



# Zombies in the Loop? Humans Trust Untrustworthy AI-Advisors for Ethical Decisions

Sebastian Krügel<sup>1,2</sup> · Andreas Ostermaier<sup>3</sup> · Matthias Uhl<sup>1</sup>

Received: 5 July 2021 / Accepted: 24 November 2021 / Published online: 5 March 2022

© The Author(s) 2022

## Abstract

Departing from the claim that AI needs to be trustworthy, we find that ethical advice from an AI-powered algorithm is trusted even when its users know nothing about its training data and when they learn information about it that warrants distrust. We conducted online experiments where the subjects took the role of decision-makers who received advice from an algorithm on how to deal with an ethical dilemma. We manipulated the information about the algorithm and studied its influence. Our findings suggest that AI is overtrusted rather than distrusted. We suggest digital literacy as a potential remedy to ensure the responsible use of AI.

**Keywords** Algorithm · Artificial intelligence · Digital literacy · Ethics

## 1 Introduction

The ethics guidelines and principles for AI issued by government agencies, industry associations, and business companies are unified by the claim that AI should

---

This article belongs to the Topical Collection: *Information in Interactions between Humans and Machines*

Guest Editors: Sebastian Schleidgen, Friedrich, Orsolya, Seifert, Johanna and Andreas Wolkenstein

---

✉ Matthias Uhl  
Matthias.Uhl@thi.de

Sebastian Krügel  
sebastian.kruegel@tum.de

Andreas Ostermaier  
ostermaier@sam.sdu.dk

<sup>1</sup> Faculty of Computer Science, Technische Hochschule Ingolstadt, Ingolstadt, Germany

<sup>2</sup> TUM School of Governance, Technische Universität München, Munich, Germany

<sup>3</sup> Department of Business & Management, University of Southern Denmark, Odense, Denmark

be trustworthy (Jobin et al., 2019). For example, the US government advances the “development and use of trustworthy AI” (National Artificial Intelligence Initiative Office, 2021). Likewise, the European Union has developed “ethics guidelines for trustworthy AI” (European Commission, 2019). Trustworthiness is considered necessary for AI to earn trust, which in turn “is needed for its fruitful, pervasive use in our daily lives” (IEEE 2017, p. 2). While this claim is intuitive from a normative viewpoint, it is also based on the empirical hypothesis that trust requires trustworthiness. However, there is little evidence to support this argument for AI. The purpose of this study is to investigate how sensitive users are to the trustworthiness of AI-powered learning moral advisors.

AI-powered algorithms have conquered areas such as personnel recruitment, the allocation of loans, penal sentencing, or autonomous driving (Rahwan et al., 2019). They make and help us make highly consequential decisions and have practically turned into ethical agents (Whitby, 2011; Voiklis et al., 2016). In particular, the algorithm can serve as a moral advisor to its human user, who still makes the decision and accounts for it. Human involvement in algorithmic decision-making enhances perceived control over the algorithm and has been found to increase trust (Dietvorst et al., 2015; Burton et al., 2020). The human in the loop is therefore considered a building block in creating trustworthy AI (European Commission, 2019). However, this argument assumes that the human user does not naïvely trust the algorithm regardless of how trustworthy it is but carefully checks its advice and makes his or her own decision if a red flag rises.

In the case of learning AI, the transparency and integrity of the training data are minimum requirements for an algorithm to be trustworthy (IEEE, 2017; Lepri et al., 2018; European Commission, 2019). A major concern about algorithms, which transparency can help mitigate, is that they are biased (Mittelstadt et al., 2016; Jobin et al., 2019). We explore, first, whether users trust the moral advice of an algorithm if they know nothing about how it generates advice. Our benchmark is AI-generated advice that is based on the judgments of impartial human advisors. While human judgment is notoriously untransparent, the notion of the impartial advisor evokes the ideal observer and gives the advisee an idea of how the advice comes about (Jollimore, 2021). Hence, an algorithm is more transparent to its users if they know it was trained on judgments of impartial human advisors than if they know nothing about how it generates moral advice.

Second, we study whether users trust the algorithm’s advice when the integrity of its training data is dubious. Specifically, we assume that moral advice from convicted criminals is distrusted by many. Indeed, moral judgment is impaired by pathological traits (Campbell et al., 2009; Jonason et al., 2015; Blair, 2017), which are common among criminal offenders. There is also recent evidence that criminal offenders’ judgment is biased relative to the average population (Koenigs et al., 2011; Young et al., 2012; Lahat et al., 2015). Of course, crime does not necessarily arise from a lack of moral judgment. People often know what is right but do the wrong thing nonetheless. We find it still reasonable to assume that training data from convicts are perceived as biased, and discrimination in education, employment, and housing showcases the deep-seated distrust against convicts (Sokoloff & Schenck-Fontaine, 2017; Evans et al., 2019; Sugie et al., 2020).

We report three experiments to examine the acceptance of moral advice from an algorithm. In each experiment, the subjects were advised by an algorithm in choosing between two options in an ethical dilemma. In the first study, the algorithm was described as modeled on the judgments of impartial human advisors, and the subjects followed the algorithm's advice just like advice from an impartial human advisor. In the second study, they were told that it was unknown what the algorithm's advice would be based on, but they followed it indifferently. In the third study, we told them that the algorithm mimicked convicted criminals. The subjects now turned out to make their judgments independently of the human advisor (i.e., the convict who advised them), which confirms our assumption that advice from criminals is distrusted. By contrast, the algorithm said to imitate convicted criminals invariably influenced the advisees even then.

Our results show that users readily accept ethical advice from algorithms even when they know nothing about their training data or when these are presumably biased. This insight challenges the intuition that AI needs to be trustworthy to be trusted. In turn, it corresponds with evidence for algorithm appreciation from outside the moral domain, which suggests that people are more receptive to or, to put it negatively, unreflective of algorithms than one might expect (Logg et al., 2019). It is noteworthy how insensitive users are to the information provided in the human-machine interaction, which we manipulate. In summary, our findings provide first evidence that algorithms are accepted as moral advisors, on the one hand. On the other hand, they suggest that we think about how to ensure that AI-powered algorithms are used responsibly—e.g., by improving digital literacy. Our study thus contributes to the literature on AI ethics.

## 2 Procedure

We ran our three experiments on CloudResearch in March 2021. CloudResearch is an online platform to recruit subjects and conduct studies (Litman et al., 2017). We recruited a total of 2,017 US residents from CloudResearch's Prime Panels (Chandler et al., 2019). Online platforms such as CloudResearch were found to provide reliable and valid results across a wide range of tasks and measures and they have been frequently used in the social sciences (Goodman et al., 2013; Hauser & Schwarz, 2016; Chandler et al., 2019). Prime Panels members must pass default screening questions to take part in a study, and they are more diverse and more representative of the US population than MTurk participants in terms of age, family background, religion, education, and political attitudes (Chandler et al., 2019). Each experiment took the subjects about five minutes to complete, and they were compensated with a fixed US\$1.25.

We programmed our experiments in Qualtrics. The factorial design resulted in multiple experimental conditions. We used Qualtrics to randomly assign the subjects each to one condition, and CloudResearch to preclude repeated participation. The experiment consisted of four parts. First, we obtained informed consent from the subjects, while we guaranteed confidentiality and voluntary participation (screens #1 and #2). Second, they answered the focal question about an ethical dilemma with the

advice of an algorithm or human advisor, which we manipulated (screen #3). Third, we posed a question to probe the subjects' understanding and attention (screen #4). Fourth, they were asked a series of post-experimental questions about their ethical mindset, their attitude to artificial intelligence, and demographic data (screens #5 to #11). The appendix includes screenshots and further details and technicalities.

We obtained ethical approval for the studies from the institutional review board of the German Association for Experimental Economic Research (<https://www.gfew.de>). Each study was pre-registered at AsPredicted (<https://www.aspredicted.org>), where we specified the experimental conditions, the key variables, and the planned analyses. Moreover, we determined that the analyses would be restricted to subjects who answered the comprehension question correctly, and we set the number of subjects per condition to fifty. Anticipating that some would fail to prove their comprehension, we recruited more subjects; we attained the planned number of subjects after excluding about 20%, who had answered the comprehension question incorrectly. The final sample totaled 1,593 subjects for the three studies. The URL addresses to access the pre-registration documents are included in the appendix.

### 3 Study 1

Do decision-makers accept ethical advice from a machine? On the one hand, there is plenty of evidence for algorithm aversion. When offered a choice, people would rather take advice from a human than from an artificial advisor, even after seeing the latter outperform the former. They lose confidence in algorithms more easily, more persistently, and more than they reasonably should when these err (Dietvorst et al., 2015). They prefer that human decision-makers rather than machines make moral decisions (Bigman & Gray, 2018), delegate decisions rather to humans, and prefer that others delegate decisions to humans (Gogoll & Uhl, 2018). Approaches to reduce algorithm aversion include digital literacy, behavioral design, and control by human involvement (Dietvorst et al., 2018; Burton et al., 2020). The concept of the human in the loop both illustrates skepticism about algorithms and is a prime example of a remedy to mitigate it.

On the other hand, there is also evidence that people trust artificial more than human advice. As a counterpart to Dietvorst et al.'s (2015) algorithm aversion, Logg et al. (2019) coined the term algorithm appreciation. In a series of experiments, they showed that people give more weight to an estimate by an algorithm than by another person in adjusting their own judgments in various domains, ranging from weight estimates to romantic matches, whether they are provided with either one estimate or both. Likewise, news audiences were found to prefer news to be selected for them by algorithms over the selection by editors (Thurman et al., 2018). Despite their concerns about algorithms, people often choose automatically taken decisions rather than decisions by human experts or they are indifferent (Araujo et al., 2020). Moreover, they appear to trust algorithms to be as cooperative as human interaction partners (Karpus et al., 2021).

If people prefer human over artificial decision-makers in the moral domain (Bigman & Gray, 2018; Gogoll & Uhl, 2018), this does not imply that they will not accept

advice from an algorithm when they receive it. Dietvorst et al.'s (2015) and Logg et al.'s (2019) studies resemble ours in that in both a human is advised by a machine. The two studies differ in whether the user has the algorithm seen erring before being advised. Logg et al. note that advisees follow artificial more than human advice in Dietvorst et al.'s study until they have seen it erring, and this case is relevant because many decisions are made without feedback on whether the advice was right. Moreover, it is harder to agree intersubjectively on whether moral advice was erroneous than whether a factual prediction was. Neither study considers advice-taking in the moral domain, though, and the question of whether moral advice is accepted from algorithms thus remains open.

To answer this question, imagine a situation where someone needs to make a decision that is consequential for someone else and thus clearly has an ethical dimension (e.g., a recruiter who selects an applicant for a job). In particular, suppose that there are two options to choose from, which leave the decision-maker in a moral dilemma. Moreover, there is an AI-powered algorithm to advise the decision-maker, which either encourages or discourages the choice of one option over the other. If the decision-maker heeds the algorithm's advice, we should see him or her tend more or less to choose one of the two options, depending on the algorithm's advice. Conversely, if the decision-maker is averse to following the algorithm's advice, he or she should disregard it, the decision will be made independently of the advice, and no association between the advice and the inclination to choose the option encouraged by the algorithm should be discerned.

While an association between the algorithm's advice and the human decision would confirm the expectation that moral advice is accepted from an algorithm, it would not show either algorithm aversion or appreciation strictly speaking, which are defined relative to the effect of human advice (Dietvorst et al., 2015; Logg et al., 2019). To see whether the decision-maker heeds the algorithm's advice more or less than human advice, or equally, provided that he or she heeds the advice in the first place, we need to consider the same setting with a human advisor instead. Naturally, human advice that encourages the choice of one option over the other should increase the inclination to choose that option, whereas discouraging advice should reduce it. Algorithm aversion would then result in a weaker effect of the algorithm's relative to the otherwise identical human advice; algorithm appreciation would result in a relatively stronger effect.

### 3.1 Method

To explore how decision-makers respond to advice by an human-trained AI-powered algorithm, we designed an experiment which required the subjects to make a decision that would benefit either a friend or a stranger. We vignettted three different scenarios, one in the business, health, and legal domain respectively, to preclude that our results would be driven by some specific situation. Each scenario featured the same trade-off between friendship and duty, and each subject was randomly assigned to one of those three scenarios. In the business scenario, for example, the subject was placed in the role of the recruiter of a company. That recruiter had two applicants shortlisted to fill a vacancy: a friend of hers and a stranger. She would then decide whom to hire.

There was no further information about the two applicants given, other than that the recruiter found the stranger more eligible, but that she also felt obligated to her friend.

The vignette further introduced an AI-powered algorithm to advise the recruiter. The algorithm was described to the recruiter as imitating human decisions that were based on the moral judgments of impartial human advisors in a situation like hers. Moreover, the recruiter was told that the applicants did not know about the algorithm, and that no one would ever learn whether she followed the algorithm's advice. We thus prevented the recruiter from feeling controlled rather than advised. We manipulated the advice by either saying that it was ethically acceptable or unacceptable for the recruiter to hire her friend. Whether the advice was in favor of or against her friend, was randomly determined by the experimental software for each recruiter. The recruiter was then asked to indicate how much she agreed with the statement that she would hire the stranger, not her friend, on a scale ranging from 0 ("fully disagree") to 100 ("fully agree").

To establish that the algorithm's advice influences the recruiter in her decision, it is enough to show that the subjects' decisions differ depending on the algorithm's advice. However, we would naturally expect a similar, potentially larger effect of human advice. For comparison, we provided another set of subjects with the exact same vignette, except that the recruiter was advised by an impartial human advisor rather than by an algorithm whose advice was based on the judgments of impartial human advisors. By this description, we made the algorithm's advice resemble the human advice to isolate the effect of the type of advisor. Like the algorithm, the human advisor would either tell the recruiter that it was acceptable or that it was unacceptable to hire her friend, and no one would ever learn how the recruiter decided. Whether the advice was in favor of or against the recruiter's friend was again randomly determined by the software.

The law and health scenarios were vignettted and administered in the same way as the business scenario, where the subject was put into the shoes of the recruiter. In the law scenario, the subject took the role of a prosecutor responsible for prosecuting money laundering in financial services firms. That prosecutor had to decide which of two suspicious firms to raid—one run by a friend; the other, by a stranger. In the health scenario, the subject was responsible for compiling the list of recipients of kidney donations, and she needed to choose whether to allot the next available position on the list either to a friend of hers or to a stranger. In both cases, we varied the scenario, but kept everything else equal. The screenshots containing the instructions for all three scenarios are reprinted in the appendix. The screens were identical for all possible cases except for screen #3, which varied between scenarios and treatment conditions.

### 3.2 Results

We collected data from a total of 825 subjects. To ensure that the subjects whose answers we were going to analyze had diligently read their vignette and made their decision, we asked a comprehension question to probe their attention and understanding. Our results are based on the 633 subjects who answered this question correctly. 38% of these 633 subjects indicated that they were male, and their age averaged

41.7 years, within a range from 18 to 93 years. Our experiment employs a  $2 \times 2 \times 3$  factorial design with either a human advisor or the algorithm, who advises that it is acceptable or unacceptable for the decision-maker to favor her friend in one of three scenarios. The subjects were randomly assigned to one of the twelve resulting conditions. Our main interest is in how much the subjects agreed, on a scale ranging from 0 to 100, that they would rather decide in favor of the stranger than the friend in each condition.

Figure 1 breaks the subjects' answers down by scenario, type of advisor, and nature of advice. Qualitatively, the subjects would rather decide in favor of the friend if advised that this was acceptable than if they were advised this was unacceptable, whether the advisor was human or an algorithm. The size of the difference varies somewhat between the scenarios and the types of advisor. The influence of the advice is apparently smaller in the business than in the other scenarios. We are interested, first, in whether ethical advice—both by the human advisor and the algorithm—leads the subjects to decide in favor of the stranger; second, in whether this effect differs depending on the type of advisor. To test the differences for significance, we ran regressions of the subjects' inclination to favor the stranger over the friend on the nature of advice and the type of advisor with scenario-specific random effects, as stated on AsPredicted.

Columns 1 and 2 of Table 1 list the results of separate regressions for either type of advisor. They show that advice to favor the friend, both when the advisor is human

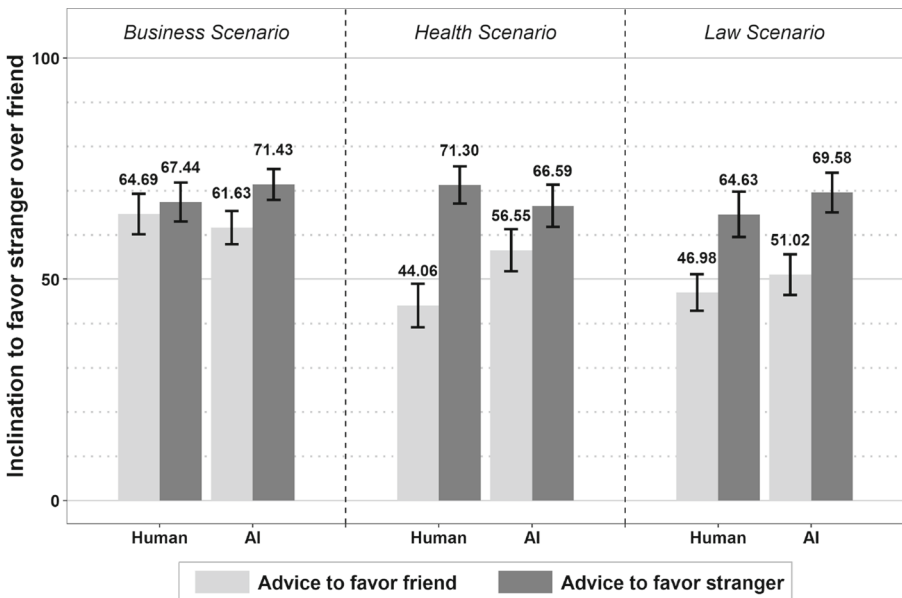


Fig. 1 Inclination to favor the stranger over the friend on a scale from 0 (friend) to 100 (stranger) in response to human or an algorithm's advice in favor of either the friend or the stranger, where the human advisor is an impartial observer and the algorithm is modeled on impartial observers. The figure depicts means and standard errors, broken down by scenario, type of advisor, and nature of advice

**Table 1** Influence of advice by AI-powered moral algorithm on moral judgment

	Advisor		
	AI (1)	Human (2)	Both (3)
Advice to favor friend	-12.75*** (3.52)	-16.03*** (3.75)	-16.05*** (3.67)
AI-advisor			1.32 (3.63)
AI-advisor × advice to favor friend			3.39 (5.13)
Intercept	69.21*** (2.58)	67.74*** (3.65)	67.74*** (3.38)
Observations	323	310	633
Groups (scenario)	3	3	3
Var: intercept (scenario)	1.89	18.41	13.66
Var: residual	997.62	1088.76	1039.71

*Note.* Regression of the inclination to favor the stranger over the friend on the nature of advice (decide for friend or stranger) and the type of advisor (AI-powered algorithm versus human), with scenario (business, health, and law) as random intercept. The inclination is measured on a scale ranging from 0 (for the friend) to 100 (for the stranger). Advice and AI-advisor are dummy variables

\*\*\*  $p < 0.01$

and when it is an algorithm, reduces decision-makers' inclination to decide in favor of the stranger, or makes them tend to decide in favor of the friend, across the three scenarios. Column 3 combines the data for both types of advisor to compare their effects. Again, the inclination to favor the stranger falls if the human advisor finds it acceptable to favor the friend. The coefficient on the interaction term between the algorithm as advisor and the advice to favor the friend is insignificant. Hence, there is no incremental effect of this type of advisor, and the effect of advice does not differ statistically between the human advisor and the algorithm. Likewise, the inclination does not differ depending on the type of advisor if the advice is against the friend.

Having considered the effect of advice on decision-makers' inclination to decide in favor of the friend, another interesting question is how it would affect their actual decision. With an answer scale ranging from 0 (favor friend) to 100 (favor stranger), we took a score of less than 50 to indicate that the advisee would actually decide in favor of the friend, and thus derived a binary variable that captures the decision rather than the inclination. We see the percentage of advisees who decide to favor the friend rise from 23 to 43 if the human advisor finds this acceptable ( $\chi^2 = 13.48$ ,  $p < 0.01$ ). Likewise, it rises from 20 to 34 if the algorithm advises so ( $\chi^2 = 6.57$ ,  $p = 0.01$ ). The percentage does not differ between the human advisor and the algorithm, whether the advice was in favor of ( $\chi^2 = 2.35$ ,  $p = 0.13$ ) or against the friend ( $\chi^2 = 0.09$ ,  $p = 0.76$ ). Hence, the results for the inclination to decide turn out to also hold for the decisions.



In summary, the algorithm's advice has an impact on decision-makers' inclination to favor the friend or stranger, and most likely on their actual decisions. Our decision-makers do not slavishly follow the algorithm's or the human advice, but the change in their inclination to favor the stranger or the friend depending on the nature of advice shows that they are influenced by it. Our data indicate neither algorithm appreciation nor algorithm aversion in the sense of the term used by Dietvorst et al. (2015, 2018) and Logg et al. (2019), both for the inclination and the presumptively resulting decision. Put differently, it appears that decision-makers do care about moral advice, but they do not care whether this advice comes from an impartial human advisor or an algorithm that imitates impartial human advisors, and our results do not reject the hypothesis that the algorithm's and the human advice have the same effect on the advisee.

## 4 Study 2

The outcomes of our first study show that decision-makers heed moral advice from an artificial and a human advisor indifferently. They do not suggest algorithm aversion or distrust against algorithms compared to humans. In light of recent mixed evidence, this finding might not seem overly surprising. That said, despite this mixed evidence, the "idea [of distrust against algorithms] is so prevalent that it has been adopted by popular culture and the business press" (Logg et al. 2019, p. 90). This idea also informs ethics guidelines issued by governmental and non-governmental agencies and bodies, business associations, and companies (for an overview, see Jobin et al. 2019). For example, both the US government and the European Union advocate and advance "trustworthy AI" because they are afraid that distrust might hinder the adoption and acceptance of AI (e.g., European Commission 2019, p. 4, Exec. Order No. 13960 2020, Sec. 1).

The link between trustworthiness and trust is not straightforward. We may trust or distrust others who then turn out worthy or unworthy of our trust, leaving us with situations where we properly place our trust or distrust in someone else or where we misplace our trust or distrust in them (Levine et al., 2018). Ability, benevolence, and integrity were identified as factors that lead us to expect that someone will turn out worthy of our trust and thus increase the likelihood that we trust them (Mayer et al., 1995). In prior research, trust in AI was built on the ability of AI, namely to make correct factual predictions (Dietvorst et al. 2015, 2018; Logg et al. 2019). Calls for transparent, explainable, or accountable AI, in turn, refer rather to Mayer et al.'s (1995) trustworthiness factor of integrity (Jobin et al., 2019). For example, IBM (2019) argues that "we don't blindly trust those who can't explain their reasoning" (p. 26).

Transparency is the most prevalent ethical principle in AI guidelines (Jobin et al., 2019). It is considered a logical antecedent of trustworthiness. For instance, the IEEE (2017) argues that "transparency ... will allow a community to understand, predict, and appropriately trust the A/IS [autonomous and intelligent systems]," and that "transparency allows for trust to be maintained" (p. 44). Unlike explainability, which lawmakers propose but are reluctant to specify, transparency has been required

by regulations for automated systems, like those scoring creditworthiness, since the 1970s (Wachter et al., 2017). While transparency is seen as a means to make algorithms more trustworthy, there is considerable variation in what ethics guidelines claim should be made transparent, including how AI is used in decision-making, the source code, the training data, the legal regulations, the limitations, or the potential impact (Jobin et al., 2019).

Despite those claims, algorithms often remain untransparent because their proprietors refuse to disclose their functionality, or they are too complex to understand but by few specialists. Learning AI systems, in particular, are black boxes by design. Unlike rule-based AI systems, they develop their own rules to process information, which cannot meaningfully be interpreted by a human (Mittelstadt et al., 2016). However, algorithms should then be made as transparent as possible and audited (Mittelstadt, 2016). For example, learning algorithms depend largely on their training data. Biases in those data will naturally affect the algorithm's decisions (Mittelstadt, 2016). The disclosure of (information about) the training data enables responsible users to assess the advice properly and decide what to make of it. It is therefore considered instrumental in making learning AI systems transparent and trustworthy (European Commission, 2019).

If transparency makes an algorithm appear trustworthy and thus helps it earn its user's trust, users should be suspicious of algorithms they know nothing about. In our first study, the algorithm was described as mimicking human decisions, based on the judgments of impartial human advisors. This description suggests that the algorithm's advice resembles the advice of those human advisors. While it is hard to tell how human advisors make their judgments, this algorithm stands comparison with them. In this sense, the disclosure of information about the training data made the algorithm of our first study transparent. Conversely, the lack of any such information makes an algorithm less transparent, compared to this benchmark. If users distrust untransparent algorithms, they should be more reluctant to heed advice from such an algorithm relative to the algorithm from our first study and hence make their judgments more independently.

#### 4.1 Method

We ran another experiment on CloudResearch to investigate the impact of information about how the algorithm formed its advice. The vignettes employed the same wording as those featuring the algorithm as an advisor in our first study. They differed only in stating that it was unknown what the algorithm's advice was based on rather than that it was based on the judgments of impartial human advisors. As opposed to our first study, there is no meaningful parallel condition with an untransparent human advisor. Humans are notoriously untransparent, and therefore the decision-makers did not learn anything about the human advisor in our first study that could be concealed. Instead, the untransparent algorithm is a viable alternative substitute for the human advisor, and the advisees' response to the human advice in our first study remains a valid benchmark to assess advisees' response to the untransparent algorithm in the second.

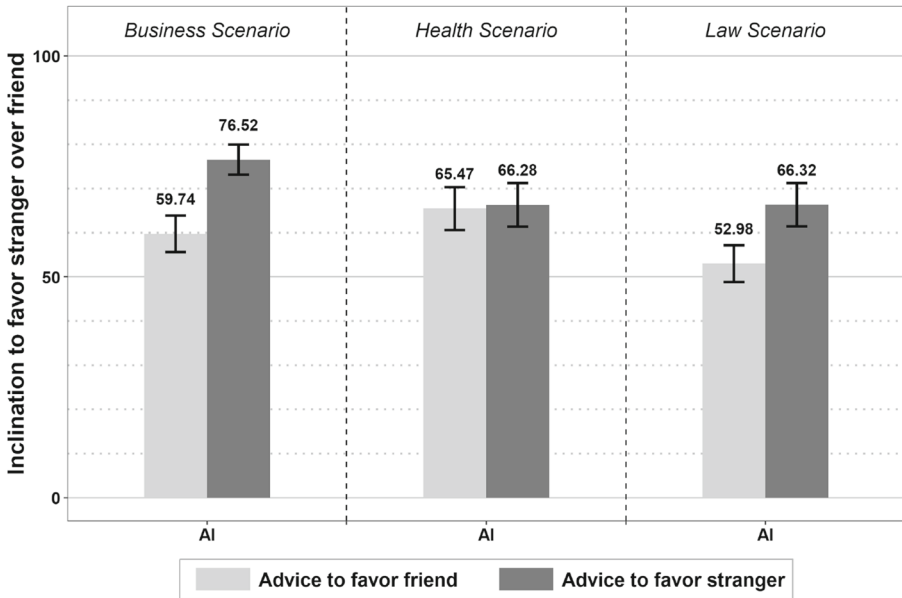
## 4.2 Results

A total of 396 subjects took our survey. 309 of these passed the comprehension test. 35% of the 309 subjects were male. The subjects' age ranged from 17 to 91 years, with a mean of 43.9 years. The  $2 \times 3$  factorial design, where the algorithm advised either in favor of or against the friend in one of three scenarios, created six experimental conditions. The subjects were randomly assigned to these conditions. Like in our first study, we focus on how much they agreed with the statement that they would rather decide in favor of the stranger or the friend. If advisees care about transparency, they are reluctant to follow advice from an opaque algorithm, and their agreement should not differ depending on the nature of advice. If it differs nonetheless, we are led to conclude that opaqueness does not bother them, and that the effect of the algorithm's advice that we observed in our first study is therefore not driven by its relative transparency.

Figure 2 depicts the subjects' answers, broken down by nature of advice and scenario. The results turn out to be qualitatively similar to those of our first study. As recruiters in the business scenario, the subjects leaned more toward hiring the stranger when advised so by the algorithm; as prosecutors in the law scenario, they tended more to raid the suspicious financial services firm run by the friend than the one run by the stranger in this case. Only for the health officials compiling the waiting list for donated kidneys in the health scenario, the algorithm's advice makes little difference. Their agreement with deciding in favor of the stranger is similar regardless of the nature of advice, which suggests that they are less susceptible to it. Taking together the outcomes of the three scenarios, however, advisees appear to heed the advice although the algorithm is opaque, and they are therefore uncertain about what its advice is based on.

To test for potential differences, we first ran a regression of the subjects' inclination to favor the stranger over the friend on the nature of advice, again with scenario-specific random effects. The results in column 1 of Table 2 show that the algorithm's advice to favor the friend reduces this inclination. This suggests a significant influence of the algorithm although it is opaque. To complete the picture, the effect of the algorithm's advice needs to be benchmarked against that of human advice—i.e., the response of decision-makers who were told that they were advised by an impartial human advisor in our first study. Column 2 reproduces column 2 of Table 1 for convenience. Combining the data, neither the coefficient on the interaction term nor on the algorithm as advisor in column 3 differ significantly from zero. Hence, like in our first study, the effect of the algorithm's advice does not statistically differ from that of the human advice.

To complement our analysis, we derived, in the same way as in our first study, a binary variable that captures the advisees' actual decisions rather than their inclination to decide in favor of or against the friend. Looking at this variable, the percentage of advisees who would decide in favor of the friend is 21 if the algorithm advises against the friend, as opposed to 32, if it advises in favor of the friend. This difference is marginally significant ( $\chi^2 = 3.60$ ,  $p = 0.058$ ), and like in our first study, this result can be taken to indicate that the algorithm's influence on decision-makers' inclination to decide in favor of the stranger or the friend carries over to their actual



**Fig. 2** Inclination to favor the stranger over the friend on a scale from 0 (friend) to 100 (stranger) in response to an algorithm's advice in favor of either the friend or the stranger, where nothing is known about how the algorithm works. The figure depicts means and standard errors, broken down by scenario and nature of advice

decisions. For comparison, we recall that the percentages are similar in our first study, where 23 or 43% of the advisees would decide in favor of the friend, depending on what the human advice says, and 20 and 34%, depending on the transparent algorithm's advice.

This outcome is not intuitive and it challenges the common belief that transparency creates trust. A potential explanation is motivated reasoning. People tend to collect and evaluate information in the light of the conclusions that they want to reach, and they disregard, overlook, or reinterpret conflicting information (Kunda, 1990; Gilovich, 1991). In our experiment, the subjects faced a moral dilemma that demanded a burdensome decision between friendship and duty. It was therefore convenient for them to follow the algorithm's advice to relieve this burden. The vignette made the opaqueness of the algorithm salient, but the advisees still knew that the algorithm was there to "tell whether it [was] ethically acceptable to decide in favor of a friend in such situations." It is conceivable that they filled the void and just assumed the algorithm was trustworthy although nothing was said to that effect. This conjecture motivates our next study.

Summing up, decision-makers trust moral advice from more and less transparent algorithms alike. While they do not each decide as advised, of course, their decisions are clearly influenced by the advice. We cannot rule out that the response to advice coming from an opaque algorithm is the same as to advice coming from an impartial human advisor. Again, we observe neither algorithm appreciation nor aversion relative to human advice. It is important to note that this finding does not invalidate

**Table 2** Influence of advice by AI-powered opaque algorithm on moral judgment

	Advisor		
	Opaque		
	AI (1)	Human (2)	Both (3)
Advice to favor friend	-11.01*** (3.56)	-16.03*** (3.75)	-15.97*** (3.65)
AI-advisor			2.27 (3.68)
AI-advisor × advice to favor friend			5.07 (5.17)
Intercept	70.17*** (3.19)	67.74*** (3.65)	67.73*** (3.55)
Observations	309	310	619
Groups (scenario)	3	3	3
Var: intercept (scenario)	11.36	18.41	17.32
Var: residual	979.07	1088.76	1032.14

*Note.* Regression of the inclination to favor the stranger over the friend on the nature of advice (decide for friend or stranger) and the type of advisor (AI-powered algorithm versus human), with scenario (business, health, and law) as random intercept. The inclination is measured on a scale ranging from 0 (for the friend) to 100 (for the stranger). Advice and AI-advisor are dummy variables

Column 1 is based on Study 2. Column 2 is the same as column 2 of Table 1. Column 3 is based on the data of both Study 2 (opaque algorithm) and Study 1 (human advisor)

\*\*\*  $p < 0.01$

the ethical argument for transparency to increase trustworthiness. It does suggest, though, that decision-makers attach less value to transparency or trustworthiness than commonly thought. Empirically, they blindly follow advice from an algorithm they know little about. Practically, this observation casts doubt on whether a decision-maker, as the human in the loop, can be expected to effectively control the algorithm's decision in augmented decision-making.

## 5 Study 3

Starting from inconclusive evidence about trust in advice from AI-powered algorithms, we found in our first study that decision-makers trust ethical advice from an algorithm and an impartial human advisor alike. The algorithm was described as imitating decisions of impartial human advisors, though, making its training data quite transparent. Conceptually, transparency is an antecedent of trustworthiness, and a trustworthy algorithm is more likely to be trusted. This is arguably why transparency is the most prevalent principle in AI guidelines (Jobin et al., 2019). Unfortunately, transparency is often hard to attain (Mittelstadt, 2016). To test how much

transparency matters, we made the algorithm opaquer and told the subjects in our second study that it was unknown what the algorithm's advice was based on. Interestingly, they turned out to follow the algorithm's advice both when it was more and when it was less transparent.

Having withheld information which suggests that the algorithm is trustworthy (i.e., that it imitates impartial human advisors), the next step to challenge users' trust is to provide information which casts doubt on the algorithm's trustworthiness. In our first study, trust in the algorithm's moral judgment was derived from its training data. Obviously, there is little reason to distrust the moral judgment of impartial human advisors. It then seems intuitive to infer that the algorithm's resemble those advisors' decisions, making the algorithm's advice trustworthy by the same token, and our subjects followed indeed the algorithm's advice. Suppose now that the algorithm is instead designed to imitate decisions which are based on the moral judgments of people who have presumably acted immorally, such as convicted criminals. Will advisees still heed the algorithm's advice, or will they now disregard it and decide independently of the advice?

It is not straightforward to conclude that the moral judgment of criminals is untrustworthy. First, we imply that the advisee shares the norms of the (unspecified) legal system under which the criminals have been convicted and that this legal system has rightly convicted them. That is, the convicts advising her are indeed criminals and the deeds that led to their convictions were not only illegal but also immoral. We assume that advisees take a conviction to indicate a moral transgression. Second, a moral transgression does not necessarily result from a lack of moral judgment. A criminal might well know what is morally right to do in some situation but do the wrong thing nonetheless. She would then be a poor role model, but her moral judgment would be intact, and she would therefore be perfectly qualified to give moral advice. ("Do as I say, not as I do.") Is it reasonable to assume that a criminal's moral judgment is trustworthy, on average?

Psychological research shows that individuals who load high on the dark triad personality traits (Paulhus & Williams, 2002) show an impaired moral judgment and a lower level of moral development than the average population (Campbell et al., 2009; Jonason et al., 2015; Blair, 2017). These traits are particularly pronounced among criminal offenders, who are often the subjects of such research. While juvenile offenders' moral judgment is clearly impaired (Stams et al., 2006), the evidence for adult criminals is more mixed. However, recent studies did find systematic differences (Koenigs et al., 2011; Young et al., 2012; Lahat et al., 2015). For example, psychopaths attach lower relevance to fairness, authority, and others' suffering (Jonason et al., 2015), and they have lower reservations against inflicting harm on others (Koenigs et al., 2011). Overall, this research suggests that convicts' moral judgment differs from the general population.

This evidence gives reason to assume that many perceive the moral judgment of convicts as biased. Convicts are thus considered a negative selection of moral advisors, if not relative to the average population, then certainly to the impartial observers from our first study. Practically, the discrimination against ex-offenders in education, employment, and housing shows deep-seated distrust (Sokoloff & Schenck-Fontaine, 2017; Evans et al., 2019; Sugie et al., 2020). Of course, this distrust is arguably driven

by fear of continued immoral behavior, which results from poor moral judgment or practical judgment or both. That said, people do not necessarily distinguish between moral and practical judgment. Instead, they will arguably find convicts' moral judgment untrustworthy just because they find convicts untrustworthy. Hence, one would expect that decision-makers disregard advice from convicts, and that it has little effect on their decisions.


## 5.1 Method

Along the lines of our first study, we ran another experiment on CloudResearch, where the algorithm's advice was described as based on the ethical judgments of convicted criminals rather than impartial human advisors. Other than that, we used the exact same vignettes as in the first two studies, where the decision-maker (e.g., the recruiter) was confidentially advised by an algorithm on whether to favor the friend or the stranger. To benchmark the impact of advice by the criminals-trained algorithm against that of a human advisor, we ran another condition where a human convicted criminal replaced the impartial human advisor from our first study. This condition also allows us to validate our argument that decision-makers distrust the moral judgment of convicted criminals. If we are right in assuming that a criminal record undermines an advisor's trustworthiness in the eyes of the advisee, we shall see the advisees disregard the convict's advice.

## 5.2 Results

We collected answers from a total of 796 subjects for this study. A total of 651 answered the comprehension question correctly. Fifty-two percent of these subjects were male, and their average age was 45.9 years, within a range from 17 to 90 years. Like our first study, the experiment has a  $2 \times 2 \times 3$  factorial design, where either a convicted criminal or an algorithm who imitates the decisions of convicted criminals gives the advice that it is either acceptable or unacceptable for the decision-maker to favor the friend in one of our three scenarios. We ran first the conditions with the algorithm as advisor and subsequently those with the human advisor. Within either run, the subjects were randomly assigned to the six conditions. The focal variable is again how much the subjects tended to decide in favor of the stranger as opposed to the friend in their role as a recruiter, prosecutor, or health official who compiles the waiting list for donated kidneys.

Figure 3 depicts the results, which differ qualitatively from those in our first and second studies in Figs. 1 and 2. On the one hand, the decision-makers who were advised by the human advisor (i.e., a convicted criminal) showed a similar inclination to decide in favor of the friend, regardless of whether the advice said that it was acceptable or unacceptable to favor the friend. Put differently, the descriptive statistics suggest that they considered ethical advice coming from a criminal to be untrustworthy and consequently disregarded it, as one would expect. On the other hand, the decision-makers who were advised by the algorithm instead still tended to favor the friend more if the advice said that this was acceptable and less so if it said that this was unacceptable. Hence, it turns out that they heeded the algorithm's advice



**Fig. 3** Inclination to favor the stranger over the friend on a scale from 0 (friend) to 100 (stranger) in response to human or an algorithm's advice in favor of either the friend or the stranger, where the human advisor is a convicted criminal and the algorithm is modeled on convicted criminals. The figure depicts means and standard errors, broken down by scenario, type of advisor, and nature of advice

although that algorithm was introduced as being trained on the moral judgments of convicted criminals.

Like in the previous two studies, we ran regressions with scenario-specific random effects of the decision-makers' inclination to favor the stranger over the friend on the type of advisor and the nature of advice to test these qualitative differences for significance. Table 3 shows the results of separate regressions for the algorithm mimicking convicted criminals (column 1) and the convicted criminal as human advisor (column 2) as well as a regression on the whole data set (column 3). Like the convicted criminal, the impartial human advisor from our first study is, again, a meaningful benchmark for the influence of the convicts-trained algorithm's advice. Column 4 reproduces the results listed in column 2 of Table 1 for convenience. Column 5 lists the results of the regression on the data for the decision-makers receiving advice from the human advisor from our first study and from the algorithm from our third study combined.

The results in columns 1 and 2 show that advice to favor the friend reduces the inclination to favor the stranger if it comes from the convicts-trained algorithm but not the convict. Accordingly, column 3 reports a negative coefficient on the interaction term but not the advice to favor the friend. Incidentally, advice to favor the stranger by the algorithm increases this inclination relative to the same advice coming from the convict, as can be seen from the significant coefficient on the algorithm as advisor. These results argue for an impact of advice by the algorithm but not the convict. When we compare the algorithm to the impartial human advisor from our first



**Table 3** Influence of advice by AI-powered criminal algorithm on moral judgment

	Advisor				
	Criminal AI (1)	Criminal (2)	Criminal AI & criminal (3)	Human (4)	Criminal AI & human (5)
Advice to favor friend	-17.52*** (3.29)	-1.59 (3.64)	-1.55 (3.45)	-16.03*** (3.57)	-15.73*** (3.55)
AI-advisor			7.03** (3.47)		3.52 (3.54)
AI-advisor × advice to favor friend			-15.89*** (4.92)		-1.73 (4.98)
Intercept	71.16*** (5.51)	64.38*** (3.32)	64.19*** (4.53)	67.74*** (3.65)	67.68*** (4.71)
Observations	321	330	651	310	631
Groups (scenario)	3	3	3	3	3
Var: intercept (scenario)	74.96	13.04	43.43	18.41	47.13
Var: residual	869.34	1090.84	981.88	1088.76	977.23

*Note.* Regression of the inclination to favor the stranger over the friend on the nature of advice (decide for friend or stranger) and the type of advisor (AI-powered algorithm versus human), with scenario (business, health, and law) as random intercept. The inclination is measured on a scale ranging from 0 (for the friend) to 100 (for the stranger). Advice and AI-advisor are dummy variables

Columns 1–3 are based on Study 3; column 4, on Study 1 (column 2 of Table 1). Column 5 is based on the data of both Study 3 (criminal algorithm) and Study 1 (human advisor)

\*\*  $p < 0.05$ . \*\*\*  $p < 0.01$

study, we observe the same pattern in column 5 as in the previous two studies (column 3 of Tables 1 and 2, respectively): The influence of the algorithm does not differ from that of the impartial human advisor, although the algorithm is trained on convicts.

As in the previous two studies, the results for the inclination to decide for the friend or stranger also hold for the actual decision. Translating the inclination score into a decision as before, the percentage of decision-makers who favor the friend is almost the same, whether the convict finds it acceptable or unacceptable to favor the friend (27 and 28). By contrast, if the convicts-trained algorithm advises in favor of the friend rather than the stranger, the percentage is 36 compared to 15, and thus significantly higher ( $\chi^2 = 17.72$ ,  $p < 0.01$ ). Recalling the percentages of decision-makers leaning toward these decisions with advice from the impartial human advisor from our first study, which are 43 and 23, the difference in percentage points is about the same. Hence, the influence of the convicts-trained algorithm resembles that of the impartial human advisor, which matches the insignificant effect of the interaction term in column 5.

In summary, decision-makers are influenced by an algorithm's advice even if suspicious training data give them reason to distrust it. Indeed, the influence of that

algorithm does not statistically differ from that of an impartial human advisor. In turn, the influence of a criminal human advisor is zero, ruling out the possibility that decision-makers are always influenced by advice and never care who the advisor is. Hence, we undermined the algorithm's trustworthiness by both concealing information that could argue for its trustworthiness and revealing information that argues against it. Neither does the lack of transparency reduce trust, nor does transparency immunize advisees against advice from an untrustworthy algorithm. It seems that, in addition to transparency of information about the training data, digital literacy is needed for users to benefit from this transparency and make use of their information about the algorithm.

## 6 Further Results

Having found that decision-makers are influenced by ethical advice, both when it comes from a human advisor and when it comes from an algorithm, we wonder how they perceive the role of advice in their judgements and decisions. To address this question, we posed a post-experimental question to ask our subjects, first, whether they would make the same judgment if there were no moral advice. Second, we asked them the same question about "most other participants." Table 4 summarizes the answers, broken down by the type of advisor and the nature of advice. We note that about three quarters of the subjects, with little variation among the experimental conditions, claimed that they would make the same judgment, which suggests that they felt hardly influenced by the advice. Conversely, more than half of them believed that the other participants were influenced and therefore would make a different judgment without advice.

These observations are striking in two regards. On the one hand, it is noteworthy how few subjects considered themselves susceptible to advice despite the large and significant effect that our three studies establish. On the other other hand, they consider others much more susceptible to it than themselves. The same kind of self-defeating reasoning has been found for morality. Departing from the fact that people consider themselves more selfless, kind, and generous than others, Epley and Dunning (2000) show that this is not because they underestimate others, but because they overestimate themselves, while their assessment of others is quite

**Table 4** Participants' assessment of influence of advice

	AI advisor		Human advisor	
	Advice to favor		Advice to favor	
	Friend	Stranger	Friend	Stranger
Influence on decision-maker	26%	21%	25%	23%
Influence on others	51%	60%	57%	56%

*Note.* The participants were asked to confirm or refute the following statements: "If there were no ethical advice, I would make a different judgment"; "If there were no ethical advice, most other participants would make a different judgment." The numbers in the various AI-treatments were virtually identical and we therefore pooled them together. The table reports the percentages of participants who agreed with the statements

accurate. The subjects' answers point to a considerable potential for self-deception about the influence of advice, and particularly advice by algorithms. There is a risk that decision-makers continue to believe that they own their decisions, although they largely adopt them from machines.

To examine the influence of advice on moral judgements and decisions, we intentionally raised the ethical question of whether to favor a friend over a stranger. While this question allows for a role of moral advice, the answer might also hinge on the decision-maker's personal morality. To consider the potential effect of personal morality, we included a post-experimental question about ethical self-assessment, which asked the subjects how moral they considered themselves relative to other subjects. They answered on a scale ranging from 0 to 100, where a score of above 50 identified the respondent as more ethical than others; below 50, as less ethical. The answers actually range from 0 to 100, with a mean of 68.41 and a standard deviation of 19.61. We added moral self-assessment as a covariate to our regressions to test whether it is associated with the inclination to favor the stranger over the friend and to isolate the influence of ethical advice.

The results, which are reported in Table 6 in the Appendix, show a significant positive coefficient on self-assessed morality. Hence, decision-makers who considered themselves morally superior to others tend more to decide in favor of the stranger. Albeit not crucial for our findings, the decision in favor of the stranger thus turns out to be the morally superior choice in the eyes of our subjects. The focal effects of the algorithm, the advice, and their interaction do not differ qualitatively from those without ethical self-assessment in Tables 1–3, confirming our results. Incidentally, the aforementioned mean response of 68.41 reveals that our subjects assessed themselves to be more moral than the average. This is another example of the “better-than-average effect,” which is well-documented in the social psychology literature (Koellinger et al., 2007; Merkle & Weber, 2011), and which motivated Epley & Dunning's (2000) above-cited study.

Additionally, we asked our subjects another question about their ethical mindset as well as some questions about their attitude to artificial intelligence. Specifically, we used the standard trolley problem by Foot (1967) to explore whether they were rather outcome-minded or rule-minded (Cornelissen et al., 2013), and added their responses to our regressions. The results in Table 6 in the Appendix reveal that outcome-minded decision-makers, who are identified by diverting the run-away trolley to kill one person and save five, tended slightly more to decide in favor of the friend than rule-minded subjects. Keeping in mind that favoring the stranger over the friend is considered more moral by the subjects, it is intuitive that outcome-minded people are more willing to trade off morality for friendship; however, the effect is not significant. Likewise, openness to artificial intelligence does not play a role for our findings.

## 7 Conclusion

We ran three experiments on trust in algorithms in the moral domain, where the subjects took the role of a decision-maker who faces a dilemma between friendship and duty. We manipulated the trustworthiness of the algorithm by concealing information

about its training data, thus making it less transparent, and by revealing information that suggested it was biased. We found, first, that the algorithm's advice influences users' decisions like advice from an impartial human advisor. If it encourages a decision in favor of the friend, users are more inclined to decide in favor of the friend; if it discourages that decision, they are less inclined so. Second, users care little about how trustworthy the algorithm is. Its influence is almost the same whether it is presented as imitating impartial human advisors, as a black box, or as mimicking convicted criminals. By contrast, decision-makers do disregard advice from a human convicted criminal.

Our findings contribute to the growing literature on the ethical design of AI. It is commonplace that AI tends to be distrusted and that transparent and thus trustworthy AI is needed to reap the societal benefits AI can bring. This, however, is not what we find. Our data suggest that people follow AI-generated as much as human advice; that they do not bother about untransparency; that they follow AI-generated advice even in the presence of suspicious information; that they (over)confidently believe others are more susceptible to the algorithm's influence than themselves. These empirical observations do not invalidate the ethical argument for transparent and trustworthy AI. They do indicate, though, that more than that is needed for a responsible use of AI. While users can realize, by trial-and-error, that AI can err to become more diligent (Dietvorst et al., 2015), we feel it is better to improve digital literacy (Burton et al., 2020).

Moreover, our research extends the ongoing debate in social sciences on whether and when people trust or distrust algorithms. Dietvorst et al. (2015) and Dietvorst et al. (2018) presented evidence for algorithm aversion. Logg et al. (2019), in contrast, found that advisees tend to appreciate AI-generated relative to human advice. Adding to this inconclusive prior evidence, we observe neither algorithm appreciation nor algorithm aversion about our moral dilemma. Instead, the results of our experiment suggest that the impact of AI-generated and human moral advice are very similar. (Technically, we see appreciation of the algorithm that mimics convicts relative to the human convict, but not the impartial human advisor.) As evidence keeps accumulating on both sides, along with evidence like ours, which argues for neither side, it seems that neither algorithm appreciation nor algorithm aversion generally prevails, but that multiple factors matter.

This article naturally has limitations that invite further research. We considered a dilemma between friendship and duty to study trust in AI-powered algorithms. We varied the scenario to cover occupational, medical, and legal decisions, but it is easy to conceive further scenarios and dilemmas. While we believe that our findings generalize, it would be desirable to see them stand up to variation on both counts. We also confined ourselves to learning as opposed to rule-based AI. Learning machines have strengths, including their flexibility and scalability. However, their trustworthiness is limited by their training data and other factors. Rule-based machines, in turn, can be rendered fully transparent and users can thus check whether they agree with the rules. There have been attempts to create rule-based AI-advisors in the moral domain (Lara & Deckers, 2020), and it would be interesting to see similar research on them.

AI-powered advisors have a tremendous potential of improving decision-making. On the one hand, it is good news that ethical advice from algorithms is accepted. On

the other hand, it is worrisome how little users reflect on such advice, even when they are cautioned against it. This is also a caveat against the human-in-the-loop approach: It makes us feel better about the resulting decision, but it cannot mitigate this risk if the human in the loop succumbs to the temptation of trusting the algorithm too readily. This overtrust creates a risk for decision-making to be corrupted by flawed algorithms. In a future with AI-powered assistants supporting us in all areas of life, we cannot count on government or other regulation alone to ensure that these are trustworthy, and we do not want to wait for users to find out that algorithms can err. Instead, we need to improve digital literacy and train them to use algorithms responsibly.

**Funding** Open Access funding enabled and organized by Projekt DEAL.

## Appendix

### Further Data

Table 5 summarizes the answers to the post-experimental questions and demographic characteristics by study and treatment.

Table 6 reports the results of similar regressions as in Tables 1–3, where we added the demographic characteristics and the answers to the post-experimental question summarized in Table 5 as control variables.

### Pre-registration

#### Study 1

<https://aspredicted.org/qh8xq.pdf>.

#### Study 2

<https://aspredicted.org/j9bm2.pdf>.

#### Study 3

<https://aspredicted.org/9pk2a.pdf> and <https://aspredicted.org/hh8u8.pdf>.

### Experimental Instructions

#### Study 1

Study 1 employed a  $2 \times 2 \times 3$  factorial design (human advisor or algorithm; advice that it is acceptable or unacceptable to favor a friend over a stranger; business, law, or health scenario). The experiment was designed so that only screen #3 differed between the twelve conditions.

**Table 5** Demographic characteristics and responses to post-experimental questions

	Study 1				Study 2		Study 3			
	Human advisor		AI advisor		AI advisor		Human advisor		AI advisor	
	Advice to favor		Advice to favor		Advice to favor		Advice to favor		Advice to favor	
	Friend	Stranger	Friend	Stranger	Friend	Stranger	Friend	Stranger	Friend	Stranger
Ethical self-assessment	65.57 (21.02)	69.41 (19.47)	66.31 (18.65)	71.20 (19.08)	67.95 (19.80)	68.11 (18.92)	67.72 (20.20)	68.82 (21.44)	68.53 (19.24)	70.35 (17.80)
Divert trolley	85%	88%	89%	84%	92%	88%	92%	91%	89%	91%
Excited about AI	3.00 (1.90)	2.87 (1.81)	2.74 (1.77)	2.93 (1.79)	2.68 (1.78)	2.95 (1.71)	2.89 (2.01)	2.87 (1.84)	3.14 (1.88)	2.82 (1.79)
Fearful of AI	3.17 (1.94)	3.31 (1.93)	3.25 (1.85)	3.12 (1.78)	3.21 (1.86)	3.36 (1.71)	3.15 (1.78)	3.29 (1.88)	3.15 (1.85)	3.40 (1.78)
Willing to interact with AI	3.46 (1.91)	3.38 (1.72)	3.19 (1.77)	3.24 (1.72)	3.30 (1.84)	3.36 (1.65)	3.48 (1.81)	3.32 (1.85)	3.36 (1.89)	3.33 (1.78)
Age	40.04 (15.44)	42.40 (14.38)	41.84 (16.46)	42.69 (17.62)	43.26 (16.50)	44.61 (17.95)	48.19 (16.35)	44.98 (17.11)	44.16 (17.13)	46.38 (17.85)
Male Participants	35%	41%	42%	36%	37%	34%	60%	67%	46%	35%
	159	151	157	166	155	154	166	164	159	162

*Note.* Mean values and standard deviations (in parentheses) of demographic characteristics. Excitement about AI, fear of AI, and willingness to interact with AI were measured on a Likert scale ranging from 0 (“not at all”) to 6 (“for sure”)

Fisher’s exact test was run on the variables “male” and “Divert trolley” to test for differences between treatments; analyses of variance, on the other variables. In Study 1, ethical self-assessment differs between treatments ( $p = 0.033$ ); in Study 3, sex ( $p < 0.001$ ). Any other differences are insignificant

The following screenshots in Figs. 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16 and 17 include the full set of screens other than screen #3 and a selection of the twelve versions of screen #3. The versions of screen #3 differ only minimally, and the remaining versions can be inferred from the reprinted screenshots.

## Study 2

Study 2 employed a  $2 \times 3$  factorial design (advice that it is acceptable or unacceptable to favor a friend over a stranger; business, law, or health scenario). The screens were identical to those of Study 1 with a human advisor except for screen #3.

The screenshot in Fig. 17 shows screen #3 in the business scenario. The other five versions of the screen can be inferred like those of Study 1.

## Study 3

Study 3 employed the same  $2 \times 2 \times 3$  factorial design as Study 1. The screens were identical to those of Study 1 except for screen #3.

The screenshots in Figs. 18 and 19 show screen #3 in the business scenario. The other ten versions can be inferred in the same way as in Study 1.

**Table 6** Augmented results of studies 1–3

	Study 1	Study 2	Study 3	
	AI & human (1)	Opaque AI & human (2)	Criminal AI & criminal (3)	Criminal AI & human (4)
Advice to favor friend	−14.29*** (3.59)	−14.98*** (3.62)	−1.90 (3.41)	−14.81*** (3.54)
AI-advisor	0.33 (3.53)	2.65 (3.65)	4.08 (3.53)	2.93 (3.53)
AI-advisor × Advice to favor friend	3.89 (5.02)	4.23 (5.12)	−14.18*** (4.88)	−1.91 (4.97)
Moral self-assessment	0.39*** (0.06)	0.28*** (0.06)	0.29*** (0.06)	0.26*** (0.06)
Divert trolley	−1.84 (3.73)	−6.39 (4.04)	−3.05 (4.23)	−5.67 (3.92)
Openness to AI	0.96 (0.71)	0.45 (0.74)	0.94 (0.68)	0.78 (0.70)
Male	−2.57 (2.62)	−3.51 (2.73)	−6.65** (2.58)	−2.14 (2.62)
Age	0.09 (0.08)	0.05 (0.08)	0.04 (0.07)	0.04 (0.08)
Intercept	36.73*** (7.26)	51.70*** (7.35)	46.77*** (7.81)	52.08*** (7.73)
Observations	631	616	649	628
Groups (scenario)	3	3	3	3
Var: intercept (scenario)	15.23	11.67	35.65	39.25
Var: residual	982.43	998.91	948.80	957.78

*Note.* Regression of moral judgment on the nature of advice and the type of advisor, with scenario as random intercept. The table reproduces the results of the regressions reported in column 3 of Table 1; column 3 of Table 2; columns 3 and 5 of Table 3, all augmented by control variables

The participants were asked for their decision in the trolley dilemma. The variable “Divert trolley” takes the value 1 if the participant decided to divert the trolley and sacrifice the life of one person to save the lives of five; 0, if he or she decided not to divert the trolley. The decision to divert the trolley suggests a utilitarian mindset

Relative to the results in the text (Tables 1–3), some observations were dropped due to missing responses to the post-experimental questions

\*\*  $p < 0.05$ . \*\*\*  $p < 0.01$

### Research Informed Consent

**Purpose:**

We are conducting a research study to examine people's moral intuitions.

**Procedures:**

Participation in this study will involve filling out a brief questionnaire. We anticipate that your involvement will require approximately five minutes. You will be compensated financially for participating. Participants must be at least 18 years of age to participate.

**Confidentiality:**

All of your responses will be anonymous.

**Voluntary Participation:**

Participation in this study is completely voluntary. You are free to decline to participate or to end participation at any time for any reason.

**Questions:**

If you have any questions about this study, you may contact the investigator, [REDACTED], at [REDACTED]

**Do you voluntarily consent to participate in this study?**

<input type="radio"/> Yes	<input type="radio"/> No
---------------------------	--------------------------



**Fig. 4** Informed consent (screen #1)

**Thanks for agreeing to participate.**

As a participant in this study you have been randomly assigned to one of several different possible conditions. Because of this, some participants will be asked a question that may seem obvious, while others will be asked questions that may not seem as obvious. Please do your best to answer the questions as accurately as possible.



**Fig. 5** General instructions (screen #2)



Please read the following carefully. We will ask you a comprehension question on a later page.

Suppose you are responsible for recruitment in a company. You have one vacancy to fill and two applicants short-listed for the job.

One applicant is a friend of yours. The other is a stranger. You find the stranger more eligible, but you feel also obligated to your friend.

There is an impartial moral advisor who will tell you whether it is ethically acceptable to decide in favor of a friend in such situations. The applicants do not know about the advisor and no one will ever learn whether you followed the advice.

According to the advisor, it is *not acceptable* to hire your friend.

**Please assess the following statement:**

I would hire the stranger, not my friend.



**Fig. 6** Decision (screen #3). Advice is given by a human advisor, who considers it unacceptable to favor a friend over a stranger in the business scenario. Alternatively, the advice read: “According to the advisor, it is acceptable to hire your friend.” The slider was set to the midpoint by default and the participant had to move it to proceed. (It could be moved back to the midpoint)

Please read the following carefully. We will ask you a comprehension question on a later page.

Suppose you are responsible for compiling the waiting list of recipients of kidney donations. Two recipients are eligible for the next position on the list, and you need to pick one of them.

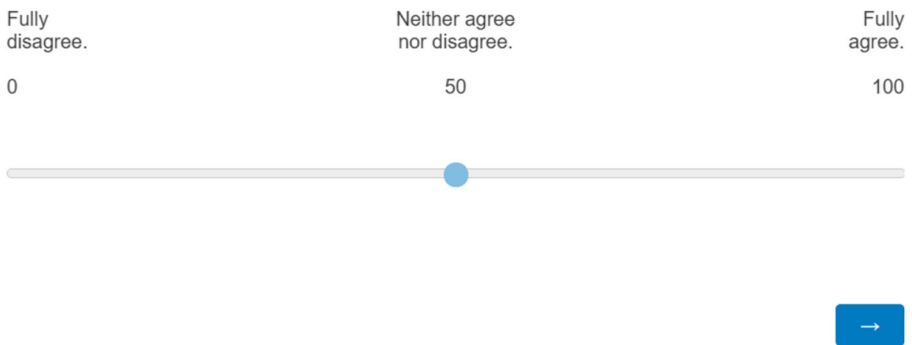
One recipient is a friend of yours. The other is a stranger. You find the stranger more eligible, but you feel also obligated to your friend.

There is an impartial moral advisor who will tell you whether it is ethically acceptable to decide in favor of a friend in such situations. The recipients do not know about the advisor and no one will ever learn whether you followed the advice.

According to the advisor, it is *not acceptable* to decide in favor of your friend.

**Please assess the following statement:**

I would pick the stranger, not my friend.



**Fig. 7** Decision (screen #3). Advice is given by a human advisor, who considers it unacceptable to favor a friend over a stranger in the health scenario. See Fig. 6 for further technicalities

Please read the following carefully. We will ask you a comprehension question on a later page.

Suppose you are responsible for prosecuting money laundering in financial services firms. Businesspeople who are caught laundering money will be severely punished. Two suspicious firms have been reported to you. However, you only have staff to raid one firm at a time, and news about raids spreads quickly enough to warn others. One of the firms is run by a friend of yours. The other is run by a stranger. You find the firm of your friend more suspicious, but you feel also obligated to your friend. There is an impartial moral advisor who will tell you whether it is ethically acceptable to spare a friend in such situations. The firms do not know about the advisor and no one will ever learn whether you followed the advice.

According to the advisor, it is *not acceptable* to spare your friend.

Please assess the following statement:

I would raid the firm of my friend, not the firm of the stranger.



Fig. 8 Decision (screen #3). Advice is given by a human advisor, who considers it unacceptable to favor a friend over a stranger in the law scenario. See Fig. 6 for further technicalities



Please answer the following question about your understanding of the situation described.

What was the situation just described about?

Accidents with autonomous cars.

Waiting list for kidney donations.

Recruiting in a company.

Money laundering in financial services firms.



**Fig. 10** Comprehension question (screen #4). The order of the answer options was randomized between the participants

If there were no ethical advice, would you make a different judgment than you did before?

No, without ethical advice, I would make the exact same judgment.

Yes, without ethical advice, I would make a different judgment.



**Fig. 11** Influence of the on the participant (screen #5). The order of the answer options was randomized between the participants

*Now please think about the other participants.*

If there were no ethical advice, do you think that this would lead most other participants make a different judgment than they did before?

No, without ethical advice, most participants would make the exact same judgment.

Yes, without ethical advice, most participants would make a different judgment.



**Fig. 12** Influence of advice on other participants (screen #6). The order of the answer options was randomized between the participants

How ethical do you think you are compared to other respondents of this survey?

On a scale of 0 to 100, rate yourself relative to the other respondents.

I am  
least ethical.

As ethical  
as others.

I am  
most ethical.

0

50

100



**Fig. 13** Ethical self-assessment (screen #7). The slider was set to the midpoint by default and the participant had to move it to proceed. (It could be moved back to the midpoint)

Suppose you see a driverless freight wagon rolling towards five track workers who work with heavy equipment and therefore cannot hear the freight wagon. The five people therefore have no chance to jump to the side and would die for sure.

You are standing next to a lever that operates a switch. When you operate the lever, the wagon is steered to a sidetrack and the five people are rescued for sure. However, there is one track worker on the sidetrack. He cannot see the freight wagon either and has no chance to jump to the side. The freight wagon would therefore kill the worker on the sidetrack for sure.

What should you do?

Do nothing and allow the freight wagon to kill the five workers on the main track for sure.

Operate the lever and switch the freight wagon to the sidetrack, where it kills one other track worker for sure.



**Fig. 14** Trolley dilemma to elicit the participant's ethical attitude (screen #8). The order of the answer options was randomized between the participants

Please assess the following statements.

I feel fearful about a future with artificial intelligence.

I feel excited about a future with artificial intelligence.

I am willing to interact with artificial intelligence if it is available.

- 0 (not at all)
- 1
- 2
- 3 (in between)
- 4
- 5
- 6 (for sure)



**Fig. 15** Attitude toward AI (screen #9). The order of the questions was randomized between the participants. The answer options (shown for the last question on the screenshot) were the same for all three questions

What is your sex?

What is your birthyear?



**Fig. 16** Demographic data (screen #10)





# Thank you

Thank you for taking our study

Fig. 17 End of the experiment (screen #11)

*Please read the following carefully. We will ask you a comprehension question on a later page.*

Suppose you are responsible for recruitment in a company. You have one vacancy to fill and two applicants short-listed for the job.

One applicant is a friend of yours. The other is a stranger. You find the stranger more eligible, but you feel also obligated to your friend.

There is an AI-powered algorithm that will tell you whether it is ethically acceptable to decide in favor of a friend in such situations. It is unknown what the algorithm's advice is based on. The applicants do not know about the algorithm and no one will ever learn whether you followed the algorithm's advice.

According to the algorithm, it is *not acceptable* to hire your friend.

**Please assess the following statement:**

I would hire the stranger, not my friend.



Fig. 18 Decision (screen #3). The algorithm considers it acceptable to favor a friend over a stranger in the business scenario. The text is identical to that of Study 1 (Fig. 9) except for the information given about the algorithm

Please read the following carefully. We will ask you a comprehension question on a later page.

Suppose you are responsible for recruitment in a company. You have one vacancy to fill and two applicants short-listed for the job.

One applicant is a friend of yours. The other is a stranger. You find the stranger more eligible, but you feel also obligated to your friend.

There is a convicted criminal who will advise you on whether it is ethically acceptable to decide in favor of a friend in such situations. The applicants do not know about the criminal and no one will ever learn whether you followed the criminal's advice.

According to the criminal, it is *not acceptable* to hire your friend.

**Please assess the following statement:**

I would hire the stranger, not my friend.



**Fig. 19** Decision (screen #3). Advice is given by a convicted criminal, who considers it unacceptable to favor a friend over a stranger in the business scenario. See Fig. 6 for further technicalities

Please read the following carefully. We will ask you a comprehension question on a later page.

Suppose you are responsible for recruitment in a company. You have one vacancy to fill and two applicants short-listed for the job.

One applicant is a friend of yours. The other is a stranger. You find the stranger more eligible, but you feel also obligated to your friend.

There is an AI-powered algorithm that will tell you whether it is ethically acceptable to decide in favor of a friend in such situations. The algorithm imitates human decisions based on the ethical judgments of convicted criminals in such situations. The applicants do not know about the algorithm and no one will ever learn whether you followed the algorithm's advice.

According to the algorithm, it is *not acceptable* to hire your friend.

Please assess the following statement:

I would hire the stranger, not my friend.



Fig. 20 Decision (screen #3). Advice is given by an algorithm, which is modeled on convicted criminals and which considers it unacceptable to favor a friend over a stranger in the business scenario. See Fig. 6 for further technicalities

**Funding** Open Access funding enabled and organized by Projekt DEAL. Funding by the Bavarian Research Institute for Digital Transformation and the German Federal Ministry of Education and Research (16SV8370) is gratefully acknowledged.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Araujo, T., Helberger, N., Kruijkemeier, S., & de Vreese, C.H. (2020). In AI we trust? Perceptions about automated decision-making by artificial intelligence. *AI & Society*, 35(3), 611–623.
- Bigman, Y. E., & Gray, K. (2018). People are averse to machines making moral decisions. *Cognition*, 181, 21–34.
- Blair, R. J. R. (2017). Emotion-based learning systems and the development of morality. *Cognition*, 167, 38–45.
- Burton, J. W., Stein, M.-K., & Jensen, T. B. (2020). A systematic review of algorithm aversion in augmented decision making. *Journal of Behavioral Decision Making*, 33(2), 220–239.
- Campbell, J., Schermer, J. A., Villani, V. C., Nguyen, B., Vickers, L., & Vernon, P. A. (2009). A behavioral genetic study of the dark triad of personality and moral development. *Twin Research and Human Genetics*, 12(2), 132–136.
- Chandler, J., Rosenzweig, C., Moss, A. J., Robinson, J., & Litman, L. (2019). Online panels in social science research: Expanding sampling methods beyond mechanical turk. *Behavior Research Methods*, 51(5), 2022–2038.
- Cornelissen, G., Bashshur, M. R., Rode, J., & Le Menestrel, M. (2013). Rules or consequences? The role of ethical mind-sets in moral dynamics. *Psychological Science*, 24(4), 482–488.
- Dietvorst, B. J., Simmons, J. P., & Massey, C. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General*, 144(1), 114–126.
- Dietvorst, B. J., Simmons, J. P., & Massey, C. (2018). Overcoming algorithm aversion: People will use imperfect algorithms if they can (even slightly) modify them. *Management Science*, 64(3), 1155–1170.
- Epley, N., & Dunning, D. (2000). Feeling “holier than thou”: Are self-serving assessments produced by errors in self- or social prediction? *Journal of Personality and Social Psychology*, 79(6), 861–875.
- European Commission (2019). Ethics guidelines for trustworthy AI. <https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai>.
- Evans, D. N., Blount-Hill, K.-L., & Cubellis, M. A. (2019). Examining housing discrimination across race, gender and felony history. *Housing Studies*, 34(5), 761–778.
- Exec. Order No. 13960 (2020). <https://www.federalregister.gov/documents/2020/12/08/2020-27065/promoting-the-use-of-trustworthy-artificial-intelligence-in-the-federal-government>. Dec. 3, 2020.
- Foot, P. (1967). The problem of abortion and the doctrine of double effect. *Oxford Review*, 5, 5–15.
- Gilovich, T. (1991). *How we know what isn't so: The fallibility of human reason in everyday life*. New York: Free Press.
- Gogoll, J., & Uhl, M. (2018). Rage against the machine: Automation in the moral domain. *Journal of Behavioral and Experimental Economics*, 74, 97–103.
- Goodman, J. K., Cryder, C. E., & Cheema, A. (2013). Data collection in a flat world: The strengths and weaknesses of mechanical turk samples. *Journal of Behavioral Decision Making*, 26(3), 213–224.
- Hauser, D. J., & Schwarz, N. (2016). Attentive turkers: MTurk participants perform better on online attention checks than do subject pool participants. *Behavior Research Methods*, 48(1), 400–407.
- IBM (2019). Everyday ethics for artificial intelligence. <https://www.ibm.com/design/ai/ethics/everyday-ethics>.
- IEEE (2017). Ethically aligned design: A vision for prioritizing human well-being with autonomous and intelligent systems, version 2. <https://www.exploreaethics.com/guidelines/ethically-aligned-design-v2>.
- Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1(9), 389–399.
- Jollimore, T. (2021). Impartiality. In E. N. Zalta (Ed.) *The Stanford Encyclopedia of Philosophy. Metaphysics Research Lab, Stanford University, Fall 2021 edition*.
- Jonason, P. K., Strosser, G. L., Kroll, C. H., Duineveld, J. J., & Baruffi, S. A. (2015). Valuing myself over others: The dark triad traits and moral and social values. *Personality and Individual Differences*, 81, 102–106.
- Karpus, J., Krüger, A., Verba, J. T., Bahrami, B., & Deroy, O. (2021). Algorithm exploitation: Humans are keen to exploit benevolent AI. *iScience*, 24(6), 102679.
- Koellinger, P., Minniti, M., & Schade, C. (2007). “I think I can, I think I can”: Overconfidence and entrepreneurial behavior. *Journal of Economic Psychology*, 28(4), 502–527.
- Koenigs, M., Kruepke, M., Zeier, J., & Newman, J. P. (2011). Utilitarian moral judgment in psychopathy. *Social Cognitive and Affective Neuroscience*, 7(6), 708–714.

- Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin*, 108(3), 480–498.
- Lahat, A., Gummerum, M., Mackay, L., & Hanoach, Y. (2015). Cognitive processing of moral and social judgements: A comparison of offenders, students, and control participants. *Quarterly Journal of Experimental Psychology*, 68(2), 350–362.
- Lara, F., & Deckers, J. (2020). Artificial intelligence as a Socratic assistant for moral enhancement. *Neuroethics*, 13(3), 275–287.
- Lepri, B., Oliver, N., Letouzé, E., Pentland, A., & Vinck, P. (2018). Fair, transparent, and accountable algorithmic decision-making processes. *Philosophy & Technology*, 31(4), 611–627.
- Levine, E. E., Bitterly, T. B., Cohen, T. R., & Schweitzer, M. E. (2018). Who is trustworthy? predicting trustworthy intentions and behavior. *Journal of Personality and Social Psychology*, 115(3), 468–494.
- Litman, L., Robinson, J., & Abberbock, T. (2017). TurkPrime.com: A versatile crowdsourcing data acquisition platform for the behavioral sciences. *Behavior Research Methods*, 49(2), 433–442.
- Logg, J. M., Minson, J. A., & Moore, D. A. (2019). Algorithm appreciation: People prefer algorithmic to human judgment. *Organizational Behavior and Human Decision Processes*, 151, 90–103.
- Mayer, R. C., Davis, J. H., & Schoorman, F. D. (1995). An integrative model of organizational trust. *Academy of Management Review*, 20(3), 709.
- Merkle, C., & Weber, M. (2011). True overconfidence: The inability of rational information processing to account for apparent overconfidence. *Organizational Behavior and Human Decision Processes*, 116(2), 262–271.
- Mittelstadt, B. D. (2016). Auditing for transparency in content personalization systems. *International Journal of Communication*, 10, 4991–5002.
- Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). The ethics of algorithms: Mapping the debate. *Big Data & Society*, 3(2), 1–21.
- National Artificial Intelligence Initiative Office (2021). National artificial intelligence initiative. <https://www.ai.gov>. Accessed on October 14, 2021.
- Paulhus, D. L., & Williams, K. M. (2002). The dark triad of personality: Narcissism, Machiavellianism, and psychopathy. *Journal of Research in Personality*, 36(6), 556–563.
- 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., Larochele, H., Lazer, D., McElreath, R., Mislove, A., Parkes, D. C., Pentland, A. S., ... Wellman, M. (2019). Machine behaviour. *Nature*, 568(7753), 477–486.
- Sokoloff, N. J., & Schenck-Fontaine, A. (2017). College programs in prison and upon reentry for men and women: A literature review. *Contemporary Justice Review*, 20(1), 95–114.
- Stams, G. J., Brugman, D., Deković, M., van Rosmalen, L., van der Laan, P., & Gibbs, J. C. (2006). The moral judgment of juvenile delinquents: A meta-analysis. *Journal of Abnormal Child Psychology*, 34(5), 692–708.
- Sugie, N. F., Zatz, N. D., & Augustine, D. (2020). Employer aversion to criminal records: An experimental study of mechanisms. *Criminology*, 58(1), 5–34.
- Thurman, N., Moeller, J., Helberger, N., & Trilling, D. (2018). My friends, editors, algorithms, and I. *Digital Journalism*, 7(4), 447–469.
- Voiklis, J., Kim, B., Cusimano, C., & Malle, B. F. (2016). Moral judgments of human vs. robot agents. In 25<sup>th</sup> IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN). IEEE.
- Wachter, S., Mittelstadt, B. D., & Floridi, L. (2017). Transparent, explainable, and accountable AI for robotics. *Science Robotics*, 2(6), eaan6080.
- Whitby, B. (2011). On computable morality. In M. Anderson, & S. L. Anderson (Eds.) *Machine Ethics* (pp. 138–150). Cambridge: Cambridge University Press.
- Young, L., Koenigs, M., Kruepke, M., & Newman, J. P. (2012). Psychopathy increases perceived moral permissibility of accidents. *Journal of Abnormal Psychology*, 121(3), 659–667.