Received: 26 April 2017 | Revised: 19 February 2018 | Accepted: 20 February 2018
DOI: 10.1002/hec.3656

RESEARCH ARTICLE

WILEY Health Economics

Reconsidering the income-health relationship using distributional regression

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

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

Correspondence
Alexander Silbersdorff, Georg August University of Göttingen, Göttingen, Germany.
Email: asilbersdorff@uni-goettingen.de

Funding information
German Research Association DFG

JEL Classification: I14; C13; C21

Abstract
We reconsider the relationship between income and health taking a distributional perspective rather than one centered on conditional expectation. Using structured additive distributional regression, we find that the association between income and 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 bad health is roughly halved when doubling the net equivalent income from 15,000 to 30,000€. This is more than tenfold of the magnitude of change found when considering expected health measures. A distributional perspective thus highlights another dimension of the income–health relation—that the poor are in particular faced with greater health risk at the lower end of the health distribution. We therefore argue that when studying health outcomes, a distributional approach that considers stochastic variation among observationally equivalent individuals is warranted.

KEYWORDS
health inequality, income–health relationship, structured additive distributional regression

1 | INTRODUCTION

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, Adler, & Dow, 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, Sorlie, & Johnson, 1996; Ettner, 1996; McDonough, Duncan, Williams, & House, 1997; Ecob & Davey Smith, 1999; Case, 2014), and recent studies exploiting exogenous variation in income have discussed causal effects of income on health (Frijters, Haisken-DeNew, & Shields, 2005; Lindahl, 2005; Case, 2014; Kuehnle, 2014; Cesarini, Lindqvist, Östling & Wallace, 2016). Scores of papers assess the relationship between income and health status using regression techniques. Both in epidemiology and health economics, however, the vast majority of these employ conventional regression methods (linear and generalized linear models) to assess the effect of variations of income and other covariates on the expected (mean) health outcome. In recent years, a growing number of papers in the health economics literature have noted the need to look beyond the expected outcome (Duclos & Échevin, 2011; Makdissi & Yazbeck, 2014; Carrieri & Jones, 2017; Heckley, Gerdtham, & Kjellsson, 2016; Schiele & Schmitz, 2016). Building on this young literature, we propose to consider, in a regression framework, the full conditional health distributions rather than just their expectations.

A simple trisection of the health distribution between those with high, medium, and low incomes shows that differences in health outcomes by income go beyond the mean. Using data from Germany in 2012, Figure 1 shows the

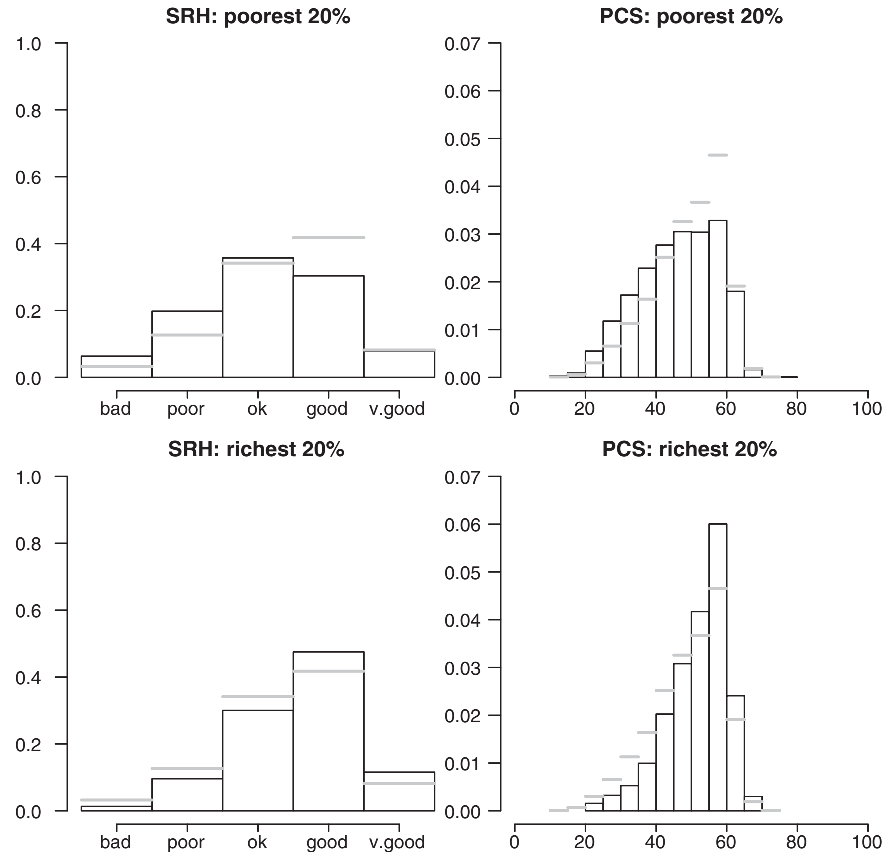

FIGURE 1 A contrast of coarsely conditioned health distributions. Top: lowest 20% of net incomes. Bottom: highest 20% of net incomes. Left: self-rated health (SRH). Right: physical component score (PCS) of the SF-12. Gray lines indicate reference lines of middle group

distribution of two measures of generalized health—self-rated health (SRH) and a physical health score (PCS)—, contrasting those with incomes that are low (bottom 20% w.r.t. net equivalent income; shown in top panels) and high (top 20% w.r.t. net equivalent income; shown in bottom panels) to those with average incomes (three inner quintiles w.r.t net equivalent income; gray lines). As can be seen, the variation in health outcomes is substantially more pronounced in the lower part of the income distribution, whereas those who are economically well off are able to practically eliminate the risk of bad health.

This simple example already shows how an assessment solely based on the conditional distributions' means may neglect or underestimate important aspects of the relationship between income and health. Although 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 leaves aside an important dimension of inequality: the additional health risk suffered by the poor due to other features of the distribution (higher variance, higher skewness, etc.). These features imply heightened probabilities of bad outcomes in the lower tail of the health distribution, above and beyond those that would be expected by a simple shift in the expectation perpetrated by standard regression.

Incorporating the additional health risks incurred by the poor into a measure of the health–income relationship is easily justified on the basis of extra-welfarist (Culyer, 1990) or capability (Sen, 1993) perspectives. In these approaches, bad health, in particular, is thought to negatively impact an individual's welfare or capabilities. By contrast, differences in the upper part of the health distribution, that is, differences between good and very good health, have only a relatively negligible impact on one's welfare or capabilities.

But from a utilitarian perspective, too, there is a compelling rationale for incorporating information based on health risks, which springs from the concavity of utility in health (Berk & Monheit, 2001; Finkelstein, Luttmer, & Notowidigdo, 2009).[1] Given a concave health–utility relationship, risk-averse individuals care not only about their expected levels of health but in particular about the risk of being in bad health. Properly accounting for these health risks is not only of academic interest but could also have ramifications for the evaluation of public health policy, the design of public and private health insurance, or the optimal allocation of funds within health care systems.

[1]For a discussion on concavity for an ordinal health variable, which we will also consider in the following, see Section A.2 in the Supporting Information.

In order to assess the changes in the health risk with respect to income and a set of control variables, the whole conditional health distribution needs to be specified. To this end, we apply the recently developed technique of structured additive distributional regression (SADR, Klein, Kneib, Lang, & Sohn, 2015). Using SADR, we are able to look at both categorical health measures (such as the standard ordered five-response format for SRH) and continuous or quasicontinuous variables (such as the SF-12). The major advantage of SADR over alternative estimation strategies such as quantile regression (Koenker, 2005), conditional transformation models (Hothorn, Kneib, & Bühlmann, 2014), or distribution regression (Chernozhukov, Fernandez-Val, & Melly, 2013) is the estimation stability that the assumption of a parametric distribution lends, especially for small sample sizes (see Sohn, 2017) and in the tails of the distribution, which are critical for the assessment of risks. SADR thus provides a generic regression framework for the different types of health outcomes used in the literature that also yields stable estimates for datasets with limited sample size.

This method also constitutes an alternative to the use of recentered influence functions (Firpo, Fortin, & Lemieux, 2009), which have recently gained traction in the health economics literature (e.g., Carrieri & Jones, 2017; Heckley et al., 2016). In contrast to recentered influence functions, which assess the impact of changes in the covariates on the marginal distribution of health, SADR allows us to focus on the conditional distribution of health for specific subpopulations defined by, for example, income, age, or education.

Using data from Germany, we show that an SADR-based approach identifies inequalities in individuals' health risk beyond differences in expected health status and thus leads to a starkly different assessment of the magnitude of the association between income and health. For example, consider the change in health for an "average Joe"/"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). The relative change in expected health is 3% for both men and women. Using the distributional approach to estimate the risk of being in bad health yields substantially larger estimates of the association between income and health, however: The relative difference in the risk of bad SRH when moving from poor to rich is 39% for men and 40% for women. The comparable risk measure for the SF-12 yields a change of 39%/42%. The proposed distributional regression approach 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.

2 | 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 $x_1,\ldots,x_K$ is related to a regression predictor η via the response function h, that is,

$$\mathbb{E}(Y|x_1,\ \ldots,x_K) = h(\eta).$$

The predictor in turn is usually modeled as a linear combination of the covariates[2] entailed in covariate vector $(x_1,\ \cdots,x_K)^T$, that is,

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

For example, in case of binary outcomes differentiating only between healthy and nonhealthy individuals, a logit or probit model is specified, in which the probability of an outcome $\pi = P(Y=1|x_1,\ \ldots,x_K) = \mathbb{E}(Y|x_1,\ \ldots,x_K)$ is related to the predictors via the cumulative distribution function of the logistic or 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), for more complex outcomes, the expectation alone generally does 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.

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

A distributional perspective is needed to allow us to not just consider the conditional expectation of the health variable of interest, $\mathbb{E}(Y|x_1, \dots, x_K)$, but also to relate the complete underlying conditional distribution, $\mathcal{D}(Y|x_1, \dots, x_K)$, 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.

Here, we will rely on SADR models as introduced in Klein et al. (2015), in which a parametric distribution type is assumed for the conditional distribution $\mathcal{D}(Y|x_1, \dots, x_K)$, 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 $\mathcal{D}(\theta_1(x_1, \dots, x_K), \theta_2(x_1, \dots, x_K), \dots, \theta_L(x_1, \dots, x_K))$ is characterized by a vector of L parameters $\theta_l(x_1,\dots,x_K)$, $l=1,\dots,L$ 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 β entails parameters not only for one predictor but for all L predictors required to specify the response distribution.

As stated above, the main advantage of SADR over alternative approaches is that it lends great estimation stability that is critical for the usually available sample sizes.[3] 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 such as quantile regression, conditional transformation models/distribution regression, or recentered influence functions, we refer the interested reader to Section A.4 in the Supporting Information.

3 | 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).

3.1 | The data

The German SOEP is a longitudinal household survey repeated annually since 1984 (Wagner, Frick & Schupp, 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 & Schupp, 2013).[4] 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 examine both SRH with five response categories and the SF-12 physical health scale, which are ordinal and (quasi) continuous health measures, respectively. The first asks respondents to rate their current health as very good, good, satisfactory, poor, or bad. The SF-12 measures SRH in eight domains (Wagner et al. 2007). We use only the physical health subscale of the SF-12.[5]

Both health measures are related to a set of sociodemographic explanatory variables. The key explanatory variable of interest is income. Here, we consider the log-transform of the annual net (disposable) equivalized household income of an individual, adjusted for household size and composition using the Organisation for Economic Cooperation and Development equivalence scale (LOGINC).

In addition to income, we consider a set of control variables described in Table 1. For further information on the variables, see Section A.1 in the Supporting Information.

[3]Given that distributional assessments are generally much more demanding than simple mean assessments, standard database sizes in the order of 10,000 observations or less may easily suffer from estimation instability as soon as a set of 10 or more covariates is introduced.

[4]In principle, a consideration of the panel dimension is also possible. However, in panel analysis, the interpretation of the conditional distribution would be based on intrapersonal, not interpersonal, variation, which is the focus of the present analysis. For further discussion of this issue, see Sohn (2017).

[5]Differential item functioning by education, age, and sex has been observed for the mental component score (MCS, Fleishman & Lawrence, 2003; Bourion-Bédès et al., 2015), and because the SOEP does not include the institutionalized population, the sample is likely to be nonrepresentative of the population with very low MCS scores.

TABLE 1 Description of variables

Variable	Definition
SRH	Ordered categorical variable for self-rated health with five levels.
PCS	Quasicontinuous measure on physical health.
AGE	The age of an individual measured in years.
AGESQ	The squared age.
LOGINC	The logarithm of the annual net equivalized household income of an individual.
EDU_1	Education level entailing individuals who have only general elementary education or less.
EDU_2	Education level entailing individuals who have completed secondary education.
EDU_3	Education level entailing individuals with higher vocational training.
EDU_4	Education level entailing individuals with completed higher education.
GER	Binary variable that is unity if the nationality is German and zero otherwise.
MAR_1	Marital status of individual living in a partnership.
MAR_2	Marital status of individual being separated or divorced.
MAR_3	Marital status of individual being single.
MAR_4	Marital status of individual being widowed.
REGION	The federal state of residence of the individual.

3.2 | Model specification

3.2.1 | Choice of the response distribution

As discussed in Section 2, a distributional regression approach requires that we specify a suitable parametric distribution that is able to approximate the empirically observed conditional health distributions.

SRH outcomes are measured on an ordinal 5-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 SRH score rather than to differentiate only between two amalgamations of the levels, which are liable to overlooking important differences (see Ziebarth, 2010). 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 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.[6]

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 contrast to the more common symmetric or positively skewed distributions for 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, whereas 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) that enhances numerical stability. Subsequently, we estimate the conditional distributions of the transformed PCS using the well-known two parameter gamma distribution.[7] Once this conditional distribution is estimated, one can easily obtain the conditional distribution of the original PCS measure

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

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

$$p(y|\mu, s) = \left(\frac{s}{\mu}\right)^s \frac{y^{s-1}}{G(s)} \exp\left(-\frac{s}{\mu}y\right), \tag{4}$$

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

by simply applying the inverse transform, g_{PCS}^{-1}. Note that the gamma distribution is invariant under scaling such that we effectively model a shifted, reversed, scaled gamma distribution for the health scores.

For both the categorical SRH scores and the (quasi) continuous SF-12, we thus specify parametric conditional health distributions that 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.

3.2.2 | Predictor specification

For the sake of simplicity, we will specify one generic predictor setup that is applied to all parameters, that is,

$$\begin{aligned}\eta^{\theta_l} &= \beta_0^{\theta_l} + \beta_1^{\theta_l}\mathrm{AGE} + \beta_2^{\theta_l}\mathrm{AGE}SQ + \beta_3^{\theta_l}\mathrm{LOGINC} + \beta_4^{\theta_l}\mathrm{GER} + \beta_5^{\theta_l}\mathrm{EDU}_2 + \beta_6^{\theta_l}\mathrm{EDU}_3 + \beta_7^{\theta_l}\mathrm{EDU}_4 \\ &\quad + \beta_8^{\theta_l}\mathrm{MAR}_2 + \beta_9^{\theta_l}\mathrm{MAR}_3 + \beta_{10}^{\theta_l}\mathrm{MAR}_4 + \beta_{11}^{\theta_l}\mathrm{EAST} + \gamma_{\mathrm{REGION}}^{\theta_l},\end{aligned} \tag{5}$$

where η_l is the predictor for the lth parameter of the response distribution (see the Supporting Information for the equations for each single parameter). The explanatory variables (defined as outlined in Section 3.1) are all included in a linear fashion bar the two last terms 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 and the Federal German Republic. The differences within the former German Democratic Republic and Federal German Republic are captured by random effects, denoted by γ_{REGION}. 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 whereas the exponential response function is used to ensure positivity of the two parameters for the gamma distribution.

3.3 | Parameter estimates

The estimation is done in the software BayesX (Belitz et al. 2015) that employs Markov Chain Monte Carlo (MCMC) simulation techniques to estimate posterior distributions in a Bayesian framework.[8] See Klein et al. (2015) for details on the estimation procedure. In the following setup, we use noninformative 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 1 million MCMC realizations that 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. Although some of the parameters are interpretable in their own right (e.g., μ for the gamma distribution), we focus on evaluating the resultant distributions rather than the single parameters' estimates.

Table 2 displays the estimates for the covariate effects on the predictors of the sequential logits for the SRH outcomes. Here, we display the medians of the posterior distributions with the 95% (symmetric) credible intervals denoted in brackets. In order to conserve space, we do not display the estimates for the random effect of the individual federal states but show them separately in Table S6.

Although 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.

Table 3 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

[8]Note that the estimation of such models is also possible in a frequentist framework. However, the Bayesian approach has proven to be not only faster but also computationally more stable. See Sohn (2017) for details.

TABLE 2 Linear effects on $\eta^{\tilde{\pi}_1}, \eta^{\tilde{\pi}_2}, \eta^{\tilde{\pi}_3}$ and $\eta^{\tilde{\pi}_4}$ for PCS

Males				
	$\eta^{\tilde{\pi}_1}$	$\eta^{\tilde{\pi}_2}$	$\eta^{\tilde{\pi}_3}$	$\eta^{\tilde{\pi}_4}$
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]
AGESQ	−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]	4−0.009[−0.120;0.100]	−0.089[−0.333;0.148]
Females				
	$\eta^{\tilde{\pi}_1}$	$\eta^{\tilde{\pi}_2}$	$\eta^{\tilde{\pi}_3}$	$\eta^{\tilde{\pi}_4}$
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]
AGESQ	−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]

Note. PCS = physical health score.

TABLE 3 Linear effects on η^{μ} and η^{s} for PCS

	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]
AGESQ	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]

Note. PCS = physical health score.

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 and other variables have a significant impact that indicates complex changes across the covariate space that go beyond the changes in the conditional mean on which standard regression techniques focus.

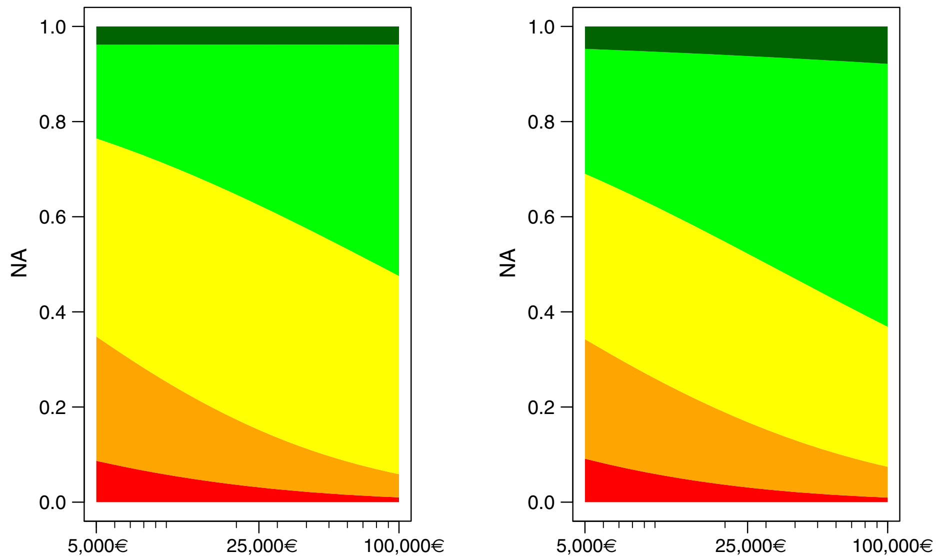

FIGURE 2 Conditional income effects for self-rated health for average Joe (left) and average Jane (right). From bad health (red) to very good health (dark green) [Colour figure can be viewed at wileyonlinelibrary.com]

3.4 | Considering the distributional changes

Because we employ nonlinear 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 & 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," that is, 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.[9]

In assessing the health differentials associated with different income levels, we focus on relative rather than absolute differences.[10] This choice is based on the recommendation to use relative inequality measures when mainly concerned about assessing health inequality rather than the absolute level of a health risk (see Harper et al., 2010). Although the absolute levels of any health measure are clearly of importance in any intertemporal 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 easily be compared for an assessment of inequality.

Figure 2 makes visible how the distribution of SRH changes with income, displaying the change in the probability of falling into one of five SRH 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 2. 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 of 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. Let us, for example, consider the conditional health distributions of males. Although the share of respondents in "very good" health is nearly constant across the income range, the

[9]See Section A.1 in the Supporting Information for the covariates' distributions underlying this choice. For the continuous variable age, we consider the arithmetic mean in our sample, whereas for the other categorical variables, we consider the mode. See Section A.7 in the Supporting Information 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 following.

[10]The distinction between relative and absolute inequality has been discussed extensively in the health inequalities literature (see Mechanic, 2002; Oliver, Healey, & Le Grand, 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).

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 "bad" health, the decrease is even larger at 88% (from 0.09 to 0.01). This shows that dichotomizing the outcome, for example, 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 25th percentile, i.e., the median for the poorer half of the population) versus 30,000€ (roughly corresponds to the 75th 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, that is, 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 (very) 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€ vs. 30,000€). We can define the following three health measures, which dichotomize the distribution in three different ways:

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

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, $\mathbf{x}$.

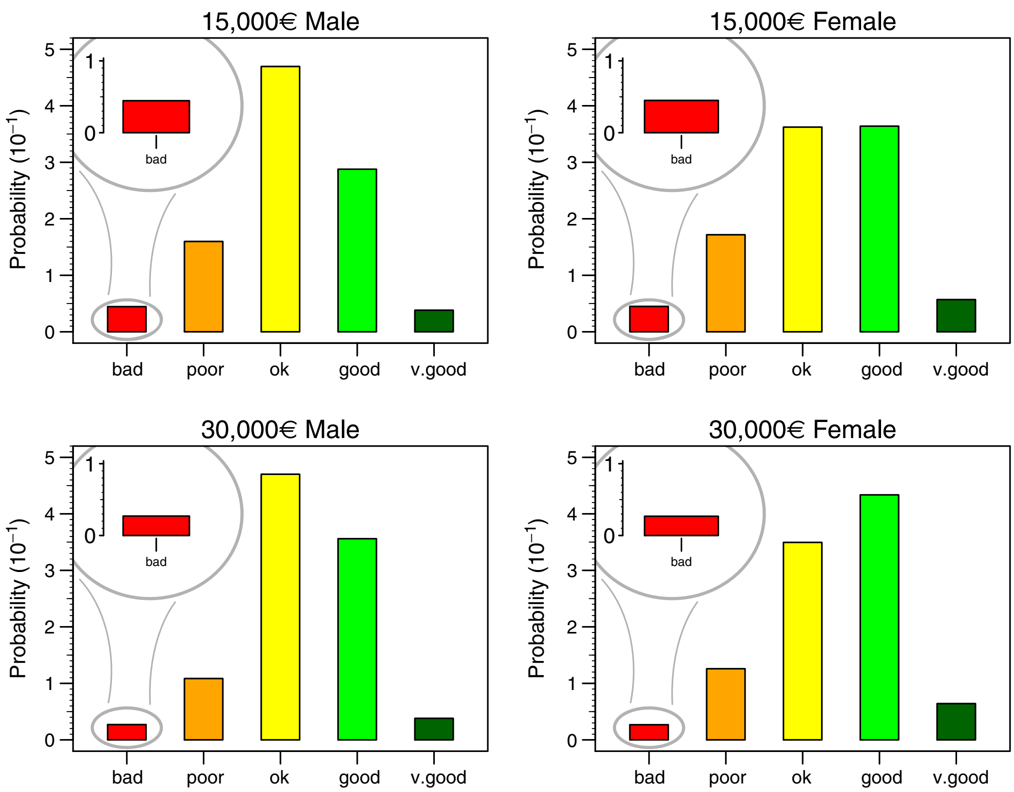

FIGURE 3 Conditional health distributions (self-rated health) for 15,000€ (top) and 30,000€ (bottom) for average Joe (left) and average Jane (right). With added focus on $\mathcal{R}_{M1}$ by magnification [Colour figure can be viewed at wileyonlinelibrary.com]

$\mathcal{R}_{M3}$ 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 large.

Second, 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, whereas 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 that focusses on those who self-report a truly bad health. Thus, it expresses the risk of "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 SRH variable. Not surprisingly, this is also true when we consider the quasicontinuous PCS health score.

To characterize the relationship between income and the risk of low health using the SF-12, 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 Figure 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

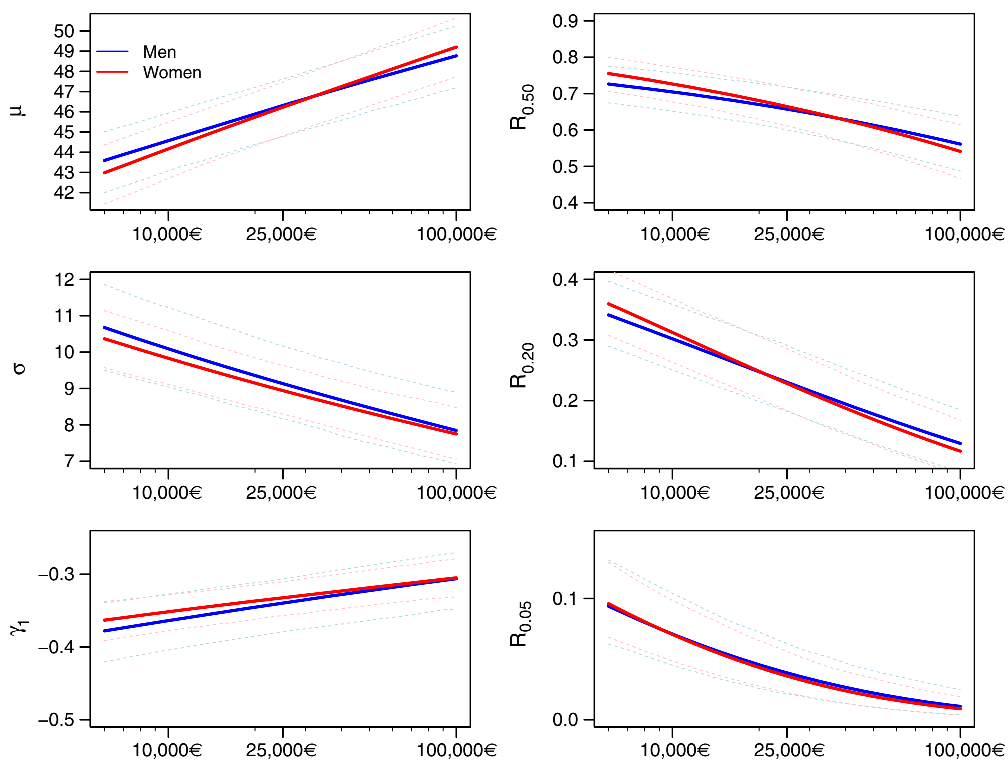

FIGURE 4 Left: effect of income on mean, standard deviation, and skewness of PCS. Right: effect of income on risk of falling below lowest quintile, decile, and vingtile of PCS. PCS = physical health score [Colour figure can be viewed at wileyonlinelibrary.com]

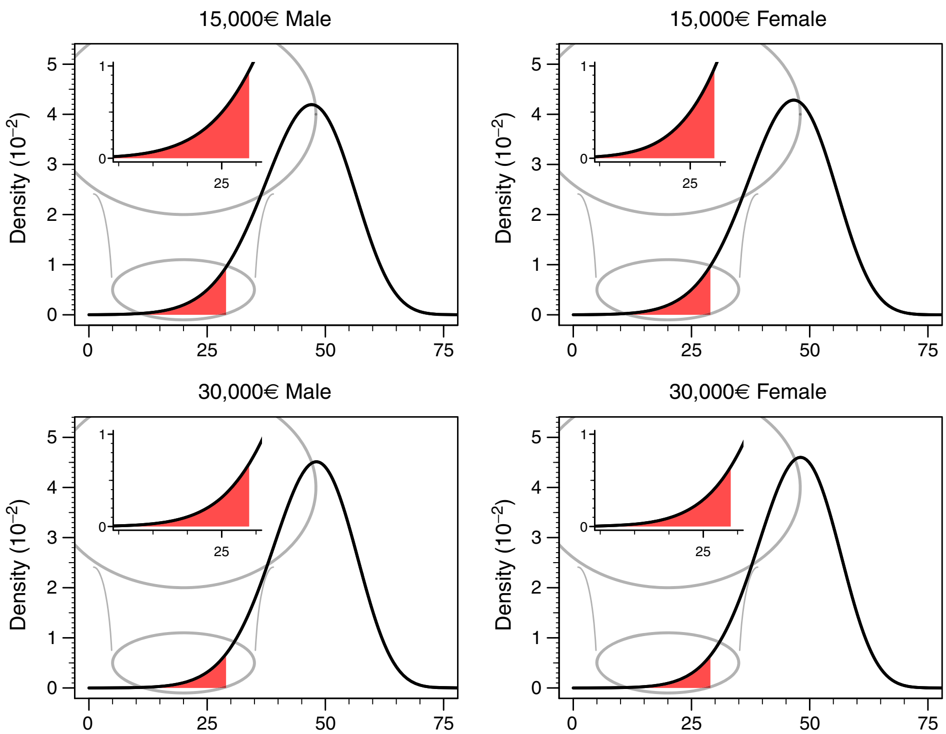

FIGURE 5 Conditional health distributions (physical health score) for average Joe (left) and average Jane (right) for 15,000€ (top) and 30,000€ (bottom). With added focus on $\mathcal{R}_{C0.05}$ by magnification [Colour figure can be viewed at wileyonlinelibrary.com]

aggregate health distribution, depending on their income.[11] 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 "ok." Lastly, $\mathcal{R}_{C0.05}$ denotes the risk of falling into the 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}. Note that in Figure 5, we display the integral yielding the corresponding risk measure only for $\mathcal{R}_{C0.05}$. The estimated full conditional distributions for the SF-12 for "average Joes" and "average Janes" 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. Although 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 PCS 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.

[11]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 SRH responses. More research is needed concerning the use of adequate scalar measures to assess this and other aspects of conditional health distributions.

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—is not only driven 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 is 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, that is, a fall 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). When considering $\mathcal{R}_{C0.20}$, the relative difference is 20% (0.27 to 0.22) and 23% (0.27 to 0.21) for men and women, respectively. The magnitude of the difference has thus already increased. 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 Figure 5, the risks and the differences thereof are highlighted by the much shaded area of the integral yielding the risk for $\mathcal{R}_{C0.05}$. 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 these effects is structurally different from the observed 3% increase observed for expected health.

3.5 | 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 4. 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%, whereas at the lower end of the health spectrum, the relative differences are in the order of 39–42%—that is, more than tenfold greater—when considering $\mathcal{R}_{M1}$ and $\mathcal{R}_{C0.05}$.

TABLE 4 Seven measures on the health-income association

Males			
	15,000€	**30,000€**	**Relative difference**
$\mathcal{R}_{M3}$	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%]

The conventional perspective generates significant results that allow us to infer the existence of a systematic 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 although average population health is an important construct for many purposes, we cannot properly calculate the utility of alternative distributions of health using only 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 that lessen the number of people in bad health would nevertheless be preferable at a societal level. At a policy level, this greater emphasis at the lower end would imply a greater need for resources dedicated to caring for those who are ill, rather than focussing on improving the health of the already healthy even further. However, to provide empirical evidence to guide such a policy, a distributional perspective that contemplates health risks beyond the conventional mean-based approach is needed. Such a perspective is provided by SADR. The results we have shown here demonstrate that although the relationship between income and health may be of modest magnitude if we focus only on average health, the relationship between income and illness (i.e., concentrating on poor and bad health states) is considerable larger.

4 | CONCLUSION

In this paper, we follow other recent publications that have pointed to the shortcomings of regression-based assessments of the income–health relationships that focus solely on the expected outcome. In order to look beyond the mean, we propose the use of SADR. These models allow for the estimation of full conditional health distributions for both multicategorical and continuous measures of health outcomes. Using health data from the German SOEP, 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 bad health is much more strongly related to income than is average health status. We find that the risk for the “average Joe” and “average Jane” of belonging to the severely sick population decreases 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 not only assess mean reported health (or the probability of a dichotomized health measure in an income group) but also employ risk measures focussing on bad health outcomes such as the ones used in this paper. Not only would this put more emphasis on the lower end of the spectrum, where we argue it is merited. In addition, it would address problems associated with nonlinearities with respect to well-being and/or health care and mean regression. 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 SRH 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 such as SRH and odds ratios) and the health economics literature (which tends to employ continuous measures such as 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 SADR, which uses copula structures to model the interrelations of the dependent variables (see Klein & Kneib, 2016) and would allow for the construction of conditional concentration curves across the covariate space. Although 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.

ACKNOWLEDGEMENTS

The authors would like to thank the editor and the two anonymous referees for their helpful comments and Hendryke Scholl for her suggestions. In addition, Alexander Silbersdorff and Thomas Kneib would like to thank the German Research Association DFG for their financial support.

ORCID

Alexander Silbersdorff http://orcid.org/0000-0002-3453-9536

REFERENCES

Backlund, E., Sorlie, P. D., & Johnson, N. J. (1996). The shape of the relationship between income and mortality in the United States. *Annals of Epidemiology*, *6*, 12–20.

Belitz, C., Brezger, A., Klein, N., Kneib, T., Lang, S., & Umlauf, N. (2015). BayesX—Software for Bayesian inference in structured additive regression models: Version 3.0. Version 3.0. URL http://www.bayesx.org.

Berk, M. L., & Monheit, A. C. (2001). The concentration of health care expenditures, revisited. *Health Affairs*, *20*, 9–18.

Bourion-Bédès, S., Schwan, R., Laprevote, V., Bédès, A., Bonnet, J.-L., & Baumann, C. (2015). Differential item functioning (DIF) of SF-12 and Q-LES-Q-SF items among french substance users. *Health and Quality of Life Outcomes*, *13*, 172.

Carrieri, V., & Jones, A. M. (2017). The income-health relationship 'beyond the mean': New evidence from biomarkers. *Health Economics*, *26*, 937–956.

Case, A. (2014). Does money protect health status? Evidence from South African pensions. In D. A. Wise (Ed.), *Perspectives on the economics of aging* (pp. 287–311). Chicago: University of Chicago Press.

Cesarini, D., Lindqvist, E., Östling, R., & Wallace, B. (2016). Wealth, health, and child development: Evidence from administrative data on Swedish lottery players. *Quarterly Journal of Economics*, *131*, 687–738.

Chernozhukov, V., Fernandez-Val, I., & Melly, B. (2013). Inference on counterfactual distributions. *Econometrica*, *81*, 2205–2268.

Culyer, A. J. (1990). Commodities, characteristics of commodities, characteristics of people, utilities, and the quality of life. In S. Baldwin, C. Godfrey & C. Propper (Eds.), *Quality of life* (pp. 9–27). London: Routledge.

Duclos, J.-Y., & Échevin, D. (2011). Health and income: A robust comparison of Canada and the US. *Journal of Health Economics*, *30*, 293–302.

Ecob, R., & Davey Smith, G. (1999). Income and health: What is the nature of the relationship?. *Social Science & Medicine*, *48*, 693–705.

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

Ettner, S. L. (1996). New evidence on the relationship between income and health. *Journal of Health Economics*, *15*, 67–85.

Fahrmeir, L., Kneib, T., & Lang, S. (2004). Penalized structured additive regression for space-time data: A Bayesian perspective. *Statistica Sinica*, *14*, 731–761.

Finkelstein, A., Luttmer, E. F. P., & Notowidigdo, M. J. (2009). Approaches to estimating the health state dependence of the utility function. *American Economic Review*, *99*, 116–121.

Firpo, S., Fortin, N. M., & Lemieux, T. (2009). Unconditional quantile regressions. *Econometrica*, *77*, 953–973.

Fleishman, J. A., & Lawrence, W. F. (2003). Demographic variation in SF-12 scores: True differences or differential item functioning?. *Medical Care*, *41*, III75–III86.

Fox, J. (1987). Effect displays for generalized linear models. *Sociological Methodology*, *17*, 347–361.

Frijters, P., Haisken-DeNew, J. P., & Shields, M. A. (2005). The causal effect of income on health: Evidence from German reunification. *Journal of Health Economics*, *24*(5), 997–1017.

Harper, S., King, N. B., Meersman, S. C., Reichman, M. E., Breen, N., & Lynch, J. (2010). Implicit value judgments in the measurement of health inequalities. *The Milbank Quarterly*, *88*, 4–29.

Hastie, T. J., & Tibshirani, R. J. (1990). *Generalized additive models*. New York: Chapman & Hall.

Heckley, G., Gerdtham, U.-G., & Kjellsson, G. (2016). A general method for decomposing the causes of socioeconomic inequality in health. *Journal of Health Economics*, *48*, 89–106.

Hothorn, T., Kneib, T., & Bühlmann, P. (2014). Conditional transformation models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, *76*, 3–27.

Kawachi, I., Adler, N. E., & Dow, W. H. (2010). Money, schooling, and health: Mechanisms and causal evidence. *Annals of the New York Academy of Sciences*, *1186*, 56–68.

Klein, N., & Kneib, T. (2016). Scale-dependent priors for variance parameters in structured additive distributional regression. *Bayesian Analysis*, *11*, 1071–1106.

Klein, N., Kneib, T., Lang, S., & Sohn, A. (2015). Bayesian structured additive distributional regression with an application to regional income inequality in Germany. *Annals of Applied Statistics*, *9*, 1024–1052.

Koenker, R. (2005). *Quantile regression*. Cambridge: Cambridge University Press.

Kuehnle, D. (2014). The causal effect of family income on child health in the UK. *Journal of Health Economics*, *36*, 137–150.

Lindahl, M. (2005). Estimating the effect of income on health and mortality using lottery prizes as an exogenous source of variation in income. *Journal of Human Resources*, *40*, 144–168.

Makdissi, P., & Yazbeck, M. (2014). Measuring socioeconomic health inequalities in presence of multiple categorical information. *Journal of Health Economics*, *34*, 84–95.

Marmot, M. (2002). The influence of income on health: Views of an epidemiologist. *Health Affairs*, *21*, 31–46.

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

Mechanic, D. (2002). Disadvantage, inequality, and social policy. *Health Affairs*, *21*, 48–59.

Nelder, J. A., & Wedderburn, R. W. M. (1972). Generalized linear models. *Journal of the Royal Statistical Society: Series A (General)*, *135*, 370–384.

Oliver, A., Healey, A., & Le Grand, J. (2002). Addressing health inequalities. *The Lancet*, *360*, 565–567.

Piketty, T. (2014). *Capital in the twenty-first century: The dynamics of inequality, wealth, and growth*. Cambridge: Harvard University Press.

Rahmann, U., & Schupp, J. (2013). *SOEP wave report 2012*. Berlin: DIW Berlin.

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

Sen, A. K. (1993). Capability and well-being. In M. C. Nussbaum (Ed.), *The quality of life* (pp. 30–53). London: Clarendon Press.

SOEP. (2014). Socio-Economic Panel study (SOEP), data of the years 1984–2013, version 30, https://doi.org/10.5684/soep.v30.

Sohn, A. (2017). The regression of an equitable market economy. (Ph.D. Thesis), Universität Göttingen, Göttingen.

Wagner, G. G., Frick, J. R., & Schupp, J. (2007). The German Socio-Economic Panel study (SOEP)—Scope evolution and enhancements. *Schmollers Jahrbuch*, *127*, 139–169.

Ziebarth, N. (2010). Measurement of health, health inequality, and reporting heterogeneity. *Social Science & Medicine*, *71*, 116–124.

SUPPORTING INFORMATION

Additional Supporting Information may be found online in the supporting information tab for this article.

How to cite this article: Silbersdorff A, Lynch J, Klasen S, Kneib T. Reconsidering the income-health relationship using distributional regression. *Health Economics*. 2018;27:1074–1088. https://doi.org/10.1002/hec.3656