

Is there a link between intelligence and lying?

Drouvelis, Michalis; Pearce, Graeme

Journal of Economic Behavior and Organization

DOI: https://doi.org/10.1016/j.jebo.2022.12.008

Published: 01/02/2023

Publisher's PDF, also known as Version of record

Cyswllt i'r cyhoeddiad / Link to publication

Dyfyniad o'r fersiwn a gyhoeddwyd / Citation for published version (APA): Drouvelis, M., & Pearce, G. (2023). Is there a link between intelligence and lying? Journal of Economic Behavior and Organization, 206, 182-203. https://doi.org/10.1016/j.jebo.2022.12.008

Hawliau Cyffredinol / General rights Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

· Users may download and print one copy of any publication from the public portal for the purpose of private study or research.

- You may not further distribute the material or use it for any profit-making activity or commercial gain
 You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Contents lists available at ScienceDirect

Journal of Economic Behavior and Organization

journal homepage: www.elsevier.com/locate/jebo

Is there a link between intelligence and lying?[☆]

Michalis Drouvelis^{a,b,*}, Graeme Pearce^c

^a University of Birmingham, United Kingdom ^b CESifo, Munich, Germany ^c Bangor University, United Kingdom

ARTICLE INFO

Article history: Received 27 October 2021 Revised 10 November 2022 Accepted 12 December 2022

JEL classifications: C90 Z13

Keywords: Intelligence levels Die roll paradigm Honesty

1. Introduction

ABSTRACT

In this paper we study the link between intelligence and lying. This is done under two distinct experimental paradigms: the die rolling paradigm, and the mind game paradigm. We compare the behaviour of subjects conditional on their intelligence in each of these paradigms when lying benefits the subject who lies (Selfish treatment) or a charitable cause (Charity treatment). We report some evidence that lying increases when it benefits a charitable cause, and limited evidence that more intelligent individuals increase lying when doing so benefits a charity. Our results are consistent with models of self-image concerns, and could have important implications for the existing lying literature.

© 2022 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license (http://creativecommons.org/licenses/by-nc-nd/4.0/)

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 extent 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 there a link between intelligence and lying?

A growing literature in behavioural economics has recognised the role of intelligence in various aspects of strategic decision making. For example, Gill and Prowse (2016) show that cognitive skills determine the evolution of Nash equilibrium

* Corresponding author at: University of Birmingham, United Kingdom.

E-mail address: m.drouvelis@bham.ac.uk (M. Drouvelis).

https://doi.org/10.1016/j.jebo.2022.12.008







^{*} 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.

^{0167-2681/© 2022} The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license (http://creativecommons.org/licenses/by-nc-nd/4.0/)

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 the area of the brain associated with cognitive control is used by participants observed to be 'cheaters' to behave honestly, whereas the same area of the brain is used by participants observed to be honest in order to cheat. 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 and 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 and 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 and Schulz (2016) which they combine with country level estimates of cognitive ability obtained from student assessments. They report that country level measures of cognition positively correlates with misbehavior.

Our paper contributes to this literature by examining how individual–level cognitive ability determines individual decisions to misbehave or not. We are concerned with the pure effects of intelligence on lying, and we study misbehaviour in a one-shot game. This provides us with cleaner tests for the effects of intelligence on lying than repeated games and allows us to rule out behaviour being confounded with strategic considerations arising from repetitions of the game. For our purposes, we employ two paradigms in two separate experiments. First, we consider behaviour in the standard die roll paradigm as introduced by Fischbacher and Föllmi-Heusi (2013) subjects roll a six sided die in private, and the number they report to have rolled determines how much they earn. Second, we examine the decisions of subjects in the mind game (see Jiang 2013, Potters and Stoop 2016, Kajackaite and Gneezy 2017 for examples of other implementations): subjects are asked to think of a number between 1 and 6, and then generate a random number between 1 and 6 using a computer. If the subject reports *Yes*, that the random number is the same as the one they were thinking of, they receive a payoff. If they report *No*, they receive 0. Our focus is therefore on cheating games that do not involve strategic uncertainty (unlike deception games, such as Gneezy (2005)). In addition, we conduct online experiments, making it impossible for the experimenter to know the true outcome of the die roll and for anyone to observe if a subject misbehaves or not. This generates an even more transparent environment for the subject, where the probability of getting caught is effectively zero.

In both experiments, we examine 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 and 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).^{1,2} Using standard parametric techniques, we then estimate the impact that Raven score has on the probability of certain types of behaviour.

A novel aspect of our study is the inclusion of an 'opt-out' decision at the end of the second experiment involving the mind game. Once all decisions have been made, and all questionnaire questions answered, subjects are asked if they want to opt-out of having their decision implemented. If they do decide to opt-out, they receive the payoff they would have received if they had reported *Yes* in the mind game. If not, their decision from the mind game is implemented and they receive the payoff associated with that decision. The decision to opt-out can therefore reveal something about any intrinsic costs and benefits incurred by the subject. This question was originally implemented by Tonin and Vlassopoulos (2013) in order to examine the self-image concerns of subjects in dictator games. In our experiment, just as in Tonin and Vlassopoulos (2013), all decisions are anonymous and the experimenter has no way of identifying a particular participant. Further, due to the online setting of our experiment, the experimenter is completely unable to observe if the subject is lying or not, and the only person who actually observes if the subject misbehaved or not, is the participant themselves. Therefore, we interpret any intrinsic costs incurred by the subject as being the result of self-image concerns.

¹ Throughout the paper, we use the terms intelligence, fluid intelligence and cognitive ability interchange- ably.

² 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.

The definition of self-image is 'the idea one has of one's abilities, appearance, and personality,' and is a central assumption in a variety of signaling models that can be used to explain prosociality, cooperation (Tonin and Vlassopoulos, 2013; Van't Veer et al., 2014) and the general notion of identity (Akerlof and Kranton, 2000). As Akerlof and Kranton (2000) suggest, it is assumed that individuals like to think positively about themselves, and therefore that there is a disutility associated with perceiving oneself negatively. Previous work supports this, with research suggesting that more self-aware individuals are fairer and more honest if moral standards are salient (Batson et al., 1999), less aggressive (Froming et al., 1982), less likely to cheat in exams (Diener and Wallbom, 1976; Vallacher and Solodky, 1979), and are less likely to take decisions that harm others (Falk, 2021). As forming a coherent self-image relies on an individual having the capacity for 'reflexive thinking', and the ability to focus on themselves as the object of attention (Leary and Tangney, 2011), it seems entirely reasonable to expect that self-image concerns will be positively correlated with higher levels of intelligence. This hypothesis is supported by previous work on social- image concerns, which highlights how individuals base their decisions on their predictions about the probability of being perceived as a liar by an outside observer (Dufwenberg and Dufwenberg, 2018; Gneezy et al., 2018; Abeler et al., 2019; Khalmetski and Sliwka, 2019), something that requires cognitive sophistication.

We report the following observations. Under the die rolling paradigm subjects in the *Charity* treatment are around 22% more likely to report a five, and 5% less likely to report a one than subjects in the *Selfish* treatment. In contrast, when examining the mind game, we find no significant differences in behaviour stemming from the treatments, with the percentage of subjects reporting *No* and *Yes* being identical across treatments. One explanation for this is that cheating, and the self-image costs associated with it, may be context dependent as suggested by Pascual-Ezama et al. (2015). We do however find that opt-out rates in the mind game are significantly different between treatments, with subjects in the *Charity* treatment more inclined to opt-out of reporting a *No*. We interpret this as being a consequence of the self-image benefit associated with being charitable outweighing that of being as honest.

With respect to intelligence we report two novel observations. First, under the die rolling paradigm we find that higher levels of intelligence are associated with a lower probability of reporting a roll of 1, but no other numbers. However, we find that this is potentially a consequence of multiple hypothesis testing. In the mind game we find that subjects with higher levels of intelligence are significantly more likely to report *Yes*, and only in the *Charity* treatment: a one point increase in Raven score is estimated to increase the probability of reporting a *Yes* by 4.3%. Finally, we find that more intelligent individuals are less likely to opt–out of reporting a *Yes* when assigned to the *Charity* treatment, with a one point increase in Raven score estimated to reduce the probability of opting–out by 6%. One interpretation is therefore, that the link between intelligence and lying depends on the context in which it is being studied: both who benefits, and the way in which lying is examined appear to determine the extent of lying. This is consistent with the previous literature that examines how the decision context impacts the decision to lie (Mazar and Ariely, 2006). Therefore, we find the link between intelligence and lying to be small, and dependent on the context in which it is being studied: both who benefits, and the way in which lying is examined determine the extent of lying.

Our experiment contributes to the existing literature in at least two respects. First, our results indicate that individuals deviate from what standard economic theory might predict: not all subjects lie to the greatest extent. Although this finding itself is not novel, what is new is that we observe this regardless of treatment manipulations and the subjects' level of intelligence. Second, 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 and Gneezy (2012), we examine a case of "white lie" where lying helps others without harming the liar. We find that in such a situation, intelligence does appear to play a role, although small, and potentially a consequence of differences in self-image concerns. Although we acknowledge that we do not identify a causal relationship between intelligence and decision making, something the literature has already identified as posing a significant challenge for a variety of reasons (Dohmen et al., 2018), our findings suggest that it is still important to generate more systematic evidence on the interplay between intelligence and individual decision making and the role played by self-image concerns. Understanding these links will contribute to more sophisticated theories of decision making, and the motivations underpinning human behaviour.

Our paper is organised as follows. Section 2 outlines the two experimental designs and each procedure. Section 3 details our hypotheses relating to our treatments and intelligence. Section 4 presents the experimental results. Section 5 presents some robustness checks, and Section 6 concludes.

2. Experimental design and procedure

2.1. Experiment 1

The first experiment utilises a design in which we examine two variables. Along 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 adopt the die roll paradigm as introduced by Fischbacher and 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

Table 1	
Experiment 1 - experimenta	al design.

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

from lying. Choices in this game have no impact on other participants' payoffs and are not affected by beliefs about others' choices and payoff considerations.³

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 the subject's decision only benefits 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 created an environment in which it was easy for subjects to not tell the truth.

Along a second dimension, we examine 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 and Raven, 2003; Proto et al., 2019, 2020). For each item, subjects were given 30 s, and correct answers were incentivised.⁴ For descriptive purposes, 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. For the analysis, we focus our attention on individual Raven scores.

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 and Müller, 2020). There is also now a rich literature that shows that gender correlates with various economic behaviours (Croson and 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 and 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.⁵ Table 1 provides an overview of the 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 approximately 15 min. Additional payments varied from \$0 up to \$6.35. We paid participants for one problem from one task at random.

³ 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.

⁴ In particular, one pattern was randomly selected at the end of the experiment, and subjects whose response was accurate received an extra \$1.

 $^{^{5}}$ It is worth mentioning at this stage that we find no evidence of sequence effects in reported die rolls (p > 0.1 in all cases, Robust Rank Order tests).

Table 2	
Experiment 2 - experimental de	sign.

	Total	Low Raven	High Raven
Selfish Treatment	167	94	73
Charity Treatment	153	78	75

2.2. Experiment 2

In order to explore the extent to which self-image concerns might explain behavioural differences between *Low* and *High* subjects in Experiment 1, we conduct a complementary second experiment. The experiment uses the same design as Experiment 1, as well as identical wording and presentation (all experimental materials are given in the Appendix). Rather than the die rolling paradigm of Fischbacher and Föllmi-Heusi (2013), in Experiment 2 we instead adopt the one-shot *mind game* paradigm, as recently used by Kajackaite and Gneezy (2017). Under this paradigm subjects are asked to think of a number from 1 to 6, and then roll a die. Subjects then report the number that was rolled, and if the number rolled is the same as the one they were thinking of. If subjects report *Yes*, that the number rolled on the die is the same as the one they were thinking of \$3.50. If they report *No*, the numbers were not the same, they receive \$0. We chose \$3.50 as the positive payoff as this is the average payoff from Experiment 1, and thus keeps stake sizes between experiments comparable. Identical to Experiment 1, in a *Selfish* treatment, the subject's choice impacts their own payoff, and in a *Charity* treatment, the subject's decision instead determines the pay off of a charity (Wildlife Conservation Society).

Rather than roll a die, we instead asked subjects to generate a random number using Google number generator, and provided them with a direct link they could click in order to generate a number from 1 to 6. We used Google to generate random numbers, rather than other smaller websites, as Google has no explicit or implicit link or affiliation to academic research. Thus, this signals to subjects the experimenter cannot observe the numbers they generate and ensured the probability of being caught is essentially zero. We used this instead of a die roll to account for the fact that subjects may not have a die at hand.

In order to shed light on the role of self-image in determining any observed differences between participants, we asked participants if they would like to 'opt-out' of implementing their decision from the mind game. This was asked as a surprise, and was the final question of the entire experiment, after all decisions were made and the demographic questionnaire completed: subjects did not know they would be asked this question when making any of their other responses. If participants chose to opt-out in the *Selfish* treatment they received \$3.50, otherwise their original decision was implemented. Similarly, in the *Charity* treatment, if participants chose to opt-out, the charity received \$3.50, otherwise their original decision was implemented. Therefore, if a subject reports *Yes* in the mind game, opting-out has no impact on monetary payoffs.

Subjects assigned to the Selfish treatment were asked the following:⁶

"..Before carrying out the payment associated with your choice you are now given the opportunity to opt out of the decision... This means that whatever decision you made will not be implemented and you will instead receive \$3.50 at the end of the HIT. If you decide not to opt out then you will receive the payment associated with the decision that was previously selected. To express your choice, tick the corresponding box below."

The inclusion of this question is motivated by the literature on audience effects. Specifi- cally, this question was implemented by Tonin and Vlassopoulos (2013) in order to examine self-image concerns. In our experiment, all decisions are anonymous and the experimenter has no way of identifying a particular participant. The only person who actually observes the participant's choice, and if they misbehaved or not, is the participant themselves.

The interpretation of a subject's choice to opt-out depends on if they reported *No* or *Yes* in the mind game. Reporting *No* implies the subject reported honestly, and that the intrinsic costs from lying are high: the utility from reporting *No* and receiving \$0 is greater than the utility from \$3.50 *and lying*. Opting-out of *No* simply implies that the utility from \$3.50 is greater than the utility from being honest and getting \$0.

In contrast, reporting a Yes could be a lie. This means the decision to opt-out of implementing a choice of Yes implies there is an intrinsic cost associated with the original decision that is incurred if it is implemented: the utility from opting-out and receiving \$3.50 must outweigh the utility from receiving \$3.50 and implementing Yes. This suggests there is an intrinsic cost associated with implementing Yes. Not opting-out of Yes implies it was true, or that the subject suffers no self-image costs from lying.

The experiment was conducted on MTurk in February 2022, using participants located in the USA and completed using Qualtrics. Participants that had taken part in the previous experiment were excluded from taking part. The experiment lasted approximately 11 min. Table 2 presents the experimental design, and the number of observations for each treatment.

⁶ An analogous question was asked in the *Charity* treatment.

3. Hypotheses

We form the following hypotheses with respect to our treatments based on the assumption that subjects like to think positively about themselves, and suffer disutility from perceiving themselves negatively (Akerlof and Kranton, 2000). Our first hypothesis relates to the extent of misbehaviour in each treatment.

Hypothesis 1. Subjects will behave such that they obtain higher payoffs in the charity treatment in comparison to the selfish treatment.

We anticipate that when lying benefits only the person who lies (as in our *Selfish* treatment) the intrinsic self-image cost associated with lying will be higher than if the lying benefits a charity (as in our *Charity* treatment). This is because we expect the utility cost of a subject perceiving themselves as misbehaving for selfish reasons to be larger than the self-image cost of misbehaving for charitable reasons: being a charitable person is likely a more positive self-image than being a selfish person. The literature on why individuals give in laboratory experiments supports this idea (Levitt and List, 2007).

In our first experiment under the die rolling paradigm, we therefore expect more reports of high payoff die rolls in the *Charity* treatment in comparison to the *Selfish* treatment. In our second experiment that utilises the mind game, we therefore predict a larger percentage of subjects in the *Charity* treatment to report *Yes* in comparison to the *Selfish* treatment. Our second hypothesis relates to the opt–out decision under the mind game paradigm.

Hypothesis 2. Subjects are more likely to opt-out of implementing both *No* and *Yes* decisions from the mind game in the *Charity* treatment in comparison to the *Selfish* treatment.

We anticipate that subjects will opt-out of *No* reports more often in the *Charity* treatment in comparison to the *Selfish* because a self-image of honesty will be more important when determining their own payoff. In the *Charity* treatment, when a subject opts-out of an honest *No* decision, they trade-off appearing honest with a self-image benefit from being charitable, thus increasing the likelihood of opting-out.

With regard to intelligence, we hypothesise higher levels of intelligence to be associated with lower reported die rolls in the *Selfish* treatment, and higher reported die rolls in the *Charity* treatment under the die rolling paradigm. Similarly, under the mind game paradigm, we hypothesise that higher levels of intelligence will positively correlate with reports of *No* in the *Selfish* treatment, and *Yes* in the *Charity*. This is our third hypothesis:

Hypothesis 3. More intelligent subjects will report lower payoffs in the *Selfish* treat- ment, and higher payoffs in the *Charity* treatment, than less intelligent subjects.

Hypothesis 3 is justified as follows. We expect self-image concerns to be more important to those who most accurately understand their motives, desires and personality: those most self-aware.⁷ Previous work has shown that self-awareness correlates positively with intelligence across species (Gallup Ir, 1998), and there is also evidence of this in humans specifically. For example, Demetriou and Kazi (2006) finds that intelligence and self-awareness are positively correlated in children from various age groups. Others have found that the accuracy of self-evaluation, as well as the evaluation of the cognitive processes underlying performance and decision making, is more accurate as intelligence increases (Demetriou and Kazi, 2006; Freund and Kasten, 2012; Demetriou et al., 2020). In the behavioural economics literature, the cognitive ability of decision makers has been studied from a number of view points, but the majority are based on the accuracy of decisions and beliefs. Most recently, differences in cognitive ability has been examined by the number of mistakes individuals make, or how accurately they are able to assess information, possibly resulting from slower learning processes in subjects with a lower cognitive ability (Gill and Prowse, 2016; Proto et al., 2020). Other work models an individuals cognitive ability as how accurately they are able to simulate the decision making process of others (Camerer et al., 2004; Gill and Prowse, 2016). It seems reasonable that if more intelligent individuals can more accurately assess others, that they can also more accurately assess themselves. As forming a coherent self-image and thinking reflexively is cognitively demanding, in much the same way that intelligence is needed to predict how actions impact your social-image, we expect individuals with higher intelligence to have a more well defined sense of identity (Akerlof and Kranton, 2000) and place greater importance on how they perceive themselves.

Therefore, we predict that more intelligent individuals will lie less often than low intelligence individuals in the *Selfish* treatment due to them having higher self-image costs from lying. Further, we expect them to lie *more often* in the *Charity* treatment, due to the positive self-image generated from a self-image of being charitable. This implies that we predict the image benefit from being charitable will outweigh the image cost from lying, although it is possible that the opposite is true, or that they are exactly equal.

Our final hypothesis is about opt-out decisions, and leads directly from the prediction about how self-image costs correlate with intelligence.

Hypothesis 4. Individuals with higher intelligence will opt-out of implementing a *Yes* in the mind game less often than lower intelligence individuals. However, this will only hold in the *Charity* treatment.

⁷ Self-aware, as defined in the Oxford English Dictionary is someone who is, "Conscious of one's own feelings, character, etc".

Table 3				
The percentage	of	reported	die	rolls.

	Reported	d Die Roll				
Treatment	1	2	3	4	5	6
Selfish	4.4***	7.1***	15.4	32.7***	34***	6.4***
Charity	0***	5.3***	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.

We anticipate that more intelligent individuals will place greater importance on a charitable self-image than lower intelligence individuals. This means the benefit from acting charitably when they lie will be larger as intelligence increases, reducing the motivation to opt-out for higher intelligence individuals.

Our hypotheses complement the findings of Erat and Gneezy (2012), who find that subjects lie more when both the individual and the other person benefit from doing so.⁸ Those studies that have examined this link previously 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 may impact self-image in different ways. 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.⁹

4. Results

In this section, we outline the results from our online experiments. 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 of the variable of interest.

Throughout the results section, we present descriptive statistics in terms of *Low* and *High* Raven subjects. This is done for ease of exposition, as it simplifies the figures and tables. However, none of our observations rely on a binary *Low/High* distinction, and all our results relating to intelligence are determined by Raven score *per se*.

4.1. Experiment 1

Although we do not observe the true die rolls reported by our subjects in Experiment 1, 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 3 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 3 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).

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 find that subjects report rolling a 3 in the *Selfish* treatment *as if* they were reporting honestly (p > 0.1 in both cases, Binomial Test). Second, even though our subjects are drawn from an

⁸ The findings of Vanberg (2017) suggest those of Erat & Gneezy (2012) are not internally valid.

⁹ It's possible that, if more intelligent people are less sensitive to irrelevant details or framing, they should opt out less often irrespective of self-image concerns.

Table 4	
Marginal effect	of the Charity treatment on earnings.

	Model 1	2	3
Charity	0.655*** (0.156)	0.659*** (0.157)	0.66*** (0.162)
Observations	314	314	314

Note: The estimates are the marginal effect of the *Charity* treatment on earnings, estimated from Tobit regressions. 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. Model 1 controls for the date of the experiment. Model 2 additionally controls for income. Model 3 controls for ethnicity, gender and political views, as well as those of the previous models.

Table 5				
Treatment effects	on die	rolls -	Marginal	effects.

T-1-1- F

	Marginal effect of 'Charity Treatment' on each die roll					
	1	2	3	4	5	6
Model 1	-0.045***	-0.018	-0.056	-0.065	0.222***	-0.038*
	(0.017)	(0.027)	(0.038)	(0.052)	(0.055)	(0.023)
Model 2	-0.045***	-0.018	-0.056	-0.066	0.223***	-0.038
	(0.017)	(0.028)	(0.038)	(0.052)	(0.055)	(0.024)
Model 3	-0.047***	-0.018	-0.037	-0.075	0.217***	-0.041
	(0.017)	(0.028)	(0.04)	(0.053)	(0.057)	(0.025)

Note: The figures show the marginal effect of the *Charity* treatment on the probability that a particular die roll is reported, estimated from multinomial logit regressions. 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. 314 observations in each model.

entirely different population to those of Fischbacher and Föllmi-Heusi (2013), the reported percentages relating to the *Selfish* treatment in Table 3 closely replicates their findings.¹⁰

We now consider how behaviour compares across treatments. This is done in two ways. First, in order to examine average reported die rolls between treatment, we estimate three Tobit regressions and the corresponding marginal effects. In each regression, the dependent variable is earnings, where earnings corresponds to the subjects payoff in the *Selfish* treatment, and the payoff to the charity in the *Charity* treatment. 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 *Charity* 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. In all models we control for the date the experiment took place. This is done in order to control for possible differences in MTurker populations. Second, in order to examine individual reported die rolls, we estimate three multinomial logit (MNL) regressions, and estimate the subsequent marginal effects coefficients. In each regression, the reported die roll is the dependent variable. We use MNL regressions rather than ordered probits because it is not clear which order the reported die rolls should be placed. Although there might be a natural order in terms of extrinsic payoffs, this is not necessarily so in terms of utility. It's likely that intrinsic motivations play a significant role in which number is reported. We use the same set of controls in each of the MNL regressions as we do in the Tobit regressions.

Table 4 present the marginal effect of the *Charity* treatment dummy from the Tobit regressions. Table 5 presents the estimated marginal effect of the *Charity* treatment dummy on the probability of each die roll being reported.

Observation 2. Subjects report die rolls that result in significantly higher earnings in Charity treatment in comparison to the Selfish treatment. Subjects are 4% less likely to report a roll of 1, but more than 20% more likely to report a 5, when they receive the Charity treatment in comparison to the Selfish treatment.

¹⁰ It's also possible to calculate the number of unconditionally honest subjects following Fischbacher & Föllmi-Heusi (2013). This is done assuming that (1) those who report a 6 are not lying and (2) unconditionally honest people roll a uniform distribution of numbers. We are unable to determine the extent to which these assumptions hold. Nevertheless, these assumptions put the percentage of unconditionally honest subjects as 6*6.4%=38.4% in the *Selfish* treatment and 6*2.6%=15.6% in the *Charity* treatment.



Support. Table 4 outlines how the Charity treatment has a positive and significant marginal effect on earnings (p < 0.01, in all cases, T-tests), suggesting subjects report significantly higher die rolls on average. Table 5 shows that, across all model specifications, that the marginal effect of the Charity treatment on the probability of reporting a 1, 2, 3, 4 and 6 is negative, being between 1% and 5% less likely to be reported respectively (Model 3 columns 1, 2, 3, 4). However, the marginal effect of Charity is only significant for rolls of 1. In contrast subjects are approximately 20% more likely to report a 5 (p < 0.01, Model 3 Column 5, T-tests) when they receive the Charity treatment compared to the Selfish treatment.

Observation 2 provides strong evidence in support of Hypothesis 1, and highlights how subjects are more willing to misreport dies rolls when their misreporting benefits a third party. It is likely that, when there is a moral justification for lying, it is more prevalent. This is similar to previous findings of Erat and Gneezy (2012), who show that subjects lie more frequently when their lies increase the payoffs of others.

4.1.1. Intelligence

We now turn to examine 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*. We use this split only for the descriptive statistics, and focus on the impact of Raven score in the analysis.¹¹ 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.¹²

Fig. 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.

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).¹³

4.1.2. Intelligence and lying behaviour

To examine if there are any differences in reported die rolls formally, Table 6 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 three Tobit and three multinomial logit models in the same way as in Section 4.1. In the Tobit regressions, earnings is the dependent variable. In each regression we include the subjects Raven score as an explanatory variable, and

¹¹ 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.

¹² 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*.

¹³ The vast majority of subjects have multiple switching points in the lottery choice lists, which suggests indifference (Anderson & Mellor, 2009). This is regardless of Raven score. As a result, we do not use this measure in our regressions.

Table 6				
The percentage	of reported	die rolls b	y Raven	score.

_ . . .

	Report	Reported Die Roll					
Raven - Treatment	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 7Raven score effect on earnings - Marginal effects.

Selfish Treatment			Charity Treatment			
Model	1	2	3	1	2	3
Raven Score	0.006 (0.042)	0.01 (0.043)	0.02 (0.042)	0.006 (0.03)	0.01 (0.031)	0.016 (0.034)

Note: The figures show the marginal effect of Raven score on earnings, estimated from Tobit regressions. ***, ** and * indicate significance at the 1%, 5% and 10% levels. *p*-values calculated using standard t-tests. Model 1 includes no controls. Model 2 controls for income. Model 3 controls for age, social preferences, race, gender, political attitudes and education.

this is our variable of interest. We include this rather than a dummy for *High* because the cut off point between *High* and *Low* is subjective, and including multiple cut points is also subjective. In the multinomial logits, the reported die roll is the dependent variable, and we always include Raven score as an explanatory variable.

Regardless of regression type, in all models we include a dummy variable controlling for the date the experiment took place. In Model 1, we include no other controls. In Model 2, we add an additional control for subjects' level of income. In Model 3 we add all subsequent controls as well as age, social preferences, race, gender, political attitudes, and their level of education. We include these variables to control for differences between individuals that might otherwise explain our results. We estimate the effect of Raven score for the *Selfish* and *Charity* treatments separately, and as such do not interact any variables.

The marginal effect of Raven score on earnings is presented in Table 7. The marginal effects of Raven score on reported die rolls are given in Table 8, with the estimates from the *Selfish* treatment given in Panel A, and those from the *Charity* treatment in Panel B. In both Tables, we estimate the marginal effects using the observations from each treatment individually.¹⁴

Observation 3. A one point increase in Raven score reduces the probability of reporting a 1 by 3%. However, this is only true for the Selfish treatment, and has no significant impact on earnings.

Support. The estimates in Table 7 are never significant at conventional levels, highlighting how earnings do not differ conditional on Raven score. Table 8, Panel A, highlights how the marginal effect of Raven score on the probability that a roll of 1 is reported in the *Selfish* treatment is significant at the 5% level (p < 0.05, Wald Tests, across models). No other estimates in Panel A are significant at conventional levels. This suggests that subjects with a higher Raven score behave differently to those with a lower Raven score in the *Selfish* treatment. Further, the estimates in Table 8, Panel B highlight how the marginal effect of the Raven score 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.

Observation 3 highlights how subjects with a higher Raven score are more likely to report a die that results in a low payoff than those with a lower Raven score, when the decision impacts their themselves. This provides some indication that the costs and benefits associated with lying may depend on an individuals level of intelligence, or that the motivations driving honest or dishonest behaviour may differ between individuals conditional on their level of intelligence. However, it should be noted that Observation 3 provides only limited support in favour of Hypothesis 3, as we observe no significant

¹⁴ As not a single roll of 1 was reported in the *Charity* treatment, it is impossible to estimate the marginal effect of Raven score on this die roll.

Table 8The effect of intelligence on reported die rolls.

Marginal effect of Raven score on each die roll						
	1	2	3	4	5	6
Panel A						
Model 1	0.028**	0.003	0.001	-0.035	0.015	-0.013
	(0.013)	(0.013)	(0.018)	(0.024)	(0.024)	(0.014)
Model 2	0.031**	0.003	-0.001	-0.025	0.01	-0.017
	(0.014)	(0.013)	(0.018)	(0.023)	(0.024)	(0.015)
Model 3	0.036**	0.001	-0.004	-0.026	0.016	-0.023
	(0.016)	(0.013)	(0.018)	(0.023)	(0.023)	(0.015)
Panel B						
Model 1	-	0.012	-0.015	0.007	0.002	-0.006
	-	(0.011)	(0.016)	(0.021)	(0.024)	(0.009)
Model 2	-	0.012	-0.022	0.008	0.005	-0.003
	-	(0.011)	(0.016)	(0.022)	(0.025)	(0.009)
Model 3	-	0.006	-0.012	0.003	0.004	00
	-	(0.013)	(0.017)	(0.024)	(0.028)	(0.011)

Note: The figures show the marginal effect of Raven score on the proba- bility that a particular die roll is reported, estimated from multinomial logit regressions. A Raven score of 0 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. The first column in Panel B has no estimates because we do not observe a single report of 1 in the *Charity* treatment.



(a) Reported numbers - Selfish treatment (b) Reported numbers - Charity treatment

Fig. 2. Reported die rolls in the mind game.

differences for any other die rolls, no differences in the *Charity* treatment, and the difference in reporting behaviour has no impact on earnings.¹⁵,¹⁶

4.2. Experiment 2

To begin, we examine the distribution of numbers that subjects report as being generated by the random number generator used in the mind game. The distribution, if numbers are reported truthfully, should be indistinguishable from the uniform distribution. Fig. 2 presents this distribution graphically.

The distributions are both significantly different to the uniform distribution (p < 0.01 in both cases, Kolmogorov Smirnov tests). There are also no significant differences in the reported rolls between treatments (p = 0.25, Robust Rank Order Test). This suggests that some subjects are not truthfully reporting the number they generated, or that they did not generate a number at all.

We now examine the percentage of subjects that report that *Yes* the number they thought of in the mind game is identical to the number they generated. Regardless of how the number held in the subject's mind is produced, any one

¹⁵ It's important to note that, in Table 8, we conduct multiple hypotheses simultaneously, which may cause problems of multiplicity. In Appendix B, we correct the calculated *p*-values using the Holm-Bonferroni procedure in order to account for this. The analysis suggests that the significance of the coefficients in Table 8 cannot be distinguished from Type I error (false positive).

¹⁶ Again, under the assumptions used by Fischbacher & Föllmi-Heusi (2013), we can calculate the number of unconditionally honest people by intelligence, with 42.6% for *Low* and 28.8% for *High* Raven subjects in the *Selfish* treatment and 18.6% for *Low* and 10.8% for *High* Raven subjects in the *Charity* treatment.

Table 9			
Results from	the	mind	game

Report	Treatment Selfish	Charity
No	18.6% (31)	13.1% (20)
Yes	81.4% (136)	86.9% (133)

Note: The number of subjects is given in parentheses.

Table 10

Opt-out decisions, conditional on behaviour in the mind game.

Report	Selfish Treatment Do not opt–out	Opt—out	Report	Charity Treatment Do not opt—out	Opt–out
No	64.5% (20)	35.5% (11)	No	30% (6)	70% (14)
Yes	31.6% (43)	68.4% (93)	Yes	31.9% (36)	68.1% (97)

Note: The number of subjects is given in parentheses.

number they think of has a probability of just 16.7% of being generated. Therefore, reports of *Yes* above 16.7% suggests some subjects are lying. Table 9 presents the percentage of subjects reporting *Yes* and *No*, disaggregated by treatment.

As can be seen, a large percentage of subjects in both treatments report Yes - more than 80%. This is a very large percentage of subjects, especially in comparison to Kajackaite and Gneezy (2017) who report only 49% when payoffs were significantly higher (\$50). One can speculate as to why this is the case. First, it could be that our experiment was run online, whereas theirs was conducted in a laboratory. Another could be differences in subject pools, as Kajackaite and Gneezy (2017) use undergraduate students in the United States, whereas ours uses MTurkers.

Although we do not observe either the true generated number, or the number held in the subject's mind, this is indicative that a large proportion of subjects are lying. However, a non-negligible number of participants also report *No*. Although a slightly larger percentage of subjects report *Yes* in the *Charity* in comparison to the *Selfish* treatment, we find no significant differences in the behaviour from the mind game between treatments (p = 0.22, Fisher's Exact Test).

In order to shed light on the motivation for individuals' choices in the mind game, we now consider the percentage of subjects deciding to opt–out of implementing their decision. Table 10 presents the percentage of participants that chose to opt–out, disaggregated by treatment, and also by the report the subject made.

Observation 4. There are no differences in reporting in the mind game when comparing the Selfish and Charity treatments. However, opting-out of reporting No in the mind game differs between treatments.

Support. Comparing percentages in Table 9, there are no significant differences between the treatments (p > 0.1, Fisher's Exact test). However, as can be seen in Table 10, around 70% of subjects that reported Yes opt-out of implementing that decision, regardless of the treatment. For the Selfish treatment, the opt-out rate is significantly different between subjects who reported Yes and No (p < 0.01, Fisher's Exact Test). However, for the Charity treatment, there is no significant difference (p = 0.79, Fisher's Exact Test). Comparing opt- out rates between treatments, there are no significant differences when the subjects report a Yes (p = 0.43, Fisher's Exact Test), however, a significantly larger percentage opt-out of a No in the Charity treatment in comparison to the Selfish (p = 0.02, Fisher's Exact Test).

Observation 4 does not support Hypothesis 1: that higher payoffs will be more prevalent in the *Charity* treatment. That we observe no differences in behaviour between the treatments, but we do in Experiment 1, could be interpreted in a number of ways. One interpretation is that image concerns, and how the treatments impact them, is context dependent. A general limitation of the paper is that we cannot identify precisely what about context might drive these difference between experiments. A second interpretation is that the way in which the subject generates the number in their mind, rather than with a die roll, may affect any costs associated with lying through self-deception (Smith et al., 2017). A third interpretation is that there is a ceiling effect: as dishonesty is already very high, there is little scope for an effect to manifest.

The opt-out results can be interpreted in the following way. First, opting-out of implementing a report of *Yes* has no impact on a subject's financial payoff. However, if there is an intrinsic cost of implementing a *Yes*, potentially derived from lying, then opting-out of implementing avoids that cost: the utility from receiving a payoff of \$3.50 is greater than the utility from \$3.50 and lying. There are no differences between the *Selfish* and *Charity* treatments in this respect, which implies the intrinsic cost for those who reported a *Yes* is similar across treatments. In contrast, reporting a *No* implies a preference for honesty, and opting-out simply reveals that the utility from receiving \$3.50 is greater than the utility from receiving \$ 0 honestly. As subjects are more inclined to opt-out of reporting a *No* in the *Charity* treatment in comparison to the *Selfish* treatment, this implies that appearing honest is more important in the *Selfish* treatment than in the *Charity* treatment. Observation 4 lends some support to Hypothesis 2.





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

(b) Boxplot of Raven scores

Fig. 3. Raven scores.

ladie II	
Reporting in the mind game and intelligence.	

Report	Low Raven Treatment Selfish	Charity	Report	High Raven Treatment Selfish	Charity
No	23.4% (22)	20.5% (16)	No	12.3% (9)	5.3% (4)
Yes	76.6% (72)	79.5% (62)	Yes	87.7% (64)	94.7% (71)

Note: The number of subjects is given in parentheses.

4.2.1. Intelligence

We now examine the extent to which the participants' intelligence correlates with their choices and opt-out decisions. This is done to examine the hypothesis that intrinsic costs associated with behaviour, or self-image concerns, differ conditional on intelligence. Fig. 3 presents the distribution of Raven Scores, as well as a boxplot diagram of Raven scores. As can be seen, the median participant managed to obtain a Raven score of 1, with a mean average of around 2. The distribution is consistent with that from Experiment 1. We divide participants into *Low* and *High* based on their Raven score, exactly as in Experiment 1.

To determine if there exists differences in reporting between *Low* and *High* Raven subjects, we examine the reported numbers and subject choices in the mind game, disaggreged by treatment. We present this in Table 11.

In order to examine any potential differences parametrically, we conduct a number of Logit regressions, where the dependent variable is a dummy variable that takes a value of 0 if the subject reported *No*, and 1 if they reported *Yes*.¹⁷ In each regressions, we include the subject's Raven score as an explanatory variable, along with a dummy that takes a value of 1 if the treatment is *Charity* (and 0 otherwise). We conduct the regressions using the *Selfish* and *Charity* treatment data separately, following the analysis used for Experiment 1. In Model 1, we include no additional control variables. In Model 2, we add education and age. In Model 3, we further control for gender, social preferences and political affiliation. Using these three models we estimate the marginal effect of Raven score for each treatment. The estimates are given in Table 12.

Observation 5. . Subjects with a higher Raven score are more likely to report Yes in the mind game when assigned to the Charity treatment.

Support. As can be seen in Table 11, the proportion of subjects reporting Yes is larger for High Raven subjects than it is for Low subjects. This difference is significant at the 1% level in the Charity treatment (p = 0.001, Fisher's Exact Test), however is only weakly significant in the Selfish treatment (p = 0.08, Fisher's Exact Test).

Table 12 presents the estimated marginal effects of Raven score. As can be seen, Raven score has a positive and significant effect in the *Selfish* treatment, but only when no other controls are accounted for (p < 0.05, Model 1, T-test); Model 1 suggests a unit increase in Raven score increases the probability of reporting a Yes by 4.7%. In the *Charity* treatment, Raven score has a significant positive marginal effect on reports of Yes in all three models, increasing in significance as controls are added (p = 0.07, Model 1, p < 0.05 Model 2 and 3, T-tests). This suggests that a unit increase in Raven scores are increases the probability of a Yes by around 3-4% in the *Charity* treatment.

Observation 5 supports Hypothesis 3, but only partially, as more intelligence subjects do not report *Yes* less often in the *Selfish* treatment. It suggests that the image benefit from being charitable is larger for more intelligent individuals.

¹⁷ We use Logits, rather than Probits, to keep the analysis consistent across experiments. The results are near identical if the Probit specification is used.

Table 12

Marginal effect of Raven score on reporting Yes in the mind game.

	Model	1	2	3
Treatment	Selfish Charity	0.057** (0.02) 0.036*	0.003 (0.014) 0.043**	0.006 (0.013) 0.041**
	Observations	(0.021) 300	(0.019) 300	(0.021) 300

Note: Dependent variable is a dummy that takes a value of 1 if the subject reported *Yes* in the mind game, and 0 otherwise. Estimates are the marginal effect of Raven score, estimated from Logit regressions, by each treatment. ***, ** and * indicate significance at the 1%, 5% and 10% levels. Standard errors presented in parentheses. Model 1 includes no controls. Model 2 controls for education and age. Model 3 adds additional controls for gender, social preferences and political attitudes. Some observations are dropped due to missing entries.

Table 13

Opt-out decisions, conditional on behaviour in the mind game.

Panel A: L	ow Raven				
Report	Selfish Treatment		Report	Charity Treatmer	ıt
	Do not opt-out	Opt-out		Do not opt-out	Opt-out
No	72.7% (16)	27.3% (6)	No	25% (4)	75% (12)
Yes	30.1% (22)	69.9% (50)	Yes	17.7% (11)	82.3% (51)
Panel B:	High Raven				
Report	Selfish Treatment		Report	Charity Treatment	
	Do not opt-out	Opt–out		Do not opt-out	Opt—out
No	44.4% (4)	55.6% (5)	No	50% (2)	50% (2)
Yes	32.8% (21)	67.2% (43)	Yes	35.2% (25)	64.8% (46)

Note: The number of subjects is given in parentheses.

Table 14 Marginal effect of Raven score on opt-out rates.

	Selfish			Charity		
Model	1	2	3	1	2	3
Reported No Reported Yes	0.07 (0.044) -0.004 (0.028)	0.003 (0.064) -0.009 (0.029)	0.005 (0.065) -0.014 (-0.029)	-0.009 (0.062) -0.067*** (0.023)	-0.005 (0.064) -0.064*** (0.023)	-0.006 (0.063) -0.064*** (0.023)
Observations	162	148	148	153	152	152

Note: Dependent variable is a dummy that takes a value of 1 if the subject opted- out of implementing their decision in the mind game, and 0 otherwise. Estimates are the marginal effect of Raven score, estimated from Logit regressions, by each treatment, conditional on what they reported. ***, ** and * indicate significance at the 1%, 5% and 10% levels. Standard errors presented in parentheses. Model 1 includes no controls. Model 2 controls for education and age. Model 3 adds additional controls for race, gender and political attitudes. Observations are dropped due to missing entries.

To examine both misbehaviour and the motivations driving the reports of subjects, we now examine opt-out rates. These are given in Table 13. As can be seen, there are some dissimilarities between *Low* and *High* Raven subjects. We examine these differences formally using Logit regressions. In each case, the opt-out decision is the dependent variable, taking a value of 1 if *Opt-out* (and 0 otherwise). In each model, we include a variable for what was reported in the mind game, 1 if *Yes* (0 otherwise), as well as the subject's Raven score. In all regressions we also include the interaction of these two variables. In model 1, no additional controls are added. In model 2, we add income. In model 3, we add controls for gender and ethnicity. We estimate the marginal effect of Raven score first for the *Selfish* treatment, and then the *Charity* treatment, each conditional on the subject having reporting *Yes* and *No*. The estimates are given in Table 14

Observation 6. More intelligent subjects are less likely to opt-out of reporting a Yes in Charity treatment: a 1 point increase in Raven score reduces the probability of opting-out of Yes by 6%.



Fig. 4. Response times.

Support. Across all three model specifications in Table 14, the marginal effect of Raven score on opt-out rates is estimated to be negative and significant at the 1% level (p < 0.01 in all cases, t-tests), if the subject reported Yes. Its magnitude is also robust to specification changes, suggesting a 1 point increase in Raven score reduces the probability of opting-out by around 6%. In all other cases the estimated marginal effect is not significant at conventional levels (p > 0.1 in all cases, t-tests).

Observation 6 suggests there are differences in the intrinsic costs associated with be- haviour conditional on intelligence and that these costs interact with our treatment. This supports our Hypothesis 4, and implies that lying, when it results in helping a charity, produces a positive self-image and this is more important as intelligence increases.

5. 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 15 presents the average completion duration of subjects, disaggregated by Raven score. This is done using observations from both experiments.¹⁸

Table 15 outlines how subjects take a very similar amount of time to complete the experiment (around 10 min), and there is no significant difference between groups (p > 0.1, in both experiments, Robust Rank Order tests). There is also no significant correlation between duration of the experiment and the number of correct answers on the Raven Test ($\rho = -0.05$ Experiment 1, $\rho = 0.05$ Experiment 2, p > 0.1 in both cases, Correlation Coefficient).

A second possible interpretation is that the reported die roll in Experiment 1 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. For completeness, we also report response times for Experiment 2. Fig. 4 presents the response times for each reported die roll in Experiment 1, and the response time by report in Experiment 2.

As can be seen in Fig. 4a, 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.

In contrast, Fig. 4b highlights that there is a clear and significant difference in response times, conditional on reports in Experiment 2 (p < 0.001, Robust rank order test). The interpretation that those who understood the experimental instructions most (therefore taking longer) reported *Yes* in the mind game seems unreasonable, given that higher Raven scores are

¹⁸ Further analyses are included in Appendix B.

not associated with more Yes reports in the Selfish treatment. We take this as evidence that our reported impact of Raven score relates to intelligence, rather than understanding.

6. Conclusion

Our paper sheds light on the conceptual link between intelligence and decision making. We report evidence from two distinct experiments, and find, under one paradigm, only weak evidence that individuals with higher intelligence (or cognitive ability) behave differently to those with a lower cognitive ability when lying benefits themselves. The evidence we do find of behavioural differences has no impact on the earnings of subjects, and cannot be distinguished from Type I error. Under a second paradigm, we find that higher intelligence individuals increase the extent of their dishonesty when lying benefits a third party. In our case, the third party is a well known international wildlife protection charity.

Our findings are robust to controlling for a wide range of individual and behavioural measures that previous work has deemed important for decision making. Our results suggest that a link between intelligence and lying may be context dependent, as discussed in Abeler et al. (2014), and this resonates with a growing literature that examines how personality and intelligence impact on decision making, and opens a number of avenues for future research. We unfortunately are not able to determine what it is that may be driving the differences between experiments and this is a general limitation of the study.

The behavioural disparity in lying observed between 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. In an experiment in which only the subject is able to view their decisions, and using a novel opt-out decision at the end of one of our experiments, we are able to shed light on the intrinsic motivations behind subjects' choices. Our findings suggest that self-image concerns may vary with intelligence, and that this is an important driver of behaviour.

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 (Childs, 2012), reputation (Feess and Kerzenmacher, 2018), and social norms (Mann et al., 2014) may all be 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 the link between intelligence and lying in one particular environ- ment, we use two distinct paradigms. The 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 and List, 2007). Observations from our own study, and these future studies, suggest important questions for the theory of why people lie.

Declaration of Competing Interest

We declare that there is no relevant or material financial interests that relate to the research described in the paper titled "Is there a link between intelligence and lying?".

Data Availability

Data will be made available on request.

Appendix - For online publication

A. Experimental instructions

Introduction

Welcome

Thank you for participating in this HIT. It should take around 14 min 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

Table A1 Task A		
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

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.

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 s 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.

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.

Table A3 Task C		
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

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.

Part B - selfish treatment (Experiment 1). In Part B you will 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 (Experiment 1). 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.

Part B - selfish treatment (Experiment 2). In Part B you will receive a bonus. However, this bonus is not the same for every participant. You determine your own bonus by doing the following.

- 1 Think of a number from 1 to 6 and remember it.
- 2 Generate a random number from 1 to 6 using Google. You can do that by clicking here: https://www.google.com/search? q=random+number+1+to+6
 - You should generate a number only once. If the number you generated is the number you thought of your payment will be \$3.50. If the number you generated is not the number you thought of your payment will be \$0.

(new page)

Please enter the number you generated:

If the number you generated is the number you thought of, please click "Yes", and your bonus will be \$3.50. If the number you generated is not the number you thought of, please click "No", and your bonus will be \$0.

Part B - charity treatment (Experiment 2). 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 doing the following.

1 Think of a number from 1 to 6 and remember it.

2 Generate a random number from 1 to 6 using Google. You can do that by clicking here: https://www.google.com/search? q=random+number+1+to+6

You should generate a number only once. If the number you generated is the number you thought of your bonus will be \$3.50. If the number you generated is not the number you thought of your payment will be \$0.

(new page)

Please enter the number you generated:

If the number you generated is the number you thought of, please click "Yes", and your bonus will be \$3.50. If the number you generated is not the number you thought of, please click "No", and your bonus will be \$0.

Final question - selfish treatment (Experiment 2 only). In Part B you were asked to think of a number. If you reported that this number was the same as the number you generated, you will receive \$3.50. If not, you will received \$0.

Before carrying out the payment associated with your choice you are now given the opportunity to opt out of the decision from Part B. This means that whatever decision you made will not be implemented and you will instead receive \$3.50 at the end of the HIT.

If you decide not to opt out then you will receive the payment associated with the decision that was previously selected. To express your choice, tick the corresponding box below.

Final question - charity treatment (Experiment 2 only). In Part B you were asked to think of a number. If you reported that this number was the same as the number you generated, the charity will receive \$3.50. If not, they will received \$0.

Before carrying out the payment associated with your choice you are now given the opportunity to opt out of the decision from Part B. This means that whatever decision you made will not be implemented and instead the charity will receive \$3.50 at the end of the HIT.

If you decide not to opt out then the charity will receive the payment associated with the decision that was previously selected. To express your choice, tick the corresponding box below.

B. Appendix

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

B.1. Summary stats

Table A4		
Summary	of subjects'	characteristics

	Low Raven	High Raven	H_0 :Low =High
Incomet	6.645	5.864	0.013
Education*	Bachelor's Degree	Bachelor's Degree	0.154
Risk Preferences**	3.224	5.025	0.00
Race	White	White	0.027
Gender	0.688	0.61	0.163
Party***	Republican	Republican	0.047
Social Preferences****	3.168	3.314	0.017

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.

** Average switching point in the Holt and Laury choice lists. There are, however, a majority of subjects with multiple switching points which may suggest indifference (Anderson & Mellor, 2009) - regardless of Raven score. We report the 'earliest' switching point.

*** Party is which American political party they support. The modal average is reported.

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

B.2. The impact of intelligence

To check the robustness of the impact of intelligence on behaviour that we observe, we estimate the marginal effect of Raven score on behaviour, conditional on Raven score. For brevity, we focus only on those estimates in Section 3 that are found to be significantly different to zero. Table A5 Panel A presents the marginal effect of Raven score on the probability of reporting a roll of one in Experiment 1. These estimates are estimated from the multinomial logit regression in Table 8, model 1. Panel B presents the marginal effect of Raven score, on the probability of reporting a Yes in the *Charity* treatment of Experiment 2. These estimates are estimated from the multinomial logit regression in Table 12, model 1. Panel C presents the marginal effect of Raven score on the probability of opting out of reporting a Yes in the *Charity* treatment in Experiment 2. These estimates are estimated from the multinomial logit regression in Table 14, model 1 of the *Charity* treatment. All estimates are conditional on the subject's Raven score, and we evaluate the marginal effect at the bottom 25% of scores, the median score and the top 25% of all Raven scores.

As can be seen, in Panel A the marginal effect of Raven score increases as we evaluate its impact at greater Raven scores, and the effect at the Top 25% is significantly larger than at the Bottom 25% (p < 0.05, Wald Test). This means intelligence has a greater impact on behaviour as intelligence increases.

In Panel B, the marginal effects are all estimated to be positive but there are no significant differences between estimated effects, suggesting that intelligence has a constant marginal effect on behaviour (p > 0.1 in all cases, Wald Tests). Similarly, the estimates in Panel C are also found to be constant across Raven scores (p > 0.1 in all cases, Wald Tests) and all estimates are negative. This robustness check supports our results, and suggests that intelligence does indeed play a role: the difference in behaviour is estimated to be largest between the those who obtain a top 25% Raven score and those who end up in the bottom 25%.

B.3. Multiple hypothesis testing

To examine the robustness of our results from Experiment 1, we correct the p-values from the analysis in order to account for the possibility of multiplicity. Table **?** presents Panel A from Table **?** in the main text, with the p-values corrected using the Holm–Bonferroni procedure.

As can be seen, none of the coefficients are significant once corrected for the potential issue of multiple hypothesis testing.(Table A1, Table A2, Appendix A, Appendix B2, Appendix B3)(Table A3, Table A4, Table A5)(Table A6)

Table A5

The marginal effect of Raven score conditional on Raven score

Panel A	Raven Score Bottom 25%	Median	Top 25%	
Marginal Effect	0.008** (0.004)	0.016*** (0.006)	0.032** (0.0139)	
Panel B	Raven Score Bottom 25%	Median	Top 25%	
Marginal Effect	0.037 (0.023)	0.032* (0.019)	0.025** (0.01)	
Panel C	Raven Score Bottom 25%	Median	Top 25%	
Marginal Effect	-0.06*** (0.019)	-0.07*** (0.026)	-0.08** (0.032)	

Note: Marginal effect of Raven score. Standard errors in parentheses. ***,**,* denote significance at the 1%, 5% and 10% level. Panel A presents the marginal effect of Raven score, evaluated at different Raven scores, on the probability of reporting a roll of one in Experiment 1 Panel B presents the marginal effect of Raven score, evaluated at different Raven scores, on the probability of reporting a Yes in the *Charity* treatment of Experiment 2. Panel C presents presents the marginal effect of Raven score, evaluated at different Raven scores, on the probability of opting out of reporting a Yes in the *Charity* treatment in Experiment 2.

Table A6

The effect of intelligence on reported die rolls

	Marginal effect of Raven score on each die roll					
	1	2	3	4	5	6
Panel A						
Model 1	0.028	0.003	0.001	-0.035	0.015	-0.013
	(0.013)	(0.013)	(0.018)	(0.024)	(0.024)	(0.014)
Model 2	0.031	0.003	-0.001	-0.025	0.01	-0.017
	(0.014)	(0.013)	(0.018)	(0.023)	(0.024)	(0.015)
Model 3	0.036	0.001	-0.004	-0.026	0.016	-0.023
	(0.016)	(0.013)	(0.018)	(0.023)	(0.023)	(0.015)

Note: The figures show the marginal effect of Raven score on the proba- bility that a particular die roll is reported, estimated from multinomial logit regressions. A Raven score of 0 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. The first column in Panel B has no estimates because we do not observe a single report of 1 in the *Charity* treatment.

References

Abeler, J., Becker, A., Falk, A., 2014. Representative evidence on lying costs. Journal of Public Economics 113, 96-104.

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

Akerlof, G.A., Kranton, R.E., 2000. Economics and identity. The quarterly journal of economics 115 (3), 715-753.

Anderson, L.R., Mellor, J.M., 2009. Are risk preferences stable? comparing an experimental measure with a validated survey-based measure. Journal of Risk and Uncertainty 39 (2), 137-160.

Batson, C.D., Thompson, E.R., Seuferling, G., Whitney, H., Strongman, J.A., 1999. Moral hypocrisy: appearing moral to oneself without being so. J. Pers. Soc. Psychol. 77 (3), 525.

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.

Camerer, C.F., Ho, T.-H., Chong, J.-K., 2004. A cognitive hierarchy model of games. The Quarterly Journal of Economics 119 (3), 861-898.

Cattell, R.B., 1963. Theory of fluid and crystallized intelligence: A critical experiment. J. Educ. Psychol. 54 (1), 1.

Childs, J., 2012. Gender differences in lying. Economics Letters 114 (2), 147-149.

Croson, R., Gneezy, U., 2009. Gender differences in preferences. J. Econ. Lit. 47 (2), 448-474.

Demetriou, A., Kazi, S., 2006. Self-awareness in g (with processing efficiency and reasoning). Intelligence 34 (3), 297-317.

Demetriou, A., Kazi, S., Makris, N., Spanoudis, G., 2020. Cognitive ability, cognitive self-awareness, and school performance: From childhood to adolescence. Intelligence 79, 101432.

Diener, E., Wallbom, M., 1976. Effects of self-awareness on antinormative behavior. J. Res. Personal. 10 (1), 107-111.

Dohmen, T., Falk, A., Huffman, D., Sunde, U., 2018. On the relationship between cognitive ability and risk preference. J. Econ. Perspect. 32 (2), 115–134.

Dufwenberg, M., Dufwenberg, M.A., 2018. Lies in disguise-a theoretical analysis of cheating. J. Econom. Theory 175, 248-264.

Erat, S., Gneezy, U., 2012. White lies. Management Science 58 (4), 723-733.

Falk, A., 2021. Facing yourself-a note on self-image. Journal of Economic Behavior & Organization 186, 724-734.

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

Feess, E., Kerzenmacher, F., 2018. Lying opportunities and incentives to lie: Reference dependence versus reputation. Games and Economic Behavior 111, 274–288.

Fischbacher, U., Föllmi-Heusi, F., 2013. Lies in disguise—an experimental study on cheating. Journal of the European Economic Association 11 (3), 525–547. Freund, P.A., Kasten, N., 2012. How smart do you think you are? a meta-analysis on the validity of self-estimates of cognitive ability. Psychol. Bull. 138 (2), 296.

Froming, W.J., Walker, G.R., Lopyan, K.J., 1982. Public and private self-awareness: When personal attitudes conflict with societal expectations. J. Exp. Soc. Psychol. 18 (5), 476–487.

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

Jr, Gallup, G, G., 1998. Self-awareness and the evolution of social intelligence. Behav. Processes. 42 (2-3), 239-247.

Gerlach, P., Teodorescu, K., Hertwig, R., 2019. The truth about lies: A meta-analysis on dishonest behavior. Psychol. Bull. 145 (1), 1.

Gill, D., Prowse, V., 2016. Cognitive ability, character skills, and learning to play equilibrium: A level-k analysis. J. Polit. Econ. 124 (6), 1619–1676.

Gino, F., Ariely, D., 2012. The dark side of creativity: original thinkers can be more dishonest. J. Pers. Soc. Psychol. 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. Psychol. Sci. 20 (3), 393–398.

Gneezy, U., 2005. Deception: The role of consequences. Am. Econ. Rev. 95 (1), 384-394.

Gneezy, U., Kajackaite, A., Sobel, J., 2018. Lying aversion and the size of the lie. Am. Econ. Rev. 108 (2), 419-453.

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.

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.

Jiang, T., 2013. Cheating in mind games: The subtlety of rules matters, Journal of Economic Behavior & Organization 93, 328–336.

Kajackaite, A., Gneezy, U., 2017. Incentives and cheating. Games and Economic Behavior 102, 433–444.

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. J. Public Econ. 182, 104076.

Khalmetski, K., Sliwka, D., 2019. Disguising lies-image concerns and partial lying in cheating games. Am. Econ. J. Microecon. 11 (4), 79-110.

Leary, M.R., Tangney, J.P., 2011. Handbook of Self and Identity. Guilford Press.

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

Mann, H., Garcia-Rada, X., Houser, D., Ariely, D., 2014. Everybody else is doing it: exploring social transmission of lying behavior. PLoS One 9 (10), e109591. Mazar, N., Amir, O., Ariely, D., 2008. The dishonesty of honest people: a theory of self-concept maintenance. J. Mark. Res. 45 (6), 633–644.

Mazar, N., Ariely, D., 2006. Dishonesty in everyday life and its policy implications. J. Public Policy Mark. 25 (1), 117–126.

Pascual-Ezama, D., Fosgaard, T.R., Cardenas, J.C., Kujal, P., Veszteg, R., de Liaño, B.G.-G., Gunia, B., Weichselbaumer, D., Hilken, K., Antinyan, A., et al., 2015. Context- dependent cheating: experimental evidence from 16 countries. Journal of Economic Behavior & Organization 116, 379–386.

Paulhus, D.L., Dubois, P.J., 2015. The link between cognitive ability and scholastic cheating: A meta-analysis. Rev. Gen. Psychol. 19 (2), 183–190.

Potters, J., Stoop, J., 2016. Do cheaters in the lab also cheat in the field? European Economic Review 87, 26–33.

Proto, E., Rustichini, A., Sofianos, A., 2019. Intelligence, personality, and gains from cooperation in repeated interactions. J. Polit. Econ. 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.

Smith, M.K., Trivers, R., Von Hippel, W., 2017. Self-deception facilitates interpersonal persuasion. Journal of Economic Psychology 63, 93–101.

Speer, S.P.H., Smidts, A., Boksem, M.A.S., 2020. Cognitive control increases honesty in cheaters but cheating in those who are honest. Proc. Natl Acad. Sci. 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. J. Econ. Behav. Organ. 90, 19–27.

Vallacher, R.R., Solodky, M., 1979. Objective self-awareness, standards of evaluation, and moral behavior. J. Exp. Soc. Psychol. 15 (3), 254-262.

Vanberg, C., 2017. Who never tells a lie? Exp. Econ. 20 (2), 448-459.

Van't Veer, A., Stel, M., van Beest, I., 2014. Limited capacity to lie: cognitive load interferes with being dishonest. Judgm. Decis. Mak. 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. Pers. Individ. Differ. 81, 107–111.