



ELSEVIER

Available online at www.sciencedirect.com

SCIENCE @ DIRECT®

Journal of Health Economics 25 (2006) 131–145

JOURNAL OF
HEALTH
ECONOMICS

www.elsevier.com/locate/econbase

Reforming health care: Evidence from quantile regressions for counts

Rainer Winkelmann*

University of Zurich, Socioeconomic Institute, Ziirichbergstr. 14, CH-8032 Ziirich, Switzerland

Received 24 August 2004; received in revised form 23 March 2005; accepted 31 March 2005

Available online 22 June 2005

Abstract

I consider the problem of estimating the effect of a health care reform on the frequency of individual doctor visits when the reform effect is potentially different in different parts of the outcome distribution. Quantile regression is a powerful method for studying such heterogeneous treatment effects. Only recently has this method been extended to situations where the dependent variable is a (non-negative integer) count. An analysis of a 1997 health care reform in Germany shows that lower quantiles, such as the first quartile, fell by substantially larger amounts than what would have been predicted based on Poisson or negative binomial models.

© 2005 Elsevier B.V. All rights reserved.

JEL classification: I11; I18; C25

Keywords: Co-payments; Prescription drugs; Count data; Quantile regression; Poisson model

1. Introduction

Suppose survey data are available to estimate the effect of a health care reform, such as an increase in the out-of-pocket expense for prescription drugs, on an individual's utilization of health services. Typical empirical strategies include pre-reform/post-reform comparisons or differences-in-differences where one compares the changes in utilization between affected

* Tel.: +41 44 634 2292; fax: +41 44 634 4996.

E-mail address: winkelmann@sts.unizh.ch.

and unaffected sub-populations. Whether the “treatment effect” is assumed constant or heterogeneous, it is typically defined as the change in expected utilization that can be attributed to the reform.

Alternatively, one can explore a broader view, going beyond the first moments, by studying the effect of the reform on the whole outcome distribution. The novelty of this paper is to conduct such an analysis, where the reform effect is potentially different in different parts of the distribution of the outcome of interest, within the context of count data modelling. This becomes necessary since the outcome variable considered in this paper is a count—the number of doctor visits during the 3 months period prior to the interview. By comparing the distribution of visits with and without reform, we can for example determine whether the policy response is relatively larger among low users than among high users. In this case, the policy effect differs depending on the realization of the *dependent* variable.

The two benchmark count data models are the Poisson and negative binomial regression models with log-linear conditional expectation function. These models, and their two-part counterparts, have been used quite extensively in the analysis of health care utilization (Cameron and Trivedi, 1986; Deb and Trivedi, 2002; Gerdtham, 1997; Riphahn et al., 2003; Jimenez-Martin et al., 2000; Santos Silva and Windmeijer, 2001; Schellhorn, 2000; Windmeijer and Santos Silva, 1997; Winkelmann, 2004a). However, these previous applications have been concerned with estimating mean effects rather than full distributional responses. When it comes to estimating distributional responses, the standard models are of little use since the distributional response is determined entirely by functional form once the conditional mean response is known.

Given this problem, there are a couple of ways to proceed and analyse the data using more general models. The approach pursued in this paper is based on quantile regression methods for count data, applying a recently developed method by Machado and Santos Silva (in press). Basically, the approach transforms the discrete data problem into a continuous data problem by adding a random uniform variable to each count. The quantile regression functions of the transformed variable can then be estimated using standard quantile regression software. To interpret the results, one can compare the freely estimated quantile functions to those implied by the respective Poisson or negative binomial estimates in order to detect excess sensitivity in specific parts of the distribution, such as the lower or upper tails.

This methodology is applied to an evaluation of a German health care reform in 1997, using data from the *German Socio-Economic Panel*. The main result is that the reform effect was relatively more pronounced in the left part of the distribution: lower quantiles, such as the 25% quantile, fell by substantially larger amounts than what would have been predicted based on Poisson or negative binomial models. This finding has important policy implications. It is compatible with the notion that the demand for more frequent users of health services, among them the chronically sick, is relatively inelastic.

2. The German health care reform of 1997

The German health care sector is largely public. More than 90% of the German population is covered by a statutory health insurance system that is funded through mandatory payroll

deductions. Naturally, the sector is highly regulated. What services are offered, by whom and for what price, and how much the user pays are all questions subject to periodic review and adaptation. A change in 1997 dealt specifically with co-payments for prescription drugs.

Prescription drugs are dispensed by retail pharmacies which charge the insurance companies for the uniform price of the prescription, minus a co-payment that is required of the patient. The amount of the co-payment varies by package size. It increased substantially on 1 July 1997, by a fixed amount of DM 6 relative to a year earlier. Since the absolute amount of the co-payment is a function of the package size, after the reform DM 9 for small, DM 11 for medium and DM 13 for large sizes, the relative effect of the 1997 reform was largest for small sizes, where it amounted to a 200% increase.

How large was the effect of the increased co-payment on the demand for prescription drugs and other aspects of medical care utilization? In assessing the effects of the reform on the demand for health services, one can usefully distinguish between a direct and an indirect effect. The direct effect is a movement up the demand curve for prescription drugs, i.e., a reduced number of drug purchases after the reform, as the increased co-payment directly increased the patient's out-of-pocket expenses for drug purchases. The indirect effect is a potential inward shift of the demand curve for doctor visits. Since prescriptions have to be issued by physicians, the demand for doctor visits and the demand for prescription drugs are close complements and one can expect a negative cross-price elasticity.

Alternatively, one can think of the problem as demand for medical care in general. A potential patient may not know whether or not the doctor will issue a prescription. Raising the price of pharmaceuticals therefore raises the total expected price of a treatment for some condition, when the total price includes costs for the doctor visit plus costs for any pharmaceuticals.

The main idea of the paper is that the response to the reform may be different for different types of people. Specifically, I am interested in the hypothesis that the reform effect differs between frequent users and occasional users. The next section demonstrates that the existing econometric models are ill suited to address this specific question. The limitation of these models can be overcome by using quantile regression methods instead, as demonstrated in the following sections of the paper.

3. Marginal probability effects

The starting point of this paper is the recognition that regression models based on a single parameter Poisson distribution imply very restrictive probability changes in response to a change in a regressor. The log-linear Poisson regression model has probability function

$$f(y; \lambda) = \frac{\exp(-\lambda)\lambda^y}{y!} \quad (1)$$

where $\lambda = \exp(x'\beta)$.

Now consider the change in the probability $f(y; \lambda)$, $y = 0, 1, \dots$ induced by an infinitesimal change in the j th regressor x_j , keeping all other regressor constant. This is the marginal

probability effect which, for the Poisson model, can be written as

$$\frac{\partial f(y; \lambda)}{\partial x_j} = \frac{\partial f(y; \lambda)}{\partial \lambda} \frac{\partial \lambda}{\partial x_j} = f(y; \lambda)(y - \lambda)\beta_j \tag{2}$$

It follows that

$$\text{sgn} \frac{\partial f(y; \lambda)}{\partial x_j} = -\text{sgn}(\beta_j) \quad \text{iff } y < \lambda, \quad \text{sgn} \frac{\partial f(y; \lambda)}{\partial x_j} = \text{sgn}(\beta_j) \quad \text{iff } y > \lambda$$

Fig. 1 illustrates the situation. It is based on the conditional expectation function $\lambda = \exp(0.5 + 0.1x)$ and shows the probability changes as x increases from 1 to 2.

We notice the “single crossing” property of the probability changes. Based on the Poisson probability distribution, only a single switch between positive and negative marginal effects is possible as the counts increase from 0 to 1, from 1 to 2 and so forth. Also, the relative magnitudes of the effects are fully determined by functional form. We have to conclude that the Poisson model is not very well suited when the interest lies in modelling the full probability response to a change in a regressor. Note that this is *not* a problem of the particular conditional expectation function.

One could choose the most general parameterization of λ possible, such as a generalized additive model, any arbitrary link function, or even a fully saturated model, and the problem would remain the same. All these approaches will translate into a specific response $\partial \lambda / \partial x$, which in turn will induce the very restrictive probability changes of the Poisson distribution (1).

The situation would also not improve if one were to choose the negative binomial model rather than the Poisson model as the basis for analysis. In fact, the sign rule and the single crossing property remain exactly the same as in the Poisson case. For example, for the




Fig. 1. Example for marginal probability effects in Poisson regression model.

Negbin model with quadratic variance function (see e.g. Winkelmann, 2003) we obtain

$$\frac{\partial f_{\text{NB}}(y; \lambda, \xi)}{\partial x_j} = f_{\text{NB}}(y; \lambda, \xi) \left(\frac{\xi}{\xi + \lambda} \right) (y - \lambda) \beta_j \quad (3)$$

where

$$f_{\text{NB}}(y; \lambda, \xi) = \frac{\Gamma(\xi + y)}{\Gamma(\xi)\Gamma(y + 1)} \left(\frac{\xi}{\xi + \lambda} \right)^\xi \left(\frac{\lambda}{\xi + \lambda} \right)^y \quad (4)$$

ξ is a positive dispersion parameter, and λ is the conditional expectation as before. Again, the marginal probability effects have opposing signs to β_j for all realizations below the mean, and equal signs for all realizations above it.

To summarize, if one wants to model the probability response of a counted outcome more flexibly, and in particular allow for different responses in different parts of the distribution (relative to the benchmark Poisson or negative binomial models) one needs to turn to alternative modelling approaches. A first possibility is to abandon the rigid single index structure of the conventional approaches. The prime example is the hurdle model (Mullahy, 1986). In most applications, the hurdle is set at zero. In such models, the probability response of the zero outcome is entirely unrelated to the probability response in the strictly positive part of the distribution. Winkelmann (2004a) has applied such models in an evaluation of the effect of the aforementioned reform. He found that the response to the reform was significantly stronger in the left tail of the distribution (relative to the Poisson or negative binomial benchmarks) than elsewhere.

Here, I will analyse the same issue using a different approach that, rather than focussing on the probability function, concentrates on the dual problem of modelling the distribution function through quantile regression. This approach is developed in the next section.

4. Quantile regression for counts

The use of quantile regression for continuous random variables is by now quite standard. Since such regressions can be performed for arbitrary quantiles of a distribution, they provide a flexible tool for modelling the effect of regressors on the full distribution of the outcome variable.

In the context of count data, the main problem is that the distribution function of a discrete random variable is not continuous. Hence, the quantiles are not continuous either, and they cannot be modelled directly as a continuous function of the regressors. However, this difficulty can be overcome, as shown by Machado and Santos Silva (in press). Let y be the count variable. The α -quantile of y is defined by

$$Q_y(\alpha) = \min(\eta | P(y \leq \eta) \geq \alpha) \quad (5)$$

where $0 \leq \alpha < 1$. The object of interest is the conditional quantile $Q_y(\alpha|x)$. Since $Q_y(\alpha|x)$ has the same support as y , it is discrete and cannot be a continuous function of x (such as $\exp(x'\beta)$). Therefore, Machado and Santos Silva suggest to introduce “jittering”: consider

a new variable z , obtained by adding a uniform random variable to the count variable

$$z = y + u, \quad u \sim \text{uniform}[0, 1) \tag{6}$$

where y and u are independent. Hence, z has density function

$$f(z) = \begin{cases} p_0 & \text{for } 0 \leq z < 1 \\ p_1 & \text{for } 1 \leq z < 2 \\ \vdots & \vdots \end{cases} \tag{7}$$

(using notation $P(Y=k) = p_k$). Moreover, the distribution function of z can be written as

$$f(z) = \begin{cases} p_0 z & \text{for } 0 \leq z < 1 \\ p_0 + p_1(z - 1) & \text{for } 1 \leq z < 2 \\ \vdots & \vdots \end{cases} \tag{8}$$

The quantiles of z are continuous. For example,

$$Q_z(\alpha) = \frac{\alpha}{p_0} \quad \text{for } \alpha < p_0 \tag{9}$$

$$Q_z(\alpha) = 1 + \frac{\alpha - p_0}{p_1} \quad \text{for } p_0 \leq \alpha < p_0 + p_1 \tag{10}$$

From (9), we see that the z_α -quantiles can never be smaller than α . This needs to be taken into account in the econometric specification below. If the underlying count variable has either a Poisson or a negative binomial distribution, the z_α -quantiles can be easily computed and plotted as a function of the parameters. In the Poisson case, $Q_z(\alpha)$ depends on λ only whereas in the negative binomial case, it depends on λ and ξ . Fig. 2 displays the $z_{0.25}$ and the $z_{0.75}$ quantiles for the Poisson and negative binomial distributions as a function of λ (for $\xi = 1$). We find that the negative binomial distribution is more spread out than the Poisson distribution (overdispersion) and that the difference between the two distributions is an increasing function of the mean.

Of course, the main advantage of this approach is that the quantiles can now be estimated freely, without imposing any arbitrary and restrictive distributional form assumption. Following Machado and Santos Silva (in press), let

$$Q_z(\alpha|x) = \alpha + \exp(x' \gamma_\alpha), \quad \alpha \in (0, 1) \tag{11}$$

where α is added on the right side in order to impose the aforementioned lower bound of $Q_z(\alpha|x)$. Next, transform z such that the transformed quantile function is linear in the parameters:

$$Q_{T(z;\alpha)}(\alpha|x) = x' \gamma_\alpha$$

where

$$T(z; \alpha) = \begin{cases} \log(z - \alpha) & \text{for } z > \alpha \\ \log(\zeta) & \text{for } z \leq \alpha \end{cases} \tag{12}$$

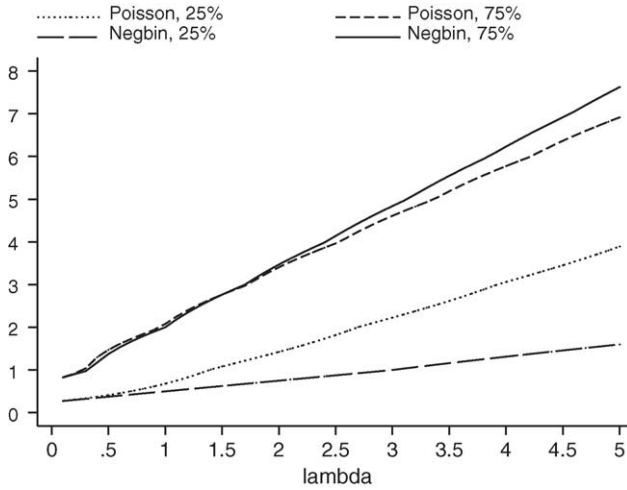


Fig. 2. Quantiles of Poisson and negative binomial distributions with continuity correction.

and $0 < \zeta < \alpha$. This can be done, since quantiles are invariant both to monotonic transformations and to censoring from below up to the quantile of interest. The censoring is required whenever $y = 0$ and the added uniform random variable u is smaller than α .

The model suggests the following empirical implementation. First, one adds uniformly distributed pseudo random numbers to the observed counts. Second, one transforms the resulting data. Third and finally, the parameter estimates are obtained as solution to

$$\min \sum_{i=1}^n \rho_{\alpha}(T(z_i; \alpha) - x'_i \gamma_{\alpha})$$

where $\rho_{\alpha}(v) = v(\alpha - I(v < 0))$.

Although the quantile function is not differentiable everywhere (the distribution function has corners), these corner points have measure zero as long as there is at least one continuous regressor. Machado and Santos Silva (in press) prove consistency and asymptotic normality of this estimator. Therefore, inferences about $Q_z(\alpha|x)$ can be based on conventional methods. For example, a Wald-test can be used to test the hypothesis that a regressor has no effect on a selected quantile.

Of course, the z_{α} -quantiles are only a means to an end, and we ultimately want to learn about the y_{α} -quantiles and the conditional distribution of the counts. The following considerations apply. First, the y_{α} -quantiles can be recovered from the z_{α} -quantiles based on the relation

$$Q_y(\alpha|x) = \begin{cases} \text{int}[Q_z(\alpha|x)] & \text{if } Q_z(\alpha|x) \text{ is not an integer} \\ Q_z(\alpha|x) - 1 & \text{if } Q_z(\alpha|x) \text{ is an integer} \end{cases} \tag{13}$$

For example, if $\alpha = p_0$ then, from Eqs. (4) and (11) we obtain that $Q_z(\alpha) = 1$ and $Q_y(\alpha) = 0$. However, this situation occurs with probability zero, and in all other cases, the y_α -quantiles are simply the integer part of the z_α -quantiles.

Second, therefore, if a variable is found to have no effect on $Q_z(\alpha|x)$, then we can conclude that it has also no effect on $Q_y(\alpha|x)$. If an effect for the z_α -quantile is observed, a given change in a regressor may or may not be sufficient to also change the y_α -quantile. This needs to be evaluated on a case by case basis.

Third, a direct analysis of the z_α -quantiles is informative, when one wants to study how the freely estimated quantiles differ from those implied by standard count data models. For example, in the context of evaluating the effect of the health reform on the distribution of the number of doctor visits per quarter, one can estimate a Negbin model and predict selected quantiles before and after the reform, separately for treatment and control group if available, where all other variables are held constant at their mean values. From these predictions, the relative response of the various quantiles can be computed and compared to those from the freely estimates quantile regressions. The comparison shows whether the reform had unusual effects in selected parts of the distribution, relative to what would have been predicted on the basis of the Negbin benchmark estimates.

Finally, there is the problem of how to choose the quantiles α . In theory, the number of quantiles that one could consider is unlimited. In practice, one needs to select a few quantiles of interest. An intelligent choice depends both on the type of question one wants to address, and on the data at hand. For example, if the marginal distribution of the data has 33% zeros so that the marginal quantiles are zero for all $\alpha < 0.33$, as will be the case in the present application, it makes little sense to compute conditional quantile functions for values as low as $\alpha = 0.1$. The variation in the conditional z_α is then mostly due to the random noise that has been added, and the quantiles will be flat and not vary as a function of the regressors. In the following application, we therefore look at four quantiles, $\alpha = 0.25, 0.50, 0.75$, and 0.90 .

5. Data and empirical strategy

The analysis is based on data from the *German Socio-Economic Panel* (SOEP). The SOEP was initiated in 1984 (SOEP Group, 2001). It is an annual survey that is ongoing. For the purpose of this study, I selected data for men and women from the so-called Sample A, which includes persons with non-guestworker status from the territory of former West Germany, for the 2 years centered around the year of the reform, i.e., 1996 and 1998. Moreover, I distinguish between treatment and control group. The control group includes all individuals with private insurance (mostly the self-employed plus workers with earnings above a certain threshold) plus those covered by statutory health insurance but explicitly exempt from co-payments (the youth and low income families, identified here by the receipt of welfare payments in the current year). Deleting observations with missing values on any of the dependent or independent variables, the sample comprises 18,683 observations.

The full descriptive statistics for the dependent and independent variables that are used in the analysis are given in Table 1. As the table shows, the majority of observations (90%) are for the treatment group. The dependent variable is the utilization of health services, as

measured by the individual number of visits to a doctor during the 3 months prior to the interview. Thus, for the treatment group, we observe the utilization in a low co-payment regime in 1996, whereas the 1998 survey provides information on utilization in the high co-payment regime. The average number of visits dropped by about 10%, from 2.96 to 2.66 visits per quarter. The control group has fewer visits overall, and the decline between pre- and post-reform period is much smaller (by 2% from 2.51 to 2.46). Table 1 also shows for the selected quantiles (the 25, 50, 75, and 90% quantiles) that there is some movement for the treatment group, but no change for the control group. Specifically, the median of the number of visits in the treatment group drops from 2 to 1, and the 90% quantile from 7 to 6.

As far as the explanatory variables are concerned, we find that an individual from the treatment group is on average about 7 years older than an individual from the control group. The health status from a subjective self-assessment is somewhat worse in the treatment group, and the employment and marital rates are higher than in the control group. There are some minor differences in the pre- and post-reform periods, but they tend to be insignificant.

The basic empirical strategy is to implement a differences-in-differences estimator by pooling the data over the two observation points and groups. We know from an earlier study (Winkelmann, 2004b) that the change in the expected demand for doctor visits, conditional on covariates, before and after the reform was practically zero for the control group, vindicating the use of simple pre–post comparisons. However, it is unclear whether such a result also holds for the conditional quantile functions considered here. Therefore, we retain the full differences-in-differences specification. In standard notation, define a linear

Table 1
Descriptive statistics ($N = 18,683$)

	Treatment group		Control group	
	Pre-reform	Post-reform	Pre-reform	Post-reform
Doctor consultations	2.96 (4.74)	2.66 (4.17)	2.51 (4.06)	2.46 (4.30)
0.25 quantile	0	0	0	0
0.50 quantile	2	1	1	1
0.75 quantile	3	3	3	3
0.90 quantile	7	6	6	6
Age	43.7 (16.4)	45.0 (16.5)	36.1 (17.0)	38.7 (16.8)
Male	0.48 (0.50)	0.48 (0.50)	0.52 (0.50)	0.59 (0.49)
Years of schooling	11.1 (2.4)	11.3 (2.4)	11.1 (3.3)	11.6 (3.0)
Married	0.66 (0.47)	0.66 (0.47)	0.42 (0.49)	0.47 (0.50)
Household size	2.97 (1.38)	2.88 (1.33)	3.32 (1.49)	3.19 (1.47)
Active sport	0.23 (0.42)	0.28 (0.45)	0.32 (0.47)	0.33 (0.47)
Good health	0.51 (0.50)	0.53 (0.50)	0.62 (0.49)	0.62 (0.49)
Bad health	0.17 (0.38)	0.16 (0.37)	0.12 (0.33)	0.13 (0.34)
Logarithmic income	7.52 (0.44)	7.54 (0.42)	7.50 (0.61)	7.49 (0.61)
Full-time employed	0.46 (0.50)	0.44 (0.50)	0.36 (0.48)	0.44 (0.50)
Part-time employed	0.09 (0.29)	0.09 (0.29)	0.04 (0.20)	0.04 (0.19)
Unemployed	0.06 (0.24)	0.06 (0.24)	0.11 (0.31)	0.10 (0.30)
Observations	8130	8666	1002	885

Notes. Standard deviations in parentheses.

index variable η_{it} as follows:

$$\eta_{it} = \beta_0 + \beta_1 \text{treat}_i + \beta_2 \text{post}_t + \beta_3 \text{treat}_i \times \text{post}_t + z'_{it}\gamma, \quad t = 96, 98 \quad (14)$$

Here, post_t is an indicator for the post-reform period, treat_i is an indicator whether person i is subject to the increased co-payment, and the interaction between treat_i and post_t marks an observation of a treated person after the reform, z_{it} includes a constant and all other characteristics controlled for in the regression, among the, a second order polynomial in age, three indicators for the quarter of the interview, three indicators of employment status (*full-time*, *part-time*, *unemployed*) plus the variables *years of education*, *married*, *logarithmic income*, *household size*, *active sport*, *good health*, *bad health*. In this set-up, β_3 measures the ceteris paribus change in the index between the two period in the treatment group above and beyond the change in the control group, for otherwise similar individuals.

As laid out in the previous section, we are interested in a comparison of two types of approaches to estimate the effect of the 1997 health care reform on the demand for health services. In a first approach, we estimate standard count data models, the Poisson and the Negbin model, with $E(y_{it}|\eta_{it}) = \lambda_{it} = \exp(\eta_{it})$. In a second approach, we freely estimate selected conditional quantiles of the continuity corrected counts by using the specification $Q_z(\alpha|\eta_{it}) = \alpha + \exp(\eta_{it})$ for $\alpha = 0.25, 0.50, 0.75, 0.90$.

The interpretation of differences-in-differences estimators in such non-linear models is not entirely straightforward. In count data models with log-linear conditional expectation function, $[\exp(\beta_2) - 1] \times 100$ gives the 96–98 ceteris paribus percentage change in the expected number of doctor visits for the control group. $[\exp(\beta_2 + \beta_3) - 1] \times 100$ gives the corresponding change for the treatment group. If β_3 is negative, the demand for doctor visits fell in the treatment group relative to the control group after the imposition of the increased co-payments. Absolute changes depend on the value of the covariates at which the effect is evaluated, as do the effects on the various quantiles of the standard count data models.

In the quantile regression models, the interpretation is more complicated because the additive constant α means that regression parameters do not give proportional effects. Let $\tilde{\eta}_i$ denote the value of the index function for given covariates \tilde{z}_i before the reform for the control group. The double difference for the α -quantile is then

$$\text{dd}_i = [\exp(\tilde{\eta}_i + \beta_1 + \beta_2 + \beta_3) - \exp(\tilde{\eta}_i + \beta_1)] - [\exp(\tilde{\eta}_i + \beta_2) - \exp(\tilde{\eta}_i)]$$

An estimate of the conditional reform effect $\widehat{\text{dd}}_i$ is obtained by replacing the parameters by their estimates. Finally, to obtain the marginal reform effect, we can estimate the unconditional expected value by averaging over the individual effects:

$$\widehat{\text{dd}} = \frac{1}{n} \sum_{i=1}^n \widehat{\text{dd}}_i$$

These estimates are directly comparable to the marginal differences-in-differences quantile effects from the Negbin model. In this way, we can assess whether or not the distributional impact of the reform differed in any systematic way from what would have been expected from using a standard Negbin model.

6. Results

The Poisson and Negbin estimates are displayed in Table 2. In addition to the differences-in-differences (dd) results, the table also shows the results from the simple pre–post-reform comparisons using a reduced sample of observations on the treatment group only. A comparison of the Poisson and negative binomial models confirms the superiority of the latter, as expected. The estimated dispersion parameter in the Negbin model is about 1, with standard error 0.02—it is zero under the Poisson restriction. Wald and likelihood ratio tests lead to a clear rejection of the Poisson model. The estimated effects of the control variables are very similar across the specifications and they agree with those reported in the previous literature (see e.g. Cameron and Trivedi, 1986). For example, the number of visits is inverse u-shaped in age. Men have fewer doctor visits than women. Labor market participants have fewer visits than non-participants. The number of visits is higher for married people and those with a poor self-assessed health status, and lower for those in good health and with lower income, *ceteris paribus*.

The point estimate of the treatment effect, a 10% reduction in the number of visits, is the same regardless of whether the Poisson or Negbin model is used and regardless of whether single or double differences are taken. The only change is that in the dd model, the effect for the control group (a relatively small and heterogeneous group) is measured with low precision, which translates into an imprecise estimate of the dd estimator. Nevertheless, we retain the dd formulation for the quantile regressions, as we cannot exclude that the

Table 2
Results for Poisson and negative binomial regressions

	Simple differences ($N = 16,796$)		Differences-in-differences ($N = 18,683$)	
	Poisson	Negbin	Poisson	Negbin
Age/10	0.035 (0.044)	-0.030 (0.044)	0.039 (0.041)	-0.028 (0.042)
Age squared/1000	-0.036 (0.043)	0.032 (0.044)	-0.040 (0.041)	0.031 (0.042)
Male	-0.142 [†] (0.028)	-0.199 [†] (0.026)	-0.159 [†] (0.027)	-0.226 [†] (0.025)
Married	0.090 [†] (0.032)	0.112 [†] (0.030)	0.083 [†] (0.030)	0.104 [†] (0.028)
Active sport	0.079 [†] (0.027)	0.099 [†] (0.026)	0.090 [†] (0.026)	0.115 [†] (0.025)
Good health	-0.628 [†] (0.027)	-0.633 [†] (0.027)	-0.628 [†] (0.026)	-0.636 [†] (0.026)
Bad health	0.780 [†] (0.029)	0.794 [†] (0.030)	0.785 [†] (0.028)	0.797 [†] (0.029)
Logarithmic income	0.124 [†] (0.029)	0.135 [†] (0.028)	0.095 [†] (0.027)	0.096 [†] (0.027)
Full-time employed	-0.278 [†] (0.034)	-0.281 [†] (0.032)	-0.278 [†] (0.033)	-0.276 [†] (0.031)
Part-time employed	-0.251 [†] (0.045)	-0.253 [†] (0.043)	-0.258 [†] (0.043)	-0.258 [†] (0.041)
Unemployed	-0.142 [†] (0.052)	-0.143 [†] (0.049)	-0.132 [†] (0.048)	-0.125 [†] (0.045)
Post-reform	-0.101 [†] (0.020)	-0.100 [†] (0.020)	-0.005 (0.065)	-0.002 (0.065)
Treatment			0.030 (0.049)	0.050 (0.049)
Post × treatment			-0.096 (0.069)	-0.098 (0.069)
α		0.992 (0.023)		1.012 (0.022)
Log-likelihood	-45659.5	-34364.5	-50583.9	-38024.2

Notes. Source: German Socio-Economic Panel, years 1996 and 1998. Dependent variable: number of doctor visits during previous quarter. Models include furthermore a constant, three indicator variable for the quarter of the interview (winter, spring, fall) and the variables *Years of schooling* and *Household size*. Robust standard errors adjusted for heteroscedasticity and for clustering observations over 2 years in parentheses. Coefficients marked with [†] are significant at the 10% level.

simple differences estimator may not work well for some quantiles. Therefore, the Negbin estimates in the fourth column of Table 2 will be the benchmark, on which the simulation and comparison of quantiles are based.

Table 3 shows the results for the quantile regressions. The regressors that are significant in the Negbin model also tend to be significant in the quantile regressions. It is possible that the signs of the effects differ at the different quantiles and this happens indeed in two instances: The effect of age at the 90% quantile is u-shaped rather than inverse u-shaped (albeit insignificant), and years of schooling has a positive effect at the two lower quantiles and a negative effect at the two upper quantiles. In all other cases, the signs do not switch. For example, all four estimated z_α -quantiles are lower for men than for women. While the point estimates decrease with increasing quantiles, a direct comparison is misleading. Instead, we can use the estimates to predict quantiles for men and women, *ceteris paribus*. For example, if all other variables are set to their sample means, the 90 % z_α -quantile is 5.39 for men and 6.44 for women. By contrast, the 25 % z_α -quantile is 0.73 for men and 1.01 for women. Thus, while the relative effect of being male – a change of -28% – is larger at the 0.25 quantile, the absolute effect – a change of -1.05 – is larger at the 0.90 quantile.

The last row of Table 3 shows the parameters associated with the treatment effect. They are negative for all quantiles, but statistically significant only for the 75 and the 25% quantiles. The relative low precision induced by the small and heterogeneous control group was mentioned before. How should the point estimates be interpreted? First, we use the point estimates to predict the y_α -quantiles for each person in the treatment group, with and without reform, using Eq. (13). In this simulation, all explanatory variables are set to their actual values, except for the *post-reform* and *post* \times *treatment* indicators that are either zero or one for all individuals. Table 4 shows the relative frequencies of the thus obtained estimated quantiles for the 16,796 observations in the treatment group. The change in the distribution

Table 3
Count data quantile regressions

	$Q_z(0.25 x)$	$Q_z(0.50 x)$	$Q_z(0.75 x)$	$Q_z(0.90 x)$
Age/10	-0.171^\dagger (0.080)	-0.169^\dagger (0.060)	-0.092^\dagger (0.037)	0.026 (0.051)
Age squared/1000	0.206^\dagger (0.083)	0.182^\dagger (0.062)	0.097^\dagger (0.038)	-0.029 (0.052)
Male	-0.427^\dagger (0.043)	-0.414^\dagger (0.033)	-0.280^\dagger (0.021)	-0.192^\dagger (0.029)
Years of schooling/10	0.261^\dagger (0.087)	0.134^\dagger (0.066)	-0.041 (0.041)	-0.115^\dagger (0.055)
Married	0.245^\dagger (0.050)	0.229^\dagger (0.038)	0.103^\dagger (0.024)	0.098^\dagger (0.033)
Household size	-0.072^\dagger (0.016)	-0.074^\dagger (0.013)	-0.052^\dagger (0.008)	-0.041^\dagger (0.011)
Active sport	0.342^\dagger (0.047)	0.217^\dagger (0.036)	0.119^\dagger (0.023)	0.089^\dagger (0.030)
Good health	-0.853^\dagger (0.047)	-0.838^\dagger (0.036)	-0.640^\dagger (0.023)	-0.614^\dagger (0.030)
Bad health	1.035^\dagger (0.061)	0.795^\dagger (0.047)	0.818^\dagger (0.029)	0.725^\dagger (0.040)
Logarithmic income	0.139^\dagger (0.051)	0.147^\dagger (0.038)	0.095^\dagger (0.024)	0.120^\dagger (0.031)
Full-time employed	-0.377^\dagger (0.055)	-0.373^\dagger (0.042)	-0.260^\dagger (0.027)	-0.265^\dagger (0.036)
Part-time employed	-0.316^\dagger (0.079)	-0.228^\dagger (0.060)	-0.210^\dagger (0.038)	-0.337^\dagger (0.051)
Unemployed	-0.323^\dagger (0.087)	-0.194^\dagger (0.066)	-0.094^\dagger (0.042)	-0.004 (0.056)
Post-reform	0.221^\dagger (0.092)	0.109 (0.070)	0.102^\dagger (0.044)	0.008 (0.059)
Treatment	0.068 (0.125)	-0.016 (0.095)	0.061 (0.060)	0.006 (0.080)
Post \times treatment	-0.231^\dagger (0.132)	-0.127 (0.101)	-0.145^\dagger (0.064)	-0.081 (0.085)

Notes. See Table 2.

Table 4

Relative frequencies of estimated y_{α} -quantiles for treatment group before and after reform

	0	1	2	3	4	5	6	7	8	>9
Before reform										
$\hat{Q}_y(0.25 x)$	67.54	21.50	10.70	0.26	0	0	0	0	0	0
$\hat{Q}_y(0.50 x)$	0.89	57.85	19.77	7.12	4.38	5.03	4.38	0.55	0.02	0
$\hat{Q}_y(0.75 x)$	0.89	36.23	21.62	10.81	8.96	5.07	2.05	2.00	2.39	12.36
$\hat{Q}_y(0.90 x)$	0.17	14.30	23.04	14.45	9.50	6.05	6.95	4.13	3.96	25.55
After reform										
$\hat{Q}_y(0.25 x)$	72.77	18.95	8.26	0.02	0	0	0	0	0	0
$\hat{Q}_y(0.50 x)$	3.62	59.39	19.29	4.88	5.80	4.94	2.04	0.05	0	0
$\hat{Q}_y(0.75 x)$	3.62	41.07	18.32	11.79	7.50	2.62	2.26	2.52	3.28	10.31
$\hat{Q}_y(0.90 x)$	0.50	22.11	22.34	12.74	7.57	7.50	5.33	4.32	1.69	21.92

– the simple difference – is one indication of the effect of the reform on the distribution of demand for health services. For example, we see that the proportion of individuals with a predicted 0.25-quantile of zero increased from 68 to 73%. Similarly, the relative frequencies at which the quantiles fall in the higher counts, such as 9 or above, decreased, in the case for the 0.90-quantile from 26 to 22%.

Of course, a more comprehensive analysis of the results has to take the distributional effects for the control group into account. Moreover, we want to compare the quantile effects of Table 3 with those implied by the Negbin estimates in Table 2. Table 5 provides the relevant information. The first two columns display the absolute and relative effects for the quantile regression, respectively. The effects are computed as $dd(\alpha) = [\hat{Q}_z(\alpha|treat, post) - \hat{Q}_z(\alpha|treat, pre)] - [\hat{Q}_z(\alpha|control, post) - \hat{Q}_z(\alpha|control, pre)]$ where $\hat{Q}_z(\alpha|\cdot)$ is the average quantile over all individuals with appropriately recoded treatment and reform indicator variables. For example, $\hat{Q}_z(0.25|treat, post) = 1.04$, $\hat{Q}_z(0.25|treat, pre) = 1.18$, $\hat{Q}_z(0.25|control, post) = 1.05$, and $\hat{Q}_z(0.25|control, pre) = 1.00$. From this, we obtain a double difference of -0.19 , or 16.3% of 1.18. Importantly, we see, from Table 5, that the estimated relative reform effect is largest for the 0.25 quantile and smallest for the 0.9 quantile.

Table 5

Predicted average treatment effects (differences-in-differences) at various z_{α} -quantiles for quantile regressions and Negbin model

Quantile	Quantile regression		Negbin model	
	Absolute	Relative	Absolute	Relative
25%	-0.194	-16.34%	-0.075	-7.40%
50%	-0.230	-9.51%	-0.193	-8.06%
75%	-0.507	-11.45%	-0.383	-8.06%
90%	-0.487	-6.87%	-0.633	-8.05%

Notes. The effects are computed as $dd(\alpha) = [\hat{Q}_z(\alpha|treat, post) - \hat{Q}_z(\alpha|treat, pre)] - [\hat{Q}_z(\alpha|control, post) - \hat{Q}_z(\alpha|control, pre)]$, where $\hat{Q}_z(\alpha|\cdot)$ is the average predicted quantile over all individuals.

This evidence can now be contrasted with the results from the Negbin model in columns 3 and 4 of Table 5. For each individual, a linear predictor was computed based on the four possible combinations of treatment and pre-/post-reform status. This was converted to the corresponding quantiles, using the Negbin probability function (4) and the definition of the distribution function in (8), and then averaged. The results show that the effect in the Negbin model is approximately proportional. The double difference of the quantiles amount to a 7–8% decrease in the number of doctor visits due to the reform.

The results are thus unequivocal. A comparison of the Negbin quantiles and the freely estimated quantiles reveals the following patterns: in the Negbin model, the relative reform effect does not depend on α ; the relative reform effect implicit in the freely estimated quantiles, on the other hand, is a decreasing function of α : the largest effect is recorded for the smallest quantile, here the 25% quantile. Hence, the quantile regression result show what an analysis based on conventional count data models would definitely miss, namely that the sensitivity of the demand for health services to the reform of 1997 was disproportionately high in the left part of the distribution. The drop in demand for health services at the 25% quantile that is attributable to the reform amounts to more than 16%, whereas the 90% quantile, representing individuals who are relatively frequent users, decreased by only 6.9%. In other words, the demand for more frequent users of health services, among them the chronically sick, reacted relatively inelastically.

7. Discussion

The German health care reform of 1997 was associated with an average decline in the number of doctor visits by 10%. In this paper it was shown that a sole focus on averages misses an important part of the story. A more detailed analysis using a novel quantile regression method for count data confirmed that the reform effect was quite heterogeneous indeed, defined here as being different in different parts of the distribution of the outcome of interest. Rare users responded more to the increased co-payment than frequent users, in relative terms. This finding corroborates an earlier analysis based on generalized parametric count data models, such as hurdle Poisson, hurdle Negbin and finite mixture models (Winkelmann, 2004a). However, the present approach based on quantiles provides a more robust tool for detecting departures from the benchmark models. It has the great advantage that it does not require the estimation of an alternative parametric, and possibly misspecified, generalized count data model.

An interesting substantive result of this paper is that it helps reconciling the present findings from the *German Socio-Economic Panel* with those reported in an earlier study of the same reform by Lauterbach et al. (2000), using a different survey. In that earlier study, the estimated reduction in the average number of doctor visits was just 4.4%, falling short of the 10% found here. The likely explanation for this discrepancy is that Lauterbach et al. based their analysis on a survey of pharmacy customers. Clearly, this approach produces a heavily selected sample in which frequent users of health services are overrepresented. Hence, a relatively small response is to be expected. However, as has been demonstrated in this paper, such a pharmacy-based survey is inappropriate to predict the effect of the reform on rare visitors and, by implication, on the population at large.

Acknowledgements

A great thank you goes to Joao Santos Silva for extended discussions and for showing me how to implement his estimator in Stata. Furthermore, I am grateful to two anonymous referees and to seminar participants at the Zurich Colloquium in Applied Statistics, the University of Konstanz and at the SFB-386 Workshop on *Statistical Modelling of Discrete Structures in Economics* in Munich for valuable comments.

References

- Cameron, A.C., Trivedi, P.K., 1986. Econometric models based on count data: comparisons and applications of some estimators and tests. *Journal of Applied Econometrics* 1, 29–53.
- Deb, P., Trivedi, P.K., 2002. The structure of demand for health care: latent class versus two-part models. *Journal of Health Economics* 21, 601–625.
- Gerdtham, U., 1997. Equity in health care utilization: further tests based on hurdle models and Swedish micro data. *Health Economics* 6, 303–319.
- Jimenez-Martin, S., Labeaga, J.M., Martinez-Granado, M., 2000. Latent class versus two-part models in the demand for physician services across the European union. *Health Economics* 11, 301–321.
- Lauterbach, K.W., Gandjour, A., Schnell, G., 2000. Zuzahlungen bei Arzneimitteln., <http://www.medizin.uni-koeln.de/kai/igmg/stellungnahme/zuzahlungen/> (25 February 2000).
- Machado, J.A.F., Santos Silva, J.M.C., in press. Quantiles for counts. *Journal of the American Statistical Association*.
- Mullahy, J., 1986. Specification and testing in some modified count data models. *Journal of Econometrics* 33, 341–365.
- Riphahn, R.T., Wambach, A., Million, A., 2003. Incentive effects in the demand for health care: a bivariate panel count data estimation. *Journal of Applied Econometrics* 18, 387–405.
- Santos Silva, J.M.C., Windmeijer, F., 2001. Two-part multiple spell models for health care demand. *Journal of Econometrics* 104, 67–89.
- Schellhorn, M., 2000. Health services utilization of elderly Swiss: evidence from panel data. *Health Economics* 9, 533–545.
- SOEP Group, 2001. The German Socio-Economic Panel (SOEP) after more than 15 years—overview. *Vierteljahrshefte zur Wirtschaftsforschung* 70, 7–14.
- Windmeijer, F., Santos Silva, J.M.C., 1997. Endogeneity in count data models: an application to demand for health care. *Journal of Applied Econometrics* 12, 281–294.
- Winkelmann, R., 2003. *Econometric Analysis of Count Data*, 4th ed. Springer, Heidelberg.
- Winkelmann, R., 2004a. Health care reform and the number of doctor visits—an econometric analysis. *Journal of Applied Econometrics* 19, 455–472.
- Winkelmann, R., 2004b. Co-payments for prescription drugs and the demand for doctor visits—evidence from a natural experiment. *Health Economics* 13, 1081–1089.