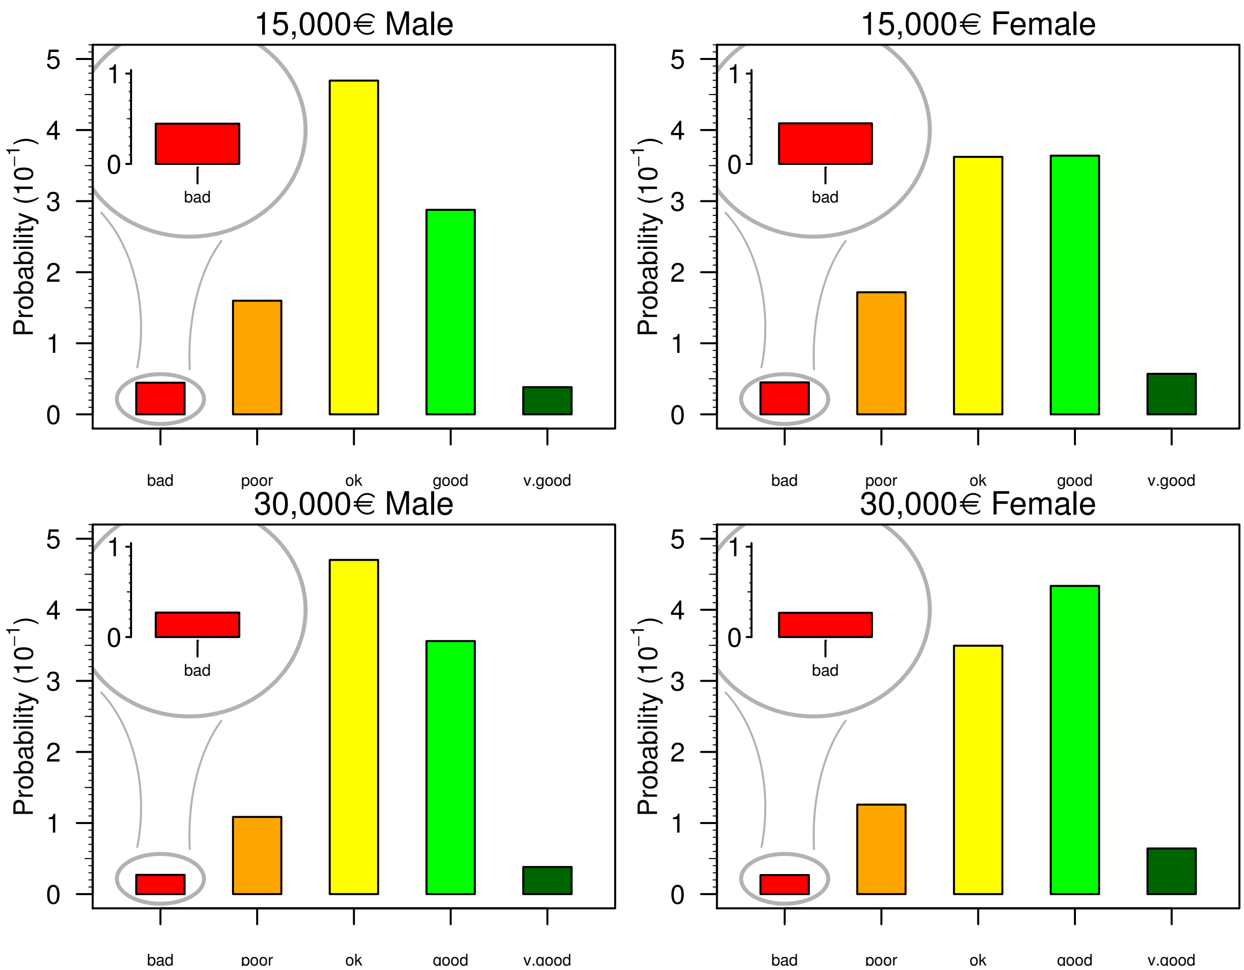

Figure 3: Conditional Health Distributions (SRH) for 15,000€ (top) and 30,000€ (bottom) for average Joe (left) and average Jane (right). With added focus on $\mathcal{R}_{M1}$ by magnification.

measures, which dichotomise the distribution in three different ways:

$$\begin{aligned}
\mathcal{R}_{M1} = &\ P(H^{\mathrm{M}} \leq \text{bad health}) \quad \text{with } H^{\mathrm{M}} \sim \mathcal{D}^{\mathrm{M}}_{\mathbf{x}}, \\
\mathcal{R}_{M2} = &\ P(H^{\mathrm{M}} \leq \text{poor health}) \quad \text{with } H^{\mathrm{M}} \sim \mathcal{D}^{\mathrm{M}}_{\mathbf{x}}, \\
\mathcal{R}_{M3} = &\ P(H^{\mathrm{M}} \leq \text{ok health}) \quad \text{with } H^{\mathrm{M}} \sim \mathcal{D}^{\mathrm{M}}_{\mathbf{x}},
\end{aligned}$$

where the health measures $\mathcal{R}_{M1}, \mathcal{R}_{M2}, \mathcal{R}_{M3}$ simply denote the risk of falling in one of the lowest response categories as given by the multinomial health distribution $\mathcal{D}^{\mathrm{M}}_{\mathbf{x}}$ which is dependent on the covariate combination under consideration, $\mathbf{x}$.

$\mathcal{R}_{M3}$ subsumes all health statuses below good into one category, thus representing the risk of "not feeling good about one's health". The probability of falling into one of these three lowest categories changes from 0.67 to 0.61 among men when moving from the conditional distribution for 15,000€ to that for 30,000€ – a change of 10%. For women, the probability falls from 0.58

to 0.50, a change of 13%. Although these differences are statistically significant, the magnitude is not substantively grave.
Secondly, we consider $\mathcal{R}_{M2}$, which by construction directs the attention towards those who are in poor or bad health (the bottom two health categories). This measure can therefore be seen as the risk of not only "not feeling good" but as "not even feeling ok". The change is of similar magnitude in absolute numbers, but much greater in relative terms. When income is doubled for men, the risk of low health status decreases by 34%, from 0.20 to 0.14, for men, while for women it falls 30%, from 0.22 to 0.15. The relative income-related change in risk of low health status is thus roughly 2-3 times as great when we aggregate the bottom two health categories as when we consider the bottom three categories together.
The third measure, $\mathcal{R}_{M1}$, is the most extreme measure which focusses on those who self-report a truly bad health. Thus it expresses the risk of positively "feeling bad about one's health". For this measure, the relative numbers are even more striking, with the probability of low health status decreasing by 39% and 40% for men and women respectively (from .04 to .03) as income doubles. The comparison of the three measures thus shows that the impact of household income on health seems to be much more drastic at the lower end of the self-rated health variable. Not surprisingly, this is also true when we consider the quasi-continuous PCS health score.

To characterize the relationship between income and the risk of low health using the SF12, we display six distributional measures in Figure 4. The blue line denotes males and the red line females, with the dashed lines denoting the 95% pointwise credible intervals.

The left-hand panels in 4 show the expectation (μ), the standard deviation (σ) and the skewness (γ_1) of the conditional distribution of the SF-12 across the full range of income. Note that we display these measures for the untransformed, original PCS variable, so that the effects are directly interpretable. The right-hand panels depict three measures of the risk of low health analogous to the ones used above. We portray the conditional probability that a person will fall below threshold values on the PCS scale representing the lower half (i.e. in the lowest 50%, denoted $T_{0.50}$), the lowest quintile (i.e. the lowest 20%, denoted $T_{0.20}$) and the lowest vingtile (i.e. the lowest 5%, denoted $T_{0.05}$) of the aggregate health distribution, depending on their income.[7] These measures can thus be seen as analogous variants of the risk measures $\mathcal{R}_{M1}$, $\mathcal{R}_{M2}$ and $\mathcal{R}_{M3}$ from above, indicating the risk of bad health. The measure $\mathcal{R}_{C0.50}$ thus yields the level of risk of belonging to the lower half of the health distribution, which can be seen as roughly equivalent to "not feeling good about one's health". Accordingly, $\mathcal{R}_{C0.20}$ yields the level of risk of belonging to the "sickest" 20% of the population, which can be seen as roughly equivalent to people associating the health status as slightly sick, that is no longer "o.k.". Lastly, $\mathcal{R}_{C0.05}$ denotes the risk of falling into the

[7] These values are obviously not the only viable options but chosen on the grounds as to provide roughly analogous risk measures to the risk measures based on the self-rated health responses. More research is needed concerning the use of adequate scalar measures to assess this and other aspects of conditional health distributions.

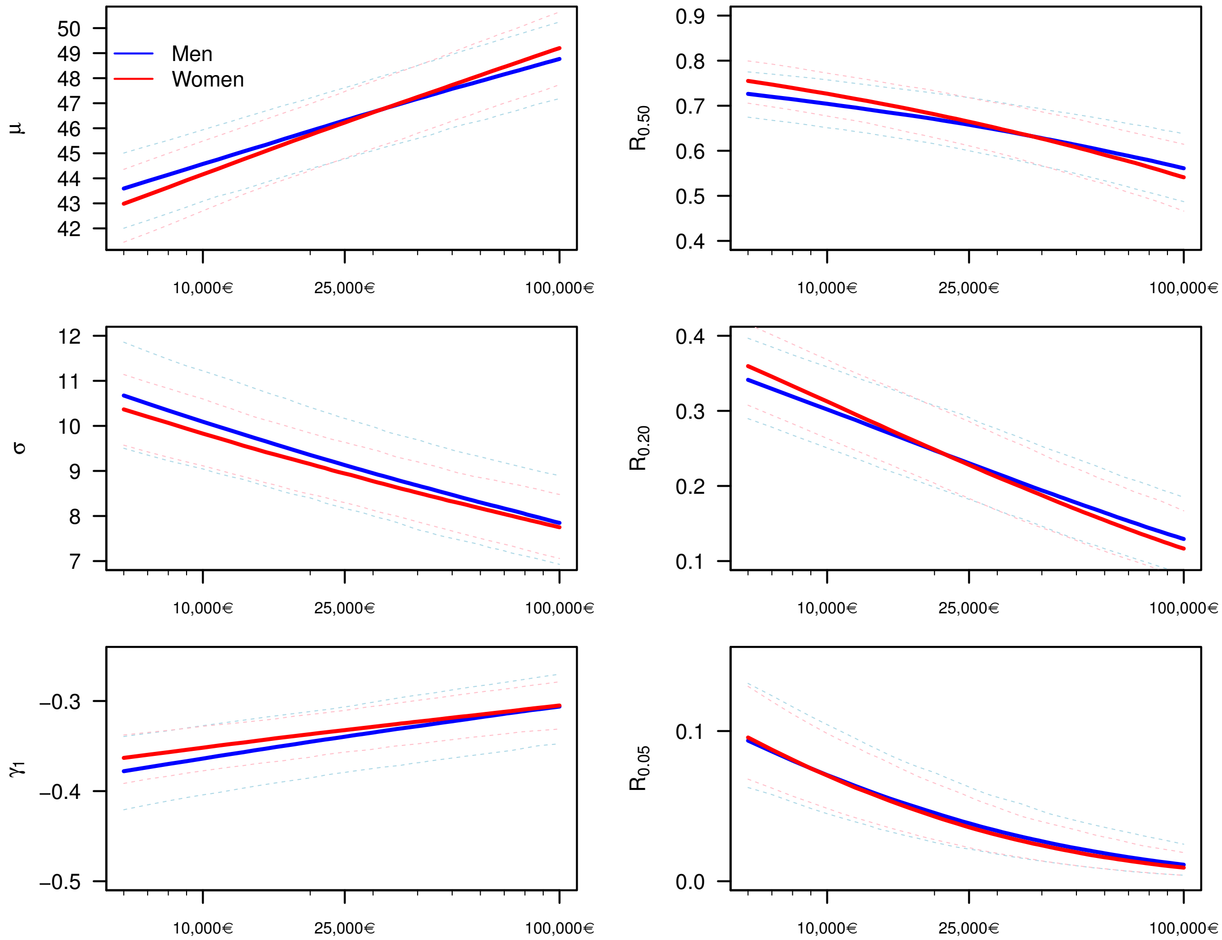

Figure 4: Left: Effect of income on mean, standard deviation and skewness of PCS. Right: Effect of income on risk of falling below lowest quintile, decile and vingtile of PCS.

lowest 5% of the health distribution, which would be associated with severe sickness and thus can roughly be seen as the equivalent to a person positively "feeling bad about one's health". More formally, the second set of risk measures can be defined as

$$\begin{aligned}
\mathcal{R}_{C0.05} &= P(H^{\mathrm{C}} \leq T_{0.05}) \quad \text{with } H^{\mathrm{C}} \sim \mathcal{D}^{\mathrm{C}}_{\mathbf{x}}, \\
\mathcal{R}_{C0.20} &= P(H^{\mathrm{C}} \leq T_{0.20}) \quad \text{with } H^{\mathrm{C}} \sim \mathcal{D}^{\mathrm{C}}_{\mathbf{x}}, \\
\mathcal{R}_{C0.50} &= P(H^{\mathrm{C}} \leq T_{0.50}) \quad \text{with } H^{\mathrm{C}} \sim \mathcal{D}^{\mathrm{C}}_{\mathbf{x}},
\end{aligned}$$

where the health variable H^{C} is now considered as continuous. The risk is thus given by the conditional distribution, $\mathcal{D}^{\mathrm{C}}_{\mathbf{x}}$, which the variable is thought to follow for an individual with characteristics $\mathbf{x}$, and the threshold value T_{α}, which we take to be a quintile from the aggregate distribution of H^{C}.

The estimated full conditional distributions for the SF-12 for "average Joes" and "average Janes"

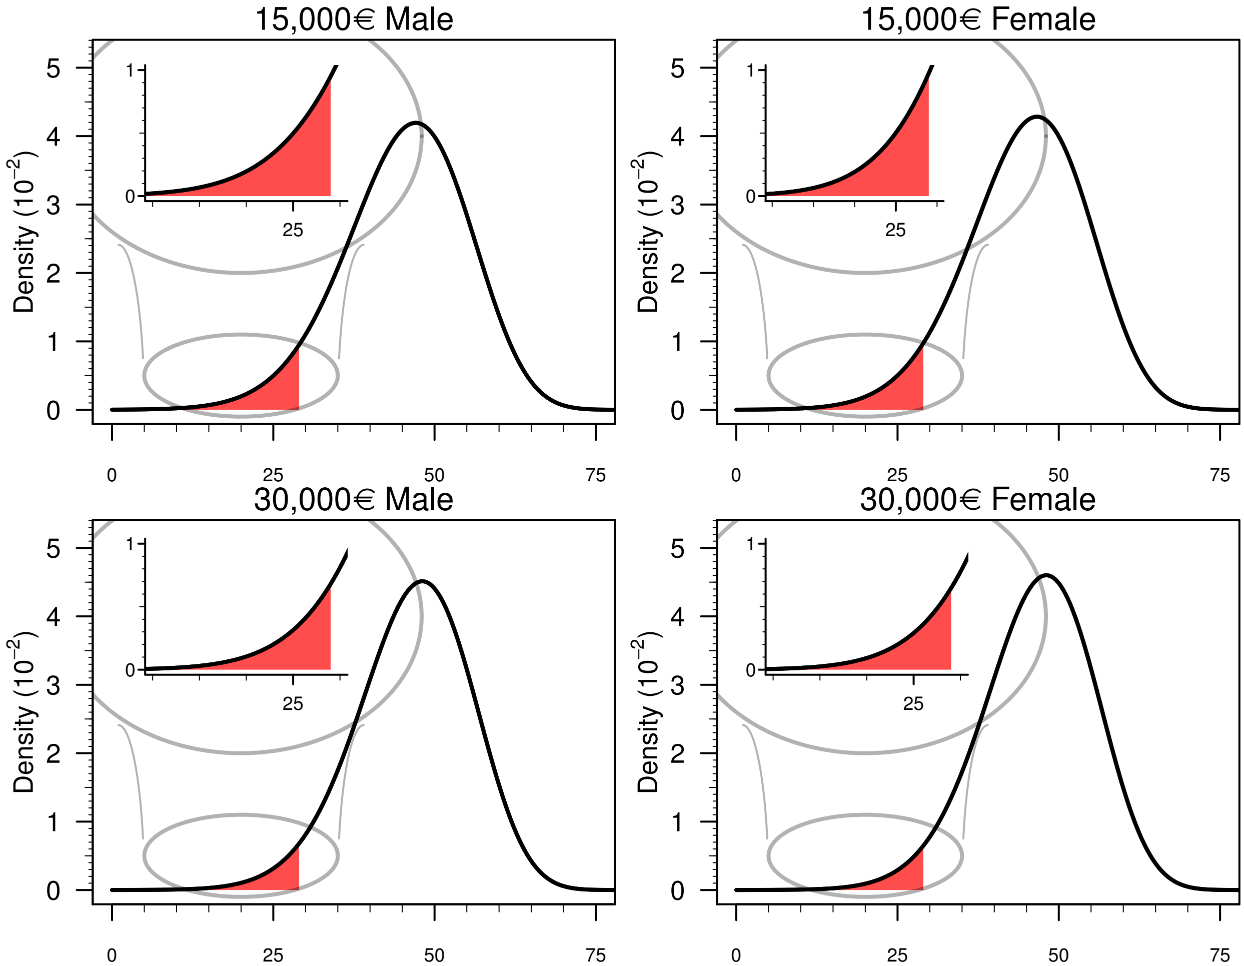

Figure 5: Conditional health distributions (PCS) for average Joe (left) and average Jane (right) for 15,000€ (top) and 30,000€ (bottom). With added focus on $\mathcal{R}_{C0.05}$ by magnification.

are displayed in Figure 5. Again we focus on the contrast between 15,000€, representing the median income level of the poorer half of the sample population, and 30,000€, representing the median income level of the richer half of the sample population. While the displayed distributions appear rather similar at the first glance, a closer look at the different distributions' attributes reveals some substantial differences. For an annual net equivalent income of 15,000€ the average physical health value is 45.3 and 45.1 for men and women respectively. In contrast, for an income of 30,000€ , we obtain 46.7 and 46.6. Thus the average male described above with a net equivalent income of 30,000€ roughly has a 3% higher expected physical health score as an otherwise equivalent male with a net equivalent income of 15,000€. For a female the difference is also roughly 3%. This effect is well known and discussed extensively in the literature.

Next to the mean, the standard deviation also decreases from 9.5 to 8.8 and 9.3 to 8.7 for men and women respectively. This 7% decrease means that men and women with higher income face a lower risk to experience very low health outcomes for a given mean. Additionally, the distribution becomes slightly more right skewed, with the skewness increasing from -0.4 to -0.3 for both men and women. This constitutes a 4% and 5% increase respectively. This change in skewness also increases the probability of an individual finding himself on the lower outskirts of the health distribution. These results thus indicate that the nature of the association of income with health beyond the mean, with the risk of very low health scores - indicating severe sickness - driven not only by a deteriorating mean but also by a higher standard deviation and a less left-skewed distribution.

As indicated by the higher order moments, the increase in the health risks are higher when directing the focus further towards the lower end of the health spectrum. Considering $\mathcal{R}_{C0.50}$ for males, we still find a moderate change in the risk from 0.70 to 0.65, constituting a decrease of 6%. For women, we see a decrease from 0.70 to 0.65, i.e. by 7%. This change can be seen as of a similar magnitude as R_{M3} and also similar to the relative change observed for the expected outcome (see above). The relative difference increases to 20% (0.27 to 0.22) and 23% (0.27 to 0.21) for men and women respectively, when considering $\mathcal{R}_{C0.20}$. The greatest relative effect is seen for $\mathcal{R}_{C0.05}$, which sees the risk of falling into the lowest health quintile of the population at 0.06 for "average Joe" and 0.05 for "average Jane" at 15,000€, whereas having an income twice as high reduces that risk down to 0.03 for both, a decrease of 39% and 42% respectively. In other words, the risk of extremely bad health can be roughly halved by doubling the net equivalent income from 15,000€to 30,000€. Obviously, the magnitude of this effects is structurally different from the observed 3% increase observed for expected health.

4.5 Implications for Health Assessment

Considering the whole conditional health distribution and changes thereto over the covariate space thus yields potentially starkly different magnitudes for the assessment of the association between income and health. The relative differences are summarized in Table 5. The relative difference is the absolute difference divided by the measure for 15,000€.

The association between income and health becomes significantly greater if we focus on the lower end of the health spectrum. The mean relative difference is around 3%, while at the lower end of the health spectrum $\mathcal{R}_{M1}$ and $\mathcal{R}_{C0.05}$ relative differences in health by income are on the order of 39%-42% – i.e. more than tenfold greater.

The conventional perspective generates significant results that allow us to infer the existence of a relationship between income and health. How does our more complicated statistical artillery help us go beyond the results more easily generated using well-established mean-based analyses? The answer lies in the fact that while average population health is an important construct for many purposes, we cannot properly calculate the utility of alternative distributions of health using only

	15,000€	30,000€	Relative Difference
males			
$\mathcal{R}_{M3}$	0.67[0.67; 0.68]	0.61[0.60; 0.61]	10.11%[10.01%;10.20%]
$\mathcal{R}_{M2}$	0.20[0.20; 0.21]	0.14[0.13; 0.14]	33.64%[33.41%;33.87%]
$\mathcal{R}_{M1}$	0.04[0.04; 0.05]	0.03[0.03; 0.03]	39.21%[38.77%;39.71%]
μ	45.33[43.84;46.69]	46.65[45.11;47.97]	2.90%[2.27%; 3.40%]
$\mathcal{R}_{C0.50}$	0.68[0.63; 0.74]	0.65[0.59; 0.71]	5.53%[3.75%; 7.83%]
$\mathcal{R}_{C0.20}$	0.27[0.22; 0.33]	0.22[0.17; 0.28]	19.92%[15.14%;25.91%]
$\mathcal{R}_{C0.05}$	0.06[0.03; 0.08]	0.03[0.02; 0.06]	38.98%[30.05%;48.95%]
females			
	15,000€	30,000€	Relative Difference
$\mathcal{R}_{M3}$	0.58[0.58; 0.58]	0.50[0.50; 0.50]	13.27%[13.15%;13.39%]
$\mathcal{R}_{M2}$	0.22[0.21; 0.22]	0.15[0.15; 0.15]	29.54%[29.32%;29.80%]
$\mathcal{R}_{M1}$	0.04[0.04; 0.05]	0.03[0.03; 0.03]	40.38%[39.96%;40.89%]
μ	45.09[43.69;46.36]	46.63[45.19;47.87]	3.43%[2.82%: 4.01%]
$\mathcal{R}_{C0.50}$	0.70[0.65; 0.75]	0.65[0.60; 0.71]	7.24%[5.14%; 9.48%]
$\mathcal{R}_{C0.20}$	0.27[0.23; 0.33]	0.21[0.17; 0.27]	22.83%[17.88%;28.13%]
$\mathcal{R}_{C0.05}$	0.05[0.04; 0.08]	0.03[0.02; 0.05]	42.36%[33.98%;50.58%]

Table 3: Seven measures on the health-income association.

this summary statistic because the utility function for health is generally thought to be concave. If the utility gain from increases in health status at the low end of the health spectrum is greater than at the high end, changes to the distribution of income that do not affect the mean health of the population but lessen the number of people in very poor health would nevertheless be preferable at a societal level. At a policy level, too, there are good reasons to care at least as much about the risk of people being in poor health as about average health achievement in the population, since the primary purpose of public or private health insurance is to cover the cost of caring for those who are ill, rather than focussing on improving the health of the already healthy even further. When we think about the relationship between health and income, then, we want to be able to pay attention not only to the average effect of income on health, but also to the where in the health distribution people of various incomes are more likely to fall. That is what SADR allows us to do. The results we have shown here demonstrate that while the income-health relation may not be of great magnitude if we focus only on average health, the income-illness relationship (i.e., concentrating on the ill) is considerable larger.

5 Conclusion

In this paper, we follow other recent publications that have pointed to the shortcomings of regression based assessments focussing solely on the expected outcome. In order to look beyond the mean, we propose the use of structured additive distributional regression (SADR). These models allow for the estimation of full conditional health distributions for both multicategorical and con-

tinuous measures of health outcomes.
Using health data from the German Socio-Economic Panel, we apply SADR and find that the standard expectation-based perspective may neglect potentially important aspects of the relationship between health and income. In particular, we show that the risk of being in very poor health is much more strongly related to income than the average health status. We find that the risk for the "average Joe" and "average Jane" of belonging to the severely sick population increases between 39% and 42% when the net equivalent household income is changed from the median income of the poorer half of the population (15,000€) to the median income of the richer half (30,000€) in Germany. This exceeds the income-related change in average health status that is estimated using standard estimation techniques by more than tenfold. This suggests that mean-based perspectives may underestimate the effect of changes in the income distribution on well-being (given a concave health-utility relationship) and/or on health care expenditures (given that health care is more cost intensive at the lower end of the health distribution).

Based on the findings of this paper, we propose that future estimates of the health-income relationship take into account not only the mean reported health (or the probability of a dichotomized health measure in an income group), but also employ risk measures focusing on very poor health outcomes like the the ones used in this paper. Not only would this put more emphasis on the the lower end of the spectrum, where we argue it is merited. In addition, it would address problems associated with non-linearities with respect to well-being and/or health care and mean regression (see above). A distributional approach and risk-based measures such as the ones we propose may also unify the interpretation of the otherwise starkly different results that can arise depending on whether discrete data (and odds-ratios) or continuous data (and arithmetic means) are used for the assessment. We find that using SADR, the estimated magnitude of the income-health relationship is very similar for the single-item self-rated health measure and the SF-12. The distributional approach thus may contribute to the convergence of findings from the epidemiological literature (which mainly employs discrete measures like self-rated health and odds ratios) and the health economics literature (which tends to employ continuous measures like the SF-12 and arithmetic means).

Several extensions to the present approach might be considered. One particularly interesting modification would be to model the full joint distribution of health and income with respect to other covariates such as age and education. This would be feasible applying bivariate structured additive distributional regression, which uses copula structures to model the interrelations of the dependent variables (see Klein and Kneib, 2016) and would allow for the construction of conditional concentration curves across the covariate space. While technically challenging, this approach would not only incorporate the workhorse method in the health economics literature into the proposed framework, but would also allow researchers to consider distributional aspects beyond the mean without the need to define threshold values. Such advancements are needed because, to paraphrase Thomas Piketty (2014), failing to deal with the distributional nature of the health-income relationship rarely serves the interests of the least well-off.

References

H. H. Andersen, A. Mühlbacher, M. Nübling, J. Schupp and G. G. Wagner (2007): Computation of Standard Values for Physical and Mental Health Scale Scores: Using the SOEP Version of SF-12v2, in: Schmollers Jahrbuch, 127, pp. 171–182.

E. Backlund, P. D. Sorlie and N. J. Johnson (1996): The Shape of the Relationship between Income and Mortality in the United States, in: Annals of Epidemiology, 6(1), pp. 12–20.

C. Belitz, A. Brezger, N. Klein, T. Kneib, S. Lang and N. Umlauf (2015): BayesX - Software for Bayesian Inference in Structured Additive Regression Models: Version 3.0, URL `http://www.bayesx.org`.

Y. Benyamini (2011): Why does Self-rated Health Predict Mortality? An Update on Current Knowledge and a Research Agenda for psychologists, in: Psychology & Health, 26(11), pp. 1407–1413.

M. Benzeval, J. Taylor and K. Judge (2000): Evidence on the Relationship between Low Income and Poor Health: Is the Government Doing Enough?, in: Fiscal Studies, 21(3), pp. 375–399.

M. L. Berk and A. C. Monheit (2001): The Concentration Of Health Care Expenditures, Revisited, in: Health Affairs, 20(2), pp. 9–18.

E. E. S. Bjornstrom and D. C. Kuhl (2014): A Different Look at the Epidemiological Paradox: Self-Related Health, Perceived Social Cohesion, and Neighborhood Immigrant Context, in: Social Science & Medicine, 120, pp. 118–125.

R. W. Bohannon, R. Maljanian and M. Landes (2004): Test–retest Reliability of Short Form (SF)-12 Component Scores of Patients with Stroke, in: International Journal of Rehabilitation Research, 27(2), pp. 149–150.

H. D. Bondell, B. J. Reich and H. Wang (2010): Noncrossing Quantile Regression Curve Estimation, in: Biometrika, 97(4), pp. 825–838.

S. Bourion-Bédès, R. Schwan, V. Laprevote, A. Bédès, J.-L. Bonnet and C. Baumann (2015): Differential Item Functioning (DIF) of SF-12 and Q-LES-Q-SF Items among French Substance Users, in: Health and Quality of Life Outcomes, 13, p. 172.

J. E. Brazier and J. Roberts (2004): The Estimation of a Preference-based Measure of Health from the SF-12, in: Medical Care, 42(9), pp. 851–859.

J. N. Burdine, M. R. Felix, A. T. Abel, C. J. Wiltraut and Y. J. Musselman (2000): The SF-12 as a Population Health Measure: An Exploratory Examination of Potential for Application, in: Health Services Research, 35(4), pp. 885–904.

V. Carrieri and A. M. Jones (2016): The Income-Health Relationship 'Beyond the Mean': New Evidence from Biomarkers, in: Health Economics.

A. Case (2001): Does Money Protect Health Status? Evidence from South African Pensions, in: NBER Working Paper, No. 8495.

D. Cesarini, E. Lindqvist, R. Östling and B. Wallace (2016): Wealth, Health, and Child Development: Evidence from Administrative Data on Swedish Lottery Players, in: Quarterly Journal of Economics, 131(2), pp. 687–738.

G. Chamberlain (1980): Analysis of Covariance with Qualitative Data, in: The Review of Economic Studies, 47(1), p. 225.

V. Chernozhukov, I. Fernandez-Val and B. Melly (2013): Inference on Counterfactual Distributions, in: Econometrica, 81(6), pp. 2205–2268.

P. Contoyannis, A. M. Jones and N. Rice (2004): The Dynamics of Health in the British Household Panel Survey, in: Journal of Applied Econometrics, 19(4), pp. 473–503.

O. Cunillera, R. Tresserras, L. Rajmil, G. Vilagut, P. Brugulat, M. Herdman, A. Mompart, A. Medina, Y. Pardo, J. Alonso, J. Brazier and M. Ferrer (2010): Discriminative capacity of the EQ-5D, SF-6D, and SF-12 as measures of health status in population health survey, in: Quality of Life Research, 19(6), pp. 853–864.

K. B. DeSalvo, N. Bloser, K. Reynolds, J. He and P. Muntner (2006): Mortality prediction with a single general self-rated health question. A meta-analysis, in: Journal of general internal medicine, 21(3), pp. 267–275.

K. B. DeSalvo, V. S. Fan, M. B. McDonell and S. D. Fihn (2005): Predicting Mortality and Healthcare Utilization with a Single Question, in: Health Services Research, 40(4), pp. 1234–1246.

K. B. DeSalvo, T. M. Jones, J. Peabody, J. McDonald, S. Fihn, V. Fan, J. He and P. Muntner (2009): Health Care Expenditure Prediction with a Single Item, Self-rated Health Measure, in: Medical Care, 47(4), pp. 440–447.

J.-Y. Duclos and D. Échevin (2011): Health and Income: a Robust Comparison of Canada and the US, in: Journal of Health Economics, 30(2), pp. 293–302.

R. Ecob and G. Davey Smith (1999): Income and Health: What is the Nature of the Relationship?, in: Social Science & Medicine, 48(5), pp. 693–705.

P. Eibich and N. R. Ziebarth (2014): Examining the Structure of Spatial Health Effects in Germany Using Hierarchical Bayes Models, in: Regional Science and Urban Economics, 49, pp. 305–320.

G. Erreygers and T. van Ourti (2011): Measuring socioeconomic inequality in health, health care and health financing by means of rank-dependent indices: A recipe for good practice, in: Journal of Health Economics, 30(4), pp. 685–694.

S. L. Ettner (1996): New Evidence on the Relationship between Income and Health, in: Journal of Health Economics, 15(1), pp. 67–85.

L. Fahrmeir, T. Kneib and S. Lang (2004): Penalized Structured Additive Regression for Space-Time data: A Bayesian Perspective, in: Statistica Sinica, 14, pp. 731–761.

B. Fayissa and P. Gutema (2005): Estimating a Health Production Function for Sub-Saharan Africa (SSA), in: Applied Economics, 37(2), pp. 155–164.

A. Finkelstein, E. F. P. Luttmer and M. J. Notowidigdo (2009): Approaches to Estimating the Health State Dependence of the Utility Function, in: American Economic Review, 99(2), pp. 116–121.

S. Firpo, N. M. Fortin and T. Lemieux (2009): Unconditional Quantile Regressions, in: Econometrica, 77(3), pp. 953–973.

J. A. Fleishman, J. W. Cohen, W. G. Manning and M. Kosinski (2006): Using the SF-12 health status measure to improve predictions of medical expenditures, in: Medical Care, 44(5 Suppl), pp. I54–63.

J. A. Fleishman and W. F. Lawrence (2003): Demographic Variation in SF-12 scores: True Differences or Differential Item Functioning?, in: Medical Care, 41(7 Suppl), pp. III75–III86.

M. Fleurbaey and E. Schokkaert (2009): Unfair Inequalities in Health and Health Care, in: Journal of Health Economics, 28(1), pp. 73–90.

J. Fox (1987): Effect Displays for Generalized Linear Models, in: C. C. Clogg (ed.), Sociological Methodology, vol. 17, pp. 347–361, American Sociological Association, Washington D.C.

P. Franks, E. I. Lubetkin, M. R. Gold and D. J. Tancredi (2003): Mapping the SF-12 to preference-based instruments: convergent validity in a low-income, minority population, in: Medical Care, 41(11), pp. 1277–1283.

P. Frijters, J. P. Haisken-DeNew and M. A. Shields (2005): The Causal Effect of Income on Health: Evidence from German Reunification, in: Journal of Health Economics, 24(5), pp. 997–1017.

A. F. Galvao, C. Lamarche and L. R. Lima (2013): Estimation of Censored Quantile Regression for Panel Data with Fixed Effects, in: Journal of the American Statistical Association, 108(503), pp. 1075–1089.

B. Gandek, J. E. Ware, N. K. Aaronson, G. Apolone, J. B. Bjorner, J. E. Brazier, M. Bullinger, S. Kaasa, A. Leplege, L. Prieto and M. Sullivan (1998): Cross-validation of Item Selection and Scoring for the SF-12 Health Survey in Nine Countries: Results from the IQOLA Project, in: Journal of Clinical Epidemiology, 51(11), pp. 1171–1178.

H. Gravelle (1998): How Much of the Relation between Population Mortality and Unequal Distribution of Income is a Statistical Artefact?, in: BMJ, 316(7128), pp. 382–385.

— (2003): Measuring Income Related Inequality in Health: Standardisation and the Partial Concentration Index, in: Health Economics, 12(10), pp. 803–819.

A. M. Gray, O. Rivero-Arias and P. M. Clarke (2006): Estimating the association between SF-12 responses and EQ-5D utility values by response mapping, in: Medical Decision Making, 26(1), pp. 18–29.

S. Harper, N. B. King, S. C. Meersman, M. E. Reichman, N. Breen and J. Lynch (2010): Implicit Value Judgments in the Measurement of Health Inequalities, in: The Milbank quarterly, 88(1), pp. 4–29.

T. J. Hastie and R. J. Tibshirani (1990): Generalized Additive Models, Chapman & Hall, New York.

G. Heckley, U.-G. Gerdtham and G. Kjellsson (2016): A General Method for Decomposing the Causes of Socioeconomic Inequality in Health, in: Journal of Health Economics, 48, pp. 89–106.

T. Hothorn, T. Kneib and P. Bühlmann (2014): Conditional Transformation Models, in: Journal of the Royal Statistical Society: Series B (Statistical Methodology), 76(1), pp. 3–27.

K. H. Humphries and E. van Doorslaer (2000): Income-related Health Inequality in Canada, in: Social Science & Medicine, 50(5), pp. 663–671.

E. L. Idler and Y. Benyamini (1997): Self-Rated Health and Mortality: A Review of Twenty-Seven Community Studies, in: Journal of Health and Social Behavior, 38(1), pp. 21–37.

A. M. Jones and J. Wildman (2005): Disentangling the Relationship between Health and Income, in: ECUITY III, Working Paper No. 4.

N. Kakwani, A. Wagstaff and E. van Doorslaer (1997): Socioeconomic Inequalities in Health: Measurement, Computation, and Statistical Inference, in: Journal of Econometrics, 77(1), pp. 87–103.

I. Kawachi, N. E. Adler and W. H. Dow (2010): Money, Schooling, and Health: Mechanisms and Causal Evidence, in: Annals of the New York Academy of Sciences, 1186, pp. 56–68.

O. Kiuila and P. Mieszkowsko (2007): The Effects of Income, Education and Age on Health, in: Health Economics, 16(8), pp. 781–798.

N. Klein and T. Kneib (2016): Scale-Dependent Priors for Variance Parameters in Structured Additive Distributional Regression, in: Bayesian Analysis, 11, pp. 1071–1106.

N. Klein, T. Kneib, S. Lang and A. Sohn (2015): Bayesian Structured Additive Distributional Regression with an Application to Regional Income Inequality in Germany, in: Annals of Applied Statistics, 9(2), pp. 1024–1052.

R. Koenker and G. Bassett (1978): Regression Quantiles, in: Econometrica, 46(1), pp. 33–50.

H.-H. König, D. Heider, T. Lehnert, S. G. Riedel-Heller, M. C. Angermeyer, H. Matschinger, G. Vilagut, R. Bruffaerts, J. M. Haro, G. de Girolamo, R. de Graaf, V. Kovess and J. Alonso (2010): Health Status of the Advanced Elderly in six European Countries: Results from a Representative Survey using EQ-5D and SF-12, in: Health and Quality of Life Outcomes, 8, p. 143.

D. Kuehnle (2014): The Causal Effect of Family Income on Child Health in the U.K, in: Journal of Health Economics, 36, pp. 137–150.

A. D. LaMontagne, A. Milner, L. Krnjacki, A. M. Kavanagh, T. A. Blakely and R. Bentley (2014): Employment Arrangements and Mental Health in a Cohort of Working Australians: Are Transitions From Permanent to Temporary Employment Associated With Changes in Mental Health?, in: American Journal of Epidemiology, 179(12), pp. 1467–1476.

W. F. Lawrence and J. A. Fleishman (2004): Predicting EuroQoL EQ-5D preference scores from the SF-12 Health Survey in a nationally representative sample, in: Medical Decision Making, 24(2), pp. 160–169.

M. Lindahl (2005): Estimating the Effect of Income on Health and Mortality Using Lottery Prizes as an Exogenous Source of Variation in Income, in: Journal of Human Resources, 40(1), pp. 144–168.

P. K. Lorgelly and J. Lindley (2008): What is the Relationship Between Income Inequality and Health? Evidence from the BHPS, in: Health Economics, 17(2), pp. 249–265.

J. W. Lynch and G. A. Kaplan (1997): Understanding how inequality in the distribution of income affects health, in: Journal of health psychology, 2(3), pp. 297–314.

P. Makdissi and M. Yazbeck (2014): Measuring Socioeconomic Health Inequalities in Presence of Multiple Categorical Information, in: Journal of Health Economics, 34, pp. 84–95.

J. Marcus (2013): The Effect of Unemployment on the Mental Health of Spouses – Evidence from Plant Closures in Germany, in: Journal of Health Economics, 32(3), pp. 546–558.

M. Marmot (2002): The Influence of Income on Health: Views of an Epidemiologist, in: Health Affairs, 21(2), pp. 31–46.

P. McDonough, G. J. Duncan, D. Williams and J. House (1997): Income dynamics and adult mortality in the United States, 1972 through 1989, in: American Journal of Public Health, 87(9), pp. 1476–1483.

D. Mechanic (2002): Disadvantage, Inequality, And Social Policy, in: Health Affairs, 21(2), pp. 48–59.

J. A. Nelder and R. W. M. Wedderburn (1972): Generalized Linear Models, in: Journal of the Royal Statistical Society: Series A (General), 135(3), pp. 370–384.

A. Oliver, A. Healey and J. Le Grand (2002): Addressing Health Inequalities, in: The Lancet, 360(9332), pp. 565–567.

OPTUM (2015): SF Health Surveys, URL `https://www.optum.com/optum-outcomes/what-we-do/health-surveys.html`.

T. Piketty (2014): Capital in the Twenty-First Century: The Dynamics of Inequality, Wealth, and Growth, Harvard University Press, Cambridge.

U. Rahmann and J. Schupp (2013): SOEP Wave Report 2012, DIW Berlin, Berlin.

B. Ravesteijn, H. van Kippersluis and E. van Doorslaer (2013): The Contribution of Occupation to Health Inequality, in: P. R. Dias and O. O'Donnell (eds.), Health and Inequality, pp. 311–332, Emerald, Bingley.

E. Rogot, P. D. Sorlie and N. J. Johnson (1992): Life expectancy by employment status, income, and education in the National Longitudinal Mortality Study, in: Public Health Reports, 107(4), pp. 457–461.

V. Schiele and H. Schmitz (2016): Quantile treatment effects of job loss on health, in: Journal of Health Economics, 49, pp. 59–69.

J. Schnittker (2004): Education and the Changing Shape of the Income Gradient in Health, in: Journal of Health and Social Behavior, 45(3), pp. 286–305.

SOEP (2014): Socio-oeconomic Panel Study (SOEP), data of the years 1984-2013, version 30, doi: 10.5684/soep.v30.

J. Thornton (2010): Estimating a Health Production Function for the US: Some new evidence, in: Applied Economics, 34(1), pp. 59–62.

G. G. Wagner, J. R. Frick and J. Schupp (2007): The German Socio-Economic Panel Study (SOEP) - Scope Evolution and Enhancements, in: Schmollers Jahrbuch, 127(1), pp. 139–169.

A. Wagstaff (2005): The Bounds of the Concentration Index when the Variable of Interest is Binary, with an Application to Immunization Inequality, in: Health Economics, 14(4), pp. 429–432.

— (2011): The Concentration Index of a Binary Outcome Revisited, in: Health Economics, 20(10), pp. 1155–1160.

A. Wagstaff and E. van Doorslaer (1994): Measuring Inequalities in Health in the Presence of Multiple-category Morbidity Indicators, in: Health Economics, 3(4), pp. 281–291.

E. Waldmann, T. Kneib, Y. R. Yue, S. Lang and C. Flexeder (2013): Bayesian Semiparametric Additive Quantile Regression, in: Statistical Modelling, 13(3), pp. 223–252.

J. E. Ware Jr., M. Kosinski and S. D. Keller (1996): A 12-Item Short-Form Health Survey: Construction of Scales and Preliminary Tests of Reliability and Validity, in: Medical Care, 34(3), pp. 220–233.

A Appendix

A.1 Data

As primary source for our data, we use the SOEP database (Wagner et al., 2007). We use all available samples in 2012, i.e. samples A-L. Concerning the wave, we only consider the wave from 2012, i.e. wave BC, for all questions on current status. For questions asked with respect to the whole last year, we consider questions from 2013, referring back to 2012, i.e. questions from wave BD. Only taking those values for which we have the full set of variables, as described below, this yields 16,732 observations (7,820 males and 8,903 females).

As for the dependent variable we simply take the single item self-rate health response (`bcp91`) on the one hand. On the other hand, we consider the physical condition score from the SF-12 (`PCS`) which is directly available via the `HEALTH` file in the SOEP.

As a variable for income we use the household's net income as the base (`i1110213` from the `BCPEQUIV` file) and divide it by the equivalised household size, based on the OECD equivalence scale (using the variables `bchhgr` and `bckzahl` from the `BCKIND` file). Thereby the first adult is given a weight of 1, whereas to every additional person aged 14 and over is given the weight of 0.5. Each child aged 13 and under is given a weight of 0.3. Each individual living in the household is then given the household's net equivalent income.

For the explanatory variable age, we simply use the year of birth (`gebjahr`) and subtract it from 2012, while the sex is determined by the variable `bcsex`.

The education level is taken on grounds of the variable `ISCED12` (from the person-related status and generated variables `PGEN`). All observations equal or lower than 5 (higher vocational training) are put in the category no higher education with only those persons with a value of 6 (higher education) considered for the category higher education.

The nationality is obtained directly from the SOEP based on the person's contemporary status (`BCP139`).

The marrital status is taken from the 6 item response to the family status available in the SOEP (`BCP129`), which is reduced to four categories as described in the text.

For the spatial effect we use the variable `bcbula` with the variable *east* set to unity for all federal states formerly belonging to the German Democratic Republic, including the whole of Berlin. West Berlin (as defined prior to 1990) is not accounted for in our sample and treated like a state from the former East.

A.2 Alternatives to SADR

As pointed out in the text, regression approaches beyond the mean are receiving increasing attention in the health-related literature. In this section, we compare SADR to three prominent alternative estimation approaches that allow for the consideration of the outcome in a distributional perspective rather than focussing solely on the expectation of the response variable.

A.2.1 Distribution regression

Chernozhukov et al. (2013) recently proposed the estimation of conditional distributions by representing the distribution's cdf through a finite set of ordered response outcomes whose transition probabilities are modelled by regression coefficients that are allowed to vary across the discretized space of the response. Thus a flexible non-parametric estimate of the conditional distribution can be obtained. This form of model is termed 'distribution regression' in the econometric literature and is closely related to the so-called 'conditional transformation models' (Hothorn et al., 2014) in the statistical literature.

Similarly to quantile regression, the advantage of this approach over SADR is the avoidance of a parametric assumption for the conditional distribution. However, the approach suffers from the drawback that it is generally less stable than SADR, especially if the sample size is limited and/or the model becomes highly complex due to very fine discretization of the response variable's support. Under the (without a doubt strong) assumption that a suitable parametric representation of the conditional distribution can be found, SADR provides a much more parsimonious and stable model, especially for continuous distributions that require a fine discretisation.

A.2.2 Quantile regression

Since their proposition by Koenker and Bassett (1978), quantiles have enjoyed a growing popularity and numerous extensions have been proposed (see among others Bondell et al., 2010; Galvao et al., 2013; Waldmann et al., 2013). Quantile regression allows for the estimation of conditional quantiles, i.e. $q_\tau(Y|x1, \dots, x_K)$, which can be estimated in the usual regression set-up by using an asymmetrically weighted $\ell 1$-norm as a loss function rather than the $\ell 2$-norm used for the least-squares estimation. Using these conditional quantile estimates q_τ for a sufficiently rich grid across $\tau \in (0, 1)$, it is in principle straight forward to retrieve the whole conditional distribution.

Like distribution regression, the quantile-based approach has the advantage that it doesn't require a parametric assumption for the response distribution. However, it also has the drawback that it is generally to be seen as less stable, especially in the tails of the distribution and without additional precautions it may suffer from quantile-crossing. Furthermore, the quantile-based based approach is ill-suited for ordered categorical responses which are of great importance in the health literature. Under the assumption that a sufficiently good parametric approximation to the conditional distribution of the response can be found, SADR thus generally provides more stable estimates for the conditional distributions than an approach based on quantile regression, especially for distributions of non-continuous outcomes, such as SRH.

In summary, SADR arguably provides a somewhat assumptive but very stable model framework for the estimation of conditional health distributions that works for both ordered categorical and continuous outcome distributions.

A.2.3 Recentred influence functions

Another distributional perspective which has gained a lot of traction in recent years is the use recentred influence functions (RIF), put forward by Firpo et al. (2009). Also aiming at an analysis beyond the mean, RIF-based analysis focusses on the marginal rather than the conditional distribution of response. In the linear model one can easily show that both interpretations relate to the same regression effects due to the linearity of the relation between the expectation of the response and the regression coefficients. However, when moving towards distributional regression, all regression specifications are inherently depending on a nonlinear relation between the response and the regression coefficients, even if a purely parametric specification is chosen for the predictor. As a consequence, a model fitted for the conditional, independent perspective does not easily allow to transfer results to the marginal perspective. To overcome this difficulty, (Firpo et al., 2009) introduced 'recentered influence functions' that rely on the notion of influence functions from robust statistics where one aims at identifying the impact of a change in the covariate distribution on marginal features of the response distribution. This is accomplished by a linearisation of the influence function such that simple regression steps can be applied to a transformed response variable.

The main difference between SADR and RIF from a conceptual perspective then relates to the interpretation of changes in the covariates and their induced changes in the response distribution. In the conditional setting, we are interested in changes in the covariates of the individual (or, equivalently, the group of individuals that shares the same covariate information) and the associated changes in the response distribution given this set of covariates. In the marginal setting, one rather evaluates the change in the population distribution of the response when the population distribution of the covariate changes.

Carrieri and Jones (2016) utilise recentered influence functions to study the effect of changes in income on the population distribution of certain biomarkers associated with different types of diseases. While approaching a conceptually related question as we do, our focus is different in the sense that we are exactly interested in who is benefitting from larger income or, vice versa, who suffers from poverty in terms of health-related outcomes. This is a question that requires a conditional perspective which can not be taken based on recentered influence functions.

A.3 Model Selection

For model selection, we use the DIC which has been shown to work well for distributional regression models (Klein et al., 2015). As an exemplary application, we consider the comparison of three regression models for the PCS with differing distributional assumptions. The issue of variable selection could be addressed analogously but is left aside here for reasons of brevity. For the model selection, we thus confine ourselves to a comparison of the following three models:

$M1$ As benchmark model for the PCS, we consider a homoskedastic, gaussian model. In this model the focus is solely directed towards the expectation (μ), with the other parameter (σ^2) considered a nuisance parameter and set as a constant. This is the standard assumption used for most generalised linear models employed in the literature on the health-income relation.

$M2$ As a second model, we consider a heteroskedastic, gaussian model. In this model the variance is now no longer considered a constant as we explicitly allow the standard deviation of the normal distribution to vary across the covariate space. While this already considerably enhances flexibility, the normal distribution is by definition symmetric such that it does not allow for the modelling of changing skewness over the covariate space.

$M3$ As a third model, we consider a two parameter gamma distribution with both parameters allowed to vary across the covariate space. In contrast to the normal distribution, the gamma distribution is not confined to a symmetric form and varies its skewness in relation to its scale parameter.

	Modelassumption	male	female
$M1$	$\mathcal{N}(\mu$ varying, σ^2 constant)	19871.1	23154.2
$M2$	$\mathcal{N}(\mu$ varying, σ^2 varying)	19484.2	22881.6
$M3$	Γ (μ varying, σ varying)	18977.6	22492.3

Table 4: DIC results on distributional assumptions for PCS.

The resultant DICs for these three models are displayed in Table 4. The DICs displayed in Table 4 indicate that out of the three distributions the gamma distribution ($M3$) is the best suited distribution, as it has the lowest DIC both for males and females.

A.4 Random Effects

Table 5 displays the random effects for the individual federal states for the multinomial. The federal states are are abbreviated according to the abbreviations for regions at the EU level[8].
As one can observe there are significant changes associated with the different states, which is in line with the literature on regional health differences (Eibich and Ziebarth, 2014).

	males			
	$\eta^{\tilde{\pi}_1}$	$\eta^{\tilde{\pi}_2}$	$\eta^{\tilde{\pi}_3}$	$\eta^{\tilde{\pi}_4}$
SH	-0.229[-0.413;-0.036]	0.113[0.008; 0.225]	0.001[-0.120; 0.135]	-0.057[-0.323; 0.228]
HH	-0.212[-0.390;-0.020]	-0.247[-0.352;-0.135]	-0.058[-0.177; 0.074]	0.082[-0.178; 0.359]
NI	-0.111[-0.296; 0.080]	-0.106[-0.207; 0.002]	0.347[0.230; 0.480]	-0.437[-0.689;-0.151]
HB	0.846[0.664; 1.036]	0.122[0.018; 0.233]	-0.498[-0.618;-0.363]	0.558[0.303; 0.850]
NW	0.091[-0.092; 0.283]	-0.005[-0.105; 0.103]	0.097[-0.019; 0.226]	-0.180[-0.430; 0.100]
HE	-0.170[-0.353; 0.021]	-0.043[-0.145; 0.065]	0.076[-0.041; 0.206]	-0.043[-0.294; 0.233]
RP	0.180[-0.001; 0.374]	0.076[-0.024; 0.185]	-0.059[-0.181; 0.073]	-0.226[-0.486; 0.056]
BW	-0.225[-0.407;-0.033]	0.003[-0.098; 0.112]	0.014[-0.102; 0.144]	-0.546[-0.798;-0.271]
BY	-0.169[-0.350; 0.021]	-0.124[-0.225;-0.015]	0.179[0.059; 0.309]	0.351[0.101; 0.631]
SL	0.558[0.381; 0.751]	0.478[0.380; 0.590]	0.143[0.026; 0.269]	1.114[0.858; 1.390]
BE	0.085[-0.107; 0.272]	0.135[0.017; 0.245]	0.063[-0.078; 0.199]	0.082[-0.211; 0.364]
BB	0.185[-0.016; 0.375]	-0.131[-0.250;-0.020]	0.245[0.101; 0.381]	0.085[-0.211; 0.373]
MV	0.223[0.019; 0.416]	0.103[-0.016; 0.216]	-0.206[-0.346;-0.068]	0.060[-0.231; 0.348]
SN	0.209[0.010; 0.398]	0.038[-0.079; 0.152]	0.051[-0.092; 0.186]	-0.402[-0.691;-0.119]
ST	-0.210[-0.413;-0.013]	0.237[0.118; 0.349]	0.173[0.032; 0.308]	0.541[0.249; 0.825]
TH	0.124[-0.079; 0.321]	-0.076[-0.194; 0.037]	-0.048[-0.186; 0.090]	0.329[0.042; 0.618]
	females			
	$\eta^{\tilde{\pi}_1}$	$\eta^{\tilde{\pi}_2}$	$\eta^{\tilde{\pi}_3}$	$\eta^{\tilde{\pi}_4}$
SH	-0.285[-0.473;-0.091]	-0.168[-0.271;-0.056]	0.001[-0.078; 0.082]	0.078[-0.113; 0.276]
HH	-0.005[-0.190; 0.195]	-0.309[-0.413;-0.197]	-0.041[-0.118; 0.043]	-0.014[-0.202; 0.190]
NI	0.216[0.026; 0.412]	0.170[0.069; 0.277]	-0.086[-0.163;-0.005]	0.215[0.031; 0.416]
HB	0.275[0.087; 0.471]	0.051[-0.051; 0.165]	0.182[0.104; 0.267]	0.306[0.114; 0.518]
NW	0.175[-0.013; 0.375]	0.171[0.070; 0.277]	-0.050[-0.128; 0.033]	-0.036[-0.225; 0.162]
HE	0.126[-0.061; 0.324]	0.095[-0.007; 0.202]	-0.044[-0.120; 0.038]	-0.246[-0.432;-0.050]
RP	-0.040[-0.226; 0.160]	0.094[-0.007; 0.204]	0.169[0.093; 0.253]	-0.376[-0.567;-0.168]
BW	0.317[0.131; 0.514]	-0.095[-0.195; 0.015]	-0.075[-0.152; 0.006]	-0.046[-0.234; 0.157]
BY	-0.130[-0.316; 0.066]	0.072[-0.029; 0.181]	0.000[-0.079; 0.082]	0.135[-0.055; 0.335]
SL	-0.125[-0.306; 0.074]	0.154[0.055; 0.265]	0.066[-0.013; 0.146]	0.358[0.169; 0.554]
BE	-0.320[-0.517;-0.126]	0.313[0.195; 0.424]	-0.093[-0.189;-0.003]	0.755[0.528; 0.968]
BB	0.230[0.024; 0.421]	-0.177[-0.295;-0.068]	0.037[-0.066; 0.128]	0.151[-0.081; 0.365]
MV	-0.182[-0.389; 0.017]	0.230[0.114; 0.341]	0.267[0.164; 0.359]	-0.520[-0.751;-0.300]
SN	-0.131[-0.332; 0.068]	0.128[0.011; 0.238]	0.118[0.017; 0.209]	-0.005[-0.232; 0.210]
ST	0.997[0.792; 1.200]	-0.130[-0.247;-0.019]	-0.004[-0.104; 0.086]	-0.141[-0.372; 0.070]
TH	-0.024[-0.230; 0.174]	-0.092[-0.209; 0.018]	-0.178[-0.278;-0.087]	0.188[-0.038; 0.400]

Table 5: Random effects for federal states on $\eta^{\tilde{\pi}_1}$,$\eta^{\tilde{\pi}_2}$,$\eta^{\tilde{\pi}_3}$ and $\eta^{\tilde{\pi}_4}$ for SRH.

Table 6 displays the random effects for the parameters of the gamma distribution.
As for the multinomial case, we can observe significant health variations for both parameters

[8]See http://www.bmelv-statistik.de/de/daten-tabellen-suche/abkuerzungen-der-bundeslaender/.

for some federal states, although in case of σ significance is restricted to Hesse and Baden-Wuerttemberg for males alone.

	males		**females**	
	η^{μ}	η^{σ}	η^{μ}	η^{σ}
SH	0.011[-0.010; 0.031]	-0.009[-0.161; 0.160]	0.008[-0.012; 0.028]	-0.008[-0.092; 0.070]
HH	-0.025[-0.049;-0.003]	-0.045[-0.221; 0.116]	-0.017[-0.039; 0.003]	0.019[-0.055; 0.128]
NI	0.000[-0.019; 0.017]	-0.046[-0.162; 0.080]	0.005[-0.013; 0.021]	-0.025[-0.110; 0.036]
HB	-0.003[-0.025; 0.022]	-0.119[-0.335; 0.046]	-0.003[-0.026; 0.020]	-0.009[-0.102; 0.071]
NW	0.007[-0.011; 0.020]	-0.022[-0.123; 0.095]	0.012[-0.004; 0.026]	-0.007[-0.070; 0.056]
HE	0.000[-0.019; 0.016]	0.135[0.013; 0.294]	0.000[-0.017; 0.016]	0.029[-0.033; 0.124]
RP	0.009[-0.010; 0.028]	0.024[-0.111; 0.164]	0.003[-0.015; 0.021]	-0.017[-0.099; 0.046]
BW	-0.010[-0.026; 0.004]	0.219[0.104; 0.352]	-0.007[-0.023; 0.006]	0.025[-0.036; 0.109]
BY	-0.006[-0.023; 0.009]	0.017[-0.095; 0.134]	-0.002[-0.017; 0.013]	0.013[-0.047; 0.081]
SL	0.019[-0.006; 0.047]	-0.147[-0.369; 0.027]	0.002[-0.021; 0.026]	-0.017[-0.125; 0.062]
BE	-0.002[-0.024; 0.020]	-0.127[-0.297; 0.043]	0.001[-0.021; 0.020]	-0.001[-0.078; 0.082]
BB	0.002[-0.019; 0.023]	0.095[-0.057; 0.263]	0.002[-0.017; 0.022]	0.016[-0.061; 0.115]
MV	0.003[-0.021; 0.025]	0.068[-0.095; 0.272]	0.011[-0.012; 0.033]	-0.001[-0.082; 0.082]
SN	-0.001[-0.022; 0.019]	-0.058[-0.203; 0.074]	-0.001[-0.022; 0.019]	-0.006[-0.088; 0.064]
ST	0.007[-0.014; 0.031]	-0.013[-0.182; 0.149]	0.003[-0.018; 0.026]	-0.017[-0.121; 0.061]
TH	-0.008[-0.030; 0.015]	0.034[-0.134; 0.210]	-0.015[-0.036; 0.007]	0.010[-0.068; 0.107]

Table 6: Random effects for federal states on η^{μ} and η^{σ} for PCS.

A.5 Other Covariate Combinations

In this section we show the seven health measures displayed in Section 4.5 for "average Joe" and "average Jane" for a different set of characteristics. For the sake of brevity and simplicity, we constrain the sets considered to 7 different sets, always varying only one covariate while keeping all the other covariates at the values used for "average Joe" and "average Jane".

Tables 7 and 8 display the seven health measures for two other ages, namely the first and the third quartile of the ages in the sample: 40 years and 66 years, respectively.
As can be seen from the tables, the general structure persists, whereby the differences between the health measures for the two different income levels becomes more pronounced as the focus is shifted towards the lower end of the health spectrum. Moreover, one may note that the health situation is generally better for younger individuals than for older individuals, which is to be expected given the physical deterioration of the body as part of the ageing process.

Table 9 displays the seven health measures for non-German nationals. Again the basic pattern remains such that income related differences are more more pronounced (in relative terms) in the lower end of the health spectrum. One other thing which can be observed from the table is the lower health risks and slightly better average health enjoyed by non-German nationals, which is in line with the healthy-migrant effect found in the literature (see Bjornstrom and Kuhl, 2014).

Table 10 displays the seven health measures for individuals in the fourth education level, i.e. those individuals with higher education. Again, the basic pattern remains in the sense that income related differences are significantly higher when focussing on the lower end of the health spectrum. As one would expect, higher education levels additionally mean a higher average health outcome and lower risk measures for both men and women, c.p.

Table 11 displays the seven health measures for individuals in the third marrital status, i.e. those individuals who are sinlge, which is the second most frequent observed marrital status in our sample. As above, the basic pattern is such that income related differences are significantly higher when looking at the lower end of the health spectrum. In terms of the absolute levels, we observe a slightly lower average health and slightly elevated risks, which is reasonable given the positive health effects that stable relationships are thought to have.

Tables 12 and 13 display the seven health measures for two other federal states in Germany, namely BadenWurttemberg and Mecklenburg-Western Pomerania. BadenWurttemberg is a very wealth state in the South-West of Germany while Mecklenburg-Western Pomerania is an economically rather depressed state in the North-East of Germany. As can be seen from the tables, the general structure persists again, i.e. the differences between the health measures for the two different income levels becomes more pronounced as the focus is shifted towards the lower end of the health spectrum. As one would expect the wealthier federal state BadenWurttemberg also features better health measures than the poorer Mecklenburg-Western Pomerania, which is to be expected given the positive effects of gdp on state finances and thus available funds for the health infrastructure in these regions.

	males		
	15,000€	30,000€	Difference
$\mathcal{R}_{M3}$	0.53[0.53; 0.53]	0.46[0.45; 0.46]	13.94%[13.81%;14.07%]
$\mathcal{R}_{M2}$	0.14[0.14; 0.14]	0.09[0.09; 0.09]	37.39%[37.14%;37.63%]
$\mathcal{R}_{M1}$	0.03[0.03; 0.03]	0.02[0.01; 0.02]	42.79%[42.34%;43.30%]
μ	48.95[47.54;50.15]	50.18[48.72;51.36]	2.51%[2.94%; 1.96%]
$\mathcal{R}_{C0.50}$	0.55[0.49; 0.61]	0.49[0.43; 0.56]	10.30%[7.13%;13.42%]
$\mathcal{R}_{C0.20}$	0.14[0.10; 0.19]	0.10[0.07; 0.14]	30.14%[22.74%;37.99%]
$\mathcal{R}_{C0.05}$	0.02[0.01; 0.03]	0.01[0.00; 0.02]	52.95%[40.94%;63.24%]
	females		
	15,000€	30,000€	Difference
$\mathcal{R}_{M3}$	0.46[0.46; 0.46]	0.38[0.38; 0.39]	16.65%[16.51%;16.79%]
$\mathcal{R}_{M2}$	0.15[0.15; 0.15]	0.10[0.10; 0.10]	33.38%[33.13%;33.64%]
$\mathcal{R}_{M1}$	0.02[0.02; 0.03]	0.01[0.01; 0.01]	44.05%[43.60%;44.56%]
μ	48.50[47.12;49.65]	49.95[48.57;51.09]	2.99%[3.50%; 2.45%]
$\mathcal{R}_{C0.50}$	0.57[0.51; 0.63]	0.50[0.44; 0.57]	11.61%[8.65%;14.64%]
$\mathcal{R}_{C0.20}$	0.15[0.12; 0.20]	0.11[0.08; 0.15]	31.01%[24.75%;37.99%]
$\mathcal{R}_{C0.05}$	0.02[0.01; 0.03]	0.01[0.00; 0.02]	52.86%[42.70%;61.48%]

Table 7: Seven measures on the health-income association for 40yrs of age (all other covariates the same).

	males		
	15,000€	30,000€	Difference
$\mathcal{R}_{M3}$	0.77[0.77; 0.77]	0.72[0.71; 0.72]	7.21%[7.13%; 7.28%]
$\mathcal{R}_{M2}$	0.27[0.27; 0.27]	0.19[0.18; 0.19]	30.32%[30.12%;30.56%]
$\mathcal{R}_{M1}$	0.06[0.06; 0.06]	0.04[0.04; 0.04]	36.05%[35.63%;36.56%]
μ	41.10[39.48;42.51]	42.49[40.91;43.93]	3.38%[4.04%; 2.70%]
$\mathcal{R}_{C0.50}$	0.81[0.76; 0.85]	0.78[0.73; 0.84]	2.71%[1.47%; 4.24%]
$\mathcal{R}_{C0.20}$	0.43[0.37; 0.49]	0.37[0.32; 0.44]	12.99%[9.53%;17.13%]
$\mathcal{R}_{C0.05}$	0.13[0.09; 0.18]	0.09[0.06; 0.13]	29.29%[22.27%;37.06%]
	females		
	15,000€	30,000€	Difference
$\mathcal{R}_{M3}$	0.70[0.70; 0.70]	0.63[0.63; 0.63]	9.66%[9.56%; 9.76%]
$\mathcal{R}_{M2}$	0.31[0.30; 0.31]	0.23[0.23; 0.23]	25.22%[25.02%;25.46%]
$\mathcal{R}_{M1}$	0.07[0.07; 0.07]	0.04[0.04; 0.05]	36.48%[36.05%;36.99%]
μ	40.83[39.25;42.18]	42.49[40.95;43.83]	4.07%[4.77%; 3.36%]
$\mathcal{R}_{C0.50}$	0.83[0.79; 0.87]	0.80[0.75; 0.84]	3.73%[2.43%; 5.35%]
$\mathcal{R}_{C0.20}$	0.44[0.38; 0.50]	0.37[0.31; 0.43]	15.70%[12.28%;19.74%]
$\mathcal{R}_{C0.05}$	0.12[0.09; 0.16]	0.08[0.06; 0.12]	33.37%[26.40%;40.95%]

Table 8: Seven measures on the health-income association for 66yrs of age (all other covariates the same).

	males		
	15,000€	30,000€	Difference
$\mathcal{R}_{M3}$	0.55[0.55; 0.55]	0.48[0.47; 0.48]	13.67%[13.54%;13.80%]
$\mathcal{R}_{M2}$	0.17[0.17; 0.17]	0.11[0.11; 0.11]	36.18%[35.92%;36.45%]
$\mathcal{R}_{M1}$	0.03[0.03; 0.03]	0.02[0.02; 0.02]	41.70%[41.21%;42.19%]
μ	46.39[44.77;47.91]	47.67[45.96;49.12]	2.75%[3.24%; 2.17%]
$\mathcal{R}_{C0.50}$	0.65[0.59; 0.72]	0.61[0.54; 0.69]	6.63%[4.53%; 9.52%]
$\mathcal{R}_{C0.20}$	0.23[0.17; 0.29]	0.17[0.12; 0.23]	23.44%[17.16%;30.42%]
$\mathcal{R}_{C0.05}$	0.04[0.02; 0.06]	0.02[0.01; 0.04]	43.76%[33.40%;55.07%]
	females		
	15,000€	30,000€	Difference
$\mathcal{R}_{M3}$	0.56[0.56; 0.56]	0.48[0.48; 0.49]	13.62%[13.49%;13.76%]
$\mathcal{R}_{M2}$	0.28[0.28; 0.29]	0.21[0.20; 0.21]	26.99%[26.74%;27.24%]
$\mathcal{R}_{M1}$	0.05[0.05; 0.05]	0.03[0.03; 0.03]	38.68%[38.21%;39.21%]
μ	44.34[42.71;45.83]	45.92[44.27;47.32]	3.55%[4.11%; 2.90%]
$\mathcal{R}_{C0.50}$	0.74[0.68; 0.79]	0.69[0.62; 0.75]	6.52%[4.50%; 8.94%]
$\mathcal{R}_{C0.20}$	0.30[0.24; 0.36]	0.23[0.17; 0.30]	22.57%[17.10%;28.81%]
$\mathcal{R}_{C0.05}$	0.06[0.03; 0.09]	0.03[0.02; 0.06]	43.04%[33.41%;52.54%]

Table 9: Seven measures on the health-income association for non-German nationals (all other covariates the same).

	15,000€	30,000€	Difference
		males	
$\mathcal{R}_{M3}$	0.62[0.62; 0.62]	0.55[0.54; 0.55]	11.60%[11.48%;11.69%]
$\mathcal{R}_{M2}$	0.20[0.20; 0.20]	0.13[0.13; 0.13]	34.19%[33.94%;34.41%]
$\mathcal{R}_{M1}$	0.05[0.05; 0.05]	0.03[0.03; 0.03]	39.59%[39.14%;40.04%]
μ	46.16[44.83;47.44]	47.41[46.12;48.71]	2.71%[3.29%; 2.19%]
$\mathcal{R}_{C0.50}$	0.65[0.60; 0.70]	0.61[0.55; 0.67]	6.47%[4.36%; 8.67%]
$\mathcal{R}_{C0.20}$	0.24[0.20; 0.29]	0.19[0.15; 0.24]	21.43%[16.60%;27.13%]
$\mathcal{R}_{C0.05}$	0.05[0.03; 0.07]	0.03[0.02; 0.04]	41.04%[32.26%;50.20%]
		females	
	15,000€	30,000€	Difference
$\mathcal{R}_{M3}$	0.51[0.51; 0.52]	0.44[0.43; 0.44]	15.13%[15.00%;15.24%]
$\mathcal{R}_{M2}$	0.16[0.16; 0.17]	0.11[0.11; 0.11]	32.17%[31.93%;32.40%]
$\mathcal{R}_{M1}$	0.03[0.03; 0.04]	0.02[0.02; 0.02]	42.60%[42.15%;43.06%]
μ	46.34[45.08;47.53]	47.83[46.60;49.05]	3.22%[3.81%; 2.66%]
$\mathcal{R}_{C0.50}$	0.64[0.60; 0.69]	0.59[0.54; 0.64]	8.21%[6.20%;10.52%]
$\mathcal{R}_{C0.20}$	0.24[0.20; 0.28]	0.18[0.15; 0.22]	23.80%[19.28%;28.84%]
$\mathcal{R}_{C0.05}$	0.05[0.03; 0.06]	0.03[0.02; 0.04]	42.52%[34.90%;50.27%]

Table 10: Seven measures on the health-income association for the fourth education level (all other covariates the same).

	15,000€	30,000€	Difference
		males	
$\mathcal{R}_{M3}$	0.73[0.73; 0.73]	0.67[0.67; 0.67]	8.20%[8.11%; 8.29%]
$\mathcal{R}_{M2}$	0.23[0.23; 0.24]	0.16[0.16; 0.16]	31.79%[31.57%;32.03%]
$\mathcal{R}_{M1}$	0.07[0.06; 0.07]	0.04[0.04; 0.04]	37.03%[36.59%;37.53%]
μ	45.14[43.59;46.61]	46.42[44.85;47.91]	2.85%[3.45%; 2.27%]
$\mathcal{R}_{C0.50}$	0.68[0.62; 0.73]	0.65[0.58; 0.70]	5.32%[3.47%; 7.46%]
$\mathcal{R}_{C0.20}$	0.29[0.23; 0.34]	0.23[0.18; 0.29]	18.82%[13.69%;24.39%]
$\mathcal{R}_{C0.05}$	0.07[0.04; 0.10]	0.04[0.02; 0.07]	36.79%[27.47%;46.40%]
		females	
	15,000€	30,000€	Difference
$\mathcal{R}_{M3}$	0.64[0.64; 0.64]	0.57[0.56; 0.57]	11.34%[11.23%;11.46%]
$\mathcal{R}_{M2}$	0.27[0.26; 0.27]	0.19[0.19; 0.20]	27.08%[26.86%;27.34%]
$\mathcal{R}_{M1}$	0.06[0.06; 0.07]	0.04[0.04; 0.04]	37.87%[37.41%;38.37%]
μ	45.07[43.65;46.43]	46.60[45.13;48.03]	3.39%[3.97%; 2.83%]
$\mathcal{R}_{C0.50}$	0.69[0.64; 0.74]	0.65[0.58; 0.70]	6.94%[5.02%; 9.37%]
$\mathcal{R}_{C0.20}$	0.28[0.23; 0.33]	0.22[0.17; 0.28]	21.91%[16.96%;27.61%]
$\mathcal{R}_{C0.05}$	0.06[0.04; 0.08]	0.03[0.02; 0.06]	41.03%[32.39%;49.66%]

Table 11: Seven measures on the health-income association for the third marrital status (all other covariates the same).

	males		
	15,000€	30,000€	Difference
$\mathcal{R}_{M3}$	0.60[0.59; 0.60]	0.52[0.52; 0.53]	12.50%[12.37%;12.62%]
$\mathcal{R}_{M2}$	0.17[0.17; 0.18]	0.11[0.11; 0.11]	35.70%[35.44%;35.98%]
$\mathcal{R}_{M1}$	0.04[0.04; 0.04]	0.02[0.02; 0.03]	40.93%[40.47%;41.44%]
μ	47.24[45.82;48.69]	48.48[47.06;49.95]	2.64%[3.19%; 2.06%]
$\mathcal{R}_{C0.50}$	0.61[0.55; 0.67]	0.56[0.50; 0.63]	7.33%[5.11%;10.10%]
$\mathcal{R}_{C0.20}$	0.21[0.16; 0.26]	0.16[0.11; 0.21]	23.18%[17.35%;29.86%]
$\mathcal{R}_{C0.05}$	0.04[0.02; 0.06]	0.02[0.01; 0.04]	42.91%[33.07%;53.63%]
	females		
	15,000€	30,000€	Difference
$\mathcal{R}_{M3}$	0.55[0.55; 0.56]	0.48[0.47; 0.48]	14.07%[13.95%;14.22%]
$\mathcal{R}_{M2}$	0.20[0.20; 0.21]	0.14[0.14; 0.14]	30.39%[30.17%;30.69%]
$\mathcal{R}_{M1}$	0.04[0.04; 0.04]	0.02[0.02; 0.02]	41.24%[40.84%;41.82%]
μ	46.46[45.10;47.83]	47.95[46.61;49.34]	3.22%[3.83%; 2.66%]
$\mathcal{R}_{C0.50}$	0.65[0.59; 0.71]	0.59[0.53; 0.66]	8.68%[6.41%;11.35%]
$\mathcal{R}_{C0.20}$	0.22[0.18; 0.27]	0.17[0.12; 0.21]	25.73%[20.76%;31.82%]
$\mathcal{R}_{C0.05}$	0.04[0.02; 0.05]	0.02[0.01; 0.03]	45.89%[37.79%;54.90%]

Table 12: Seven measures on the health-income association for BadenWurttemberg (all other covariates the same).

	males		
	15,000€	30,000€	Difference
$\mathcal{R}_{M3}$	0.63[0.63; 0.63]	0.56[0.56; 0.56]	11.53%[11.45%;11.64%]
$\mathcal{R}_{M2}$	0.24[0.24; 0.24]	0.16[0.16; 0.16]	32.72%[32.52%;32.95%]
$\mathcal{R}_{M1}$	0.04[0.04; 0.04]	0.02[0.02; 0.02]	38.73%[38.32%;39.26%]
μ	45.89[44.67;47.15]	47.17[45.91;48.38]	2.78%[3.33%; 2.23%]
$\mathcal{R}_{C0.50}$	0.66[0.61; 0.71]	0.62[0.57; 0.67]	6.05%[4.19%; 8.06%]
$\mathcal{R}_{C0.20}$	0.25[0.21; 0.30]	0.20[0.16; 0.25]	20.59%[15.86%;25.92%]
$\mathcal{R}_{C0.05}$	0.05[0.03; 0.08]	0.03[0.02; 0.05]	39.65%[31.28%;48.80%]
	females		
	15,000€	30,000€	Difference
$\mathcal{R}_{M3}$	0.67[0.66; 0.67]	0.59[0.59; 0.60]	10.80%[10.71%;10.89%]
$\mathcal{R}_{M2}$	0.24[0.23; 0.24]	0.17[0.17; 0.17]	27.97%[27.77%;28.22%]
$\mathcal{R}_{M1}$	0.05[0.05; 0.06]	0.03[0.03; 0.03]	38.76%[38.35%;39.27%]
μ	45.24[44.03;46.44]	46.78[45.54;47.95]	3.40%[4.01%; 2.79%]
$\mathcal{R}_{C0.50}$	0.69[0.65; 0.74]	0.64[0.59; 0.69]	7.27%[5.35%; 9.29%]
$\mathcal{R}_{C0.20}$	0.27[0.23; 0.32]	0.21[0.17; 0.25]	22.87%[18.54%;27.65%]
$\mathcal{R}_{C0.05}$	0.05[0.04; 0.07]	0.03[0.02; 0.05]	41.92%[34.65%;50.19%]

Table 13: Seven measures on the health-income association for Mecklenburg-Western Pomerania (all other covariates the same).