


Diskussionspapiere des
Europäischen Instituts für Sozioökonomie e. V.

Working Papers of the
European Institute for Socioeconomics



**Public goods, private consumption,
and human-capital formation:
On the economics of volunteer labour supply**

Eike Emrich & Christian Pierdzioch

Nr. 14

Jahr 2015

Prof. Dr. Eike Emrich

Fachbereich: Sportökonomie und Sportsoziologie

Adresse: Universität des Saarlandes
Sportwissenschaftliches Institut
Campus
66123 Saarbrücken

Tel: 049(0)681 302 4170

E-Mail: e.emrich@mx.uni-saarland.de

Prof. Dr. Christian Pierdzioch

Fachbereich: Makroökonomie und Internationale Wirtschaftsbeziehungen

Adresse: Helmut-Schmidt-Universität/Universität der Bundeswehr
Holstenhofweg 85
22043 Hamburg

Tel: 049(0)40 6541 2879

E-Mail: c.pierdzioch@hsu-hh.de

© 2015

Europäisches Institut für Sozioökonomie e. V. / European Institute for Socioeconomics

c/o Universität des Saarlandes

Sportwissenschaftliches Institut
Arbeitsbereich Sportökonomie und Sportsoziologie
Campus, Gebäude 8.1

66123 Saarbrücken

<http://www.soziooekonomie.org>

Public goods, private consumption, and human-capital formation: On the economics of volunteer labour supply

Eike Emrich^a and Christian Pierdzioch ^{b*}

July 2015

Abstract

Economists use three types of models to describe volunteer labour supply: the public-goods model, the private-consumption model, and the human-capital model. We used data from an online survey questionnaire of volunteers working for the German Red Cross to study the extent to which each type helps to explain volunteer labour supply. To this end, we empirically studied various correlates of volunteer labour supply, including the components of utility that agents receive from volunteering. We used boosted regression trees to trace out the main correlates of volunteer labour supply, to study the relative contributions of the utility components, and potential interaction effects between the utility components and other correlates.

JEL classification: H41, J22, L31

Keywords: Volunteer labour supply; Boosted regression trees; German Red Cross

Addresses:

^a Department of Sports Science, Saarland University, Campus Bildung B8.2 | 66123 Saarbrücken, Germany.

^bDepartment of Economics, Helmut Schmidt University, Holstenhofweg 85, P.O.B. 700822, 22008 Hamburg, Germany

* Corresponding author. E-mail: c.pierdzioch@hsu-hamburg.de.

1 Introduction

Three broad types of models form the foundation of most economic research on volunteer labour supply: the public-goods model (Roberts 1984, Bergstrom et al. 1986, Duncan 1999), the private-consumption model (Andreoni 1989, 1990, Harbough 1998), and the human-capital model (Menchink and Weisbrod 1987). In a prototype public-goods model, a volunteer contributes to and derives utility from the total provision of a public good. Because non-excludability and non-rivalry in consumption are by definition constituent characteristics of any public good, the model implies that a volunteer has altruistic preferences in the sense that a volunteer derives utility when the utility of others increases. A typical private-consumption model, in contrast, implies that a volunteer has egoistic preferences insofar as it is the act of volunteering that enters into the utility function. A volunteer may derive direct utility from the act of volunteering because volunteering gives rise to a “warm-glow” feeling, helps to spend leisure time in a worthwhile manner, and makes it possible to come together with other individuals. Finally, human-capital models focus on the aspect that, in an intertemporal setting, a volunteer benefits from doing volunteer work because volunteering helps to accumulate job-market skills, on-the-job experience, and social capital.

We used the three types of models to study how volunteer labour supply is linked to the utility derived by volunteers, controlling for several socioeconomic variables like age, sex, religiosity, and others. We also studied how the utility derived by volunteers from volunteering interacts with these socioeconomic variables in shaping volunteer labour supply. We used in our empirical analysis a micro dataset from an online questionnaire study of volunteers working for the German Red Cross (GRC). Emrich and Pierdzioch (2015a) analyse this data in recent research to trace out how volunteer labour supply is correlated with the motives for doing volunteer work. In our empirical research, in contrast, we focused on the correlation of volunteer labour supply with the components of the utility that volunteers derive from their volunteer work. To this end, we collected data on various utility components that represent the public-goods model, the private-consumption model, and the human-capital model. We then studied how these individual utility components correlate with volunteer labor supply, controlling for the influence of several other potentially important socioeconomic variables. In proceeding in this way, we had to handle a large number of predictive variables, and this is one reason why we analysed the data by means of an advanced statistical learning algorithm that renders it possible to explore in a straightforward and economically easy-to-interpret way the correlation of a large number of utility dimensions with volunteer labour supply.

The statistical learning algorithm that we applied in our research is known as boosted re-

gression trees (Friedman 2001, 2002). Regression trees recursively subdivide the range of explanatory variables into non-overlapping regions and set the reaction of volunteer labour supply (the response variable) in every region to its region-specific mean to minimize a loss function, which in our case is a standard quadratic loss function. Aggregation over regions then gives the predicted value of volunteer labour supply (on regression trees, see Breiman et al. 1983). As compared to a standard regression model, regression trees have several advantages (Breiman et al. 1983, Hastie et al. 2009, James et al. 2013). Regression trees give easy-to-interpret results that can be visualized in terms of decision trees. Regression trees capture potential interaction effects between explanatory variables, they can handle various types of explanatory variables, including qualitative predictors, and they are robust to outliers in the predictors and to the inclusion of irrelevant predictors. A drawback of regression trees is that their data sensitivity implies that they are high-variance predictors. Boosted regression trees overcome this drawback by combining predictions from a large number of trees (Friedman et al. 2000). Special techniques like partial dependence plots and specific measures of relative variable influence are available that make the economic interpretation of empirical results obtained by means of boosted regression trees straightforward. We applied these special techniques to study the potentially nonlinear and complex correlation of volunteer labour supply with the various forms of utility derived by volunteers from volunteering and our socioeconomic control variables. We also studied the relative importance of the utility components and the socioeconomic control variables for explaining volunteer labour supply, and the potential interaction effects between our predictors.

We proceed as follows. In Section 2, we use a simple theoretical model to explain our research strategy. In Section 3, we describe how boosted regression trees are computed. In Section 4, we describe our data. In Section 5, we summarize our empirical results. In Section 6, we conclude.

2 Theoretical Foundation

In order to explain our research strategy, it is useful to lay out a simple reductionist microeconomic model that embeds, in a very stylized way, the three types of economic models of volunteer labour supply. The model features a volunteer who maximizes utility, u , subject to a standard budget constraint:

$$\max u = \log c + \beta \log(g + \gamma G^o) \quad \text{s.t.} \quad c + g = wT + e, \quad (1)$$

where c = a consumption good, g = a public good, w = wage rate, T = hours worked, e = other sources of income, and G^o = contributions of other volunteers to the production of a public good (assumed by a volunteer to be exogenously given).

If we set $\gamma = 1$, only the total supply of the public good, $g + G^o$, enters into the utility function and we obtain a pure public-goods (PG) model. In the public-goods model, a volunteer accepts that the utility of others increases when his or her volunteer labour supply helps to increase the total supply of a public good. It is in this sense that the public-goods model accounts for a volunteer's altruism. If a volunteer's preferences would be different such that the positive externality of volunteering were completely unacceptable (absence of altruism) then such a volunteer would not want to volunteer at all (complete free riding). At the same time, the composition of the total supply of a public good does not matter in the public-goods model and, hence, a PG volunteer reduces his or her effort when others increase their labour supply (free riding in the presence of altruism).

If we set $\gamma = 0$, only the private contribution enters into the utility function and we obtain a pure private-consumption (PC) model. Intermediate values of γ then result in an impure altruism model of the type studied by Andreoni (1989, 1990). A human-capital (HC) model, in turn, results if we set $\gamma = 0$ and, in addition, assume that the parameter that captures the relative weight of the second term in the utility function depends on the contributions of other volunteers: $\beta = f(G^o, \mathbf{q})$ with $f > 0$ and $\partial f / \partial G^o > 0$ and \mathbf{q} = a vector of other socioeconomic control variables. While we cannot really integrate the kind of intertemporal optimisation underlying any model of human-capital formation into a static model, the assumption captures the idea that the formation of job-market skills and social networks should be easier in a larger community of volunteers. At the same time, a volunteer's effort may increase in the presence of other volunteers because signaling social skills and social competence is harder if volunteering is a social norm. Hence, if the human-capital model is at work volunteer labour supply should increase in the size of the organisational GRC unit for which a volunteer works. For a discussion of the complementarity of volunteers' labour supply in the investment model and a theoretical analysis, see Ziemek (2006, page 536). For a model of competition among volunteers, see Cugno and Ferrero (2004).

The maximisation problem given in Equation (1) is solved by a volunteer-labour-supply function of the format

$$g^* = \frac{(wT + e)\beta}{1 + \beta} - \frac{\gamma}{1 + \beta} G^o. \quad (2)$$

Optimal volunteer labour supply, g^* , should be high if a volunteer derives "much" utility from the second term in the utility function (a large weighting factor, β), irrespective of

the specific type of economic model of volunteer labour supply under consideration. Hence, volunteer labour supply should increase *ceteris paribus* for a volunteer who derives much utility due to PG, PC, or HC considerations. In our empirical analysis, we used the results of a survey questionnaire among volunteers of the GRC to find out how much of the cross-sectional variation of volunteer labour supply correlates with various utility components that are associated with the PG, PC, and HC models.

In addition, in case of the human-capital model the weighting parameter, by assumption, increases if a volunteer deems that human-capital and social networking are an important facet of volunteering ($\partial f / \partial G^0 > 0$). At the same time, however, the term γG^0 enters with a negative sign into Equation (2), giving rise to the crowding-out hypothesis that has been extensively studied in empirical and experimental research (Payne 1998, Eckel et al. 2005, Ziemek, 2006, Crumpler and Grossman 2008, Andreoni et al. 2014, Emrich and Pierdzioch 2015a). Hence, the labour supply, G^0 , of other volunteers has two effects: a negative effect due to the cross-sectional substitutability of volunteer labour supply in case of the PG model, and a positive effect due to the formation of human capital in larger GRC communities. The contributions of others, γG^0 , may also have a negative effect if a volunteer feels that his or her contribution no longer “makes a difference” if others’ contributions increase, as in Duncan’s (2004) model of impact philanthropy.

3 Boosted Regression Trees

Various boosting techniques have been developed and extensively studied in the statistical learning literature. For recent surveys, see Bühlmann and Hothorn (2007) and Mayr et al. (2014). Applications of boosting techniques in economics, however, are relatively scarce. Economists mainly have used boosting techniques to study macroeconomic fluctuations (Bai and Ng 2009, Buchen and Wohlrabe 2014) and to model the returns and the volatility of financial market prices (Audrino and Trojani 2007, Berge 2014, among others). The list of applications of boosted regression trees in economics includes the research by Rossi and Timmermann (2010) and Mittnik et al. (2015), who study the determinants of stock returns and financial market volatility. For applications of boosted regression trees in other disciplines like ecology and medicine, see Elith et al. (2008) and Friedman and Meulman (2003), who also discuss issues relevant for the practical application of boosted regression trees. A recent application of regression trees in economics can be found in the research by Bajari et al. (2015). For a non-technical introduction with a focus on economic applications of regression trees, see Varian (2014). For an introduction, see James et al. (2013, Chapter 8).

Adopting the notation used by Hastie et al. (2009, Chapter 9), a regression tree, T , with J terminal nodes partitions the space of the explanatory variables, $\mathbf{x} = (x_1, x_2, \dots)$, into l non-overlapping regions, R_l , where a greedy algorithm is used to determine partitions in a top-down and binary way. At the top level of a tree, the first partition is chosen such that the partitioning predictor variable, s , and the partitioning point, p , define the half-planes $R_1(s, p) = \{x_s | x_s \leq p\}$ and $R_2(s, p) = \{x_s | x_s > p\}$ that minimize the loss function

$$\min_{s,p} \left\{ \min_{\bar{y}_1} \sum_{x_s \in R_1(s,p)} (y_i - \bar{y}_1)^2 + \min_{\bar{y}_2} \sum_{x_s \in R_2(s,p)} (y_i - \bar{y}_2)^2 \right\}, \quad (3)$$

where y_i denotes observation i of the response variable, and the inner minimization is done by choosing constants, $\bar{y}_k, k = 1, 2$ to minimize the region-specific squared error loss. For a quadratic loss function, the constants are given by $\bar{y}_k = \text{mean}\{y_i | x_s \in R_k(s, p)\}$. The outer minimization then involves a simple search over all s and p that gives the first optimal partitioning predictor, the first optimal partitioning point, and the corresponding two region-specific means of the response variable. This search gives a slightly more complex new tree that now has two terminal nodes. For this new tree, the minimization given in Equation (3) is repeated separately for the optimal two half-planes, $R_1(s, p)$ and $R_2(s, p)$, identified at the top level. This gives up to two additional optimal partitioning predictors, two additional optimal partitioning points, and four additional region-specific means of the response variable. Upon continuing this partitioning process until a preset maximal tree size is reached or, for example, every terminal node has a minimum number of observations, we get a final full-fledged hierarchical regression tree. The final regression tree allocates the predictors into the R_l optimal regions of the tree and then predicts the i -th observation of the response variable, y_i , by its constant region-specific mean. Upon using $\mathbf{1}$ to denote the indicator function, this can be formally expressed by writing

$$T(\mathbf{x}_i, \{R_l\}_1^L) = \sum_{l=1}^L \bar{y}_l \mathbf{1}(\mathbf{x}_i \in R_l). \quad (4)$$

While simple to estimate and easy to interpret, its hierarchical structure implies that a regression tree can be sensitive to small changes in the data. If such a small change in the data alters one of the top-level splits this affects all splits lower in the tree hierarchy. Several approaches have been advocated in the statistical learning literature to mitigate the data sensitivity of regression trees. One such approach is bootstrap aggregation (bagging, Breiman 1996), which requires drawing a large number of bootstrap samples from the data and fitting a regression tree to every sample. The response variable is then predicted by averaging over the bootstrapped regression trees.

Growing random forests is another approach (Breimann 2001). Similar to the bootstrap approach, growing a random forest requires sampling from the data to build a large number of simulated datasets. On every simulated dataset, a random regression tree is estimated. A random regression tree differs from a standard tree in that for every partition only a random subset of the explanatory variables is used. Using only a random subset of the explanatory variables for tree growing reduces the influence of individual explanatory variable and, thus, reduces the variance of the predictions of the response variable.

While random forecasts consist of a large number of regression trees estimated on simulated data, a boosting algorithm involves growing a large number of trees on the same original data. Various variants of boosting algorithms have been proposed in the statistical learning literature. The boosting algorithm that we used in our empirical research is known as gradient least-squares boosting (Friedman 2001, 2002). Gradient least-squares boosting aims at minimizing a standard loss function that is quadratic in the prediction error, $L = (y - F(\mathbf{x}))^2$, where $F(\mathbf{x})$ is an unknown and potentially complex function that maps the predictors into predictions of the response variable. The boosting algorithm aims at approximating this function, $F(\mathbf{x})$, using a sequence of regression trees, $T_m, i = 0, 1, 2, \dots, M$. In the boosting literature, the elements of this sequence are known as base learners. The first base learner in this sequence, T_0 , initializes the boosting algorithm and, once the first base learner has been fixed, the other base learners are determined recursively by the boosting algorithm.

Adopting a notation that closely follows the one used by Friedman (2001), the specific algorithm underlying gradient least-squares boosting can be described as follows. The algorithm is initialized by setting $T_0 = \bar{y}$, which is the prediction of a zero-node tree. In the next step, T_0 is used to predict the response variable and to compute a vector of prediction errors with elements $r_{i,0} = y_i - T_0$. The vector of prediction errors gives the negative gradient of the loss function, $-\frac{\partial L}{\partial F(\mathbf{x})}|_{F(\mathbf{x})=F_0(\mathbf{x})}$. The negative gradient defines the direction in which to search in the next step to minimize the loss function. Hence, the predictors are used to fit the second base learner, T_1 , to the negative gradient vector using Equation (3). The new base learner, T_1 , is of the format given in Equation (4). In the next step, the second base learner, T_1 , is used to construct a new function approximation given by $F_1(\mathbf{x}) = F_0(\mathbf{x}) + \sum_{l=1}^{L,1} \bar{y}_{l,1} \mathbf{1}(\mathbf{x}_i \in R_{l,1})$. This new function approximation defines a new gradient vector, $-\frac{\partial L}{\partial F(\mathbf{x})}|_{F(\mathbf{x})=F_1(\mathbf{x})}$, which is then used to fit T_2 , and so on. Hence, the function, $F(\mathbf{x})$, is a “strong learner” that is successively approximated by means of a growing number of recursively estimated simple base learners. The final approximation of the function, $F(\mathbf{x})$, is then given by an ensemble of base learners:

$$F(\mathbf{x}) = \sum_{m=0}^M T_m \left(\mathbf{x}, \{R_{l,m}\}_1^{L,m} \right). \quad (5)$$

In order to prevent the algorithm from overfitting, two parameters must be set. The first parameter is the total number of base learners, M . Common techniques to determine M are splitting the dataset into a training set and a prediction set, and cross validation. The second parameter is the learning rate $0 < \zeta \leq 1$. Introducing the learning rate entails a slight modification of the gradient boosting algorithm (Friedman 2001) and gives the recursion

$$F_m(\mathbf{x}) = F_{m-1}(\mathbf{x}) + \zeta T_m \left(\mathbf{x}, \{R_{l,m}\}_1^{L,m} \right), \quad m = 1, \dots, M. \quad (6)$$

Hence, the learning rate curbs the contribution of every added single base learner to the strong learner. The learning rate is inversely linked to the optimal total number of base learners, M . A lower learning rate implies that more base learners should be used to approximate $F(\mathbf{x})$. In the boosting literature, it is common to set the learning rate to some small value, $\zeta < 0.1$, where the total number of trees typically is set to approximately $M = 1,000$ or larger.


Friedman (2002) suggests to introduce a bootstrapping element into gradient least-squares boosting. The resulting stochastic gradient boosting algorithm requires to sample without replacement, before fitting a base learner, a subset from the data. Only the sampled elements of the gradient vector and the sampled observations of the predictors are then used to estimate the next base learner, T_m , in the recursion given in Equation (6).

4 The Data

After some pre-sample testing based on interviews with 32 volunteers had been done, an online questionnaire study was conducted from April to May 2013. The link to the online questionnaire study was distributed among GRC volunteers by means of a top-down snowball approach. Filling in the online questionnaire took approximately 20–30 minutes of time, where the volunteers could participate anonymously. In total, $N = 4,611$ volunteers participated in the survey. Data on volunteer labour supply, as measured in hours per week, are available for $N = 4,199$ volunteers.

Figure 1 plots the distribution of volunteer labour supply. Apparently, many volunteers work a few hours per week, while a few volunteers answered that they work 40 hours per week or more. For this reason, we studied log transformed volunteer labour supply in our empirical research (see Section 5). Moreover, we deleted data (i) if a volunteer answered that he or she does not volunteer at least one hour per week ($N = 15$), and, (ii) if a volunteer answered that he or she would work more than 80 hours per week ($N = 7$), leaving $N = 4,177$ data for the

Figure 1: Volunteer labour Supply ($N = 4,177$)



empirical analysis. The mean of volunteer labour supply is 9.65 hours per week, the median is 7 hours per week, and the standard deviation is 9.45 hours per week.

Table 1 shows the utility components for which we collected data along with a short description, summary statistics, and the model type to which the utility components belong. Volunteers could rank every utility component on a 5-point scale from “do not agree” to “totally agree” (“When you think about the volunteering in the GRC, has volunteering helped you to...”). The summary statistics show that volunteers do not attach the same level of importance to all utility components. For example, the utility components “to bring about changes in politics” and “to improve job-market prospects” received on average low ranks. In contrast, the utility components “to meet other members of the GRC” and “to help others” received on average high ranks, where the first utility component belongs to the human-capital/private-consumption model type while the second utility component belongs to the public-goods model type.

Table 1: Utility Components

Utility component	N	Mean	Median	Std. Dev.	Model type	Description
Utility 1	4,056	4.53	5	0.76	HC/PC	to meet other members of the GRC
Utility 2	4,040	3.43	4	1.28	HC/PC	to meet individuals not members of the GRC
Utility 3	4,013	3.40	3	1.06	HC	to come to other's attention within the GRC
Utility 4	4,021	3.03	3	1.15	HC	to come to other's attention in society
Utility 5	4,025	3.34	4	1.35	HC	to acquire job-market skills
Utility 6	4,058	4.40	5	0.79	PC	to work together with other individuals
Utility 7	3,998	1.62	1	0.94	PG	to bring about changes in politics
Utility 8	4,029	3.38	3	1.13	PG	to contribute to small-scale developments in society
Utility 9	4,021	3.24	3	1.16	HC/PC	to improve one's standing within the GRC
Utility 10	4,032	2.97	3	1.18	HC/PC	to improve one's standing in other parts of society
Utility 11	4,000	2.51	2	1.37	HC	to improve job-market prospects
Utility 12	4,051	4.30	4	0.84	PC	to spend leisure time in a worthwhile manner
Utility 13	4,057	4.52	5	0.73	PG	to help others
Utility 14	4,053	4.41	5	0.76	PC	to have fun
Utility 15	4,019	3.17	3	1.17	PC	to defend one's interests
Utility 16	4,037	3.20	3	1.23	PG	to strengthen the GRC

Note: PG = public-goods model. PC = private-consumption model. HC = human-capital model. All utility components measured using a 5-point scale.

While the allocation to economic models is quite clear for most utility components, there is some ambiguity with regard to the four components “to meet other members of the GRC”, “to meet individuals not members of the GRC”, “to improve one’s standing with the GRC”, and “to improve one’s standing in other parts of society”. These four utility components may express that volunteering is a form of conspicuous consumption (Veblen 2011[1899], for a model of conspicuous giving, see Glazer and Konrad 1996), in which case the utility components belong to the private-consumption model. However, meeting other volunteers and people who are not members of the GRC, and improving one’s standing in the GRC and in society, may also help to build social networks and to strengthen “weak ties” (Granovetter 1973). In case a volunteer’s networking activities mainly aim at improving job-market prospects, the four components belong to the human-capital model. Accordingly, Emrich and Pierdzioch (2015a) treat four motives for doing volunteer work that correspond to the four utility components studied here in their empirical analysis as indicators of the human-capital model (but reallocating motives across models leaves results qualitatively unchanged).

The summary statistics given in Table 1 describe the utility components, but do not inform about the strength of their correlation with volunteer labour supply, nor do the summary statistics recover how exactly the utility components are linked to volunteer labour supply and how the utility components interact with other socioeconomic control variables. Table 2 depicts summary statistics of the socioeconomic control variables that we used in our empirical analysis. Our socioeconomic control variables capture the mechanics of the theoretical model laid out in Section 2. The model highlights that it is important to control for the influence of a volunteer’s time constraints, T , and income and financial well-being, w and e , the size of the GRC unit for which a volunteer works, G^0 , and other factors summarized in the vector \mathbf{q} . For a survey of the large and significant literature on the correlates of volunteering, see Musick and Wilson (2008). The following discussion of the correlations between our socioeconomic control variables and volunteer labour supply in part builds on arguments they lay out and survey in their book.

Table 2: Socioeconomic Control Variables

Control variable	N	Mean	Median	Std. Dev.	Description
Hours worked	3,233	39.65	40.00	12.08	hours worked on an income-generating job (in hours per week)
Friends	4,165	2.84	3.00	1.07	proportion of friends working for the GRC (5-point scale from low to high)
Religiosity	2,924	2.73	3.00	1.03	religiosity (5-point scale from unimportant to very important)
Interest in politics	4,144	3.85	4.00	0.86	interest in politics (5-point scale from uninterested to very interested)
Sex	4,134	—	—	—	sex (male = 0, female = 1, 35% female)
Age	4,087	38.14	36.00	14.43	age (in years)
Children	4,154	—	—	—	volunteer has a child / has children (no = 0, yes = 1, 42% have children)
Financial situation	4,095	3.37	3.00	0.92	perceived financial situation (5-point scale from very poor to very good)
Household size	3,139	3.18	3.00	1.25	number of household members (integer 1,2,...)
Income	3,779	2.72	3.00	1.26	income in euros measured in six broad categories (<1,000 €–2,250 €,...)
Size of unit	2,472	4.95	4.38	1.77	size of the organisational GRC unit (in logs)
Proportion of volunteers	3,794	55.46	60.00	34.66	proportion of volunteers in the GRC unit for which a volunteer works (in percent)
Membership in clubs	4,122	0.64	0.00	0.95	number of other club memberships
low-threshold activities	4,177	1.37	1.00	1.39	other activities (health sector, labour union, sport, etc.; in total 15 sectors)

The control variable hours worked accounts for the fact that volunteering competes with the time spent working on an income-generating job. While working on a job may reduce volunteer labour supply due to time constraints, the correlation with volunteer labour supply may also be positive if having a job improves social integration. We also asked volunteers whether they are members of other clubs, and whether they allocate time for low-threshold activities in other sectors (health sector, sports, etc.). Other leisure-time activities reduce the time that a volunteer can spend on volunteering but, as with working on a job, spending time with friends doing sport can enhance social integration and, thereby, lead to an increase in volunteer labour supply. Social integration most likely also depends on the proportion of friends in the GRC, and so we also collected data on this control variable.

Having many friends in the GRC does not necessarily imply that “to have fun” and “to meet other members of the GRC” receive a high ranking. However, the utility components and the proportion of friends may interact in shaping volunteer labour supply, and boosted regression trees capture such interactions. As Equation (2) shows, the utility components, on the one hand, shed light on which proportion of the cross-sectional variation of volunteer labour supply can be attributed to differences in preferences (a direct effect) and, on the other hand, the interactions (indirect effects) highlight how the utility components shape the response of volunteer labour supply to, for example, time constraints.

In order to capture whether a volunteer can afford “working for nothing” we asked volunteers to describe their overall financial situation. If a volunteer perceives his or her financial situation as rather tight, the labour supplied to the GRC may be low. At the same time, if a volunteer perceives his or her financial situation as precarious or unsatisfactory volunteer labour supply may increase simply because, as assumed by the human-capital model, volunteering can be a means to acquire job-market skills. To further measure the financial situation of a volunteer, we asked volunteers to provide data on their income, where volunteers could choose between six broad categories (less than 1,000 €, between 1,000 € and 2,250 €, ..., 5,500 € or more).

We also collected data on the size of the organisational GRC unit for which a volunteer works, G^0 . In this respect, it may also matter how many members of an organisational GRC unit do volunteer work. If the public-goods model is at work, a higher proportion of other volunteers is likely to reduce volunteer labour supply. In the human-capital model, in contrast, volunteer labour supply is increasing in the size of the community of volunteers.

With regard to the variables stacked in the vector \mathbf{q} , we collected data on religiosity, interest in political issues, sex, and age. The control variables religiosity and interest in political issues may indicate humanitarian values and a general norm of social responsibility, resulting

in a positive association with volunteer labour supply. At the same time, a volunteer who is religious and interested in politics may prefer to spend his or her leisure time doing volunteer work for a religious or a political organisation rather than the GRC. Hence, the overall correlation between religiosity and interest in political issue, on the one hand, and volunteer labour supply, on the other hand, can be positive or negative.

As for the control variable sex, it is well-known from empirical work (Andreoni et al. 2003, Emrich and Pierdzioch 2015b) and experimental research (Eckel and Grossman 1998, Andreoni and Vesterlund 2001) that there are gender differences with regard to charitable giving. Such gender differences may carry over to volunteer labour supply. In our sample, there are 35% female volunteers. Gender differences with regard to volunteer labour supply may depend on whether a volunteer has a child or children, and so we asked volunteers whether they have a child or children. In this respect, the total number of individuals forming a household may also matter. For example, the time spent on child care and, thus, a potentially important time constraint for doing volunteer work may depend on whether volunteers live with their parents.

Finally, we controlled for a volunteer's age. Age may have a first-order effect on volunteer labour supply because social roles and health conditions change with age, especially for the elderly. From a purely economic perspective, age may correlate with volunteer labour supply due to life-cycle considerations. For example, Glaeser et al. (2002) use a life-cycle model to argue that individuals will accumulate social capital in their early years to increase market returns (higher wages and better employment prospects) and non-market returns (for example, improving relationships and happiness) while they will let social capital depreciate in their later years. Glaeser et al. (2002) use the number of organisation memberships as a measure of the stock of an individual's social capital and find, in line with their life-cycle model, an inverted u-shaped empirical nonlinear link of memberships with age. Hence, the correlation with age and volunteer labour supply may turn out to be highly nonlinear. Boosted regression trees capture such nonlinearities in a straightforward way.

5 Empirical Analysis

We used the R programming environment for statistical computing (R Core Team 2015) to carry out our empirical analysis. For estimation of boosted regression trees, we used the add-on package “gbm” (Ridgeway 2015).

The “gbm” package makes use of the fact that regression trees allow missing data on the predictor variables to be handled in a natural way by means of surrogate tree splits. Our data feature several missing data on the predictor variables because, as witnessed by the varying number of observations reported in Tables 1 and 2, not all predictors are available for all volunteers who participated in the online questionnaire. Surrogate splits use predictors which trigger splits highly correlated with the splits triggered by the variable for which data are missing to allocate data down a tree (for the case of a classifier model, see Breimann et al. 1983, page 142).

Before estimation of boosted regression trees, a few parameters have to be fixed. Two important parameters, the learning rate and the total number of base learners, were already introduced in Section 3. Setting the learning rate to a small value requires that the ensemble of base learners that optimises the overall performance of the boosted regression-tree model gets larger. Overall model performance can be measured in terms of the expected squared error loss that results when a boosted model is applied to a test set:

$$E(y - F(\mathbf{x}))^2 = \text{Var}(F(\mathbf{x})) + [\text{Bias}(F(\mathbf{x}))]^2 + \text{Var}(e), \quad (7)$$


where $\text{Var}(F(\mathbf{x}))$ = variance of the estimated boosted predictor that results from variation in the data, $\text{Bias}(F(\mathbf{x}))$ = bias of the estimated boosted predictor that results from approximating the true functional dependence of y on \mathbf{x} by means of the boosted function $F(\mathbf{x})$, and $\text{Var}(e)$ = variance of the remaining statistical error (see, for example, James et al. 2013, page 34–36).

Adding more base learners reduces bias because the model more closely tracks even a complex functional dependence of volunteer labour supply on its predictors. At the same time, the variance of the estimated boosted predictor increases because the resulting complex model closely (and perhaps too closely) tracks the training set. Estimating only a few base learners, in contrast, leads to a large bias, while the danger of overfitting the data is low and, hence, the variance of the predictions of the resulting simple model when applied to a test set is relatively low. The resulting trade-off between variance and bias implies that it is possible to detect an optimal number, M^* , of base learners.

The optimal number of base learners increases if the numerical value assumed for the learning rate gets smaller. As Equation (6) witnesses, a small learning rate limits the influence of every single base learner on the boosted estimate of the function $F(\mathbf{x})$, and, thereby, curbs the dependence of the boosted regression-tree model on the training set. Hence, variance decreases. At the same time, however, bias increases because, for a given number of base

learners, the approximation of the true functional form gets worse. As a result, it is possible to reduce bias by adding additional base learners to the model. In doing so, however, variance increases again, so that we can identify a new, and now larger, optimal number, M^* , of base learners.

Figure 2: Optimal Number of Base Learners



Note: Learning rate = 0.01. The green line denotes the squared error loss for the test set obtained by means of five-fold cross-validation. The black line denotes the squared error loss for the training set. The dashed vertical line denotes the optimal number of base learners. Both panels were computed for the same random seed.

Figure 2 shows results for the optimal number of trees, where we used five-fold cross validation to define training and test sets. The black line denotes the squared error loss for the training set. The squared error loss for the training data decreases in the number of base learners. The green line, which represents the squared error loss for the test set, first rapidly decreases as the complexity of the boosted regression-tree model increases. Once the model features a substantial number of base learners, however, the green line stabilises and eventually reaches a minimum, represented by the dashed vertical line. For a learning rate of $\zeta = 0.01$, the minimum is reached when the model includes approximately 1,100 base learners. Further experimentation showed that, when we set the learning rate to $\zeta = 0.005$, it takes roughly 2,000 base learners to reach a minimum. Because of faster computations and because results do not change much, we set $\zeta = 0.01$ in our empirical analysis.

Table 3: Relative Influence of Predictors (in %)

Predictor	Mean	SD
Friends	19.77	1.15
Size of unit	11.88	0.34
Age	7.58	0.25
Utility 13	5.60	0.29
Utility 9	5.28	0.28
Working hours	4.86	0.32
Utility 16	4.07	0.19
Financial situation	4.06	0.15
Sex	3.91	0.20
Proportion of volunteers	3.58	0.34
Religiosity	2.99	0.14
Utility 2	2.15	0.13
Income	1.80	0.18
Utility 10	1.75	0.15
Utility 1	1.67	0.14
Utility 7	1.66	0.11
Utility 3	1.58	0.14
Utility 15	1.57	0.14
Utility 12	1.56	0.17
Household size	1.45	0.19
Utility 5	1.44	0.17
Utility 14	1.39	0.15
Interest in politics	1.30	0.11
Utility 6	1.25	0.17
Children	1.21	0.15
Utility 4	1.19	0.16
Utility 8	1.06	0.17
Low-threshold activities	0.83	0.12
Utility 11	0.80	0.14
Membership in clubs	0.76	0.14

Note: Relative influence is defined as the improvement in squared error resulting from using a predictor to form splits, averaged across base learners. Mean = mean relative influence. SD = standard deviation of relative influence. Results are based on 250 model simulation runs (50% sampling rate, learning rate 0.01, 5-fold cross-validation, mean number of optimal base learners = 949.33, SD of the optimal number of base learners = 129.62).

Another parameter that has to be fixed before a boosted regression-tree model can be estimated is tree size, that is, the maximum number of terminal nodes per base learner. Setting tree size to one gives stumps and so we would get an additive model that neglects potential interactions between the predictors. In order to capture potential interactions, we set the tree size to 5. Increasing tree size to, say, 10 leaves our results qualitatively unchanged. Furthermore, we set the minimum number of observations per terminal node to 10.

Table 3 summarizes the relative influence of the predictors (for a model estimated on the full sample of data). Relative influence is measured by the improvement in the squared error resulting from using a predictor to form splits (Breimann et al. 1983), averaged across base learners (Friedman 2001). We report the mean and the standard deviation of the relative influence of the predictors based on 250 simulation runs of our model. We report simulation results because we introduced a stochastic element into our model by assuming that 50% of the data are selected at random in every step of the recursion given in Equation (6) to build an ensemble of base learners.

The proportion of friends in the GRC, the size of a organizational GRC unit, a volunteer’s age, and time constraints imposed by working on an income-generating job are among the influential predictors. Among the various utility components, the component “to help others” (Utility 13) has the largest relative influence followed by the utility components “to improve one’s standing within the GRC” (Utility 9) and “to strengthen the GRC” (Utility 16). The utility components Utility 13 and Utility 16 are associated with the public-goods model, while the utility component Utility 9 is associated with the human-capital and/or private-consumption model. It should also be noted that the utility components are not among the “top-three” predictors, which shows that some of our control variables are relatively more important than differences in preferences for explaining the cross-sectional variation of volunteer labour supply.

At the same time, the relative importance of the predictors shows that the utility components are not unimportant. Table 4 summarizes aggregates of the relative influence of the utility components across the three model types. The relative influence of the four utility categories adds up to approximately 34%. The aggregation results further show that the public-goods model is about twice as influential (mean relative influence about 12%) than the private-consumption model and the human-capital model (relative influence each about 5–6%). The four utility components that can be either allocated to the private-consumption or the human-capital model help to improve the squared error by about 11%. If we add the mean relative influence of the mixed PC/HC utility category in full either to the PM or the HC category, we get a mean relative influence of about 16–17%. This mean relative influence, when divided

Table 4: Relative Influence by Type of Economic Model (in %)

Utility category	Mean	SD	Components	Component loading
PGM	12.39	0.35	4	3.10
PCM	5.77	0.45	4	1.44
HCM	5.00	0.45	4	1.25
PCM/HCM	10.85	0.33	4	2.71

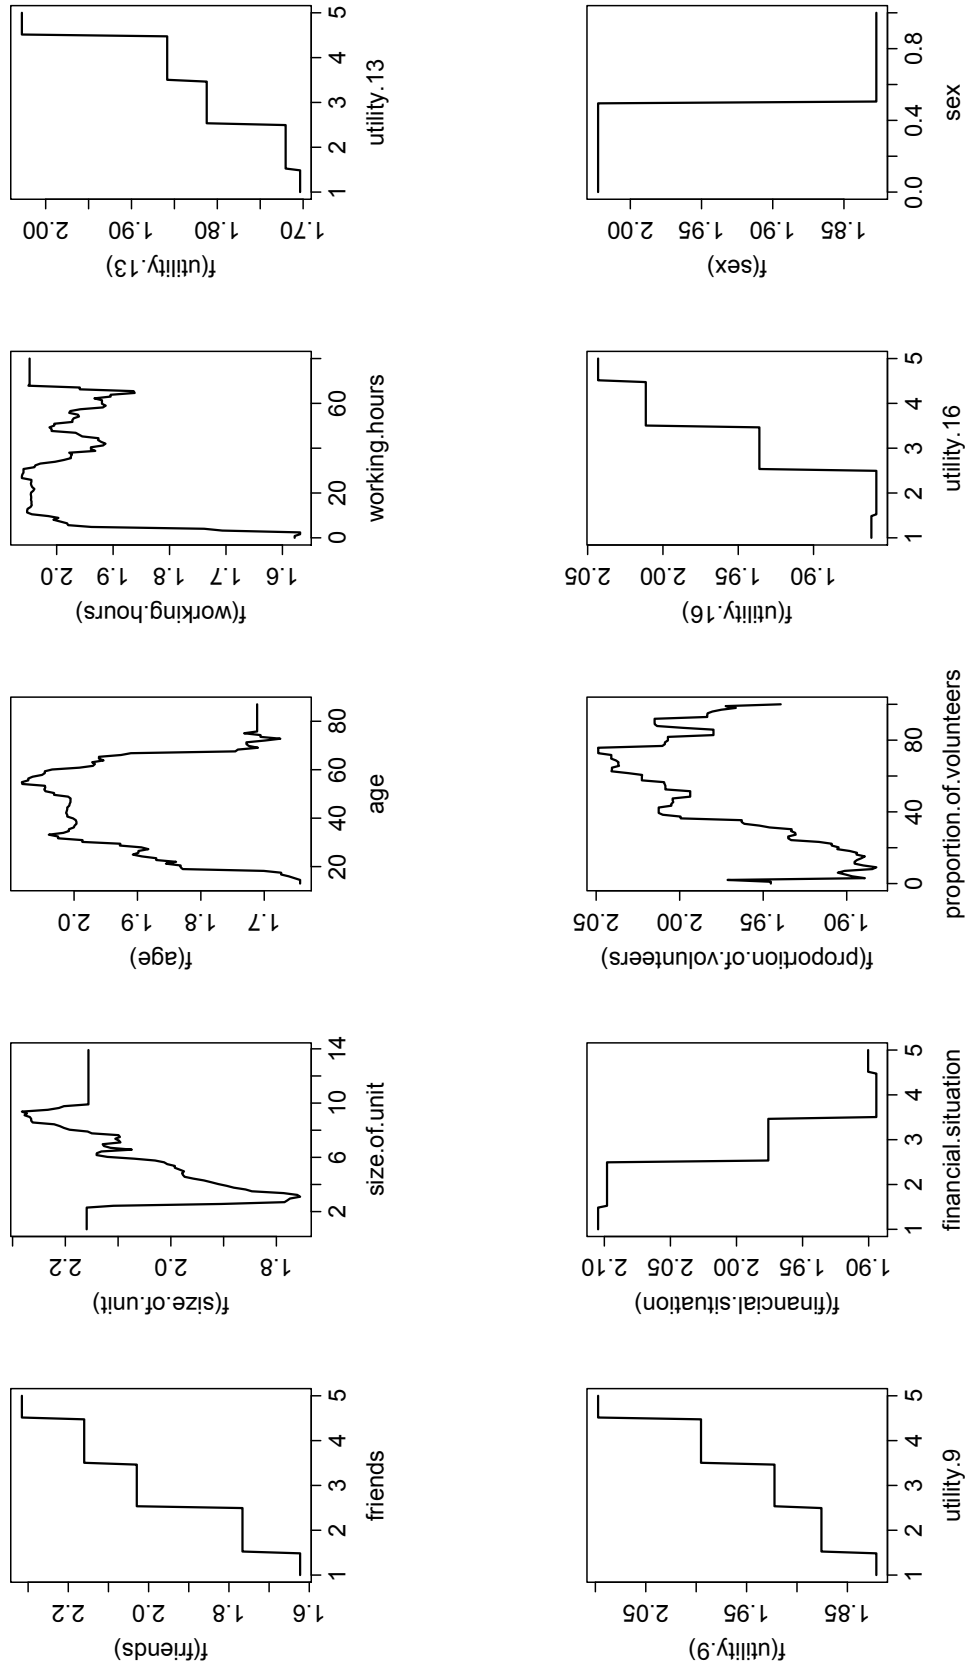
Note: Relative influence is defined as the improvement in squared error resulting from using a predictor to form splits, averaged across base learners. PGM = public-goods model. PCM = private-consumption model. HCM = human-capital model. Mean = mean relative influence. SD = standard deviation of relative influence. Results are based on the same 250 model simulation runs used to compute the results summarized in Table 3. Components = number of utility components in a utility category. Component loading = mean relative influence divided by the number of components.

by the number of utility components forming a utility category, corresponds to a component loading of about 2.0%. The component loading for the public-goods model, based on four utility components, is 3.1%.

Because the relative contribution of the predictors does not answer how exactly volunteer labour supply correlates with the predictors, we plot in Figure 3 the partial dependence of volunteer labour supply to 10 selected predictors. The partial dependence plots show the predictor on the horizontal axis and the predicted volunteer labour supply (in natural logarithm) on the vertical axis, where the effects of the other predictors are taken into account using the weighted traversal technique described by Friedman (2001).

The partial dependence plots show that volunteer labour supply is higher if a volunteer has a large proportion of friends in the GRC. For the lowest rank category volunteer labour supply is approximately $\exp(1.6) \approx 5$ hours per week, while volunteer labour supply is about 10 hours per week for the highest rank category. The partial dependence plot further show that volunteer labour supply first drops as the size of the the organisational GRC unit for which a volunteer works increases. However, when the organiational unit reaches approximately $\exp(3.5) \approx 33$ members, volunteer labour supply starts increasing in unit size. In other words, the PG model (or a model of impact philanthropy) seems to be applicable to small-scale units, while the HC model is useful to organise thinking about volunteer labour supply in larger GRC units. The partial dependence plot for the proportion of volunteers confirms this interpretation insofar as volunteer labour supply first is low if the participation rate is rather low, but then starts increasing. Interestingly, the correlation between volunteer labour supply and the participation rate of other volunteers turns negative again if the participation rate of others gets large.

Figure 3: Partial Dependence Plots



The partial dependence plot for the predictor age shows an inverted u-shaped pattern that is, in principle, consistent with a life-cycle model. The interpretation of the partial dependence plot in terms of a life-cycle model, however, should not be stretched too far. The drop in volunteer labour supply occurs relatively late at an age of about 65–70 years, that is, at retirement age. Moreover, volunteer labour supply not only peaks when volunteers are roughly 30–40 years old, but there is also a second peak when volunteers are 55–60 years old. Moreover, a further inspection of the data showed that low-threshold activities start decreasing relatively late in life, and the number of senior volunteers aged 70 years or older is relatively small.

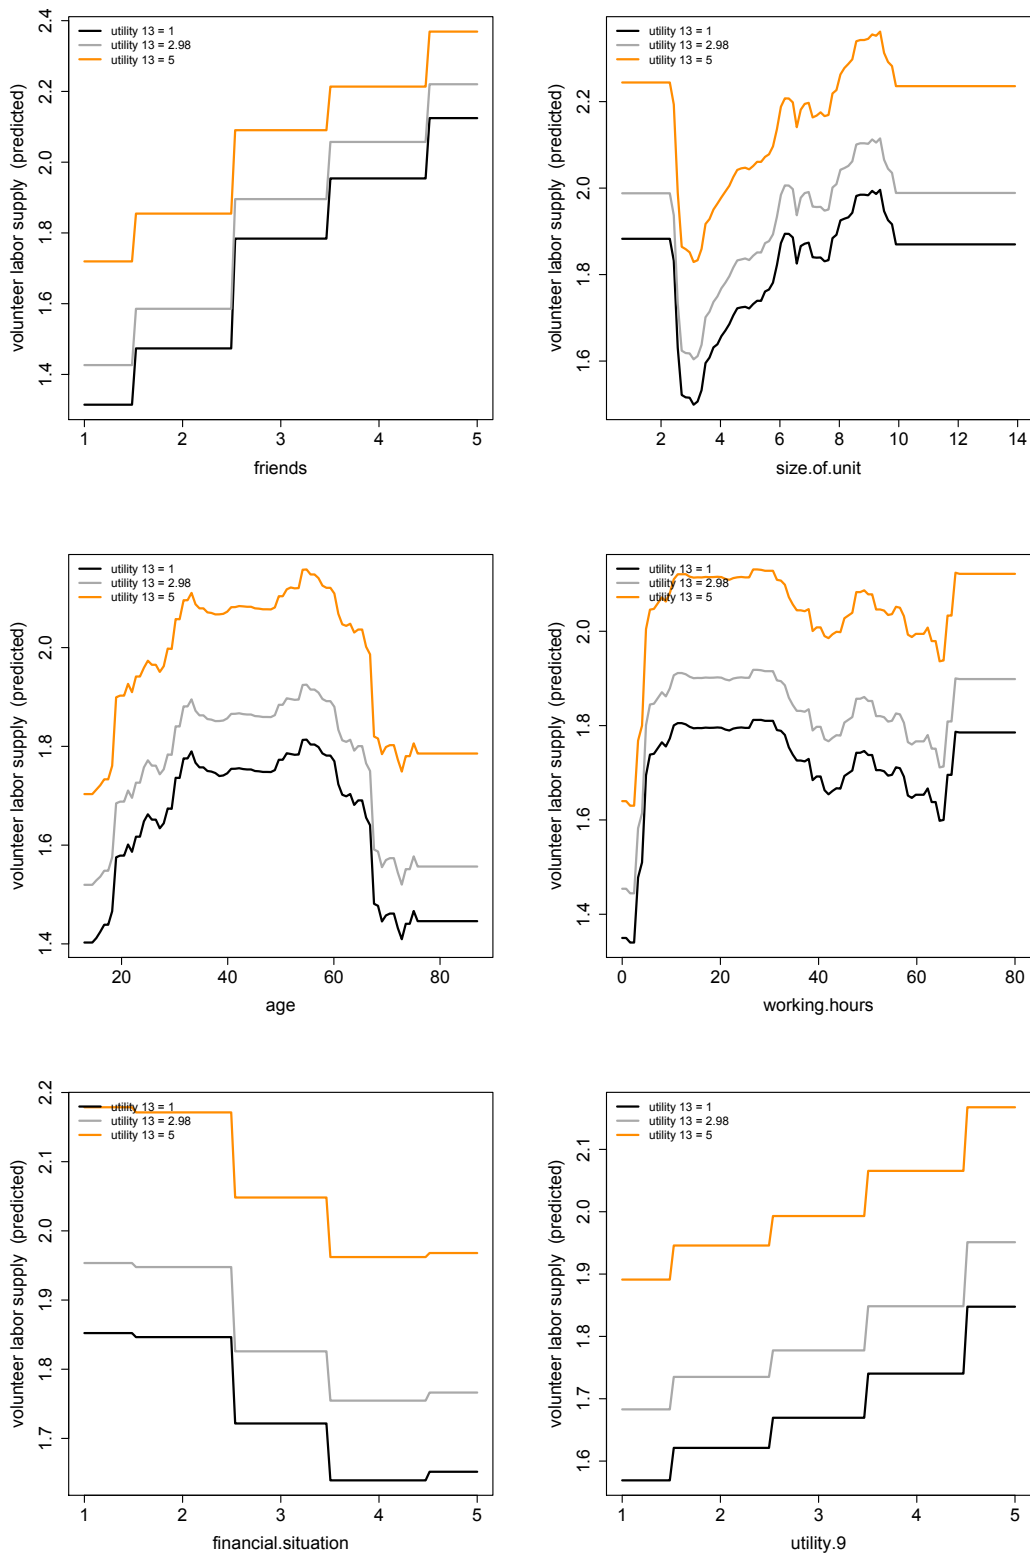
The partial dependence plots further show that volunteer labour supply first increases from approximately 4–5 hours per week to roughly 7–8 hours per week if a volunteer spends a few hours working on an income-generating job. If hours worked reach approximately 25 hours per week, volunteer labour supply starts decreasing until hours worked reaches around 40 hours per week. As hours worked further increase above 60 hours per week volunteer labour supply starts increasing again, but the number of observations in this region is very small.

Male volunteers tend to supply more labour than female volunteers, but the difference is only approximately 1–2 hour per week. Similarly, the correlation of a volunteer’s self-reported financial situation and volunteer labour supply is negative, but the effect is small in quantitative terms.

The partial dependence plots show a positive correlation between volunteer labour supply and the two utility components “to help others” (Utility 13) and “to improve one’s standing within the GRC” (Utility 9) that, according to the relative influence measures reported in Table 4, have the largest relative influence on volunteer labour supply among the various utility components.

Figure 4 illustrates how the utility components “to help others” (Utility 13) interacts with six other predictors. The interaction effects capture how the correlation of volunteer labour supply with, for example, the control variable “friends” changes when we vary the level of the utility component from its minimum, to an intermediate, and then to its maximum value. Three results stand out. First, if the utility component assumes a larger value the partial responses of the other plotted predictor variables shift upward, that is, the cross-derivatives are positive. Second, the upward shift is small when the utility component switches from its minimum to an intermediate value, but the shift gets larger when the switch is from the intermediate value to the maximum value of the utility component, suggesting that a certain extent of nonlinearity characterizes the interaction effects. Third, the model captures interaction effects between

Figure 4: Interaction Effects



utility components. Volunteer labour supply tends to be higher for a given rank of Utility 9 if Utility 13 is also ranked higher.

In order to inspect the fit and data sensitivity of the model, we subdivided the data into an in-sample training set and an out-of-sample test set containing 20% (that is, more than 830 observations) of the data. We then reestimated the boosted regression-tree model on the training set (learning rate 0.01, random sampling rate 50%, 5-fold cross validation). Equipped with the new estimates, we found a satisfactory coefficient of determination for the in-sample set of $R^2 = 0.31$ and of $R^2 = 0.16$ for the out-of-sample test set (the specific values assumed by the R^2 statistics fluctuate to some extent randomly due to the stochasticity of the model).

6 Concluding Remarks

We applied boosted regression trees to study the correlates of volunteer labour supply using a micro dataset for volunteers working for the GRC. Our results suggest that, in terms of the relative-influence measure studied in this research, utility components associated with the public-goods model are somewhat more important correlates of volunteer labour supply than utility components associated with the human-capital model and the private-consumption model. We also found some hints that volunteer labour supply follows a life-cycle pattern, although the interpretation of this finding should not be pushed too far. Furthermore, the public-goods model seems to describe volunteer labour supply in relatively small organisational units. In contrast, a cross-sectional complementarity of volunteer labour supply seems to be at work in larger organisational units, which is in line with the human-capital model.

As for the limitations of our research, a concern is that a social-desirability bias (Bertrand and Mullainathan 2001, among others) might distort our results. A social-desirability bias arises if a volunteer reports that he or she mainly derives utility from helping others simply because altruism is socially highly valued. If interpreted in this way, a social-desirability bias would explain that the PG utility components have a larger relative influence on volunteer labour supply than the PC and HC utility components. While we could not control in this study for a social-desirability bias, it is reassuring that Emrich and Pierdzioch (2015a) in their study of the motives for doing volunteer work find strong evidence supporting the public-goods model. They controlled for a potential social-desirability bias insofar as their empirical model features both self-attributed and other-regarding motives for doing volunteer work. While we deem it unlikely that a social-desirability bias completely reverses our results regarding the relative importance of utility categories, the strength of such a bias deserves further attention in future research.

It is important to emphasize that many of the empirical results reported in this research should not be interpreted in terms of causal effects running from variations in the predictors to a variation in volunteer labour supply. A causal interpretation is justified as far as truly exogenous variables like age, sex, and household size are considered. Exogeneity can also be assumed with regard to the size of an organisational GRC unit and the proportion of volunteers in that organisational unit, and perhaps with regard to religiosity and interest in politics if the latter two represent invariant broader humanitarian values. Other control variables like the proportion of friends in the GRC are likely to be at least in part endogenous. Endogeneity almost surely also holds with respect to the utility derived from volunteering. For example, after having volunteered for a while a volunteer may realise that volunteering more hours per week makes it easier to meet other volunteers and to make a contribution that strengthens the GRC. The potential endogeneity of the utility derived from volunteering can be interpreted to reflect a search process that arises because volunteering changes a volunteer's perception of himself or herself, but also of the GRC (for a survey of the relevant literature, see Wilson 2012; for volunteering as a search process, see Schiff 1980). While our cross-sectional data do not capture such search effects, the interpretation that the utility components recover those part of the variation of volunteer labour supply that can be attributed to cross-sectional differences in preferences (at the time the survey was conducted) is still valid.

It should also be mentioned that, while the sample that we could use for our empirical analysis was relatively large as it contained data for more than 4,000 volunteers, we can be no claim that our sample is representative of the entire population of GRC volunteers. The GRC has more than 400,000 volunteers (<http://www.drk.de/angebote/engagement/ehrenamt.html>). Also, the sample we analysed does not contain data on non-volunteers, implying that the results reported in this research have nothing to say about the dimensions in which volunteers differ from non-volunteers. Finally, it is important to emphasize that the GRC is a very large volunteer organisation, and that it is an altruistic resource pooling. The relative importance of the three types of economic models of volunteer labour supply may be quite different in small volunteer organisation and in group-egoistic volunteer organisations like, for example, sports clubs.

References

- Andreoni, J. (1989). Giving With Impure Altruism: Applications to Charity and Ricardian Equivalence. *Journal of Political Economy*, **97**: 1447–1458.
- Andreoni, J. (1990). Impure Altruism and Donations to Public Goods: A Theory of Warm-Glow Giving. *Economic Journal*, **100**: 464–477.
- Andreoni, J., Brown, E., and Rischall, I. (2003). Charitable Giving by Married Couples: Who Decides and Why Does It Matter?. *Journal of Human Resources*, **38**: 111–133.
- Andreoni, J., and Vesterlund, L. (2001). Which Is The Fair Sex? Gender Differences in Altruism. *Quarterly Journal of Economics*, **116**: 293–312.
- Andreoni, J., Payne, A., and Smith, S. (2014). Do Grants to Charities Crowd Out Other Income? Evidence From the UK. *Journal of Public Economics*, **114**: 75–86.
- Audrino, F., and Trojani, F. (2007). Accurate Short-Term Yield Curve Forecasting using Functional Gradient Descent. *Journal of Financial Econometrics*, **5**, 591–623.
- Bai, J., and Ng, S. (2009). Boosting Diffusion Indices. *Journal of Applied Econometrics*, **24**, 607–629.
- Bajari, P., Nekipelov, D., Ryan, S.P., and Yang, M. (2005). Machine Learning Methods for Demand Estimation. *American Economic Review Papers and Proceedings*, **105**: 481–485.
- Berge, T.J. (2013). Forecasting Disconnected Exchange Rates. *Journal of Applied Econometrics*, **29**, 713–735.
- Bergstrom, T., Blume, L., and H. Varian (1986). On the Private Provision of Public Goods. *Journal of Public Economics*, **29**: 25–49.
- Bertrand, M., and S. Mullainathan (2001). Do People Mean What They Say? Implications for Subjective Survey Data. *American Economic Review Papers and Proceedings*, **91**: 67–72.
- Buchen, T., and Wohlrabe, K. (2014). Assessing the Macroeconomic Forecasting Performance of Boosting: Evidence for the United States, the Euro Area, and Germany. *Journal of Forecasting*, **33**, 231–242.

- Bühlmann, P., and Hothorn, T. (2007). Boosting Algorithms: Regularization, Prediction, and Model Fitting. *Statistical Science*, **22**, 477–505.
- Breiman, L. (1996). Bagging Predictors. *Machine Learning*, **24**: 123–140.
- Breiman, L. (2001). Random Forests. *Machine Learning*, **45**: 5–32.
- Breiman, L., Friedman, J.H., Oshen, R., and Stone, C. (1983). *Classification and Regression Trees*. New York: Chapman & Hall.
- Crumpler, H., and Grossman, P.J. (2008). An Experimental Test of Warm Glow Giving. *Journal of Public Economics* **92**: 1011–1021.
- Cugno, F., and Ferrero, M. (2004). Competition Among Volunteers. *European Journal of Political Economy*, **20**: 637–654.
- Duncan, B. (1999). Modeling Charitable Contributions of Time and Money. *Journal of Public Economics*, **72**: 213–242.
- Duncan, B. (2004). A Theory of Impact Philanthropy. *Journal of Public Economics*, **88**: 2159–2180.
- Eckel, C. C., and Grossman, P. J. (1998). Are Women Less Selfish Than Men?: Evidence From Dictator Experiments. *Economic Journal*, **108**: 726–735.
- Eckel, C.C., Grossman, P.J., and Johnston, M.R. (2005). An Experimental Test of the Crowding Out Hypothesis. *Journal of Public Economics* **89**: 1543–1560.
- Elith, J., Leathwick, J.R., and Hastie, T. (2008). A Working Guide to Boosted Regression Trees. *Journal of Animal Ecology*, **77**: 802–813.
- Emrich, E., and Pierdzioch, C. (2015a). Testing Economic Models of Volunteer Labour Supply: Some Empirical Evidence for the German Red Cross. *Applied Economics*, **47**: 4247–4259.
- Emrich, E., and Pierdzioch, C. (2015b). Gender and Generosity in Charitable Giving: Empirical Evidence for the German Red Cross. *Applied Economics Letters*, **22**: 1041–1045.
- Friedman, J.H. (2001). Greedy Function Approximation: A Gradient Boosting Machine. *Annals of Statistics*, **29**: 1189–1232.
- Friedman, J.H. (2002). Stochastic Gradient Boosting. *Computational Statistics and Data Analysis*, **38**: 367–378.

- Friedman, J.H., and Meulman, J.J. (2003). Multiple Additive Regression Trees With Application in Epidemiology. *Statistics in Medicine*, **22**, 365–1381.
- Glaeser, E.L., Laibson, D., and Sacerdote, B. (2002). An Economic Approach to Social Capital. *Economic Journal*, **112**: F437–F458.
- Glazer, A., and Konrad, K. (1996). A Signaling Explanation for Charity. *American Economic Review*, **86**, 1019–1028.
- Granovetter, M.S: (1973). The Strength of Weak Ties. *American Journal of Sociology*, **78**: 1360–1380.
- Hastie, T., Tibshirani, R. and Friedman, J. (2009) *The Elements of Statistical Learning. Data Mining, Inference, and Prediction*, Springer Science and Business Media, LLC.
- Harbaugh, W.T. (1998). What Do Donations Buy? A Model of Philanthropy Based on Prestige and Warm Glow. *Journal of Public Economics*, **67**: 269–284.
- James, G., Witten. D., Hastie, T., and Tibshirani, R. (2013). *An Introduction to Statistical Learning: With Applications in R*. New York: Springer.
- Mayr, A., Binder, H., Gefeller, O., and Schmid, M. (2014). The Evolution of Boosting Algorithms – From Machine Learning to Statistical Modelling. *Methods of Information in Medicine*, **53**: 419–427.
- Menchink, P.L., and B.A. Weisbrod (1987). Volunteer labour Supply. *Journal of Public Economics*, **32**: 159–183.
- Mitnik, S., Robinzonov, N., and Spindler, M. (2015). Stock Market Volatility: Identifying Major Drivers and the Nature of Their Impact. *Journal of Banking and Finance*, **58**, 1–14.
- Musick, M. A., and J. Wilson (2008). *Volunteers: A Social Profile*. Bloomington: Indiana University Press.
- Payne, A.A. (1998). Does the Government Crowd-Out Private Donations? New Evidence from a Sample of Non-Profit Firms. *Journal of Public Economics* **69**: 323–345.
- R Core Team (2015). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria. <http://www.R-project.org/>.
- Ridgeway, G., with contributions from others (2015). *gbm: Generalized Boosted Regression Models*. R package version 2.1.1. <http://CRAN.R-project.org/package=gbm>.

- Roberts, R.D. (1984). A Positive Model of Private Charity and Public Transfers. *Journal of Political Economy*, **92**: 136–148.
- Rossi, A., and Timmermann, A. (2010). What is the Shape of the Risk-Return Relation?. Unpublished manuscript, available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1364750.
- Schiff, G. (1980). *Charitable Giving and Government Policy: An Economic Analysis*, Greenwood Press: Westport, Connecticut.
- Varian, H.R. (2014). Big Data: New Tricks for Econometrics. *Journal of Economic Perspectives*, **28**: 3–28.
- Veblen, T. (2011)[1899]. *The Theory of the Leisure Class: An Economic Study of Institutions*, Second German Edition. Frankfurt a.M.: Fischer Verlag.
- Wilson, J. (2012). Volunteerism Research: A Review Essay. *Nonprofit and Voluntary Sector Quarterly*, **41**: 176–212.
- Ziemek, S. (2006). Economic Analysis of Volunteers' Motivations – A Cross-Country Study. *Journal of Socio-Economics*, **35**: 532–555.