



CENTRE FOR DECISION RESEARCH & EXPERIMENTAL ECONOMICS



The University of
Nottingham

UNITED KINGDOM • CHINA • MALAYSIA

Discussion Paper No. 2018-13

Hanna Fromell, Daniele
Nosenzo and Trudy Owens

December 2018

**Altruism, Fast and Slow?
Evidence from a Meta-Analysis
and a New Experiment**

CeDEx Discussion Paper Series

ISSN 1749 - 3293



CENTRE FOR DECISION RESEARCH & EXPERIMENTAL ECONOMICS

The Centre for Decision Research and Experimental Economics was founded in 2000, and is based in the School of Economics at the University of Nottingham.

The focus for the Centre is research into individual and strategic decision-making using a combination of theoretical and experimental methods. On the theory side, members of the Centre investigate individual choice under uncertainty, cooperative and non-cooperative game theory, as well as theories of psychology, bounded rationality and evolutionary game theory. Members of the Centre have applied experimental methods in the fields of public economics, individual choice under risk and uncertainty, strategic interaction, and the performance of auctions, markets and other economic institutions. Much of the Centre's research involves collaborative projects with researchers from other departments in the UK and overseas.

Please visit <http://www.nottingham.ac.uk/cedex> for more information about the Centre or contact

Suzanne Robey
Centre for Decision Research and Experimental Economics
School of Economics
University of Nottingham
University Park
Nottingham
NG7 2RD
Tel: +44 (0)115 95 14763
suzanne.robey@nottingham.ac.uk

The full list of CeDEX Discussion Papers is available at

<http://www.nottingham.ac.uk/cedex/publications/discussion-papers/index.aspx>

Altruism, Fast and Slow?

Evidence from a Meta-Analysis and a New Experiment

Hanna Fromell^a, Daniele Nosenzo^{b,c,*} and Trudy Owens^c

^a Department of Economics, Econometrics, and Finance, University of Groningen

^b Luxembourg Institute of Socio-Economic Research (LISER)

^c School of Economics, University of Nottingham

30 November 2018

Abstract:

Can we use the lens of dual-system theories to explain altruistic behavior? In recent years this question has attracted the interest of both economists and psychologists. We contribute to this emerging literature, by reporting both the results of a meta-study of the literature and a new experiment. Our meta-study is based on 19 experimental studies conducted with nearly 11,000 subjects. We show that the overall effect of manipulating cognitive resources to promote the “intuitive” system at the expense of the “deliberative” system is very close to zero. We argue that this null result could be because the interventions used in the existing literature to manipulate cognitive resources are vulnerable to the presence of heterogeneity in the direction of the effect of the intervention. We design a new experiment that is not vulnerable to this potential heterogeneity. We still fail to find support for the notion that altruistic choices are the result of a conflict between the intuitive and deliberative systems. Taken together, the findings of our meta-study and the new experiment offer little support for dual-system theories of altruistic behavior.

Keywords: altruism; giving; dictator game; dual-system model; intuition; deliberation; self-control; willpower; depletion; Stroop task

JEL Classification: C91

Acknowledgements: We thank Valerio Capraro, Dirk Engelmann, Fabio Tufano and participants at the CeDEx Brownbag seminar, the 2015 Inaugural RES Symposium of Junior Researchers, and the 2015 ESA European meeting, for helpful comments. We thank Valerio Capraro and Jonathan Schulz for sharing the data of their studies with us. We acknowledge support from the University of Nottingham and the British Academy & Leverhulme Trust (SG140436). The paper was previously circulated under the title: “Do Tradeoffs between Self-interest and Other-Regarding Preferences Cause Willpower Depletion?”.

* Corresponding author at: Luxembourg Institute of Socio-Economic Research (LISER), 11 Porte des Sciences, 4366 Esch-sur-Alzette, Luxembourg.

E-mail addresses: h.c.fromell@rug.nl (H. Fromell), daniele.nosenzo@liser.lu; daniele.nosenzo@nottingham.ac.uk (D. Nosenzo), trudy.owens@nottingham.ac.uk (T. Owens).

1. INTRODUCTION

A recent literature in economics and psychology argues that human behavior can be understood as the interaction between two different decision systems (Kahneman, 2002; 2011): one that is fast, intuitive, automatic and largely effortless (“System 1”), and one that is slower, more deliberate, and requires some level of reflection and cognitive effort (“System 2”). For some decisions, the two systems may diverge in the choices they favor: while deliberation may pull the individual towards a certain choice (e.g., keeping a healthy diet), intuition may pull her towards a different choice (e.g., eating tasty but highly caloric food). In these cases, an individual must use willpower and spend cognitive resources to override the intuitive impulse and take the choice favored by the deliberative system.

It has been suggested that altruistic behavior can also be rationalized by this dual-system framework (Loewenstein and O’Donoghue, 2007; Moore and Loewenstein, 2004; Zaki and Mitchell, 2013; Deck and Jahedi, 2015; Dreber et al., 2016). One of the most debated questions in this literature is whether altruism (and pro-social behavior more generally) is a spontaneous and intuitive response, where individuals must use willpower if they wish to act in their self-interest; or whether instead self-interest is intuitive, and individuals must use willpower to behave pro-socially.

A number of experiments have been designed to address this question. These experiments rely on different types of manipulations, designed to inhibit one of the systems and promote the other. For example, subjects in an experiment may be asked to distribute an amount of money between themselves and another participant (e.g., in a dictator game), while at the same time performing another task that is more or less cognitively-taxing (e.g., holding a 7-digit or 3-digit number in their memory). Compared to subjects who perform the easier task, subjects who perform the hard task are under “cognitive load”: they have fewer cognitive resources to devote to the dictator game decision, and are therefore less able to use their deliberative system when deciding how much to give to the other participant. Other commonly used types of manipulations include: “ego depletion” where subjects participate in a sequence of two tasks, the first to deplete cognitive resources, and the second to measure how the consequent reduced ability to use deliberation affects behavior; “time pressure” where subjects are forced to make a decision either

quickly or after having deliberated for some time; and “priming” where subjects are encouraged to decide using either their intuitive or deliberative system.

In the first part of this paper we show that the existing literature does not provide a clear answer to the question of whether intuition favors altruism or self-interest - or indeed whether a dual-system framework is at all suitable for rationalizing altruistic behavior. We conduct a meta-analysis of the experimental studies that have used cognitive load, ego depletion, time pressure, and priming to study the effects of promoting intuition on altruistic behavior. Our meta-study covers 19 papers involving a total of 55 experiments and nearly 11,000 subjects. We find that in approximately 58% of the experiments promoting intuition leads to more self-interested behavior, suggesting that self-interest is an intuitive response. In the other 42% of experiments promoting intuition encourages more altruistic behavior, suggesting that altruism, and not self-interest, is intuitive. These effects, however, are statistically significant only in a minority of studies. In the large majority of cases (78% of experiments), the effect of promoting intuition is small and the 95% confidence interval of the effect size includes zero. The overall effect size estimated across all 55 experiments is only -0.030, and we cannot reject the null hypothesis that this is zero.

It is unclear how to interpret this evidence. On the one hand, taken at face value, the fact that the literature has found effects in either direction, and an average estimated effect close to zero, may suggest that the true underlying effect is actually very small or non-existent. That is, altruism is neither fast nor slow: this type of behavior simply escapes the logic of the dual-system framework. On the other hand, some researchers have argued that whether intuition favors altruism or self-interest may depend on a variety of individual, social, and contextual factors. While for some subgroups of individuals in some specific situations intuition may favor altruism, for other subgroups or situations intuition may favor self-interest (e.g., Hauge et al., 2016; Rand et al., 2016; Grossman and Van der Weele, 2017; Balafoutas et al., 2018).¹ That is, the heterogeneous effects reported in the literature could simply reflect a genuine heterogeneity in

¹ Indeed, in our meta-study we find that promoting intuition seems to have a different effect on men and women. This is in line with Rand et al. (2016), who also conducted a meta-study of the literature on role of intuition and deliberation for altruism. They found that promoting intuition does not significantly affect men, but significantly increases altruism among women. However, our meta-study, which subsumes most of their data, finds the opposite: in our data promoting intuition does not significantly affect women, but marginally significantly decreases altruism among men. See Section 2 for further details.

the underlying decision processes that different individuals use in different decision situations – and since this heterogeneity is typically unanticipated and unaccounted for in most existing studies, this could explain why the reported effect is small and close to zero in most of the literature.

In the second part of the paper, we propose a new experimental paradigm to resolve this conundrum. Our starting point is the observation that a key tenet of dual-system theories of pro-social behavior is that choices involving trade-offs between altruism and self-interest trigger a conflict between the intuitive and deliberative systems, and that the individual must consume cognitive resources and willpower to regulate this conflict. Note that this is true regardless of whether one's theory proposes that altruism is intuitive and self-interest deliberate, or whether it proposes instead that self-interest is intuitive and altruism deliberate. In the former case, willpower is required to rein in the pro-social impulse; in the latter it serves to override the selfish impulse. In either case, the individual will have to use willpower when she faces decisions that involve a trade-off between altruism and self-interest. In other words, these trade-offs are *willpower-depleting*. This implies a straightforward prediction: individuals exposed to trade-offs between altruism and self-interest will have less willpower available for *subsequent* tasks that also require the use of willpower. A crucial advantage of this prediction is that it holds irrespective of whether altruism or self-interest are intuitive: it only requires that the intuitive process diverges from the deliberate process so that a conflict between the two systems arises.

We test this prediction by reversing the order of tasks typically used in ego depletion experiments: we use a dictator game as the first, willpower-depleting task, and a cognitively-demanding and willpower-requiring task as the second task, to measure the effect of participating in the dictator game on residual willpower. Across two between-subject treatments we manipulate the extent to which the dictator game requires exertion of willpower by varying the structure of payoffs. In our Conflict treatment, the option that maximizes the dictator's payoff always minimizes the recipient's payoff. Thus, in this treatment dictators face a trade-off between self-interest and altruistic choices, as is customary in standard dictator games. In our NoConflict treatment, we remove this trade-off by manipulating payoffs so that now the option that maximizes the dictator's payoff also maximizes the recipient's payoff. In both treatments, subjects then participate in a version of the Stroop (1935) color-word task, which is often used to

measure willpower depletion. If trade-offs between self-interested and other-regarding choices require the use of willpower, we would expect dictators in Conflict to use on average more willpower in the dictator game than those in NoConflict, and to be therefore comparatively less able to expend further willpower in the subsequent Stroop task.

We find small and statistically insignificant differences in Stroop task performance between dictators who were exposed to trade-offs between altruism and self-interest and those who were not exposed to such trade-offs. These results suggest that dictator game choices are not willpower-depleting, which questions the notion that pro-social behavior is the result of a competition between intuitive and deliberative choice processes. Together with the mixed evidence reported in the previous literature, our results offer little support for a dual-system theory of altruistic behavior.

The remainder of this paper is structured as follows. In Section 2 we present the findings of our meta-analysis of the existing literature. Section 3 describes the design and results of our new experiment. Section 4 concludes.

2. META-STUDY

2.1 Design

We searched the literature for experimental studies investigating the effect of promoting intuition on altruistic behavior. We include studies that use one of the four standard types of interventions to promote intuition at the expense of deliberation (cognitive load, ego depletion, time pressure, or priming), and that assess the effect of these interventions on individuals' decisions to distribute wealth between themselves and another passive player.² The passive player can be another participant in the experiment (as in standard dictator games), or a charitable organization (as in donation experiments). In all cases, the decisions must involve a trade-off between the decision-maker's and passive player's payoffs (i.e., we exclude settings in which the decision-maker can increase the passive player's payoff at no cost for themselves; or cases where the choice that maximizes the decision-maker's payoff also maximizes the passive player's payoff). We require

² The proliferation of studies using cognitive resources manipulations has triggered a number of recent meta-studies of these literatures. For example, Rand (2016) assess the role of intuition and deliberation in cooperation decisions. Verschuere et al. (2018) study the role of cognitive load interventions on lying. Rand et al. (2016) – discussed in more detail below – study how the interaction between gender and intuition shapes dictator game decisions.

that the study adheres to the methodology of experimental economics (i.e., no deception), and that decisions have real monetary consequences for the parties involved (i.e., no hypothetical studies). For more details on the selection process, see Appendix A in the Online Supplementary Materials (OSM).

Based on our inclusion criteria, the meta-study covers 19 studies involving 55 experiments conducted with a total of 10,898 subjects across 11 countries.³ About half of the experiments were run with university students, a third with Amazon Mechanical Turk (AMT) workers, and the rest with other specific non-student samples, for instance junior school students or members of the general population. About two-thirds of experiments involved some type of dictator game decision (where the passive players are other experiment participants), while the rest involved a charitable donation decision. Table A.1 lists all included studies and the number of experiments each study contributed to the meta-study.⁴

For each experiment we quantified the effect that promoting intuition had on altruistic behavior by calculating the standardized mean difference (Cohen's d) in altruism between the experimental condition that promoted intuition and the condition that promoted deliberation.⁵ In most cases, we measured altruism as the monetary amount (or fraction of endowment) that the decision-maker gives to the passive player; in studies involving binary dictator decisions, we used the fraction of decision-makers sacrificing own payoff to increase the passive player's payoff.⁶ In all cases, a positive effect size indicates that individuals became more altruistic when intuition was promoted, relative to the condition that promoted deliberation. In contrast, a negative effect size indicates that promoting intuition led individuals to become more self-regarding.

³ Rand et al. (2016) also conducted a meta-analysis of the literature on the role of intuition in dictator game decisions, with the special focus on gender as mediator of the main effect. Their study covers 22 experiments involving 4,366 subjects. Rand and colleagues kindly shared their dataset with us, and we therefore include in our meta-study all experiments covered in their analysis, with the exception of 4 experiments from 2 studies that involve deception. In addition, our meta-analysis includes another 16 studies (37 experiments, 6,988 subjects) that were not included in Rand et al's meta-study. See Appendix A in the OSM for details.

⁴ In most cases, a single study contains more than one treatment (or was run in more than one location), hence most studies contribute more than one experiment each to the meta-study.

⁵ We apply Hedges' correction for small sample bias to calculate effect sizes and associated standard errors (see (Lipsey and Wilson 2001)).

⁶ The information required to perform these calculations was retrieved either from the statistics, graphs and tables reported in the papers, or, in a few cases, from the raw data made available by the authors (studies for which the full raw data are available are marked by * in Table A.1).

2.2 Results

Figure 1 contains a forest plot showing, for each of the 55 experiments included in our study, the associated effect size and 95% confidence interval. The figure is divided into four panels, one for each type of intervention included in the meta-study. In each panel, the bottom row reports the average effect size and associated confidence interval of each intervention, estimated using a random-effects meta-analysis model. The last row of the figure reports the overall effect size estimated across all four types of interventions.


For each type of intervention there are experiments associated with both negative and positive effect sizes. Across the four types of interventions, 58% of experiments report a negative effect of promoting intuition, while the remaining 42% report a positive effect. In most cases, however, the 95% confidence interval of the effect size includes zero. Only in 12 out of 55 experiments (22% of cases) does the confidence interval not contain zero: in 8 cases the estimated effect size is negative, and in the remaining 4 cases it is positive.

As a consequence, the average effect size of each type of intervention is close to zero: 0.005 for cognitive load studies, -0.054 for time pressure studies, -0.074 for ego depletion studies, and 0.012 for priming studies. In all cases, we cannot reject the null that the average effect size is zero (all $p > 0.185$). Across all types of intervention, the overall effect size is -0.030, with an associated 95% confidence interval of [-0.091, 0.030]. Similarly, we cannot reject the null that the overall effect size is actually zero ($z = 0.99$, $p = 0.322$).

Based on these results, what should we conclude about the role of intuition and deliberation for altruistic behavior? The face-value interpretation of our findings is that promoting intuition has only a very small (if any) effect on altruistic choices. Thus, one may conclude that the logic of the dual-system model does not seem to extend to altruistic behavior.

However, several researchers have suggested that the mixed evidence concerning the *overall* effect of promoting intuition on altruism may reflect a genuine heterogeneity in the size and direction of the effect across different subgroups of the population, or across different decision settings. That is, the researcher may observe no aggregate effect of the intervention on altruistic behavior, when in fact the manipulation may have led a subgroup of subjects to become more altruistic and another subgroup to become more self-regarding – with the two opposing effects cancelling each other out in the aggregate.

FIGURE 1
Results of the random-effects meta-analysis of promoting intuition on altruism



Note: Effect sizes (ES) measured as standardized mean difference in altruism between conditions where intuition or deliberation were promoted. Positive values imply more altruism in the intuitive condition. Error bars indicate 95% confidence intervals. The size of the grey boxes indicates the weight of the effect size in the meta-analysis (the relative weights are also reported in the last column of the figure). In each panel of the figure, the row labeled “Subtotal” reports the average effect size for each type of intervention, estimated by the random-effects meta-analysis model. The last row of the figure (labeled “overall”) reports the average effect size across all experiments and its associated confidence interval. For a legend of the study IDs refer to Table A.1 in the OSM

Rand et al. (2016) for instance, propose that what becomes automatized as an intuitive response depends on the strategies that are typically advantageous in one's daily social interactions. They argue that what constitutes a socially advantageous strategy may vary across individuals or groups, and propose that gender may be an important moderating factor in the case of altruism: altruism may be an intuitive social response for women, but less so for men. Other researchers have argued that the underlying pro-social inclinations of the individual may moderate the effect of intuition on altruism: while altruism may be an intuitive response for pro-social types, the opposite may be true for self-interested types (e.g., Balafoutas et al., 2018; Chen and Krajbich, 2018). Moreover, some researchers have focused on contextual and situational factors as possible moderators of the effect of promoting intuition on altruism. Andersen et al. (2018) and Mrkva (2017), for instance, examine the role of experimental stakes. Banker et al. (2017) propose that there is an interaction between the effect of promoting intuition and whether the decision situation uses a "giving" or "taking" frame: in the former case, manipulations that promote intuition may decrease altruism, while in the latter case they may increase altruism.

We use the meta-study to assess the role of individual and situation factors in explaining the variance in effect sizes across experiments included in our analysis. In Appendix A in the OSM we report moderator analysis examining a long list of factors discussed in the literature, namely: the role of gender; the type of intervention used to manipulate cognitive resources; the frame of the game (give or take); the nature of the passive player (another participant or a charity); the stakes used in the experiment, the type of subject pool used to conduct the study (e.g., students, AMT workers, etc.); and the location where the experiment was run.

From this list of factors we find only one significant effect. In line with the argument put forward by Rand et al. (2016), we observe a significant difference in effect sizes between men and women ($p = 0.020$)⁷, however, our results are in contrast. Rand et al.'s meta-study found that promoting intuition has a positive and significant effect for women, but an insignificant effect for men. Instead, we find that promoting intuition has no significant effect among women (average

⁷ This result is based on a subset of studies and experiments for which we can compute separate effect sizes for men and women. We can do so for 8 studies involving 26 experiments and 6,378 individuals.

effect size = 0.033; $p = 0.461$), but has a negative and marginally significant effect among men (average effect size = -0.076; $p = 0.068$).⁸

This result lends some support to the argument that there may be a genuine heterogeneity in the size and direction of the effect of intuition on altruism across different subsamples, and this may explain the mixed evidence reported in the literature about the overall effect of intuition on altruism. However, the differences between our findings and Rand et al.'s (2016) results also highlight that we still lack a clear understanding of the role of moderators in explaining differences in effect sizes. Moreover, and perhaps because of this, it is often difficult for researchers to anticipate how the direction of the effect may depend on subject pool characteristics or the details of the decision environment. Taken together, these considerations point to an important limitation of the existing experimental designs that rely on cognitive load, time pressure, ego depletion or priming to study whether intuition affects altruistic behavior.

Since these designs are vulnerable to the presence of (unanticipated and/or unobserved) heterogeneity in the *direction* of the effect triggered by the manipulation of cognitive resources, it is difficult to use these interventions to establish whether, if at all, the logic of dual-system theories extends to altruistic behavior. Given the current state of literature this seems a question of primary importance. In the next section, we propose a new experimental paradigm that allows us to address this question while circumventing the issue of heterogeneity that plagues existing experimental designs.

3. NEW EXPERIMENT

The starting point of our experimental design is the observation that, if choices that involve trade-offs between altruism and self-interest trigger a conflict between intuition and deliberation, then individuals exposed to such trade-offs will have to consume cognitive resources and willpower to regulate this conflict. This observation also underlies the rationale of existing experimental design used in previous studies. However, while these designs assess the importance of cognitive resources for altruistic choices by manipulating the extent to which

⁸ There are two differences between our sample and Rand et al.'s: 1) our sample size is larger because it contains studies that were not included in Rand et al.'s meta-analysis; 2) our inclusion criteria differ, and we exclude 2 studies involving deception which Rand et al. had instead used in their meta-analysis. In Appendix A, we show that what drives the difference in results is the larger sample size, and not the exclusion of the studies involving deception.

subjects have access to such resources *when making these choices*, we take the opposite approach and measure how making altruistic choices affects the cognitive resources that subjects have access to *in subsequent tasks*. That is, we measure whether facing trade-offs between altruism and self-interest depletes subjects of cognitive resources and willpower.

Note that our approach circumvents the issue of heterogeneity in the role of intuition and deliberation for altruistic behavior. The prediction that exposure to trade-offs between altruism and self-interest is willpower-depleting holds regardless of whether altruism or self-interest are intuitive. If altruism is intuitive and self-interest requires deliberation, then individuals will have to use willpower to take choices that maximize their self-interest. If instead self-interest is intuitive, and pro-sociality requires deliberation, then willpower will be spent to override the selfish impulse and reach an altruistic choice. In either case, as long as the logic of dual-system theories extends to altruistic behavior, the individual will have to use cognitive resources and willpower to regulate the conflict between the intuitive and deliberative systems. Hence, the presence of unobserved heterogeneity in the direction of the effect does not constitute a problem in our design, thus allowing us to use our experiment to probe the validity of dual-system theories of pro-social behavior.⁹

3.1 Experimental Design

The experiment is based on a dual-task paradigm akin to that used in ego depletion studies. However, in contrast to ego depletion experiments, we *reverse* the order of tasks that subjects take part in during the experiment. We use a series of 16 binary dictator games as the first task. We vary the structure of the 16 games across two treatments (Conflict and NoConflict) to manipulate whether or not subjects are exposed to a conflict between self-interest and altruism, and thus the extent to which they may need to consume willpower. To measure subjects' ability to exert further willpower after the games conducted in the first task, we use a version of the Stroop (1935) color-word task as our second task.

At the beginning of the first task subjects were randomly matched into pairs. Within each pair one subject was allocated the role of dictator and the other the role of recipient. Pairs and

⁹ Our design does not allow to study the direction of the effect – if we find one. However, at this stage, and given the mixed findings reported in the literature, our main interest is in establishing whether an effect exists at all.

roles were kept fixed across all 16 games. In each game, dictators could choose between two possible options, A or B, each implying a different distribution of money between themselves and the recipient. Our two between-subject treatments only differed in the payoff structure of the 16 games. Table 1 shows the different versions of the games used in the two treatments.¹⁰

TABLE 1
Payoffs in the binary dictator games

Game	Conflict		NoConflict	
	Option A	Option B	Option A	Option B
1	8; 8	12; 0	8; 0	12; 8
2	11; 1	8; 8	11; 8	8; 1
3	9; 1	7; 5	9; 5	7; 1
4	7; 5	8; 2	7; 2	8; 5
5	10; 6	12; 0	10; 0	12; 6
6	8; 8	16; 0	8; 0	16; 8
7	6; 6	8; 4	6; 4	8; 6
8	10; 2	9; 7	10; 7	9; 2
9	12; 0	6; 6	12; 6	6; 0
10	11; 5	15; 1	11; 1	15; 5
11	8; 4	10; 2	8; 2	10; 4
12	9; 3	8; 4	9; 4	8; 3
13	10; 6	8; 8	10; 8	8; 6
14	13; 3	10; 6	13; 3	10; 3
15	6; 6	8; 2	6; 2	8; 6
16	10; 0	6; 4	10; 4	6; 0

Note: in each cell, the first number indicates the dictator's payoff and the second number the recipient's payoff (both displayed in GBP). In both treatments, the 16 games were presented to subjects in the same order as shown in the Table.

In each game of the Conflict treatment, the option that gave the highest possible payoff to the dictator gave the lowest possible payoff to the recipient. For example, in game 3 of Table 1 option A gave a payoff of £9 to the dictator and a payoff of £1 to the recipient, whereas option B

¹⁰ We had dictators face 16 decision problems because there is evidence in the ego depletion literature that the extent to which willpower is depleted is positively related to the duration of the depleting task (Hagger et al. 2010). We varied the payoffs from the two alternatives across the 16 games so that dictators were forced to reconsider their choice each time a new game was presented to them. Payoffs were manipulated to vary: i) the dictator's endowment (either £10, 12 or 16); ii) the opportunity cost of making the choice that most benefited the recipient; iii) whether an equal split of money was available in the choice set; and iv) whether making the own payoff-maximizing choice reduced joint payoffs.

gave payoffs of £7 and £5 to the dictator and the recipient respectively. In contrast, in the NoConflict treatment the recipient's payoffs were flipped across options A and B of each game, while the dictator's payoffs were the same as in Conflict. In the version of game 3 used in NoConflict, the recipient received £5 from option A (and the dictator £9) and £1 from option B (and the dictator £7). Thus, in each game of the NoConflict treatment, the option that gave the highest possible payoff to the dictator also gave the highest possible payoff to the recipient.

Hence, our two treatments differ in whether dictators face a trade-off between their own payoff and the recipient's payoff in each of the 16 games of the first task. In the Conflict treatment, dictators whose intuitive and deliberative motives differ in how much they care about others' well-being are exposed to a motivational conflict in each of the 16 games. Note that this holds regardless of how the intuitive and deliberative motives differ: for some dictators, the conflict may arise because their intuition is self-regarding, while for others conflict may arise because their intuition is altruistic. Either way, these dictators will face a motivational conflict in the first task of the experiment. On the other hand, such motivational conflict cannot arise for any of the dictators in the NoConflict treatment, because in this treatment there is no trade-off between the dictator's payoff and the recipient's payoff.¹¹ If the logic of dual-system theories extends to altruistic behavior, this implies that dictators will use more willpower to regulate their behavior across the 16 games of the Conflict treatment compared to the NoConflict treatment. Since willpower is considered as a resource that can be depleted with use (Muraven and Baumeister, 2000; Baumeister et al., 2007), this leads to the following hypothesis:

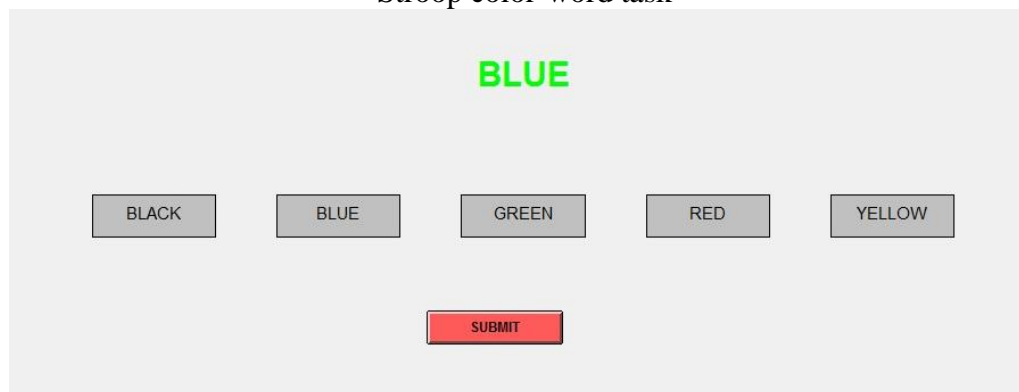
Hypothesis 1: *Dictators in Conflict will have on average less willpower available after the first task compared to dictators in NoConflict.*

We test our hypothesis by measuring how our treatments affect dictators' ability to exert willpower in a computerized version of the Stroop (1935) color-word task that was administered in the second part of the experiment. In the task subjects were shown a series of words of color

¹¹ In this sense, the games of the Conflict treatment are less "trivial" than those in NoConflict, because the former involve trade-offs between different motivations while the latter do not. Note, however, that this definition of task triviality relies on the concept of trade-offs, i.e. a task becomes less trivial if it involves a larger trade-off between self-interest and altruistic concerns. In terms of pure computational complexity, both tasks are equally "trivial" in the sense that they require very simple numerical comparisons between alternative options.

names printed in various font colors. Both color names and font colors could be either black, blue, green, red, or yellow. The font colors and color names were randomly matched so that, for each word, the font color did not correspond to the color name. For example, as shown in Figure 2, the word “blue” could have been displayed to subjects in green. Subjects had to indicate, for each word, which color it was printed in, and not the color that the word read. For example, for the word shown in Figure 2, the correct answer was “green” and not “blue”. If a correct answer was given, a new word appeared on the screen; if the answer was incorrect the same word remained on the screen. Subjects performed the Stroop task for five minutes and were rewarded with £0.10 for each correct answer.

FIGURE 2
Stroop color-word task



The Stroop color-word task has been used frequently in the literature to measure or induce willpower depletion.¹² The Stroop task requires regulation of choice because, in order to submit a correct answer, subjects must override their intuitive and automatic impulse to respond by reading the color name of the word, and look at its font color instead. Thus, the task demands continuous exertion of willpower by the participant to regulate the conflict that arises between their intuitive and deliberative systems. We will thus test our hypothesis that dictators in the

¹² See MacLeod (1991) for a review of studies using the Stroop task. In dual-task studies, the Stroop task has been used both as a task to measure willpower depletion and as a willpower-depleting task. Hagger et al. (2010) show that the task is particularly sensitive to willpower depletion relative to other commonly used tasks (e.g. math or mental arithmetic or solvable anagrams) and thus is well suited as a dependent measure. There is some evidence that the task is instead less effective when used as a willpower-depleting instrument (Hagger et al., 2010; Carter et al., 2015). In the economics literature, the Stroop task has been predominantly used as a willpower-depleting task to study, for example, the effects of willpower depletion on procrastination (Burger et al., 2011), sensitivity to framing effects (De Haan and Van Veldhuizen 2015) and time preferences (Kuhn et al., 2014).

Conflict treatment have lower willpower after part one than dictators in NoConflict, by comparing the number of correct answers in the Stroop task between the two treatments.

3.2 Experimental Procedures

The experiment was programmed in z-Tree (Fischbacher, 2007) and was conducted at the University of Nottingham using students from a wide range of disciplines recruited through the online recruitment system ORSEE (Greiner, 2015). We have 396 subjects in total, equally divided across the two treatments. Thus, our analysis is based on the choices of 99 dictators per treatment. This implies that we can detect effects of size 0.40 or larger with 80% power and 5% probability of a type-I error, using a two-sided Wilcoxon-Mann-Whitney test.

We conducted thirteen sessions in total. In each session, we randomly allocated subjects to either the Conflict treatment or the NoConflict treatment. Whether a subject was allocated to the Conflict or NoConflict treatment depended on the computer terminal he or she was randomly assigned to at the beginning of the session. All instructions and procedures during the session were kept identical across treatments. The only difference between treatments was in the payoffs of the 16 dictator games, which were shown to subjects privately on their computer screens.

At the start of the session, the experimenter distributed and read aloud preliminary instructions explaining the general two-part structure of the experiment (see Appendix B in the OSM for a copy of the instructions). Participants were then provided with detailed instructions for part one, which were again read aloud. These instructions included a series of control questions aimed at testing subjects' understanding of the task. Part one was started once all subjects had answered all questions correctly.

At the beginning of part one, subjects were asked to type their first name on their computer screen. This was then shown to the person they were paired with throughout part one. Subjects were then randomly assigned either the role of dictator or recipient. Dictators made their choices in the 16 games of part one. Recipients did not have any choices to make but were informed about the dictator's decision in each of the games, and had to confirm that they had seen each decision before a new game was presented to the dictator.¹³

¹³ We chose to decrease dictator-recipient anonymity and to have recipients interact with dictators throughout part one as a way of reducing social distance between dictators and recipients and thus increase the pull of other-regarding concerns and create a starker conflict with the self-interest motive. Decreasing anonymity has been shown

Once everyone had completed part one, instructions for part two appeared on subjects' computer screens. These instructions explained the Stroop task and illustrated how to submit answers on the computer. After performing the Stroop task, which lasted for five minutes, subjects completed a post-experimental questionnaire collecting standard socio-demographic measurements (such as gender, age, nationality).

After completing the questionnaire subjects were shown their payoffs from part one and two of the experiment. We used a random incentive lottery system to pay subjects. At the end of each session, one of the two parts of the experiment was selected at random and subjects were paid according to their earnings from the selected part. If part one was selected, one of the 16 games was chosen at random and subjects were paid according to their earnings in this game. Sessions lasted approximately 50 minutes and earnings ranged between £0 and £19, averaging £11.81 (s.d. 4.77).

3.3 Manipulation Check

The 16 games of part one were designed so that dictators in Conflict would face a conflict between self-interest and altruistic motives, whereas dictators in NoConflict would face no such conflict.¹⁴ We included two questions in the questionnaire to probe the extent to which subjects (self-reportedly) perceived such a motivational conflict. These questions asked dictators to rate the extent to which they found the dictator game choices “hard” and “uncomfortable”.¹⁵

Dictators in the Conflict treatment reported that choices in the dictator games were harder and more uncomfortable relative to dictators in NoConflict (hard: mean 3.07, s.d. 2.64 vs. mean 0.84, s.d. 1.75; discomfort: mean 4.12, s.d. 3.26 vs. mean 1.05, s.d. 1.90). In both cases, we find that responses are significantly different across the two treatments (Wilcoxon-Mann-Whitney tests; hard: $p < 0.001$; discomfort: $p < 0.001$).¹⁶ This suggests that dictators in the Conflict treatment may have perceived a stronger motivational conflict, relative to dictators in the

to significantly increase giving in previous dictator-game experiments (e.g., Bohnet and Frey, 1999; Charness and Gneezy, 2008). Dictator giving has also been shown to increase when recipients interact with dictators (e.g., Yamamori et al., 2007; and Yamamori et al., 2008).

¹⁴ Table C.1, Appendix C in the OSM, reports summary statistics of the choices made by dictators in these games.


¹⁵ These items were: “Overall, how hard was it to choose between option A and option B in the 16 situations of Part 1?”, and “Overall, to what extent did you experience discomfort in making your choices in the 16 situations of Part 1?”. Responses were collected on a scale from 0 (“not at all”) to 10 (“very much”).

¹⁶All tests reported in the paper use the individual as the independent unit of observation and therefore are based on 99 dictators per treatment, unless otherwise stated.

NoConflict treatment. We next check whether this (perceived) stronger conflict resulted in a more severe depletion of willpower by comparing dictators' Stroop task performance across the two treatments.

3.4 Main Results: Dictators' Performance in the Stroop Task


Figure 3 shows the average number of correct answers in the Stroop task by dictators in the Conflict and NoConflict treatments. Dictators in Conflict gave fewer correct answers compared to NoConflict (Conflict: mean 146.21, s.d. 17.67; NoConflict: mean 149.44, s.d. 15.82). The resulting effect size (measured as Cohen's d) is -0.193. Its associated 95% confidence interval is [-0.472, 0.086]. Note that it includes zero. The difference in performance is not statistically significant according to a Wilcoxon-Mann-Whitney test ($p = 0.157$). An OLS regression (reported in Appendix C in the OSM) confirms that the difference in performance remains insignificant after controlling for a number of observable characteristics of the subjects and of the decision setting. Overall, we do not find support for Hypothesis 1.



Rand et al. (2016) and our meta-analysis suggests that trade-offs between self-interest and altruism may have different effects on men and women. Although our design is robust to heterogeneity in the direction of the effect, it is possible that the true effect is only present among men, or only among women, while it is zero in the other subgroup. This heterogeneity would result in a weak aggregate effect and could explain why we fail to find support for Hypothesis 1

in the aggregate data. To examine whether this is the case, in Figure 4 we show the average number of correct answers in the Stroop task by treatment, separately for men and women.¹⁷

FIGURE 4
Performance in the Stroop task, by treatment and gender



Note: Bars indicate 95% confidence intervals.

For each gender the performance in the Stroop task is lower in the Conflict treatment than in the NoConflict treatment. The treatment difference in performance is larger for women (-4.51 correct answers) than for men (-1.23 correct answers). In either case, however, the difference is not significant at the 5% level according to Wilcoxon-Mann-Whitney tests (men: $p = 0.892$; women: $p = 0.093$). These results are confirmed by OLS regressions with controls (see Appendix C in the OSM).¹⁸

¹⁷ This part of the analysis is exploratory, since we did not set out to detect differences in the effect between men and women. The analysis and tests reported below are based on the subsamples of 78 men and 120 women who participated in the experiment. For men, we can detect effects of size 0.58 or larger with 80% power and 5% probability of a type-I error. For women, we can detect effects of size 0.47 or larger.

¹⁸ Another possible source of heterogeneity, that may weaken the aggregate effect, is the strength of the conflict between self-interest and altruistic concerns that subjects may have experienced during the dictator games. Subjects who have no concern for others and are exclusively motivated by self-interest (or vice versa subjects who have no concern for their self-interest and are exclusively motivated by altruism) may have experienced little or no motivational conflict and depletion of willpower during the dictator games. It is difficult, however, to identify this subgroup of dictators in our sample, since the choices in the dictator games are not indicative of whether or not the decision-maker has experienced a motivational conflict in making those choices. For example, observing an own payoff-maximizing choice in a game does not reveal whether this has been made by an entirely self-interested dictator, who experienced no conflict in making such choice, or by an other-regarding dictator who was conflicted in her decision but ultimately opted for the selfish option. In Appendix C we report an analysis based on the self-reported perception of conflict experienced in the dictator games, and find no evidence that this moderates the treatment effect in our experiment.

4. DISCUSSION AND CONCLUSIONS

In this paper, our aim was to probe the validity of a dual-system approach to altruistic behavior, by testing the hypothesis that trade-offs between self-interest and altruism trigger a conflict between our “fast”, intuitive System 1 and our “slower”, more deliberative System 2. Our empirical strategy relies on two elements. First, we performed a meta-analysis of the existing literature on the relation between altruism and intuition/deliberation. Second, we designed and conducted a new experimental study.

The meta-study shows that the manipulations used in previous experiments to promote the intuitive system at the expense of the deliberative system (using ego depletion tasks, cognitive load, time pressure or priming) have an overall effect on altruism that is very close to, and not significantly different from, zero.

Can this aggregate null result be due to the fact that the relation between altruism and our two systems is genuinely heterogeneous across different subgroups and decision situations, as some researchers have suggested? It is difficult to give a definite answer to this question. Our meta-study shows that gender may be a mediator of this relation, in line with the argument proposed by Rand et al. (2016). However, our findings are in contrast with the findings of the meta-study by Rand et al. (2016): we find a (marginally significant) negative relation between intuition and altruism for men and no relation for women, while Rand et al. (2016) find no relation for men and a positive and significant relation for women. Moreover, none of the other potential mediators that we consider (including mediators that had been previously suggested in the literature, like the stakes of the experiment, or the frame of the task) plays a significant role in explaining the heterogeneity in existing results.

The results from the meta-study suggest that we still do not have a clear understanding of what mediates the effect of intuition and deliberation on altruism. This, combined with a lack of theory to guide us in understanding the possible sources of heterogeneity, highlights the importance of designing new experiments that allow us to probe the validity of dual-system theories of altruism without being vulnerable to the presence of unobserved and/or unanticipated heterogeneity in the relation between altruism and the intuitive and deliberative systems.

In the second part of the paper, we propose a new experimental design that achieves this objective. We argue that, if trade-offs between altruism and self-interest trigger a conflict

between the intuitive and deliberative systems, then being exposed to these trade-offs should be willpower-depleting. We do not find evidence that this is the case. Directionally, the effect is in line with the hypothesis: subjects who were exposed to trade-offs between altruism and self-interest perform worse on a task that requires willpower than subjects who were not exposed to such trade-offs. However, we cannot reject the null that the two groups perform similarly.¹⁹ We also do not find evidence that the effect is present only among men, or only among women – reinforcing our impression that the relation between gender, intuition and altruism is more controversial than what has previously been suggested.

Overall, the combination of evidence from the meta-study and the new experiment suggest that choices that involve trade-offs between altruism and self-interest do not trigger any strong conflict between intuition and deliberation. This could be because, in the realm of altruistic behavior, the decision processes governed by the intuitive and deliberative systems may actually lead to the same outcome for any given individual. That is, our choices may be governed by either intuition or deliberation, depending on the circumstances, but the two systems are not actually in conflict when it comes to making these type of decisions, and hence their outcome is the same. Alternatively, the null results reported in the literature and in our new experiment could mean that the lens of dual-system models does not extend to altruistic behavior. In either case, our study offers little support for the notion that, in the domain of altruistic choices, the individual must spend cognitive resources to override the intuitive impulse and take the choice favored by the deliberative system.

¹⁹ Of course, it is possible that an effect is actually present in our data – but it is smaller than our study is powered to detect (we are powered to detect effects of size 0.4 or larger). The estimated size of the effect in our new experiment is -0.193. Detecting such an effect with 80% power would require re-running our experiment with a sample of about 450 dictators per treatment.

REFERENCES

- Achtziger, Anja, Carlos Alós-Ferrer, and Alexander K. Wagner. 2015. "Money, Depletion, and Prosociality in the Dictator Game." *Journal of Neuroscience, Psychology, and Economics* 8 (1): 1–14.
- Andersen, Steffen, Uri Gneezy, Agne Kajackaite, and Julie Marx. 2018. "Allowing for Reflection Time Does Not Change Behavior in Dictator and Cheating Games." *Journal of Economic Behavior & Organization* 145: 24–33.
- Balafoutas, Loukas, Rudolf Kerschbamer, and Regine Oexl. 2018. "Distributional Preferences and Ego Depletion." *Journal of Neuroscience, Psychology, and Economics* 11 (3): 147–65.
- Banker, Sachin, Sarah E. Ainsworth, Roy F. Baumeister, Dan Ariely, and Kathleen D. Vohs. 2017. "The Sticky Anchor Hypothesis: Ego Depletion Increases Susceptibility to Situational Cues." *Journal of Behavioral Decision Making* 30 (5): 1027–1040.
- Baumeister, Roy F., Kathleen D Vohs, and Dianne M Tice. 2007. "The Strength Model of Self-Control." *Current Directions in Psychological Science* 16 (6): 351–55.
- Benjamin, Daniel J., Sebastian A. Brown, and Jesse M. Shapiro. 2013. "Who Is 'Behavioral'? Cognitive Ability and Anomalous Preferences." *Journal of the European Economic Association* 11 (6): 1231–55.
- Bohnet, Iris, and Bruno S. Frey. 1999. "Social Distance and Other-Regarding Behavior in Dictator Games: Comment." *American Economic Review* 89 (1): 335–39.
- Burger, Nicholas, Gary Charness, and John Lynham. 2011. "Field and Online Experiments on Self-Control." *Journal of Economic Behavior & Organization* 77 (3): 393–404.
- Capraro, Valerio, Brice Corgnet, Antonio M. Espín, and Roberto Hernán-González. 2017. "Deliberation Favours Social Efficiency by Making People Disregard Their Relative Shares: Evidence from USA and India." *Royal Society Open Science* 4 (2).
- Carter, Evan C, Lilly M Kofler, Daniel E Forster, and Michael E McCullough. 2015. "A Series of Meta-Analytic Tests of the Depletion Effect: Self-Control Does Not Seem to Rely on a Limited Resource." *Journal of Experimental Psychology: General* 144 (4): 796–815.
- Charness, Gary, and Uri Gneezy. 2008. "What's in a Name? Anonymity and Social Distance in Dictator and Ultimatum Games." *Journal of Economic Behavior & Organization* 68 (1): 29–35.
- Chen, Fadong, and Ian Krajbich. 2018. "Biased Sequential Sampling Underlies the Effects of Time Pressure and Delay in Social Decision Making." *Nature Communications* 9 (3557).
- Cornelissen, Gert, Siegfried Dewitte, and Luk Warlop. 2011. "Are Social Value Orientations Expressed Automatically? Decision Making in the Dictator Game." *Personality and Social Psychology Bulletin* 37 (8): 1080–90.
- De Haan, Thomas, and Roel Van Veldhuizen. 2015. "Willpower Depletion and Framing Effects." *Journal of Economic Behavior & Organization* 117: 47–61.
- Deck, Cary, and Salar Jahedi. 2015. "The Effect of Cognitive Load on Economic Decision Making: A Survey and New Experiments." *European Economic Review* 78: 97–119.
- Dreber, Anna, Drew Fudenberg, David K Levine, and David G Rand. 2016. "Self-Control, Social Preferences and the Effect of Delayed Payments." *Working Paper, Washington University in St. Louis, St. Louis* Available at SSRN: <http://ssrn.com/abstract=2477454>.

- Fischbacher, Urs. 2007. "Z-Tree: Zurich Toolbox for Ready-Made Economic Experiments." *Experimental Economics* 10 (2): 171–78.
- Gärtner, Manja. 2018. "The Prosociality of Intuitive Decisions Depends on the Status Quo." *Journal of Behavioral and Experimental Economics* 74: 127–138.
- Greiner, Ben. 2015. "Subject Pool Recruitment Procedures: Organizing Experiments with ORSEE." *Journal of the Economic Science Association* 1 (1): 114–25.
- Grossman, Zachary, and Joël J. Van der Weele. 2017. "Dual-Process Reasoning in Charitable Giving: Learning from Non-Results." *Games* 8 (3): 36.
- Hagger, Martin S, Chantelle Wood, Chris Stiff, and Nikos LD Chatzisarantis. 2010. "Ego Depletion and the Strength Model of Self-Control: A Meta-Analysis." *Psychological Bulletin* 136 (4): 495.
- Halali, Eliran, Yoella Bereby-Meyer, and Axel Ockenfels. 2013. "Is It All about the Self? The Effect of Self-Control Depletion on Ultimatum Game Proposers." *Frontiers in Human Neuroscience* 7.
- Hauge, Karen Evelyn, Kjell Arne Brekke, Lars-Olof Johansson, Olof Johansson-Stenman, and Henrik Svedsäter. 2016. "Keeping Others in Our Mind or in Our Heart? Distribution Games under Cognitive Load." *Experimental Economics* 19: 562–76.
- Itzchakov, Guy, Liad Uziel, and Wendy Wood. 2018. "When Attitudes and Habits Don't Correspond: Self-Control Depletion Increases Persuasion but Not Behavior." *Journal of Experimental Social Psychology* 75: 1–10.
- Janssen, Loes, Bob M. Fennis, Ad Th H. Pruyn, and Kathleen D. Vohs. 2008. "The Path of Least Resistance: Regulatory Resource Depletion and the Effectiveness of Social Influence Techniques." *Journal of Business Research* 61 (10): 1041–1045.
- Jarke, Johannes, and Johannes Lohse. 2016. "I'm in a Hurry, I Don't Want to Know! The Effects of Time Pressure and Transparency on Self-Serving Behavior." *WiSo-HH Working Paper Series* No. 32.
- Kahneman, Daniel. 2002. "Maps of Bounded Rationality: A Perspective on Intuitive Judgment and Choice." *Nobel Prize Lecture* 8: 351–401.
- . 2011. *Thinking, Fast and Slow*. Farrar, Straus and Giroux.
- Kessler, Judd B., and Stephan Meier. 2014. "Learning from (Failed) Replications: Cognitive Load Manipulations and Charitable Giving." *Journal of Economic Behavior & Organization* 102: 10–13.
- Kinnunen, Suna Pirita, and Sabine Windmann. 2013. "Dual-Processing Altruism." *Frontiers in Psychology* 4: 193.
- Krawczyk, Micha, and Marta Sylwestrzak. 2018. "Exploring the Role of Deliberation Time in Non-Selfish Behavior: The Double Response Method." *Journal of Behavioral and Experimental Economics* 72: 121–134.
- Kuhn, Michael A, Peter Kuhn, and Marie Claire Villeval. 2014. "Self Control and Intertemporal Choice: Evidence from Glucose and Depletion Interventions." *CESifo Working Paper* No. 4609.
- Lipsey, Mark W., and David B. Wilson. 2001. *Practical Meta-Analysis*. Sage Publications, Inc.
- Loewenstein, George, and Ted O'Donoghue. 2007. "The Heat of the Moment: Modeling Interactions between Affect and Deliberation." *Carnegie Mellon University Department of Social and Decision Sciences Unpublished Working Paper* Available at <http://econweb.arts.cornell.edu/edo1/heat.pdf>.

- MacLeod, Colin M. 1991. "Half a Century of Research on the Stroop Effect: An Integrative Review." *Psychological Bulletin* 109 (2): 163–203.
- Merkel, Anna Louisa, and Johannes Lohse. 2018. "Is Fairness Intuitive? An Experiment Accounting for Subjective Utility Differences under Time Pressure." *Experimental Economics*, 1–27.
- Moore, Don A., and George Loewenstein. 2004. "Self-Interest, Automaticity, and the Psychology of Conflict of Interest." *Social Justice Research* 17 (2): 189–202.
- Mrkva, Kellen. 2017. "Giving, Fast and Slow: Reflection Increases Costly (but Not Uncostly) Charitable Giving." *Journal of Behavioral Decision Making* 30 (5): 1052–1065.
- Muraven, Mark, and Roy F. Baumeister. 2000. "Self-Regulation and Depletion of Limited Resources: Does Self-Control Resemble a Muscle?" *Psychological Bulletin* 126 (2): 247.
- Rand, David G. 2016. "Cooperation, Fast and Slow: Meta-Analytic Evidence for a Theory of Social Heuristics and Self-Interested Deliberation." *Psychological Science* 27 (9): 1192–1206.
- Rand, David G., Victoria L. Brescoll, Jim AC Everett, Valerio Capraro, and Hélène Barcelo. 2016. "Social Heuristics and Social Roles: Intuition Favors Altruism for Women but Not for Men." *Journal of Experimental Psychology: General* 145 (4): 389–96.
- Schulz, Jonathan F., Urs Fischbacher, Christian Thöni, and Verena Utikal. 2014. "Affect and Fairness: Dictator Games under Cognitive Load." *Journal of Economic Psychology* 41: 77–87.
- Strombach, Tina, Zsófia Margittai, Barbara Górczyca, and Tobias Kalenscher. 2016. "Gender-Specific Effects of Cognitive Load on Social Discounting." *PLoS One* 11 (10): e0165289.
- Stroop, Ridley J. 1935. "Studies of Interference in Serial Verbal Reactions." *Journal of Experimental Psychology* 18: 643–62.
- Tinghög, Gustav, David Andersson, Caroline Bonn, Magnus Johannesson, Michael Kirchler, Lina Koppel, and Daniel Västfjäll. 2016. "Intuition and Moral Decision-Making—the Effect of Time Pressure and Cognitive Load on Moral Judgment and Altruistic Behavior." *PloS One* 11 (10): e0164012.
- Verschuere, Bruno, Nils C. Köbis, Yoella Bereby-Meyer, David Rand, and Shaul Shalvi. 2018. "Taxing the Brain to Uncover Lying? Meta-Analyzing the Effect of Imposing Cognitive Load on the Reaction-Time Costs of Lying." *Journal of Applied Research in Memory and Cognition*.
- Xu, Hanyi, Laurent Bègue, and Brad J. Bushman. 2012. "Too Fatigued to Care: Ego Depletion, Guilt, and Prosocial Behavior." *Journal of Experimental Social Psychology* 48 (5): 1183–1186.
- Yamamori, Tetsuo, Kazuhiko Kato, Toshiji Kawagoe, and Akihiko Matsui. 2007. "Online Chat Leads to Fairness in Dictator Games." Available at <http://homepage3.nifty.com/tetsu-ym/ChatDG.pdf>.
- . 2008. "Voice Matters in a Dictator Game." *Experimental Economics* 11 (4): 336–43.
- Zaki, Jamil, and Jason P. Mitchell. 2013. "Intuitive Prosociality." *Current Directions in Psychological Science* 22 (6): 466–470.

ONLINE SUPPLEMENTARY MATERIALS

FOR THE PAPER

“ALTRUISM, FAST AND SLOW? EVIDENCE FROM A META-ANALYSIS AND A NEW EXPERIMENT”

BY HANNA FROMELL, DANIELE NOSENZO, AND TRUDY OWENS

- Appendix A contains details of the design and results of the meta-analysis discussed in Section 2 of the paper.
- Appendix B contains the instructions used in the new experiment described in Section 3.
- Appendix C contains regression analysis in support of the analysis of Section 3.

Appendix A: Meta-Study

In this appendix, we discuss additional design details and results of the meta-study. Table A.1 provides a list of studies included in the meta-study. For each study, the table reports the number of experiments contained in the study and the number of subjects who participated in those experiments, the method used to manipulate cognitive resources (cognitive load, priming, ego depletion or time pressure), and the type of game used to measure altruism (dictator game or donation game). Table A.2 contains descriptive statistics of the variables used in the mediator analysis.

A.1 Design

We searched the literature for studies to include in the meta-study, initially using Google Scholar and New Economic Papers (NEP) alerts, and then following up on relevant papers cited in the set of papers that we initially found in our literature search. We looked for experimental studies that manipulated cognitive resources of participants using one of the four standard types of interventions, i.e. cognitive load, priming, ego depletion, or time pressure. Since our focus is on altruistic behavior, we restricted the search to experiments in which subjects were asked to allocate money between themselves and another passive player (i.e., dictator games or donation games). We excluded all types of experimental games in which the recipient of the money allocated by the decision-maker was not completely passive (i.e., we excluded ultimatum games; ring measures of Social Value Orientation; trust games; etc.). We allow designs that use role uncertainty (both the decision-maker and the passive player make one decision, but only one decision, randomly selected at the end of the experiment, counts for payment).

We require that the allocation decisions have real monetary consequences for participants, i.e. we only include non-hypothetical studies. We also require that the studies follow the methodology of experimental economics (although we do not restrict the analysis to economics studies), in particular that they do not use deception in their designs. This excludes a handful of studies, including two studies previously incorporated in the meta-study by Rand et al. (2016).¹

¹ The excluded studies due to use of hypothetical stakes or deception are: Janssen et al. (2008); Cornelissen et al. (2011); Xu et al. (2012); Kinnunen and Windmann (2013); and one treatment of Banker et al. (2017).

TABLE A.1
List of studies included in the meta-study

Study	Acronym	N	Method	N. experiments contributed to study	Game
Benjamin et al. 2013 (*)	BBS	97	CL	2 (pilot / main study)	DG
Grossman and v.d. Weele 2017	GVdW	224	CL	2 (high / low affect)	Donation
Hauge et al. 2016	HBJJS	253	CL	5 (exp 1 / exp 2 give / exp 2 take / exp 3 give / exp 3 take)	DG / Donation
Kessler and Meier. 2014	KM	405	CL	5 (study 1 / replication 1-4)	Donation
Schulz et al. 2014 (*)	SFTU	135	CL	1	DG
Banker et al. 2017	BABAV	459	PR	2 (exp 2 / exp 4b)	DG
Rand et al. 2016 (*)	RBECB	2,788	PR	7 (D / E / F / G / H / V / study 2)	DG
Achtziger et al. 2015	ACW	128	DPL	1	DG
Banker et al. 2017	BABAV	795	DPL	3 (exp 1 / exp 3 low / high anchor)	DG
Halali et al. 2013	HBMO	55	DPL	1	DG
Itzchakov et al. 2018	IUW	104	DPL	2 (with / without persuasion message)	Donation
Tinghög et al. 2016 (*)	TABJKKW	309	DPL/CL	1	Donation
Andersen et al. 2018	AGKM	208	TP	2 (high / low stakes)	DG
Capraro et al. 2017 (*)	CCEHG	460	TP	2 (India / USA)	DG
Chen and Krajbich 2018 (*)	CK	204	TP	1	DG
Gärtner 2018	G	1349	TP	3 (no / selfish / prosocial status quo)	DG
Jarke and Lohse 2016	JL	58	TP	1	DG
Krawczyk and Sylwestrzak 2018	KS	136	TP	1	DG
Mrkva 2017	M	354	TP	3 (study 1 / study2 low stakes / study2 high stakes)	Donation
Merkel and Lohse 2018	ML	250	TP	2 (weak / strong time pressure)	DG
Rand et al 2016 (*)	RBECB	1,025	TP	5 (study 1 A / B / C / I / J)	DG
Tinghög et al. 2016 (*)	TABJKKW	1,102	TP	3 (Austria / Sweden / USA)	Donation

Note: In a few cases, the same study reports data from interventions using different manipulations of cognitive resources (e.g. time pressure and priming, Rand et al., 2016). In those cases, we report the same study more than once in the table. N = number of subjects involved in the experiment; CL = cognitive load; PR = priming; DPL = ego depletion; TP = time pressure. DG = dictator game, Donation = donation experiment. “Acronym” refers to the acronym used in Figure 1 in the main paper. (*) indicates that the raw data of the study are available.

TABLE A.2
Descriptive statistics of variables used in mediator analysis

Variable	Mean	N of experiments for which data is available	N of subjects
Stakes of the experiment (in 2017 USD)	5.12	48	9,745
Sample		55	10,898
<i>1 if students</i>	0.53		
<i>1 if AMT workers</i>	0.36		
Location of experiment		55	10,898
<i>1 if US</i>	0.56		
<i>1 if Europe</i>	0.36		
Type of game		55	10,898
<i>1 if donation game</i>	0.31		
<i>1 if dictator game</i>	0.69		
Frame of the game		55	10,898
<i>1 if give</i>	0.87		
<i>1 if take</i>	0.13		
Type of intervention		55	10,898
<i>1 if cognitive load</i>	0.29		
<i>1 if time pressure</i>	0.42		
<i>1 if ego depletion</i>	0.16		
<i>1 if priming</i>	0.16		
Gender of participants		26	6,378
<i>1 if female</i>	0.50		

Note: The means are computed at the experiment level, except for gender where we report the fraction of female participants in the 26 experiments for which we have data. Stakes are computed as the maximum nominal payoff available to the decision-maker multiplied by the probability a subject is actually paid, converted to 2017 USD using PPP. In a handful of cases (6 experiments) we could not compute this based on the information available in the paper.

We refer to the unit of observation in the meta-analysis as an “experiment”. Many studies contribute more than one experiment to our meta-study. This is because some of the included studies investigate the interaction between the cognitive resource manipulation and some other treatment variable (e.g., the frame of the game, Banker et al., 2017; the size of the stakes, Andersen et al., 2018; etc.), and so report results from more than one experiment. Other studies report several independent experiments conducted at different points in time, with different subject pools, or in different countries. Finally, in a few cases, the same study reports different experiments using different types of manipulations of cognitive resources. In a few cases, the

studies are based on multiple rounds or on different versions of an underlying game (e.g. games with different payoff configurations). In these cases we aggregate across the rounds of the experiment (e.g., Achtziger, et al., 2015) or across the different versions of the game (e.g., Schulz et al., 2014) and use only one observation per study in the analysis.

The meta-study focuses on the effect of manipulating cognitive resources on altruism. We measure altruism as the amount (or fraction of endowment) that the decision-maker gives to the passive player. In games where decision-makers face binary choices (give / not give) we measure altruism as the fraction of decision-makers sacrificing own payoff to increase the passive player's payoff. In order for a study to be included in the meta-analysis we must be able to compute this measure of altruism, either from the statistics, tables and graphs reported in the paper, or from its raw data. We exclude two studies because we could not retrieve this information.²

For each study, we derive two measurements of altruism, one for the treatment condition where cognitive resources were manipulated to promote intuitive responses (e.g., treatments where decisions had to be taken quickly in time pressure studies), and one for the condition that promoted deliberation. To quantify the effect of promoting intuition vs. deliberation on altruism, we calculate the standardized mean difference (Cohen's d) in altruism between these two conditions. To account for small sample bias, we apply Hedges' correction to the computed Cohen's d and associated standard error. In all cases, a positive value of Cohen's d indicates that promoting intuition triggered more altruistic behavior relative to the deliberative condition, while a negative value indicates the opposite.

In some of the analysis below, we will analyze heterogeneity in effect sizes between male and female participants. We do not have sufficient information to compute separate effect sizes for men and women for all 19 studies and 55 experiments included in the main analysis. We can compute this for 8 studies involving 26 experiments and 6,378 participants. This is about 45% of the sample in terms of number of studies/experiments and 58% of the sample in terms of number of participants. In a previous meta-analysis, Rand et al. (2016) also focused on gender heterogeneity in effect sizes. We obtained their data and so our meta-analysis subsumes theirs, except that we exclude two of the studies that they had included because they do not meet our

² The two studies are: Balafoutas et al. (2018) and Strombach et al. (2016). Both studies report estimates of parameters of utility functions, from which it is not possible to reconstruct choice behavior.

inclusion criteria (the studies are Cornelissen et al., 2011 and Kinnunen and Windmann, 2013 – both involve deception). In addition, we can compute separate effect sizes for men and women for another 5 studies (10 experiments, 2,425 participants), which were not included in the meta-study by Rand et al.

A.2 Mediator Analysis

In this section, we explore the role of a number of mediating factors in explaining potential heterogeneity in effect sizes across the experiments included in the meta-study. We focus on the mediators listed in Table A.2. For each variable, we conduct a random-effects meta-regression where the dependent variable is the effect size detected in the experiment, and where each experiment is weighted by the inverse of its variance so that more precise studies have more influence in the analysis. The regressions for these variables (with the exception of gender, discussed below) are reported in Table A.3. The last column of Table A.3 reports a specification where all mediators are simultaneously included in the regression. In all regressions, we include study dummies (not reported in the table) to account for study fixed effects.

Stakes: Some authors have argued that whether intuition favors altruism or self-interest depends on the amount of money at stake in the decision situation. For instance, Mrkva (2017) argues that when stakes are high, an individual's impulsive response is to reject a request for money, and that generosity may thus be fostered by deliberation, while the reverse may happen when stakes are low. We can compute stake levels for 48 of the 55 experiments included in the meta-study. The stake level is defined as the maximum payoff available to the decision-maker (e.g. in a dictator game, the endowment received by the dictator), converted in 2017 USD PPP and multiplied by the probability that this amount is actually paid to the subject (in several studies not all choices made by subjects are paid with certainty, either because the study uses role uncertainty or because the giving decision is part of a set of tasks and the random lottery incentive system is used).³ Column (1) of Table A.3 reports the regression results. We do not find any significant association between stake level and effect size ($p = 0.411$). We find the same

³ The studies for which we could not compute a stake level typically did not report sufficient information about the probability with which each choice made in the experiment was paid.

result in the regression of column (7), where we control for other mediators. We conclude that stakes are not a significant mediator of the effect of intuition on altruism.

TABLE A.3
Mediator analysis: random-effects meta-regressions

Variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Stakes	0.004 (0.005)						0.004 (0.006)
1 if students		0.118 (0.153)					-0.027 (0.283)
1 if AMT workers		-0.149 (0.126)					-0.045 (0.160)
1 if exp. run in USA			-0.142 (0.236)				-0.134 (0.235)
1 if exp. run in Europe			0.063 (0.276)				0.100 (0.375)
1 if donation game				0.050 (0.361)			0.153 (0.357)
1 if take frame					0.278* (0.143)		0.157 (0.171)
1 if cognitive load						-0.045 (0.325)	0.072 (0.468)
1 if ego depletion						0.013 (0.235)	-0.186 (0.453)
1 if priming						-0.050 (0.125)	-0.066 (0.120)
Constant	-0.030 (0.074)	-0.099 (0.137)	-0.017 (0.270)	-0.068 (0.372)	-0.021 (0.082)	-0.007 (0.111)	-0.157 (0.442)
Study FE?	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N. experiments	48	55	55	55	55	55	48
N. subjects	9,745	10,898	10,898	10,898	10,898	10,898	9,745

Note: Dependent variable is the effect size associated to an experiment. * = 10% significance level.

Sample: About half of the experiments were run with university students as participants, about a third with Amazon Mechanical Turk workers, and the rest with other specific non-student samples. In column (2) we investigate whether there is heterogeneity in the effect sizes associated with different types of samples. The base category in the regression are studies run with non-student samples that are not AMT workers. We do not detect any significant differences between the base category and AMT workers or students. We also do not find any significant difference between AMT workers and students (F-test: $p = 0.134$). This is also true for the regression in column (7).

Location of the experiment: In column (3) we explore whether the location where the experiment was run affects the estimate of the effect size. About half of the experiments were run in the US, about a third in Europe, while the rest in countries such as Chile, India or Israel (the base category in the regression). We do not detect any difference between experiments run in Chile, Israel or India and those run in the US or Europe, nor do we find a difference between Europe and the US (F-test: $p = 0.161$). The same holds for the regression in column (7).

Charity as passive player: About a third of experiments were based on donation games, where decision-makers decide how much money to donate to a charitable organization. We do not find that the effect sizes in these studies differ from those reported in the dictator game experiments, as shown in columns (4) and (7).

Frame of the game: Banker et al. (2017) argue that the framing of the decision situation plays an important role in determining whether altruism or self-interest are an intuitive response. This is because intuition may favor choices that are salient. When the choice involves taking money that has been initially allocated to the passive player, the salient cue is to leave the amount with the passive player, and therefore promoting intuition will lead to more altruism. The reverse may happen when the choice involves giving money that has been initially allocated to the decision-maker. Gärtner (2018) entertains a similar hypothesis. Hauge et al. (2016) also report games with either give or take frames. In column (5) we test whether experiments using a take frame are associated with larger effect sizes (indicating more altruism in the intuitive condition) than experiments using a give frame. The coefficient of the take frame dummy is positive and marginally significant at the 10% level ($p = 0.060$) in column (5). However, as shown in column (7), this result is not robust to the inclusion of controls ($p = 0.368$).

Type of intervention: In column (6) we test whether there are systematic differences between the effect sizes reported in studies that use different types of manipulation of cognitive resources. We do not find any difference between types of manipulation, in any of the possible bilateral comparisons (all $p > 0.695$). The same holds for the regression in column (7).

Gender: Rand et al. (2016) propose that gender is an important mediator of the role of intuition and deliberation on altruism. They argue that altruism may be an intuitive response for women more than for men. In their meta-study they indeed find a significant difference between effect

sizes for men and women. Moreover, in line with their hypothesis, they find that the estimated effect size is positive and significant for women, while negative and insignificant for men (they use 5% as the probability of type-I error in their study).

For a subset of the studies in our meta-analysis we can compute separate effect sizes for men and women and thus re-test the hypothesis put forward by Rand et al. (2016) with a larger sample than they had available in their meta-study (26 experiments involving 6,378 subjects compared to 22 experiments and 4,366 subjects in Rand et al.). Table A.4 reports random-effect meta-regressions based on the subsample of experiments for which we can compute gender-specific effect sizes.⁴

TABLE A.4
Mediator analysis: the role of gender

Variable	(1)	(2)
1 if female	0.125** (.051)	0.131** (.053)
Constant	-0.103* (0.060)	-0.149 (0.226)
Study FE?	Yes	Yes
Additional controls?	No	Yes
N. effect sizes	50	48
N. experiments	26	25
N. subjects	6,378	6,174

Note: Random-effects meta-regressions. Dependent variable is the effect size associated to an experiment. The controls are the other mediators listed in Table A.2. ** = 5% significance level.


In column (1) we only include study dummies to control for study fixed effects. We detect a difference between effect sizes for men and women that is significant at the 5% level ($p = 0.020$). The regression estimate indicates that the effect size is on average larger for women than for men, in line with Rand et al.'s hypothesis. In column (2) we add the mediators listed in Table A.2 as additional controls (with the exception of the take frame dummy, since all experiments for

⁴ For one study (Grossman and Van der Wee 2017) we can only compute effect size for women, but not for men. Thus, while for all other experiments we include in the regression two effect sizes (one for men and one for women), for Grossman and van der Wee (2017) we only include the effect size for women. This is why in column (1) of Table A.4 we have 50 effect sizes for 26 experiments.

which we can compute gender-specific effect sizes involve a give frame).⁵ The difference in effect sizes between men and women remains significant at the 5% level ($p = 0.018$).

Figure A.1 further illustrates this result. It contains forest plots showing the effect sizes and 95% confidence intervals of the 26 experiments used in the gender analysis. The left panel of the figure shows the forest plot for men, while the right panel shows the plot for women. The last row of each panel reports the overall effect size estimated across all experiments.

FIGURE A.1
Effect sizes for men (left panel) and women (right panel)



For men, the overall effect size is negative (-0.076). The 95% confidence interval, however, includes zero, and we can only reject the null that the effect size is not zero at the 10% level ($p = 0.068$). For women, the overall effect size is positive (0.033). However, we cannot reject that this is not different from zero at any conventional level of significance ($p = 0.461$).

Thus, although the meta-regressions show that the effect sizes are significantly different between men and women, in either case the gender-specific effects sizes are quite small. Moreover, our results stand in contrast with what Rand et al. (2016) found in their meta-study. They found a positive and significant aggregate effect size for women, and a negative effect size, which was not significantly different from zero, for men. They concluded that promoting

⁵ Note that we lose one study between columns (1) and (2) due to the fact that we do not have information on stakes size for one experiment.

intuition increases altruistic behavior among women, but has no effect among men. If anything, our study suggests the opposite: promoting intuition leads men to become somewhat less altruistic, but has no effect on women's behavior.

Note that the sample we used to conduct the gender analysis differs in two ways from Rand et al.'s sample: we have added 10 experiments from 5 studies that were not available to Rand et al. at the time of their meta-analysis, and we have excluded 4 experiments from 2 studies that Rand et al. had included, because they involved deception. Is the difference in results between our meta-analysis and Rand et al.'s due to the addition of new studies, or to the exclusion of the studies that involved deception? We repeated the analysis above using only the studies included in Rand et al. except those that involved deception. We can replicate Rand et al.'s results in this subsample: for men the overall effect size is negative (-0.078), but only significant at the 10% level ($p = 0.087$), while for women the overall effect size is positive (0.135) and significant ($p = 0.003$). Thus, the difference in results is due to the addition of the 5 studies that were not included in Rand et al.⁶

⁶ Two of the new studies are based on donation games between a decision-maker and a charity. We checked whether the difference in results is driven by the inclusion of these two studies, but found no evidence that this is the case.

Appendix B: Experimental Instructions and Informed Consent

PRELIMINARY INSTRUCTIONS

Welcome! You are about to take part in a decision-making experiment. This experiment is run by the “Centre for Decision Research and Experimental Economics” and has been financed by various research foundations.

There are other people in this room, who are also participating in the experiment. Everyone is participating for the first time, and all participants are reading the same instructions. It is important that you **do not communicate with any of the other participants** during the experiment. If you have a question at any time, raise your hand and an experimenter will come to your desk to answer it.

The experiment consists of two parts: **Part 1 and Part 2.**

In each part of the experiment you will be asked to make decisions, and will have a chance to earn money. Decisions that will be made in one part of the experiment will not affect decisions or earnings in the other part of the experiment.

You will be informed of your earnings from Part 1 and Part 2 of the experiment once everyone in the room has completed Part 2. Therefore everyone will make their decisions in Part 2 without knowing their earnings from Part 1.

Only **one part of the experiment will be taken into account in determining your final earnings** from today’s experiment. At the end, we will toss a fair coin. If the coin lands heads, all participants in today’s experiment will be paid according to their earnings from Part 1. If the coin lands tails, all participants will be paid according to their earnings from Part 2. Your earnings will be paid out to you in private and in cash.

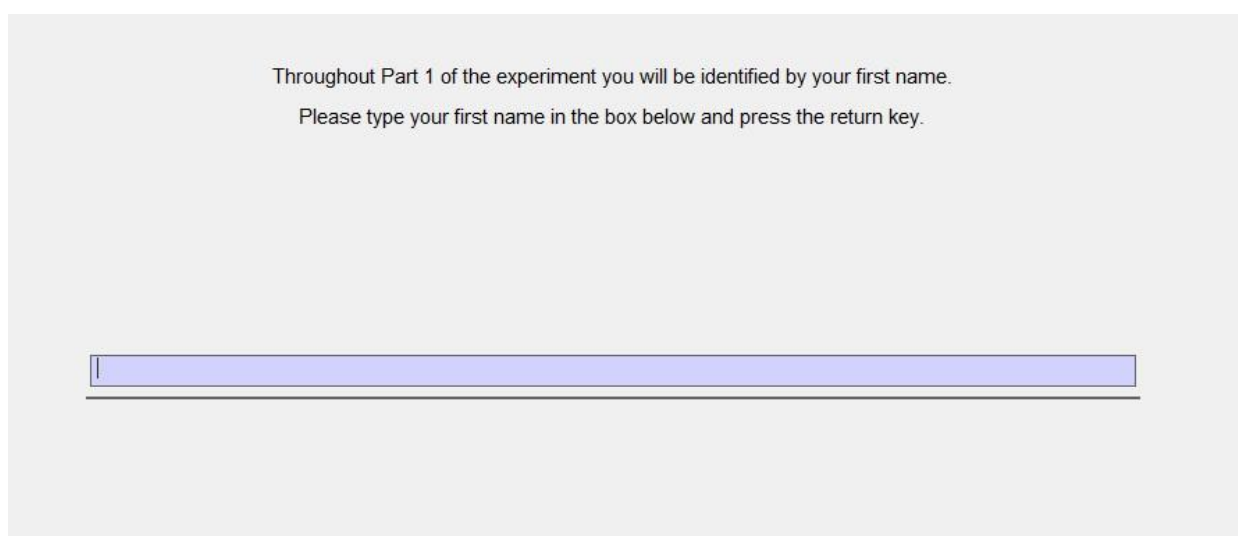
Shortly, you will receive detailed instructions about Part 1 of the experiment. You will receive detailed instructions about Part 2 directly on your computer screen once you have completed Part 1.

If you have a question now, please raise your hand and an experimenter will come to your desk to answer it.

PART 1 INSTRUCTIONS

At the beginning of PART 1 you will be **randomly paired with another person** in this room. You will remain paired with this person for the whole duration of PART 1. At the end of PART 1 the pair will be dissolved, and you will not be matched with this person again during today's experiment.

Throughout PART 1 you and the person you are paired with will be **identified by your first names**. At the beginning of PART 1 you will be asked to type your first name in a screen like the one below. Please enter your first name exactly as shown on your ID card and then press the return/enter key on your keyboard. On the following screen you will be shown the name of the person you are paired with in PART 1 of the experiment.



Throughout Part 1 of the experiment you will be identified by your first name.
Please type your first name in the box below and press the return key.

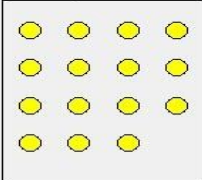
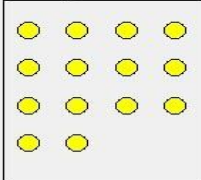

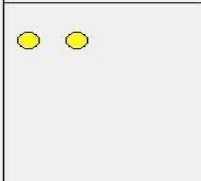
Each person in the pair will then be **randomly assigned a role**, either 'Person 1' or 'Person 2'. The computer will inform you of your role, which will stay the same throughout PART 1 of the experiment.

The participant in the role of **Person 1** will then be asked to make choices in a series of **16 situations**. For each situation he/she will have to choose between two options: Option A or Option B. Each option specifies an amount of money that Person 1 will receive and an amount of money that Person 2 will receive (all amounts are in British pounds).

For example, a possible situation may look as follows:

Situation XXX

You have to choose between Option A and Option B.
Once you have made your decision, click the SUBMIT button.

Option A		Option B	
xxx (Person 1)			xxx (Person 1)
xxx (Person 2)			xxx (Person 2)

Which do you choose?

In this example situation, if Person 1 chooses Option A, Person 1 receives £15 and Person 2 receives £1. If Person 1 chooses Option B, Person 1 receives £14 and Person 2 receives £2.

The participant in the role of **Person 2** will have no choices to make in PART 1 of the experiment. However, Person 2 will be informed in real time of the choices made by Person 1 in each of the 16 situations, and will have to confirm that he/she has seen each choice before a new situation will be presented to Person 1.

At the end of the experiment **one of the 16 situations will be selected at random** by the computer. Your final earnings in PART 1 of the experiment will be based on this randomly selected situation. Each situation has an equal chance of being selected, so please consider each situation carefully. If PART 1 is selected for payment you will be paid this amount in private and in cash.

Please raise your hand if you have any questions.

To make sure that everyone understands the instructions, please complete the questions below. In a couple of minutes an experimenter will come to your desk to check the answers.

Questions about PART 1:

1. How many choices will Person 1 make in total in PART 1? _____
2. Is the following statement true or false: you will be paired with the same person throughout PART 1 of the experiment _____
3. Suppose you are randomly assigned the role of Person 1. What will be the role of the other person in your pair? _____
4. Is the following statement true or false: if you are Person 2, you will not learn the choices made by Person 1 during PART 1 _____
5. Is the following statement true or false: your final earnings in PART 1 will be based on one of the 16 situations, which will be randomly selected at the end of the experiment _____

PART 2 INSTRUCTIONS

Part 2 of the experiment

In Part 2 of the experiment you will be shown a sequence of words on your screen. These words will be printed in black, blue, green, red or yellow.

Your task is to indicate which colour the word is printed in.

For each correct answer you will earn **10 pence**. Thus, the more colours you name correctly, the more points you earn.

This task will last for **5 minutes**. The time you have left will be displayed (in seconds) in the upper right corner of the screen.

The screen will also show you the number of correct answers so far.

The next screen will illustrate how to submit your answers in the task. Please click the CONTINUE button.

CONTINUE

Appendix C: Additional Analyses

In this Appendix we report additional analysis of the data from the new experiment. We start by reporting, in Table C.1, the choices of dictators in the 16 dictator games of the experiment. The table reports the fraction of dictators who chose the own payoff maximizing option in the Conflict and NoConflict treatments. In the NoConflict treatment nearly all dictators chose the option that maximized the own (and the recipient's) payoff in all of the games. In the Conflict treatment the proportion of own payoff maximizing choices is lower and varies between 18.2 and 67.9 percent across games, showing that a sizeable fraction of dictators took both the recipient's and their own interests into consideration when making their choices.

TABLE C.1
Percentage of subjects choosing the own payoff maximizing option

Game	Conflict		NoConflict	
	Payoffs from A vs. B	% own payoff maximizing choices	Payoffs from A vs. B	% own payoff maximizing choices
1	(8; 8) vs. (12; 0)	60.6	(8; 0) vs. (12; 8)	99.0
2	(11; 1) vs. (8; 8)	55.6	(11; 8) vs. (8; 1)	99.0
3	(9; 1) vs. (7; 5)	46.5	(9; 5) vs. (7; 1)	98.0
4	(7; 5) vs. (8; 2)	44.4	(7; 2) vs. (8; 5)	97.0
5	(10; 6) vs. (12; 0)	60.6	(10; 0) vs. (12; 6)	100
6	(8; 8) vs. (16; 0)	37.4	(8; 0) vs. (16; 8)	97.0
7	(6; 6) vs. (8; 4)	41.4	(6; 4) vs. (8; 6)	98.0
8	(10; 2) vs. (9; 7)	27.3	(10; 7) vs. (9; 2)	96.0
9	(12; 0) vs. (6; 6)	61.6	(12; 6) vs. (6; 0)	99.0
10	(11; 5) vs. (15; 1)	67.7	(11; 1) vs. (15; 5)	98.0
11	(8; 4) vs. (10; 2)	53.5	(8; 2) vs. (10; 4)	99.0
12	(9; 3) vs. (8; 4)	52.5	(9; 4) vs. (8; 3)	100
13	(10; 6) vs. (8; 8)	67.9	(10; 8) vs. (8; 6)	97.0
14	(13; 3) vs. (10; 6)	40.4	(13; 6) vs. (10; 3)	99.0
15	(6; 6) vs. (8; 2)	29.3	(6; 2) vs. (8; 6)	99.0
16	(10; 0) vs. (6; 4)	18.2	(10; 4) vs. (6; 0)	99.0

Note: in each cell of the column "Payoffs from A vs. B", the first number indicates the dictator's payoff and the second number the recipient's payoff (both displayed in GBP).

We next report regression analysis to support the results presented in Section 3.4 of the main paper. Table C.2 reports OLS regressions of dictators' performance in the Stroop task.

TABLE C.2
Performance in the Stroop task

Sample Model	All (1)	Men (2)	Women (3)
Conflict	-3.679 (2.339)	-3.354 (4.364)	-4.074 (2.981)
Male	0.407 (2.535)		
Age	-1.290*** (0.427)	-1.630*** (0.564)	-1.175* (0.564)
Economics	1.533 (3.528)	-1.607 (5.056)	6.532 (5.362)
Native	0.689 (3.350)	7.247 (5.503)	-3.388 (4.123)
Time elapsed (in seconds)	-0.093 (0.096)	-0.003 (0.148)	-0.177 (0.096)
Constant	249.161*** (72.362)	181.594 (110.926)	313.413*** (72.171)
Session dummies	Yes	Yes	Yes
R ²	0.14	0.26	0.19
N	198	78	120

Note: OLS regressions, robust standard errors in parentheses. In all models the dependent variable is the number of correct answers in the Stroop task. All models include 13 session dummies, not reported in the table. *** 1% significance level; * 10% significance level.

In Model (1) we use the whole sample and regress dictators' performance on a treatment dummy ("Conflict") taking value 1 for observations from the Conflict treatment and 0 for observations from the NoConflict treatment. We also control for observable characteristics of the subjects by including the following variables: *Male* is a dummy variable taking value 1 for male subjects and 0 for female subjects; *Age* is the subjects' age; *Economics* is a dummy assuming value 1 if the subject is an economics student and 0 otherwise (the experiment was run in the School of Economics); *Native* is a dummy variable taking value 1 if a subject is a native English speaker and 0 otherwise⁷; *Time elapsed* measures the amount of time (in seconds) subjects spent

⁷ We identify this by checking subjects' nationality and considering as "native" only those who are citizens of a country with English as its official language. We include this control to assess whether language proficiency may

in the lab between the beginning of part one and the beginning of the Stroop task to control for the impact on performance of any time-related factors such as boredom or opportunity to rest; The model also includes session dummies to control for session-specific effects, although these are not reported in Table C.2. Models (2) and (3) replicate the regression of Model (1), but for the subsample of male and female dictators, respectively.

In all models, the coefficient of the Conflict dummy is not significantly different from zero at any conventional level, indicating that performance in the Conflict treatment is indistinguishable from performance in the NoConflict treatment. This is also true when we focus on the subsample of male participants or the subsample of only female participants. Among the control variables, age has a negative and significant effect on performance, with older participants performing worse than younger ones.⁸ None of the other variables significantly affect performance.

Finally, we report an analysis of the effect of treatment on Stroop task performance distinguishing between subjects who reported to have experienced a more or less strong motivational conflict in the dictator games of part one. To distinguish between these two subgroups, we use subjects' responses in the questionnaire to the two questions about the extent to which they found the dictator game choices "hard" and "uncomfortable".⁹ We take the average of these two questions as our measure of the strength of the motivational conflict experienced by the subject, and perform a median split of the sample in the Conflict treatment. Subjects with an average response below the median are classified as "low conflict" (49 subjects), while those with a response above the median are classified as "high conflict" (50 subjects). Figure C.1 shows the Stroop task performance of subjects in each subgroup, as well as those in NoConflict.


High-conflict subjects gave on average 147.78 correct answers in the Stroop task (s.d. 16.69), while low-conflict subjects gave 144.61 correct answers (s.d. 18.65). Using Wilcoxon-Mann-Whitney tests, we do not find any significant difference in Stroop task performance between these two subgroups ($p = 0.763$), or between each subgroup and the subjects in the NoConflict

have affected performance in the task, which contains an element of language specificity since it uses incongruence between word meaning and font color as its manipulation.

⁸ The mean and median age in our sample are 21 and 20, respectively. The variable, which is self-reported, ranges from 18 to 42 years.

⁹ Recall that these items were: "Overall, how hard was it to choose between option A and option B in the 16 situations of Part 1?", and "Overall, to what extent did you experience discomfort in making your choices in the 16 situations of Part 1?". In both cases, responses were collected on a scale from 0 ("not at all") to 10 ("very much").

treatment ($p > 0.172$). **FIGURE C.1**
Performance in the Stroop task, by treatment and strength of conflict



Note: Bars indicate 95% confidence intervals.