

Does revealing personality data affect prosocial behavior?

Abstract

Many modern organisations collect data on individuals' personality traits as part of their human resource selection processes. We test experimentally whether revealing information on personality data impacts on pro-social behaviour as measured in a one-shot modified dictator game and a public goods game. Our focus is on the personality trait of agreeableness which has been shown to be a significant determinant of pro-sociality. We provide new evidence that revealing personality information for disagreeable individuals has detrimental effects on their pro-social behaviour as compared to the baseline no-information benchmark. This is not the case, however, for agreeable individuals when they are matched with agreeable individuals. Agreeable individuals become less pro-social when matched with disagreeable individuals and are aware of this. Our results suggest that information cues about personality significantly affect economic behaviour and have implications for employees' personality assessments as part of standard hiring processes.

JEL-Codes: C910, D700, H410.

Keywords: personality, social preferences, inequity aversion, cooperation, laboratory experiment.

Michalis Drouvelis
Department of Economics
University of Birmingham, Edgbaston
United Kingdom – Birmingham, B15 2TT
m.drouvelis@bham.ac.uk

Nikolaos Georgantzis
Burgundy School of Business
School of Wine and Spirits Business
France – 21006 Dijon Cedex
nick.georgantzis@gmail.com

February 2019

Forthcoming Journal of Economic Behavior and Organization.

We thank Loukas Balafoutas, Aurora Garcia-Gallego, Martin Kocher and Benjamin Marx, seminar participants at the University of Cologne, Universitat Jaume I Castellon, University of Michigan, the University of Illinois at Urbana-Champaign, the University of Kent and the University of Warwick for helpful comments and Jose Vicente Guinot for programming assistance. Michalis Drouvelis thanks C-SEB at the University of Cologne for the hospitality he enjoyed while working on this paper. Financial support from the University of Birmingham, Ministerio de Economía y Competitividad (ECO2015-68469-R) and UJI (project P1.1B2015-48) is gratefully acknowledged.

1. Introduction

In controlled laboratory experiments, interactions often take place without individuals having information about whom they are matched with. However, this is not typically observed in real-life settings where people can acquire information about others' individual characteristics. For example, in many modern corporations, human resource selection processes involve a personality assessment of the employees.¹ These processes are increasingly based on the elicitation of the candidates' personality traits (Campion et al., 1997). Such personality assessments rely on self-reported attitudes of the individuals, who respond to structured questionnaires. The reasons for using personality inventories in human resource selection processes include comparability across professional profiles within a firm but also the relevance of personality data in predicting job behaviour. Among such personality tests, the most well-known is the Big Five questionnaire which has been developed, validated and refined by Goldberg (1981, 1990). As noted by Sackett and Walmsley (2014), personality tests (such as the Big Five questionnaire) can be used by human resource managers to select their employees as personality traits are important for success across many jobs. Relatedly, a meta-analysis by Barrick and Mount (1991) provides evidence that the Big Five personality dimensions show significant relations to job performance in several occupational groups (see also Vinchur et al. (1998)). In addition to this, social relationships are often shaped on the basis of such characteristics and thus, having access to this information may affect economic behaviour. Of particular importance in shaping social interactions is the role of personality. This has been recognised by personality psychologists (e.g., see Roberts, (2006) for a theoretical framework and Almlund et al., (2011) for an empirical review) as well as economists (for a review, see Becker et al., 2012). While existing research shows that personality plays a central role in social and labour relationships, it still remains unclear whether and if so, how revealing information about individuals' personality traits affects economic decision making. In this paper, we specifically ask: does revealing personality data have any impact on pro-sociality?

An established literature on personality psychology starting from the 1970s suggests that the Big Five framework posits a hierarchical organisation for personality traits. The Big Five traits are classified as: openness to experience, conscientiousness, extraversion, agreeableness and neuroticism. Previous research in psychology (e.g., Costa and McCrae,

¹ In particular, there is evidence for the use of personality tests in firms' hiring decisions (see Autor and Scarborough, 2008).

1992) indicates that most of the variables used to measure personality traits can be mapped into one of the above dimensions of the Big Five framework. In our study, we concentrate on the personality trait of *agreeableness*, defined as “the tendency to act in a cooperative, unselfish manner” (see Almlund et al., 2011). Our motivation to focus on agreeableness stems from existing research in psychology as well as in economics which provides evidence that agreeableness is a good predictor of behavioural outcomes. For example, Dohmen et al. (2008) find that agreeableness is positively correlated with trustworthiness in a trust game and Ben-Ner et al., (2004a) find that the more agreeable the subjects are the greater is the amount they send in response to the amount they receive. Dohmen et al. (2010) show that agreeableness is negatively correlated with risk tolerance. Volk et al. (2011) find that agreeableness emerged as the only significant predictor of cooperation preferences (see also Ben-Ner et al., 2004b). Bekkers (2006) provides evidence that agreeableness and blood donations are positively correlated, while Ashton et al. (1998) and Osinski (2009) show that agreeableness is positively correlated with altruistic behaviour. The key message from these studies is that agreeableness constitutes an important determinant of pro-social behaviour.

Yet, a common characteristic of this literature is that individuals are not informed of their own and others’ personality traits. In our experiment, we deviate from this strand of literature and analyse whether receiving information in relation to own and others’ personality type impacts pro-social behaviour. If personality data were to be revealed in workplace environments, then it is natural to assume that employees would be likely to find out information both about their own as well as other employees’ personality traits through their multiple interactions within a firm. This corresponds to a full-information environment which we have implemented in our setting. Additionally, unlike other individual characteristics, such as gender and age, which are known to oneself and can be typically identified by others, personality traits are less salient and observable (especially in the absence of previous interactions similar to our experimental setting). It is also often the case that inferences about others’ personality traits can be quite inaccurate.² Thus, revealing data about own and others’ personality types is warranted and informative for both parties. We test whether simple information cues about personality types embody crucial information, which in turn affects behaviour. Thus far, the literature lacks of evidence assessing the impact of revealing personality data on pro-sociality. We fill this gap in the literature by taking advantage of the experimental methodology and design a laboratory experiment which allows

² Note that it is also likely to be the case that individuals may not have accurate information about their own score on personality tests and in particular, on specific personality traits.

us to draw causal conclusions about the relevance of providing information about personality cues on pro-social behaviour. Relatedly, Büyükboyacı (2014) investigates how information about the risk attitude of an opponent affects own choice in a stag-hunt game and finds that although a subject's propensity to choose the risky action depends on the opponent's risk attitude, it does not depend on own risk attitude.³

In our paper, we measure economic behaviour by considering two simple experimental games which have played a major role in the social preferences literature: a modified dictator game and a public goods game. We use a within-subjects design whereby individuals participate in both games, after having completed the Big Five questionnaire at the beginning of a session. Our treatments vary depending on whether information on the personality trait of agreeableness is provided to subjects or not. In the baseline no-information treatment, subjects do not receive any such information. In the remaining three treatments, subjects are informed of their own and other's degree of agreeableness (high vs. low agreeableness within their session) and are paired with another subject who has been classified to have the same and/or different degree of agreeableness (depending on the treatment). Following this, we observe behaviour in each of our two games.

Our findings provide new evidence that revealing personality information for disagreeable individuals has detrimental effects on their pro-social behaviour as compared to the baseline no-information benchmark. This is not the case, however, for agreeable individuals when they are matched with agreeable individuals. When agreeable individuals are matched with disagreeable individuals, we find that the former become less pro-social when personality data are revealed compared to when such data are not revealed. Therefore, in homogeneous pairs with known personality types, higher agreeableness favours pro-social behaviour as opposed to lower agreeableness; however, this is not the case for heterogeneous pairs with known personality types where agreeable individuals' behaviour is also affected negatively.

Our paper broadens the existing literature in at least three important respects. First, from a theoretical perspective, in line with existing studies, we show that personality traits, and in particular, agreeableness, affect economic preferences such as advantageous inequality in a modified dictator game and cooperative behaviour in a public goods game. Our findings

³ In a separate paper, Büyükboyacı and Küçükşenel (2017) use communication (by sending costly or costless messages) as a means of endogenously determining whether subjects share their risk-taking history in a stag-hunt game. Their findings suggest that such messages have an impact on actual risk taking.

thus provide further evidence for inspiring development of theories that can integrate the relationship between economic preferences and psychological measures of personality.

Second, we add to the literature which has provided evidence that context matters for economic behaviour (for example, social cues such as smiles and handshakes for coordination in a weak-link game; see Manzini et al., 2009). We show that another factor that is relevant for economic behaviour is providing personality information. Our findings indicate that subjects' behaviour is sensitive to the provision of information on personality data. Overall, our study demonstrates that individuals attach significant informational value when they are aware of their own and others' personality types as observed in their behaviour.

Third, our results have implications in relation to hiring processes which involve a personality assessment of their employees. Many modern corporations perform such assessments when it comes to hiring, which are treated confidentially on the basis of data security norms. In our experiment, the only publicity of information relevant for behaviour is the one affecting pro-sociality negatively. Thus, since providing information about personality traits has detrimental effects on pro-social behaviour, a clear cut recommendation in favour of the confidentiality of personality information stems from our findings.

Our paper is organised as follows. Section 2 outlines the experimental design and procedures. Section 3 presents the experimental results and Section 4 concludes.

2. Experimental design

2.1 Frameworks

In all treatments, each subject participated in two 2-player games: a modified dictator game and a public goods game. Each game was played exactly once by each subject. Each game was presented separately in a different section of the experiment. The parameters we have used for each game were adopted from Blanco et al. (2011). Below we provide the structure of each game in turn.

a) The Modified Dictator Game (MDG): In this game, the allocator has to decide about how much of the initial pie of £10 (if any) s/he is at most willing to sacrifice in order to achieve an equal distribution of payoffs. Subjects were given a list of 11 pairs of payoff vectors and had to choose one of the two payoff vectors in all 11 cases. The left payoff vector was always (£10, £0), that is, if the left payoff distribution column was chosen, the allocator would receive £10 and the recipient £0. The right payoff vector contained equal payoffs

varying from (£0, £0), (£1, £1) all the way up to (£10, £10). Each subject made a choice in the role of the allocator. If this game was randomly selected to be relevant for payment at the end of the experiment, one of the 11 payoff vector pairs was chosen at random and then the corresponding decision of the (randomly chosen) allocator determined the payments for the pair of subjects. The payoff table for the MDG that subjects faced in the experiment is presented in Table 1.

Table 1 – Payoff table for the MDG

Person A's Payoff	Person B's Payoff	Decision		Person A's Payoff	Person B's Payoff
10	0	Left	Right	0	0
10	0	Left	Right	1	1
10	0	Left	Right	2	2
10	0	Left	Right	3	3
10	0	Left	Right	4	4
10	0	Left	Right	5	5
10	0	Left	Right	6	6
10	0	Left	Right	7	7
10	0	Left	Right	8	8
10	0	Left	Right	9	9
10	0	Left	Right	10	10

Note: Payoffs are expressed in terms of GBP.

b) The Public Goods Game (PGG): In this game, both subjects in a group received a fixed endowment of £10 each. Each subject had to decide simultaneously how much (if any) money from their initial endowment to contribute to a public good. Each monetary unit that the individual keeps for her/himself raises her/his payoff by exactly that amount. Both subjects receive £0.7 for each £1 contributed to the public good (this corresponds to the marginal per capita return). If the public goods game was chosen at the end of the experiment, payoffs were calculated according to the contributions of the randomly paired players.

Subjects did not receive any feedback or payment until the end of the experimental session. All decisions were to be made without any information on other subjects' choices and without any communication. At the end of the session, one game was chosen randomly and subjects were paid according to their and their counterparts' decisions in this game. In the modified dictator game that involves different roles (Person A and Person B), these were determined randomly between the two subjects of each pair. Subjects knew about these

procedures in advance. At the end of the experiment they were informed about all the random draws and about the payment-relevant decisions. These design features rule out any potential income or other confounding effects that may be present across games when feedback is provided to subjects. As subjects participated in both games, we also controlled for potential order effects by switching the sequence with which the games were played. In half of the sessions, the modified dictator game preceded the public goods game and in the other half of the sessions, the order of the games was reversed.⁴

2.2 Experimental design

We conducted an experiment which consists of four treatments. Our treatments vary with respect to subjects' matching depending on whether they are classified as being agreeable or disagreeable. Prior to playing each of the two games described above, subjects had to respond to the set of questions provided in Appendix A. Based on their answers to this questionnaire, we calculated a within-session agreeableness score for each subject using standard practices in personality research (e.g., Costa et al., 1991). We then split the session in half, consisting of subjects who were characterised by high and low agreeableness.⁵ Subjects were then informed of their own and others' types⁶ and they were randomly assigned in one of the three following treatments: 1) "*High Agreeableness*" treatment whereby a pair of subjects consists of those who have both been classified as being types of high agreeableness; 2) "*Low Agreeableness*" treatment whereby a pair of subjects consists of those who have both been classified as being types of low agreeableness; and 3) "*Mixed Agreeableness*" treatment whereby a pair of subjects consists of one subject who has been

⁴ Our procedures were adopted from Blanco et al. (2011).

⁵ An alternative design option would be to collect information on personality traits in a pre-test experiment. However, this procedure would require that all pre-test subjects come back and participate in the main experiment. This is likely not to be the case and as a result, there could be self-selection effects for those taking part in the main study. To avoid this issue, we decided to do subjects' personality grouping at the session level. Using session-specific relative agreeableness scores is easier for subjects to interpret but also is closer to real-life working environments whereby inferences about others' personality are typically made within a given group in relative terms. Our subjects and, more generally, individuals in real-world situations would have no knowledge of the meaning of an absolute score of agreeableness and as a result, giving relative information on agreeableness scores is more informative. In our study, the median of the agreeableness score was 116, and even if we look at absolute agreeableness scores across the whole sample size, we find that only 8.19% (14/171) of our subjects who were "absolutely high" in agreeableness were classified as disagreeable and only 4.64% (7/151) of our subjects who were "absolutely low" in agreeableness were classified as agreeable. Table 3 provides summary statistics of agreeableness across treatments.

⁶ More specifically, subjects were informed whether they are below or above the average in their session. For example, in the "Low (High) Agreeableness" treatment, subjects received the following information on their screens: "Based on your answers, you have been classified according to the personality trait *agreeableness*: below (above) the average in this session. You have been matched with someone who has been classified *agreeableness*: below (above) the average in this session". Together with this information, subjects also received a brief definition of the term "agreeableness". An example screenshot showing how the personality information was revealed to subjects can be found in Appendix B.

classified as being type of high agreeableness and of another subject who has been classified as being type of low agreeableness.⁷ By examining behaviour in these treatments, we are able to evaluate whether pro-sociality as measured by observing decisions in the MDG and in the PGG is a function of agreeableness in the presence of personality data information.

We compare behaviour from these treatments with a control treatment, in which subjects did not receive any personality information about their own and other subjects' type, but had to respond to the same set of questions as in the other three treatments described above. This aspect of our design allows us to address our main research questions of whether revealing personality data impacts on subjects' pro-social behaviour.⁸ Note that when subjects completed the Big Five questionnaire, they were not aware of how this information could be used later in the experiment. Had subjects known that their personality measures will be used for being allocated to specific groups, they would have an obvious incentive to misrepresent their answers to the personality questionnaire. Our design feature minimises potential misrepresentations in their responses to the Big Five questionnaire (i.e., gaming the measures collected in the personality questionnaire) which are likely to result from subjects knowing that the personality information will be shared with others. Table 2 summarises an overview of our experimental design.

Table 2 – Overview of experimental design

Treatment	Sequence 1 (MDG → PGG)	Sequence 2 (PGG → MDG)	Total
Control	68	82	150
High Agreeableness	26	24	50
Low Agreeableness	28	26	54
Mixed Agreeableness	32	36	68

Note: In the cells, we report the number of subjects participating in each treatment.

Procedures: In total, 322 subjects participated in our experiment.⁹ All subjects were recruited at the University of Birmingham, using the ORSEE software (Greiner, 2015) and the

⁷ We have used an ex-post matching protocol based on the agreeableness score. The reason for doing so was to generate a relatively equal number of observations across personality types in each of the three conditions (“High Agreeableness”, “Low Agreeableness” and “Mixed Agreeableness”). Importantly, we note that our focus is on the causal impact that revealing information about personality data has on pro-social behaviour.

⁸ In our experiment, subjects respond to the Big Five questionnaire in both treatments (irrespective of whether personality information was revealed to them) before playing the games. This means any potential priming effects stemming from the responses to the questionnaire may be present in both of our treatments; however, we are concerned with the differential effects of revealing personality information on behaviour across treatments.

⁹ We provide detailed information about the number of sessions we conducted as well as the number of subjects participating in each session and sequence across treatments in Table C.1 (see Appendix C).

experiment was computerized and programmed with the software z-Tree (Fischbacher, 2007). At the end of each session, subjects were privately paid according to their decisions in one of the randomly selected game. Average earnings (including a show-up fee of £2.50) were £10.72. Sessions lasted 50 minutes, on average. Before subjects played the game, they received the instructions reproduced in Appendix D. Only when all participants had made their decisions in one game were the instructions for the following game distributed.

3. Results

We start our data analysis by presenting how agreeableness varied across our treatments, followed by an exploration of how subjects behaved, first, in the MDG and then, in the PGG. We note that average behaviour is not statistically significantly different based on the sequence with which the games were played (Wilcoxon ranksum test, $p=0.961$ for the MDG and $p=0.176$ for the PGG).¹⁰ We therefore pool the data across the two sequences for each game.

3.1. Levels of agreeableness

Before we examine our findings in relation to how subjects behaved in each of the four treatments, we provide an analysis of the personality trait in which we are interested. Table 3 reports summary statistics for the levels of agreeableness in each treatment separately. Across all four treatments, the minimum value of agreeableness is 45 and its maximum value is 172 (with the median value being 116). On average, the level of agreeableness is highest in the “High Agreeableness” treatment (128.58) and lowest in the “Low Agreeableness” treatment (101.33).¹¹ The difference between these two treatments is statistically significant (Wilcoxon ranksum test, $p<0.001$). The average level of agreeableness lies in between the other two treatments (“Control”¹² and “Mixed Agreeableness”) and the difference between these two treatments is not significant (Wilcoxon ranksum test, $p=0.693$). However, when we compare either the “Control” or the “Mixed Agreeableness” treatments with the “High Agreeableness” treatment, we find significant differences in average agreeableness (for both pair-wise

¹⁰ All p-values reported in this section correspond to two-sided tests. In our analyses, each subject is treated as our unit of independent observation.

¹¹ We find that the agreeableness scores of 11 subjects in the “Low Agreeableness” treatment falls above the minimum agreeableness score of subjects in the “High Agreeableness” treatments. When we rule these 11 subjects out, we obtain similar results in behaviour described in Sections 3.2 and 3.3.

¹² Looking at the Control treatment, we observe that the average level for high (low) agreeableness subjects is equal to 130.32 (102.78). When we compare the average agreeableness of the high agreeableness subjects in the Control with that of subjects in the “High Agreeableness” treatment (130.32 vs. 128.58), we find statistically insignificant differences ($p=0.674$). This is also the case when we compare the corresponding average agreeableness levels for the low agreeableness subjects (102.78 vs. 101.33, $p=0.334$).

comparisons, $p < 0.001$). Similarly, we observe significant differences when we compare the average level of agreeableness in the “Low Agreeableness” treatment with the corresponding levels in the “Control” or the “Mixed Agreeableness” treatments (for both pair-wise comparisons, $p < 0.001$).

Table 3 – Summary statistics of agreeableness across treatments

Treatment	Average Agreeableness	Standard deviation	Minimum value	Maximum value
Control	116.37	19.10	45	172
High Agreeableness	128.58	11.12	111	158
Low Agreeableness	101.33	12.47	65	116
Mixed Agreeableness	116.09	16.46	70	152

Following our analysis about the average level of agreeableness across treatments, we next examine whether there are any resulting differences in behaviour as measured in the modified dictator game and in the public goods game.

3.2 Behaviour in the modified dictator game

We first look at the average choices of the right payoff vector as shown in Table 1. We refer to this payoff vector as the “fair payoff vector” as it yields an equal allocation of money between two subjects in a pair. Figure 1 shows the average choice of the fair payoff vector relative to the unequal (selfish) option where the dictator keeps everything for him/herself. As we move towards the right hand side of the x-axis (i.e. from (£0, £0) to (£10, £10)), subjects increasingly choose the fair payoff distribution in all four treatments. However, we observe striking differences about the average choice of each fair payoff vector across treatments. For example, in almost all eleven payoff vector, subjects in the “High Agreeableness” treatment choose, on average, the fair option more frequently in relation to those in the “Low Agreeableness” and the “Mixed Agreeableness” treatments, for which average frequencies overlap across the eleven payoff vectors. As far as the “Control” treatment is concerned, subjects’ choices of the fair payoff distribution lie in between the other treatments.

Figure 1 – Average choice of the fair payoff distribution across treatments




Table 4 provides the average choice of fair payoff distribution across all eleven payoff vector pairs in each treatment separately. Recall that in the modified dictator game subjects had to decide between two options about how much of £10, they are willing to sacrifice in order to achieve an equal distribution of payoffs. The left payoff corresponds to an unequal payoff distribution (£10, £0), whereas, the right payoff distribution corresponds to equal payoff distributions starting from (£0, £0) up to (£10, £10). We therefore use the switch point from the left to the right payoff distribution as an indicator of subjects' fair behaviour (aversion to advantageous inequality). If a subject always chooses the left payoff distribution, the maximum average value is equal to 0. If a subject always chooses the right payoff distribution, the maximum average value is equal to 1. This implies that the lower (higher) the average value is, the more (less) selfish or unfair a subject is.

Table 4 – Summary statistics of behaviour in the MDG across treatments

Treatment	Average choice of the fair payoff distribution	Standard deviation
Control	0.52	0.24
High Agreeableness	0.57	0.28
Low Agreeableness	0.39	0.23
Mixed Agreeableness	0.39	0.25

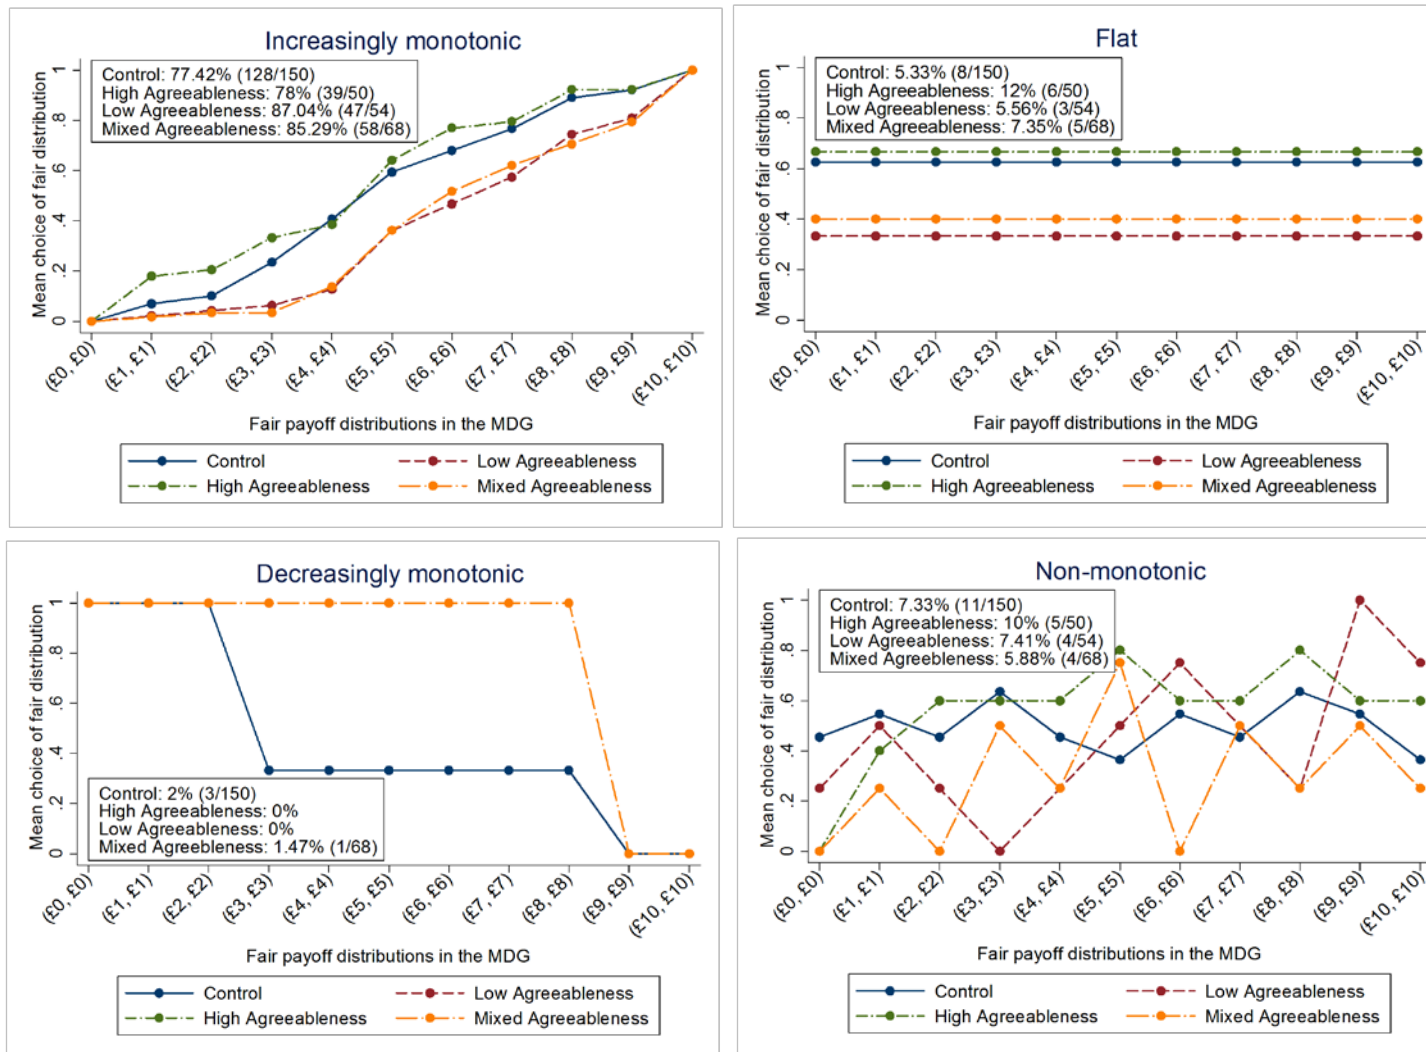
We observe that, in the “Control” treatment, the average choice of the fair payoff distribution is equal to 0.52, implying that subjects move from the left to the right payoff

distribution approximately in the middle of the allocation table where the right payoff distribution is equal to (£5, £5).¹³ In the other three treatments, the average choice of the fair distribution is either higher or lower than the corresponding value in the “Control” treatment. In particular, it is highest (0.57) in the “High Agreeableness” treatment and lowest (0.39) in the “Low Agreeableness” and “Mixed Agreeableness” treatments. By using a Wilcoxon ranksum test, we find that subjects in the “High Agreeableness” treatment are significantly more inequity averse than subjects in the “Low Agreeableness” treatment ($p=0.001$) and in the “Mixed Agreeableness” treatment ($p<0.001$), but not in the “Control” treatment ($p=0.235$). In relation to the “Low Agreeableness” treatment, we find that subjects behave similarly when compared to the “Mixed Agreeableness” ($p=0.967$) but they are less inequity averse when compared to the “Control” treatment ($p=0.001$). Finally, subjects in the “Mixed Agreeableness” treatment are less inequity averse than in the “Control” treatment ($p<0.001$).

To understand the observed pattern of allocations better, we provide further analysis on how agreeableness affects subjects’ choices at the individual level. We divide subjects into four categories: (1) subjects with no switching point (“flat” category), (2) subjects with one switching point from the left to the right payoff distribution (“increasingly monotonic” category), (3) subjects with one switching point from the right to the left payoff distribution (“decreasingly monotonic” category) and (4) subjects with more than one switching point (“non-monotonic” category). Figure 2 shows the average choice of the fair payoff vector relative to the selfish option in each of the four categories above across treatments.

¹³ This is similar to the finding by Blanco et al. (2011) who report that the average switching point was roughly (£11, £11). Note that in their experiment the maximum value of the right payoff distribution is equal to (£20, £20).

Figure 2 – Average choice of the fair payoff distribution in each category across treatments



Note: The numbers reported inside the plot area of each graph indicate the percentages (relative frequencies in parentheses) of subjects in each relevant category per treatment.

Our statistical analysis shows that the significant behaviour differences across treatments are stemming from those subjects who belong to the “increasingly monotonic” category. This is also the largest category within which subjects’ choices fall, accounting between 77% and 87% of subjects (depending on the treatment). In particular, we find that subjects in the “High Agreeableness” treatment are more inequity averse than those in the “Low Agreeableness” treatment ($p=0.001$) and in the “Mixed Agreeableness” treatment ($p=0.001$), but not in the “Control” treatment ($p=0.390$). For the “Low Agreeableness” treatment we observe significant differences when compared with the “Control” treatment ($p=0.001$), meaning that subjects are more inequity averse in the latter relative to the former treatment. This is not the case for the comparison between the “Low Agreeableness” treatment and the “Mixed Agreeableness” treatment ($p=0.925$). Finally, subjects in the “Mixed Agreeableness” treatment are less inequity averse than in the “Control” treatment ($p<0.001$). By looking at the average choices of the fair payoff distribution between all subjects (as reported in Table 4) and the increasingly monotonic subjects (as reported in Table 5), we observe very similar choices for a given treatment. This implies that the overall pattern we obtain is driven by the increasingly monotonic subjects.

Turning to the “Flat” category (where subjects had no switching point across all eleven payoff distributions), we find insignificant differences for any pairwise treatment comparisons ($p>0.371$). This is the case for the “non-monotonic” category ($p>0.146$), with the only exceptions being the comparison between the “High Agreeableness” and the “Mixed Agreeableness” treatment ($p=0.081$) and between the “Control” and the “Mixed Agreeableness” treatment ($p=0.068$) where we record weakly significant differences at the 10% level. However, we note that the number of observations in these two categories is quite low. The average choices of the fair payoff distribution of the subjects in each category separately across treatments are reported in Table 5. Our analysis indicates that subjects in the “High Agreeableness” treatment are more averse to advantageous inequality than subjects in the “Low Agreeableness” treatment.

Table 5 – Average choice of the fair distribution in the MDG for each category across treatments

Treatment	Flat	Increasingly monotonic	Decreasingly monotonic	Non-monotonic
Control	0.63 (0.52)	0.51 (0.22)	0.45 (0.31)	0.50 (0.17)
High Agreeableness	0.67 (0.52)	0.56 (0.24)	–	0.56 (0.17)
Low Agreeableness	0.33 (0.58)	0.38 (0.21)	–	0.45 (0.17)
Mixed Agreeableness	0.40 (0.55)	0.38 (0.21)	0.82 (–)	0.30 (0.19)

Note: Numbers in parentheses indicate standard deviations.

We next examine whether the observed behavioural differences in the modified dictator game are affected when information about personality data is available (that is, when subjects are informed of their own and their counterpart’s personality type). To do so, we analyse how pairs of subjects behaved in the “Control” treatment and compare them with the behaviour of the respective subjects in the other three treatments. Recall that in the “Control” treatment, subjects can still be classified as being agreeable or disagreeable (low vs. high agreeableness) as they also responded to the Big Five questionnaire. However, they were not provided with such information, as opposed to the main three treatments. Table 6 disentangles each type of pair of subjects, showing the average value of fair payoff distribution in each sub-category for the “Control” treatment.

Table 6 – Summary statistics of behaviour in the MDG for the “Control” treatment

Control Treatment	Average choice of the fair payoff distribution	Standard Deviation	Obs.
High Agreeableness pairs	0.57	0.21	44
Low Agreeableness pairs	0.49	0.22	46
Mixed Agreeableness pairs	0.50	0.28	60

We observe that the ordering of average choices of the fair payoff distribution is similar to our findings from Table 4. In particular, the average choice of the fair payoff distribution is highest for the high agreeableness subjects (0.57) and lowest for the low agreeableness subjects (0.49). The average choice of the fair payoff distribution of the mixed agreeableness subjects (0.50) is only slightly higher than that of the low agreeableness subjects. Using a Wilcoxon ranksum test, we find that, within the “Control” treatment, differences in behaviour are statistically significant when we compare high agreeableness subjects with low agreeableness subjects ($p=0.057$), but there are not statistically significant

for the other two pairwise comparisons (high agreeableness subjects vs. mixed agreeableness subjects, $p=0.336$; low agreeableness subjects vs. mixed agreeableness subjects, $p=0.457$).

We further compare the behaviour each sub-category of subjects in the “Control” treatment with the behaviour of subjects in the corresponding main treatments. Our findings indicate that subjects become significantly more unfair when they are informed of their own and their co-player’s personality type with respect to both low agreeableness pairs (0.49 for the low agreeableness pairs in the “Control” treatment vs. 0.39 in the “Low Agreeableness” treatment, $p=0.037$) and mixed agreeableness pairs (0.50 for the mixed agreeableness subjects in the “Control” treatment vs. 0.39 in the “Mixed Agreeableness” treatment, $p=0.006$). For the high agreeableness pairs, we find that subjects behave similarly irrespective of whether they were provided with relevant information about the personality type of others and themselves ($p=0.909$).

Looking at the mixed agreeableness pairs, we find that in the “Control” treatment the disagreeable subjects’ average choice of the fair payoff distribution in the MDG is 0.51; while, the agreeable subjects’ average choice is 0.50. This difference is not statistically significant ($p=0.946$). This is also the case when we compare disagreeable with agreeable subjects’ average choice of the fair payoff distribution in the “Mixed Agreeableness” treatment (0.38 vs. 0.39, respectively; $p=1.000$). We next test whether revealing personality data affects behaviour in the MDG for agreeable and disagreeable subjects when they are matched together and find that both disagreeable and agreeable subjects become less inequity averse (for disagreeable subjects: 0.51 vs. 0.38; $p=0.056$ and for agreeable subjects: 0.50 vs. 0.39; $p=0.041$).

Result 1: *In the MDG, revealing personality data does not affect the behaviour of agreeable pairs as opposed to disagreeable pairs, who become less inequity averse in the presence of personality data information. For mixed agreeableness pairs, we find that agreeable and disagreeable subjects become less inequity averse when personality data are revealed.*

As a robustness check, we test whether our main findings hold when we take into account subjects’ observable characteristics. We have collected data on subjects’ gender, nationality and field of study. Tables C.2-C.4 in Appendix C provide summary statistics (and p-values from Wilcoxon ranksum tests) of subjects’ behaviour in the MDG, split by men vs. women (Table C.2), UK vs. non-UK (Table C.3) and economists vs. non-economists (Table C.4). We

find that revealing or not personality data does not change significantly the behaviour of the high agreeableness subjects for all pairwise comparisons ($p > 0.592$). However, we find significant differences for the behaviour of disagreeable pairs in that revealing personality data makes them less inequity averse for all pairwise comparisons ($p < 0.068$), with the only exceptions being that of men ($p = 0.222$), UK subjects ($p = 0.130$) and non-economists ($p = 0.101$). We also find significant differences for the behaviour of mixed agreeableness pairs in that revealing personality data makes subjects less inequity averse for all pairwise comparisons ($p < 0.047$), with the only exceptions being that of men ($p = 0.432$) and economists ($p = 0.186$).

3.3 Behaviour in the public goods game

In this section, we analyse subjects' cooperative behaviour in the public goods game across treatments. In this game, subjects were given £10 and were asked to make a contribution decision towards the common resource. Figure 3 shows the distribution of the absolute levels of contributions in each of the four treatments, separately. We observe that the majority of subjects (53.70%) in the "Low Agreeableness" treatment are full free-riders (i.e. contribute £0 to the public good). For the "Mixed Agreeableness" treatment, 42.65% of the subjects are free-riders.¹⁴ This percentage dramatically drops to 18% in the "High Agreeableness" treatment. In the "Control" treatment, we find that the percentage of free-riders is equal to 25.33%. A non-parametric Kruskal-Wallis test suggests significant differences among treatments ($\chi^2(3) = 22.679$, $p < 0.001$).

¹⁴ When we look separately for the subjects with low agreeableness and those with high agreeableness in this treatment, we find that the percentage of full free riders is similar in these two categories: 41.48% and 44.12%, respectively.

Figure 3 – Distribution of contribution levels across treatments




Table 7 shows the average contribution behaviour of subjects across all treatments. The average contribution is largest in the “High Agreeableness” treatment (i.e. £4.50) and smallest in the “Low Agreeableness” treatment (i.e. £2.19). In the other two treatments, average contributions lie within this range: £4.39 for the “Control” treatment¹⁵ and £3.09 for the “Mixed Agreeableness” treatment. Using a Wilcoxon rank-sum test, we find that subjects in the “High Agreeableness” treatment contribute significantly more compared to those in the “Low Agreeableness” treatment ($p < 0.001$), in the “Mixed Agreeableness” treatment ($p = 0.010$).¹⁶ Similarly, subjects contribute significantly less in the “Low Agreeableness” treatment compared to the “Control” treatment ($p < 0.001$) but not to the “Mixed Agreeableness” treatment ($p = 0.150$). Our analysis indicates that subjects in the “High Agreeableness” treatment are more cooperative than subjects in the “Low Agreeableness” treatment.

¹⁵ We find similar average contributions compared to Blanco et al. (2011) who report that subjects contributed 47% of the endowment in their experiment.

¹⁶ The difference between the “High Agreeableness” treatment and the “Control” treatment is not statistically significant, although subjects contribute more in the former than in the latter treatment.

Table 7 – Summary statistics of behaviour in the PGG across treatments

Treatment	Average contribution	Standard Deviation
Control	4.39	3.68
High Agreeableness	4.50	3.37
Low Agreeableness	2.19	3.14
Mixed Agreeableness	3.09	3.57

Following this analysis, we next examine how subjects' behaviour varies depending on whether they receive information about their own and the other player's personality type. We perform a similar exercise as in Section 4.2 and first look at how pairs of subjects behaved in the "Control" treatment. We then compare them with the respective behaviour of those subjects in the three main treatments whereby subjects were provided with information with their personality type. Table 8 reports the average contribution behaviour for each sub-category of subjects in the "Control" treatment.

Table 8 – Summary statistics of behaviour in the PGG for the "Control" treatment

Control Treatment	Average contribution	Standard Deviation	Obs.
High Agreeableness pairs	4.91	3.76	44
Low Agreeableness pairs	3.91	3.37	46
Mixed Agreeableness pairs	4.37	3.85	60

In terms of the ordering of average contribution, we observe a similar picture relative to our findings in Table 7. In particular, the average contribution is highest for the high agreeableness subjects (£4.91) and lowest for the low agreeableness subjects (£3.91). The average contribution for the mixed agreeableness subjects (£4.37) lies within these two values. Using a Wilcoxon ranksum test, we find insignificant differences in contribution behaviour for any pairwise comparison in each of the three categories within the "Control" treatment ($p > 0.2250$).

We further compare the behaviour for each sub-category of subjects in the "Control" treatment with the behaviour of respective subjects in the three relevant treatments. Our findings indicate that subjects become significantly less cooperative when they are informed of their own and their co-player's personality type with respect to both the low agreeableness subjects (£3.91 for the low agreeableness subjects in the "Control" treatment vs. £2.19 in the "Low Agreeableness" treatment, $p = 0.005$) and the mixed agreeableness subjects (£4.37 for the mixed agreeableness subjects in the "Control" treatment vs. £3.09 in the "Mixed Agreeableness" treatment, $p = 0.044$). For high agreeableness pairs, we find that subjects

behave similarly irrespective of whether they were provided with relevant information about their personality type ($p=0.652$). These findings are similar to our analysis from the modified dictator game.

Looking at the mixed agreeableness pairs, we find that in the “Control” treatment disagreeable subjects contribute £3.83; while, agreeable subjects contribute £4.9. This difference is not statistically significant ($p=0.241$). In the “Mixed Agreeableness” treatment, where personality information is revealed, we find that disagreeable subjects contribute £3.35; while, agreeable subjects contribute £2.82. This difference is not statistically significant ($p=0.497$). We next test whether revealing personality data affects the cooperative behaviour of agreeable and disagreeable subjects when they are matched together and find that while disagreeable subjects’ behaviour does not differ across treatments (£3.35 vs. £3.83; $p=0.585$), agreeable subjects become less cooperative when matched with disagreeable subjects and are aware of this (£2.82 vs. £4.9; $p=0.027$).

Taken together, our findings indicate that low agreeableness subjects are sensitive to the information about the personality type of the subjects they are matched with. Specifically, in homogeneous pairs with known personality types, both in the modified dictator game and in the public goods game, disagreeable subjects become less pro-social. However, providing information on personality does not change significantly the behaviour of the high agreeableness subjects, with the only exception being when they are matched with disagreeable subjects and this is known. This is the case for both games.

Result 2: *In the PGG, revealing personality data does not affect the behaviour of agreeable pairs as opposed to disagreeable pairs, who become less cooperative in the presence of personality data information. For mixed agreeableness pairs, we find that agreeable subjects become less cooperative when personality data are revealed.*

As a robustness check, we test whether our main findings hold when we take into account subjects’ observable characteristics. Similar to our earlier analysis in the MDG, we report in Tables C.5-C.7 in Appendix C summary statistics (and p-values from Wilcoxon ranksum tests) of subjects’ behaviour in the PGG, split by men vs. women (Table C.5), UK vs. non-UK (Table C.6) and economists vs. non-economists (Table C.7). We find that revealing or not personality data does not change significantly the behaviour of the high agreeableness subjects for all pairwise comparisons ($p>0.286$). However, we find significant differences for

the behaviour of disagreeable pairs in that revealing personality data makes them less cooperative for all pairwise comparisons ($p < 0.069$), with the only exceptions being that of non-UK subjects ($p = 0.168$) and economists ($p = 0.128$). We also find significant differences for the behaviour of mixed agreeableness pairs in that revealing personality data makes subjects less cooperative for all pairwise comparisons ($p < 0.069$), with the only exceptions being that of men ($p = 0.570$), non-UK subjects ($p = 0.154$) and economists ($p = 0.284$).

As a final step for our analysis, we explore whether and if so, how subjects' behaviour correlate in both games. By performing a Pearson correlation test, we find that, in all four treatments, subjects' choices in the MDG and the PGG are significantly and positively correlated ($p < 0.007$). This indicates that those subjects who are more inequity averse in the MDG are also more cooperative in the PGG. Our results hold in all four treatments, implying that revealing or not personality information does not play a key role in determining this relationship.

4. Conclusions

Humans are biological organisms with significant individual variation stemming from their personality traits. As a result, their personality traits are likely determinants of economic behaviour. While there is research recognising the importance of personality in shaping our behaviour and motives, we know little about the informational value that individuals attach to personality traits in terms of their pro-social behaviour: does revealing data on personality traits influence such behaviour? This paper presents an experimental investigation which fills this gap in the literature. In particular, we measure pro-social behaviour using as frameworks two simple experimental games: the modified dictator game and the public goods game. Based on an established psychological literature, we use the Big Five framework to measure subjects' personality traits. The main focus of our study is to understand the behavioural effects of agreeableness. Our findings provide new evidence that revealing personality data plays a key role in shaping subjects' social preferences. We find a positive effect of agreeableness on pro-social behaviour: "high agreeableness" pairs are more averse to advantageous inequality and more cooperative in relation to "low agreeable" pairs. More importantly, we show that revealing information about personality data affects this relationship. While information about personality types does not shift the social preferences of agreeable individuals when matched together, this is not the case for disagreeable ones. Specifically, we find that disagreeable individuals become less inequity averse in the modified dictator game and less cooperative in the public goods game when they are

themselves disagreeable and are informed that others are of the same personality types.¹⁷ We also find that agreeable subjects become less pro-social when they are matched together with disagreeable subjects and are informed of this. Overall, this provides evidence that cues on personality traits embody significant informational value, which may have detrimental effects on pro-sociality.

Our study contributes to the existing literature by highlighting the need to develop theories which account for personality differences in economic preferences. We provide additional evidence for inspiring theory development which may incorporate the role of personality traits into economic analysis. Furthermore, we show that individuals attach negative informational value when information about own and other's personality is available, leading them to behave in a more selfish way (relative to when such information is not available). Our findings, thus, reinforce existing policies relating to hiring processes which treat data on employees' personality assessment confidentially. Since our findings indicate that revealing information about personality traits alters individuals' pro-social behaviour in a negative manner, our implication from our experiment is in favour of such processes.

Our results are therefore relevant regarding the use of personality traits in hiring processes, but mainly on the way personality information should be handled by employers and managers within a company. For example, pairing employees whose agreeableness score is higher will lead to a more cooperative corporate environment, but informing employees about this will not have any further positive effect, because knowledge of one's own and others' high agreeableness scores does not affect cooperative behaviour. On the contrary, providing information about disagreeableness has detrimental effects on behaviour as it affects both disagreeable individuals but also agreeable individuals, who become less pro-social when matched with disagreeable individuals. In fact, our results highlight the crucial role of revealing personality information – especially, revelation of disagreeableness which impacts negatively on behaviour – and as a result, personality data should be processed with care by employers.

Our paper also gives rise to future research avenues. First, it is important to examine whether other personality traits (as measured by the Big Five model) are relevant determinants of social preferences. This will provide a more comprehensive picture of the sensitivity of social preferences on personality. Second, the persistence of personality effects

¹⁷ The asymmetric effects observed in subjects' behaviour in the "High Agreeableness" versus the "Low Agreeableness" treatments also provide evidence against potential priming effects in our study.

on economic behaviour is a separate issue that needs to be explored further. Third, the role of personality in shaping other aspects of social preferences is necessary in order for a unified framework to be established in economic thinking. While additional research is needed to understand the interplay between other-regarding preferences and measures of personality, we hope that our study further bridges the gap between personality psychology and behavioural economics by providing evidence that personality cues are an important determinant of economic preferences.

References

- Almlund, M., Duckworth, A. L., Heckman, J. J., and Kautz, T. (2011). Personality psychology and economics. In E. A. Hanushek, S. Machin, & L. Wößmann (Eds.), *Handbook of the economics of education* Amsterdam, Elsevier.
- Ashton, M. C., Lee, K., Perugini, M., Szarota, P., deVries, R.E., Di Blas, L., Boies, K., and De Raad, B., (2004). A Six-Factor Structure of Personality-Descriptive Adjectives: Solutions from Psycholexical Studies in Seven Languages, *Journal of Personality and Social Psychology*, 86: 356-366.
- Autor, D.H., Scarborough, D. (2008). Does job testing harm minority workers? Evidence from retail establishments. *Quarterly Journal of Economics*, 123: 219–277.
- Barrick, M. R. and Mount, M. K. (1991), The Big Five personality dimensions and job performance: A meta- analysis. *Personnel Psychology*, 44: 1-26.
- Becker, A., Deckers, T., Dohmen, T., Kosse, F., and Falk, A. (2012). The Relationship Between Economic Preferences and Psychological Personality Measures, *Annual Review of Economics*, 4: 453-478.
- Bekkers, R., (2006). Traditional and Health-Related Philanthropy: The Role of Resources and Personality, *Social Psychology Quarterly*, 69: 349-366.
- Ben-Ner, A., Putterman, L., Kong, F., Magan, D., (2004a). Reciprocity in a two-part dictator game, *Journal of Economic Behavior & Organization*, 53:333-352.
- Ben-Ner, A., Kong, F., Putterman, L. (2004b). Share and share alike? Gender-pairing, personality, and cognitive ability as determinants of giving, *Journal of Economic Psychology* 25:581–589.
- Blanco, M., Engelmann, D., and Normann, H. T. (2011). A within-subject analysis of other-regarding preferences. *Games and Economic Behavior*, 72, 321-338.
- Büyükboyacı, M. (2014). Risk attitudes and the stag-hunt game, *Economics Letters*, 124:323-325.
- Büyükboyacı, M. and Küçükşenel, S. (2017). Costly Pre-Play Communication and Coordination in Stag-Hunt Games, *Managerial and Decision Economics*, 38:845-856.

- Campion, M. A., Palmer, D. K., and Campion, J. E. (1997). A review of structure in the selection interview. *Personnel Psychology*, 50:655–702.
- Costa, P., and McCrae, R., (1992). Revised NEO Personality Inventory (NEO PI-R) and NEO Five-Factor Inventory (NEO-FFI). Odessa, FL, Psychological Assessment Resources.
- Costa, P. T., McCrae, R., and Dye, D. A. (1991). Domains and facets scales for agreeableness and conscientiousness: a revision of the NEO personality inventory. *Personality and Individual Differences*, 12, 887-898.
- Dohmen, T., Falk, A., Huffman, D., and Sunde, U., (2008). Representative Trust and Reciprocity: Prevalence and Determinants, *Economic Inquiry*, 46: 84-90.
- Dohmen, T., Falk, A., Huffman, D., and Sunde, U., (2010). Are Risk Aversion and Impatience Related to Cognitive Ability? *American Economic Review*, 100: 1238-1260.
- Fischbacher, U. (2007). z-Tree: Zurich Toolbox for Ready-made Economic Experiments, *Experimental Economics*, 10: 171-178.
- Goldberg, L. R. (1981). Language and individual differences: The search for universals in personality lexicons. In L. Wheeler (Ed.), *Review of personality and social psychology*, (Vol. 2, pp. 141-165). Beverly Hills, CA: Sage.
- Goldberg, L. R. (1990). An alternative "description of personality": The Big-Five factor structure. *Journal of Personality and Social Psychology*, 59:1216-1229.
- Greiner, B. (2015). Subject pool recruitment procedures: organizing experiments with ORSEE, *Journal of the Economic Science Association*, 1:114-125.
- Manzini, P., Sadrieh, A., Vriend, N.J., 2009. On smiles, winks and handshakes as coordination devices. *Economic Journal*, 119, 826–854.
- Osinski, J., (2009). Kin Altruism, Reciprocal Altruism and Social Discounting. *Personality and Individual Differences*, 47: 374-378.
- Roberts, B. W., (2006). Personality Development and Organizational Behavior. In: Staw, B., 27, (Ed.), *Research on Organizational Behavior*. Elsevier Science / JAI Press, Oxford, 1-41.
- Sackett, P. R. and Walmsley, P. T. (2014). Which Personality Attributes Are Most Important in the Workplace?, *Perspectives on Psychological Science*, 9:538 –551.

- Vinchur, A. J., Schippmann, J. S., Switzer, F. S., and Roth, P. L. (1998). A meta-analytic review of predictors of job performance for salespeople. *Journal of Applied Psychology*, 83, 586–597.
- Volk, S., Thöni, C., and Ruigrok, W., (2011). Personality, personal values and cooperation preferences in public goods games: A longitudinal study, *Personality and Individual Differences*, 50: 810-815.

Supplementary material (not intended for publication)

Appendix A – Additional Material

[Note: Onscreen instructions for the Big Five questionnaire which subjects had to respond at the beginning of the experiment]

The questionnaire which follows contains statements concerning your perception about yourself in a variety of situations. Your task is to indicate the strength of your agreement with each statement, utilizing a scale consisting of five radio buttons, in which the top button denotes strong disagreement, the bottom button denotes strong agreement, and the other three buttons represent intermediate judgments.

I tend to be cynical and sceptical concerning others' intentions.

I am not very crafty or pretender.

Some people think I am selfish and egocentric.

I would prefer to cooperate with others rather than to compete with them.

I do not mind to boast of my talents and achievements.

Political leaders should be more conscious of the human side of their policies.

I believe that the majority of people have in general good intentions.

If it is necessary, I am prepared to manipulate people in order to get what I want.

I try to be nice to everyone I know.

I can be sarcastic and scathing if necessary.

I prefer not to talk about myself and my success.

I am tough and hard-headed in my behaviour.

I believe that the majority of the people would take advantage of someone who permits it.

Even if I wanted, I would not be able to cheat anyone.

Some people think that I am cold and calculating.

I find it difficult to express my anger even when I am right.

I am better than the majority of people and I am sure about it.

No matter how much we do for poor and old people, it will never be too much.

I think that the majority of people I interact with are honest and trustworthy.

Being absolutely honest is not good for doing business.

Generally speaking, I try to think of others and be considerate.

If I do not like someone, I tell them.

I try to be humble.

Beggars do not inspire me any sympathy.

If someone is nice to me, I become suspicious.

I would hate it if someone considered me a hypocrite.

I am not known for my generosity.

When someone offends me, I try to forgive and forget.

I have a great opinion of myself.

Human needs should be always above economic considerations.

My first reaction is to trust others.

Sometimes, I trick others in order to make them do what I want.

The majority of people I know like me.

If someone starts a fight with me, I am prepared to fight back.

I think I am not better than others independent of their condition.

I believe that all human beings deserve respect.

I tend to think the best of others.

Sometimes I intimidate or flatter others in order to make them do what I want.

I consider myself as a charitable person.

I am stubborn.

I would rather praise others than to be praised by them.

I am sympathetic with people who have been less lucky than me.

I have a lot of faith in human nature.

I am proud of my craftiness in dealing with people.

I give way to others when I can.

Frequently, I argue with my family and colleagues.

I feel superior.

I would prefer to be known as a merciful rather than as a strict person.

Appendix B – Screenshot about the personality information revelation



Appendix C – Additional Analyses

Table C.1 – Number of sessions and subjects participating in each session and sequence across treatments

Treatments	Number of subjects per session (Sequence 1)	Number of subjects per session (Sequence 2)	Total number of observations
Control	S1: 14 S2: 14 S3: 20 S4: 20	S1: 16 S2: 18 S3: 20 S4: 12 S5: 16	150
High Agreeableness	S1: 4 S2: 6 S3: 4 S4: 6 S5: 6	S1: 4 S2: 6 S3: 6 S4: 4 S5: 4	50
Low Agreeableness	S1: 4 S2: 6 S3: 6 S4: 6 S5: 6	S1: 6 S2: 6 S3: 6 S4: 4 S5: 4	54
Mixed Agreeableness	S1: 6 S2: 6 S3: 8 S4: 6 S5: 6	S1: 8 S2: 6 S3: 8 S4: 8 S5: 6	68

Note: In the “High Agreeableness”, “Low Agreeableness” and “Mixed Agreeableness” treatments, the total number of observations in Sequence 1 (Sequence 2) per session are as follows: S1: 14 (18); S2: 18 (18); S3: 18 (20); S4: 18 (16); and S5: 18 (14).

Table C.2 – Summary statistics of behaviour in the MDG across treatments, split by gender

Treatment	Average choice of the fair payoff distribution of men	Average choice of the fair payoff distribution of women
Control	0.47 (0.24)	0.58 (0.24)
High Agreeableness	0.55 (0.19)	0.60 (0.23)
Low Agreeableness pairs	0.45 (0.21)	0.52 (0.23)
Mixed Agreeableness pairs	0.43 (0.29)	0.61 (0.24)
High Agreeableness	0.47 (0.25)	0.64 (0.28)
Low Agreeableness	0.37 (0.25)	0.40 (0.24)
Mixed Agreeableness	0.36 (0.26)	0.42 (0.23)

Note: Standard deviations are reported in parentheses.

P-values from Wilcoxon ranksum tests:

- For men: High agreeableness pairs vs. High Agreeableness treatment: $p = 0.592$; Low agreeableness pairs vs. Low Agreeableness treatment: $p = 0.222$; Mixed agreeableness pairs vs. Mixed Agreeableness treatment: $p = 0.432$.
- For women: High agreeableness pairs vs. High Agreeableness treatment: $p = 0.725$; Low agreeableness pairs vs. Low Agreeableness treatment: $p = 0.029$; Mixed agreeableness pairs vs. Mixed Agreeableness treatment: $p = 0.001$.

Table C.3 – Summary statistics of behaviour in the MDG across treatments, , split by nationality

Treatment	Average choice of the fair payoff distribution of UK subjects	Average choice of the fair payoff distribution of non-UK subjects
Control	0.50 (0.25)	0.54 (0.23)
High Agreeableness	0.56 (0.23)	0.58 (0.19)
Low Agreeableness pairs	0.47 (0.22)	0.51 (0.23)
Mixed Agreeableness pairs	0.49 (0.28)	0.53 (0.27)
High Agreeableness	0.60 (0.28)	0.53 (0.28)
Low Agreeableness	0.35 (0.19)	0.43 (0.27)
Mixed Agreeableness	0.35 (0.25)	0.43 (0.23)

Note: Standard deviations are reported in parentheses.

P-values from Wilcoxon ranksum tests:

- For UK subjects: High agreeableness pairs vs. High Agreeableness treatment: $p = 0.858$; Low agreeableness pairs vs. Low Agreeableness treatment: $p = 0.130$; Mixed agreeableness pairs vs. Mixed Agreeableness treatment: $p = 0.047$.
- For non-UK subjects: High agreeableness pairs vs. High Agreeableness treatment: $p = 1.000$; Low agreeableness pairs vs. Low Agreeableness treatment: $p = 0.066$; Mixed agreeableness pairs vs. Mixed Agreeableness treatment: $p = 0.012$.

Table C.4 – Summary statistics of behaviour in the MDG across treatment, split by field of study

Treatment	Average choice of the fair payoff distribution of economists	Average choice of the fair payoff distribution of non-economists
Control	0.52 (0.27)	0.52 (0.24)
High Agreeableness pairs	0.58 (0.11)	0.57 (0.22)
Low Agreeableness pairs	0.55 (0.26)	0.46 (0.21)
Mixed Agreeableness pairs	0.49 (0.31)	0.51 (0.27)
High Agreeableness	0.57 (0.29)	0.57 (0.27)
Low Agreeableness	0.34 (0.22)	0.46 (0.23)
Mixed Agreeableness	0.39 (0.28)	0.38 (0.22)

Note: Standard deviations are reported in parentheses.

P-values from Wilcoxon ranksum tests:

- For economists: High agreeableness pairs vs. High Agreeableness treatment: $p = 0.862$; Low agreeableness pairs vs. Low Agreeableness treatment: $p = 0.068$; Mixed agreeableness pairs vs. Mixed Agreeableness treatment: $p = 0.186$.
- For non-economists: High agreeableness pairs vs. High Agreeableness treatment: $p = 0.855$; Low agreeableness pairs vs. Low Agreeableness treatment: $p = 0.101$; Mixed agreeableness pairs vs. Mixed Agreeableness treatment: $p = 0.006$.

Table C.5 – Summary statistics of behaviour in the PGG across treatments, split by gender

Treatment	Average Contribution of men	Average Contribution of women
Control	4.10 (3.99)	4.74 (3.25)
High Agreeableness pairs	5.05 (4.11)	4.77 (3.46)
Low Agreeableness pairs	3.60 (3.5)	4.29 (3.24)
Mixed Agreeableness pairs	3.86 (4.25)	5.08 (3.15)
High Agreeableness	3.95 (3.24)	4.87 (3.46)
Low Agreeableness	1.62 (2.62)	2.84 (3.60)
Mixed Agreeableness	3.11 (3.98)	3.06 (3.14)

Note: Standard deviations are reported in parentheses.

P-values from Wilcoxon ranksum tests:

- For men: High agreeableness pairs vs. High Agreeableness treatment: $p = 0.679$; Low agreeableness pairs vs. Low Agreeableness treatment: $p = 0.069$; Mixed agreeableness pairs vs. Mixed Agreeableness treatment: $p = 0.570$.
- For women: High agreeableness pairs vs. High Agreeableness treatment: $p = 0.753$; Low agreeableness pairs vs. Low Agreeableness treatment: $p = 0.004$; Mixed agreeableness pairs vs. Mixed Agreeableness treatment: $p = 0.003$.

Table C.6 – Summary statistics of behaviour in the PGG across treatments, split by nationality

Treatment	Average Contribution of UK subjects	Average Contribution of non-UK subjects
Control	4.79 (3.77)	3.93 (3.53)
High Agreeableness pairs	5.35 (4.11)	4.43 (3.37)
Low Agreeableness pairs	4.70 (3.38)	3.13 (3.24)
Mixed Agreeableness pairs	4.47 (3.86)	4.23 (3.70)
High Agreeableness	4.59 (3.49)	4.33 (3.24)
Low Agreeableness	1.32 (2.47)	3.12 (3.56)
Mixed Agreeableness	3.08 (3.77)	3.10 (3.35)

Note: Standard deviations are reported in parentheses.

P-values from Wilcoxon ranksum tests:

- For UK subjects: High agreeableness pairs vs. High Agreeableness treatment: $p = 0.446$; Low agreeableness pairs vs. Low Agreeableness treatment: $p = 0.001$; Mixed agreeableness pairs vs. Mixed Agreeableness treatment: $p = 0.069$.
- For non-UK subjects: High agreeableness pairs vs. High Agreeableness treatment: $p = 0.954$; Low agreeableness pairs vs. Low Agreeableness treatment: $p = 0.168$; Mixed agreeableness pairs vs. Mixed Agreeableness treatment: $p = 0.154$.

Table C.7 – Summary statistics of behaviour in the PGG across treatments, split by field of study

Treatment	Average Contribution of economists	Average Contribution of non- economists
Control	4.06 (3.41)	4.49 (3.76)
High Agreeableness pairs	6 (2.90)	4.74 (3.88)
Low Agreeableness pairs	3 (2.45)	4.24 (3.61)
Mixed Agreeableness pairs	4.12 (3.97)	4.47 (3.84)
High Agreeableness	3.94 (2.49)	4.76 (3.72)
Low Agreeableness	2.69 (3.41)	1.45 (2.61)
Mixed Agreeableness	2.55 (3.34)	3.49 (3.72)

Note: Standard deviations are reported in parentheses.

P-values from Wilcoxon ranksum tests:

- For economists: High agreeableness pairs vs. High Agreeableness treatment: $p = 0.286$; Low agreeableness pairs vs. Low Agreeableness treatment: $p = 0.128$; Mixed agreeableness pairs vs. Mixed Agreeableness treatment: $p = 0.284$.
- For non-economists: High agreeableness pairs vs. High Agreeableness treatment: $p = 0.872$; Low agreeableness pairs vs. Low Agreeableness treatment: $p = 0.006$; Mixed agreeableness pairs vs. Mixed Agreeableness treatment: $p = 0.048$.

Appendix D – Experimental Instructions

[Note: These are the written instructions for the subjects facing Sequence 1 of our experiment. The order of the two sections below was reversed for Sequence 2.]

Instructions

Welcome to the Birmingham Experimental Economics Laboratory. This is an experiment in decision making. The University of Birmingham has provided funds for this research. Just for showing up you have already earned £2.50. You can earn additional money depending on the decisions made by you and other participants. It is therefore very important that you read these instructions with care.

It is important that you remain silent and do not look at other people's work. If you have any questions, or need assistance of any kind, please raise your hand and an experimenter will come to you. If you talk, laugh, exclaim out loud, etc., you will be asked to leave and you will not be paid. We expect and appreciate your following of these rules.

We will first jointly go over the instructions. After we have read the instructions, you will have time to ask clarifying questions. Each of you will then need to answer a few brief questions to ensure everybody understands. We would like to stress that any choices you make in this experiment are entirely anonymous. Please do not touch the computer or its mouse until you are instructed to do so. Thank you.

The experiment consists of two different sections. At the beginning of the experiment, we would like you to answer a brief questionnaire. Once everybody has completed this questionnaire, Section 1 of the experiment will begin. In each section you will be asked to make one or more decisions. You will have to make your decisions without knowing other participants' decisions in each section. Note also that the other participants will not know your decisions either. At no point during the experiment, nor afterwards will you be informed about the identity of the other participants and the other participants will never be informed about your identity.

Only one of the two sections will be taken into account in determining your final payoff. This will be randomly determined by the computer program at the end of the experiment. Each section has the same probability of being selected. Your earnings will be paid to you in cash at the end of the experiment.

In the next page, you will find the instructions for Section 1 of the experiment. You will receive new instructions for Section 2 once everyone in the room has completed Section 1.

Section 1

In this section the situation is as follows:

Person A is asked to choose between two possible distributions of money between herself and Person B in eleven different decision problems. Person B knows that A has been asked to make those decisions, and there is nothing he can do but accept them. The roles of Person A and Person B will be determined randomly and will remain anonymous.

Before making your decisions please read carefully the following paragraphs. The decision problems will be presented in a table. Each decision problem will look like the following:

Person A's Payoff	Person B's Payoff	Decision	Person A's Payoff	Person B's Payoff
10	0	Left Right	5	5

You will have to decide as Person A; hence if in the above decision problem you choose Left, you decide to keep the £10 for you so Person B's payoff will be £0. Similarly, if you choose Right, you and Person B will earn £5 each.

You will need to choose one distribution (Left or Right) in each of the eleven rows you will have in the screen.

If this is chosen as the payoff relevant section, the computer will randomly choose one of the eleven decisions. The outcome in the chosen decision will then determine your earnings.

The computer will pair you with another participant in the room and will assign the roles. The matching and the roles' assignment will remain anonymous.

Please note that you will make all decisions as Person A but the computer might assign you Person B's role.

If you are assigned the role of Person A, you will earn the amount that you have chosen for Person A in the relevant situation and the person paired with you will earn the amount that you have chosen for Person B.

In the case that you are assigned the role of Person B, you will earn the amount that Person A whom you are paired with has chosen for Person B in the relevant situation.

Section 2

In this section the situation is as follows:

Person A and Person B are given £10 each. In what follows we call this the “endowment”. Each person’s task is to decide how to use their endowment. Each person has to decide how much of the £10 s/he wants to contribute to a project (from 0 to 10) and how much to keep for her/himself.

The consequences of their decision are explained in detail below.

Each person’s earnings will be determined as follows:

Earnings = 10 – (your contribution to the project) + 0.7 * (total contribution to the project)

This formula shows that your “Earnings” in this section consists of two parts:

- 1) The part of the endowment which you have kept for yourself.
- 2) The income from the project, which equals to the 70% of the pair’s total contribution.

The income of each person from the project is calculated in the same way. This means that each person receives the same income from the project. Suppose that the total contribution to the project of both members is £20. In this case, each person receives an income from the project of: $0.7 * 20 = 14$ pounds. If the total contribution to the project is 4 pounds, then each member of the group receives an income of: $0.7 * 4 = 2.8$ pounds from the project.

Person A, as well as Person B, always has the option of keeping the endowment for her/himself or contributing part, or all of it, to the project. Each pound that s/he keeps raises her/his earnings by 1 pound. If he had contributed this unit to the project instead, then the total contribution to the project would increase by 1 pound. Her/his income from the project would thus increase by $0.7 * 1 = 0.7$ pounds. However, the earnings of the other person would also increase by 0.7 pounds, so that the total income of the pair from the project would be 1.4 pounds. Person A’s contribution to the project therefore also raises the earnings of Person B. In the same way, Person A also earns an income for each unit contributed by Person B to the project. In particular, for each unit contributed by either of them, both Person A and Person B earn 0.7 pounds.

Control questionnaire

Before making your decision please answer the following questionnaire. In a couple of minutes someone will come to your desk to check your answers. Once everybody answers the following questions correctly, the experiment will start.

1. Each person has an endowment of £10. Nobody contributes any pound to the project. What is:

a. Person A's earnings?

b. Person B's earnings?

2. Each person has an endowment of £10. Person A contributes £10 to the project. Person B contributes £10 to the project. What is:

a. Person A's earnings?

b. Person B's earnings?

3. Each person has an endowment of £10. Person A contributes £10 to the project. Person B contributes £4 to the project. What is:

a. Person A's earnings?

b. Person B's earnings?

4. Each person has an endowment of £10. Person A contributes £6 to the project. What is:

a. Person A's earnings if Person B contributes £8 to the project?

b. Person A's earnings if Person B contributes £6 to the project?

c. Person A's earnings if Person B contributes £2 to the project?