Estimating Grouped Patterns of Heterogeneity in Repeated Public Goods Experiments

Sebastian Bordt

Theory of Machine Learning Group,

Max Planck Institute for Intelligent Systems and University of Tübingen, Germany

Helmut Farbmacher*

Department of Economics, University of Munich, Germany

Munich Center for the Economics of Aging, Max Planck Society, Germany

This version: September 8, 2019

Abstract

We revisit the role of social preferences and beliefs in voluntary cooperation. It is well documented that people exhibit heterogeneous patterns of behavior in public goods experiments. We estimate individual heterogeneity in data from two such experiments by making use of recently developed methods to detect unobserved parameter heterogeneity in panel data. Comparing our results with those of standard linear regression demonstrates how the latter can yield inconclusive findings. Our results suggest that, in the repeated game, a rather large proportion of players are willing to invest strongly in cooperation, matching their beliefs essentially one to one.

Keywords: experimental economics, public goods game, parameter heterogeneity
JEL Codes: C33, C38, H41

*farbmacher@econ.lmu.de

1 Introduction

In this study we revisit the role of social preferences and beliefs in voluntary cooperation. There is a vast literature discussing potential explanations for two key observations in repeated public goods games: non-negligible contributions at the beginning of the game followed by a decline in cooperation as the game continues. Using an innovative experimental design, Fischbacher and Gächter (2010) provide evidence that this decline is due to beliefs about other people's contributions and differences in cooperation preferences. In the experiment, their subjects contributed, on average, a weighted average of their elicited contribution schedule and exhibited perfect conditional cooperation. Making use of recent advancements in econometrics, we are able to separate both of these aspects and provide evidence that a rather large proportion of players are actually willing to invest strongly in cooperation, matching their beliefs essentially one to one.

It is well known in the literature that different players in public goods experiments exhibit heterogeneous patterns of behavior (Chaudhuri, 2011). These patterns may result from differences in, for example, social preferences, fairness preferences, or sensitivity to social norms, all of which are variables that are generally not observed by the researcher. There have been many successful attempts to measure and control for certain aspects of this heterogeneity. For instance, deOliveira et al. (2015) investigated the extent to which heterogeneity in social preferences affects cooperation. Kimbrough and Vostroknutov (2016) measured heterogeneity in sensitivity to social norms and investigated its impact on cooperation in public goods games. Quantifying the heterogeneity along all relevant dimensions simultaneously, however, is fundamentally difficult, if not impossible, in such research. We approach this problem by pursuing an econometric approach to account for unobserved heterogeneity in public goods experiments.

We report heterogeneous behavior in two independent public goods experiments, making use of recent advancements in the estimation of unobserved parameter heterogeneity in panel data (Bonhomme and Manresa, 2015; Su et al., 2016; Ando and Bai, 2016; Ke et al., 2016).

We apply the grouped fixed-effects estimator developed by Bonhomme and Manresa (2015) and the classifier-Lasso of Su et al. (2016). Both estimators assume a finite number of groups within the data and allow for heterogeneous regression coefficients across these groups. With our approaches, neither the behavioral type nor the behavioral patterns (for instance, social preferences) that shape the relation between outcome and explanatory variables within a type have to be known to the researcher, but rather are determined in a data-driven fashion using the observed contribution behavior.

To our knowledge, this is the first time that such novel estimation methods have been applied to public goods experiments or experimental economics in general. However, they have already been applied in other fields of economics. For instance, Bonhomme and Manresa (2015) used their grouped fixed-effects estimator to analyze the association between income and democracy. Lu and Su (2017) employed the classifier-Lasso to examine this same relationship. Oberlander et al. (2017) and Guner et al. (2018) used the grouped fixed-effects estimator in health economics to explore the health effects of nutrition and marriage, respectively. Wang et al. (2019), in turn, used the classifier-Lasso estimator in the field of labor economics to analyze heterogeneous effects of minimum wages on employment.

In our study we use data from public goods experiments conducted by Fischbacher and Gächter (2010) and deOliveira et al. (2015), both of which contain information about beliefs. Fischbacher and Gächter (2010) presented an innovative experimental design in which participants' preferences were elicited with the strategy method and participants played a public goods game that was repeated 10 times. In the game, players also stated their beliefs about the average contributions made by others. In addition to the conditional contribution schedules elicited via the strategy method, this makes it possible to analyze the relation between contributions, beliefs, and social preferences. deOliveira et al. (2015) modified this experimental design to study the role of group composition and information. In a public goods game repeated 15 times, participants played in fixed groups of three, simultaneously stating their beliefs and contributions.

Our main result is a classification of participants into two groups. The heterogeneity we detect is broadly consistent with classifications based on social preferences elicited by the strategy method. In our study, we provide evidence that participants in one group behaved, on average, as perfect conditional cooperators, matching their contributions and beliefs one to one; once beliefs have been accounted for, social preferences elicited with the strategy method do not help explain contributions. Participants in the other group, however, showed a pronounced self-serving bias and contributed significantly less than they believed others were going to contribute. Interestingly, the parameter heterogeneity suggests the existence of a rather large proportion of participants who were willing to invest strongly in cooperation.

We descriptively relate our two groups to classifications based on the strategy method and find that a non-negligible proportion of players did not behave in the repeated game as predicted by this method. Because our approach is based on the actual behavior in the game, we can shed light on these deviations. However, future research is needed to confirm the predictive power of our data-driven classification. Other studies have already reported non-declining contributions in repeated public goods experiments once subjects were sorted (with or without their knowledge) into groups according to their preference type (Burlando and Guala, 2005; Gunnthorsdottir et al., 2007; Kimbrough and Vostroknutov, 2016). This is an interesting avenue for future research.

Our paper is organized as follows. Section 2 outlines our estimation methods. Section 3 contains our empirical results, and Section 4 relates them to classification methods based on the strategy method. Section 5 concludes.

2 Model

Consider the following standard linear regression model

$$y_{it} = x'_{it}\beta + u_{it}, \tag{1}$$

where y_{it} is individual i's contribution to a public good during period t and x_{it} a vector of explanatory variables (containing beliefs and predicted contributions). The model above averages the relation between outcome and explanatory variables across all individuals, summarizing it in a single regression coefficient β. In public goods experiments, there is an agreement that people behave in ways that are qualitatively different (Fischbacher et al., 2001; deOliveira et al., 2015). Qualitatively different types of behavior should be reflected by different relations between outcome and explanatory variables. Hence, a single vector of regression coefficients (β) may not be sufficient to describe individual behavior in the repeated game accurately. In fact, the individual who behaves like the "average" player whose behavior is analyzed by standard linear regressions may not even exist (see also Moffatt, 2016, for a thorough discussion of this issue).

To overcome this problem, we use the grouped-fixed effects (GFE) estimator by Bonhomme and Manresa (2015) to estimate the following model

$$y_{it} = \begin{cases} x'_{it}\beta_1 + \alpha_{1\,t} + \epsilon_{it} & \text{if } g_i = 1 \\ x'_{it}\beta_2 + \alpha_{2\,t} + \epsilon_{it} & \text{if } g_i = 2\,, \end{cases} \tag{2}$$

which has two different slope coefficients β_1 and β_2. Individual i's contribution relates to the explanatory variables through either β_1 or β_2. We say that those individuals for whom the relationship is described by β_1 are in group 1, and those for whom it is described by β_2 are in group 2. The fact that there are two group-specific slope coefficients allows us to capture two qualitatively different relations between outcome and explanatory variables. Group membership of individual i is denoted by the indicator variable $g_i \in \{1, 2\}$. Lastly, $\alpha_{1\,t}$ and $\alpha_{2\,t}$ are two group-specific unobserved time trends. These time trends are a powerful way to control for additional unobserved heterogeneity.[1]

In addition to the GFE estimator, we employ the classifier-Lasso (c-Lasso) estimator by

[1]Computation was done in Python. Our implementation is available at `https://github.com/sbordt/gfe-py`.

Su et al. (2016) to estimate the following model

$$y_{it} = \begin{cases} x'_{it}\beta_1 + c_i + \epsilon_{it} & \text{if } g_i = 1 \\ x'_{it}\beta_2 + c_i + \epsilon_{it} & \text{if } g_i = 2, \end{cases} \tag{3}$$

which is similar to the model in eq. (2), except that the group-specific time trends $\alpha_{1\,t}$ and $\alpha_{2\,t}$ have been replaced by individual fixed-effects (c_i). The c-Lasso estimator serves as a robustness check regarding the a priori unknown structure of unobserved heterogeneity.[2]

A key advantage of our approach is that estimation and classification are perfomed in a data-driven manner and require no prior description of different behavioral types. In both models, group membership, $g_i \in \{1, 2\}$, and the group-specific regression coefficients, β_1 and β_2, are jointly estimated from the data. This implies that among rival groupings of individuals, our methods impartially detect the ones that are best supported by the data. Note that because our estimation assigns each individual i to group g_i, we also perform a classification, which can be used to analyze further the individual behavior of a single player or compare the classification with other approaches based, for example, on the strategy method.

An important issue is that the researcher has to choose the number of potential groups in both models. Ideally, one would like to have a separate set of regression coefficients for each meaningful behavioral pattern. In practice, however, the size of our experimental datasets is too limited to attempt this. Hence, we resort to determining the behavioral patterns that can actually be estimated with the available data. The choice of two different regression coefficients is supported by an AIC criterion. Additionally, the results with three and four groups reveal that there is no additional meaningful group.

[2] Computation was done using the MATLAB code provided by Zhentao Shi at `https://github.com/zhentaoshi/C-Lasso`.

Dependent Variable: Contribution

Model	OLS	FE	Grouped Fixed-Effects		Post c-Lasso	
	All (1)	All (2)	Group 1 (3a)	Group 2 (3b)	Group 1 (4a)	Group 2 (4b)
Panel A. All Players, Periods 1-10						
Belief	0.666*** (0.059)	0.620*** (0.053)	1.005*** (0.180)	0.145 (0.177)	1.041*** (0.102)	0.070* (0.028)
Predicted Contribution	0.242** (0.069)	0.266** (0.046)	-0.071 (0.216)	0.359 (0.575)	0.003 (0.057)	0.118 (0.075)
Constant	-0.473 (0.244)	-0.270 (0.351)			-0.935 (0.487)	1.756*** (0.187)
Avg. of Periods 1-5			1.439	0.474		
Avg. of Periods 6-10			0.547	0.086		
Observations	1400	1400	700	700	900	500
R^2	0.34	0.38	0.50	0.24	0.55	0.02

Table 1: Columns 1, 2, 4a, 4b: Robust standard errors in parenthesis. Columns 3a, 3b: Bootstrapped standard errors in parenthesis. Data from Fischbacher and Gächter (2010).
*** 1% significant, ** 5% significant, * 10% significant.

3 Results

We first analyze the Fischbacher and Gächter (2010) data. The regressions cover all 140 players who took part in the experiment. Our main results are depicted in Table 1. This shows the results of the regression of contributions on beliefs and the predicted contributions for four different models. The estimators are, as depicted from left to right, ordinary least squares (OLS), fixed-effects (FE), grouped fixed-effects and c-Lasso.[3] The OLS results – depicted in column 1 – replicate model 3 in table 2 of Fischbacher and Gächter (2010). The coefficients on beliefs and predicted contributions are both positive and significant. Accounting for fixed effects hardly affects the coefficients (see column 2).

Columns 3a-4b display the results of regressions that allow for unobserved parameter heterogeneity between two groups. The estimation results are presented in two columns per estimator, and each column shows the regression coefficients in one group. Both the GFE and c-Lasso estimators identify one group in which the coefficient on beliefs is significant and almost exactly equal to 1, whereas the coefficient on predicted contributions is close to 0 and statistically insignificant. We will call it group 1 from now on. The constant term in column

[3] In case of the c-Lasso, the regression coefficients are obtained by post-Lasso regressions (i.e., they are unshrinked). This does not affect the groups. See Su et al. (2016).

Mean results (by group)

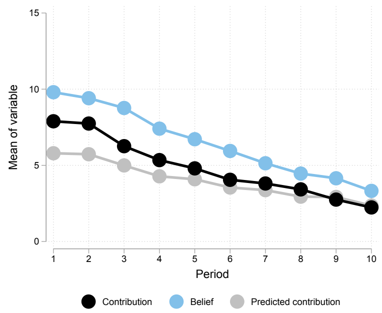

(a) Experiment

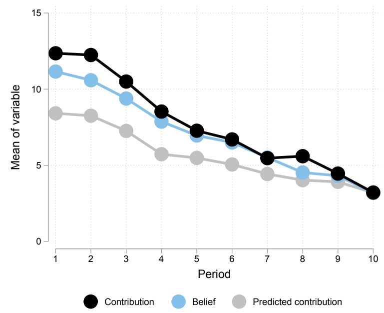

(b) Group 1

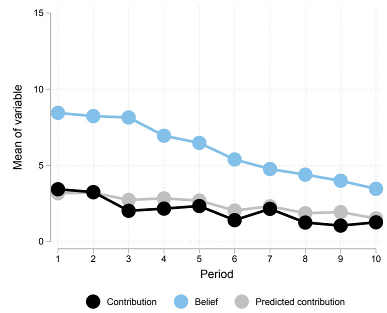

(c) Group 2

Classifications of three individuals (as examples)

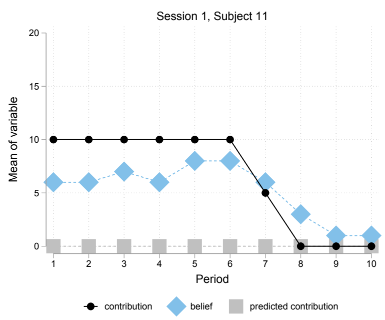

(d) GFE: 1, c-Lasso: 1

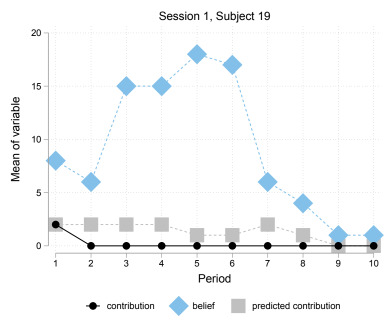

(e) GFE: 2, c-Lasso: 2

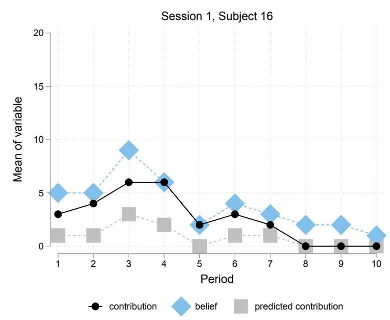

(f) GFE: 2, c-Lasso: 1

Figure 1: *Top row:* Mean contributions, mean beliefs, and mean predicted contributions in the overall experiment and the groups identified by the grouped fixed-effects estimator. *Bottom row:* Three subjects and their classification according to the grouped fixed-effects estimator and the c-Lasso. All figures are based on the data of Fischbacher and Gächter (2010).

4a is small and insignificant, which suggests that players in group 1 played, on average, in alignment with their beliefs. In columns 3a and 3b, we summarize the estimated group-specific time trends by their averages over periods 1-5 and 6-10. In both groups, the time trends are positive but small and decline over time, which is in line with the well-documented decline in cooperation observed in repeated public goods games.

The number of players in one group can be calculated by dividing the number of observations by the number of periods. The GFE estimator assigns 70 players to group 1 and 70 players to group 2. In comparison, the c-Lasso assigns 90 players to group 1 and 50 players to group 2. Both classifications overlap to a large extent. There are 63 players who belong to group 1 regardless of the estimation method.

The summary statistics depicted in Figure 1 (upper panel) give further insights into the nature of the groups estimated by GFE.[4] Figure 1(a) depicts average contributions, beliefs, and predicted contributions over all subjects. Consistent with the OLS results in column 1 of Table 1, contributions are a weighted average of beliefs and predicted contributions. However, this relation does not hold within any of the two estimated subgroups. On the one hand, players in group 1 played, on average, in alignment with their beliefs (see Figure 1b). On the other hand, players in group 2 – depicted in part (c) – played, on average, in alignment with their predicted contributions. Moreover, the mean contributions were considerably higher in group 1 than in group 2, although beliefs about the contributions of the other players were just slightly different across both groups.

To analyze whether the estimated group-specific behaviors hold not only on average but also at the individual level, we visualize the experimental data at the latter. Figure 1(d)-(f) depicts the relation between contributions, beliefs, and predicted contributions for three differently classified individuals. The subject depicted in part (d), although classified as a free rider based on the strategy method, invested strongly in cooperation in the first six periods. Afterwards, her contributions declined along with her beliefs about the others' contributions. She is classified by both estimators as belonging to group 1 – that is, individuals who matched their beliefs essentially one to one. The subject depicted in part (e) played the dominant strategy except in the first period. She is classified by both estimators as belonging to group 2. The subject depicted in part (f) often contributed a weighted average of what she believed others had contributed and her predicted contributions. She is classified as belonging to group 2 of the GFE estimator and group 1 of the c-Lasso. Hence, there is indeed evidence that the average behavior within groups holds at the individual level as well. Figures for all subjects are available at `https://github.com/sbordt/PublicGoodsHeterogeneity/`.

We now turn to the data of deOliveira et al. (2015). Table 2 shows regression results covering all 102 players who took part in the experiment. As in Table 1, the estimators

[4]Similar results hold for the groups estimated by the C-Lasso (see Appendix).

Dependent Variable: Contribution

Model	OLS	FE	Grouped Fixed-Effects		Post c-Lasso	
	All	All	Group 1	Group 2	Group 1	Group 2
	(1)	(2)	(3a)	(3b)	(4a)	(4b)
Panel A. All Players, Periods 1-15						
Belief	0.912***	0.624***	1.015***	0.187	1.031***	0.177***
	(0.024)	(0.108)	(0.135)	(0.140)	(0.153)	(0.060)
Constant	-0.502***	1.573*			-0.752	3.936***
	(0.136)	(0.778)			(1.371)	(0.322)
Avg. of Periods 1-5			1.235	0.428		
Avg. of Periods 6-10			0.230	0.036		
Observations	1530	1530	960	570	750	780
R^2	0.49	0.18	0.57	0.14	0.35	0.03

Table 2: Columns 1, 2, 4a, 4b: Robust standard errors in parenthesis. Columns 3a, 3b: Bootstrapped standard errors in parenthesis. Data from deOliveira et al. (2015).
*** 1% significant, ** 5% significant, * 10% significant.

are OLS, FE, GFE, and c-Lasso. Note that predicted contributions are not present in this dataset. Nevertheless, the results obtained from this independent experiment reveal remarkable similarities to the results discussed above. Both GFE and c-Lasso estimate one group in which the coefficient on beliefs is significant and almost exactly equal to 1, and a second group in which the coefficient on beliefs is also positive but much smaller compared to OLS or the other group. As above, we label the group with the larger coefficient on beliefs as "group 1". The constant term in column 4a is insignificant, which shows that players in group 1, again, played on average in alignment with their beliefs. The overlap between both classifications is slightly smaller compared to the Fischbacher and Gächter (2010) results. There are 45 players who are classified into group 1 by both estimation methods.

4 Discussion

Fischbacher and Gächter (2010) label subjects who match contributions and beliefs about other players' contributions one to one as perfect conditional cooperators. It is striking that our estimation methods can detect subjects who, on average, exhibited precisely such a behavior although we do not set up behavioral types a priori. As outlined in Section 2, these estimates are based solely on the relation between outcome and explanatory variables

Estimated player groups and player types as elicited with the strategy method

Model	Grouped Fixed-Effects		Post c-Lasso	
	Group 1	Group 2	Group 1	Group 2
Panel A. Fischbacher et al. (2001)				
Conditional Cooperators	44 [38.5]	33 [38.5]	53 [49.5]	24 [27.5]
Free Riders	9 [16.0]	23 [16.0]	14 [20.6]	18 [11.4]
Triangle Contributors	9 [8.5]	8 [8.5]	14 [10.9]	3 [6.1]
Other	8 [7.0]	6 [7.0]	9 [9.0]	5 [5.0]
	Fisher's Exact Test p = 0.0449		Fisher's Exact Test p = 0.0326	
Panel B. Fallucci et al. (2017)				
Unconditional Cooperators	3 [2.5]	2 [2.5]	2 [3.2]	3 [1.8]
Strong Conditional Cooperators	35 [25.5]	16 [25.5]	40 [32.8]	11 [18.2]
Weak Conditional Cooperators	14 [12.0]	10 [12.0]	18 [15.4]	6 [8.6]
Own Maximisers	14 [27.5]	41 [27.5]	26 [35.4]	29 [19.6]
Various	4 [2.5]	1 [2.5]	4 [3.2]	1 [1.8]
	Fisher's Exact Test p = 0.0000		Fisher's Exact Test p = 0.0042	

Table 3: Relation between estimated player groups and preference types as elicited with the strategy method. Depicted are the number of players belonging to each of the player types deterimend by Fischbacher et al. (2001) and Fallucci et al. (2017) within each group. Groups are those from table 1. Expected frequencies under independence are in brackets.

observed in the data.

It is natural to ask how the estimated groups relate to the well-known player types as proposed by Fischbacher et al. (2001). Note that our method is based on behavior in the repeated game, whereas the classification by Fischbacher et al. (2001) is based on the strategy method. This may explain differences between contribution preferences identified by the strategy method and our classification, which builds upon the actual contribution decisions. The relation is tabulated in panel A of Table 3. As one might expect, conditional cooperators are more likely to be in group 1 than in group 2. Similarly, free riders are more likely to be in group 2 than in group 1. A test of independence rejects in all cases at conventional levels. However, there are nine free riders who decided to match their contributions and beliefs closely in the repeated game (one of them is Figure 1d). Similar behavior was observed by Gächter and Thöni (2005), who found that players in sorted groups who had previously been classified as "low contributors" decided to contribute more than some "cooperative" players.

Panel B of Table 3 compares our classification with the typology of behavior recently proposed by Fallucchi et al. (2018). Fallucchi et al. (2018) used hierarchical clustering analysis to identify four main behaviors under the strategy method. Panel B shows that strong conditional cooperators are more likely to be in group 1, whereas own maximizers are more likely to be in group 2.

Another dimension of heterogeneity is stressed in Chaudhuri et al. (2017), who show that heterogeneity in the distribution of initial beliefs among conditional cooperators can lead to a decay of contributions. Because initial beliefs were higher in group 1 than in group 2 (grouped-fixed effects, Mann-Whitney $p < 0.001$), we also find that initial beliefs correlate with behavior in the repeated game.

5 Conclusion

We estimate individual heterogeneity in repeated public goods experiments using data from two such experiments conducted independently. Our study reveals substantial heterogeneity in the relation between contributions, beliefs, and predicted contributions. More specifically, the regression results imply that predicted contributions have little explanatory power for a substantial proportion of players once the heterogeneous effects of beliefs have been controlled for. Although predicted contributions are clearly related to behavior in the repeated game, our results suggest that there is substantial heterogeneity within public goods games that standard linear regressions cannot account for.

Our main result is that in the repeated game, a substantial proportion of players behaved as perfect conditional cooperators. We find some overlap with classifications based on the strategy method. For instance, conditional cooperators are more likely to be in the group in which the coefficient on beliefs is almost exactly equal to 1. However, the observed behavior of some individuals was opposed to their social preferences as elicited with the strategy method.

Our data-driven classification is entirely descriptive insofar as it remains silent on the causes of belief play. A natural distinction would be between social-preference-based explanations and strategic considerations. Probably both explanations hold true for different players (see Reuben and Suetens, 2012; Gächter and Thöni, 2005; Weber et al., 2018). Strategic considerations seem especially plausible for free riders (as identified by the strategy method) who decide to contribute in the repeated game.

Further research is needed to reveal the structural and behavioral nature of the classification, and its power to predict future behavior. For instance, it would be useful to verify whether a sorting according to the classification can be used to overcome the decline in contributions.

References

Ando, Tomohiro and Jushan Bai (2016), "Panel data models with grouped factor structure under unknown group membership." *Journal of Applied Econometrics*, 31, 163–191.

Bonhomme, Stéphane and Elena Manresa (2015), "Grouped patterns of heterogeneity in panel data." *Econometrica*, 83, 1147–1184.

Burlando, Roberto M. and Francesco Guala (2005), "Heterogeneous agents in public goods experiments." *Experimental Economics*, 8, 35–54.

Chaudhuri, Ananish (2011), "Sustaining cooperation in laboratory public goods experiments: a selective survey of the literature." *Experimental Economics*, 14, 47–83.

Chaudhuri, Ananish, Tirnud Paichayontvijit, and Alexander Smith (2017), "Belief heterogeneity and contributions decay among conditional cooperators in public goods games." *Journal of Economic Psychology*, 58, 15–30.

deOliveira, Angela C. M., Rachel T. A. Croson, and Catherine Eckel (2015), "One bad apple? Heterogeneity and information in public good provision." *Experimental Economics*, 18, 116–135.

Fallucchi, Francesco, R. Andrew Luccasen, and Theodore L. Turocy (2018), "Identifying discrete behavioural types: a re-analysis of public goods game contributions by hierarchical clustering." *Journal of the Economic Science Association*, forthcoming.

Fischbacher, Urs and Simon Gächter (2010), "Social preferences, beliefs, and the dynamics of free riding in public goods experiments." *American Economic Review*, 100, 541–556.

Fischbacher, Urs, Simon Gächter, and Ernst Fehr (2001), "Are people conditionally cooperative? Evidence from a public goods experiment." *Economics Letters*, 71, 397–404.

Gächter, Simon and Christian Thöni (2005), "Social learning and voluntary cooperation among like-minded people." *Journal of the European Economic Association*, 3, 303–314.

Guner, Nezih, Yuliya Kulikova, and Joan Llull (2018), "Reprint of: Marriage and health: Selection, protection, and assortative mating." *European Economic Review*, 109, 162–190.

Gunnthorsdottir, Anna, Daniel Houser, and Kevin McCabe (2007), "Disposition, history and contributions in public goods experiments." *Journal of Economic Behavior & Organization*, 62, 304–315.

Ke, Yuan, Jialiang Li, and Wenyang Zhang (2016), "Structure identification in panel data analysis." *Annals of Statistics*, 44, 1193–1233.

Kimbrough, Erik O. and Alexander Vostroknutov (2016), "Norms make preferences social." *Journal of the European Economic Association*, 14, 608–638.

Lu, Xun and Liangjun Su (2017), "Determining the number of groups in latent panel structures with an application to income and democracy." *Quantitative Economics*, 8, 729–760.

Moffatt, Peter G (2016), *Experimetrics: Econometrics for Experimental Economics*. Palgrave, London.

Oberlander, Lisa, Anne-Célia Disdier, and Fabrice Etilé (2017), "Globalisation and national trends in nutrition and health: A grouped fixed-effects approach to intercountry heterogeneity." *Health Economics*, 26, 1146–1161.

Reuben, Ernesto and Sigrid Suetens (2012), "Revisiting strategic versus non-strategic cooperation." *Experimental Economics*, 15, 24–43.

Su, Liangjun, Zhentao Shi, and Peter C. B. Phillips (2016), "Identifying latent structures in panel data." *Econometrica*, 84, 2215–2264.

Wang, Wuyi, Peter CB Phillips, and Liangjun Su (2019), "The heterogeneous effects of the minimum wage on employment across states." *Economics Letters*, 174, 179–185.

Weber, Till O., Ori Weisel, and Simon Gächter (2018), "Dispositional free riders do not free ride on punishment." *Nature Communications*, 9, 2390.

A Appendix

A.1 Additional Figures

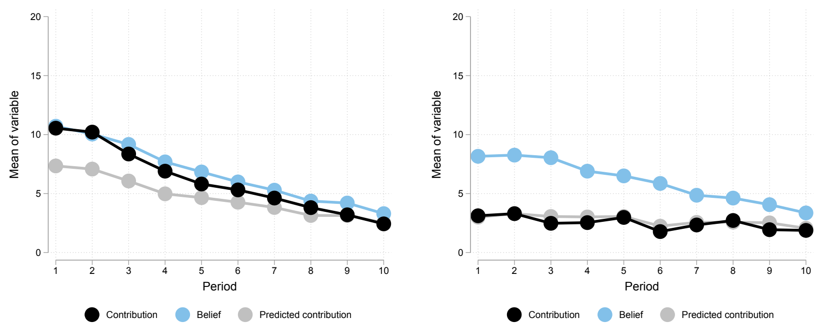

(a) Group 1 **(b)** Group 2

Figure A.1: Mean contributions, mean beliefs, mean predicted contributions and mean contributions made by other players in the groups identified by the C-Lasso.