

Pelligra, Vittorio; Reggiani, Tommaso; Zizzo, Daniel John

Working Paper

Responding to (un)reasonable requests by an authority

Cardiff Economics Working Papers, No. E2020/7

Provided in Cooperation with:

Cardiff Business School, Cardiff University

Suggested Citation: Pelligra, Vittorio; Reggiani, Tommaso; Zizzo, Daniel John (2020) : Responding to (un)reasonable requests by an authority, Cardiff Economics Working Papers, No. E2020/7, Cardiff University, Cardiff Business School, Cardiff

This Version is available at:

<https://hdl.handle.net/10419/250317>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Cardiff Economics Working Papers



Working Paper No. E2020/7

Responding to (Un)Reasonable Requests by an Authority

Vittorio Pelligra, Tommaso Reggiani and Daniel John Zizzo

May 2020

ISSN 1749-6010

Cardiff Business School
Cardiff University
Colum Drive
Cardiff CF10 3EU
United Kingdom
t: +44 (0)29 2087 4000
f: +44 (0)29 2087 4419
business.cardiff.ac.uk

This working paper is produced for discussion purpose only. These working papers are expected to be published in due course, in revised form, and should not be quoted or cited without the author's written permission.

Cardiff Economics Working Papers are available online from:

<http://econpapers.repec.org/paper/cdfwpaper/> and

business.cardiff.ac.uk/research/academic-sections/economics/working-papers

Enquiries: EconWP@cardiff.ac.uk

Responding to (Un)Reasonable Requests by an Authority

Vittorio Pelligra

University of Cagliari [1]
CRENoS [2]

Tommaso Reggiani

Cardiff University [3]
Masaryk University [4]
IZA [5]

Daniel John Zizzo

University of Queensland [6]

This version: April 30, 2020

Abstract

We consider the notions of static and dynamic reasonableness of requests by an authority in a trust game experiment. The authority, modelled as the experimenter, systematically varies the experimental norm of what is expected from trustees to return to trustors, both in terms of the level of each request and in terms of the sequence of the requests. Static reasonableness matters in a self-biased way, in the sense that low requests justify returning less, but high requests tend to be ignored. Dynamic reasonableness also matters, in the sense that, if requests keep increasing, trustees return less compared to the same requests presented in random or decreasing order. Requests never systematically increase trustworthiness but may decrease it.

Keywords: trust; trustworthiness; authority; reasonableness; moral wiggle room; moral licensing.

JEL classification codes: C91; D01; D03; D63.

Vittorio Pelligra: [1] Department of Economics and Business, University of Cagliari, V.le S. Ignazio 72, 09121 Cagliari, Italy; [2] CRENoS, Cagliari, Italy. Financial support from the University of Cagliari is gratefully acknowledged.

E-mail: pelligra@unica.it - <https://orcid.org/0000-0002-8012-4243>

Tommaso Reggiani: [3] Cardiff Business School, Cardiff University, Aberconway Building, Colum Drive, CF10 3EU Cardiff, United Kingdom; [4] Masaryk University – MUEEL lab, Brno, Czech Republic; [5] IZA, Bonn, Germany. Financial support from the Czech Science Foundation - GACR: GA18-19492S - is gratefully acknowledged.

E-mail: reggianit@cardiff.ac.uk - <https://orcid.org/0000-0002-3134-1049>

Daniel John Zizzo: [6] School of Economics, University of Queensland, Level 6, Colin Clark Building, St Lucia QLD 4072, Australia.

E-mail: d.zizzo@uq.edu.au - <https://orcid.org/0000-0003-4830-4364>

Acknowledgements: We thank Diya Elizabeth Abraham, Giuseppe Attanasi, Rym Ayadi, Nick Bardsley, Stefania Bortolotti, Eric Cardella, Gary Charness, Martin Dufwenberg, Giovanni Ferri, Milos Fisar, Felix Höfler, Kiryl Khalmetski, Bernd Irlenbusch, Michael Niekamp, Rainer Michael Rilke, Matteo Rizzolli, Stefania Sitzia, Dirk Sliwka, Alessandro Sontuoso, Jiri Spalek, two anonymous reviewers and participants to presentations at Cagliari, CERGE-EI, Cologne, Durham, Heidelberg, Lille, LUMSA, Maryland, Masaryk University, Middlesex, Norwich UEA, Pittsburgh, Sophia University Institute, Vienna and at the ESA 2016 European conference in Bergen for useful feedback. The usual disclaimer applies. Online appendix: shorturl.at/iAESY.

Final version forthcoming in *Theory and Decision*.

1. Introduction

It is commonly believed that compliance is ubiquitous in social life. People may respond to explicit and implicit requests by modifying their behavior according to what they are requested to do. Managers in organizations may find this particularly helpful, and there is a variety of other contexts where it can also be useful, such as tax compliance and public good contribution (Cadsby, 2006; Silverman et al., 2014). An overlooked factor that may influence compliance is the reasonableness of the request received by an authority. When a manager explicitly or implicitly asks someone to do something, it is likely that her willingness to fulfil your request depends on how reasonable she perceives such a request. We try to operationalize the idea of reasonableness and to study its effect on compliance in the context of a fiduciary relationship.

Our baseline is a simple trust game in which the trustor has to decide whether or not to send her entire monetary endowment and the trustee, in turn, has to decide what proportion, if any, to send back. We investigate whether and to what extent if someone in a position of authority, such as a manager and in our study the experimenter, asks the trustees to send back positive amounts, this leads to an increase of trustworthiness and trust, and whether this depends on the reasonableness of the request.

The request is framed in the form of a message to the trustee saying that the experimenters expect him or her to send a specific share of what she received back to the trustor.¹ We make clear that the subjects are free to do whatever they want and we vary whether or not the trustor is aware of this message, with having this knowledge (or lack thereof) being known to the trustee. Our information manipulation relates to Ellingsen et al.'s (2010) manipulation of disclosing the first order expectation by the trustor to the trustee and determining whether this correlates with the return rate.² In the context

¹ Our external non-binding message manipulation, nicely complements the series of studies analyzing the role played by pre-play communication and cheap talk in trust games (Charness and Dufwenberg, 2006; Ben-Ner and Putterman 2009; Bracht and Feltovich, 2009; Ben-Ner et al., 2011; Bicchieri et al., 2010).

² Based on the finding of a lack of correlation, Ellingsen et al. (2010) argue that a false consensus effect drives evidence for guilt aversion. Khalmetski et al. (2015) show that heterogeneity in the responses to such knowledge may be an explanation for such lack of correlation. In different experimental settings, Vanberg (2008), Kawagoe and Narita (2014),

of our experiment, it works as a robustness test of our findings. This simple situation is intended to describe the basic dynamic underlying agents' decision to comply with what is perceived to be a changing or an externally imposed norm from someone in a position of authority. Consider a manager requesting a subordinate to fulfil trust to a specific degree towards an external contractor or buyer or colleague. We are interested in describing and testing whether and to what extent, trust and trustworthiness could be fostered or hampered by external intervention, for instance, in this case, by the signals coming from a manager. In real world settings, trustors may not precisely be aware of the specific request by a manager. However, it is plausible that in many cases a customer or external contractor may know, for example, about a company's expectation for customer service or about a company's reputation or ethical code of conduct, which imply a managerial expectation about how the customer or external contractor's company contact will behave towards them.³

As in Cadsby (2006), Silverman et al. (2014) and Sonntag and Zizzo (2015), and as discussed methodologically in Zizzo (2010), we deliberately use experimental demand as a treatment manipulation, i.e., in our study experimental demand is not a confound but rather a tool of the experimental design – in the case of our paper, to study the reasonableness of demands by the manager and alternative managerial strategies. This tool is additionally useful in our context is because it enables an exogenous and systematic manipulation in both the level and the order of the requests. By doing this, we can have possible requests across the whole range (from none to the whole of the pie) and have them systematically for all of the trustees; and we can have systematically different orders in which we present the requests.

The real world examples that we capture are ones where the cooperative behavior by the trustees usually benefits the manager. This could take place directly, as in the example of greater cooperation

and Conrads and Reggiani (2017) also argue against guilt aversion relative to alternative explanations. Evidence in support of guilt aversion relative to alternative explanations includes Reuben et al. (2009), Bellemare et al. (2011), Khalmetski (2016), Attanasi et al. (2016), Ederer and Stremitzner (2017), Khalmetski et al. (2018). In a dictator game setting, Hauge (2016) finds evidence both of an aversion to letting down others' belief and an aversion to letting down moral standards.

³ Companies and organizations often advertise their codes of conduct and standards of consumer service on their website.

within an organizational unit yielding to greater productivity, that the manager can take credit for e.g. for performance bonuses. It could also take place indirectly, as in the example of a manager requiring subordinates to behave ethically towards an external contractor, where compliance may simply improve the reputation of the company and only indirectly benefit the manager, or perhaps not at all. That said, there may be cases where the manager may be better off, in terms of financial unit performance that he or she can take credit for, by cutting corners and nudge the employee towards being less trustworthy with an outside contractor or buyer. By having exogenous requests, this paper abstracts from whatever motivation may lead the manager to recommend a higher or a lower level of trust fulfilling, though it considers both cases.

We argue that the experimentally induced norm is taken into account if “reasonable,” along two different dimensions, and find evidence that this is the case. Specifically, the *level* of the request is taken into account by trustees but in a self-biased way: requests may lower returns back to the trustors but do not raise them. The *dynamics* of requests matters as well: if requests keep ratcheting up, trustees return less than if requests of different size are presented in a random or decreasing order. We also review the extent to which trustors take the static and dynamic reasonableness of the requests into account when deciding whether to trust, and we find evidence to support that they take into account the static reasonableness of the requests, as well as their priors of what is reasonable. The rest of the paper proceeds as follows. Section 2 defines more precisely static and dynamic reasonableness and our experimental hypotheses in the context of our trust games. Section 3 presents the experimental design, while section 4 presents the results, section 5 provides a discussion, and section 6 concludes.

2. Reasonableness and Behavioral Hypotheses

We are interested in studying how and to what extent the individual's perception of the reasonableness of a request may affect his or her response to that request. This question is independent from the reasons why the agent respond to the request. Therefore we are agnostic on this latter point. We mention two possible explanations why here, and refer the reader to the working paper version of this paper for more (Pelligra et al., 2016). A first explanation could be framed in terms of guilt aversion (Dufwenberg, 2002; Battigalli and Dufwenberg, 2007; Bacharach et al., 2007); another could be a preference for compliance towards an experimentally defined norm, triggered by the experimental cues provided (Zizzo, 2010; Ellingsen et al., 2010; Karakostas and Zizzo, 2016).

2.1 Static Reasonableness

We define reasonableness both along a *static* dimension and a *dynamic* dimension. The static dimension relates to the absolute size of the request being made. Specifically, we assume that a request's reasonableness coincides with the trustee's prior, e.g., her or his understanding of the social norm to be followed. We define this subjective level, reasonable trustworthiness. By 'reasonable' trustworthiness we refer to the spontaneous idiosyncratic trustworthiness level of a given agent in absence of any specific request. This is likely to be at least partially socially determined, e.g. long-term social learning, as well as partially genetic (e.g., Zizzo, 2003). Assume that a request is lower than the trustee's prior understanding of the social norm to be followed. This means that the request is lower than the rate of that in our experimental setup is represented by the individual level of trustworthiness when there is no request. We would expect this to lead to a down-adjustment of the amount returned along the lines of the request. Now assume that the request is higher than the trustee's reasonable trustworthiness level. Such a request will be deemed unreasonable, and the trustee will stick to her or his prior, for two possible reasons. First, 'self-serving' biases (e.g., Babcock and Loewenstein, 1997; Konow, 2005; Valdesolo and DeSteno, 2008; Kocher et al., 2017; Charness et

al., 2019) could be at work.⁴ Second, there could be anger or annoyance at a request higher than the reasonable trustworthiness level (Bolle et al., 2014; Battigalli et al., 2019).⁵ The two key behavioral predictions that we derive from the introduction of the concept of static reasonableness are:

H1. There is a positive relationship between request levels and trustworthiness. For a sufficiently high request, the relationship becomes weaker.

H2. Request levels that are lower than the individual rate of reasonable trustworthiness reduce trustworthiness, but requests higher than the baseline priors do not increase it.

Our analysis so far does not make the strength of the requests a function of whether the requests are communicated to trustor and trustee knows this. However, it may be that if guilt aversion is one of the driver of the trustee's willingness to comply, requests should have a greater impact under a public information condition:

H3. The relationship between requests and trustworthiness is stronger if the trustor is informed of the request and the trustee is aware of that.

2.2 Dynamic Reasonableness

Dynamic reasonableness reflects the fact that the reasonableness of the request may be shaped by experience with previous requests. Intuitively, if the requests keep becoming higher, this may be seen as more unreasonable than if requests keep becoming lower. In the second case, the trustees may be all too happy to comply with the lower request; in the first case, and insofar as previous requests are partially used as anchors for what is reasonable, the subject may find requests more unreasonable.

⁴ A version of this is that agents tend to exploit what 'moral wiggle room' is available to them (e.g., Dana et al., 2007; Grossman, 2014; Rustichini and Villeval, 2014, Khalmetski, 2016)

⁵ Though we note that Battigalli et al. (2019) believe that anger is not relevant in trust games, on the ground that any return in this game requires players to place a positive weight on a co-player's material payoffs.

The idea that an increasing sequence of requests may be seen less favorably than a decreasing sequence of requests is compatible with experimental evidence that an increasing sequence of prices leads to lower sales than a decreasing sequence of prices (Sitzia and Zizzo, 2012). It is also consistent with evidence that in bargaining settings, better outcomes may be obtained by starting tough and then becoming softer than vice versa (Hilty and Carnevale, 1993). Siegel and Fouraker (1960) noted the importance of aspiration levels in bargaining settings and a substantial literature has looked at the effects of aspiration levels, anchors or reference points – however labeled – and their adjustment, in determining behavior in a bargaining context (e.g., Esser and Komorita, 1975; Levy, 1997; Korobkin and Guthrie, 2003-2004; Schotter and Sopher, 2007). Many recent experimental studies on negotiation have shown the significant effect of changing reference points, that may assume the form of pre-existing contracts, informal agreements or fairness norms, among the others, in influencing the bargaining process and the final agreement (see Blount et al. 1996; Kristensen and Gärling 2000; Gächter and Riedl 2005; Bartling and Schmidt 2015; Karagözoğlu and Riedl 2015; Bolton and Karagözoğlu 2016 among others). Intuitively, reasonableness may, at least in part, be evaluated in a reference-dependent manner, where previous experiences of requests may shape what is considered more or less reasonable. The initial request may influence the initial anchor or reference point and, due to the perception of a higher subsequent request as entailing a comparative loss or bad deal (Isoni, 2011), trustees will be more reluctant to adjust their reference point to a higher level of return than they would be to adjust it to a lower level. Put it simply, they accommodate to perceived gains more easily than they accommodate to perceived losses (Levy, 1997).⁶ Further related literature can be found in Pelligra et al. (2016).

⁶ Sliwka and Werner (2017) present an experiment focusing only on an increasing (wage profile) case, and specifically on the effect of the granularity of the wage increase profile: they find that gradually increasing wage profiles are better in achieving compliance than having large infrequent jumps. Marginal changes may be easier to adjust to than large changes in requests.

The following hypothesis follows from the notion of dynamic reasonableness:

H4. Under increasing requests, trustworthiness will be lower than if requests are presented in decreasing or random order in successive rounds.

2.3 Trustor Behavior

Two final hypotheses refer to the expected trustor's behavior. Trustors are obviously more likely to trust the more they expect trustees to fulfil trust. Rational expectations on the part of the trustor would imply that he or she should correctly anticipate how trustees would on average react to the static and dynamic reasonableness of requests. Specifically:

H5. Trust will be lower (a) when requests are low and trustors know they are low and (b) under increasing requests.

Furthermore, supported by their own experience as trustees in the first stage, trustors may use their own sense of what is reasonable as a guide to what others may think as reasonable in terms of amount that will be returned (a form of common consensus effect: e.g., Gilovich et al. 1983). This leads to the following hypothesis:

H6. Trust will be higher the higher the reasonable trustworthiness rate of the trustor.

3. Experimental Design

A total of 120 subjects (mean age of 25 years, 43% male) participated in the experiment. Paper-and-pencil sessions took place at the BERG Lab of the University of Cagliari, Italy in May 2014. On average, the experimental subjects received €16.85 (including a €5 show-up fee) for a 1-hour session. Subjects were students enrolled at the University of Cagliari (Schools of Business and Law). The conversion rate between experimental points and Euro was 1 point = €0.25. The experimental instructions are reported in Appendix A.

Each experimental session had two stages, to maximize sample power. In the first stage, each subject first played eight rounds of the trust game in the role of the trustor “Player A”. In the second stage, each subject played eight rounds of the trust game in the same experimental treatment but this time in the role of the trustee “Player B”.⁷ Subject interact under a partner matching protocol without any feedback provision, about actions or outcomes, before the end of the experimental session. To implement an incentive-compatible payment mechanism, at the end of each experimental session half of the subjects taking part in the experiment were randomly actually assigned to the trustor role and the other half to the trustee role. Each trustor was then randomly matched with a trustee, one out of the eight rounds was randomly drawn. The outcomes generated by the correspondent individual decisions were implemented and paid out in cash.

In our version of the trust game, the trustor had a binary choice whether to send 50 points or 0 points to the trustee (intermediate values were not allowed). If the trustor transferred his or her 50 points to the trustee, they were multiplied by 3 and became 150 points. Then the trustee determined how to share between himself or herself and the trustor. Excluding rounds 1 and 8, in which no requests were made, in each of the other rounds a request ranging between 0% and 100% of B’s points (0, 30, 60, 90, 120 and 150 points) was made to the trustee as the proportion of points to be returned to the trustor.⁸ We employed the strategy method to be able to collect a full profile of trust responses (i.e., for each level of request) from trustees in terms of how much they would be willing to return conditionally on the trustor having chosen to transfer his or her 50 points.⁹ The additional advantage

⁷ We explain the reason for this specific sequence below.

⁸ Player B: “In this scenario you are free to decide how many points to return to Player A, however if Player A has sent to you 50 points, we request you to return # % (# points)”. The full experimental instructions are contained in Appendix A.

⁹ As a potential limitation of this approach, we acknowledge that Burks et al. (2003) showed that playing both roles increases selfish behavior, leading to a reduction in trust as well as trustworthiness. That said, Brandts and Charness (2011) compiled a survey of the literature focusing on the implications generated by the adoption of the strategy method, including Burks et al. (2003). They conclude that in no cases is it found that treatment effects from an experiment with the strategy method are not observed if the direct-response method is implemented. We also note that none of our hypotheses depend on the absolute level of trust and trustworthiness but rather on how these differ across rounds and treatments. As such, the potential confound from a downward shift of the absolute levels observed in the trust interaction is not an issue for the conclusions of our experiment.

of the strategy method in our setting is to minimize learning opportunities between rounds (though an ad-hoc check for learning effects is provided comparing rounds 1 and 8, which are identical other than for sequencing).

We employed a between-subjects 2 x 3 factorial design (see Table 1) based on two different information settings (*Communication* vs. *NoCommunication*) and three different orders of the trust games and associated requests (*Increasing*, *Decreasing* and *Random*).

[Insert Table 1 about here.]

The first experimental manipulation is about whether or not the trustor is aware of the experimenters' request to the trustee. In the *Communication* (*Com*) experimental condition, the trustor is aware of the experimenters' request to the trustee, and the latter knows that this information is common knowledge.¹⁰ Under the *NoCommunication* (*NoCom*) condition, the trustor is not informed about the requests of the experimenters to the trustee, and the latter is aware that the former is not informed about these requests.¹¹ In order for the *NoCom* condition to work as a test of the effect of the trustor being unaware of the request and the trustee knowing it, subjects played first as trustors in the first stage of the experiment and then as trustees in the second stage. If they had played in the reverse order, they could have naturally inferred when playing as trustors in stage 2 that the trustee was being made a request by the experimenters. This fact could have altered the response in the second stage.¹² The rest of the structure of the game is public information under both conditions.

¹⁰ To the text in footnote 14, the following was added for players B (trustees): "Player A is aware of this request". Player A's corresponding instructions read: "Player B receives the following message: 'In this scenario you are free to decide how many points to return to Player A, however if Player A has sent to you 50 points, we request you to return # % (# points)'".

¹¹ To the text in footnote 14, the following was added for players B: "Player A is not aware of this request".

¹² A secondary reason for our stage sequence is that our interest in this experiment is mainly about Player B's behavior. By playing in the role of Player A first, we ensured a greater understanding of the game being played when they had to choose as players B.

The sequence in which the different requests are made in successive rounds of the trust game, in each one of the two player-role stages, constitutes the second experimental manipulation. This experimental variation aims at investigating how *dynamic* considerations affect reasonableness-related concerns. In the *Increasing (Inc)* experimental condition, the sequence of the requests is organized following an escalation from round 2 to round 7 (0, 30, 60, 90, 120, 150 points in successive rounds). In the *Decreasing (Dec)* condition, the sequence of requests follows a decreasing trend from rounds 2 to 7 (150, 120, 90, 60, 30, 0 points in successive rounds). In the *Random (Rand)* experimental condition, the ordering of the requests is randomly determined between rounds 2 and 7. In all conditions, and as a control for any potential effect that is just having a request (whatever that may be) may have, there are no requests in rounds 1 and 8. Based on the above abbreviations, we label the experimental treatments as *NoCom-Inc*, *Com-Inc*, *NoCom-Dec*, *Com-Dec*, *NoCom-Rand*, and *Com-Rand*. Each subject played the Player A and Player B roles in the same treatment, e.g. if a subject played *Com-Dec* as Player A, he or she also played *Com-Dec* as Player B.

4. Results

Define the **trust rate** as the share of trustors that sent their 50 points to the corresponding trustees and the **trustworthiness rate** as the number of points (out of 150 points) returned by trustees to trustors. The amount of points returned by the trustee in round 1 identifies the *a priori* level of reasonableness in the context of our trust game setup. It represents the individual rate of reasonable trustworthiness. The mean round 1 trustworthiness rate was 57.57 points, with some statistically insignificant sample heterogeneity across treatments (Table 2) that will be controlled for in the regression analysis described later.¹³ The mean round 8 trustworthiness rate was 56.18 points, which

¹³ When points returned at individual level in round 1 are regressed against manipulation dummies, no coefficient turns out to be statistically significant at any conventional level (see Appendix B, Tables B1 and B2). A non-parametric Kruskal-Wallis test fails to reject the joint null hypothesis of equality of the means ($p=0.32$). In the *Dec* treatment (both *Com* and *NoCom*) we observe a slightly higher natural trustworthiness than in the other treatments. The qualitative difference is

is not statistically significantly different from round 1 (signed-Mann Whitney $p = 0.95$),¹⁴ suggesting that learning across rounds (which should be minimal due to the lack of feedback between rounds) is not an issue in identifying the level of reasonable trustworthiness.

[Insert Table 2 about here.]

Table 2 shows, in square brackets, mean points returned in the different treatments under the different levels of request; the corresponding trust rates are reported in standard brackets.

[Insert Figure 1 about here.]

Figure 1 summarizes the Table 2 data in terms of the order sequence:¹⁵ under *Dec* and *Rand*, trustworthiness rates increase with the request, albeit not on a one-to-one basis, and reach around 70 points when the request is for 150 points. The story is different under *Inc*: Figure 1 shows that mean trustworthiness rates are always lower than under *Dec* and *Rand* (Mann-Whitney $p = 0.03$ and 0.06 , respectively).¹⁶ That could be in part explained by different priors in terms of reasonable trustworthiness levels, which will be controlled for in the regression analysis. Figure 1 compares round 1 baseline trustworthiness rates (reasonable trustworthiness) with trustworthiness rates when there are requests. It shows a clear asymmetry. i) Decreasing requests tend to lower trustworthiness. ii) This negative effect is stronger than the positive effect generated by increasing ones.

[Insert Tables 3 and 4 about here.]

Table 3 reports a battery of Tobit regressions to assess the causal effect generated by the experimental manipulations on trustworthiness (measured as the number of points that the trustee returns to the

mainly generated by a higher number of ‘unconditional egalitarian’ participants in that treatment (see Appendix C for participants’ classification).

¹⁴ All p values provided in this paper are two sided.

¹⁵ Figure B1, in the Appendix B, provides a further graphical stratification by information setting (*Com* / *NoCom*).

¹⁶ In this and later between-subjects bivariate tests with multiple observations per subject, we use mean values by subject as the independent observation to avoid within-subject non-independence of the observations.

trustor).¹⁷ The regressions include dummy variables for the experimental manipulations (*Inc*, *Dec* and *Com*, equal to 1 in the respective treatments, else 0) and their interactions,¹⁸ dummies for the value of the request (Request_0 for a request of 0 points, Request_30 for one of 30 points, and so on) and depending on the model, for round dummies (Round 8 = 1 for round 8 decisions) and demographic covariates (age, gender, religion and economics background).

Models 1-4 include observations from all rounds and given the value of the request as well as round dummies, has as a baseline the mean rate of reasonable trustworthiness. Models 5-8 use the individual rate of reasonable trustworthiness, as revealed by trustee behavior in round 1, as a control variable, therefore controlling for any sample heterogeneity across treatments.¹⁹ These models are especially useful to test H2 (on the effect of requests relative to the rate of reasonable trustworthiness).

However, as there is no experimental manipulation in rounds 1 and 8, and indeed round 1 takes place before any experimental manipulation takes place, experimental treatment effects are likely to be noisy and diluted in these regressions. We therefore also estimate models 9-16 in Table 4, which include only observations from rounds 2-7 when the experimental manipulations took place, and which more accurately test for the effects of our experimental manipulations and therefore for H3 (on the effect of knowledge of trustor being informed of the request) and H4 (on the effect of requests order). The baseline for these regressions is the case where the request is 150. While the baselines are different, we can use both sets of models to look at H1 (on the relationship between requests and trustworthiness).

¹⁷ Since the outcome variable ranges in the truncated interval between 0 and 150, the adoption of Tobit models is a natural choice. Robust standard errors of the estimates (clustered at individual level since each subject plays repeatedly) are reported in brackets. Comparable results are delivered by GLS estimations with random effects (see Appendix B). We also ran regressions on round 2 return rate (see Appendix C).as well as regressions without treatment variables, again leading to comparable results.

¹⁸ We have also tried to run the regressions with no interaction terms and the qualitative picture remains the same.

¹⁹ We have also tried to use the difference of the individual round 1 trustworthiness rate from the average round 1 rate instead, in which case the constant takes values similar to those estimated in models 1-4. None of the other results changes. A parallel analysis for the regression in Table 4 leads to the same picture.

Round and demographic coefficients are never statistically significant and do not affect the other coefficients. The small and not significant coefficient on round 8 supports our earlier finding of lack of difference between round 1 or 8 trustworthiness, though that for round 1 is a cleaner measure.

H1 is supported. There is evidence of a positive relationship between requests and return rate, but this tends to disappear for a sufficiently high request.

Visual inspection of Figure 1 and Table 2 suggests the existence of a positive aggregate relationship between request levels and trustworthiness.²⁰ The regression analysis of Table 3 supports this. It shows that the point coefficients are consistently increasing in the size of the request up to a request of 120, though with strong effects particularly evident for low request values and coefficients for requests of 90 and above not statistically different from the rates of reasonable trustworthiness as defined in models 1-4 and 5-8. Focusing on the most encompassing models 4 and 8, the coefficients for Request_0 are statistically lower than those on Request_60 and higher request values (Wald test $p < 0.1$ or better);²¹ those on Request_30 are statistically lower than those on Request_90 and higher request values (Wald test $p < 0.05$ or better). Conversely, coefficients associated with Request_90 are not different from the ones of Request_150 in regressions 1-4 (Wald test $p > 0.1$); and we have the same picture from the lack of significance of Request_90 in models 9-16, as this denotes the lack of significant difference relative to a request of 150 that constitutes the baseline for these regression models.

H2 is supported. Trustworthiness rates are lower than the rate of reasonable trustworthiness when requests are lower than the baseline priors, but the converse is not true: requests higher than the baseline priors do not push trustworthiness rates up.

²⁰ A simple test just comparing the mean return rate when the first request (round 2) is received, shows that it is equal to 34.5 when the request is 0 and twice as much (68.2) when it is 150 (Mann-Whitney $p < 0.001$). Further supportive evidence is provided by the round 2 only regression reproduced in Appendix C, Table C2.

²¹ Here and below, we focus on tests of requests that differ by 60 points or more, as we do not have the statistical power to detect differences for requests that differ by only 30 points.

While Figure 1 shows preliminary evidence for this, models 1-8 enable us to test this and corroborate this result. Focusing on models 5-8 that control for individual level of reasonable trustworthiness, requests of 0, 30, and 60 significantly decrease the trustworthiness, by around 18-20, 13 and 9-11 points respectively. High requests instead do not lead to statistically significant increases, with positive point values no larger than 4 points for requests of 120 and 150.²²

Appendix C shows the result of a regression analysis where we introduce both the request level and a dummy variable equal to 1 when the request is higher than the individual level of reasonable trustworthiness, and else equal to 0. We find that, while the coefficient on the request is positive and significant, that on the dummy variable is negative and offsetting it ($p < 0.001$).

H3 is not supported. The relationship between requests and trustworthiness is not stronger if trustor is informed of the request, and the trustee knows this.

Table 2 does not show much difference in terms of the effect of communication on trustworthiness in terms of point returned: the average is 56.45 points in the *NoCom* treatments versus 48.01 points in the *Com* treatments (Mann-Whitney $p > 0.1$). The regressions reported in Table 3 and 4 also confirm this pattern: the *Com* dummy is never significant at conventional test levels, even in models 9-16 that should be especially suited to detect such effects (and removing the interaction dummies does not help).

H4 is broadly supported. Under increasing requests, trustworthiness tends to be lower than if requests are presented in random or decreasing order in successive rounds.

²² H2 is also supported by the round 2 only regressions reported in Appendix C. It also provides some additional analysis based on reasonable trustworthiness levels. It shows that the evidence for H1 and H2 is likely to be at least partially driven by trustees with an intermediate level of reasonable trustworthiness, i.e. round 1 trustworthiness rate of 50 points (33.3% of the amount received) or 75 (50% of the amount received). Note that Figure 1 also shows that there is no systematic *decrease* in trustworthiness when requests are too high, e.g. out of anger or annoyance.

Visual inspection of Figure 1 and Table 2 suggests that the increasing order of requests reduces trustworthiness. This point received support from the regression analysis of Table 4. It shows that coefficients on *Inc* are always negative and statistically significant at the 5% level, once we focus on models 5-8 that are suited to test H4 to exclude rounds 1 and 8 which have no requests.²³ The effects appear economically meaningful in their magnitudes, with values of 19-20 points once one controls for sample heterogeneity across treatments (models 13-16). The coefficients on *Dec* are not statistically significant and are consistently lower at the $p < 0.1$ level than those on *Inc*.²⁴

We now move on to consider trustor behavior and to test H5 and H6. In doing so, we need to treat separately treatments with communication and treatments without it. The treatments with communication are ones where the trustor know that the requests have been made to the trustee and can, therefore, decide whether to send or not based on his/her expectations about trustees' reactions to the requests, hence enabling a test of H5a.²⁵

Table 2 suggests that, in all three treatments with communication, mean trust rates follow an inverted U-shaped relationship in the size of the requests. Specifically, in all three main treatments, they are in an upward trend as requests go from 0 up to a peak between 60 and 120 points and then, in all three cases, seem to go down. This suggests that, on average, trustors believe that trustees adapt their behavior to the requests, though the effects are seen to level out with a sufficiently high request. This pattern is not found in the treatments without communication.²⁶ Correspondingly, aggregate earnings and therefore efficiency follows an inverted U pattern, ranging from an average of 41% for a request of 0 to 61% for a request of 90 and back to 50% for a request of 150 in the *Com* treatments.

²³ That said, Table 3 shows still shows $p < 0.1$ or 0.05 depending on the regression.

²⁴ Wald test $p = 0.088$ in models 5 and 6 and 0.082 in models 7 and 8.

²⁵ Put it differently, players' A perceptions of how reasonable the requests may shape their offers.

²⁶ There is still some random variation in choices in these cases. This might be due to some combination of noise and mixed strategies in trustors' choices.

[Insert Tables 5 and 6 about here.]

We formalize and extend the analysis using Probit regression models. Table 5 considers observations from all rounds in the second stage (with the rate of reasonable trustworthiness estimated from the first stage as for Tables 4 and 5). As predicted by H5a, the trust rate is low when requests are low (0 and 30) and trustors know when they are low. Table 6 considers observations from rounds 2-7 only of the second stage, with the round 150 request observation as the baseline (i.e. it is the equivalent of Table 5), so as to enable us to test whether trustors predict the impact of dynamic reasonableness as per H5b. We do not find evidence of a statistically significant effect in support of H5b, a finding that we turn back to in the discussion section.²⁷

H5 is partially supported. In the presence of communication, the trust rate is low for low requests. We do not find evidence of a statistically significant effect of an increasing pattern of request on trustors' behavior.

Table 5 shows that the coefficient on the individual level of reasonable trustworthiness is positive and statistically significant, though it is over twice as large in *NoCom* treatments than in *Com* treatments. In the *Com* treatments, the coefficient is small enough that, with the smaller sample of Table 6, it is no longer significant, but it is significant in the larger sample of Table 5. Under *NoCom*, trustors do not have the requests to rely on, and so, in deciding whether to trust, they place larger weight on their priors on what is reasonable to return.

H6 is supported. Particularly in the lack of communication, subjects' trust rate is a positive function of their own level of reasonable trustworthiness.

Probit analysis reported in Table 5 confirms the sizable and statistically significant decline in trust rates for low requests (0 and 30 points), though not for higher requests. The trustors just fail to

²⁷ We have also conducted the regression models in Table 5 and 6 without the reasonable trustworthiness rate variable, and these findings are unchanged.

anticipate the mild negative effect, generated by the 60 points request, in the trustworthiness rates. The outcome variable is a dummy that has value 1 when the trustor sends to the trustee her endowment of 50 points, 0 otherwise. Under the *Communication* experimental condition (Models 1-4) the negative effect induced by low request levels (0 and 30 points) is sizable and highly significant in all the different specifications. As expected, this is not the case under the *NoCommunication* condition (Models 5-8) in which we control for the ‘virtual’ level of request (for any given round, the ‘virtual’ request relates to the actual request submitted to the trustee but not revealed to the trustor).²⁸

5. Discussion

We found a positive relationship between requests by the authority and return, but one largely relying on low requests becoming lower than the baseline priors in terms of reasonable return rates, as opposed to high requests leading to higher trustworthiness. Such requests are not more effective when the trustee knows that the trustor is informed about the requests. They are however less effective when presented in increasing order, compared with random or decreasing order.

To evaluate the static and dynamic reasonableness of the requests, we had to vary such requests systematically, and the exogeneity of the manipulation makes such requests coming from the experimenter such a useful tool in our context.²⁹ Our regression analysis controlled for possible sample heterogeneity across treatments, by relying on individual rates of reasonable trustworthiness as measured by round 1 return rates (models 1-4 and 9-16 in Tables 3 and 4 respectively). Results 1 and 2 shows that static reasonableness matters in a self-biased way: they are consistent with a preference for compliance to an authority or an experimentally defined norm being adjusted by the request, but only in a self-biased direction. Consistently with Ellingsen et al. (2010), the relationship

²⁸ See also Tables B3, B4, B5 in Appendix B. We also checked whether round 1 trust rates are statistically different across treatments and they are not.

²⁹ For an example of use of cheap talk communication *between* trustor and trustee, see Charness and Dufwenberg (2006).

is not strengthened by the trustee knowing that the trustor knows about the request. We cannot rule out some underlying response heterogeneity (Khalmetski et al., 2015), and it is plausible that the experimenter request is sufficient to guide the trustee or induce the experimentally defined norm, without any further effect from trustor knowledge of it.³⁰

There is a significant literature showing how moral wiggle room is exploited in self-serving ways (e.g., Dana et al., 2007; Grossman, 2014) and that fairness is perceived in self-biased ways (e.g., Babcock and Loewenstein, 1997; Konow, 2005; Valdesolo and DeSteno, 2008; Kocher et al., 2017), and our findings are consistent with this research. In a sender-receiver game, Khalmetski (2016) find that, when the incentives for lying are high, the trustees tend not to take into account the trustors' high expectations. Charness et al. (2019) consider moral wiggle room in the context of trust games and find evidence that conformist participants are so in a self-serving way. Regner and Matthey (2015) contain a good recent discussion of the moral wiggle room literature and also in the context of trust games, find that 40% of reciprocators in their experiment exploited moral wiggle room.

Our experiment employs the experimental methodology of Cadsby et al. (2006), Silverman et al. (2014), Karakostas and Zizzo (2016) and Sonntag and Zizzo (2015) of using experimenter demand as a tool to define the norm by which behavior is expected in the experiment.³¹ Such requests are effective in inducing greater compliance in a tax payment game (Cadsby et al., 2006), in line with previous work showing that a 'tax frame' induces greater tax payment than a neutral gamble frame of a tax decision (Alm et al., 1992); in a public good game, in terms of inducing greater contribution (Silverman et al., 2014; Sally, 1995); in an obedience game, where subjects are asked to destroy the

³⁰ We also cannot rule out that the trustees did not believe that the trustors' expectations would be affected by the requests. This is however somewhat implausible since, as previously noted, on average trustors did change their behavior as a result of knowing about the requests, and trustees previously played as trustors. A related scenario, which has been suggested to us, would be one where trustees thought that the expectations of the trustors at the end of the experiment mattered, and that at the end of the experiment trustors would simply think of the sequence of requests as opposed to remembering the request specific to any given round. Of course this specific argument is based on a process of fairly sophisticated inference by the trustees, which is unlikely to have held when *NoCom* trustees first saw a request in round 2. However, when we compare mean return rates in round 2 between *Com* and *NoCom*, we still find no difference (Mann-Whitney $p = 0.31$).

³¹ See Zizzo (2010) for a methodological discussion of deliberately using experimenter demand as an experimental tool.

money of others (Karakostas and Zizzo, 2016); in a dictator game, where asking increases giving (Andreoni and Rao, 2011); and in a Cournot oligopoly setting, where greater collusion is induced (Sonntag and Zizzo, 2015). A non-experimental example of the tendency to defer to authority is Harrington (1988), who provides evidence that firms tend to comply to environmental regulation to a much greater extent than theoretically predicted, that is, even when monitoring is rare, the punishment of the transgressors is unlikely and fines negligible. While it would require a more complex design given the loss of exogeneity of the request manipulation, one possible and natural extension of our work would be to compare requests by an authority with requests by peers; Silverman et al. (2012) have already looked at this in their public good contribution setting. While it is natural to claim that the experimenter is perceived to be in a position of authority with respect to subjects in an experimenter laboratory (Zizzo, 2010), one possible avenue for future research would be to elicit subjects' perception of the request or make its source more transparent in order to verify whether this is indeed the case. It would also be interesting to vary how the degree of compliance by the subordinate affects the agent making the request.³²

Result 4 shows that dynamic reasonableness matters. A “foot in the door” approach (Cialdini and Goldstein, 2004) of starting with a low request and then go higher can be problematic since the recipients of the requests can consider the ratcheting up of requests unreasonable and therefore discount these requests. This is again consistent with a self-biased perception of reasonableness. Once initially shifted downward by the initial very low requests and once a clear ratcheting up pattern is identified, reasonableness perceptions are hard to shift to the levels that we observe with a random order or a decreasing order in terms of return rates.

The dynamic reasonableness effect for the *Inc* treatment is not driven just by the initial anchoring at 0%. Under the *Rand* condition with a round 2 request of 0%, we observed significantly higher

³² Another direction for future research would be to provide request only as absolute numbers or only as percentages of the total amount, rather than as both as we did in our experiment.

trustworthiness in the subsequent rounds 3-7 than what we detected in the *Inc* condition (Mann-Whitney $p < 0.001$) in the analogous rounds. It is also not driven just by taking pairs of consecutive rounds: a regression on the return rate in the *Rand* treatment having the request from the previous round as an independent variable shows that this is not statistically significant. An alternative potential dynamic effect is that later requests lose credibility relative to early requests. This can be tested in the *Rand* treatment, wherein additional regression analysis we generally find that requests are still responsive to (low) requests in rounds 5-7, against this possibility. Still, more research is needed.

We did not elicit second-order beliefs of the trustee in this experiment. Due to the adoption of the strategy method in correspondence of different levels of the requests, it would have been cumbersome to operationalize and possibly distorting the trustee's behavior. Nevertheless, this could be an interesting avenue for future research, particularly if it were to separate out second-order beliefs of the trustor and second-order beliefs with respect to the experimenter. Equally, it would be interesting to tie this research with that on cheap talk between trustor and trustee in trust games as illustrated by Charness and Dufwenberg (2006); or to extend it to other settings such as gift exchange games. A further possible extension would be to run a control treatment with constant requests.

Our H5 posits that trustors are capable of correctly anticipating trustees' average responses. We found that trustors correctly take static reasonableness into account, but do not take into dynamic reasonableness into account. One conjecture for this is that, for simplicity in a first experiment, we did not tell trustors that trustees had seen requests in the same increasing order. We expect that trustees will have seen quickly enough that requests were in the same increasing order as they experienced, but nevertheless we may not rule out the possibility that this may have weakened their ability to take the dynamic reasonableness into account; further research providing explicit guidance to trustors on this point is needed. Another conjecture is that dynamic reasonableness is harder for trustors to appreciate, particularly in a setting such as ours where there is not opportunity for learning

feedback. We do find that, particularly when they do not know about what request has come to them, trustors project their own belief on what constitutes a reasonable request into whether trustors are likely to fulfil trust and, therefore, on whether they should be trusted in the first place; this also may change with experience. Repeated play with feedback would clearly be a useful avenue for further research.

Overall, there is a pessimistic message to our paper, in that overall requests do not help but may harm. Given the managerial importance of requests, it is important to identify the determinants of reasonableness and under what circumstances requests can be used to elicit greater trustworthiness. For example, the natural language phrasing of the requests – e.g., whether they are framed more or less politely – and the extent to which persuasive arguments are made to justify the requests could matter.

6. Conclusion

There are two key messages from this paper. First, when requests are received from an authority, their reasonableness is taken into account in determining how to respond to them, but in a self-biased way. We have tested this intuition in the context of trust games where we found that both static and dynamic reasonableness of the requests matter when trustees decide how much to return-back to trustors. Static reasonableness refers to the level of the requests: when the request is too high, it does not generate additional return back to the trustors. Dynamic reasonableness refers to the order of successive requests: if requests keep ratcheting up, trustees return less than if requests of different size are presented in a random or decreasing order. In our experiment, trustors take the static reasonableness of the requests into account, while further research is needed with regard to whether they may take dynamic reasonableness into account. When deciding whether to trust, trustors also take into account their priors on what is reasonable, particularly when they do not know about the requests.

The second message of our paper is that, in our trust game setting, requests never systematically increase return rates, but may decrease them. We interpret this in terms of moral wiggle room and, taken literally, it would imply that managerial neglect may not be detrimental to trustworthiness, but managerial activism could lead to lower rates of trustworthiness. Given the importance of requests as a managerial tool for organizations, further research should identify ways of mitigating such a detrimental effect and of making requests more effective.

References


- Alm, J., McClelland, G.H., Schulze, W.D. (1992). Why do people pay taxes? *Journal of Public Economics* 48(1), 21-38.
- Andreoni, J., Rao, J.M. (2011). The power of asking: How communication affects selfishness, empathy, and altruism. *Journal of Public Economics* 95, 513–520.
- Attanasi, G., Battigalli, P., Nagel, R. (2016). Disclosure of belief-dependent preferences in a trust game. IGIER Working Paper 506, Bocconi University.
- Babcock, L., Loewenstein, G. (1997). Explaining bargaining impasse: The role of self-serving biases. *Journal of Economic Perspectives* 11(1), 109-126.
- Bacharach, M., Guerra, G., Zizzo, D.J. (2007). The self-fulfilling property of trust: An experimental study. *Theory and Decision* 63(4), 349-388.
- Bartling, B., Schmidt, K.M. (2015). Reference points, social norms, and fairness in contract renegotiation. *Journal of the European Economic Association* 13, 98–129.
- Battigalli, P., Dufwenberg, M. (2007). Guilt in games. *American Economic Review* 97(2), 170-176.
- Battigalli, P., Dufwenberg, M., Smith, A. (2019). Frustration, aggression, and anger in leader-follower games. *Games and Economic Behavior*, 117, 15-39.
- Bellemare, C., Sebald, A., Strobel, M. (2011). Measuring the willingness to pay to avoid guilt: estimation using equilibrium and stated belief models. *Journal of Applied Econometrics* 26(3), 437-453.
- Ben-Ner, A., Putterman, L. (2009). Trust, communication and contracts: An experiment. *Journal of Economic Behavior & Organization*, 70(1-2), 106-121.
- Ben-Ner, A., Putterman, L., Ren, T. (2011). Lavish returns on cheap talk: Two-way communication in trust games. *The Journal of Socio-Economics*, 40(1), 1-13.
- Bicchieri, C., Lev-On, A., Chavez, A. (2010). The medium or the message? Communication relevance and richness in trust games. *Synthese*, 176(1), 125-147.
- Blount, S., Thomas-Hunt, M.C., Neale, M.A. (1996). The price is right—or is it? A reference point model of two-party price negotiations. *Organizational Behaviour and Human Decision Process* 68, 1–12.

- Bolle, F., Tan, J.H.W., Zizzo, D.J. (2014). Vendettas. *American Economic Journal: Microeconomics* 6(2), 93-130.
- Bolton, G.E., Karagözoğlu, E. (2016). On the influence of hard leverage in a soft leverage bargaining game: the importance of credible claims. *Games and Economic Behavior* 99, 164–179.
- Bracht, J., Feltovich, N. (2009). Whatever you say, your reputation precedes you: Observation and cheap talk in the trust game. *Journal of Public Economics*, 93(9-10), 1036-1044.
- Brandts, J., Charness, G. (2011). The strategy versus the direct-response method: A first survey of experimental comparisons. *Experimental Economics* 14, 375-398.
- Burks, S.V., Carpenter, J.P., Verhoogen, E. (2003). Playing both roles in the trust game. *Journal of Economic Behavior and Organization* 51, 195-216.
- Cadsby, C.B., Maynes, E., Trivedi, V.U. (2006). Tax compliance and obedience to authority at home and in the lab: A new experimental approach. *Experimental Economics* 9, 343–359.
- Charness, G., Dufwenberg, M. (2006). Promises and partnership. *Econometrica* 74, 1579-1601.
- Charness, G., Naef, M., Sontuoso, A. (2019). Opportunistic conformism. *Journal of Economic Theory* 180, 100-134.
- Cialdini, R., Goldstein, N. (2004). Social influence: Compliance and conformity. *Annual Review of Psychology* 55, 591-622.
- Conrads, J., Reggiani, T. (2017). The effect of communication channels on promise-making and promise-keeping: Experimental evidence. *Journal of Economic Interaction and Coordination* 12(3): 595-611.
- Dana J., Weber R., Kuang J. X. (2007). Exploiting moral wiggle room: Behavior inconsistent with a preference for fair outcomes. *Economic Theory* 33, 67-80
- Dufwenberg, M. (2002). Marital investment, time consistency and emotions. *Journal of Economic Behavior and Organization* 48, 57-69.
- Ederer, F., Stremitzer, A. (2017). Promises and expectations. *Games and Economic Behavior*. 106, 161-178
- Ellingsen, T., Johannesson, M., Torsvik, G., Tjotta, S. (2010). Testing guilt aversion. *Games and Economic Behavior* 68(1), 95-107.
- Esser, J. K., Komorita, S. S. (1975). Reciprocity and concession-making in bargaining. *Journal of Personality and Social Psychology* 31, 864-872.
- Gächter, S., Riedl, A. (2005). Moral property rights in bargaining with infeasible claims. *Management Science* 51, 249–263.
- Gilovich T., Jennings D. L., Jennings S. (1983). Causal focus and estimates of consensus: an examination of the false-consensus effect. *Journal of Personality and Social Psychology* 45, 550–559.
- Grossman, Z. (2014). Strategic ignorance and the robustness of social preferences. *Management Science* 60(11), 2659-2665.
- Harrington, W. (1988). Enforcement leverage when penalties are restricted. *Journal of Public Economics* 37, 29-53.
- Hauge, K.E. (2016). Generosity and guilt: The role of beliefs and moral standards of others. *Journal of Economic Psychology* 54, 35-43.

- Hilty, J.H., Carnevale, P.J. (1993). Black-hat/white-hat strategy in bilateral negotiation. *Organizational Behavior and Human Decision Processes* 55, 444-469.
- Isoni, A. (2011). The willingness-to-accept/willingness-to-pay disparity in repeated markets: loss aversion or 'bad-deal' aversion? *Theory and Decision*, 71(3), 409-430.
- Karakostas, A., Zizzo, D.J. (2016). Compliance and the power of authority. *Journal of Economic Behavior and Organization* 124, 67-80.
- Kawagoe, T., Narita, Y. (2014). Guilt aversion revisited: an experimental test of a new model. *Journal of Economic Behavior and Organization* 102, 1-9.
- Khalmetski, K. (2016). Testing guilt aversion with an exogenous shift in beliefs. *Games and Economic Behavior* 97, 110–119.
- Khalmetski, K., Danilov, A., Sliwka, D. (2018). Norms and guilt. *CESifo Working Paper Series* No. 6999
- Khalmetski, K., Ockenfels, A., Werner, P. (2015). Surprising gifts: Theory and laboratory evidence. *Journal of Economic Theory* 159, 163-208.
- Kocher, M.G., Poulsen, O., Zizzo D.J. (2017). Social preferences, accountability, and wage bargaining. *Social Choice and Welfare* 48, 659-678.
- Konow, J. (2005). Blind spots: The effects of information and stakes on fairness bias and dispersion. *Social Justice Research* 18, 349-390.
- Korobkin, R., Guthrie, C. (2003-2004). Heuristics and biases at the bargaining table. *Marquette Law Review* 87, 795-808.
- Kristensen, H. Gärling, T. (2000). Anchor points, reference points, and counteroffers in negotiations. *Group Decision and Negotiation* 9, 493–505.
- Levy, J.S. (1997). Prospect theory, rational choice, and international relations. *International Studies Quarterly* 41, 87-112.
- Pelligra, V., Reggiani, T., Zizzo, D.J. (2016). *Responding to (Un)Reasonable Requests*. IZA Discussion Paper No. 10189.
- Regner, T., Matthey, A. (2015). Do reciprocators exploit or resist moral wiggle room? An experimental analysis. Jena Economics Research Paper 2015-27.
- Reuben, E., Sapienza, P., Zingales, L. (2009). Is mistrust self-fulfilling? *Economics Letters* 104, 89-91.
- Rustichini, A., Villeval, M. C. (2014). Moral hypocrisy, power and social preferences. *Journal of Economic Behavior & Organization*, 107, 10-24.
- Sally, D. (1995). Communication and cooperation in social dilemmas: A meta-analysis of experiments from 1958 to 1992. *Rationality and Society* 7(1), 58-92.
- Schotter, A., Sopher, B. (2007). Advice and behavior in intergenerational ultimatum games: An experimental approach. *Games and Economic Behavior*, 58(2), 365-393.
- Siegel, S., Fouraker, L. E. (1960). *Bargaining and group decision making*. New York: McGraw-Hill.
- Silverman, D., Slemrod, J., Uler, N. (2014). Distinguishing the role of authority in and authority to. *Journal of Public Economics* 113(1), 32-42.
- Sitzia, S., Zizzo, D.J. (2012). Price lower and then higher or price higher and then lower? *Journal of Economic Psychology* 33(6), 1084-1099.

- Sliwka, D., Werner, P. (2017). Wage increases and the dynamics of reciprocity. *Journal of Labor Economics*, 35(2), 299-344.
- Sonntag, A., Zizzo, D.J. (2015). Institutional authority and collusion. *Southern Economic Journal*.
- Valdesolo, P., DeSteno, D. (2008). The Duality of Virtue: Deconstructing the Moral Hypocrite. *Journal of Experimental Social Psychology* 44, 1334-1338.
- Vanberg, C. (2008). Why do people keep their promises? An experimental test of two explanations *Econometrica*, 76, 1467-1480
- Zizzo, D.J. (2003). Empirical Evidence on Interdependent Preferences: Nature or Nurture? *Cambridge Journal of Economics* 27(6), 867-880.
- Zizzo, D.J. (2010). Experimenter demand effects in economic experiments. *Experimental Economics* 13, 75–98.

Figure 1: Mean Trustworthiness by Increasing / Decreasing / Random order



Notes: This figure contains the mean points returned for each request level in the *Inc* (panel A), *Dec* (panel B) and *Rand* (panel C) treatments. The horizontal dotted red line indicates the corresponding round 1 baseline return rate.

Table 1: Factorial design, number of subjects by treatments

	<i>NoCom</i>	<i>Com</i>	(total)
<hr/>			
<i>Inc</i>	20	20	40
<i>Dec</i>	20	20	40
<i>Rand</i>	20	20	40
(total)	60	60	120

Table 2: Descriptives

	<i>NoCom Inc</i>	<i>Com Inc</i>	<i>NoCom Dec</i>	<i>Com Dec</i>	<i>NoCom Rand</i>	<i>Com Rand</i>
Round 1- baseline						
(no request)	[50.5] (45%)	[50.5] (45%)	[63.5] (65%)	[65] (65%)	[63] (60%)	[53.35] (45%)
Request 0%						
(0 pts)	[31.75] (50%)	[35.3] (20%)	[56.5] (65%)	[47.75] (30%)	[56.85] (60%)	[29.25] (20%)
Request 20%						
(30 pts)	[42.5] (50%)	[39.75] (40%)	[54] (60%)	[41.25] (45%)	[60.4] (55%)	[32.5] (35%)
Request 40%						
(60 pts)	[40] (25%)	[48] (70%)	[56] (60%)	[52.65] (50%)	[51.55] (55%)	[45.75] (60%)
Request 60%						
(90 pts)	[42.5] (60%)	[39.75] (60%)	[71.75] (45%)	[48.85] (70%)	[71.8] (70%)	[58.1] (60%)
Request 80%						
(120 pts)	[48.75] (40%)	[52.5] (60%)	[67.25] (50%)	[59] (60%)	[77.05] (60%)	[59.5] (75%)
Request 100% (150 pts)	[46.5] (35%)	[39.75] (35%)	[70] (55%)	[67.25] (55%)	[70.95] (55%)	[67.25] (65%)
Round 8						
(no request)	[43] (50%)	[59.25] (40%)	[63] (55%)	[55] (80%)	[64.8] (50%)	[52] (45%)

Notes: In square brackets: Mean trustworthiness rate - points returned by the trustee to the trustor out of 150 points, by request/treatment. In standard brackets: Trust rate - share of trustors sending 50 points to trustees.

Table 3: Tobit regressions on trustworthiness (rounds 1-8)

Outcome: Points returned	[1]	[2]	[3]	[4]	[5]	[6]	[7]	[8]
Reasonable trustworthiness rate					1.037*** (0.08)	1.036*** (0.08)	1.047*** (0.08)	1.046*** (0.08)
<i>Inc</i>	-29.116* (14.94)	-29.102* (14.93)	-31.730* (16.20)	-31.707* (16.19)	-17.022** (7.72)	-17.010** (7.71)	-16.611** (8.29)	-16.589** (8.27)
<i>Dec</i>	-5.016 (15.11)	-5.030 (15.11)	-6.812 (15.34)	-6.822 (15.34)	-7.123 (7.15)	-7.162 (7.15)	-6.363 (7.36)	-6.400 (7.36)
<i>Com</i>	-21.202 (14.89)	-21.240 (14.89)	-23.915 (16.59)	-23.944 (16.58)	-10.331 (8.79)	-10.381 (8.80)	-11.370 (9.78)	-11.410 (9.77)
<i>Com*Inc</i>	24.784 (21.08)	24.791 (21.07)	31.429 (24.92)	31.416 (24.91)	15.916 (11.60)	15.931 (11.58)	16.479 (12.74)	16.475 (12.72)
<i>Com*Dec</i>	12.107 (20.81)	12.146 (20.81)	13.552 (21.00)	13.586 (21.00)	0.649 (11.28)	0.706 (11.28)	0.746 (11.19)	0.799 (11.19)
Request_0	-19.631*** (3.87)	-21.466*** (5.29)	-19.621*** (3.89)	-21.470*** (5.31)	-18.620*** (3.72)	-20.442*** (5.14)	-18.567*** (3.72)	-20.384*** (5.15)
Request_30	-14.185*** (4.08)	-13.876** (5.70)	-14.150*** (4.07)	-13.888** (5.71)	-13.060*** (3.94)	-13.339** (5.60)	-13.103*** (3.94)	-13.374** (5.59)
Request_60	-8.959*** (2.74)	-11.215* (6.44)	-8.941*** (2.73)	-11.215* (6.45)	-8.899*** (2.76)	-11.396* (6.37)	-8.954*** (2.75)	-11.436* (6.35)
Request_90	-2.256 (3.79)	-4.727 (5.92)	-2.210 (3.79)	-4.699 (5.93)	-1.514 (3.73)	-4.218 (5.84)	-1.532 (3.72)	-4.211 (5.83)
Request_120	2.946 (4.05)	2.586 (5.18)	2.972 (4.06)	2.566 (5.20)	4.093 (3.96)	3.176 (5.09)	4.059 (3.96)	3.156 (5.09)
Request_150	3.481 (4.76)	1.515 (5.74)	3.470 (4.76)	1.492 (5.75)	4.293 (4.59)	2.348 (5.60)	4.291 (4.59)	2.348 (5.60)
round 8		-2.027 (3.07)		-2.040 (3.09)		-1.667 (3.15)		-1.742 (3.12)
round dummies	NO	YES	NO	YES	NO	YES	NO	YES
demographics	NO	NO	YES	YES	NO	NO	YES	YES
Constant	66.794*** (10.63)	67.829*** (10.81)	68.885*** (15.57)	69.908*** (15.78)	1.326 (8.23)	2.242 (7.74)	3.302 (9.51)	4.233 (9.09)
Pseudo R ²	0.006	0.007	0.007	0.007	0.088	0.088	0.089	0.089
Obs.	960	960	960	960	960	960	960	960

Notes: Outcome, points returned. Tobit regressions. Robust standard errors clustered at the individual level are reported in brackets. Models 1 - 8 include observations from all rounds in the first stage; Random and Round 1 without request are captured in the Constant term. Three stars, two stars, and one star refer to significant effects at the 1%, 5%, and 10% level respectively.

Table 4: Tobit regressions on trustworthiness (rounds 2-7)

Outcome: Points returned	[9]	[10]	[11]	[12]	[13]	[14]	[15]	[16]
Reasonable trustworthiness rate					0.986*** (0.09)	0.986*** (0.09)	0.998*** (0.09)	0.998*** (0.09)
<i>Inc</i>	-31.419** (14.90)	-31.392** (14.89)	-33.493** (16.00)	-33.453** (15.98)	-20.174** (9.10)	-20.142** (9.09)	-19.400** (9.67)	-19.355** (9.66)
<i>Dec</i>	-5.976 (15.21)	-5.990 (15.21)	-7.284 (15.38)	-7.293 (15.38)	-8.217 (8.28)	-8.254 (8.28)	-7.134 (8.51)	-7.168 (8.52)
<i>Com</i>	-24.276 (15.42)	-24.318 (15.42)	-26.568 (17.03)	-26.597 (17.02)	-13.920 (10.65)	-13.970 (10.66)	-14.631 (11.71)	-14.664 (11.72)
<i>Com*Inc</i>	24.967 (21.78)	24.961 (21.77)	30.563 (25.29)	30.526 (25.28)	16.777 (14.32)	16.775 (14.32)	16.663 (15.43)	16.629 (15.43)
<i>Com*Dec</i>	13.003 (21.48)	13.049 (21.48)	14.295 (21.62)	14.333 (21.62)	2.353 (13.58)	2.412 (13.59)	2.331 (13.44)	2.382 (13.46)
Request_0	-23.355*** (5.91)	-23.216*** (5.87)	-23.327*** (5.92)	-23.190*** (5.89)	-23.296*** (5.75)	-23.163*** (5.71)	-23.234*** (5.74)	-23.098*** (5.71)
Request_30	-17.759*** (5.58)	-15.437** (5.99)	-17.729*** (5.59)	-15.436** (6.02)	-17.481*** (5.45)	-15.706*** (5.91)	-17.533*** (5.46)	-15.748*** (5.91)
Request_60	-12.454** (5.03)	-12.745** (6.44)	-12.430** (5.03)	-12.720** (6.45)	-13.186*** (4.93)	-13.734** (6.33)	-13.242*** (4.92)	-13.770** (6.31)
Request_90	-5.771 (4.94)	-6.282 (5.65)	-5.725 (4.94)	-6.235 (5.65)	-5.848 (4.75)	-6.612 (5.45)	-5.871 (4.75)	-6.606 (5.45)
Request_120	-0.571 (4.45)	1.070 (4.87)	-0.535 (4.46)	1.079 (4.88)	-0.294 (4.29)	0.826 (4.70)	-0.325 (4.28)	0.811 (4.69)
round dummies	NO	YES	NO	YES	NO	YES	NO	YES
demographics	NO	NO	YES	YES	NO	NO	YES	YES
Constant	72.433*** (10.72)	73.313*** (11.08)	75.003*** (16.04)	75.855*** (16.06)	10.733 (10.55)	11.664 (11.03)	13.090 (12.00)	13.956 (12.36)
Pseudo R ²	0.008	0.008	0.008	0.008	0.072	0.073	0.073	0.073
Obs.	720	720	720	720	720	720	720	720

Notes: Outcome, points returned. Tobit regressions. Robust standard errors clustered at the individual level are reported in brackets. Models 9 - 16 include observations from round 2 to 7 in the first stage. Robust standard errors clustered at the individual level are reported in brackets. Random and Request_150 are captured in the Constant term. Three stars, two stars, and one star refer to significant effects at the 1%, 5%, and 10% level respectively.

Table 5: Probit regressions on trust (rounds 1-8)

Outcome: Trust probability	[1]	[2]	[3]	[4]	[5]	[6]	[7]	[8]
	<i>Communication</i>				<i>NoCommunication</i>			
Reasonable trustworthiness rate	0.003*** (0.001)	0.003*** (0.001)	0.003*** (0.001)	0.003*** (0.001)	0.008*** (0.002)	0.008*** (0.002)	0.009*** (0.002)	0.009*** (0.002)
<i>Inc</i>	-0.039 (0.11)	-0.038 (0.11)	-0.122 (0.11)	-0.121 (0.11)	-0.083 (0.12)	-0.087 (0.12)	-0.012 (0.12)	-0.014 (0.12)
<i>Dec</i>	0.032 (0.11)	0.031 (0.11)	0.019 (0.11)	0.019 (0.11)	-0.055 (0.12)	-0.059 (0.12)	-0.006 (0.12)	-0.009 (0.12)
Request_0	-0.316*** (0.06)	-0.343*** (0.07)	-0.321*** (0.06)	-0.345*** (0.08)	0.007 (0.04)	-0.092 (0.07)	0.005 (0.05)	-0.104 (0.08)
Request_30	-0.138** (0.06)	-0.198** (0.08)	-0.140** (0.06)	-0.199** (0.08)	-0.021 (0.07)	-0.08 (0.10)	-0.024 (0.07)	-0.088 (0.10)
Request_60	0.067 (0.07)	-0.006 (0.11)	0.069 (0.07)	-0.003 (0.12)	-0.018 (0.07)	-0.032 (0.09)	-0.020 (0.07)	-0.035 (0.09)
Request_90	0.102 (0.07)	0.029 (0.11)	0.105 (0.07)	0.033 (0.12)	0.071 (0.06)	-0.081 (0.09)	0.069 (0.07)	-0.09 (0.1)
Request_120	0.119 (0.08)	0.05 (0.10)	0.121 (0.08)	0.053 (0.10)	0.033 (0.07)	-0.021 (0.15)	0.032 (0.07)	-0.023 (0.12)
Request_150	-0.018 (0.08)	-0.049 (0.10)	-0.018 (0.08)	-0.047 (0.10)	-0.039 (0.07)	-0.136 (0.10)	-0.045 (0.08)	-0.150 (0.10)
round 8		0.035 (0.06)		0.036 (0.06)		-0.070 (0.07)		-0.078 (0.07)
round dummies	NO	YES	NO	YES	NO	YES	NO	YES
demographics	NO	NO	YES	YES	NO	NO	YES	YES
Constant (Prediction)	0.510	0.511	0.513	0.514	0.538	0.540	0.555	0.557
Pseudo R ²	0.083	0.088	0.103	0.123	0.228	0.237	0.269	0.280
Obs.	480	480	480	480	480	480	480	480

Notes: Outcome, trust=1, not trust=0. Probit regressions, marginal effects reported; Robust standard errors clustered at individual level are reported in brackets. Three stars, two stars and one star for significant effects at the 1%, 5%, and 10% level respectively. Models 1-8 employ data from all rounds of stage 2. Models 1-4 employ *Com* treatments data; Models 5-8 employ *NoCom* treatments data with ‘virtual’ request dummies (for any given round, the ‘virtual’ request relates to the actual request submitted to the trustee but not revealed to the trustor).

Table 6: Probit regressions on trust (rounds 2-7)

Outcome: Trust probability	[1]	[2]	[3]	[4]	[5]	[6]	[7]	[8]
	<i>Communication</i>				<i>NoCommunication</i>			
Reasonable trustworthiness rate	0.002 (0.001)	0.002* (0.001)	0.002 (0.001)	0.002 (0.001)	0.007*** (0.002)	0.007*** (0.002)	0.008*** (0.002)	0.009*** (0.002)
<i>Inc</i>	-0.048 (0.10)	-0.047 (0.10)	-0.113 (0.13)	-0.112 (0.13)	-0.111 (0.12)	-0.116 (0.12)	-0.045 (0.13)	-0.05 (0.13)
<i>Dec</i>	-0.032 (0.11)	-0.034 (0.11)	-0.04 (0.12)	-0.04 (0.12)	-0.078 (0.12)	-0.083 (0.12)	-0.049 (0.12)	-0.054 (0.12)
Request_0	-0.297*** (0.07)	-0.300*** (0.08)	-0.299*** (0.07)	-0.301*** (0.08)	0.045 (0.07)	0.044 (0.07)	0.049 (0.07)	0.048 (0.07)
Request_30	-0.119 (0.08)	-0.149* (0.08)	-0.119 (0.08)	-0.151* (0.08)	0.018 (0.09)	0.055 (0.09)	0.021 (0.09)	0.061 (0.10)
Request_60	0.085 (0.07)	0.044 (0.1)	0.086 (0.08)	0.045 (0.1)	-0.014 (0.07)	-0.02 (0.07)	-0.015 (0.07)	-0.021 (0.08)
Request_90	0.119 (0.08)	0.077 (0.1)	0.121 (0.08)	0.078 (0.1)	0.107 (0.08)	0.053 (0.08)	0.113 (0.08)	0.059 (0.09)
Request_120	0.136* (0.07)	0.099 (0.07)	0.138* (0.07)	0.101 (0.07)	0.069 (0.07)	0.111 (0.09)	0.073 (0.07)	0.120 (0.09)
round dummies	NO	YES	NO	YES	NO	YES	NO	YES
demographics	NO	NO	YES	YES	NO	NO	YES	YES
Constant (Prediction)	0.502	0.503	0.504	0.505	0.534	0.535	0.550	0.553
Pseudo R ²	0.079	0.086	0.090	0.096	0.214	0.214	0.214	0.214
Obs.	360	360	360	360	360	360	360	360

Notes: Outcome, trust=1, not trust=0. Probit regressions, marginal effects reported; Robust standard errors clustered at individual level are reported in brackets. Three stars, two stars and one star for significant effects at the 1%, 5%, and 10% level respectively. Models 9-16 employ data from rounds 2-7 of stage 2. Models 9-12 employ *Com* treatments data; Models 13-16 employ *NoCom* treatments data with ‘virtual’ request dummies (for any given round, the ‘virtual’ request relates to the actual request submitted to the trustee but not revealed to the trustor).