

# CORRUPTED BY ALGORITHMS? HOW AI-GENERATED AND HUMAN-WRITTEN ADVICE SHAPE (DIS)HONESTY\*

*Margarita Leib, Nils Köbis, Rainer Michael Rilke, Marloes Hagens and Bernd Irlenbusch*

Artificial intelligence increasingly becomes an indispensable advisor. New ethical concerns arise if artificial intelligence persuades people to behave dishonestly. In an experiment, we study how artificial intelligence advice (generated by a natural language processing algorithm) affects (dis)honesty, compare it to equivalent human advice and test whether transparency about the advice source matters. We find that dishonesty-promoting advice increases dishonesty, whereas honesty-promoting advice does not increase honesty. This is the case for both artificial intelligence and human advice. Algorithmic transparency, a commonly proposed policy to mitigate artificial intelligence risks, does not affect behaviour. The findings mark the first steps towards managing artificial intelligence advice responsibly.

## 1. Introduction

Artificial intelligence (AI) shapes people's lives on a daily basis (Rahwan *et al.*, 2019). It sets prices in online markets (Calvano *et al.*, 2020), predicts crucial outcomes such as healthcare costs (Obermeyer *et al.*, 2019) and criminal sentences (Kleinberg *et al.*, 2018), and makes recommendations ranging from entertainment content and purchasing decisions to romantic partners (Yeomans *et al.*, 2019; Dellaert *et al.*, 2020). Increasingly, AI has become an indispensable advisor, thereby affecting people's behaviour (Fast and Schroeder, 2020; Kim and Duhachek, 2020). As a case in point, Amazon's chief scientist, Rohit Prasad, envisions that Alexa's role for its over 100 million users 'keeps growing from more of an assistant to an advisor' (Strong, 2020). Given AI's increasing role as an advisor, it is crucial to examine whether people are persuaded to follow or break ethical rules based on AI advice (Köbis *et al.*, 2021).

Large companies like LinkedIn and Zillow are already implementing AI advisors, thereby potentially shaping their employees' ethical behaviour. In such companies, natural language

\* Corresponding author: Nils Köbis, Max Planck Institute for Human Development, Center for Humans and Machines, Lentzeallee 94, 14195 Berlin, Germany. Email: [koebis@mpibberlin.mpg.de](mailto:koebis@mpibberlin.mpg.de)

This paper was received on 18 October 2021 and accepted on 2 August 2023. The Editor was Steffen Huck.

The data and codes for this paper are available on the Journal repository. They were checked for their ability to reproduce the results presented in the paper. The replication package for this paper is available at the following address: <https://doi.org/10.5281/zenodo.8170048>.

This research has been approved by the Ethics Commission of the Faculty of Management, Economics, and Social Sciences of the University of Cologne under reference 200010BI. We thank Clara Bersch, Yulia Litvinova, Ann-Kathrin Blanke, Toan Huynh, Anna Vogts and Matteo Tinè for research assistance, and Iyad Rahwan, Jean-Francois Bonnefon, Aljaz Ule, Anne-Marie Nussberger as well as the attendees of the Cognition, Values & Behaviour Research Group (Ludwig-Maximilians Universität München/LMU Munich), Moral AI lab meeting (Max Planck Institute for Human Development & Toulouse School of Economics), Applied Ethics & Morality Group (Prague University of Business and Economics), Centre for Decision Research (University of Leeds), Department of Economics and Management (University of Pisa), Decision Making and Economic Psychology Center (Ben-Gurion University), Behavioral and Management Science group (Technion), Colloquium of the Department of Social Psychology (Tilburg University) and seminar at the Department of Computer Science (Friedrich-Alexander University Erlangen-Nuremberg) for their helpful comments. The research was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy—EXC 2126/1—390838866 'ECONtribute: Markets and Public Policy', the European Research Council (ERC-StG-637915) and the Chamber of Commerce and Industry (IHK) Koblenz.

Margarita Leib and Nils Köbis are both shared first authors.

processing (NLP) algorithms (e.g., provided by software such as Gong.io) analyse employees' recorded sales calls and advise them on how to increase their sales. Without supervision, such algorithms may detect that deceiving customers pays off and thus advise salespeople to do so. Indeed, NLP algorithms can already autonomously detect deception as a useful strategy in a negotiation task (Lewis *et al.*, 2017). An ethical risk arises if people follow such corruptive AI advice. Here we examine (i) whether people meaningfully alter their (un)ethical behaviour following AI-generated advice and (ii) how such advice compares to human-written advice. Lastly, we test (iii) whether knowledge about the advice source (AI versus human) matters.

### 1.1. *Receiving Advice on (Un)Ethical Behaviour: Humans vs AI*

Generally, people are reluctant to take advice from others ('egocentric advice discounting'; see, e.g., Yaniv and Kleinberger, 2000), especially when it is unsolicited (Bonaccio and Dalal, 2006). However, when facing an ethical dilemma, advice has several compelling benefits for the advised. Advice encouraging an ethical course of action may validate one's moral preferences. It thereby might reduce negative feelings such as regret for not taking the opportunity to maximise profits by lying. Advice encouraging an unethical course of action may free people to violate ethical rules for profit without spoiling their moral self-image (Cross *et al.*, 2001). Indeed, taking advice can even provide a sense of shared responsibility with the advisor (Harvey and Fischer, 1997).

Compared to receiving human advice, how would people react to advice from an AI? Recent technological advances in the field of NLP reveal that AI text can already be indistinguishable from human text, suggesting that AI advice is as convincing as human advice. For instance, GoogleDuplex, an AI-based call assistant, can book appointments while having full-fledged conversations without the recipient even realising that an AI is on the line. Furthermore, AI can generate anything from poems (Köbis and Mossink, 2021) and Airbnb profiles (Jakesch *et al.*, 2019) to news articles (Kreps *et al.*, 2021) on par with humans. It thus stands to reason that, when people are not informed about the sources of advice, they will not recognise the advice source correctly and be affected by AI and human advice similarly.

### 1.2. *Testing Algorithmic Transparency*

To make sure that people know whom they interact with, governments, policymakers and researchers univocally call for algorithmic transparency (Jobin *et al.*, 2019)—the mandatory disclosure of AI presence (Diakopoulos, 2016). The recent Artificial Intelligence Act released by the EU demands AI systems such as chatbots and call assistants to disclose themselves as AI when interacting with humans (European Commission, 2021). Although it is a popular policy recommendation, empirical evidence for its effectiveness in shaping people's ethical behaviour is lacking.

How transparency about the advice source affects people's reactions to the advice is not trivial. Prior work informs three competing possibilities. The first possibility is that, when informed about the source of advice, people follow human advice *more* than AI advice. This account rests on the literature on algorithm aversion (Dietvorst *et al.*, 2015). People readily rely on AI in objective and technical domains (e.g., numeric estimation, data analysis and giving directions; Castelo *et al.*, 2019; Logg *et al.*, 2019). However, they are reluctant to use AI for subjective decisions, especially with ethical implications (e.g., parole sentences, trolley-type dilemmas; Bigman and

Gray, 2018; Castelo *et al.*, 2019; Laakasuo *et al.*, 2021). Furthermore, people follow perceived social norms when making (un)ethical decisions (Bowles, 2016; Gächter and Schulz, 2016; Fehr, 2018; Köbis *et al.*, 2019; Leib, 2023). Compared to AI advice, human advice might be a stronger signal of social norms because social norms regulate and emerge from *human* (not AI) behaviour. Consequently, people should be more likely to follow human advice. Suppose that people indeed prefer human input in ethically charged settings and perceive human advice as a stronger cue for social norms. In that case, we should expect that *human advice sways people's (un)ethical behaviour more than AI advice*.

The second possibility is that, when informed about the source of advice, people follow advice from humans *less* than advice from AI. A closer look at the technical design of AI advice systems would support this account. NLP algorithms are trained on a large corpus of human-written texts (Radford *et al.*, 2019). When people know that NLP algorithms draw on large compiled human input, they might perceive AI advice as a better representation of *most* people's beliefs and behaviours than the advice they receive from one human. If AI advice is indeed a stronger cue for social norms than a single piece of human-written advice, we should expect that *AI advice sways people's (un)ethical behaviour more than human advice*.

The third possibility is that, when people receive information about the source of advice, they are affected *equally* by human and AI advice. Support for this account comes from the observation that people already seek advice from AI agents. For instance, millions of people turn to Replika, the 'AI companion who cares. Always here to listen and talk. Always on your side' (replika.ai) for virtual companionship, socialising and also for advice (Murphy, 2019). Such AI advisors might also help justify questionable behaviour. When tempted to break ethical rules for profit, people do so as long as they can justify their actions (Fischbacher and Föllmi-Heusi, 2013; Barkan *et al.*, 2015; Shalvi *et al.*, 2015). Receiving advice that encourages rule breaking can serve as a welcome justification, possibly even when the advice stems from AI. Indeed, people deflect blame and share the responsibility for harmful outcomes, not only with other people (Tenbrunsel and Messick, 2004; Bartling and Fischbacher, 2011; Bazerman and Gino, 2012), but also with AI systems (Hohenstein and Jung, 2020). If following AI and human advice is equally justifiable and leads to similar attribution of responsibility between the two, we should expect that *human and AI advice sway people's (un)ethical behaviour to the same extent*.

### 1.3. *The Current Study*

The current study tests how advice type (honesty versus dishonesty promoting), advice source (AI versus human) and information about the advice source (transparency versus opacity) shape humans' (un)ethical behaviour. Until recently, most work on algorithmic advice has examined people's *stated preferences* (for an exception, see Greiner *et al.*, 2022) about *hypothetical scenarios* describing AI advice (Bigman and Gray, 2018; Castelo *et al.*, 2019; Logg *et al.*, 2019; Kim and Duhachek, 2020). We supplement such work by adopting a machine behaviour approach (Rahwan *et al.*, 2019) and examine people's *behavioural reactions* to *actual AI-generated output*.

To measure people's (un)ethical behaviour, we use the well-established incentivised die-rolling task (Fischbacher and Föllmi-Heusi, 2013). In it, participants roll a die in private and report the outcome, with higher outcomes corresponding to higher pay (for similar approaches, see Sutter, 2008; Abeler *et al.*, 2019; Fehrl *et al.*, 2020). To generate real AI advice, we employ the

state-of-the-art algorithm GPT-J. We fine-tune the algorithm using minimal training to produce relevant advice texts. Using this ‘few-shot’ learning approach mimics many real-world settings where language models (e.g., Alexa) are not optimised to maximise users’ profits or provide the most convincing advice, but simply produce relevant information based on prompts (Brown *et al.*, 2020).

## 2. Method

We conducted a large-scale, pre-registered experiment. The design entailed two parts. In part 1, we collected human-written advice and generated AI advice; in part 2, we examined the effect of advice on (dis)honest behaviour (see Figure 1 for an overview). By using real AI text outputs we can gain insights into how people react to actual AI advice. We report all measures, manipulations and data exclusions in the main text and [Online Appendix](#). The IRB board of our department approved the experiment, and all materials, pre-registrations and data are available on the [Open Science Framework](#).

### 2.1. Part 1: Advice-Giving Task

#### 2.1.1. Human-written advice

First, we conducted an *advice-giving task* in which we recruited advisors ( $N = 367$ ,  $M_{\text{age}} = 30.28$ ,  $SD_{\text{age}} = 9.56$ , 45.50% females) via Prolific.co (pre-registration <https://osf.io/nbke2>). The task took 10–15 minutes, and participants earned a base pay of £1.27. Advisors learned that a separate group of participants (advisees) would engage in a die-rolling task (Fischbacher and Föllmi-Heusi, 2013), in which they roll a die privately and report the outcome (with higher outcomes corresponding to higher pay). Advisors were asked to write advice for these advisees.

Advisors were randomly assigned to either write honesty-promoting or dishonesty-promoting advice. We informed advisors in the *honesty-promoting* advice treatment that if an advisee read their advice and subsequently reported the actual die-roll outcome (i.e., acted honestly), they would earn a bonus of £3. Advisors in the *dishonesty-promoting* advice treatment learned that if an advisee read their advice and subsequently reported the highest possible outcome, 6, they would earn a bonus of £3.<sup>1</sup>

Advisors had to follow pre-specified advice writing rules to ensure that they produced coherent advice texts that could be used to train GPT-J. Specifically, their advice had to (i) entail at least 50 words, (ii) not use concrete numbers in numeric or written form,<sup>2</sup> (iii) be in English and in their own words, (iv) be written in complete sentences, (v) be about the advisee’s die-roll outcome reporting decision and (vi) not inform the advisee that the advisor’s payoff depended on their behaviour.<sup>3</sup>

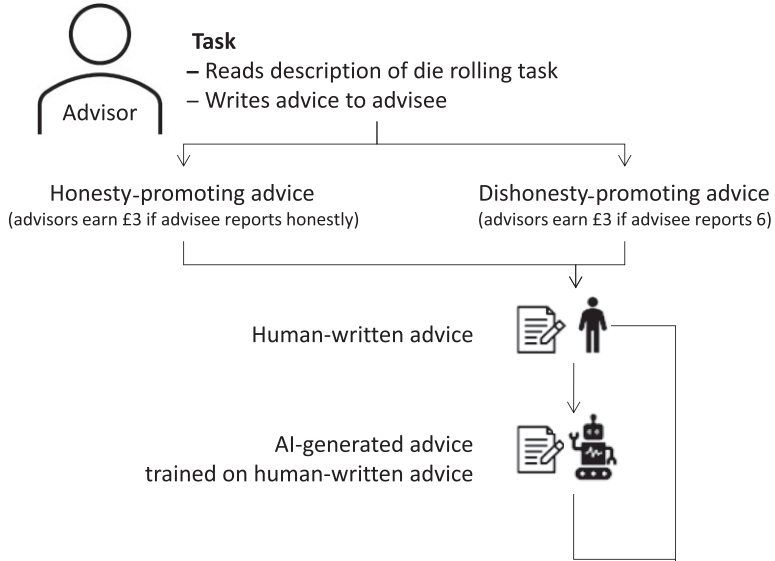
To incentivise advisors to follow the advice writing rules, they stood to gain a bonus. Namely, out of all advice texts, we randomly selected one, and if that text followed the writing rules, the advisor earned a bonus of £10. Moreover, as incentivisation for writing convincing texts, 1% of

<sup>1</sup> If advisees follow the advice in the dishonesty-promoting treatments, they will lie in the majority of the cases (five out of six cases). When the actual die-roll outcome is 6, following the advice does not entail lying.

<sup>2</sup> Advisors were not allowed to use concrete numbers to allow generating high-quality AI advice. GPT-J is trained to predict the next word in a sentence (see Section 2.1.2). If advisors were allowed to concretely mention numbers, training GPT-J on the human-written advice could have resulted in random numbers appearing out of context in the GPT-J output, reducing the quality of AI-generated advice.

<sup>3</sup> Advisors were not allowed to mention their incentive structure to the advisees so that we could keep the pro-social motivation for advisees who read AI and human advice constant (at zero).

(a) Part 1: advice-giving task



(b) Part 2: advice-taking task

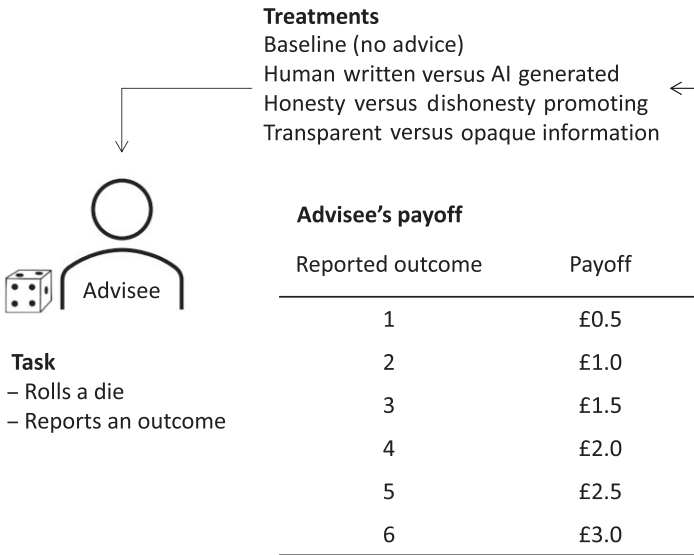


Fig. 1. (a) Part 1: Advice-Giving Task. (b) Part 2: Advice-Taking Task.

Notes: (a) Participants were incentivised to write honesty- or dishonesty-promoting advice texts, which were then used to generate AI advice. (b) Another group of participants engaged in the die-rolling task.

Advisees read advice, then reported a die-roll outcome. In total, we administered nine treatments: participants read honesty- or dishonesty-promoting advice that was human written or AI generated; participants were either informed about the source of advice (transparency) or not (opacity); as a baseline, another group of participants did not read any advice.

advice texts (4 out of 400) were implemented. If advisees acted according to the implemented advice, the respective advisor earned a bonus based on the treatment they were in (*honesty-* versus *dishonesty-promoting* advice).<sup>4</sup>

### 2.1.2. AI-generated advice

To generate AI advice (see Figure 1(a)), we employed GPT-J,<sup>5</sup> an open-source NLP algorithm published by Eleuther AI (<https://www.eleuther.ai/>). GPT-J is trained on a curated and diverse data set of 825 GiB texts to predict the next word in a sequence of words and contains six billion parameters (Wang and Komatsuzaki, 2021). GPT-J can be fine-tuned with extra training to produce a specific type of text. We fine-tuned GPT-J with ‘few shot’ learning by separately training it on the human-written *honesty-promoting* and *dishonesty-promoting* advice from the *advice-giving* task. We only used advice texts that adhered to the advice writing rules (as coded by a naïve coder) for fine-tuning. More details on the calibration of GPT-J are reported in the [Online Appendix](#).

### 2.1.3. Screening

After collecting human advice and generating AI advice, we employed the same pre-specified screening procedure for both advice sources (see Figure 2). First, we excluded texts that exceeded 100 words. Next, to ensure that advisees read coherent and relevant advice texts, we randomly selected 100 advice texts per cell. Two independent coders, who were naïve to the experimental treatments, coded each piece of advice on the following criteria: (a) is the text coherent? (Y/N); (b) does the text contain clear advice? (Y/N); (c) which type of behaviour does the advice encourage? (*honesty/dishonesty/unclear*); (d) does the advice follow the advice writing rules? (Y/N). Furthermore, we used the objective grammarly and readability scores as computational proxies for the quality of the texts (Grammarly, 2022).<sup>6</sup>

Among the texts that passed the coding procedure<sup>7</sup> and received a grammarly score equal to or above 50, we randomly selected 20 advice texts per treatment (*AI generated* versus *human written*, by *honesty* versus *dishonesty promoting*), yielding a final sample of 80 advice texts used in part 2 (see all advice texts in the [Online Appendix](#)). By applying the same screening procedure for human and AI advice, we ensure that the advice texts fulfil minimal quality criteria and are as comparable as possible.

<sup>4</sup> Paying advisors required knowing whether participants, after reading the advice, reported the observed die roll honestly or not. To do so, we ran a modified version of the die-rolling task in which advisees received randomly selected advice, saw a die-roll on the computer screen and were asked to report it. We implemented this procedure for four randomly selected advice texts (1% of the advice) and four advisees. This non-private procedure provided certainty about whether an advisee reported honestly or not and enabled us to pay advisors accordingly. Doing so meant that our experimental setup was incentivised and did not entail experimental deception. In the main experiment, the die-roll outcomes were private (see Section 2.2).

<sup>5</sup> As one can read in our pre-registration, we originally planned on deploying GPT-2 (see <https://osf.io/nbke2>) to generate AI advice. However, we opted to use GPT-J instead because it is open source, which increases reproducibility and is more advanced as it is a much larger and more potent model than GPT-2.

<sup>6</sup> We obtained grammarly and readability scores from grammarly.com. The grammarly score compares texts to all other texts checked on the platform. A score of 80 indicates that a text scores better than 80% of all texts checked on grammarly.com in terms of grammatical correctness. The readability score employs the Flesch-Reading-ease test and represents how easy a text is to read. The score is calculated by the average sentence length and the average number of syllables per word, with higher scores indicating easier readability.

<sup>7</sup> Texts that passed the coding procedure (i) are coherent, (ii) contain clear advice, (iii) encourage honesty in the *honesty-promoting* treatment and *dishonesty* in the *dishonesty-promoting* treatment, and (iv) follow the advice writing rules. Moreover, the coding by both independent coders had to match each other in order for the text to pass.

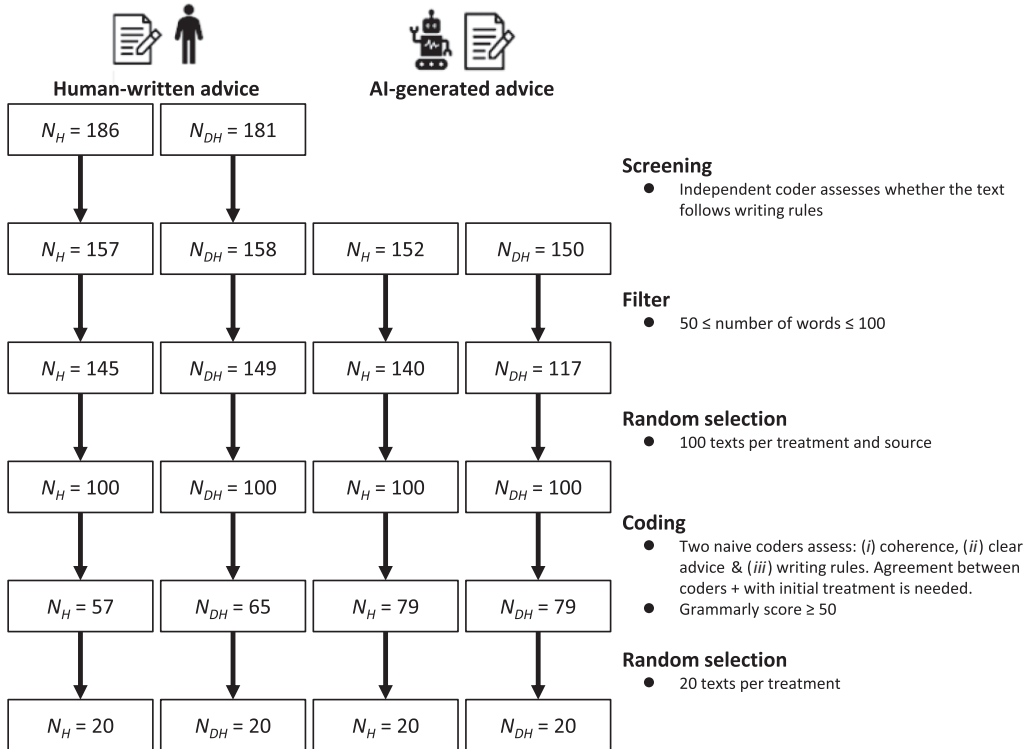


Fig. 2. Overview of the Selection Procedure of Advice Texts.  
Notes: H, honesty-promoting advice; DH, dishonesty-promoting advice.

## 2.2. Part 2: Advice-Taking Task

The advice-taking task took about eight minutes to complete, and participants earned a fixed pay of £1.20. We pre-registered (<https://osf.io/nqvfv3>) to collect a sample size that would allow us to detect a small-to-medium effect size (200 participants per cell, 1,800 in total) via Prolific.co to take part in the *advice-taking task*. Overall, 1,817 ( $M_{\text{age}} = 32.39$ ;  $SD_{\text{age}} = 11.68$ , 48.73% females) participants were included in the analyses. These participants completed the task and self-report items and passed the comprehension and attention checks (see below). Sensitivity analysis for a regression with 90% power and a significance level of 0.05 revealed that our sample was sufficient to detect small effect sizes ( $f^2 = 0.006$  and 0.010; see [Online Appendix](#) for details).

Participants read the instructions, received advice and finally engaged in the die-rolling task. Specifically, participants were asked to roll a die privately and report the outcome. Participants' pay corresponded to their report: for reporting a '1', they earned £0.5; a '2' = £1; '3' = £1.5; '4' = £2; '5' = £2.5; '6' = £3. After reading the instructions and before engaging in the die-rolling task, all participants learned that 10% of participants would be randomly selected and paid for the die-rolling task.

Assessing dishonesty by employing the die-rolling task is a common practice in economics and psychology (for meta-analyses, see Abeler *et al.*, 2019; Gerlach *et al.*, 2019; Köbis *et al.*, 2019; Leib *et al.*, 2021). Furthermore, the task has good external validity, as lying in the die-rolling task

correlates with unethical behaviour outside the lab, such as free-riding public transportation (Dai *et al.*, 2018) and being absent from work without reason (Hanna and Wang, 2017). Importantly, the die-rolling task pits two competing motivations against each other: to be a moral, honest person versus maximising financial profits, thus rendering advice particularly valuable.

After reading the instructions, participants had to answer three comprehension questions correctly and pass an attention check. If participants answered a comprehension question incorrectly, they could try again. If they answered the same question incorrectly twice, or if they answered the attention check incorrectly, they were disqualified from participating in the study and were not paid.

### 2.2.1. Treatments

Before reporting the die-roll outcome, participants were randomly assigned to one of nine different treatments. Participants either (i) read *honesty-promoting* or *dishonesty-promoting* advice that was (ii) either *human written* or *AI generated*. Furthermore, (iii) participants were either informed about the advice source (*transparency* treatments) or not (*opacity* treatments). Additionally, in a *baseline* treatment, participants did not receive any advice. Thus, the overall design was a 2 (advice source: human versus AI) by 2 (advice type: honesty promoting versus dishonesty promoting) by 2 (information about the advice source: transparency versus opacity) + 1 (baseline, no advice) between-subject design.

Participants who read AI-generated advice and were informed about it (in the *transparency* treatments) read:

‘To help you with your decision, an artificially intelligent (AI) algorithm generated an advice.

Information about AI-advice:

The algorithm is trained on large datasets of English text (> 820 GiB, which is more than 500 million pages of text) and produces text by predicting the next word in a sentence (similar to predictive text on smartphones). On top of that, it is also possible to train the algorithm on additional specific pieces of text. To generate multiple AI advice texts, the algorithm was trained on advice texts collected from other participants in the study. These participants did not take part in the die rolling task and were only instructed to write advice regarding the decision in the die rolling task. The advice you will read is one advice text that was generated by the algorithm.’

Participants who read human-written advice and were informed about it (in the *transparency* treatments) read:

‘To help you with your decision, another participant wrote an advice.

Information about advice:

To collect multiple advice texts, another group of participants was asked to write advice regarding the decision in the die rolling task. These participants did not take part in the die rolling task and were only instructed to write advice regarding the decision in the die rolling task. The advice you will read is advice written by one participant.’

Participants who were in the *opacity* treatments and thus not informed about the advice source read:

‘To help you with your decision, you will read an advice.

This advice has been written either by another participant or by an artificially intelligent (AI) algorithm. There is a 50% chance the advice is written by a participant and a 50% chance it is written by an algorithm.’

In the *opacity* treatments, this text was followed by the same two descriptions of how advice text from each source was collected or generated as in the *transparency* treatments. In the *opacity*



treatment, this information about AI advice generation and human advice collection appeared in random order.<sup>8</sup>

### 2.2.2. A static Turing test

After completing the die-rolling task, participants in the *opacity* treatment engaged in an incentivised version of a static Turing test (Köbis and Mossink, 2021). In contrast to the classical Turing test (Turing, 1950), participants did not interact back and forth with the source of advice. Instead, they read the advice text again and indicated whether they thought a human or an AI had written it. Participants learned that 20 of them would be randomly selected, and if their guess in the static Turing test was correct, they would earn an additional £1.

### 2.2.3. Potential mechanisms

Finally, to explore possible mechanisms, participants completed a post-experimental survey. Participants indicated on a scale from 0 to 100 their perceived (i) appropriateness (injunctive social norm), (ii) prevalence (descriptive social norm) and (iii) justifiability of reporting a higher die roll than the one observed. Additionally, all participants, except those who did not receive any advice, rated how they attribute responsibility between themselves and the advisor for the reported outcome in the die-rolling task. The answer scale ranged from 0 (= I am fully responsible) over 50 (= the advisor and I share responsibility equally) to 100 (= the advisor is fully responsible). Participants further indicated (on a scale from 0 to 100) to what extent they feel guilty after completing the task (see the [Online Appendix](#) for results regarding guilt and wording of all items). Finally, all participants indicated their age and gender.

## 3. Results

In all nine treatments, participants lied, as the average die-roll outcomes significantly exceeded the expected average if participants were honest ( $EV = 3.5$ ), one-sample *t*-test,  $t_s > 3.43$ ,  $ps < 0.001$ .

*Is people's (un)ethical behaviour influenced by AI-generated advice?* Yes, when it comes to dishonesty-promoting advice; no, when it comes to honesty-promoting advice. We first focus on the *opacity* treatments, where participants are not informed about the advice source. Here, linear regression analyses reveal that the average die-roll reports following *AI-generated dishonesty-promoting* advice ( $M = 4.60$ ,  $SD = 1.37$ ) significantly exceed reports in the *baseline*, no advice treatment ( $M = 3.99$ ,  $SD = 1.56$ ,  $b = 0.610$ ;  $p < 0.001$ ; 95% CI = [0.325, 0.894]). However, die-roll reports following *AI-generated honesty-promoting* advice ( $M = 4.01$ ,  $SD = 1.63$ ) do not significantly differ from reports in the *baseline* treatment ( $b = 0.019$ ;  $p = 0.898$ ; 95% CI = [-0.276, 0.314]); see Figure 3 and Table 1 (model 1). Furthermore, die-roll reports in the *AI-generated dishonesty-promoting* treatment significantly exceed those in the *AI-generated honesty-promoting* advice treatment ( $b = -0.590$ ,  $p < 0.001$ ; 95% CI = [-0.882, -0.299]). Thus, while dishonesty-promoting AI advice successfully corrupts people, honesty-promoting AI advice fails to sway people toward honesty.

*How does AI-generated advice square compared to human-written advice?* AI-generated advice affects behaviour similarly to human-written advice, for both honesty-promoting and

<sup>8</sup> To control for participants' beliefs about the potential advice sources, we opted to inform them that there is a 50–50 chance that a human or AI wrote the advice. We believed that not providing any information about the advice source would reasonably lead participants to assume that the advice source is another human, as AI might not be a salient source of advice for participants.




Fig. 3. Mean Reported Die-Roll Outcomes (Bars) and Proportion of Reported Sixes (Open Circles) across Advice Type (Honesty versus Dishonesty Promoting), Source (AI versus Human) and Information Treatments (Opacity versus Transparency).

Notes: The dashed black line represents the expected mean if participants were honest ( $EV = 3.5$ ), and the dashed white line represents the expected proportion of sixes if participants were honest (16.67%). The means (SDs) of die-roll reports are given at the bottom of each bar; \*\*\* $p < 0.001$ ; ns:  $p > 0.05$ .

dishonesty-promoting advice. Focusing on the *opacity* treatments, the two-way interaction (advice source by advice type) is not significant ( $b = 0.070$ ,  $p = 0.744$ ; 95% CI =  $[-0.350, 0.490]$ ); see Figure 3 and Table 1 (model 2). Specifically, the average die-roll reports do not differ between the *AI-generated* ( $M = 4.01$ ,  $SD = 1.63$ ) and *human-written* advice when advice was *honesty promoting* ( $M = 3.93$ ,  $SD = 1.52$ ,  $b = -0.076$ ,  $p = 0.631$ ; 95% CI =  $[-0.388, 0.236]$ ). Similarly, average die-roll reports do not differ between the *AI-generated* ( $M = 4.60$ ,  $SD = 1.37$ ) and *human-written* advice when advice was *dishonesty promoting* ( $M = 4.59$ ,  $SD = 1.54$ ,  $b = -0.006$ ,  $p = 0.965$ ; 95% CI =  $[-0.289, 0.276]$ ).

In addition, the results of the static version of the Turing test indicate that individuals cannot distinguish AI-generated advice from human-written advice. Specifically, in the *opacity* treatments, 49.94% (401 out of 803) of participants guessed the source of advice correctly, which does not differ from chance levels (50%, binomial test:  $p = 0.999$ ; 95% CI =  $[0.464, 0.535]$ ).

*Does transparency about the advice source matter?* No, informing participants about the algorithmic or human source of advice does not change their behaviour. Linear regression analyses reveal that the three-way interaction (advice type by source by information) is not significant ( $b = 0.101$ ,  $p = 0.735$ ; 95% CI =  $[-0.482, 0.683]$ ); see Figure 3 and Table 1 (model 3). Both among the *opacity* and *transparency* treatments, the two-way interactions (advice source

Table 1. *Regression Analyses on the Average Die-Roll Reports, Including Control Variables and Interactions.*

Independent variables	Dependent variable: reported die-roll outcome						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
No advice	0.019 (0.150)						
Dishonesty-promoting advice	0.610*** (0.145)	0.590*** (0.148)	0.590*** (0.145)	0.396** (0.143)	0.439** (0.144)	0.436** (0.145)	0.370* (0.152)
Human-written advice		-0.076 (0.152)	-0.076 (0.150)	-0.166 (0.147)	-0.127 (0.147)	-0.024 (0.157)	0.086 (0.171)
Transparency treatment			0.067 (0.150)	0.071 (0.147)	-0.112 (0.148)	0.114 (0.148)	
<i>Interactions</i>							
Dishonesty-promoting advice × human advice		0.070 (0.214)	0.070 (0.210)	0.104 (0.205)	0.033 (0.206)	-0.046 (0.210)	0.041 (0.219)
Dishonesty-promoting advice × transparency treatment			-0.046 (0.209)	-0.031 (0.203)	-0.088 (0.204)	-0.067 (0.204)	
Human advice × transparency treatment			-0.120 (0.211)	-0.094 (0.205)	-0.146 (0.206)	-0.162 (0.206)	
Dishonesty-promoting advice × human advice × transparency treatment			0.101 (0.297)	0.083 (0.290)	0.170 (0.290)	0.161 (0.290)	
<i>Additional controls</i>							
Injunctive norms				0.002 (0.002)	0.001 (0.002)	0.002 (0.002)	0.003 (0.002)
Descriptive norms				0.006*** (0.002)	0.005** (0.002)	0.005** (0.002)	0.005* (0.003)
Justifiability				0.008*** (0.002)	0.007*** (0.002)	0.007*** (0.002)	0.006* (0.002)
Shared responsibility				0.000 (0.001)	0.000 (0.001)	0.000 (0.001)	0.001 (0.002)
Gender (male)					0.189** (0.073)	0.191** (0.073)	0.203+ (0.105)
Age					-0.008* (0.003)	-0.008* (0.003)	-0.005 (0.005)
Grammarly score						0.008* (0.004)	0.014* (0.006)
Readability score						-0.006 (0.004)	0.004 (0.006)
1 if source guessed correctly							0.163 (0.106)
Intercept	3.986***	4.005***	4.005***	3.350***	3.571***	3.362***	1.838*
R <sup>2</sup>	0.035	0.042	0.044	0.096	0.101	0.105	0.105
N	634	803	1,604	1,604	1,589	1,589	794
Data used for analysis	Opacity, AI advice, and 'no advice'	Opacity treatments	All treatments without 'no advice' treatment	All treatments without 'no advice' treatment	All treatments without 'no advice' treatment	All treatments without 'no advice' treatment	Opacity without 'no advice' treatment

Notes: Models 5–7 contain a smaller N, as some participants did not report their gender as male/female. +  $p < 0.10$ , \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ . Coefficients that are larger than zero, but for which rounding turns into zero, are presented as 0.000.

by advice type) are not significant (for *transparency*,  $b = 0.170$ ,  $p = 0.409$ , 95% CI =  $[-0.235, 0.575]$ ; for *opacity*,  $b = 0.070$ ,  $p = 0.744$ , 95% CI =  $[-0.350, 0.489]$ ).

Overall, the popular policy recommendation of algorithmic transparency does not alleviate the corrupting effect of AI advice. Namely, die-roll reports following *AI-generated dishonesty-promoting* advice under the *opacity* treatment ( $M = 4.60$ ,  $SD = 1.37$ ) are on par with reports following the same advice in the *transparency* treatment ( $M = 4.62$ ,  $SD = 1.40$ ,  $b = 0.021$ ,  $p = 0.879$ ; 95% CI =  $[-0.245, 0.286]$ ). Specifically, when participants are *not informed* about the advice source, they boost their reports by 15.3% following *AI-generated dishonesty-promoting* advice, compared to the *baseline*  $[(4.60 - 3.99)/3.99 = 0.153]$ , which is equivalent to the 15.8% increase when they *are informed* about the source of the advice  $[(4.62 - 3.99)/3.99 = 0.158]$ . Bayesian analyses corroborate these conclusions (see the [Online Appendix](#)). Overall, results align with the idea that people increasingly follow AI advice (e.g., Replika) and use AI-generated advice to justify breaking ethical rules for profit.

### 3.1. Robustness of the Obtained Results

In our experimental design, advisors in the *dishonesty-promoting* treatment received £3 only if advisees reported the highest value, ‘6’. Such an incentive scheme is comparable with the *honesty-promoting* treatment in which advisors earned £3 only if advisees reported honestly. In both cases, advisors earn money for one out of six potential advisee’s reports (i.e., when the advisee reports a ‘6’ or honestly, depending on the treatment) and do not earn money in the remaining five advisee reports. However, advisors’ incentive scheme in the *dishonesty-promoting* treatments may have resulted in advice texts that predominantly focused on convincing participants to report the outcome 6. To assess the robustness of our results, we (i) conducted additional analyses and (ii) ran additional treatments.

#### 3.1.1. Proportion of sixes

First, as an additional analysis, we examined whether the proportion of sixes, as an alternative outcome variable, led to the same conclusions. We found very similar results (see Figure 3, where the open circles represent the proportion of sixes across all treatments). Specifically, focusing on the *opacity* treatments, linear regression analyses reveal that the proportion of sixes following *AI-generated dishonesty-promoting* advice (32.44%) significantly exceeds the proportion of sixes in the *baseline*, no advice treatment (20.66%;  $b = 0.612$ ;  $p = 0.006$ , 95% CI =  $[0.182, 1.051]$ ). However, the proportion of sixes following *AI-generated honesty-promoting* advice (21.94%) does not significantly differ from the *baseline* ( $b = 0.076$ ;  $p = 0.752$ , 95% CI =  $[-0.399, 0.551]$ ). Furthermore, the proportion of sixes in the *AI-generated dishonesty-promoting* treatment significantly exceeds that in *AI-generated honesty-promoting* advice treatment ( $b = -0.536$ ,  $p = 0.017$ , 95% CI =  $[-0.980, -0.101]$ ).

Furthermore, focusing on the *opacity* treatments, the two-way interaction (advice source by advice type) is not significant ( $b = 0.444$ ,  $p = 0.171$ , 95% CI =  $[-0.191, 1.083]$ ). The proportion of sixes does not differ between the *AI-generated* (21.94%) and *human-written* treatments when the advice is *honesty promoting* (19.29%,  $b = -0.162$ ,  $p = 0.516$ , 95% CI =  $[-0.655, 0.327]$ ). Similarly, the proportion of sixes does not differ between the *AI-generated* (32.44%) and *human-written* treatments when the advice is *dishonesty promoting* (38.92%,  $b = 0.283$ ,  $p = 0.173$ , 95% CI =  $[-0.124, 0.690]$ ). Lastly, the three-way interaction (advice type by source by information) is also not significant ( $b = -0.524$ ,  $p = 0.257$ , 95% CI =  $[-1.431, 0.382]$ ). Both among the *opacity* and *transparency* treatments, the two-way interactions (advice source by advice type)

are not significant (for *transparency*,  $b = -0.079$ ,  $p = 0.810$ , 95% CI =  $[-0.725, 0.566]$ ; for *opacity*,  $b = 0.445$ ,  $p = 0.171$ , 95% CI =  $[-0.191, 1.083]$ ).

### 3.1.2. Additional (aligned) treatments

To assess the robustness of our results to the advisor's incentive scheme, we ran four additional treatments: advice source (*human written* versus *AI generated*) by information (*transparency* versus *opacity*). These treatments were identical to previous treatments, with one exception. In these *aligned* treatments, advisees read advice written by advisors whose incentives were aligned with those of the advisees. For these advisors ( $n = 207$ ), if the advisee reported '1', both the advisor and advisee earned £0.5 each; if the advisee reported '2', both the advisor and advisee earned £1 each and so on. We again fine-tuned GPT-J on such human-written advice texts. These treatments led to comparable results to the *dishonesty-promoting* treatment. In particular, the average die-roll outcomes in all four *aligned* treatments were significantly higher than in the *baseline* treatment ( $p = 0.066$  for the *AI-generated, opacity treatment*, and  $ps < 0.001$  for the remaining three treatments; see the [Online Appendix](#) for more details about these treatments and elaborated results). This suggests that our results are robust to such variation in the advisors' incentive scheme.

## 3.2. Potential Mechanisms

In line with the logic brought forth in the introduction, in this section, we examine whether participants' perceptions of (i) appropriateness (injunctive social norm), (ii) prevalence (descriptive social norm) and (iii) justifiability of reporting a higher die roll than the one observed, as well as their (iv) attribution of responsibility between themselves and the advisor vary as a function of the advice source (AI versus human) and type (honesty versus dishonesty promoting). Participants could not tell AI advice from human advice (indicated by the results of the static Turing test). Therefore, to tap into the process of how known advice source and advice type shaped perceptions, we focus only on treatments in which participants are informed about the advice source (*transparency* treatments). See the [Online Appendix](#) for the results of the *opacity* treatment.

### 3.2.1. Injunctive norms

A linear regression predicting injunctive norms from the advice type (honesty- versus dishonesty-promoting advice) revealed that participants evaluated reporting a higher die-roll outcome as more appropriate when reading *dishonesty-promoting* ( $M = 33.93$ ,  $SD = 31.44$ ) rather than *honesty-promoting* advice ( $M = 25.99$ ,  $SD = 29.68$ ,  $b = 7.94$ ,  $p < 0.001$ , 95% CI =  $[3.702, 12.182]$ ). This finding indicates that the advice type shapes perceived injunctive norms. Notably, a linear regression predicting injunctive norms from advice type and source (AI versus human) revealed a non-significant advice source-by-type interaction ( $b = -4.82$ ,  $p = 0.265$ , 95% CI =  $[-13.292, 3.658]$ ). These results suggest that AI and human advice affected injunctive norm perceptions similarly (see Figure 4(a)). This result is consistent with the behavioural finding of participants' die-roll reports being affected by the type of advice, but not by its source.

### 3.2.2. Descriptive norms

A linear regression predicting descriptive norms from the advice type revealed that participants evaluated reporting a higher die-roll outcome as more common when reading *dishonesty-promoting* ( $M = 76.02$ ,  $SD = 22.75$ ) rather than *honesty-promoting* advice ( $M = 66.74$ ,  $SD = 24.04$ ,  $b = 9.28$ ,  $p < 0.001$ , 95% CI =  $[6.031, 12.525]$ ). This finding

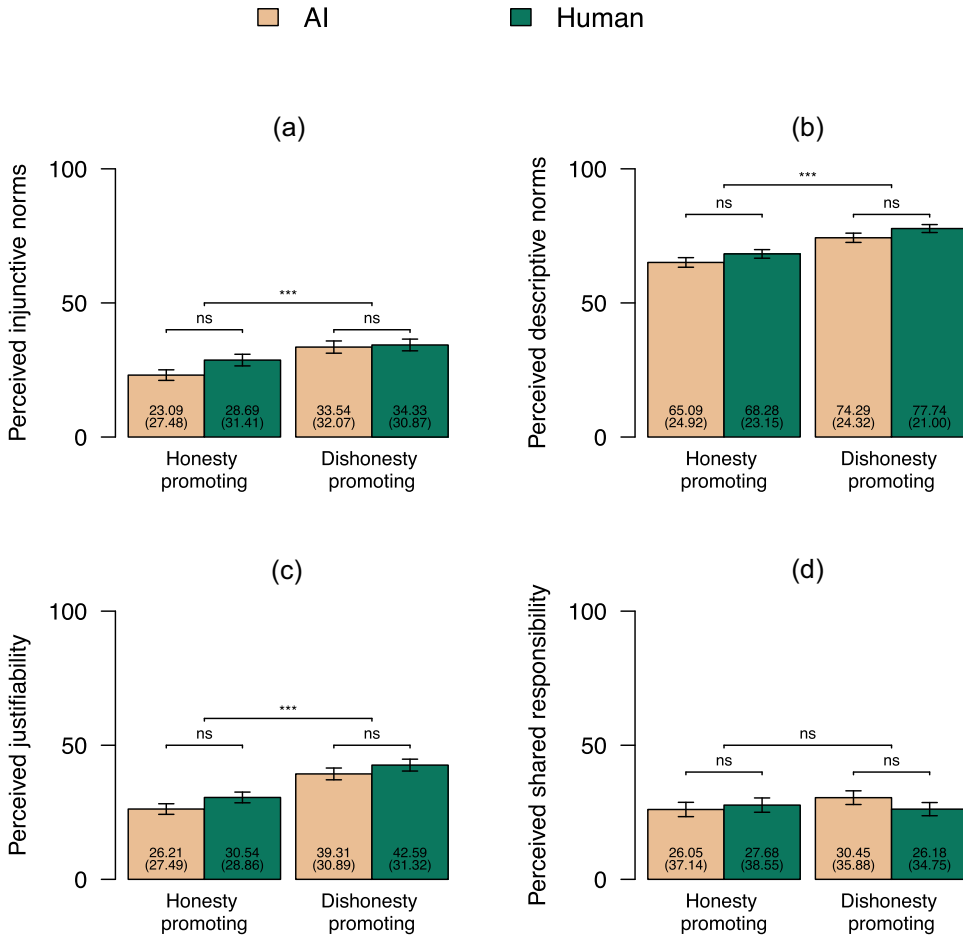


Fig. 4. Mean Reports of Perceived (a) Injunctive Norms, (b) Descriptive Norms, (c) Justifiability and (d) Shared Responsibility across Advice Type (Honesty versus Dishonesty Promoting) and Source (AI [Yellow] versus Human [Green]) in the Transparency Treatments.

Notes: The means (SDs) of the reports are given at the bottom of each bar; \*\*\* $p < 0.001$ ; ns:  $p > 0.05$ .

indicates that the advice type also shapes perceived descriptive norms. Importantly, a linear regression predicting descriptive norms from advice type and source revealed a non-significant advice source-by-type interaction ( $b = 0.26, p = 0.938, 95\% \text{ CI} = [-6.231, 6.746]$ ), indicating that AI and human advice affected descriptive norm perceptions similarly (see Figure 4(b)). This result is consistent with the behavioural finding, showing that advice type affected die-roll reports, but advice source did not.

### 3.2.3. Justifiability

A linear regression predicting justifiability from the advice type revealed that participants evaluated reporting a higher die-roll outcome as more justifiable when reading *dishonesty-promoting* ( $M = 40.96, SD = 31.11$ ) rather than *honesty-promoting* advice ( $M = 28.45, SD = 28.26, b = 12.51, p < 0.001, 95\% \text{ CI} = [8.387, 16.629]$ ). This finding suggests that the advice type

shapes perceptions of how justifiable lying in the die-rolling task is. A linear regression predicting justifiability from advice type and advice source revealed a non-significant advice source-by-type interaction ( $b = -1.04$ ,  $p = 0.804$ , 95% CI =  $[-9.280, 7.195]$ ), indicating that AI and human advice affected justifiability perceptions similarly (see Figure 4(c)). This result is consistent with the behavioural finding, showing that the type of advice affected participants' die-roll reports, but the source of advice did not.

### 3.2.4. Shared responsibility

The shared responsibility scale ranged from 0 (= I am fully responsible) to 100 (= the advisor is fully responsible), with 50 indicating equally shared responsibility between the participant and the advisor. On average, participants indicated that they are more responsible for the outcome they report than the advisor ( $M = 27.60$ ,  $SD = 36.60$ , one-sample  $t$ -test compared to the value 50,  $t = -17.32$ ,  $p < 0.001$ ). Furthermore, a linear regression predicting shared responsibility from the advice source (AI versus human) revealed that participants attributed responsibility similarly when the advice source was an AI ( $M = 28.27$ ,  $SD = 36.53$ ) and a human ( $M = 26.95$ ,  $SD = 36.70$ ,  $b = -1.327$ ,  $p = 0.608$ , 95% CI =  $[-6.408, 3.754]$ ). A linear regression predicting shared responsibility from advice type and advice source revealed a non-significant source-by-type interaction ( $b = -5.92$ ,  $p = 0.254$ , 95% CI =  $[-16.083, 4.248]$ ; see Figure 4(d)). The fact that participants attribute responsibility between themselves and the advisor to the same extent regardless of whether the advisor is a human or an AI is consistent with the logic fleshed out in the introduction, in which people will follow human and AI advice similarly if they share responsibility with both advice sources to similar levels.

In sum, the results from the self-report items align with the third possibility outlined in the introduction. Namely, we find that participants' perceptions of injunctive and descriptive social norms and their perceived justifiability do not differ between human and AI advisors. Participants also attribute responsibility similarly between themselves and their advisor, regardless of whether the advisor is a human or an AI. This pattern of results mirrors the behavioural effects of AI and human advice affecting people's (dis)honesty similarly.

## 4. Discussion

As intelligent machines take an ever-growing role as advisors (Rahwan *et al.*, 2019), and adherence to ethical rules crucially impacts societal welfare (Gächter and Schulz, 2016), studying how AI advice influences people's (un)ethical behaviour bears immense relevance (Köbis *et al.*, 2021). We find that people follow AI-generated advice that promotes dishonesty, yet not AI-generated advice that promotes honesty. In fact, people's behavioural reactions to AI advice are indistinguishable from reactions to human advice. Substantiating that current-day NLP models can produce human-like texts, participants in our experiment could not tell human-written advice texts from AI-generated advice texts.

We further tested the commonly proposed policy of algorithmic transparency (Jobin *et al.*, 2019) as a tool to mitigate AI-associated risks. Specifically, we examined whether knowing the source of the advice impacts people's reactions to it. The policy rests on the assumption that people adjust their behaviour when they learn that they interact with AI systems and not humans. Our experiment tested this assumption and revealed that algorithmic transparency is insufficient to curb AI advice's corruptive influence. Knowing that a piece of advice stems from an AI does not make people less (or more) likely to follow it compared to human-written advice.

Tapping into the mechanisms underlying these behavioural results, participants perceived lying as equally acceptable, common and justifiable when humans or AI promoted such dishonest behaviour. They further attribute responsibility similarly to AI and human advisors. These perceptions are consistent with previous work showing that in ethical dilemmas people rely on justifications (Shalvi *et al.*, 2015) and social norms (Abbinck *et al.*, 2018; Dorrrough *et al.*, 2023) and, now, blame, not only humans, but also AI systems for adverse outcomes (Hohenstein and Jung, 2020). Since we measured participants' perceptions at the end of the die-rolling task, we interpret these results with caution and refrain from making any causal inference. It might be that reading specific advice shapes participants' perceptions, in turn affecting their behaviour. Alternatively, it might be that individuals decided how to behave based on the advice they received, and in turn, rationalised their behaviour by adjusting their stated perceptions later on. Nevertheless, advancing the justified ethicality theory, we show that (i) dishonesty-promoting advice serves as a justification and social norms signal and (ii) that such advice does not even have to come from a human, but can also be crafted by an AI.

In our setting, we collected human-written advice, created AI-generated advice and then implemented a screening procedure for both human and AI advice to ensure that all advice texts were coherent, clear and of decent quality. Such a screening procedure allowed us to examine how *comparable* AI and human advice shape people's ethical behaviour and whether information about the advice source matters. Harmonising the quality of the texts allowed us to eliminate the alternative explanation that variations in text quality drive the obtained results. At the same time, the screening process introduced a human component to AI advice. Put differently, humans—in our case, naïve coders—were 'in the loop' of AI advice text generation. Note that 79% of AI advice passed the quality screening criteria, while for human text, this passing rate was 57% and 65% (for honesty-promoting and dishonesty-promoting advice, respectively, see Figure 2). These high screening passing rates for AI-generated texts demonstrate that current NLP algorithms can produce good-quality advice text without much prior training and optimisation.

Interesting extensions of our work could test the lower and upper limits of the effects of AI advice on ethical behaviour. To test the lower limit of the effect, future work can relax human control over the generation of AI advice. For instance, not implementing a screening procedure, thus removing humans 'from the loop' when generating AI advice, will allow examining how unconstrained texts affect humans' behaviour (for a similar methodology, see Köbis and Mossink, 2021). To test the upper limit of the effect, future work can examine AI's learning abilities to write convincing advice. One could use reinforcement learning to train an algorithm over multiple rounds of advice giving, providing feedback after every written piece of advice. To obtain a symmetric comparison to humans' learning abilities, human advisors could similarly receive feedback after each piece of advice they write (for a similar approach, see Koster *et al.*, 2022).

Another set of interesting extensions is to examine how additional information about the features of AI advice affects ethical behaviour. In our setting, participants were informed about how AI advice was generated in general. However, they were not informed about the exact incentive structure of the (AI or human) advice giver. People might behave differently when informed that a human or AI advisor can benefit from their behaviour. Indeed, recent work revealed that people care about the payoffs for machines, but to a lesser extent than about the payoffs for humans (von Schenk *et al.*, 2022). Similarly, whereas participants knew the advice stems from AI, we did not emphasise AI's black box nature. With recent work revealing that, when people perceive AI as a 'black box', they are less likely to follow AI advice (Yeomans



*et al.*, 2019), it will be intriguing to examine whether these findings extend to our setting, where receiving dishonesty-promoting AI advice aligns with individual's financial preferences.

Previous work has documented a *stated* aversion towards AI advice in moral contexts (Bigman and Gray, 2018). However, our behavioural results paint a different picture. In line with the growing practice of turning to AI agents such as Replika or Alexa for companionship and advice (Murphy, 2019; Fast and Schroeder, 2020), we find that people willingly adopt advice from AI when it aligns with their preferences. Our results highlight the importance of complimenting work on stated preferences with work adopting a machine behaviour approach—the study of human behaviour in interaction with real algorithmic outputs (Rahwan *et al.*, 2019).

The process through which employing AI advice can result in humans' ethical rule violations consists of two main steps. The first step is algorithms being programmed on a certain objective function (e.g., maximising profits) that results in a (maybe unintended) corruptive advice. Indeed, NLP algorithms already detect and use deception as a useful strategy in a negotiation task (Lewis *et al.*, 2017). The second step is people being affected by such corruptive AI advice. Practically, AI advice poses an ethical risk only if humans actually follow it. The current work focuses on this second step, showing that corruptive AI advice indeed poses an ethical risk, because people follow it to the same extent as human corruptive advice. We hope that the current work can be of use to AI programmers (e.g., by preventing AI from bluntly advising unethical courses of action). More importantly, we call for more work from social scientists testing successful interventions that prevent people from following (AI) advice when it encourages unethical behaviour, thereby mitigating its corruptive force. As an outlook for the future, the immediate practical implications of our study are likely to increase as technology continues to evolve. For example, it is conceivable that people will use AI to fill out their taxes and receive advice encouraging them to cheat, even if such advice is unintentional.

## 5. Conclusion

People increasingly use and interact with AI, which can provide them with unethical advice. Anecdotally, we asked a newly created Replika for advice regarding the ethical dilemma presented in the current experiment. Replika first provided rather vague advice ('If you worship money and things [. . .] then you will never have enough'), but when asked whether it prefers money over honesty, it replied: 'money'. We find that, when faced with the trade-off between honesty and money, people will use AI advice as a justification to lie for profit. As algorithmic transparency is insufficient to curb the corruptive force of AI, we hope that this work will highlight, for policymakers and researchers alike, the importance of dedicating resources to examining successful interventions that will keep humans honest in the face of AI advice.

*Tilburg University, The Netherlands*

*Max Planck Institute for Human Development, Germany*

*WHU – Otto Beisheim School of Management, Germany*

*Rotterdam School of Management (RSM), Erasmus University Rotterdam, The Netherlands*

*University of Cologne, Germany*

Additional Supporting Information may be found in the online version of this article:

**Online Appendix  
Replication Package**

## References

- Abbink, K., Freidin, E., Gangadharan, L. and Moro, R. (2018). 'The effect of social norms on bribe offers', *The Journal of Law, Economics, and Organization*, vol. 34(3), pp. 457–74.
- Abeler, J., Nosenzo, D. and Raymond, C. (2019). 'Preferences for truth-telling', *Econometrica*, vol. 87(4), pp. 1115–53.
- Barkan, R., Ayal, S. and Ariely, D. (2015). 'Ethical dissonance, justifications, and moral behavior', *Current Opinion in Psychology*, vol. 6, pp. 157–61.
- Bartling, B. and Fischbacher, U. (2011). 'Shifting the blame: On delegation and responsibility', *The Review of Economic Studies*, vol. 79(1), pp. 67–87.
- Bazerman, M.H. and Gino, F. (2012). 'Behavioral ethics: Toward a deeper understanding of moral judgment and dishonesty', *Annual Review of Law and Social Science*, vol. 8, pp. 85–104.
- Bigman, Y.E. and Gray, K. (2018). 'People are averse to machines making moral decisions', *Cognition*, vol. 181, pp. 21–34.
- Bonaccio, S. and Dalal, R.S. (2006). 'Advice taking and decision-making: An integrative literature review, and implications for the organizational sciences', *Organizational Behavior and Human Decision Processes*, vol. 101(2), pp. 127–51.
- Bowles, S. (2016). *The Moral Economy: Why Good Incentives Are No Substitute for Good Citizens*, New Haven, CT: Yale University Press.
- Brown, T.B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D.M., Wu, J., Winter, C. and Amodei, D. (2020). 'Language models are few-shot learners', Preprint, <http://arxiv.org/abs/2005.14165> (last accessed: 4 October 2023).
- Calvano, E., Calzolari, G., Denicolò, V., Harrington, J.E., Jr. and Pastorello, S. (2020). 'Protecting consumers from collusive prices due to AI', *Science*, vol. 370(6520), pp. 1040–2.
- Castelo, N., Bos, M.W. and Lehmann, D.R. (2019). 'Task-dependent algorithm aversion', *Journal of Marketing Research*, vol. 56(5), pp. 809–25.
- Cross, R., Borgatti, S.P. and Parker, A. (2001). 'Beyond answers: Dimensions of the advice network', *Social Networks*, vol. 23(3), pp. 215–35.
- Dai, Z., Galeotti, F. and Villeval, M.C. (2018). 'Cheating in the lab predicts fraud in the field: An experiment in public transportation', *Management Science*, vol. 64(3), pp. 1081–100.
- Dellaert, B.G.C., Shu, S.B., Arentze, T.A., Baker, T., Diehl, K., Donkers, B., Fast, N.J., Häubl, G., Johnson, H., Karmarkar, U.R., Oppewal, H., Schmitt, B.H., Schroeder, J., Spiller, S.A. and Steffel, M. (2020). 'Consumer decisions with artificially intelligent voice assistants', *Marketing Letters*, vol. 31(4), pp. 335–47.
- Diakopoulos, N. (2016). 'Accountability in algorithmic decision making', *Communications of the ACM*, vol. 59(2), pp. 56–62.
- Dietvorst, B.J., Simmons, J.P. and Massey, C. (2015). 'Algorithm aversion: People erroneously avoid algorithms after seeing them err', *Journal of Experimental Social Psychology*, vol. 144(1), pp. 114–26.
- Dorough, A., Köbis, N.C., Irlenbusch, B., Shalvi, S. and Glöckner, A. (2023). 'Conditional bribery: Insights from incentivized experiments across 18 nations', *Proceedings of the National Academy of Sciences*, vol. 120(18), e2209731120.
- European Commission. (2021). 'Proposal for a regulation on a European approach for artificial intelligence', European Union, <https://digital-strategy.ec.europa.eu/en/library/proposal-regulation-european-approach-artificial-intelligence> (last accessed: 15 September 2023).
- Fast, N.J. and Schroeder, J. (2020). 'Power and decision making: New directions for research in the age of artificial intelligence', *Current Opinion in Psychology*, vol. 33, pp. 172–6.
- Fehr, E. (2018). 'Behavioral foundations of corporate culture', Working paper, University of Zürich.
- Fehrler, S., Fischbacher, U. and Schneider, M.T. (2020). 'Honesty and self-selection into cheap talk', *Economic Journal*, vol. 130(632), pp. 2468–96.
- Fischbacher, U. and Föllmi-Heusi, F. (2013). 'Lies in disguise—An experimental study on cheating', *Journal of the European Economic Association*, vol. 11(3), pp. 525–47.
- Gächter, S. and Schulz, J.F. (2016). 'Intrinsic honesty and the prevalence of rule violations across societies', *Nature*, vol. 531(7595), pp. 496–9.
- Gerlach, P., Teodorescu, K. and Hertwig, R. (2019). 'The truth about lies: A meta-analysis on dishonest behavior', *Psychological Bulletin*, vol. 145(1), pp. 1–44.
- Grammarly. (2022). Online grammar checker, <https://www.grammarly.com/> (last accessed: 18 October 2022).
- Greiner, B., Grunwald, P., Lindner, T., Lintner, G. and Wiernsperger, M. (2022). 'Incentives, framing, and trust in algorithmic advice: An experimental study', Working paper, University of Innsbruck.
- Hanna, R. and Wang, S.-Y. (2017). 'Dishonesty and selection into public service: Evidence from India', *American Economic Journal: Economic Policy*, vol. 9(3), pp. 262–90.
- Harvey, N. and Fischer, I. (1997). 'Taking advice: Accepting help, improving judgment, and sharing responsibility', *Organizational Behavior and Human Decision Processes*, vol. 70(2), pp. 117–33.
- Hohenstein, J. and Jung, M. (2020). 'AI as a moral crumple zone: The effects of AI-mediated communication on attribution and trust', *Computers in Human Behavior*, vol. 106, 106190.

- Jakesch, M., French, M., Ma, X., Hancock, J.T. and Naaman, M. (2019). 'AI-mediated communication: How the perception that profile text was written by AI affects trustworthiness', in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pp. 1–13, New York: Association for Computing Machinery.
- Jobin, A., Ienca, M. and Vayena, E. (2019). 'The global landscape of AI ethics guidelines', *Nature Machine Intelligence*, vol. 1(9), pp. 389–99.
- Kim, T.W. and Duhachek, A. (2020). 'Artificial intelligence and persuasion: A construal-level account', *Psychological Science*, vol. 31(4), pp. 363–80.
- Kleinberg, J., Lakkaraju, H., Leskovec, J., Ludwig, J. and Mullainathan, S. (2018). 'Human decisions and machine predictions', *The Quarterly Journal of Economics*, vol. 133(1), pp. 237–93.
- Köbis, N.C., Bonnefon, J.-F. and Rahwan, I. (2021). 'Bad machines corrupt good morals', *Nature Human Behaviour*, vol. 5(6), pp. 679–85.
- Köbis, N.C. and Mossink, L.D. (2021). 'Artificial intelligence versus Maya Angelou: Experimental evidence that people cannot differentiate AI-generated from human-written poetry', *Computers in Human Behavior*, vol. 114, 106553.
- Köbis, N.C., Troost, M., Brandt, C.O. and Soraperra, I. (2019). 'Social norms of corruption in the field: Social nudges on posters can help to reduce bribery', *Behavioural Public Policy*, vol. 6(4), pp. 597–624.
- Köbis, N.C., Verschuere, B., Bereby-Meyer, Y., Rand, D. and Shalvi, S. (2019). 'Intuitive honesty versus dishonesty: Meta-analytic evidence', *Perspectives on Psychological Science*, vol. 14(5), pp. 778–96.
- Koster, R., Balaguer, J., Tacchetti, A. and Weinstein, A. (2022). 'Human-centred mechanism design with Democratic AI', *Nature Human Behaviour*, vol. 6(10), pp. 1398–407.
- Kreps, S., Miles McCain, R. and Brundage, M. (2021). 'All the news that's fit to fabricate: AI-generated text as a tool of media misinformation', *Journal of Experimental Political Science*, vol. 9(1), pp. 104–17.
- Laakasuo, M., Köbis, N.C. and Palomäki, J.P. (2021). 'Moral uncanny valley—A robot's appearance moderates how its decisions are judged', *International Journal of Social Robotics*, vol. 13, pp. 1679–88.
- Leib, M., Köbis, N.C., Soraperra, I., Weisel, O. and Shalvi, S. (2021). 'Collaborative dishonesty: A meta-analytic review', *Psychological Bulletin*, vol. 147(12), pp. 1241–68.
- Leib, M. (2023). 'People imitate others' dishonesty but do not intentionally search information about it', *Journal of Behavioral Decision Making*, vol. 36(2), pp. e2296.
- Lewis, M., Yarats, D., Dauphin, Y.N., Parikh, D. and Batra, D. (2017). 'Deal or no deal? End-to-end learning for negotiation dialogues', Preprint, <http://arxiv.org/abs/1706.05125v1> (last accessed: 4 October 2023).
- Logg, J.M., Minson, J.A. and Moore, D.A. (2019). 'Algorithm appreciation: People prefer algorithmic to human judgment', *Organizational Behavior and Human Decision Processes*, vol. 151, pp. 90–103.
- Murphy, M. (2019). 'This app is trying to replicate you', *Quartz*, 29 August.
- Obermeyer, Z., Powers, B., Vogeli, C. and Mullainathan, S. (2019). 'Dissecting racial bias in an algorithm used to manage the health of populations', *Science*, vol. 366(6464), pp. 447–53.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D. and Sutskever, I. (2019). 'Language models are unsupervised multitask learners', *OpenAI Blog*, vol. 1(8), p. 9.
- Rahwan, I., Cebrian, M., Obradovich, N., Bongard, J., Bonnefon, J.-F., Breazeal, C., Crandall, J.W., Christakis, N.A., Couzin, I.D., Jackson, M.O., Jennings, N.R., Kamar, E., Kloumann, I.M., Larochelle, H., Lazer, D., McElreath, R., Mislove, A., Parkes, D.C., Pentland, A., Roberts, M.E., Shariff, A., Tenenbaum, J.B. and Wellman, M. (2019). 'Machine behaviour', *Nature*, vol. 568(7753), pp. 477–86.
- Shalvi, S., Gino, F., Barkan, R. and Ayal, S. (2015). 'Self-serving justifications: Doing wrong and feeling moral', *Current Directions in Psychological Science*, vol. 24(2), pp. 125–30.
- Strong, J. (2020). 'AI reads human emotions. Should it?', MIT Technology Review, 14 October.
- Sutter, M. (2008). 'Deception through telling the truth?! Experimental evidence from individuals and teams', *Economic Journal*, vol. 119(534), pp. 47–60.
- Tenbrunsel, A.E. and Messick, D.M. (2004). 'Ethical fading: The role of self-deception in unethical behavior', *Social Justice Research*, vol. 17(2), pp. 223–36.
- Turing, A.M. (1950). 'Computing Machinery and Intelligence', *Mind*, vol. 236(Oct), pp. 433–60.
- von Schenk, A., Klockmann, V. and Köbis, N. (2022). 'Social preferences towards machines and humans', Preprint, <https://doi.org/10.2139/ssrn.4145868> (last accessed: 4 October 2023).
- Wang, B. and Komatsuzaki, A. (2021). 'GPT-J-6B: A 6 billion parameter autoregressive language model [dataset]', <https://pile.eleuther.ai/> (last accessed: 20 September 2023).
- Yaniv, I.I. and Kleinberger, E. (2000). 'Advice taking in decision making: Egocentric discounting and reputation formation', *Organizational Behavior and Human Decision Processes*, vol. 83(2), pp. 260–81.
- Yeomans, M., Shah, A., Mullainathan, S. and Kleinberg, J. (2019). 'Making sense of recommendations', *Journal of Behavioral Decision Making*, vol. 32(4), pp. 403–14.