cesifo WORKING PAPERS

9223 2021

July 2021

Understanding the Link between Intelligence and Lying

Michalis Drouvelis, Graeme Pearce

Impressum:

CESifo Working Papers
ISSN 2364-1428 (electronic version)
Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH
The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute
Poschingerstr. 5, 81679 Munich, Germany
Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de
Editor: Clemens Fuest
https://www.cesifo.org/en/wp
An electronic version of the paper may be downloaded
· from the SSRN website: www.SSRN.com
· from the RePEc website: www.RePEc.org
· from the CESifo website: https://www.cesifo.org/en/wp

CESifo *Working Paper No. 9223*

Understanding the Link between Intelligence and Lying

Abstract

Standard economic theory suggests that the decision to lie requires careful weighting of the associated economic costs and benefits, raising the question of whether intelligence matters for misbehaviour. Using the die roll paradigm, we compare behaviour between individuals who score either low or high on a Raven test when lying only benefits the subject who lies (Selfish treatment) or a charitable cause (Charity treatment). We find that high Raven individuals are honest in the Selfish treatment; however, their aversion to lying vanishes in the Charity treatment. Our results have important implications for the rapidly growing lying literature, indicating that intelligence is a key characteristic of misbehaviour.

JEL-Codes: C900, Z130.

Keywords: intelligence levels, die roll paradigm, honesty.

Michalis Drouvelis
University of Birmingham / United Kingdom
M.Drouvelis@bham.ac.uk

Graeme Pearce
University of Exeter / United Kingdom
g.pearce@exeter.ac.uk

We thank the University of Birmingham Business School, and the University of Exeter Business School for funding this research. We thank David Gill, Agne Kajackaite, John List and Shaul Shalvi for helpful comments and suggestions.

1 Introduction

In many social and economic aspects of their daily interactions, individuals are faced with decision-making situations involving misbehaviour such as not telling the truth. For example, people sometimes evade their taxes, claim welfare benefits they are not entitled to, and engage in financial frauds. Despite being considered one of the most common acts of wrong doing, lying is prevalent and people continue to mislead and behave dishonestly. Understanding the forces that determine the content of lying is therefore an important area of research in the behavioural and social sciences that has subsequently generated a rapidly growing interdisciplinary literature (see Gneezy et al. (2018); Abeler et al. (2019); Gerlach et al. (2019) for overviews). In standard economics the assumption is that, if lying is beneficial to individuals, they will refrain from misbehaving when the economic costs from doing so are high enough (Becker, 1968). An expansive literature in economics and psychology shows that the psychological disutility from lying makes some individuals tell the truth, but may cause others to lie even if they may not exploit the full benefits of it (e.g., Mazar et al. (2008); Gino et al. (2009)). It is therefore crucial to understand not only which groups of individuals are the most prolific liars, but also to examine those who understand the costs and benefits associated with lying. In this paper, we answer the following question: is lying a matter of intelligence?

A growing literature in behavioural economics has recognised the role of intelligence in various aspects of strategic decision making. For example, Gill & Prowse (2016) show that cognitive skills determine the evolution of Nash equilibrium play in repeated games. In a recent study, Fe et al. (2020) examine how childhood cognitive ability affect strategic sophistication. Others have examined the extent to which deception is cognitively demanding by considering the impact of cognitive load on dishonest behavior. For example, Van't Veer et al. (2014) show that subjects are more honest in situations of high cognitive load in comparison to low cognitive load, suggesting that 'cognitive capacity' is an important determinant of behaviour. Speer et al. (2020) explore the neuroeconomic foundations of dishonesty using MRI scanners, and find that areas of the brain associated with cognitive control help dishonest participants to be honest, whereas it enabled cheating for honest participants. Proto et al. (2019) provide evidence that differences in pro-social behaviour are mediated by differences in individuals' intelligence levels. In particular, the authors use a repeated game of cooperation and compare rates of cooperation in groups of subjects based on their IQ level. Their findings indicate that only the higher IQ subjects converge to full cooperation, understanding the benefits of long-run cooperation, as opposed to individuals with low IQ levels. In a separate series of experiments, Proto et al. (2020) show that subjects with higher IQ exhibit a lower frequency of errors in strategy implementation. Although the intelligence literature primarily focuses on strategic interactions, the issue of how intelligence impacts on individual–decision making has received limited attention.

A more specific motivation for our study stems from recent evidence showing that non-cognitive (personality) traits have been shown to affect lying behaviour. For example, Gino & Ariely (2012) show that participants with creative personalities tend to be more dishonest compared with individuals with less creative personalities. Using the HEXACO model, Zettler et al. (2015) offer evidence that the Honesty-Humility factor is the basic trait explaining individual differences in lying behaviour. In a more recent study, Heck et al. (2018) also find personality traits to be linked to unethical decision making. Paulhus & Dubois (2015) report in a meta analysis of 22 studies that higher ability students cheat less in scholastic settings. Rindermann et al. (2018) use the data collected by Gächter & Schulz (2016) which they combine with country level estimates of cognitive ability obtained from student assessments. They report that cognition correlates with misbehaviour in the die–rolling task. Our paper contributes to this literature by examining how individual–level cognitive ability determines individual decisions to misbehave or not.

For our purposes, we employ the standard die roll paradigm as introduced by Fischbacher & Föllmi-Heusi (2013): subjects roll a six sided die in private, and the number they report to have rolled determines how much they earn. Our focus is on cheating games that do not involve strategic uncertainty (unlike deception games, for example Gneezy (2005)).[1] In addition, we conduct an online experiment, making it impossible for the experimenter to know the true outcome of the die roll. This further generates an even more transparent environment where the probability of getting caught is virtually zero. We then observe behaviour in two treatments: one in which lying only benefits the decision-maker, the *Selfish* treatment, and one in which lying only benefits a charitable organisation, the *Charity* treatment. Following standard techniques in the literature, we measure fluid intelligence levels by having individuals perform a Raven test (Raven & Raven, 2003). The Raven Test has been used widely in research across the social science, as a tool in hiring, in military (Burke, 1958; Sundet et al., 2004) and educational settings in order to measure an individual's problem solving ability, or what educational psychologists often call fluid intelligence (Cattell, 1963).[2,3] We then compare the behaviour of those individuals with 'low' cognitive ability, those with scores below or equal the study population median (*Low* Raven subjects) to the behaviour of individuals with high cognitive ability, those with scores above the median (*High* Raven subjects).

We are concerned with the pure effects of intelligence on lying. We thus study behaviour in a one-shot game providing us with cleaner tests for the effects of intelligence on lying than repeated games. This allows us to rule out behaviour being confounded with strategic considerations arising from repetitions of the game. Our hypothesis is that when lying benefits only the person who lies (as in our *Selfish* treatment), individuals with high intelligence levels would refrain from lying as opposed to individuals with low intelligence levels who would actually lie. Our next hypothesis is that, in an environment where lying creates a positive externality (as in our *Charity* treatment) intelligence levels will not play a role. We predict that individuals with both a *Low* and *High* Raven score would lie and by doing so a third party benefits (without monetarily harming the liar). One explanation for this could be that there is a moral justification for lying when there is an externality, which intelligent individuals may care about more than others, as evidence by Guo et al. (2019).

[1] In order to study the origins of honesty, Houser et al. (2015) examine honest behaviour in children using the die rolling paradigm.

[2] Throughout the paper, we use the terms intelligence, fluid intelligence and cognitive ability interchangeably.

[3] Cattell (1963) provides a discussion of the definitions and differences between what are widely regarded as the two different types of intelligence, fluid and crystalised intelligence.

Another might be that the more intelligent hold a self–image about themselves as being honest, but also that they are charitable. Previous work has shown self–image concerns to be important for giving behaviours (Tonin & Vlassopoulos, 2013; Van't Veer et al., 2014).

Our hypotheses do not contradict findings that suggest lying decreases when lying for someone else. Those studies that have examined this link consider a setting where lying impacts an *individual*, rather than a charity. Motivations for giving to each of these is likely to be different, and therefore linked to different characteristics. Second, the link we draw between cognition and lying builds on the literature that links cooperation and intelligence, where social dilemmas are regarded as complex situation that require intelligence to solve. In our study, understanding how lying will impact others through a charity in situations very different from one's own is complex. Charitable donations are likely to impact others in the future, rather than instantly, something that Proto et al. (2019) find is linked to intelligence.

Our findings offer evidence in support of both our hypotheses and highlight the role of intelligence when it comes to lying behaviour. Replicating existing literature on lying, we observe that people lie across treatments, although not to the maximum extent. Novel to existing work is our finding that *High* Raven subjects do not lie to the same ectent as *Low* Raven subjects in the standard die roll paradigm (our *Selfish* treatment), as their reports do not differ from the distribution of an honest die to the same extent. *Low* Raven subjects, are found to misreport in line with previous studies; *High* Raven subjects are 9% less likely to report the profit maximising die roll of 5 than a *Low* Raven subject. When positive externalities are introduced, whereby misbehaviour benefits a charitable organisation, we observe that both *Low* and *High* Raven subjects lie: no significant differences are reported between the two groups in the *Charity* treatment. In sum, our findings show that individuals with high intelligence levels lie more only when positive externalities from lying benefit a third party, but not themselves. This effect is robust even when controlling for differences in subjects social preferences and a range of other individual characteristics, including income and gender, ruling out the possibility that one of these factors drive our results. Our findings suggest that differences in lying are due to subjects' cognitive abilities, thus indicating that honesty and intelligence are linked.

Our experiment contributes to existing literature in at least three respects. First, from a theoretical perspective, our results indicate that *High* Raven individuals deviate from what standard economic theory would predict in the *Selfish* treatment. Our findings also imply that lying aversion can be explained by models based on monetary outcomes when cognitive abilities are taken into account. In the *Selfish* treatment, this is the case for *Low* Raven individuals who decide to lie more than the *High* Raven individuals. In the *Charity* treatment, when social preferences are relevant for decision-making, *High* and *Low* Raven individuals behave similarly and choose actions that increase a third party's payoff, an outcome which standard economic theory would predict.

Second, recent evidence demonstrates that intelligence is a crucial determinant of individual behaviour, adding to the existing literature that mostly focuses on strategic behaviour. An exception is Burks et al. (2009) who show that cognitive skills are related to individuals' patience and attitudes towards risk. Together with our findings, it is warranted that it is important to generate more systematic evidence on the interplay between intelligence and individual decision making.

We expand the economics literature by looking at how "white lies" are perceived by individuals with different intelligence levels. Following the taxonomy of lies developed by Erat & Gneezy (2012), we examine a case of "white lie" where lying helps others without harming the liar. While the main purpose of this treatment was to investigate the extent to which individuals with high cognitive abilities are averse to lying, creating an environment where lying benefits a third party without monetarily harming the liar and it is morally justifiable to do so, we also further expand the discussion in the economics literature to "white lies" (i.e. lies that help others). In such situations, we observe that lying occurs significantly (benefiting the charity although it is far from being exploited to the full extent) but is not a matter of intelligence, as expected.

Our paper is organised as follows. Section 2 outlines the experimental design and procedures. Section 3 presents the experimental results. Section 4 presents some robustness checks, and Section 5 concludes.

2 Experimental design and procedure

Our experiment consists of 2x2 between-subjects design. In one dimension, we vary who benefits from subjects' misbehaviour. We refer to these treatments as the *Selfish* and the *Charity* treatments. In the *Selfish* treatment, we adopted the die roll paradigm as introduced by Fischbacher & Föllmi-Heusi (2013). Under this paradigm, subjects are asked to roll a six–sided die and the number appearing on the die determines their payoff. If the die number that came up was equal to 1, 2, 3, 4, or 5, then subjects received the corresponding payoff amount (in US dollars). If the die number that came up would equal 6, the corresponding payoff was equal to $0; subjects can potentially receive a higher payoff from lying. Choices in this game has no impact on other participants' payoffs and are not affected by beliefs about others' choices and payoff considerations.[4]

In contrast, in our *Charity* treatment, lying creates a positive externality for a third-party, a well-known charitable organisation (Wildlife Conservation Society). This particular charity was chosen as it is not associated with any particular political or religious views that may otherwise influence participants, whilst also being well known. The payoff structure in this treatment is identical to the one described for the *Selfish* treatment, with the only difference being who receives the payoff. This means that subjects' decisions only benefit the charity without monetarily harming the decision maker.

Subjects participated in a one-shot individual decision-making situation. In both treatments, it is impossible to detect lying at the individual level as no one could make a note of the actual number in the die roll. To further minimise the presence of potential experimenter demand effects that may discourage people from lying, our experiment was conducted online, using Amazon Mechanical Turk, whereby the probability of getting caught is essentially zero. All decisions were completely anonymous and this was explicitly mentioned to subjects at the beginning of the experiment. Collectively, these conditions

[4] In experiments conducted in the lab, the experimenter cannot verify if the subject has actually rolled the die or not. The same is true in our experiment.

created an environment in which it was easy for subjects to not tell the truth.

In the second dimension we vary whether lying behaviour depends on subjects' cognitive skills. To measure intelligence, subjects completed a ten item Raven test in order to measure fluid intelligence, as is standard in the literature (Raven & Raven, 2003; Proto et al., 2019, 2020). For each item, subjects were given 30 seconds, and correct answers were incentivised.[5] Depending on performance in the Raven Test, we split subjects into those whose score is above the median (*High* Raven) and those who are below or equal to the median Raven score (*Low* Raven). This differs to the methodology of Proto et al. (2019), who used session level Raven scores in order to determine if a subject is *Low* or *High*. The advantage of our approach is that a subject in the top 50% of all subjects will always be classified as *High* Raven.

As social preferences are likely to play an important role in the setting we study, we employed the Equality Equivalence Test (EET) (Kerschbamer, 2015) in order to measure them. The EET provides a score that categorises the individual into a social preference type such as inequality averse, selfish, or altruistic. The test provides subjects with a list of ten binary decisions over two allocations between themselves and a charity, and in each decision they have to select which one to implement: either LEFT or RIGHT. We employed a ten item test, with each decision presented in Table A3. These preference measures allow us to examine how social preference characteristics that may be relevant to lying vary with intelligence, and also provide important controls for our parametric analysis.

Subjects also completed a comprehensive questionnaire that elicited a range of personal characteristics, such as their income, race, gender, political party affiliation, education level and also how much they donated to charity in the previous year. We elicit these characteristics in order to act as controls in our parametric analysis, and because they may otherwise vary with our variable of interest. For example, previous work has shown that pro social behaviour, preferences and political party affiliation are closely related (Kerschbamer & Müller, 2020). There is also now a rich literature that shows that gender correlates with various economic behaviours (Croson & Gneezy, 2009). Income is an important variable in the vast majority of economic models, and previous work has shown that income and scores on fluid intelligence tests are correlated (Irwing & Lynn, 2006).

It is important to note that, in order to control for sequence effects in the procedure, the order in which subjects completed the die rolling task, the preference elicitation questions and questionnaire, was randomised. All experimental materials, showing how the behavioural and individual measures were elicited, are given in the Appendix.[6] Table 1 provides an overview of our experimental design and displays the number of observations we collected for each treatment.

The experiment was conducted using participants located in the USA and completed using Qualtrics. We ran two online sessions: one in May 2020 from which we obtained around 100 observations per treatment, and an additional session in June 2021 in which we obtain around 60 observations per treatment. Participants received a show–up fee of $1.50, and earned an additional $3.50 on average for their decisions. The experiment lasted

[5]In particular, one pattern was randomly selected at the end of the experiment, and subjects whose response was accurate received an extra $1.

[6]It is worth mentioning at this stage that we find no evidence of sequence effects.

	Total	Low Raven	High Raven
Selfish Treatment	161	98	63
Charity Treatment	153	98	55

Table 1: Experimental design

approximately 15 minutes. Additional payments varied from \$0 up to \$6.35. We paid participants for one problem from one task at random.

3 Results

In this section, we outline the results from our online experiment. We use a number of common features throughout the analysis. Where non–parametric tests are utilised, both the p–value and test statistic are presented in parentheses. Unless otherwise stated, all tests are two–sided. Where parametric analysis is used, we report the marginal effects, with full regressions provided in Appendix B.

3.1 Reported die rolls

Although we do not observe the true die rolls reported by our subjects, we can consider if the percentage of reported die rolls for each number deviates from an expected percentage. As our subjects were asked to roll a six sided die, we should expect to see each number on the die reported around 16.7% of the time if all subjects report honestly. Deviations from this percentage identified as statistically significant imply that at least some of the subjects are lying.

Table 2 presents the percentage of reported die rolls for each number, for both the *Selfish* and *Charity* treatments, and whether or not the observed percentages are significantly different from 16.7%. Formal testing reveals that the distribution of reported die rolls are significantly different between treatments ($p < 0.05$, Fisher's Exact Test).

Observation 1. *Regardless of treatment, die rolls that provide higher (lower) payoffs are reported more (less) frequently than if all subjects were reporting honestly.*

Support. Table 2 outlines how, in both the *Selfish* and *Charity* treatments, subjects report die rolls of 1 and 2 less frequently than expected ($p < 0.01$ and $p < 0.05$ in all cases, Binomial Test); in the *Selfish* treatment just 4% of subjects reported rolling a 1, and 0% of subjects reported a 1 in the *Charity* treatment. Similarly, subjects report die rolls of 4 and 5 more frequently than expected, with 32% of subjects reporting a 5 in the *Selfish* treatment, almost twice the expected amount, whereas 56% of subjects reported a 5 in the *Charity* treatment ($p < 0.01$ in all cases, Binomial Test). Although a die roll of 6 is reported significantly less often than expected, we still observe 6% of subjects in the *Selfish* treatment, and around 3% of subjects in the *Charity* treatment, reporting a roll of 6 ($p < 0.01$ in both cases, Binomial Test).

Treatment	1	2	3	4	5	6
	Reported Die Roll					
Selfish	4.3***	6.8***	14.9	31.7***	32.9***	6.2***
Charity	0***	5.2***	9.8**	26.1***	56.2***	2.6***

Note: The figures show the percentage of reported die rolls for each number in a given treatment. ***, ** and * indicate significance at the 1%, 5% and 10% levels. Significance implies the percentage of people reporting that number is significantly different to 0.167, at the respective significance level. p–values calculated using two–sided binomial tests.

Table 2: The percentage of reported die rolls

Observation 1 highlights two things. First, regardless of the treatment they receive, a significant number of subjects appear to dishonestly report their die roll. However, we still observe a significant number of honest subjects - we find no evidence that subjects misreport rolling a 3 in the *Selfish* treatment ($p > 0.1$ in both cases, Binomial Test). Second, even though our subjects are drawn from an entirely different population to those of Fischbacher & Föllmi-Heusi (2013), the reported percentages relating to the *Selfish* treatment in Table 2 closely replicates their findings.

We now consider how behaviour compares across treatments. This is done by estimating three ordered Probit regressions. In each case, we use the reported die roll as the dependent variable. We always take the *Selfish* treatment as the baseline, and include a variable that takes a value of 1 if the subject received the *Charity* treatment (and 0 otherwise). In Model 1 we only include the *High* dummy as an explanatory variable. In Model 2 we add a control for the subjects' reported income. In Model 3 we add additional controls for social preferences, race, gender and their political party affiliation.

Table 3 presents the estimated marginal effect of the *Charity* treatment dummy on the probability of each die roll being reported.

Observation 2. *Subjects are less likely to report a roll of 1, 2, 3 and 4, but more likely to report a 5, when they receive the Charity treatment in comparison to the Selfish treatment.*

Support. Table 3 shows that, across all model specifications, that the marginal effect of the *Charity* treatment on the probability of reporting a 1, 2, 3 and 4 is negative ($p < 0.05$ in all cases, T–Tests), being between 2%-5% less likely to be reported respectively (Model 3 columns 1, 2, 3, 4). In contrast subjects are approximately 10%-11% more likely to report a 5 ($p < 0.01$, Model 3 Column 5, T–Test) when they receive the *Charity* treatment compared to the *Selfish* treatment.

Observation 2 highlights how subjects are more willing to misreport dies rolls when their misreporting benefits a third party. As there is a moral justification for lying, it is more prevalent. This is similar to previous findings of Erat & Gneezy (2012), who show that subjects' lie more frequently when their lies increase the payoffs of others. Interestingly, we also find that rolls of six are around 3% more likely to be reported in the *Charity* treatment, although this is likely a consequence of the small percentage of reported sixes, as evidenced in Table 2.

	Marginal effect of 'Charity Treatment' on each die roll					
	1	2	3	4	5	6
Model 1	-0.017**	-0.033**	-0.045**	-0.038**	0.101***	0.032**
	(0.008)	(0.014)	(0.018)	(0.015)	(0.037)	(0.013)
Model 2	-0.017**	-0.034**	-0.045**	-0.037**	0.101***	0.032**
	(0.008)	(0.014)	(0.018)	(0.015)	(0.037)	(0.013)
Model 3	-0.016**	-0.032**	-0.043**	-0.035**	0.096**	0.031**
	(0.008)	(0.014)	(0.018)	(0.015)	(0.038)	(0.014)

Note: The figures show the marginal effect of the *Charity* treatment on the probability that a particular die roll is reported, estimated from an Ordered Probit regression. The *Selfish* treatment is taken as the baseline. ***, ** and * indicate significance at the 1%, 5% and 10% levels. p–values calculated using standard t-tests.

Table 3: Treatment effects on lying - Marginal effects

3.2 Intelligence

We now turn to examine our main research question, how subjects' cognitive ability correlates with misreporting behaviour, and how this interacts with the *Selfish* and *Charity* treatments. As outlined in Section 2, we divide subjects into *Low* and *High* Raven depending on their score in a ten item Raven Test - those below and equal to the median are identified as *Low*, and those above the median, *High*.[7] This is done following Proto et al. (2019), although we divide subjects into *High* and *Low* based on the entire sample, not only at the session level.[8] We focus on *Low* and *High* Raven scores, rather than the Raven score *per se*, because we are interested in how general levels of intelligence correlate with decision making.

Figure 1 outlines the distribution of correct answers achieved by subjects on the Raven test, and also highlights the mean and median number of correct answers.

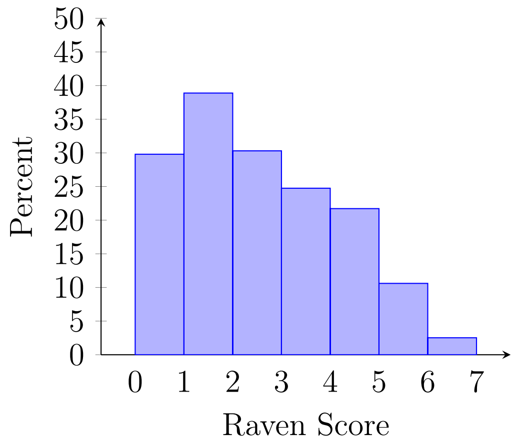

(a) Distribution of Raven scores

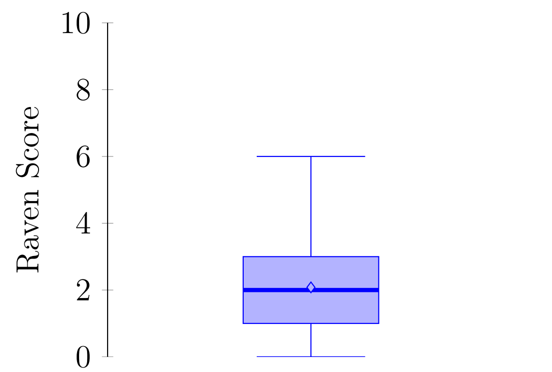

Note: Thick line represents the median Raven score, and the circle represents the mean Raven score.

(b) Boxplot of Raven scores

Figure 1: Raven scores

There are some observable differences between *Low* and *High* Raven subjects. Table A4 in the Appendix presents summary statistics about both *Low* and *High* Raven subjects. It also presents the p-values of the test of the hypothesis that the *Low* Raven subjects characteristics are equal to the *High* Raven subjects. As some variables are categorical (education and race), we present the modal response. As can be seen, *Low* Raven subjects have a significantly higher self–reported income, lower tolerance for risk and a significantly different distribution of social preference scores than *High* Raven subjects ($p < 0.05$ in all cases, various tests). Thus, we control for each of these variables within our parametric analysis.

[7] Although we divide subjects into *Low* and *High* Raven based on the median Raven score of the population of subjects, our results are robust to using the mean.

[8] We deviate slightly from Proto et al. (2019), who also assign subjects into *High* and *Low* Raven based on educational attainment if they scored equal to the median. This was done in their experiment in order to ensure an equal split of *High* and *Low* subjects in each session. As we do not require an equal split of subjects for our analyses, it seems reasonable to classify subjects that are above average as *High*.

Raven - Treatment	*Reported Die Roll* 1	2	3	4	5	6
Low - Selfish	1***	6.1***	12.2	34.7***	33.7***	7.1***
Low - Charity	0***	3.1***	11.2	25.5**	57.1***	3.1***
High - Selfish	9.5	7.9*	19	27**	31.7***	4.8***
High - Charity	0***	9.1	7.3*	27.3**	54.5***	1.8***

Note: The figures show the percentage of reported die rolls for each number in a given treatment. ***, ** and * indicate significance at the 1%, 5% and 10% levels. Significance implies the percentage of people reporting that number is significantly different to 0.167, at the respective significance level. p–values calculated using two–sided binomial tests.

Table 4: The percentage of reported die rolls by Raven score

3.2.1 Intelligence and lying behaviour

To examine if there are any differences in reported die rolls formally, Table 4 presents the percentage of reported die rolls for each number on the die, by Raven score, and by treatment. The table highlights the percentages that are significantly different to the expected 16.7% had all subjects behaved honestly.

To determine if there exists differences in misreporting between *Low* and *High* Raven subjects, and between treatments, we estimate the coefficients and marginal effects from three ordered Probit models. In each model, the reported die roll is the dependent variable, and we always include a dummy variable for *High* Raven (1 if *High*, 0 otherwise). We then increase the number of controls in each subsequent model. In Model 1, we only include the *High* Raven dummy. In Model 2, we control for subjects' level of income. In Model 3 we also control for age, social preferences, race, gender, if they report supporting the Republican party, and their level of education. We include these variables to control for differences between *Low* and *High* Raven groups that might otherwise explain our results.

The marginal effects of *High* Raven on reported die rolls for the *Selfish* treatment are given in Table 5, Panel A, and those for the *Charity* treatment are given in Panel B.[9]

Observation 3. *High Raven subjects are 9% less likely to report the payoff maximising die roll of 5, and between 3-5% more likely to report a 1, 2 and 3 in the Selfish treatment compared to Low Raven subjects. High Raven subjects' behaviour is indistinguishable from Low Raven subjects in the Charity treatment.*

Support. Table 5, Panel A, highlights how the marginal effect of *High* Raven on the probability that a roll of 1, 2, 3, 5 and 6 is reported in the *Selfish* treatment is significant at the 5% level ($p < 0.05$ across models). The estimated marginal effect on 5 is estimated to be negative and around 9% ($p < 0.05$ in all cases, T–Tests), whereas the marginal effect is positive on 1, 2 and 3. This suggests that *High* Raven subjects are significantly more likely to report 1, 2 and 3 and less likely to report a 5 and 6 than are *Low* Raven subjects in the

[9] As not a single roll of 1 was reported in the *Charity* treatment, it is impossible to estimate the marginal effect of *High* Raven on this die roll.

Selfish treatment. The estimates on rolls of 1 and 6 should be interpreted with care, given the small number of observations.

Further, the estimates in Table 5, Panel B highlight how the marginal effect of the *High* Raven dummy is never significant at conventional levels ($p > 0.1$ in all cases, T–Tests). This suggests that *Low* and *High* Raven subjects behave identically in the *Charity* treatment. The marginal effect of *High* on rolls of 1 and 6 in the *Charity* treatment cannot be computed from such a small percentage of observations, as evidenced in Table 4.

	Marginal effect of High Raven on each die roll					
	1	2	3	4	5	6
Panel A						
Model 1	0.032*	0.035*	0.047**	0.02	-0.091**	-0.044*
	(0.018)	(0.019)	(0.024)	(0.013)	(0.043)	(0.024)
Model 2	0.033*	0.035*	0.047**	0.019	-0.09**	-0.044*
	(0.019)	(0.019)	(0.024)	(0.013)	(0.043)	(0.024)
Model 3	0.031*	0.035*	0.047**	0.018	-0.088**	-0.044*
	(0.018)	(0.019)	(0.024)	(0.012)	(0.042)	(0.024)
Panel B						
Model 1	-	0.017	0.02	0.024	-0.051	-
	-	(0.021)	(0.024)	(0.029)	(0.061)	-
Model 2	-	0.008	0.009	0.011	-0.023	-
	-	(0.021)	(0.025)	(0.029)	(0.063)	-
Model 3	-	0.013	0.015	0.018	-0.038	-
	-	(0.022)	(0.025)	(0.031)	(0.065)	-

Note: The figures show the marginal effect of *High* Raven on the probability that a particular die roll is reported, estimated from Ordered Probit regressions. *Low* Raven is taken as the baseline. ***, ** and * indicate significance at the 1%, 5% and 10% levels. p–values calculated using standard t-tests. Panel A presents the estimates for the *Selfish* treatment, and Panel B for the *Charity* treatment.

Table 5: The effect of intelligence on reported die rolls

4 Robustness Checks

In this section we conduct a number of checks in order to examine the robustness of our results.

4.1 Multiple hypothesis testing and post–study probabilities

We first correct all our calculated p-values using the Holm–Bonferroni procedure in order to account for multiplicity. We present all these corrected p–values in Tables A7 in Appendix B. The reported significance remain as in Section 3. This suggests our results are robust to criticisms related to issues of multiple hypothesis testing.

Second, in order to examine the extent to which we might incorrectly reject the null hypothesis of our binomial tests presented in Table 2 and Table 4, we calculate the power associated with these tests. In Table 2, those tests associated with the *Selfish* and *Charity treatments*, all tests are powered at 99%, except for those examining reported die rolls of three. Similarly, in Table 4 any of the tests for the *High* Raven subjects in which the reported percentage of die rolls is either equal to 4% or below, or above 30% is powered at a standard 90% statistical power. For the *Low* Raven, they are powered at similar levels to those rolls in Table 2. As these are the tests we rely on to support our results, we conclude that the probability of us making Type II error is sufficiently low.

Finally, following Maniadis et al. (2014), we calculate the post-study probability estimates (PSPs) associated with our tests that have a power of 90% or more. We calculate the probability that our study will be replicated given its power of 90%, for priors of 1% to 50%, and for situations where there are 1 to 10 other teams also working on a similar study. The table can be found in Appendix B; if only we have conducted this experiment, and we had a prior belief that our effects would occur with a 30% probability, our experiment would replicate with 90% probability; if our prior was that our results would occur with 50% probability, then given the power of our experiment along with the sample size, our experiment would replicate with 95% probability.

4.2 Attention checks and compliance

One possible interpretation of our results is that the difference in Raven scores represent a difference in attention to the experiment, with those with lower scores simply not spending as much time completing the tasks. To address this, we examine the amount of time subjects took to complete the experiment. Table 6 presents the average completion duration of subjects, disaggregated by Raven score.

	Low Raven	*High* Raven
Duration	688 seconds (555)	627 seconds (294)

Note: Standard deviations in parentheses.

Table 6: Duration of experiment in seconds

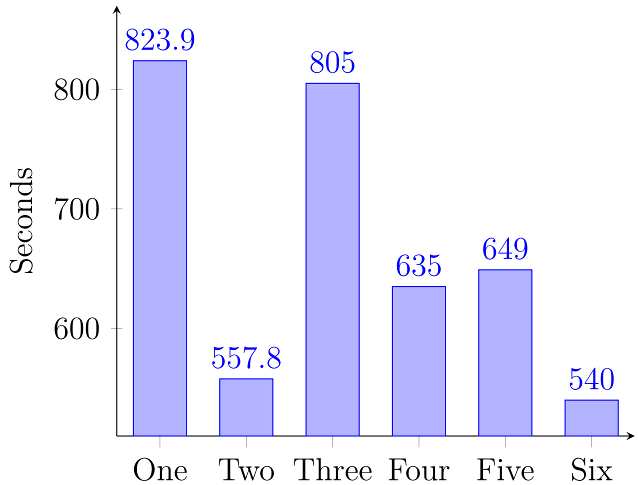

Figure 2: Duration of experiment by reported die roll

Table 6 outlines how subjects take a very similar amount of time to complete the experiment (around 10 minutes), and there is no significant difference between groups ($p > 0.1$, Robust Rank Order test). There is also no significant correlation between duration of the experiment and the number of correct answers on the Raven Test ($\rho = -0.05$, $p > 0.1$, Correlation Coefficient).

A second possible interpretation is that the reported die roll represents the extent to which a subject complied with the experimental instructions, with those subjects that reported the most profitable die roll (a five) being the subjects that simply did not follow the instructions, and therefore should take the shortest amount of time. Comparing completion times by die roll should shed some light on this possible interpretation; Figure 2 presents the response times for each reported die roll.

As can be seen in Figure 2, those who report fives (the most profitable die roll) do not take the shortest amount of time - those who report sixes and twos take very similar small amounts of time. Interestingly, those who report a 1 and three take the longest amount of time. There is no significant correlation between die rolls and duration ($\rho = -0.03$, $p > 0.1$, correlation coefficient). We take this as evidence that the report die rolls do not represent the amount of attention given to the experiment.

5 Conclusion

Our paper sheds light on the conceptual link between intelligence and decision making. We report evidence that individuals with higher intelligence (or cognitive ability) behave more honestly than those with a lower cognitive ability when lying benefits themselves. However, we find that higher intelligence individuals increase the extent of their dishonesty, and behave identically to those with low intelligence, when lying benefits a third party. In our case, the third party is a well known international wildlife protection charity. This finding is robust to controlling for a wide range of individual and behavioural measures that previous work has deemed important for decision making. Our results resonate with a growing literature that examines how personality and intelligence impact on decision making, and opens a number of avenues for future research.

The behavioural disparity in lying observed between low and high intelligence subjects has a number implications. First, it has implications for the development of theoretical models that seek to explain the motivation behind lying and honesty. Whereas previous work has shown that intelligence is important in strategic interactions, we show that even in simple decision tasks that are not cognitively demanding, cognition still plays an important role. Second, our study brings to attention the importance of intelligence for the empirical analysis of honest and dishonest behaviours. Previous work has shown that gender, reputation, and social norms are important considerations for honest behaviour, and our findings complement this work. Finally, from an organizational and managerial perspective, understanding the individual characteristics that correlate with 'desirable' behaviours in the workplace can be useful for identifying employees, and creating a more trustworthy workforce.

Although we examine this link in one particular environment, and using one particular paradigm, results from experiments such as ours provide insights upon which future work can build. Future avenues of research could further examine the link between intelligence and dishonesty in a range of interactions and contexts to determine the generalisability of the observations reported here, using various stake sizes, field contexts, decisions and under differing levels of anonymity (Levitt & List, 2007). Observations from our own study, and these future studies, suggest important questions for the theory of why people lie.

References

Abeler, J., Nosenzo, D. & Raymond, C. (2019), 'Preferences for truth-telling', *Econometrica* **87**(4), 1115–1153.

Becker, G. S. (1968), Crime and punishment: An economic approach, *in* 'The economic dimensions of crime', Springer, pp. 13–68.

Burke, H. R. (1958), 'Raven's progressive matrices: A review and critical evaluation', *The Journal of Genetic Psychology* **93**(2), 199–228.

Burks, S. V., Carpenter, J. P., Goette, L. & Rustichini, A. (2009), 'Cognitive skills affect economic preferences, strategic behavior, and job attachment', *Proceedings of the National Academy of Sciences* **106**(19), 7745–7750.

Cattell, R. B. (1963), 'Theory of fluid and crystallized intelligence: A critical experiment.', *Journal of educational psychology* **54**(1), 1.

Croson, R. & Gneezy, U. (2009), 'Gender differences in preferences', *Journal of Economic literature* **47**(2), 448–74.

Erat, S. & Gneezy, U. (2012), 'White lies', *Management Science* **58**(4), 723–733.

Fe, E., Gill, D. & Prowse, V. (2020), 'Cognitive skills, strategic sophistication, and life outcomes', *IZA Working Paper Series* .

Fischbacher, U. & Föllmi-Heusi, F. (2013), 'Lies in disguise—an experimental study on cheating', *Journal of the European Economic Association* **11**(3), 525–547.

Gächter, S. & Schulz, J. F. (2016), 'Intrinsic honesty and the prevalence of rule violations across societies', *Nature* **531**(7595), 496–499.

Gerlach, P., Teodorescu, K. & Hertwig, R. (2019), 'The truth about lies: A meta-analysis on dishonest behavior.', *Psychological bulletin* **145**(1), 1.

Gill, D. & Prowse, V. (2016), 'Cognitive ability, character skills, and learning to play equilibrium: A level-k analysis', *Journal of Political Economy* **124**(6), 1619–1676.

Gino, F. & Ariely, D. (2012), 'The dark side of creativity: original thinkers can be more dishonest.', *Journal of personality and social psychology* **102**(3), 445.

Gino, F., Ayal, S. & Ariely, D. (2009), 'Contagion and differentiation in unethical behavior: The effect of one bad apple on the barrel', *Psychological science* **20**(3), 393–398.

Gneezy, U. (2005), 'Deception: The role of consequences', *American Economic Review* **95**(1), 384–394.

Gneezy, U., Kajackaite, A. & Sobel, J. (2018), 'Lying aversion and the size of the lie', *American Economic Review* **108**(2), 419–53.
URL: *https://www.aeaweb.org/articles?id=10.1257/aer.20161553*

Guo, Q., Sun, P., Cai, M., Zhang, X. & Song, K. (2019), 'Why are smarter individuals more prosocial? a study on the mediating roles of empathy and moral identity', *Intelligence* **75**, 1–8.

Heck, D. W., Thielmann, I., Moshagen, M. & Hilbig, B. E. (2018), 'Who lies? a large-scale reanalysis linking basic personality traits to unethical decision making', *Judgment and Decision making* **13**(4), 356.

Holt, C. A. & Laury, S. K. (2002), 'Risk aversion and incentive effects', *American Economic Review* **92**(5), 1644–1655.

Houser, D., List, J. A., Piovesan, M., Samek, A. S. & Winter, J. (2015), On the Origins of Dishonesty: From Parents to Children, NBER Working Papers 20897, National Bureau of Economic Research, Inc.
URL: *https://ideas.repec.org/p/nbr/nberwo/20897.html*

Irwing, P. & Lynn, R. (2006), 'The relation between childhood iq and income in middle age', *Journal of Social Political and Economic Studies* **31**(2), 191.

Kerschbamer, R. (2015), 'The geometry of distributional preferences and a non-parametric identification approach: The equality equivalence test', *European Economic Review* **76**, 85 – 103.

Kerschbamer, R. & Müller, D. (2020), 'Social preferences and political attitudes: An online experiment on a large heterogeneous sample', *Journal of Public Economics* **182**, 104076.

Levitt, S. D. & List, J. A. (2007), 'What do laboratory experiments measuring social preferences reveal about the real world?', *Journal of Economic Perspectives* **21**(2), 153–174.
URL: *https://www.aeaweb.org/articles?id=10.1257/jep.21.2.153*

Maniadis, Z., Tufano, F. & List, J. A. (2014), 'One swallow doesn't make a summer: New evidence on anchoring effects', *American Economic Review* **104**(1), 277–90.

Mazar, N., Amir, O. & Ariely, D. (2008), 'The dishonesty of honest people: A theory of self-concept maintenance', *Journal of Marketing Research* **45**(6), 633–644.

Paulhus, D. L. & Dubois, P. J. (2015), 'The link between cognitive ability and scholastic cheating: A meta-analysis', *Review of General Psychology* **19**(2), 183–190.

Proto, E., Rustichini, A. & Sofianos, A. (2019), 'Intelligence, personality, and gains from cooperation in repeated interactions', *Journal of Political Economy* **127**(3), 1351–1390.

Proto, E., Rustichini, A. & Sofianos, A. (2020), 'Intelligence, errors and strategic choices in the repeated prisoners' dilemma interactions', *Unpublished* .

Raven, J. & Raven, J. (2003), *Raven Progressive Matrices*, Springer US, Boston, MA, pp. 223–237.

Rindermann, H., Becker, D. & Thompson, J. (2018), 'Honesty, rule violation and cognitive ability: A reply to gächter and schulz', *Intelligence* **68**, 66–69.

Speer, S. P. H., Smidts, A. & Boksem, M. A. S. (2020), 'Cognitive control increases honesty in cheaters but cheating in those who are honest', *Proceedings of the National Academy of Sciences* **117**(32), 19080–19091.

Sundet, J. M., Barlaug, D. G. & Torjussen, T. M. (2004), 'The end of the flynn effect?: A study of secular trends in mean intelligence test scores of norwegian conscripts during half a century', *Intelligence* **32**(4), 349–362.

Tonin, M. & Vlassopoulos, M. (2013), 'Experimental evidence of self-image concerns as motivation for giving', *Journal of Economic Behavior & Organization* **90**, 19–27.

Van't Veer, A., Stel, M. & van Beest, I. (2014), 'Limited capacity to lie: Cognitive load interferes with being dishonest.', *Judgment and Decision Making* **9**(3), 199–206.

Zettler, I., Hilbig, B. E., Moshagen, M. & de Vries, R. E. (2015), 'Dishonest responding or true virtue? a behavioral test of impression management', *Personality and Individual Differences* **81**, 107–111.

Appendix - For online publication

A Experimental Instructions

Introduction

Welcome

Thank you for participating in this HIT. It should take around 14 minutes to complete, and you will receive 1 dollar 50 cents for participating. In addition, there is a chance to earn additional money throughout the HIT. The maximum the additional earnings can be is 6 dollars and 35 cents. You will receive any additional money you earn within 10 days of the HIT closing.

In this HIT, we will ask you to complete some tasks and answer some questions.

After you have finished the task, you will receive a completion code. Please return to the HIT on MTurk and enter the completion code in the space provided to receive your credit.

This HIT is part of a University of Exeter and University of Birmingham scientific research project. Your decision to complete this HIT is voluntary. There is no way for us to identify you. The only information we will have, in addition to your responses, is the time at which you completed the survey. The results of the research may be presented at scientific meetings or published in scientific journals. Choosing the 'I agree' option below indicates that you are at least 18 years of age and agree to complete this HIT voluntarily.

Instructions

This HIT consists of two Parts, Part A and Part B. You will complete them in a random order.

Part A

Part A consists of three tasks, Task A, Task B and Task C. You will complete them in a random order.

You will be paid for one question from one task in Part A. The question and task will be chosen at random.

Task A

In this task, you will be asked to make nine decisions. For each decision, you will be asked to choose between two lotteries. These lotteries will be labelled either LEFT or RIGHT. Whether you select LEFT or RIGHT will determine how much money you earn.

Each lottery has two possible outcomes. The lottery on the LEFT gives $2 with some probability, and $1.60 with some probability. The lottery on the RIGHT gives $3.85 with some probability and $0.10 with some probability. The probabilities assigned to each outcome are varied in each decision.

You should think of the probabilities as follows. A probability of 50% assigned to each outcome means that each outcome is equally likely, and the same as flipping a coin - heads one outcome, tails the other. A probability of 100% means an outcome is certain, and a probability of 0% means an outcome will never happen.

If Task A is chosen for payment, one decision will be chosen at random, and the lottery you choose will be played and you will receive payment for it.

Decision	LEFT	RIGHT
1.	10% chance of $2, 90% chance of $1.80	10% chance of $3.85, 90% chance of $0.10
2.	20% chance of $2, 80% chance of $1.80	20% chance of $3.85, 80% chance of $0.10
3.	30% chance of $2, 70% chance of $1.80	30% chance of $3.85, 70% chance of $0.10
4.	40% chance of $2, 60% chance of $1.80	40% chance of $3.85, 60% chance of $0.10
5.	50% chance of $2, 50% chance of $1.80	50% chance of $3.85, 50% chance of $0.10
6.	60% chance of $2, 40% chance of $1.80	60% chance of $3.85, 40% chance of $0.10
7.	70% chance of $2, 30% chance of $1.80	70% chance of $3.85, 30% chance of $0.10
8.	80% chance of $2, 20% chance of $1.80	80% chance of $3.85, 20% chance of $0.10
9.	90% chance of $2, 10% chance of $1.80	90% chance of $3.85, 10% chance of $0.10

Table A1: Task A

Task B

In this task, you will be asked to consider 10 different patterns. Each pattern has an image left out. All you have to do is select the image that is right in order to complete the pattern. The patterns will start simple, and increase in difficulty as you go on. Each pattern has one correct answer. You will have a maximum of 30 seconds to complete each pattern.

At the end of the survey if this task is chosen for payment one pattern will be selected at random and if your answer is correct, you will be given $1.

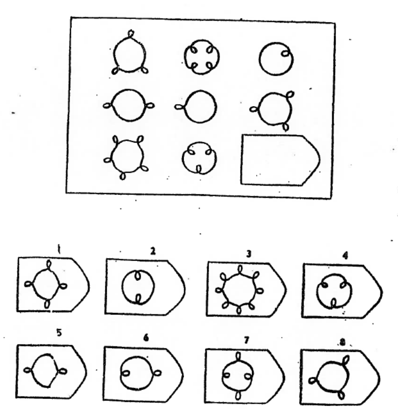

Table A2: Task B (example)

Task C

In Task C you will be asked to make ten decisions. For each decision, you will be asked to choose between two allocations. These allocations will be labelled either LEFT or RIGHT. Whether you select LEFT or RIGHT will determine how much money you earn, and how much money is donated to charity by the experimenter.

If Task C is chosen for payment, one decision will be selected at random and the charity we will donate any money to is the WORLD WILDLIFE FUND.

Decision	LEFT	RIGHT
1.	$0.50 to you, $1 to the charity	$0.70 to you, $0.70 to the charity
2.	$0.60 to you, $1 to the charity	$0.70 to you, $0.70 to the charity
3.	$0.70 to you, $1 to the charity	$0.70 to you, $0.70 to the charity
4.	$0.80 to you, $1 to the charity	$0.70 to you, $0.70 to the charity
5.	$0.90 to you, $1 to the charity	$0.70 to you, $0.70 to the charity
6.	$0.50 to you, $0.50 to the charity	$0.70 to you, $0.70 to the charity
7.	$0.60 to you, $0.50 to the charity	$0.70 to you, $0.70 to the charity
8.	$0.70 to you, $0.50 to the charity	$0.70 to you, $0.70 to the charity
9.	$0.80 to you, $0.50 to the charity	$0.70 to you, $0.70 to the charity
10.	$0.90 to you, $0.50 to the charity	$0.70 to you, $0.70 to the charity

Table A3: Task C

Part B - selfish treatment

In Part B you can receive a bonus.

However, this bonus is not the same for every participant. You determine your own bonus by throwing a six-sided die. You should throw the die once. This decides how much of a bonus you receive. You can see the exact bonus you receive from your throw from the following chart. It will remain on the screen until you have entered your throw.

Die Roll	Bonus
1	$1
2	$2
3	$3
4	$4
5	$5
6	$0

Part B - charity treatment

In Part B you will receive a bonus that will be donated to charity.

The charity is the Wildlife Conservation Society.

However, this bonus is not the same for every participant. You determine your own bonus by throwing a six-sided die. You should throw the die once. This decides how much of a donation the charity will receive. You can see the exact donation the charity will receive from your throw from the following chart. It will remain on the screen until you have entered your throw.

Die Roll	Donation
1	$1
2	$2
3	$3
4	$4
5	$5
6	$0

B Additional Analyses

In this Appendix we include any additional analyses and tables referred to in the main text.

	Low Raven	*High Raven*	H_0 :*Low=High*
Income†	6.645	5.864	0.01
*Education**	Bachelor's Degree	Bachelor's Degree	0.154
*Race**	White	White	0.027
Gender	0.688	0.61	0.16
Party	Republican	Republican	0.047

Note: The fourth column presents p–values from the test of the null hypothesis that Low and High Raven characteristics are equal (Fischer's exact test for nominal variables, Wilcoxon Tests for all other variables).

† Category 6 for income represents an annual income of $60,000-69,000, and category 7 $70,000-79,000.

* Reported values represent averages: Education and race are modal averages. Although the modal value is identical between groups, the distributions differ for race. All other variables are mean averages.

** Risk preferences is the average 'switching point' in the Holt & Laury (2002) test.

*** Social preferences is the average XY score calculated from the EET (Kerschbamer, 2015).

Table A4: Summary of subjects' characteristics

	Model 1	Model 2	Model 3
Charity treatment	0.334***	0.338***	0.323***
	(0.122)	(0.123)	(0.127)
Income		0.049**	0.049**
		(0.022)	(0.022)
Social Pref. Score			0.043
			(0.07)
Black			-0.05
			(0.15)
Female			-0.015
			(0.131)
Republican			0.136
			(0.435)

Table A5: Full ordered Probit estimates - Table 3

	Selfish Treatment - Panel A			*Charity Treatment - Panel B*		
	Model 1	Model 2	Model 3	Model 1	Model 2	Model 3
High Raven	-0.353**	-0.355**	-0.361**	-0.157	-0.072	-0.12
	(0.172)	(0.174)	(0.176)	(0.188)	(0.196)	(0.204)
Income		0.045	0.051*		0.049	0.048
		(0.028)	(0.029)		(0.035)	(0.036)
Age			0.061			-0.026
			(0.076)			(0.193)
Social Pref. Score			-0.33			0.03
			(0.22)			(0.24)
Black			0.153			-0.197
			(0.177)			(0.221)
Female			0.107**			0.00
			(0.044)			(0.042)
Republican			-0.096*			-0.063
			(0.052)			(0.052)
Education			-0.003			0.002
			(0.006)			(0.007)

Table A6: Full ordered Probit estimates - Table 5

	Reported Die Roll					
Raven - Treatment	1	2	3	4	5	6
Low - Selfish	0.00***	0.038**	0.558	0.00***	0.001***	0.103
Low - Charity	0.00***	0.001***	0.525	0.289	0.00***	0.001***
High - Selfish	0.696	0.443	0.612	0.365	0.044**	0.079*
High - Charity	0.001***	0.748	0.414	0.362	0.00***	0.013**

Note: The reported figures are p–values calculated from binomial tests examining if each probability is equal to 0.167, corrected using the Holm–Bonferroni procedure. We treat all 24 tests as being in the same 'family' of tests. ***, ** and * indicate significance at the 1%, 5% and 10% levels, after the correction has been applied. Significance implies the percentage of people reporting that number is significantly different to 0.167, at the respective significance level once the correction has been applied.

Table A7: Holm–Bonferoni Corrected p–values

Competing Studies	1	2	5	10
Prior				
0.01	0.15	0.08	0.03	0.01
0.02	0.27	0.14	0.05	0.02
0.05	0.49	0.30	0.12	0.04
0.10	0.67	0.48	0.22	0.09
0.20	0.82	0.68	0.39	0.18
0.30	0.89	0.78	0.53	0.27
0.40	0.92	0.85	0.64	0.37
0.50	0.95	0.89	0.72	0.46

Note: Numbers represent probabilities. Prior is the prior probability that the effects will be reported, and competing studies is the number of potential competing studies answering the same research questions.

Table A8: Post–study probabilities