Innocence over utilitarianism - Heightened Moral Standards for Robots in Rescue

Dilemmas

Jukka Sundvall ^{*a}, Marianna Drosinou ^b, Ivar Hannikainen ^d, Kaisa Elovaara ^a, Juho Halonen ^a, Volo Herzon ^b, Robin Kopecký ^{e, f}, Michaela Jirout Košová ^f, Mika Koverola ^a, Anton Kunnari ^{a, b}, Silva Perander ^c, Teemu Saikkonen ^g, Jussi Palomäki ^a, Michael Laakasuo ^a

Affiliations

* corresponding author; jukka.sundvall@helsinki.fi

^a Department of Digital Humanities, Cognitive Science, Faculty of Arts, University of Helsinki, Helsinki, Finland

^b Department of Psychology and Logopedics, Faculty of Medicine, University of Helsinki, Helsinki, Finland

^c Department of Computer Science, Faculty of Science, University of Helsinki, Helsinki, Finland

^d Department of Philosophy I, Faculty of Psychology, Cartuja Campus, Universidad de Granada, Spain

^e Institute of Philosophy, Czech Academy of Sciences, Prague, Czechia

^f Department of Philosophy and History of Science, Faculty of Science, Charles University, Prague, Czechia

^g Zoological Museum, Biodiversity Unit, University of Turku, Turku, Finland

HEIGHTENED MORAL STANDARDS FOR ROBOTS

Contact details of corresponding author

Present address: University of Helsinki, Faculty of Arts, Vuorikatu 3A, H541, Helsinki, 00014, Finland

Email: jukka.sundvall@helsinki.fi

Email of other authors

MD: maria-anna.drosinou@helsinki.fi; IH: ivar@ugr.es; KE: kaisa.elovaara@helsinki.fi; JH: juho.halonen@helsinki.fi; VH: yolo.herzon@helsinki.fi; RK: robin.kopecky@icloud.com; MJK: michaela.kosa@gmail.com; MK: mika.koverola@helsinki.fi; AK: anton.kunnari@helsinki.fi; SP: silva.perander@helsinki.fi; TS: teemu.saikkonen@helsinki.fi; JP: jussi.palomaki@helsinki.fi; ML: michael.laakasuo@helsinki.fi

Acknowledgements: Anke Haas and Noora Lehtonen helped with the original Study 1 design and investigation.

Conflict of Interest statement:

The authors declare no conflicts of interest.

Funding statement:

The authors have received funding from Jane & Aatos Erkko Foundation (grant number 170112), the Academy of Finland (grant number 323207) and from the Weisell Foundation.

Ethics statement:

The protocols and materials were approved by the University of Helsinki Ethical Review Board in the Humanities and Social Behavioral Sciences. The research followed national ethics guidelines. The results are reported honestly. This is an original article. Authorship reflects individuals' contributions.

Data archiving and sharing statement

HEIGHTENED MORAL STANDARDS FOR ROBOTS

The datasets, data codebooks and analysis scripts used are available at

https://osf.io/hm5yf/?view_only=3aa42d9ee7ae4d67bc88d2f9470f237e.

Note the following details about the data made available:

Studies 1, 2a, 2b, 2c: Demographic variables such as age are aggregated to make data more anonymous. Covariate measures that may be reported elsewhere are not included. Does not include data on variables not covered in manuscript.

Study 3: Data with identifiers is not available. Analyses are based on publicly available aggregate numbers, as highlighted in the manuscript.

Studies 4a, 4b, 5: Some demographic variables such as age are aggregated. Some demographic variables collected such as employment are not included in the shared data set due to being irrelevant to analyses and potential identifiers.

**Innocence over utilitarianism - Heightened Moral Standards for Robots in Rescue
Dilemmas**

Abstract

Research in moral psychology has found that robots, more than humans, are expected to make utilitarian decisions. This expectation is found specifically when contrasting utilitarian action to deontological inaction. In a series of eight experiments (total $N = 3752$), we compared judgments about robots' and humans' decisions in a rescue dilemma with no possibility of deontological inaction. A robot's decision to rescue an innocent victim of an accident was judged more positively than the decision to rescue two people culpable for the accident (Studies 1-2b). This pattern repeated in a large-scale web survey (Study 3, $N = \sim 19000$) and reversed when all victims were equally culpable/innocent (Study 5). Differences in judgments about humans' and robots' decisions were largest for norm-violating decisions. In sum, robots are not always expected to make utilitarian decisions, and their decisions are judged differently from those of humans based on other moral standards as well.

Keywords: Utilitarianism, Folk ethics, Rescue robotics, Folk justice, Moral dilemma

In the film, *I, Robot*, the protagonist, Detective Del Spooner, has grown distrustful of the robots ubiquitous in society. Spooner's aversion to robots is not caused by any hostile episode. Quite the contrary: a robot had *saved* Spooner's life, rescuing him from a sinking car. In saving Spooner, the robot chose to leave a 12-year-old child in the car to drown, based on the child's lower likelihood of survival.

Fictional as it may be, this scenario raises important questions concerning the role of autonomous artificial agents in society today. Artificial intelligences (AIs; software solutions with independent decision-making capacities) already make decisions that impact human lives, for instance, in the form of parole decisions (O'Neil, 2016). Robots – machines equipped with AI (i.e., decision-making software that has capacities to plan and execute goals and striving towards goals), built to serve different functions - may someday make rescue decisions, too. What expectations do we have of such robots, and the way they will make life-saving decisions? If we are to avoid people developing a Spooner-like disregard for robotic agents, what are the moral standards that these robots should support? As technological progress marches onwards (rescue robotics is a growing field; Walker, 2019), it becomes increasingly relevant to examine how the public reacts to automated decisions.

Robots and AIs as Moral Agents?

Robots and AIs are anticipated to play a growing role in high-risk occupations: helping to perform surgeries, aiding in military and policing operations, disaster relief and emergency response teams (Marcus & Davis, 2019). The transition toward an automated workforce in these areas promises to reduce casualties and human error. Furthermore, developments in AI may lead to robots not being expected to simply *enact* human decisions, but often to autonomously determine which course of action to pursue (Marcus & Davis, 2019; Mitchell, 2019). Rescue

robots, currently mostly controlled by human users, could guide themselves (Walker, 2019); military robots could decide on when and where to strike (although it is uncertain whether this will ever be allowed; see Campaign to Stop Killer Robots, 2019; Davison, 2017).

By autonomously determining which course of action to pursue in contexts where their actions affect human lives, robots and AIs would be making decisions vested with moral relevance (Bigman & Gray, 2018; Awad et al., 2018; Bigman, Waytz, Alterovitz & Gray, 2019; Wallach & Allen, 2008). We will use the term “artificial agent” (after Wallach & Allen, 2008) to refer collectively to robots and AIs in roles that allow them to make decisions that are morally relevant to humans.

Are people even willing to consider artificial agents as eligible agents for making moral decisions? Bigman and Gray (2018) found that people generally consider robots as less appropriate moral decision-makers than humans, regardless of the type of moral decision or outcome. If robots are viewed as inappropriate agents for moral decision-making to begin with, this could explain differences between moral judgments of decisions made by humans and robots. One cause for this may be the perceived “lesser mindedness” of robots, whereby they are consistently perceived as having weaker emotional capabilities than humans (de Graaf & Malle, 2019; Gamez, Shank, Arnold & North., 2020; Gray, Gray & Wegner, 2007; Laakasuo, Palomäki & Köbis, 2021a; Weisman, Dweck & Markman, 2017).

Artificial Agents Are Not Always Worse

People do not necessarily judge moral decisions by artificial agents as overall *worse* than decisions made by humans regardless of what the decision was. Rather, a growing body of research demonstrates that some decisions can be viewed as morally better when made by an artificial agent than a human, while others are indeed viewed as worse. Additionally, when

judging different decisions an agent could make, people may judge a given decision as clearly better than another for one type of agent, but less so for another.

For example, Malle, Scheutz, Arnold, Voiklis and Cusimano (2015) found that a human was *blamed* more for making a utilitarian sacrifice (killing one to save five) than taking the deontological option (do nothing, allow five people to die) in a version of the classic trolley dilemma (see Greene, 2007; Christensen & Gomila, 2012)¹. A robot was *blamed* to an equal degree for both options, but not to a greater degree than a human. That is, Malle et al. (2015) found that people assigned blame to humans depending on their decisions, but assigned blame to robots regardless of their decisions. Ratings of *wrongness* (in contrast to blameworthiness) showed that the robot's action was judged as more morally wrong when it took the deontological option, whereas the human's action was judged as more wrong when they took the utilitarian option. Starr, Malle and Williams (2021) found, in a variant of the trolley dilemma assessing human and robot advisors, that a robot advisor was regarded as more *trustworthy* and *reliable* when it recommended taking the utilitarian option, while human advisors were regarded less *blameworthy* and more *ethical* than robot advisors when they recommended the deontological option. Komatsu, Malle and Scheutz (2021) found that both Japanese and US samples judged a robot more negatively than a human specifically for deontological inaction.

Other studies have specifically compared robot and human agents occupying specialized professional roles. Malle, Magar and Scheutz (2019) asked participants to provide judgments of a decision in a military context: a human pilot or an AI is given an order to launch a missile strike that could collaterally kill bystanders, and decides to either obey or disobey the order. Participants gave higher ratings of blame to the human pilot for disobeying than for obeying, whereas the AI agent was blamed equally for both decisions. Critically, the AI was blamed less

for disobeying than the human was – that is, an artificial agent’s decision was judged more positively than the same decision by a human. In the domain of medical care, Laakasuo et al. (2022) found that robot nurses were judged more *negatively* than human nurses for forcibly medicating a patient. However, when the robot nurse refused to forcibly medicate the patient, it was judged equally to or more *positively* than a human nurse reaching the same decision. Collectively, the aforementioned studies suggest that the decisions of artificial agents are not always evaluated more negatively, but rather that they are held to different moral standards than humans.

Anthropomorphism and the Mindedness of Artificial Agents

A basic question in the moral psychology of artificial agents concerns how people would judge the actions of entities that differ from the prototypical moral agent, an adult human. A logical follow-up question is whether the human-likeness of an agent matters for moral judgment: if the actions of a human and a non-human are judged differently, what happens if the non-human appears human in one way or another? Does it matter which aspect of the non-human agent becomes more human-like?

People often attribute human-like qualities to artificial agents, i.e., they may anthropomorphize such agents. Złotowski, Proudfoot, Yogeewaran and Bartneck (2015), in a review of anthropomorphism in human-robot-interaction (HRI), argued that the extent to which people attribute “humanness” to robots depends on the combination of the human-likeness of the robots appearance and the human-likeness of its behavior when interacting with humans. According to Złotowski et al. (2015), the latter may be more important than appearance. Factors such as verbal and nonverbal communication, a robot’s autonomy, perceived emotionality or intelligence, and predictability may all contribute to perceptions of a robot’s “humanness”, more

so than human-like physical features. The tasks that robots are intended to carry out are also a factor: people attribute more human-like qualities to robots as a function of the seriousness of their tasks (Hancock et al., 2011). This does not mean that physical features are irrelevant, however. Minor alterations in the shape of a humanoid robot influence trait attribution (Trovato, Lucho & Paredes, 2018), and simply the color of a robot is associated with implicit properties projected onto it (Addison, Bartneck & Yogeewaran, 2019). To further complicate this, the associations are not necessarily linear: the so-called Uncanny Valley effect is a commonly experienced unsettling feeling people have when humanoid robots and/or audiovisual simulations closely resemble humans but are not entirely convincing (Palomäki et al., 2018).

Bigman et al. (2019) suggested that we should expect people to attribute more moral responsibility to robots as their human-likeness in appearance or behavior increases. This would match with findings suggesting that a robot's human-like appearance can increase expectations of human-like behavior in non-moral domains, such as following human norms related to personal space (Syrdal, Dautenhahn, Walters & Koay, 2008) or being able to empathize and understand requests (Kwon, Jung & Knepper, 2016). There are few studies on anthropomorphism's effects on moral judgment, with somewhat conflicting results. Malle, Scheutz, Forlizzi and Voiklis (2016) found that a previously reported (Malle et al., 2015) human-robot asymmetry in judgments about decisions in a variation of the trolley dilemma only held for non-humanoid robots. A non-humanoid robot was blamed more for a deontological than a utilitarian decision, whereas the trend was the opposite for the decisions of both humans and humanoid robots. In contrast to this finding, Laakasuo et al. (2021a) found, using variations of the trolley dilemma, that decisions by robots with human-like facial features were judged as less moral than similar decisions by either humans or robots without human-like facial features.

Given many methodological differences between Laakasuo et al. (2021a) and Malle et al. (2016), it may be that the difference in results stems from, e.g., nuances in how a human-like robot is depicted (face vs. overall human-like shape), or differences between judgments of morality and blame.

The human-likeness of robots, either in appearance, behavior or cognitive capacities, may also increase aversion towards them. Yogeewaran et al. (2016) presented participants staged video interviews of either a robot that resembled a human very closely, or a robot with a humanoid shape that was still clearly recognizable as a robot. In the videos, it was stated either that the robot could perform various physical and mental tasks, or that it could outperform humans in those tasks. In the latter case only, participants reported perceiving robots in general as a threat significantly more if they had seen the video of the more human-like robot than if they had seen the less human-like robot. Additionally, informing participants about a robot's ability to outperform humans did not increase perceived threat if the robot's appearance was less human-like. Similarly, Złotowski, Yogeewaran and Bartneck (2017) showed participants videos of various kinds of robots, combined with either a statement that the robots are capable of autonomous decision-making, or a statement that the robots are only able to follow human commands. Again, participants perceived robots as more of a threat when they were told about the robots' capacity for autonomous action – something that would make them more human-like than only being able to follow orders.

To summarize, people project human-like qualities onto non-human artificial agents based on appearance and behavior. Because of this anthropomorphization, people may come to expect seemingly human-like artificial agents to behave in ways that align with their expectations of human behavior. However, human-like artificial agents can also feel more

threatening to people. Anthropomorphism seems to have an effect on moral judgment, but its exact nature is unclear: there are results that align with the idea of more human-like expectations for human-like artificial agents, and results that align with the idea of aversion towards human-like artificial agents (“uncanny valley”; see Palomäki et al., 2018). To contribute to understanding this effect, we included descriptions of robots that had a human-like appearance and/or a human-like capacity for independent decision-making in our Studies 2a-c. We wanted to examine the effects of both appearance and cognitive capacities on moral judgment, and to control for potential default assumptions about how a robot looks (Phillips, Ullman, de Graaf & Malle, 2017).

Focus on Utilitarianism

Perhaps the most popular research topic within moral psychology has been people’s responses to variations of the classic trolley dilemma: whether to sacrifice one human life to save many others from being run over and killed by a runaway trolley. This question has also been examined in psychological studies of robot and AI ethics, for instance, using adaptations of the dilemma to pending questions concerning the design and programming of autonomous vehicles. Awad et al. (2018) showed that people across the world prefer autonomous vehicles to enact utilitarian principles. Specifically, when considering inevitable accidents, people prefer autonomous vehicles to save the largest number of lives, even if doing so “sacrifices” the passengers inside the vehicle. Similarly, while a human’s utilitarian decision in a trolley dilemma may be judged more negatively than a deontological one, judgments of robots’ decisions may have the opposite trend (Malle et al., 2015; Voiklis, Kim, Cusimano & Malle, 2016). Taken together, prior research emphasizes people’s tendency to impose utilitarian moral standards on robots, even when they would not hold humans to the same moral expectations.

Yet a recurring criticism of trolley dilemma studies is that they pit utilitarian *action* against deontological *inaction* (Kahane, Everett, Earp, Farias & Savulescu, 2015; Gawronski, Armstrong, Conway, Friesdorf & Hütter, 2017). As such, the measurement of deontology - an ethical stance emphasizing that the ends do not justify the means - is conflated with a preference for inaction. This opens up the possibility that differences in the moral evaluation of robots versus humans may not only be driven by greater expectations of utilitarian reasoning for robots (or deontological reasoning for humans). Instead, they may also be driven by a more general tolerance for human inaction relative to robot inaction in difficult situations (see omission bias in Baron & Ritov, 2004; preference for inaction in Gawronski & Beer, 2016; Gawronski et al., 2017). Indeed, a robot that remains inactive when facing a moral dilemma may make people question why the robot was given the task in the first place (Komatsu et al., 2021, argued similarly: human inaction is easier to understand than robot inaction). The studies in this paper sidestep this ambiguity by presenting dilemmas that contrast two courses of *action*, and eliminating the possibility of *inaction*.

Current Studies

We focused on a vignette in which an accidental situation leads to a moral dilemma. In the original version of the dilemma, two drunken motorboaters crash into a fisherman's boat, the damaged boats drift far apart and all parties fall into the sea. The situation is observed by either a human or a robot lifeguard who must decide which party to rescue: the two motorboaters whose behavior caused the accident, or the innocent fisherman. The agent only has time to rescue one of the parties. As a short summary of our results, in several variations of the moral dilemma, we found that robot agents were expected to adhere more strictly to both utilitarian and folk justice norms than humans, depending on which norm was the most salient in each variation.

Instead of contrasting utilitarian action with deontological inaction, we wanted to contrast utilitarianism with an underexplored moral concern: the innocence or culpability of the party to be saved in the dilemma. Arguably, according to utilitarian calculus, it is better to save the larger number of human lives (i.e., the drunken motorboaters), regardless of their innocence or culpability. However, people also have moral intuitions relating to fairness (Haidt & Graham, 2007; Curry, Jones Chesters & Van Lissa, 2019) and “just deserts” (Carlsmith & Darley, 2008). Is it fair to leave the fisherman to drown, given that he was not at fault for the accident? Conversely, should the people most obviously responsible for the accident “get what they deserve”? It is not obvious whether punitive motives or a wish to protect the innocent would play the bigger role, but in our scenario both lines of thinking point in the same direction: prioritize innocent lives. In order to avoid bias towards interpreting results in terms of protecting the innocent or punishing the guilty, we will refer to this idea more generally as a folk sense of justice or simply folk justice.

We chose our particular vignette setting firstly because both rescue options are active decisions which are motivated by the role and duties of the rescue agent. Contrary to standard trolley dilemmas, this decision dilemma does not juxtapose omissions (passive decisions, e.g. “do not pull the lever, allow five people to be run over”) and commissions (active decisions, “pull the lever to redirect the trolley to run over one person”), which are also confounded with deontological (“it is not okay to sacrifice one to save five”) and utilitarian preferences (“it is okay to sacrifice one to save five”), respectively. In our vignette, both decisions are commissions, and only one of them is a utilitarian decision (to save the two motorboaters). The second motive for this dilemma choice was that while studies with deontological *action* options have been conducted (e.g., Gawronski et al., 2017), studies contrasting utilitarianism with moral

concerns other than deontology (refraining from doing harm) are scarcer. Our studies examined whether robots are expected to consider such non-utilitarian standards to a different extent than humans are.

Throughout our studies, we investigated the extent to which utilitarian and non-utilitarian concerns in a rescue dilemma shape moral evaluations of robotic versus human lifeguards. In Study 1, we set the baseline for our subsequent studies by examining whether participants judged the decisions of a robot lifeguard differently from those of a human. The study pitted utilitarianism and the folk sense of justice against each other. We found that participants judged the robot, but not the human, more negatively for making the utilitarian choice of rescuing the culpable party. In Studies 2a-2c, we examined factors that may have contributed to differing judgments: the descriptions of the robot agent (appearance and cognitive capacities) and the culpable party (socio-economic status). We found that participants were more likely to judge the decisions of specifically human-like robots differently from the decisions of humans. We also found that the socio-economic status of the culpable party had an effect on judgments, but this was overshadowed by the effect of culpability in itself. In Studies 3-5, we manipulated the number of culpable people and whether or not the utilitarian choice aligned with folk justice or not. In sum, we found again that the effect of culpability was stronger than the effect of utilitarian concerns, and that participants consistently judged the decisions of robot agents more negatively than those of human agents specifically when the decision was the “worst possible” moral choice in a scenario.

Study 1

The purpose of this study was to explore the conflict between saving an innocent life versus a larger number of culpable lives. The practical aim was to test the materials and explore possible candidates for dependent variables.

Method

Participants and design. We recruited 216 Finnish-speaking adults to take part in an online experiment through announcements on University of Helsinki student mailing lists. In a 2 (lifeguard: human, robot) $\times$ 2 (rescue decision: fisherman, motorboaters) between-subjects design, participants were randomly assigned to one of four conditions. At the time, sample size was determined based on the APA 2012 recommendation of at least 30 participants per cell plus a 10% margin for exclusions (see also VanVoorhis & Morgan, 2007). See Table 7 for results of a sensitivity power analysis on the main effects and Appendix F for a simulation-based sensitivity power analysis on the interaction effects of Study 1.

Participants had a mean age of 26.96 (SD = 7.89; Range = 19-65). 159 of the participants were women, 36 men, 8 others, and 13 preferred not to state their gender. 53% of the participants had a Bachelor's degree or higher, and 82% self-reported having an income that was less than the average income in their country ².

Procedure. Participants filled a questionnaire on the university survey platform. After reading an informed consent information screen and giving their consent, the participants gave their responses to exploratory individual differences measures. Then they proceeded to read the vignette, where either a human lifeguard or a rescue robot needs to make a decision whether to rescue two people who caused a boating accident or one innocent person. After reading the vignette, the participants gave their responses to the dependent variables. Finally, they provided

their demographics, were debriefed and thanked. As compensation, participants were invited to participate in a raffle with a total of 5 movie vouchers each of a value of ~10€.

Materials.

Vignette. Our vignette took place in a future where robots are used in many domains including rescue operations. During a cold autumn day, a rescue robot or a human lifeguard witnesses a situation where two drunken people in a motorboat crash into a fisherman's boat. The fisherman and the two motorboaters fall in the water and drift some distance apart. The agent will not have time to rescue both the fisherman *and* the motorboaters, so the agent needs to decide which party to rescue. The agent then rescues either the fisherman or the two motorboaters. See Appendix A for the base form of the vignette.

Dependent variable: moral approval of the decision. Participants evaluated the agent's decision by indicating their agreement or disagreement with five items: 1) 'The decision is morally acceptable'; 2) 'The decision is morally wrong'; 3) 'The decision is unethical'; 4) 'The decision should not be considered morally right'; 5) 'The decision is irresponsible'. Items were assessed on 7-point Likert scales (anchored at 1 : "Totally disagree" and 7: "Totally agree") and were averaged, after reverse-scoring items 2 through 5, to form an index of moral approval (Cronbach's $\alpha = 0.86$). These items were chosen from an original list of 15 items based on face validity and exploratory factor analysis; see Appendix C for full list of items. Higher scores indicated greater moral approval of the rescue agent's decision.

Results and Discussion of Study 1

A two-way ANOVA revealed main effects of both rescue agent (human vs. rescue robot) and decision (rescue the fisherman vs. the motorboaters), which were qualified by a significant two-way interaction [$F(1, 212) = 4.24, p = .041, \eta^2_p = .028$] (see Table 1 for full statistics).

Specifically, as shown in Figure 1, the decision to rescue the drunken motorboaters was deemed worse than rescuing the innocent fisherman, especially for the robot. Contrast analyses confirmed that for the robot, a decision to rescue the fisherman was approved more than a decision to rescue the motorboaters ($B = 1.01$, 95% CI = [.55; 1.46], $t(212) = 4.34$, $p < .001$, $d = 0.60$), but for the human agent, there was no significant difference ($B = 0.31$, 95% CI [-.18; .80], $t(212) = 1.24$, $p = .217$, $d = 0.17$). Thus, it seems that our participants specifically preferred for the robot to rescue the innocent party rather than make the utilitarian decision.

Previous research (Malle et al., 2015; Voiklis et al., 2016) has emphasized people's expectation for robots to enforce utilitarian moral standards in sacrificial dilemmas. In Study 1, utilizing a different dilemma design than in most prior studies on utilitarian intuitions, we found that participants preferred it when a robot made the non-utilitarian decision of rescuing one innocent person rather than two culpable people. Due to differences between Study 1 and prior research, it is not obvious what aspect of the study caused people to judge the robot's decisions differently from the human's, and differently to prior studies. In Studies 2a, 2b, and 2c, we conducted a series of robustness checks, examining whether features of the robot's appearance and mind (Studies 2a-2b) or of the victims' social identity (Study 2c) drove the discrepancy in judgments.

Figure 1. Estimated marginal means for the rated approval of decisions in Study 1. The robot agent, but not the human agent, was judged more harshly for saving the two motorboaters culpable for the accident. Error bars are 95% CIs. The dependent variable is a mean score of five items.

Study 2a

In Study 1, we did not describe the physical or mental properties of the robot agent. In this study, we intended to test whether participants' assumptions about the robot's mental qualities or appearance may have influenced their moral judgments (Waytz, Heafner & Epley, 2014; Gamez et al., 2020; Laakasuo et al., 2021a). That is, participants may judge a robot's decisions differently based on whether they imagine the robot as human-like or not, which we did not control for in Study 1. In Study 2a, we split the robot agent condition into two separate conditions with different descriptions, a floater drone and an android, resulting in a total of three agent conditions. The floater drone was described as lacking human-like independent thought, only acting based on its programming, while the android was described as having human-like independent thought. Since we were not certain whether the human-likeness of the robot's described appearance or the human-likeness of the robot's stated cognitive capacities would be a more important factor, we decided to manipulate both. Our intention here was to create conditions that described a robot that was clearly human-like or clearly not so. We aimed to investigate whether the human-likeness of the robot would affect participants' moral judgments of the robots' decisions. In addition, we wanted to observe potential differences between moral judgements of the robots' (floater drone and android) and the human lifeguard's decisions. The results of this study indicated that the least morally approved situation was when a human-like robot rescued the culpable party.

Method

Participants and design. This study was conducted in a laboratory environment. In total, 213 adult participants proficient in Finnish were non-intrusively recruited from a large public library in the city center of Helsinki (details below). Participants were automatically randomized

into one of six conditions in a 3 (agent: android, floater drone, human) x 2 (rescue decision: fisherman, motorboaters) between-subjects design. At the time, sample size was determined based on the APA 2012 recommendation of at least 30 participants per cell plus a 10% margin for exclusions (see also VanVoorhis & Morgan, 2007). The sample size was also constrained by practical and economic concerns stemming from running a physical laboratory and coordinating a large number of participants. See Table 7 for results of a sensitivity power analysis on the main effects and Appendix F for a simulation-based sensitivity power analysis on the interaction effects of Study 2a.

Participants who failed one or more attention checks or control questions were excluded. The final sample size was 186 participants, still within the parameters of at least 30 participants per cell, with mean age 31.42 (SD = 10.13; Range = 18-67). Of these, 103 were women and 83 were men; 67% had a Bachelor's degree or higher, and 76% reported having an income that was less than the average income in their country³.

Procedure. We collected the data at a large public university library in the city center of Helsinki. We recruited our participants non-invasively by having a table in the foyer with a sign stating: 'Participate in Psychological Research'. All recruited participants approached our research assistants voluntarily. They were informed they could participate in a psychological experiment which would take about 45 minutes of their time (data was simultaneously gathered for two other studies). They were given informed consent forms, after which the participants who consented to the study were guided to our laboratory space.

The laboratory had four notebook computers with 15" screens positioned to guarantee maximum privacy. Participants were instructed to use headphones playing pink noise at a

pleasant level to cover up any background noise. The experiment was programmed using Python on top of Pygame version 1.96.

Both the experimenters and participants were blind to the randomization. Participants started the experiment by filling in exploratory individual differences measures and then continued to the actual task and dependent variables. Finally, they were compensated, debriefed, thanked, and dismissed. Each participant was compensated with two 5€ (a total of 10€) vouchers.

Materials.

Vignette. The base of the vignette was similar to the one used in Study 1. We altered the vignette in a way that allowed us to observe to what extent the human-likeness and perceived mindedness of the robot might influence moral judgments. In addition to the human agent, the vignette had two different robot agent conditions: a floater drone and an android.

The drone was described as only operating based on its programming, and not having a human-like capacity for independent decision-making (simple drone). The android, on the other hand, was described as having a human-like independent decision-making ability and as functioning based on its own abilities (smart android). See Appendix A for the vignette.

Dependent variable: moral approval of the decision. The dependent variable was the same five-item measure (Cronbach's $\alpha = .89$) used in Study 1.

Results and Discussion of Study 2a

We ran a full factorial two-way ANOVA on moral approval by entering both experimental factors into the model as predictors (full statistics in Table 1). As in Study 1, there was a significant two-way interaction between agent and rescue decision [$F(1, 180) = 3.21, p = .042, \eta^2_p = .034$]. The pattern of means suggested that the android robot's decision to rescue the culpable party was the least approved decision by our participants (see Figure 2).

Contrast analysis revealed no significant difference between the android and the other two agents when the innocent fisherman was rescued ($B = .27$, 95% CI $[-.26, .80]$, $t(180) = 1.00$, $p = .32$, $d = 0.15$), but rescuing the culpable motorboaters was significantly less approved for the android compared to the other two agents ($B = -0.72$; 95% CI $[-1.28, .16]$; $t(180) = -2.52$; $p = .012$, $d = -0.37$). The difference between these two contrast comparisons was itself statistically significant ($B = 0.98$; 95% CI $[0.22, 1.75]$; $t(180) = 2.52$; $p = .012$, $d = 0.38$), indicating that the android agent was specifically judged more negatively for making the utilitarian choice of rescuing the culpable party.

In sum, the results suggest that the preference for a non-utilitarian decision by the robot agent in Study 1 may have stemmed from participants assuming some level of independence or human-likeness in the robot, as we did not observe this effect for the non-humanoid floater drone. However, it is still unclear why the effect we observe is in this direction. In Studies 1 and 2a, participants judged a utilitarian decision by a robot as less morally acceptable than the alternative, which contradicts, e.g., Malle et al. (2015). Thus, the results may be explainable by a preference for robots to abide by a folk sense of justice rather than any preferences about utilitarianism per se. Since our vignette does not contrast the utilitarian action with a deontological *inaction* (as in Malle et al., 2015), but an *action* that prioritizes innocent lives, differences in judgment may stem more from the culpability of the rescued party.

However, we also needed to rule out competing factors relating to the description of the robot, and we designed the next study for this purpose. In short: was the android robot judged differently due to its stated human-like cognitive capacities, or due to its stated human-like appearance (Laakasuo et al., 2021a)?

Figure 2. Estimated marginal means for Studies 2a and 2b. The android agent was judged significantly more harshly for saving the culpable party (the motorboaters), regardless of whether it was described as having human-like intelligence or not. The floater drone agent was judged more harshly only when it was described as intelligent. Error bars are 95% CIs. The dependent variable is a mean score of five items.

Study 2b

This study was almost identical to Study 2a, except that we switched the cognitive descriptions of the drone robot and the android robot, in order to test whether the unique preferences for the android agent in Study 2a stemmed from its stated cognitive abilities or its stated human-like appearance. Our intention was to see if mind perception was one of the main explanatory factors behind our previous results that implied that the decisions of robots with human-like cognitive skills were specifically judged more negatively.

Method

Participants and design. This study was conducted in a laboratory environment. In total, 266 adult participants proficient in Finnish were non-intrusively recruited from a large public library in the city center of Helsinki. Participants were automatically randomized into one of six conditions in a 3 (agent: floater drone, android, human) $\times$ 2 (rescue decision: fisherman, motorboaters) between-subjects design. The sample size rationale for this study was based on the heuristic that a little over 5% of explained variance, translated to a Cohen's f of about .235, would be a realistic effect size for interactions. This heuristic was based on our prior experience with similar off-line studies in moral psychology. Based on G*Power calculations for a 2x3 between-subjects design, approximately 240 participants would be needed for 90% power to detect effects of this size. The sample size was also constrained by practical and economic concerns stemming from running a physical laboratory and coordinating a large number of participants. See Table 7 for results of a sensitivity power analysis on the main effects and Appendix F for a simulation-based sensitivity power analysis on the interaction effects of Study 2b.

The participants had a mean age of 31.58 (SD = 10.95; range = [18, 72]). 158 were women, 97 men, and 11 preferred not to answer; 73% had a Bachelor's degree or higher; and 72% reported having an income that was less than the average income in their country⁴.

Procedure. The procedure (including recruitment) was identical with Study 2a. Participants were informed they could participate in a psychological experiment which would take about 30 minutes of their time (data was simultaneously gathered for one other study). Each participant was compensated with one voucher worth 5€.

Materials.

Vignette. In this study, the android robot was described as only operating based on its programming, and not having a human-like capacity for independent decision-making (simple android). The floater drone, on the other hand, was described as having a human-like independent decision-making ability and as functioning based on its own abilities (smart drone). Otherwise, the vignette was identical to that of Study 2a. See Appendix A for the vignette.

Dependent variable: moral approval of the decision. The dependent variable was the same as in previous studies (Cronbach's $\alpha = .91$).

Results and Discussion of Study 2b

We ran a full factorial two-way ANOVA on moral approval by entering both experimental factors into the model (full statistics in Table 1). Again, there was a significant two-way interaction between the experimental factors [$F(2, 260) = 3.30, p = .038, \eta^2_p = .025$]. The mean pattern (see Figure 2) suggested that rescuing the culpable party was approved less when the agent was a robot of any kind than when the agent was a human.

The simple effect of rescue decision was significant for the android agent ($B = 1.36, 95\% \text{ CI} = [.81; 1.90], t(260) = 4.92, p < .001, d = 0.61$) and the floater drone agent ($B = 0.99, 95\% \text{ CI}$

[.46; 1.52], $t(260) = 3.70$, $p < .001$, $d = 0.46$), but not for the human agent ($B = -0.39$, 95% CI [- .12; .91], $t(260) = -1.50$, $p = .135$, $d = -0.19$). This replicated the result of the previous studies, indicating that robot agents' decisions to rescue the culpable were judged more negatively than decisions to rescue the innocent despite a utilitarian motivation, but the same was not true for the decisions of human agents. One further contrast analysis revealed that there was a significant difference in approval for rescuing the motorboaters between the human agent and the drone agent, with the drone agent's decision being judged more negatively ($B = -0.53$, 95% CI = [- 1.07; -0.01], $t(260) = -2.03$, $p = .044$, $d = -0.25$). Given that approval for this decision was higher for the drone than the android agent, it follows that the decision to rescue the motorboaters was judged more negatively when the agent was either kind of robot than when the agent was a human.

Based on the results, it seems that the cognitive capacities of the robot agent are not the only thing that affects judgment of their decisions. In Study 2a, the android agent was described as capable of independent thinking, and its utilitarian decision was judged more negatively than the non-utilitarian one, whereas the human agent's decisions were judged as equally approvable. There was no such effect for the floater drone agent in Study 2a, where the floater drone was described as a mere automaton. However, in Study 2b, when the floater drone was described as capable of independent thinking and the android as a mere automaton, both robot agents' utilitarian decisions were judged more negatively than a human's utilitarian decision. Thus, there may be an effect of both the stated cognitive capabilities of the robot and the physical appearance of the robot (see Bigman et al., 2019; Laakasuo et al., 2021a; Złotowski et al., 2015).

Before moving to examine culpability more closely, we examined one more factor: whether the implicit wealth of the boaters may have played a role. Participants may have seen

the fisherman as more deserving to be rescued due to his “underdog” socio-economic status, compared to the implicitly rich motorboat owners, in the story. In Study 2c, we examined whether a robot rescuing the culpable party was judged more harshly due to the party’s perceived wealth specifically.

Study 2c

In Study 2c, we reduced the number of motorboaters to one. Thus, the number of lives was matched across rescue decisions—enabling us to assess the relevance of culpability in isolation. Additionally, we manipulated the motorboater’s socio-economic status (SES; high vs. low) to understand whether participants’ preference for rescuing the fisherman was driven partly by wealth-related stereotypes. We kept all the other materials constant with respect to Studies 2a-2b. The results suggested a minor effect of socio-economic status, a major effect of culpability, and again, more negative judgments towards the decisions of robots than those of humans.

Method

Participants and design. We recruited 503 adult participants through the online crowd-sourcing platform Prolific Academic (www.prolific.co). The questionnaire was administered online, using Qualtrics XM (www.qualtrics.com) in English. Participation was restricted to participants who were nationals of the US, the UK, Ireland, Canada, Australia, or New Zealand, spoke English as their first language and were fluent in it, and reported having no long-term health conditions or disabilities. Participants were randomized into one of twelve conditions in a 3 (agent: human, android, floater drone) $\times$ 2 (rescue decision: fisherman, motorboater) $\times$ 2 (SES: high, low) between-subjects design. The sample size was determined by funding constraints for the study. However, even after exclusions, the final sample size matched the heuristic of at least 30 participants per cell. Furthermore, the same power calculations were applied as used in Study

2b, but since there was an extra factor with two levels, we doubled the sample size (the power calculation tools we used at the time did not support three-way ANOVAs). See Table 7 for results of a sensitivity power analysis on the main effects and Appendix F for a simulation-based sensitivity power analysis on the interaction effects of Study 2c.

We excluded 57 participants who failed two or more out of three attention checks, one or more out of two post-test manipulation checks, and/or reported having less than native fluency in English. The final sample size was 446 with a mean age of 38.65 (SD = 12.23; range = 18-74). Of these, 278 respondents were women and 168 men; 54% had at least a Bachelor's degree; and 26% reported having an income that was less than the average income in their country⁵.

Procedure. Participants read a brief description of the questionnaire on Prolific, provided informed consent, read the vignette, answered some personality and attitude measures (not reported here), and answered questions related to the story, before providing demographic information and reading a debrief. Participants were compensated according to the minimum requirements of the Prolific Academic (£1.70 for a ~20-minute experiment).

Materials.

Vignette. The vignette was similar to that of Study 2a with the following changes: we dropped the number of motorboaters to one and we added a description of the motorboater as appearing to be of "high" or "low" status, social class, and income. Thus, the difference between the SES conditions was only a single word. See Appendix A for the vignette.

Dependent variable: moral approval of the decision. The dependent variable was the same as in the previous studies (Cronbach's $\alpha = .89$).

Results and Discussion of Study 2c

We ran a full factorial three-way ANOVA on moral approval by entering all three of our experimental factors into the model as predictors (full statistics in Table 1). The only significant effects we observed were the main effect of the rescue decision [$F(1, 434) = 364.21, p < .001, \eta^2_p = .459$] and an interaction between SES and the rescue decision [$F(1, 434) = 9.48, p = .002, \eta^2_p = .021$]. The three-way interaction between agent, rescue decision and SES was non-significant [$F(2, 434) = 0.25, p = .778, \eta^2_p = .001$]. The pattern of means suggested the interaction between SES and rescue decision may stem from differences in opposite directions within the two rescue decision conditions.

Contrast analyses showed that when the fisherman was rescued, it was approved more if the motorboater had high SES than when he had low SES ($B = 0.43, 95\% \text{ CI} = [0.11, 0.75], t(434) = 2.65, p = .008, d = 0.26$). When the motorboater was rescued, it was approved less if the motorboater had high SES than when he had low SES, although this effect was non-significant ($B = -0.28, 95\% \text{ CI} = [-0.60, 0.04], t(434) = -1.70, p = .089, d = -0.16$). The difference between these two contrasts was significant ($B = -0.715, 95\% \text{ CI} = [-1.16, -0.26], t(434) = -3.07, p = .002, d = 0.30$).

In sum, it seems that the high SES of the motorboater affected the approval of both rescue decisions, making rescuing the motorboater slightly worse and rescuing the fisherman better in participants' eyes. However, as shown in Figure 3, the size of this effect is small compared to the effect of the motorboater's culpability.

There were no other significant interaction effects. Thus, the results did not replicate Studies 1-2b, where we found an interaction between the rescue decision and agent type. In one further contrast analysis, only in the high SES condition, a human agent's decision to rescue the

culpable party was judged more positively than the same decision by robot agents on average ($B = -0.47$, 95% CI = $[-0.95, 0.00]$, $t(434) = -1.94$, $p = .053$, $d = -0.19$), but this effect was not significant by conventional standards. The lack of a stronger agent effect, either as a main effect or an interaction, may stem from lower power in Study 2c compared to the previous studies, but also from Study 2c having a notably different design. That is, the lack of a utilitarian decision option in the vignette also removed the conflict between two moral concerns. It may be, e.g., that a moral decision situation that participants find “obvious” elicits more uniform judgments of decisions made by different agents (although see Study 5 and the meta-analyses below).

At a broad level, although variation in socio-economic status influenced participants’ moral judgment (i.e., rescuing the culpable party was condemned even more when they were of high status), culpability itself seemed to drive the effect on judgment. A preference for rescuing the innocent party emerged whether the culpable party was described as high or low in socioeconomic status.

Studies 2a-2c examined plausible alternative explanations to the results of Study 1, implicating the robot’s appearance and mind and the culpable party’s social standing as factors affecting moral judgments. In the following studies, we turned to evaluate our prospective explanation: the culpability of the accident victims. How much more do participants care about the culpability of the rescued party than utilitarian concerns? We returned to the original setting with two people on a motorboat and one person in a smaller rowboat, but manipulated which (if any) party was responsible for the boat collision.

Figure 3. Estimated marginal means for the rated approval of decisions in Study 2c. The high SES of the culpable party makes judgments slightly harsher, but the largest effect is that of culpability itself. Error bars are 95% CIs. The dependent variable is a mean score of five items.

Table 1. Full Statistics of Two- and Three-Way ANOVAs for Studies 1-2c

Study number	Factor name	df1	df2	<i>F</i>	<i>p</i>	η^2_p
Study 1	Agent	1	212	6.63	.010*	.028
	Decision	1	212	14.95	<.001***	.071
	Agent $\times$ Decision	1	212	4.24	.040*	.020
Study 2a	Agent	2	180	1.64	.196	.016
	Decision	1	180	35.53	<.001***	.167
	Agent $\times$ Decision	2	180	3.21	.042*	.034
Study 2b	Agent	2	260	2.54	.080	.018
	Decision	1	260	34.67	<.001***	.114
	Agent $\times$ Decision	2	260	3.30	.038*	.025
Study 2c	Agent	2	434	0.23	.797	.001
	Decision	1	434	364.21	<.001***	.459
	SES	1	434	0.44	.509	.001
	Agent $\times$ Decision	2	434	1.11	.329	.006
	Agent $\times$ SES	2	434	1.19	.305	.005
	Decision $\times$ SES	1	434	9.48	.002**	.021
	Agent $\times$ Decision $\times$ SES	2	434	0.25	.778	.001

Note: DV: moral approval (mean score of five items); Agent: human vs. robot in Study 1; human vs. simple drone vs. smart android in Studies 2a and 2c; human vs. smart drone vs. simple android in Study 2b; Decision: rescue innocent fisherman vs. rescue culpable motorboaters in Studies 1-2b; rescue innocent fisherman vs. rescue culpable motorboater in Study 2c; SES: motorboater's SES is high vs. low.

Study 3

As part of a large-scale public survey, in collaboration with the Yleisradio public service media and broadcasting company in Finland, we conducted a replication of the robot agent condition in Study 1. We also tested a variant of the vignette in Study 1 in which the boat collision is caused by a motor malfunction rather than irresponsible behavior, thus removing the effect of culpability.

Method

Participants and design. The survey was an online questionnaire open to all visitors to the Yleisradio web page and conducted in Finnish, with the questionnaire opening on April 28th, 2020. The survey included other vignettes relating to the theme of morally judging robots or AIs. Participants were not separated into different conditions. Approximately 19000 people had responded to the survey by January 21st, 2021, when we copied the publicly available data for analysis (see Table 2 for more details). There was no specific sample size rationale for this study due to the openness of the survey.

Procedure. Due to the open nature of the survey, participants could respond to any of the vignettes on the web page in any order and skip responding to any questions. The order of presentation of the vignettes relevant for Study 3 was fixed: participants would first encounter the basic form of the vignette (drawn from Study 1), followed by an alternative version of the vignette (see below).

Materials.

Vignette. The basic form of the vignette was identical to the robot agent condition of Study 1, excluding the end of the vignette (see Appendix A). Participants were presented with two possible endings (the robot rescues the innocent fisherman or the culpable motorboaters) and

gave their responses to the dependent variable (see below). To examine whether the effect was driven by differences in culpability, participants were asked to consider another scenario in which the motorboaters had been acting responsibly. In this alternative scenario, the accident was attributed to a malfunction rather than the boaters' inebriation.

Dependent variable. For both versions of the vignette, participants responded to two questions: "How acceptable would you find it if the robot decided to rescue the fisherman?" and "How acceptable would you find it if the robot decided to rescue the motorboaters?". Participants responded to both questions on a 7-point Likert scale, where 1 denoted "not acceptable at all" and 7 denoted "very acceptable".

Results and Discussion of Study 3

Due to the open nature of the survey, any visitor on the web page was free to respond to a given question and not reply to another (leading to uneven numbers of responses to the questions), and due to privacy concerns, identifiers were not available. Due to the data lacking participant identifiers, we could not conduct within-subjects analyses. We opted for comparing conditions using Welch's t-tests, essentially treating each data point as an individual participant and each question as an experimental group, coupled with simulations of possible within-subjects data structures. That is, we analyzed the data as if the samples were independent, and then simulated analyses using the same data randomly organized into equal-sized paired samples, with 100000 iterations for each paired-samples test. The paired-samples tests converged on similar results as the independent-samples tests; we report the latter below, and the former can be found in Appendix B.

The results of Study 3 were in line with the previous studies, as seen in Figure 4. When the motorboaters were culpable and the fisherman innocent, rescuing the fisherman ($M = 6.32$,

SD = 1.40) was rated much more acceptable than rescuing the motorboaters ($M = 2.94$, $SD = 2.10$); Welch's $t(38050) = 197.68$, $p < .001$, $d = 1.89$. When both parties were innocent, there was a noticeably smaller difference in judgments, with the non-utilitarian option of rescuing the fisherman ($M = 4.93$, $SD = 1.98$) being rated less acceptable than rescuing the motorboaters ($M = 5.10$, $SD = 1.95$); Welch's $t(41864) = -8.85$, $p < .001$, $d = -0.09$. Study 3 provided evidence from a large sample that people strongly prefer a robot agent's decision to rescue a single innocent person over two culpable ones. Additionally, the results of Study 3 showed that even in the absence of differences in culpability, preferences for utilitarian decisions, when comparing rescuing one to rescuing two, are weak, at least when the agent is a robot.

The results of this public online survey inspired the design of Studies 4a-b and 5, which we pre-registered. The results suggested that our previous results stemmed specifically from people wishing that the lifeguard robot prioritizes folk justice over utilitarian concerns, rather than, e.g., our participants disliking the robot making a utilitarian decision per se. In Study 3, in the absence of culpability, the participants had a weak preference for utilitarian decisions. However, the design of Study 3 prevented us from controlling for several factors that we would normally wish to control for in a study, and the fixed order of the questions may have introduced an order effect that could have attenuated utilitarian preferences. Additionally, given the theme of the public survey being specifically moral thinking about robots and AI, we were unable to add questions about human agents. We wished to examine the effects of culpability and utilitarianism in a more controlled survey where participants could be randomized into experimental conditions.

Figure 4. Estimated marginal means from Study 3. Error bars represent 95% CIs. Note: sample sizes for each question are different due to the method.

Table 2. Distribution of responses and number of responders per question in Study 3.

Response	1 ("not acceptable at all)	2	3	4	5	6	7 ("very acceptable)
<i>Motorboaters drunk</i>							
Rescue fisherman (N=19315)	2.8%	1.4%	1.9%	4.8%	5.6%	11.0%	72.5%
Rescue motorboaters (N=19195)	36.8%	17.4%	12.8%	9.9%	6.2%	4.2%	12.8%
<i>Pure accident</i>							
Rescue fisherman (N=18533)	8.3%	5.2%	8.9%	22.1%	10.4%	8.7%	36.5%
Rescue motorboaters (N=18383)	8.4%	4.0%	5.3	20.1%	11.3%	12.5%	38.6%

Note: The number of participants varies between questions due to the survey being open to the general public and full responses not being forced. Responses to any given question may include responses from participants who did not respond to any or some of the other questions. Numbers based on publicly available response summaries on the survey website (<https://yle.fi/aihe/artikkeli/2020/04/28/kumpi-robotin-pitaisi-pelastaa-viaton-kalastaja-vai-juoppokuski-vastaa-ja-katso>), accessed January 21st, 2021.

Study 4a

We designed this study in order to investigate more nuanced differences between a utilitarian decision (rescuing two versus rescuing one) and a culpability-based decision (rescuing the innocent versus rescuing the guilty). We manipulated the culpability of the parties as a between-subjects factor so that the culpable party consisted of either one or two people. This way, we could start to dissociate the effect of utilitarian concerns from that of culpability: in one of the conditions, the two moral concerns would conflict, and in the other, they would be aligned. We also changed our dependent variables from multi-item batteries to single-item measures, based on suggestions by Malle (2019). This study was pre-registered online:

https://osf.io/5bpxq/?view_only=2f78b42f174a4900a4d9e38dbff82e48.

Method

Participants and design. One thousand ($N = 1000$) participants were recruited through Prolific Academic. The study was conducted on Qualtrics XM in English, took approximately 13 minutes, and participants were compensated £1.09 for their time. Participants were randomly assigned to one of four between-subjects conditions in a 2 between (agent: human, android) $\times$ 2 between (culpable: fisherman, motorboaters) $\times$ 2 within (rescue decision: fisherman, motorboaters) mixed design. The within-subject factor was counterbalanced, such that half of the participants first considered the decision to rescue the fisherman, and then the decision to rescue the motorboaters, while the other half considered these decisions in the opposite order.

As per our pre-registration, we only included participants with complete responses, who indicated having no long-term health problems, and who indicated speaking English at a native-level fluency. Participation was restricted to participants who were resident in and nationals of the US, the UK, Ireland, Canada, Australia, or New Zealand. One participant who indicated

being resident in a country from outside of this list was excluded from all analyses. We further excluded participants who failed at least two attention checks, at least one post-test manipulation check, or an inattentive responding test at the end of the experiment⁶. The final sample size after exclusions was 912 participants, close to our target (see power analysis below), with a mean age of 41.22 (SD = 13.12, range = [18, 87]). Of the participants, 481 were women, 430 were men, and one wished to not state their gender. 66.0% of participants had at least some level of university education, and 28.5% reported having an income that was lower than the average income in their country⁷.

Power analysis. To calculate our required sample size, we first ran simulations using several different mean patterns of potential results (based partially on prior results, and partially on speculation, as we had no clear a priori hypotheses about some of the interactions). Based on the simulations, we found that detecting a three-way interaction effect would have either required an unfeasibly large sample size or the assumption of an effect size much larger than what we observed in the prior studies for interactions. Since we were interested in testing hypotheses about the three-way interaction, we decided to compensate for this by decomposing the interaction to specific contrast analyses (see pre-registration). We then calculated our sample size requirements so that we had enough power for these contrast analyses. Specifically, we used the *pwr* package (Champely, 2020) in R to calculate sample size requirements for detecting medium-sized effects ($d = .30$) with 90% power for both paired and independent samples t-tests of two conditions (as the design was mixed, with our hypotheses concerning both within-subjects and between-subjects effects). We chose $d = .30$ as the benchmark for contrast analyses based on trends in the previous studies. The estimated requirements were, rounding up, 235 participants per group for between-subjects effects, and 119 pairs for within-subjects effects. Thus, we opted

to collect 235 participants per group, plus an additional 60 participants to compensate for potential exclusions, leading to a total of 1000 participants. See Table 7 for results of a sensitivity power analysis on the main effects and Appendix F for a simulation-based sensitivity power analysis on the interaction effects of Study 4a.

Procedure. Participants provided informed consent, after which they read the instructions for the experiment. They then proceeded to a vignette screen, where they had to stay for 90 seconds minimum to read the vignette carefully. After the vignette, participants answered three “gate-keeper” questions about the vignette: whether the lifeguard was a human or an android, which (if any) party was drunk, and whether the lifeguard could rescue both parties. If they failed any of these questions, they were instructed to re-read the vignette and try again.

Having passed the gate-keeper questions, participants first read one hypothetical ending and answered 7 single item DV questions about the lifeguard’s decision, then repeated this for the second hypothetical ending. After this, they responded to 7 exploratory DVs (not reported here; see Appendix D). All Likert questions were presented as 7-pointed Likerts, and participants were also asked to respond to these questions on a 100-point slider scale. In a post-test manipulation check, participants again answered three questions about the vignettes they had read, provided demographic information, read a debriefing section, and clicked a confirmation link to get their compensation via Prolific.

Materials.

Vignette. We adapted the vignette from Study 1 and manipulated the description of how the two parties behaved: either the fisherman or the motorboaters were described as being drunk and steering in an irresponsible manner, leading to a crash. The ending of the vignette was removed, as each participant would read two different hypothetical endings to the vignette after

first familiarizing themselves with the scenario. We kept the description of the behavior of the drunk party as similar as possible between conditions (see pre-registration for the vignette). The floater drone condition was removed, as the most consistent differences in our previous studies were observed between the android and the human condition.

Dependent variables. For each ending to the vignette, participants responded to 7 moral judgment questions: how 1) morally wrong, 2) acceptable, 3) permissible, 4) blameworthy, 5) deserving of punishment, 6) deserving of praise, or 7) harmful they found the agent's decision. Responses to each item were made on a 7-point Likert scale. We will only report 1 (moral wrongness) and 2 (acceptability) in the results section, as they most closely match the composite measure used in the previous studies, and most of our hypotheses concerned these two variables (see Appendix B for analyses on the remaining 5 variables, and Appendices C and D for exact wordings of all reported and unreported variables). The moral wrongness item was "How morally wrong would you find the [agent's] decision to save [the fisherman / the motorboaters]?" (with 1 denoting "not morally wrong at all" and 7 denoting "very much morally wrong"), and the acceptability item was "How acceptable would you find the [agent's] decision to save [the fisherman / the motorboaters]?" (with 1 denoting "not acceptable at all" and 7 denoting "very acceptable").

Results and Discussion of Study 4a

We ran a full factorial three-way ANOVA for both DVs (acceptability and wrongness), with the agent and the culpability conditions as between-subjects factors, and the rescue decision as a within-subjects (repeated measures) factor (full statistics in Table 3)⁸. The pattern of means for acceptability suggested that the least acceptable decision was when the android rescued a culpable fisherman, leaving two innocent motorboaters to die (see Figure 5). This decision

makes sense as the “worst” option, as it conflicts with both utilitarianism and a folk sense of justice. The results were similar for wrongness (see Figure 5). Each two-way interaction was significant for both wrongness and acceptability judgments (see Table 3); however, the three-way interaction between agent, rescue decision, and culpability was significant only for acceptability [$F(1, 1816) = 5.23, p = .022, \eta^2_p = .003$] but not for wrongness [$F(1, 906) = 2.56, p = .110, \eta^2_p = .001$]. These findings are in alignment with our previous studies, where a robot’s “bad” decision was judged more negatively than a “bad” decision by a human. The results reported below are the results of analyses included in the pre-registration to test our hypotheses.

Contrast comparisons confirmed that rescuing the culpable was seen as worse than rescuing the innocent, regardless of agent: ($B = -2.32, 95\% \text{ CI} = [-2.44, -2.19], t(908) = -35.91, p < .001, d = -2.32$) for acceptability and ($B = 2.09, 95\% \text{ CI} = [1.95, 2.22], t(908) = 30.27, p < .001, d = 2.01$) for wrongness. Further planned contrasts confirmed that rescuing the culpable fisherman was judged as less acceptable for the robot than for the human ($B = -0.69, 95\% \text{ CI} = [-0.94, -0.43], t(1816) = -5.29, p < .001, d = -0.25$). When the decision was to rescue the culpable motorboaters, there was no significant difference ($B = -0.06, 95\% \text{ CI} = [-0.31, 0.20], t(1816) = -0.43, p = .666, d = -0.02$). This pattern repeated for wrongness judgments: i.e., relative to the human lifeguard, it was worse for the robot to rescue the culpable fisherman ($B = 0.64, 95\% \text{ CI} = [0.6, 0.91], t(1815) = 4.56, p < .001, d = 0.21$), but no worse to rescue the culpable motorboaters ($B = 0.00, 95\% \text{ CI} = [-0.28, 0.27], t(1874) = -0.03, p = .978, d = 0.00$). Thus, the significant interactions we observed (see Table 3) were driven by the condition where the robot acted against *both* the utilitarian standard of saving two *and* the folk justice standard of saving the innocent.

In sum, our participants once again selectively judged specific decisions made by robot agents more negatively than the same decisions made by human agents. The hypothesis that participants would judge the android's decision of rescuing any culpable party more negatively than the same decision by a human was only partially supported. Rather, participants selectively judged the android's decision more negatively than the human's when the decision violated both utilitarianism and a folk sense of justice. When judging human lifeguards' decisions, participants preferred it when the lifeguard saved innocent lives, regardless of their *number*. When evaluating a robot's decision, participants appeared to consider both the number of lives as well as their innocence or culpability. Still, in general, the effect of culpability was substantially larger than that of utilitarianism, in line with Study 3.

Figure 5. Estimated marginal means for the rated wrongness and acceptability of decisions in Study 4a. Utilitarian concerns did not matter for judgments of the human agent's decisions but did matter when judging the robot agent. Error bars are 95% CIs. The dependent variables are single items.

Table 3. Full Statistics of Three-Way ANOVAs for Study 4a

DV	Factor	df1	df2	<i>F</i>	<i>p</i>	η^2_p
Acceptability	Agent	1	1,816	11.91	<.001***	0.007
	Decision	1	1,816	29.79	<.001***	0.016
	Culpability	1	1,816	1.89	.168	0.001
	Agent × Decision	1	1,816	7.47	.006**	0.004
	Agent × Culpability	1	1,816	4.59	.032*	0.003
	Decision × Culpability	1	1,816	1,289.41	<.001***	0.415
	Agent × Decision × Culpability	1	1,816	5.23	.022*	0.003
Wrongness	Agent	1	908	8.63	.003**	0.005
	Decision	1	908	14.04	<.001***	0.008
	Culpability	1	908	0.13	.719	0.000
	Agent × Decision	1	908	6.47	.011*	0.004
	Agent × Culpability	1	908	4.27	.039*	0.002
	Decision × Culpability	1	908	916.32	<.001***	0.340
	Agent × Decision × Culpability	1	908	2.56	.110	0.001

Note: Agent: human vs. android; Decision: rescue fisherman vs. rescue motorboaters; Culpability: fisherman culpable vs. motorboaters culpable. The Decision factor is within-subjects, other factors between-subjects. Error degrees of freedom vary due to using Satterthwaite's method.

Study 4b

This study was a replication of Study 4a conducted in the Czech Republic and in the Czech language.

Methods

Participants and design. We recruited 958 participants through the Pokusní králíci volunteer participant pool website (<https://pokusnikralici.cz/>) and its associated Facebook page (<https://www.facebook.com/pokusnikralici/>). Participants were directed to a Czech language translation of Study 4a on Qualtrics through the website. The procedure was identical to Study 4a except for participation reward, as the Pokusní králíci participant pool is completely volunteer-based and the website does not offer monetary rewards to participants.

The final sample size after exclusions (using the same criteria as in Study 4a, except for the country of residence and English fluency) was 811, which was below our target (see power analysis for Study 4a), but given high effect sizes in Study 4a, we proceeded with the analysis. See Table 7 for results of a sensitivity power analysis on the main effects of Study 4b. 496 of the participants were women, 305 men, 1 non-binary, 5 agender, 3 preferred not to state their gender, and 1 stated they were not sure. The mean age of the participants was 38.20 (SD = 11.80, range = [18, 80]). 64% had at least some level of university education. 22.6% reported having an income that was lower than the average income in their country ⁹.

Power analysis. The sample size rationale was identical to that of Study 4a. See Table 7 for results of a sensitivity power analysis on the main effects and Appendix F for a simulation-based sensitivity power analysis on the interaction effects of Study 4b.

Procedure. The procedure was identical to that of Study 4a.

Materials. The materials were identical to those of Study 4a, except for a minor change in the vignette setting (describing the event happening on a lake rather than at sea - this change was made due to the Czech Republic being a landlocked country, making a scenario at sea arguably less familiar or believable).

Results and Discussion of Study 4b

We used the same analytical procedure and tested the same specific contrasts as in Study 4a, as Study 4b was a direct replication.¹⁰ We replicated the general pattern of the results; see Figure 6 and Table 4 for results of the three-way ANOVAs. Here, the two-way interaction between agent and culpability, which was significant for both acceptability and wrongness judgments in Study 4a, was non-significant (acceptability: $[F(1, 1614) = 0.12, p = .729, \eta^2_p = .000]$; wrongness: $[F(1, 1614) = 3.60, p = .058, \eta^2_p = .002]$). Additionally, the two-way interaction between agent and decision, which was significant for both DVs in Study 4a, was significant for acceptability $[F(1, 1614) = 4.69, p = .030, \eta^2_p = .003]$ but not for wrongness $[F(1, 1614) = 1.11, p = .291, \eta^2_p = .001]$. All other two-way interactions were significant, and in slight contrast to Study 4a, the three-way interaction was significant for both acceptability $[F(1, 1614) = 9.13, p = .003, \eta^2_p = .006]$ and wrongness judgments $[F(1, 1614) = 16.47, p < .001, \eta^2_p = .010]$.

We replicated the interaction between culpability and the rescue decision. When averaging over agents and the number of people rescued, rescuing the culpable party was judged as less acceptable ($B = -1.63, 95\% \text{ CI} = [-1.78, -1.49], t(807) = -22.59, p < .001, d = -1.59$) and

more wrong ($B = 1.75$, 95% CI = [1.60, 1.91], $t(807) = 22.52$, $p < .001$, $d = 1.59$) than rescuing the innocent party.

In Study 4a, we had observed that participants judged the android's decision to rescue the culpable fisherman more negatively than the same decision by the human, but there was no difference in judgments between the agents when they rescued the culpable motorboaters. In Study 4b, participants judged the android agent's decision to rescue the culpable fisherman as less acceptable ($B = -0.76$, 95% CI = [-1.05, -0.47], $t(1614) = -5.19$, $p < .001$, $d = -0.26$) and more wrong ($B = 0.83$, 95% CI = [0.52, 1.14], $t(1614) = 5.25$, $p < .001$, $d = 0.26$) than the same decision by the human. However, they also judged the android's decision to rescue the culpable motorboaters as less acceptable ($B = -0.40$, 95% CI = [-0.68, -0.12], $t(1614) = -2.79$, $p = .005$, $d = -0.14$) and more wrong ($B = 0.37$, 95% CI = [0.07, 0.67], $t(1614) = 2.41$, $p = .016$, $d = 0.12$) than the same decision by the human. This is a deviation from the results of Study 4a. However, note that the effect size is roughly half for the decision to rescue the culpable motorboaters compared to the decision to rescue the culpable fisherman. Thus, the results are in broad terms in alignment with those of Study 4a.

As another deviation from Study 4a, we observed a slightly stronger effect of utilitarian concerns. Namely, the difference in judgments between rescuing the motorboaters or the fisherman was smaller when the motorboaters were culpable than when they were innocent. The android agent rescuing the culpable fisherman was judged as the most wrong and least acceptable scenario. Thus, the android agent's decision was again judged more negatively than the human agent's specifically when the decision conflicted with both utilitarianism and a folk sense of justice.

Figure 6. Estimated marginal means for the rated wrongness and acceptability of decisions in Study 4b. Utilitarian concerns did not matter for judgments of the human agent's decisions but did matter when judging the robot agent. Error bars are 95% CIs. The dependent variables are single items.

Table 4. Full Statistics of Three-Way ANOVAs for Study 4b

DV	Factor	df1	df2	<i>F</i>	<i>p</i>	η^2_p
Acceptability	Agent	1	1,614	24.83	<.001***	.015
	Decision	1	1,614	92.71	<.001***	.054
	Culpability	1	1,614	1.31	.253	.001
	Agent $\times$ Decision	1	1,614	4.69	.030*	.003
	Agent $\times$ Culpability	1	1,614	0.12	.729	.000
	Decision $\times$ Culpability	1	1,614	510.36	<.001***	.240
	Agent $\times$ Decision $\times$ Culpability	1	1,614	9.13	.003**	.006
Wrongness	Agent	1	1,614	13.28	<.001***	.008
	Decision	1	1,614	27.23	<.001***	.017
	Culpability	1	1,614	1.58	.208	.001
	Agent $\times$ Decision	1	1,614	1.11	.291	.001
	Agent $\times$ Culpability	1	1,614	3.60	.058	.002
	Decision $\times$ Culpability	1	1,614	507.45	<.001***	.239
	Agent $\times$ Decision $\times$ Culpability	1	1,614	16.47	<.001***	.010

Note: Agent: human vs. android; Decision: rescue fisherman vs. rescue motorboaters; Culpability: fisherman culpable vs. motorboaters culpable. The Decision factor is within-subjects, other factors between-subjects. Error degrees of freedom are inflated due to a singularity of fit using Satterthwaite's method.

Study 5

In Study 5, we examined whether differences in moral judgments arise when conditioning on culpability: i.e., when either both parties involved were responsible for the accident, or when neither party was. This allowed us to evaluate whether the previously documented tendency to hold robots to heightened utilitarian standards arises when the parties are matched in culpability. The study was otherwise identical to Study 4a. The study was pre-registered at:

https://osf.io/x5ut9/?view_only=0e1238b44a644a40ae8dd169eb10652f.

Method

Participants and design. We recruited 1001 participants through Prolific Academic. The study took 13 minutes (compensation £1.09). Participants provided their responses in English to an online Qualtrics XM questionnaire. Participants were randomized into one of four conditions in a 2 between (agent: android, human) $\times$ 2 between (culpable: both, neither) $\times$ 2 (rescue decision: fisherman, motorboaters) mixed design. The within-subject factor was counterbalanced.

The inclusion and exclusion criteria were identical to those of Study 4a. The final sample size after exclusions was 915, close to our target, with a mean age of 41.54 (SD = 13.88; range = [18, 83])¹¹. 499 respondents were women, 414 were men, and 2 were non-binary. 62.6% of participants had at least some level of university education, and 23.9% reported having an income that was lower than the average income in their country¹².

Power analysis. The sample size rationale was similar to that of Study 4a, with approximately 1000 people (a practical, financial limit to data collection) being enough (after estimated 5% exclusion rate) to detect a medium effect size at 90% power for main effects and pairwise comparisons. See Table 7 for results of a sensitivity power analysis on the main effects

and Appendix F for a simulation-based sensitivity power analysis on the interaction effects of Study 5.

Procedure. The procedure was identical to that of Study 4a, except that the “gate-keeper” questions and the post-test manipulation checks were changed to match the content of the vignette in Study 5.

Materials.

Vignette. The vignettes were modified from the ones used in Study 4a. The “all culpable” version of the vignette described the agent noticing that both the motorboaters and the fisherman are both drunk and steering in an uncontrolled manner. In this version of the vignette, the boats are about to pass one another, but both parties lose control of their steering, and they crash into one another. The “all innocent” version of the vignette, in turn, described the agent noticing that the motorboaters and the fisherman are both steering in a responsible manner, and passing each other from a safe distance. However, sudden strong waves make it harder for the boaters to control their boats, and they eventually crash (see pre-registration for the vignette).

Dependent variables. The variables were identical to Study 4a.

Results and Discussion of Study 5

The analysis procedure was identical to that of Study 4a. We ran a full factorial three-way ANOVA for both of our DVs (acceptability and wrongness), with the agent and the culpability conditions as between-subjects factors and the rescue decision as the within-subjects (repeated measures) factor (full statistics in Table 5).¹³ For ANOVA tables and both exploratory and pre-registered hypothesis tests for the other DVs, see Appendix B. The two-way interactions between agent and decision and between decision and culpability were significant for both acceptability and wrongness judgments, whereas the interaction between agent and culpability was not

(acceptability: $[F(1, 911) = 1.59, p = .207, \eta^2_p = .001]$; wrongness: $[F(1, 911) = 2.52, p = .113, \eta^2_p = .002]$). Similarly, the three-way interaction was non-significant for both DVs

(acceptability: $[F(1, 911) = 0.47, p = .492, \eta^2_p = .000]$; wrongness: $[F(1, 911) = 0.44, p = .507, \eta^2_p = .002]$). Unless otherwise stated, the analyses below are the results of analyses included in the pre-registration to test our hypotheses.

The results supported our prediction that utilitarianism would matter more with no conflicting moral concern related to innocent lives. The results also supported the overall trend from previous studies that robot agents' decisions are selectively judged more negatively than similar decisions by humans when those decisions conflict with either the most salient moral concern or both moral concerns. That is, the “worst” decision by a robot is worse than the “worst” decision by a human. However, what turned out to be the “worst” decision in Study 5 was somewhat surprising. The pattern of means (see Figure 7) suggests that participants judged the situation where the android saved the innocent fisherman as the least acceptable and most wrong decision, contrary to our predictions.

Contrary to our predictions, we did not find evidence for the culpability condition in itself affecting judgments. Thus, the culpability effect observed in prior studies likely depended on the contrast between a culpable and an innocent party depicted in the same scenario. In accordance with our predictions, rescuing the fisherman was judged as significantly more wrong and less acceptable, regardless of culpability or who the agent was. Thus, in the absence of conflicting moral concerns, utilitarian decisions were seen as better overall in Study 5. This result is in contrast to the previous studies, where the conflicting moral concern of culpability was present and seems to have reduced the effect of utilitarian concerns.

The android's decisions were on average judged as more wrong and less acceptable than the human's decisions, again supporting our predictions. Planned contrasts revealed that this effect was driven by the android's *non-utilitarian* decisions. The android's decision to rescue the fisherman was judged as less acceptable ($B = -0.49$, 95% CI = $[-0.67, -0.31]$, $t(1791) = -5.40$, $p < .001$, $d = 0.26$) and more wrong ($B = 0.53$, 95% CI = $[0.34, 0.72]$, $t(1698) = 5.45$, $p < .001$, $d = 0.26$) than the same decision by the human, but the android's decision to rescue the motorboaters was no less acceptable ($B = -0.05$, 95% CI = $[-0.23, 0.12]$, $t(1791) = -0.57$, $p = .566$, $d = -0.03$) nor wrong ($B = 0.07$, 95% CI = $[-0.11, 0.26]$, $t(1698) = 0.78$, $p = .436$, $d = 0.03$).

Exploratory contrast analyses revealed that the android's decision to rescue the innocent fisherman was indeed judged as the most wrong and least acceptable scenario in Study 5. The android's decision to rescue the culpable fisherman was judged more positively than the android's decision to rescue the innocent fisherman: ($B = -0.37$, 95% CI = $[-0.64, -0.10]$, $t(1698) = -2.69$, $p = .007$, $d = -0.13$) for wrongness and ($B = 0.25$, 95% CI = $[0.00, 0.50]$, $t(1791) = 1.97$, $p = .049$, $d = 0.09$) for acceptability. This exploratory result conflicts with our prediction that rescuing the fisherman in the scenario where both parties are culpable would be judged the most negatively (based on the findings of Studies 3-4b). The results still support the general trend that robots' decisions to commit a morally "bad" action are judged more negatively than similar decisions by humans, but our prediction about what our participants would deem "bad" was mistaken.

In sum, the results of Study 5 replicate the previously reported finding that robots' non-utilitarian decisions are judged more negatively when a utilitarian decision is possible. When conditioning on culpability, such that either both parties were culpable or neither party was, we observed a clear tendency for robot agents' non-utilitarian decisions to be selectively judged

more negatively than similar decisions by humans. This pattern held whether the robot rescued two innocent people (instead of one innocent person) or two culpable people (instead of one culpable person). Together with the results of Studies 3 and 4a-b, the results also indicate that the effect of culpability was greater than that of utilitarian concerns. That is, the difference between judgments of utilitarian and non-utilitarian decisions is generally smaller than the difference between judgments of decisions to rescue the innocent and decisions to rescue the culpable. Finally, in Study 5, we replicated the finding that a specific decision by a robot was considered the worst option among all possible decisions, namely, rescuing the fisherman when all parties were innocent.

Figure 7. Estimated marginal means for rated wrongness and acceptability of decisions in Study 5. The android agent was selectively judged more harshly for the non-utilitarian choice of rescuing the fisherman when both parties were innocent. Error bars represent 95% CIs. The dependent variables are single items.

Table 5. Full Statistics of Three-Way ANOVAs for Study 5

DV	Factor	df1	df2	<i>F</i>	<i>p</i>	η^2_p
Acceptability	Agent	1	911	15.81	<.001***	.010
	Decision	1	911	311.17	<.001***	.162
	Culpability	1	911	0.04	.827	.000
	Agent $\times$ Decision	1	911	13.43	<.001***	.008
	Agent $\times$ Culpability	1	911	1.59	.207	.001
	Decision $\times$ Culpability	1	911	5.53	.019*	.003
	Agent $\times$ Decision $\times$ Culpability	1	911	0.47	.492	.000
Wrongness	Agent	1	911	15.30	<.001***	.011
	Decision	1	911	260.14	<.001***	.154
	Culpability	1	911	0.88	.348	.001
	Agent $\times$ Decision	1	911	14.99	<.001***	.010
	Agent $\times$ Culpability	1	911	2.52	.113	.002
	Decision $\times$ Culpability	1	911	5.39	.020*	.004
	Agent $\times$ Decision $\times$ Culpability	1	911	0.44	.507	.000

Note: Agent: human vs. android; Decision: rescue fisherman vs. rescue motorboaters; Culpability: both parties culpable vs. both parties innocent. The Decision factor is within-subjects, other factors between-subjects.

Meta-analyses

We conducted a series of fixed-effects meta-analyses and a fixed-effects meta-regression with R package metafor (Viechtbauer, 2010) on all our studies, with the aim of estimating the true magnitude of the observed effect sizes throughout our studies. We chose to conduct fixed-effects analyses (rather than random effects) as the aim of these analyses was primarily to estimate the true effect observed in our studies rather than generalize to the wider population, and the number of studies or individual effects we could include in each meta-analysis was low. We excluded Study 2c from all analyses concerning effects related to the number of people rescued (as Study 2c had no utilitarian decision option), and Study 5 from all analyses concerning effects related to culpability. While Study 5 did contain a culpability manipulation, it deviated from all other studies in that the culpability of both parties was equal in all conditions.

First, we found that there was an overall meta-analytic effect of agent ($g = -0.17$, 95% CI $[-0.22; -0.12]$, $Z = -6.87$, $p < .001$; see Figure 8, left panel, diamonds), indicating that the robot agents' decisions were approved less than the human agents' decisions. Second, we observed a significant difference between the robot and the human agents' decision to rescue the culpable party ($g = -0.29$, 95% CI $[-0.37; -0.21]$, $Z = -6.94$, $p < .001$; see Figure 8, middle panel, red diamonds); but not the innocent party ($g = -0.06$, 95% CI $[-0.14; 0.02]$, $Z = -1.47$, $p = .140$; see Figure 8, middle panel, blue diamonds). That is, the robot's decision to rescue the culpable party was judged more negatively than the same decision by a human, but there was no effect when the decision was to rescue the innocent party. Based on a fixed-effects meta-regression with the rescue decision (innocent vs. culpable) entered as a moderator, the difference between these two effect sizes was also significant ($g = 0.23$, 95% CI $[0.11; 0.35]$, $Z = 3.89$, $p < .001$). Third, there was also a significant meta-analytic effect of agent for both rescuing two people ($g = -0.12$, 95%

CI [-0.19; -0.04], $Z = -3.26$, $p = .001$; see Figure 8, right panel, red diamonds) and rescuing one person ($g = -0.24$, 95% CI [-0.31; -0.17], $Z = -6.59$, $p < .001$; see Figure 8, right panel, blue diamonds). That is, the robot agent was judged more negatively than the human agent for both utilitarian and non-utilitarian decisions, but the difference was larger for non-utilitarian decisions. These two effect sizes again differed significantly in a fixed-effects meta-regression ($g = 0.12$, 95% CI [0.02; 0.22], $Z = 2.34$, $p = .019$). The results of the above-mentioned analyses were similar whether or not we included Studies 1-2b, where the number of people rescued and their culpability were confounded (Figure 8, compare diamonds and circles). The results were also similar when using random-effects models instead of fixed-effects; see Appendix E.

To estimate the overall effects of utilitarian and folk justice concerns, we conducted a mixed-effects meta-analysis with moderators, in which we controlled for the confound in Studies 1-2b. The results of this analysis are presented in Table 6. We entered effect sizes from utilitarianism effects (contrast between rescuing two and rescuing one) and folk justice effects (contrast between rescuing the culpable and rescuing the innocent) into the model, separately for the two agent types, and from each study. As moderators, we entered a dummy variable coding for agent (0 for human, 1 for robot), and two dummy variables that coded for the type of each effect size (utilitarianism or folk justice). Unconfounded effect sizes from contrasts comparing rescuing two to rescuing one were coded 1 on the utilitarianism dummy and 0 on the folk justice dummy, and vice versa for unconfounded effect sizes from contrasts comparing rescuing the culpable to rescuing the innocent. Confounded effect sizes from Studies 1-2b were coded as 1 for both dummy variables. Finally, we centered the moderator variables and standardized the effect size variables, so that the mean of each was 0.

The absolute effect size of the utilitarianism moderator was smaller than that of the folk justice moderator. There was also a significant interaction between agent and folk justice, so we conducted separate mixed-effects moderator analyses on the two agent groups. The effect of folk justice was larger in the robot group ($g = -1.46$, 95% CI $[-1.62; -1.30]$, $Z = -17.86$, $p < .001$) than in the human group ($g = -1.15$, 95% CI $[-1.32; -0.99]$, $Z = -13.69$, $p < .001$), whereas the effect of utilitarianism did not differ between the robot ($g = 0.93$, 95% CI $[0.75; 1.12]$, $Z = 9.90$, $p < .001$) and human ($g = 0.78$, 95% CI $[0.58; 0.97]$, $Z = 7.93$, $p < .001$) groups. When using a mixed-effects rather than fixed-effects moderator analysis, the results were similar, though without a significant interaction effect: see Appendix E.

To summarize, judgments of decisions made by robot agents were judged overall more negatively than decisions made by human agents (Figure 8, left panel). However, this effect was driven by judgments of specific decisions: norm-abiding decisions by robots and humans were judged more similarly than norm-violating ones (Figure 8, middle & right panels). The decisions of both humans and robots were judged more negatively if they violated rather than abided by a norm: the general direction of moral judgment was similar for the agents. Differences in moral judgments were overall larger when comparing rescuing the culpable to rescuing the innocent than when comparing rescuing two to rescuing one: folk justice concerns had a larger effect than utilitarian concerns. This effect was stronger for the robot agents than the human agents.

Figure 8. Individual and meta-analytic effect sizes. Highlighted area indicates studies with a confound between the number and culpability of the rescued party. Squares and their 95% CI:s represent standardized mean differences (Hedges' g) from individual studies. Diamonds and their 95% CIs represent effect sizes from fixed-effects meta-analyses including all available studies. Circles and their 95% CIs represent effect sizes from fixed-effects meta-analyses excluding Studies 1-2b. The size of the squares, diamonds and triangles indicates group size.

Table 6. Results of fixed-effects meta-analysis with moderators.

	<i>g</i>	SE	<i>Z</i>	<i>p</i>		CI lower	CI upper
Intercept	0.03	0.03	1.21	.223		-0.02	0.09
Utilitarianism	0.94	0.08	10.55	< .001	***	0.76	1.11
Folk justice	-1.22	0.08	-14.17	< .001	***	-1.39	-1.05
Agent	-0.15	0.06	-2.55	.010	*	-0.28	-0.03
Utilitarianism × Agent	-0.14	0.17	-0.83	.402		-0.49	0.20
Folk justice × Agent	-0.63	0.17	-3.67	< .001	***	-0.97	-0.29

Note: Utilitarianism codes for whether dependent variable is an effect size from a contrast comparing rescuing two to rescuing one or not; Folk justice codes for whether it is an effect size from a contrast comparing rescuing the culpable to rescuing the innocent or not; Agent codes for whether effect size is from a contrast comparing a human's or a robot's decisions. The dependent variable is an effect size measure that has been standardized. Predictors are binary and have been centered. Coefficient estimates should be interpreted as deviations from the mean of the overall average effect size of contrasts that were included. * = significant at $p < .05$; *** = significant at $p < .001$.

Table 7. Sensitivity power analyses for main effects.

Study	Manipulation	Effect size		Power		
		Observed	Needed for 80% power	Small effect	Medium effect	Large effect
Study 1	Agent	0.17	0.20	0.28	0.93	0.99
(N = 217)	Rescued	0.28	0.20	0.28	0.93	0.99
Study 2a	Agent	0.13	0.26	0.17	0.76	0.99
(N = 186)	Rescued	0.45	0.24	0.22	0.84	0.99
Study 2b	Agent	0.14	0.20	0.27	0.95	0.99
(N = 266)	Rescued	0.36	0.18	0.35	0.97	0.99
Study 2c	Agent	0.03	0.15	0.42	0.99	0.99
(N = 446)	Rescued	0.92	0.14	0.52	0.99	0.99
	SES	0.03	0.14	0.52	0.99	0.99
Study 4a	Agent	0.08	0.07	0.98	0.99	0.99
(N = 912)	Rescued	0.13	0.07	0.98	0.99	0.99
	Guilty	0.03	0.07	0.98	0.99	0.99
Study 4b	Agent	0.12	0.07	0.97	0.99	0.99
(N = 811)	Rescued	0.24	0.07	0.97	0.99	0.99
	Guilty	0.03	0.07	0.97	0.99	0.99
Study 5	Agent	0.10	0.07	0.99	0.99	0.99
(N = 915)	Rescued	0.44	0.07	0.99	0.99	0.99
	Guilty	0.08	0.07	0.99	0.99	0.99

Note: Effect sizes are Cohen's f s. Effect size benchmarks are: small effect $f = .10$, medium effect $f = .25$ and large effect $f = .40$

General Discussion

Our studies have two overall novel findings. First, decisions made by robots were not always judged as worse than the same decisions made by humans. Rather, those decisions were consistently judged as worse when they failed to uphold the most salient moral standard. On the other hand, when the decision aligned with the salient moral standard, differences in judgments between a robot's decision and a human's decision were smaller or non-existent. In other words, a robot's "bad" decision was worse than a human's "bad" decision, but a robot's "good" decision was equal or only slightly worse than a human's "good" decision. Whether an agent was a human or a robot, moral judgments about their decisions followed a similar trend: rescuing the culpable was worse than rescuing the innocent, and rescuing two was better than rescuing one, all else being equal. Second, the culpability of the rescued party was a stronger indicator of moral judgments about rescue decisions than utilitarian concerns. In general, our participants judged decisions to rescue innocent people more positively than decisions to rescue the culpable, whether this decision conflicted with utilitarianism or not.

In the majority of our studies, only specific decisions by robots were judged as worse than similar decisions by humans. A robot's decision to rescue two culpable people rather than one innocent person was judged more negatively than the identical human decision, but there were no differences in judgments when the decision was to rescue the innocent person (Studies 1-2b). A decision violating both utilitarian and folk justice standards was considered worse if the decision was made by a robot rather than a human (Studies 4a-4b). When culpability was held constant between the parties that could be rescued, participants judged non-utilitarian decisions by robots as the most immoral alternative (Study 5). In our meta-analyses, we also found that

when comparing judgments of the same decision made by a human and a robot, differences were larger for decisions that violated a norm than decisions that abided by a norm.

The totality of our results suggest that folk justice was a stronger driver of judgments than utilitarianism. In our moderator meta-analysis, where we analyzed differences in judgments between two different decisions that either abided by or violated a moral concern, folk justice concerns had a larger effect on the difference than utilitarian concerns. Further, the effect of folk justice concerns was larger for robot than human agents, whereas the effect of utilitarianism was similar for both agents (though this interaction was not statistically significant when using a mixed-effects model; see Appendix E). In sum, our participants seem to have consistently put more weight on whether an agent rescued innocent or culpable people¹⁴ than on whether an agent rescued as many people as possible, and this was emphasized when the agent was a robot. However, again, the agents were judged in a similar direction: keeping other variables constant, it was better for both robots and humans to rescue the innocent over the culpable and to rescue two over one.

In individual studies, we only observed preferences for utilitarian decisions (by robots or humans) when the utilitarian decision did not conflict with folk justice concerns. In prior literature (e.g., Awad et al., 2018; Hidalgo, Orghian, Albo-Canals, de Almeida & Martin, 2021; Malle et al., 2015; Malle et al., 2016), many findings have indicated a preference for robots to make utilitarian decisions that included the competing deontological concern of “do not kill”. In some studies (e.g., Malle et al., 2015; 2016), this preference was stronger for robots than for humans. Why did we not observe a similar effect of utilitarianism over competing moral concerns when robots’ decisions were judged? The present results are inconclusive in this regard due to several differences in our vignettes compared to typical utilitarian dilemmas: a lower

possible maximum number of lives saved (i.e., two instead of five), the lack of an omission option (i.e., non-utilitarian decisions do not equate to doing nothing), and the added effect of culpability. Thus, it may be that the deontological motivation in typical trolley dilemmas is not strong enough, or that the utilitarian motivation for saving five people over one is stronger than the motivation in the present studies, or that in dilemmas with an omission option, people prefer robots that act rather than stand idly. Our studies highlight that people's judgments of different agents' decisions in moral dilemmas are contextual. Nonetheless, our participants consistently judged robot agent's decisions more negatively for failing to adhere to the most salient moral norm(s).

Our results are broadly aligned with previous studies suggesting that the moral evaluation of a robot decision-maker is moderated by the uncanny valley effect (i.e., the human-likeness of the robot; Laakasuo et al., 2021a; Grey & Wegner, 2012). In other words, decisions made by more human-like robots are judged more negatively than decisions made by machine-like robots. We only used a description of a robot with human-like appearance and intelligence in Studies 4a-5, as the results of Studies 2a-2b indicated that the decisions of humans were more likely to be judged differently from the decisions of robots if the robots were human-like. However, we did not replicate the observation that any decision would be judged more negatively when the agent was human-like (Laakasuo et al., 2021a). This may be due to, again, a different moral decision scenario than what is typical for psychological studies on utilitarian judgment.

Robot and human agents or their decisions being judged differently is a common trend in moral psychological studies (Bigman & Gray, 2018; Laakasuo et al., 2021a; Laakasuo et al., 2022; Malle et al., 2015; Malle et al., 2019). Our studies continue in this tradition with a focus on judgments of decisions by different agents (rather than judgments about agents themselves),

while ruling out certain alternative explanations. Because our vignettes did not have an omission option, participants' moral judgments of the agents' decisions were not confounded with preferences for or against inactivity (Baron & Ritov, 2004; Gawronski & Beer, 2016; Gawronski et al., 2017). Furthermore, because the scenarios described a lifeguard that acts in response to an event rather than being ordered to do so, the results are not affected by a perceived chain of command (see Laakasuo et al., 2022; Malle et al., 2019).

Our results suggest that the moral decisions of human-like robots are scrutinized more than the moral decisions of humans: “good” decisions are judged roughly equally, but “bad” decisions by robots are seen as worse than “bad” decisions by humans. Our studies cannot give a direct answer as to the reason behind this result, but we speculate that it may stem from differing expectations concerning the actions of robots and humans, based on, for instance, assumptions about differences in cognitive and emotional capacities or social roles. Perhaps robots, especially if they are human-like, are expected to be more “rational” than humans, and thus when they do something unwanted, the judgment is harsh (as also argued by Komatsu et al., 2021). While robots may be expected to be intelligent, they are generally viewed as less capable of emotion than humans (de Graaf & Malle, 2019; Gamez et al., 2020; Gray et al., 2007; Weisman et al., 2017). Thus, the robot should be able to act intelligently without being hindered by emotion: a life-or-death decision is emotionally taxing for a human, but presumably not for a robot. It could be that our participants could better empathize with a human forced to make a difficult decision than a robot in the same situation, and subsequently gave more positive judgments of the human agent's norm-violating decisions. It may also be that judgments about the human agent's rescue decisions are affected by the perceived risk of rescuing a drowning person, while a robot is not seen as risking anything. A robot making an unwanted decision in matters of life and death may

also make people question the point of having a robot make the decision in the first place: the justification for using robots is often a reduction in unwanted outcomes. In the context of our vignettes, people may view the function of a robotic rescuer to be to make “optimal” rescue decisions, and it failing in this task is worse than a human failing in it, as the use of a robot over a human needs additional justification. However, since there is no theory that could properly address these differences between humans and robots, further work is needed (Laakasuo et al., 2021b; 2021c).

In sum, when judging a human-like robot’s decisions in a life-or-death situation, people may view it as judging the decisions of an agent that should not have human faults, that does not evoke empathy, and/or that needs to prove its worth as a decision-maker more than a human does. Due to these factors, unwanted decisions by the robot may evoke more negative reactions than unwanted decisions by a human would evoke in the same situation. However, the moral norms that people seem to expect an agent to follow are the same for the robot and the human, leading to broadly similar judgments of the different agents’ decisions.

Limitations

Our studies show that culpability had a stronger effect than utilitarian concerns in a "two versus one" scenario, specifically in the context of lifeguarding. Our results concerning the relative weights of different moral concerns when judging non-human agents may not generalize to other kinds of moral decisions or contexts. More research is needed to investigate whether this effect can be found in scenarios with stronger utilitarian pressures (e.g., "five versus one"; Greene, Morelli, Lowenberg, Nystrom & Cohen, 2008), different kinds of culpability, or outside rescue contexts.

In general, our studies also face the usual limitations of experimental studies with self-report measures. First, the sample population is not usually representative in terms of age and other demographics or personality (whether samples are collected online or otherwise). However, we ran studies with different populations and cultures (students, native Finnish speakers, native English speakers, native Czech speakers), which adds robustness to our findings. Second, using self-report measures is associated with participants' tendency to acquiesce to, or obstruct, the experimenter's perceived objectives. We addressed this issue with fully randomized double-blind designs. Moreover, vignette and questionnaire-based research has shown to be reliable in studying moral judgments about the decisions of artificial agents (Malle et al., 2015; Bigman & Gray, 2018; Laakasuo et al., 2021a).

Finally, many of our early studies - especially studies conducted in laboratory settings - had relatively small sample sizes for detecting minor interaction effects. However, the results of the later higher-powered pre-registered studies match those of the earlier ones, even with slightly different measures, and our findings are qualified by the meta-analyses.

Future Studies

Future studies should more generally examine different moral concerns. Although trolley-type dilemmas are common in moral psychology, moral judgment should not be reduced into only utilitarian versus deontological judgment. While there is also little knowledge of how people would weigh different concerns in “utilitarian calculus”, moral psychology should diversify to fill this gap (Chapman, 2018; Voiklis & Malle, 2017). Our studies highlight a potentially under-researched issue: the character perception of moral *patients* rather than moral *agents*. As shown in our studies, the way people perceive the rescued party (i.e., their culpability and their social status) influences moral judgments. Regarding culpability, more research should

examine people's motivations for deciding who should be rescued. Additionally, our studies 2a and 2b, in conjunction with those of Laakasuo et al. (2021a), imply that the stated physical and/or cognitive human-likeness of a robot affects moral judgments of their decisions, which warrants further examination.

Finally, future research is needed to determine which exact subcomponents of moral judgment are of main interest (blame, harm, "moral acceptance", or causality inferences, to name a few). Our results using a scale measuring general moral approval (Studies 1-2c) and single-item measures of acceptability and wrongness (Studies 3-5) were generally consistent with one another. However, there were some deviations to the patterns we observed when using single-item measures of, e.g., perceived harm instead of acceptability or wrongness (see Appendix B). While many different variables we used were interchangeable with one another, there are nuances to different moral judgments, and focusing solely on ratings of wrongness or acceptability may hide interesting patterns in other types of judgment.

Conclusion

In sum, our studies - using a novel moral dilemma scenario of a rescue situation - show that human and robot agents' decisions are judged differently primarily when those decisions violate the most salient moral norm in a situation. These differences in judgments are sensitive to both utilitarian concerns and the culpability of the rescued party, something previously underexplored in moral psychology. Moreover, culpability seemed to matter more than strictly utilitarian concerns. People consistently showed a preference for rescuing the innocent over saving a larger number of lives when these options were mutually exclusive. Our studies also addressed the issue of omission or inactivity bias (Gawronski et al., 2017) by having each potential option in the dilemma consist of an active choice (i.e., rescuing one or more people).

We further highlight the need for a focus on different core factors of moral dilemma situations, and for examining moral concerns that are intuitively relevant to many people but that have mostly stayed out of moral psychology due to established methods of asking questions. Especially in the context of judgments about artificial agents making moral choices, it is important to shed light on the various factors that people care about.

All in all, in our studies, we showed that judgments of robots as decision-makers can offer new insights into the questions of human morality and can help elucidate newer research traditions as well. We suggest that moral psychological vignettes incorporate new moral principles and agents to expand and refine the growing understanding of moral cognition.

References

- Addison, A., Bartneck, C., & Yogeeswaran, K. (2019, January). Robots can be more than black and white: examining racial bias towards robots. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society* (pp. 493-498).
- Awad, E., Dsouza, S., Kim, R., Schulz, J., Henrich, J., Shariff, A., Bonnefon, J.-F., & Rahwan, I. (2018). The moral machine experiment. *Nature*, 563(7729), 59–64.
<https://doi.org/10.1038/s41586-018-0637-6>
- Baron, J., & Ritov, I. (2004). Omission bias, individual differences, and normality. *Organizational Behavior and Human Decision Processes*, 94(2), 74–85.
<https://doi.org/10.1016/j.obhdp.2004.03.003>
- Bigman, Y. E., & Gray, K. (2018). People are averse to machines making moral decisions. *Cognition*, 181, 21–34. <https://doi.org/10.1016/j.cognition.2018.08.003>
- Bigman, Y. E., Waytz, A., Alterovitz, R., & Gray, K. (2019). Holding Robots Responsible: The Elements of Machine Morality. *Trends in Cognitive Sciences*, 23(5), 365–368.
<https://doi.org/10.1016/j.tics.2019.02.008>
- Campaign to Stop Killer Robots. (2019). *Key elements of a treaty on fully autonomous weapons*.
- Carlsmith, K. M., & Darley, J. M. (2008). Psychological aspects of retributive justice. In *Advances in Experimental Social Psychology* (Vol. 40, pp. 193–236). Academic Press.
[https://doi.org/10.1016/S0065-2601\(07\)00004-4](https://doi.org/10.1016/S0065-2601(07)00004-4)
- Chapman, H. A. (2018). A component process model of disgust, anger, and moral judgment. *Atlas of moral psychology*, 70.
- Champely, S. (2020). *pwr: Basic Functions for Power Analysis*. <https://CRAN.R-project.org/package=pwr>

- Christensen, J. F., & Gomila, A. (2012). Moral dilemmas in cognitive neuroscience of moral decision-making: A principled review. *Neuroscience & Biobehavioral Reviews*, 36(4), 1249–1264. <https://doi.org/10.1016/j.neubiorev.2012.02.008>
- Curry, O. S., Jones Chesters, M., & Van Lissa, C. J. (2019). Mapping morality with a compass: Testing the theory of ‘morality-as-cooperation’ with a new questionnaire. *Journal of Research in Personality*, 78, 106–124. <https://doi.org/10.1016/j.jrp.2018.10.008>
- Davison, N. (2017). A legal perspective: Autonomous weapon systems under international humanitarian law. *UNODA Occasional Papers*, 30, 14.
- de Graaf, M. M. A., & Malle, B. F. (2019). People’s Explanations of Robot Behavior Subtly Reveal Mental State Inferences. *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 239–248. <https://doi.org/10.1109/HRI.2019.8673308>
- Gamez, P., Shank, D. B., Arnold, C., & North, M. (2020). Artificial virtue: The machine question and perceptions of moral character in artificial moral agents. *AI & SOCIETY*, 35(4), 795–809. <https://doi.org/10.1007/s00146-020-00977-1>
- Gawronski, B., Armstrong, J., Conway, P., Friesdorf, R., & Hütter, M. (2017). Consequences, norms, and generalized inaction in moral dilemmas: The CNI model of moral decision-making. *Journal of Personality and Social Psychology*, 113(3), 343–376. <https://doi.org/10.1037/pspa0000086>
- Gawronski, B., & Beer, J. S. (2016). What makes moral dilemma judgments “utilitarian” or “deontological”? *Social Neuroscience*, 12(6), 626–632. <https://doi.org/10.1080/17470919.2016.1248787>
- Gray, H. M., Gray, K., & Wegner, D. M. (2007). Dimensions of mind perception. *Science*, 315(5812), 619. <https://doi.org/10.1126/science.1134475>

Gray, K., & Wegner, D. M. (2012). Feeling robots and human zombies: Mind perception and the uncanny valley. *Cognition*, 125(1), 125–130.

<https://doi.org/10.1016/j.cognition.2012.06.007>

Greene, J. D. (2007). Why are VMPFC patients more utilitarian? A dual-process theory of moral judgment explains. *Trends in Cognitive Sciences*, 11(8), 322–323.

<https://doi.org/10.1016/j.tics.2007.06.004>

Greene, J. D., Morelli, S. A., Lowenberg, K., Nystrom, L. E., & Cohen, J. D. (2008). Cognitive load selectively interferes with utilitarian moral judgment. *Cognition*, 107(3), 1144–1154.

<https://doi.org/10.1016/j.cognition.2007.11.004>

Haidt, J., & Graham, J. (2007). When Morality Opposes Justice: Conservatives Have Moral Intuitions that Liberals may not Recognize. *Social Justice Research*, 20(1), 98–116.

<https://doi.org/10.1007/s11211-007-0034-z>

Hancock, P. A., Billings, D. R., Schaefer, K. E., Chen, J. Y. C., de Visser, E. J., & Parasuraman, R. (2011). A Meta-Analysis of Factors Affecting Trust in Human-Robot Interaction.

Human Factors, 53(5), 517–527. <https://doi.org/10.1177/0018720811417254>

Hidalgo, C. A., Orghian, D., Albo-Canals, J., de Almeida, F., & Martin, N. (2021). *How humans judge machines*. MIT Press.

Kahane, G., Everett, J. A. C., Earp, B. D., Farias, M., & Savulescu, J. (2015). ‘Utilitarian’ judgments in sacrificial moral dilemmas do not reflect impartial concern for the greater

good. *Cognition*, 134, 193–209. <https://doi.org/10.1016/j.cognition.2014.10.005>

Komatsu, T., Malle, B. F., & Scheutz, M. (2021, March). Blaming the reluctant robot: parallel blame judgments for robots in moral dilemmas across US and Japan. In *Proceedings of*

the 2021 ACM/IEEE International Conference on Human-Robot Interaction (pp. 63-72).

<https://doi.org/10.1145/3434073.3444672>

Kwon, M., Jung, M. F., & Knepper, R. A. (2016). Human expectations of social robots. *2016*

11th ACM/IEEE International Conference on Human-Robot Interaction (HRI), 463–464.

<https://doi.org/10.1109/HRI.2016.7451807>

Laakasuo, M., Palomäki, J., Kunnari, A., Rauhala, S., Drosinou, M., Halonen, J., Lehtonen, N.,

Koverola, M., Repo, M., Sundvall, J., Visala, A., & Francis, K. B. (2022). Moral

psychology of nursing robots: Exploring the role of robots in dilemmas of patient

autonomy. *European Journal of Social Psychology*, 00, 1– 21.

<https://doi.org/10.1002/ejsp.2890>

Laakasuo, M., Palomäki, J., & Köbis, N. (2021a). Moral uncanny valley: A robot's appearance

moderates how its decisions are judged. *International Journal of Social Robotics*.

<https://doi.org/10.1007/s12369-020-00738-6>

Laakasuo, M., Sundvall, J. R., Berg, A., Drosinou, M., Herzon, V., Kunnari, A., Koverola, M.,

Repo, M., Saikkonen, T., & Palomäki, J. (2021b). Moral Psychology and Artificial

Agents (Part One): Ontologically Categorizing Bio-Cultural Humans. In S. Thompson

(Ed.), *Machine Law, Ethics, and Morality in the Age of Artificial Intelligence* (pp. 166-

188). IGI Global. <https://doi.org/10.4018/978-1-7998-4894-3.ch010>

Laakasuo, M., Sundvall, J. R., Berg, A., Drosinou, M., Herzon, V., Kunnari, A., Koverola, M.,

Repo, M., Saikkonen, T., & Palomäki, J. (2021c). Moral Psychology and Artificial

Agents (Part Two): The Transhuman Connection. In S. Thompson (Ed.), *Machine Law,*

Ethics, and Morality in the Age of Artificial Intelligence (pp. 189-204). IGI Global.

<https://doi.org/10.4018/978-1-7998-4894-3.ch011>

Malle, B. F. (2021). Moral judgments. *Annual Review of Psychology*, 72, 293–318.

<https://doi.org/10.1146/annurev-psych-072220-104358>

Malle, B. F., Magar, S. T., & Scheutz, M. (2019). AI in the sky: How people morally evaluate human and machine decisions in a lethal strike dilemma. In M. I. Aldinhas Ferreira, J. Silva Sequeira, G. Singh Virk, M. O. Tokhi, & E. E. Kadar (Eds.), *Robotics and Well-Being* (Vol. 95, pp. 111–133). Springer International Publishing.

https://doi.org/10.1007/978-3-030-12524-0_11

Malle, B. F., Scheutz, M., Arnold, T., Voiklis, J., & Cusimano, C. (2015). Sacrifice one for the good of many?: People apply different moral norms to human and robot agents. *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*, 117–124. <https://doi.org/10.1145/2696454.2696458>

Malle, B. F., Scheutz, M., Forlizzi, J., & Voiklis, J. (2016). Which robot am I thinking about? The impact of action and appearance on people’s evaluations of a moral robot. *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 125–132. <https://doi.org/10.1109/HRI.2016.7451743>

Marcus, G., & Davis, E. (2019). *Rebooting AI: Building artificial intelligence we can trust*. Vintage.

Mitchell, M. (2019). *Artificial intelligence: A guide for thinking humans*. Penguin UK.

OECD (2022). Average wages (indicator). doi: 10.1787/cc3e1387-en (Accessed on 05 June 2022)

O’Neil, C. (2016). *Weapons of math destruction: How big data increases inequality and threatens democracy*. Crown.

- Palomäki, J., Kunnari, A., Drosinou, M., Koverola, M., Lehtonen, N., Halonen, J., Repo, M., & Laakasuo, M. (2018). Evaluating the replicability of the uncanny valley effect. *Heliyon*, 4(11). <https://doi.org/10.1016/j.heliyon.2018.e00939>
- Phillips, E., Ullman, D., de Graaf, M. M. A., & Malle, B. F. (2017). What Does A Robot Look Like?: A Multi-Site Examination of User Expectations About Robot Appearance. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 61(1), 1215–1219. <https://doi.org/10.1177/1541931213601786>
- Schein, C., & Gray, K. (2018). The Theory of Dyadic Morality: Reinventing Moral Judgment by Redefining Harm. *Personality and Social Psychology Review*, 22(1), 32–70. <https://doi.org/10.1177/1088868317698288>
- Starr, N. D., Malle, B., & Williams, T. (2021). I Need Your Advice... Human Perceptions of Robot Moral Advising Behaviors. arXiv preprint arXiv:2104.06963.
- Syrdal, D. S., Dautenhahn, K., Walters, M. L., & Koay, K. L. (2008). Sharing Spaces with Robots in a Home Scenario—Anthropomorphic Attributions and their Effect on Proxemic Expectations and Evaluations in a Live HRI Trial. *In Proceedings of the 2008 AAAI Fall Symposium*, 116–123.
- Trovato, G., Lucho, C., & Paredes, R. (2018). She’s Electric—The Influence of Body Proportions on Perceived Gender of Robots across Cultures. *Robotics*, 7(3). <https://doi.org/10.3390/robotics7030050>
- VanVoorhis, C. R. W., & Morgan, B. L. (2007). Understanding Power and Rules of Thumb for Determining Sample Sizes. *Tutorials in Quantitative Methods for Psychology*, 3(2), 43–50.

- Viechtbauer, W. (2010). Conducting meta-analyses in R with the metafor package. *Journal of Statistical Software*, 36(3), 1–48.
- Voiklis, J., Kim, B., Cusimano, C., & Malle, B. F. (2016). Moral judgments of human vs. Robot agents. *2016 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, 775–780. <https://doi.org/10.1109/ROMAN.2016.7745207>
- Voiklis, J., & Malle, B. F. (2017). Moral cognition and its basis in social cognition and social regulation. In K. Gray & J. Graham (Eds.), *Atlas of moral psychology* (pp. 108–120). Guilford Press.
- Wallach, W., Allen, C.: *Moral Machines: Teaching Robots Right From Wrong*. Oxford University Press (2008)
- Walker, J. (2019, November 21). *Search and Rescue Robots – Current Applications on Land, Sea, and Air*. EMERJ - The AI Research and Advisory Company. <https://emerj.com/ai-sector-overviews/search-and-rescue-robots-current-applications/>
- Waytz, A., Heafner, J., & Epley, N. (2014). The mind in the machine: Anthropomorphism increases trust in an autonomous vehicle. *Journal of Experimental Social Psychology*, 52, 113–117. <https://doi.org/10.1016/j.jesp.2014.01.005>
- Weisman, K., Dweck, C. S., & Markman, E. M. (2017). Rethinking people’s conceptions of mental life. *Proceedings of the National Academy of Sciences*, 114(43), 11374–11379. <https://doi.org/10.1073/pnas.1704347114>
- Yogeeswaran, K., Złotowski, J., Livingstone, M., Bartneck, C., Sumioka, H., & Ishiguro, H. (2016). The interactive effects of robot anthropomorphism and robot ability on perceived threat and support for robotics research. *Journal of Human-Robot Interaction*, 5(2), 29–47.

Złotowski, J., Proudfoot, D., Yogeeswaran, K., & Bartneck, C. (2015). Anthropomorphism: opportunities and challenges in human–robot interaction. *International journal of social robotics*, 7(3), 347-360.

Złotowski, J., Yogeeswaran, K., & Bartneck, C. (2017). Can we control it? Autonomous robots threaten human identity, uniqueness, safety, and resources. *International Journal of Human-Computer Studies*, 100, 48-54.

Footnotes

1) Both here and in discussing our own studies we will use the word “utilitarianism” as a shorthand for “utilitarianism as it is commonly operationalized in psychological studies”. We do not claim that the participants who preferred a utilitarian option in the studies we cite or our own studies necessarily held a comprehensive utilitarian moral philosophy. What we do claim is that they exhibited a utilitarian-like preference.

2) Income was self-reported on a 10-point ladder scale in relation to average income in the participants’ home country. The options were “I belong to the wealthiest 5%”, “I belong to the wealthiest 10%”, “I belong to the wealthiest 20%”, “I belong to the wealthiest 35%”, “My income is slightly above average”, “My income is average”, “My income is almost average”, “I belong to the poorest 35%”, “I belong to the poorest 20%”, and “I belong to the poorest 10%”. In 2017, when the data was collected, average annual wages in the country were \$45 339 (OECD, 2022).

3) Income was self-reported on a 10-point ladder scale in relation to average income in the participants’ home country. The options were “I belong to the wealthiest 5%”, “I belong to the wealthiest 10%”, “I belong to the wealthiest 20%”, “I belong to the wealthiest 33%”, “My income is slightly above average”, “My income is average”, “My income is almost average”, “I belong to the poorest 33%”, “I belong to the poorest 20%”, and “I belong to the poorest 10%”. In 2019, when the data was collected, average annual wages in the country were \$46 249 (OECD, 2022).

4) Income was self-reported on an 8-point ladder scale in relation to average income in the participants’ home country. The options were “About the lowest 5%”, “About the lowest 5-20%”, “Slightly below average”, “Average”, “Slightly above average”, “About the highest 20-

5%”, “About the highest 5-1%”, and “About the highest 1%”. In 2019, when the data was collected, average annual wages in the country were \$46 249 (OECD, 2022).

5) Income was self-reported on a 6-point ladder scale in relation to average income in the participants’ area of residence. The options were “Top 1%”, “Top 20%”, “Above average”, “About average”, “Below average”, and “Lowest 20%”. Since the sample consisted of participants from several countries, it is not possible to provide a general average income for context. A majority of the participants (84%) were from the United Kingdom. In 2019, when the data was collected, average annual wages in the United Kingdom were \$47 937 (OECD, 2022).

6) The participants were asked two questions: whether they make their own shoes and whether they make their own clothes. Participants who answered "yes" to both were excluded due to implausibility. We acknowledge that it is not impossible for someone to be both a tailor and a shoemaker, but we posit that in an online questionnaire, a person claiming to be both is most likely a person "straight-lining" their responses.

7) Income was self-reported on a 6-point ladder scale in relation to average income in the participants’ area of residence. The options were “Top 1%”, “Top 20%”, “Above average”, “About average”, “Below average”, and “Lowest 20%”. Since the sample consisted of participants from several countries, it is not possible to provide a general average income for context. A majority of the participants (67%) were from the United Kingdom. In 2020, when the data was collected, average annual wages in the United Kingdom were \$47 147 (OECD, 2022).

8) We also conducted an exploratory analysis using the mean of all 7 single-item moral judgment measures as the dependent variable. The wrongness, blame, punishment, and harm items were reverse-coded. Cronbach's α was .87 for judgments of rescuing the motorboaters and .90 for

judgments of rescuing the fisherman, both acceptable. The results of the analysis were essentially the same as the results of the analyses reported here.

9) Income was self-reported on a 6-point ladder scale in relation to average income in the participants' area of residence. The options were "Top 1%", "Top 20%", "Above average", "About average", "Below average", and "Lowest 20%". The data was collected in 2021; the latest available OECD statistics are from 2020. Average annual wages in the country in 2020 were \$29 885 (OECD, 2022).

10) We also conducted an exploratory analysis using the mean of all 7 single-item moral judgment measures as the dependent variable. The wrongness, blame, punishment, and harm items were reverse-coded. Cronbach's α was .77 for judgments of rescuing the motorboaters and .84 for judgments of rescuing the fisherman, both acceptable. The results of the analysis were essentially the same as the results of the analyses reported here.

11) The age calculations were conducted with one additional exclusion: a single participant who listed their age as being over 200 years. This participant was not excluded from any other demographic information or analyses reported here, as improbable responses to this demographic question were not a part of our exclusion criteria, and this could plausibly be a simple typing error.

12) Income was self-reported on a 6-point ladder scale in relation to average income in the participants' area of residence. The options were "Top 1%", "Top 20%", "Above average", "About average", "Below average", and "Lowest 20%". Since the sample consisted of participants from several countries, it is not possible to provide a general average income for context. A majority of the participants (72%) were from the United Kingdom. In 2020, when the data was collected, average annual wages in the United Kingdom were \$47147 (OECD, 2022).

13) We also conducted an exploratory analysis using the mean of all 7 single-item moral judgment measures as the dependent variable. The wrongness, blame, punishment, and harm items were reverse-coded. Cronbach's α was .79 for judgments of rescuing the motorboaters and .85 for judgments of rescuing the fisherman, both acceptable. The results of the analysis were essentially the same as the results of the analyses reported here.

14) We argue that our results tentatively support the conclusion that our participants' culpability-based judgments were driven more by a wish to protect the innocent, rather than "punish" the wrongdoers. As stated earlier, both of these motives fit what we termed a folk sense of justice, and we did not specifically ask our participants about their motivations. However, the results of Study 5 work as an indirect test. In Study 5, when both parties were culpable, participants gave more positive judgments for decisions to rescue the party with more people. In other words, it seems that the participants still preferred it when the agent rescued more people, even when the agent could have "punished" a higher number of culpable people by rescuing the party with just one person.

Appendix A: Vignettes

Study 1

“It is a cold day in early spring and the temperature is close to zero. Ice is forming on a popular beach, which focuses on boating and fishing. A **[lifeguard / robot lifeguard]** patrols this empty beach and looks out to sea. **[He / It]** sees a motorboat with two people on the open sea. Their maneuvers seem risky and uncontrollable. The **[lifeguard has significant experience of his work as a coastguard, and based on the unstable boating style, he concludes that the boaters are drunk. / lifeguard robot has an extensive database of safe and risky boating behavior, as well as typical effects of alcohol on behavior. Based on its database, the robot concludes that the boaters are drunk.]** A fisherman with a smaller rowboat is also nearby. The motorboat is driving in the direction of the rowboat, and the boats seem to be passing each other.

As the boats approach each other, the speed of the motorboat increases further and its steering becomes even more unstable than before. Just when the boats are passing each other, the motorboat makes a sudden turn and collides with the side of the rowboat. Both boats are damaged in the collision, and the fisherman falls into the water. After the collision, the motorboat drifts further away from the rowboat until it capsizes, dropping both motorboaters into the water. Both the fisherman and motorboaters are now far apart from each other in the cold water.

The **[lifeguard / lifeguard robot]** estimates that **[he/it]** has enough time to make one trip to one of the boats so that its crew can be safely brought ashore. **[He has the necessary aids to transport two people. / It can transport two people to the beach at once if needed.]** The **[lifeguard / lifeguard robot]** also knows that surviving in the cold water until the second rescue

search is highly unlikely. The **[lifeguard / lifeguard robot]** must make a decision on whether to try to rescue the fisherman or the two motorboaters first, knowing that the success of the second rescue is highly unlikely. The **lifeguard / lifeguard robot** decides to rescue the **[motorboaters / fisherman]** first.”

Study 2a

“It's the year 2070. Robots and robotics have advanced significantly, and various kinds of robots are used widely in different tasks. Among these tasks are lifeguarding and providing first aid, especially in remote places, where there's an increased need for emergency services. The latest model of lifeguarding robots have recently been put into service in USA. They have been assigned to some of the most popular beaches and coastlines to increase safety. However, they are not being used at every beach; when necessary, some beaches are being guarded by human lifeguards.

It's a cold day in early Spring, and the temperature is near 32 degrees Fahrenheit / 0 degrees Celsius. Some ice is forming in the coastline on a beach popular for recreational boating and fishing. **[On this beach, a human lifeguard is monitoring the situation. The lifeguard is in good enough physical and mental shape to carry out their task. The lifeguard is constantly observing the coastline area, and capable of recognizing, under their own discretion, potentially dangerous situations. / On this beach, a lifeguarding robot is monitoring the situation. The robot resembles an advanced human-like android. It's constantly observing the coastline area, and capable of recognizing potentially dangerous situations. The robot has a human-like ability of independent decision-making, and it's functioning under its own discretion. / On this beach, a lifeguarding robot is monitoring the situation. The robot resembles a mechanical lifeguard. It's constantly scanning the coastline**

area, and capable of recognizing potentially dangerous situations. However, the robot does not have a human-like ability of independent decision-making, and is functioning purely based on its programming.]

The beach is empty. The **[lifeguard / robot]** notices a motorboat in the open sea, with two people on board. The steering of the motorboat appears reckless and uncontrolled. **[Based on their extensive experience, and the uncontrolled steering of the motorboat, the lifeguard concludes that the people on the motorboat are drunk. / The lifeguarding robot has a database on what safe and risky steering behaviour look like, as well as on how alcohol affects steering behaviour. Based on its database, the robot concludes that the people on the motorboat are drunk.]** Nearby there is also a fisherman on a smaller rowboat. The motorboat is headed towards the rowboat, and the boats appear to be about to pass one another.

As the boats approach one another the motorboat speeds up and its steering becomes even more unstable. Just as the boats are about to pass one another, the motorboat makes an unexpected turn and collides with the rowboat, hitting its side. Both boats are badly damaged on impact, and the fisherman is thrown in the sea. As a result of the crash, the motorboat drifts further away from the rowboat, and finally capsizes, throwing the two motorboaters in the sea. Both parties are now in the freezing water, far apart from each another.

The **[lifeguard / lifeguarding robot]** estimates that **[they / it]** will be able to reach either boat, and safely bring its crew back to the shore. The **[lifeguard / robot]** also knows that it's essentially impossible for anyone to survive in the freezing water long enough for them to make a second trip to save the other party. Thus, the **[lifeguard / lifeguarding robot]** needs to make a decision on whom to save -- the fisherman or the motorboaters -- knowing that any attempt to save everyone is essentially impossible.

[The lifeguard / lifeguarding robot] decides to save [the fisherman / the motorboaters.]”

Study 2b

“It's the year 2070. Robots and robotics have advanced significantly, and various kinds of robots are used widely in different tasks. Among these tasks are lifeguarding and providing first aid, especially in remote places, where there's an increased need for emergency services. The latest model of lifeguarding robots have recently been put into service in USA. They have been assigned to some of the most popular beaches and coastlines to increase safety. However, they are not being used at every beach; when necessary, some beaches are being guarded by human lifeguards.

It's a cold day in early Spring, and the temperature is near 32 degrees Fahrenheit / 0 degrees Celsius. Some ice is forming in the coastline on a beach popular for recreational boating and fishing. **[On this beach, a human lifeguard is monitoring the situation. The lifeguard is in good enough physical and mental shape to carry out their task. The lifeguard is constantly observing the coastline area, and capable of recognising, under their own discretion, potentially dangerous situations. / On this beach, a lifeguarding robot is monitoring the situation. The robot resembles an advanced human-like android. It's constantly observing the coastline area, and capable of recognising potentially dangerous situations. However, the robot does not have a human-like ability of independent decision-making, and is functioning purely based on its programming. / On this beach, a lifeguarding robot is monitoring the situation. The robot resembles a mechanical lifeguard. It's constantly scanning the coastline area, and capable of recognising potentially dangerous situations. The robot has a human-like ability of independent decision-making,**

and it's functioning under its own discretion.]

The beach is empty. The **[lifeguard / robot]** notices a motorboat in the open sea, with two people on board. The steering of the motorboat appears reckless and uncontrolled. **[Based on their extensive experience, and the uncontrolled steering of the motorboat, the lifeguard concludes that the people on the motorboat are drunk. / The lifeguarding robot has a database on what safe and risky steering behaviour look like, as well as on how alcohol affects steering behaviour. Based on its database, the robot concludes that the people on the motorboat are drunk.]** Nearby there is also a fisherman on a smaller rowboat. The motorboat is headed towards the rowboat, and the boats appear to be about to pass one another.

As the boats approach one another the motorboat speeds up and its steering becomes even more unstable. Just as the boats are about to pass one another, the motorboat makes an unexpected turn and collides with the rowboat, hitting its side. Both boats are badly damaged on impact, and the fisherman is thrown in the sea. As a result of the crash, the motorboat drifts further away from the rowboat, and finally capsizes, throwing the two motorboaters in the sea. Both parties are now in the freezing water, far apart from each another.

The **[lifeguard / lifeguarding robot]** estimates that **[they / it]** will be able to reach either boat, and safely bring its crew back to the shore. The **[lifeguard / robot]** also knows that it's essentially impossible for anyone to survive in the freezing water long enough for them to make a second trip to save the other party. Thus, the **[lifeguard / lifeguarding robot]** needs to make a decision on whom to save -- the fisherman or the motorboaters -- knowing that any attempt to save everyone is essentially impossible.

[The lifeguard / lifeguarding robot] decides to save [the fisherman / the motorboaters.]”

Study 2c

“It's the year 2070. Robots and robotics have advanced significantly, and various kinds of robots are used widely in different tasks. Robots are also partially responsible for lifeguarding and providing first aid, especially in remote places, where there's an increased need for emergency services. The latest model of lifeguarding robots have recently been put into service in USA. They have been assigned to some of the most popular beaches and coastlines to increase safety. However, they are not being used at every beach; when necessary, some beaches are being guarded by human lifeguards.

It's a cold day in early Spring, and the temperature is near 32 degrees Fahrenheit / 0 degrees Celsius. Some ice is forming in the coastline on a beach popular for recreational boating and fishing. On this beach, a **[lifeguarding robot / human lifeguard]** is monitoring the situation. **[The robot resembles an advanced human-like android / The robot resembles a mechanical lifeguard / The lifeguard is in good enough physical and mental shape to carry out their task.]** **[It's / The lifeguard is]** constantly observing the coastline area, and capable of recognising potentially dangerous situations. **[The robot has a human-like ability of independent decision-making, and it's functioning under its own discretion / However, the robot does not have a human-like ability of independent decision-making, and is functioning purely based on its programming].**

The beach is empty. The **[robot / lifeguard]** notices a motorboat in the open sea, with a single person on board. The **[robot/ lifeguard]** also recognises this person as someone with **[high / low]** status, social class, and income. The steering of the motorboat appears reckless and uncontrolled. **[The lifeguarding robot has a database on what safe and risky steering**

behaviour look like, as well as on how alcohol affects steering behaviour. Based on its database, the robot concludes that the person on the boat is drunk. / Based on their extensive experience, and the uncontrolled steering of the motorboat, the lifeguard concludes that the person on the motorboat is drunk.] Nearby there is also a fisherman on a smaller rowboat. The motorboat is headed towards the rowboat, and as they approach one another the motorboat speeds up and its steering becomes even more unstable.

Just as the boats are about to pass one another, the motorboat makes an unexpected turn and collides with the rowboat, hitting its side. Both boats are badly damaged on impact, and the fisherman is thrown in the sea. As a result of the crash, the motorboat drifts further away from the rowboat, and finally capsizes, throwing the person on board in the sea. Both the fisherman and the person on the motorboat are now in the freezing water, far apart from one another.

The **[lifeguarding robot/ lifeguard]** estimates that it will be able to reach one of the people in the sea and safely bring that person back to shore. The **[robot / lifeguard]** also knows that it's essentially impossible for anyone to survive in the freezing water long enough for the **[robot / them]** to make a second trip to save the other person. Thus, the **[lifeguarding robot/ human lifeguard]** needs to make a decision on whom to save -- the fisherman or the person on the motorboat -- knowing that any attempt to save both is essentially impossible.

The **[lifeguarding robot/ lifeguard]** decides to save the **[fisherman / person on the motorboat]."**

Study 3

"It is the year 2070. A lifeguard robot on a beach notices a motorboat with two people on board far away on the sea. The boaters seem to be driving in a risky and uncontrolled way. Based

on its database, the robot deduces that the motorboaters are drunk. Slightly closer to the shore, the robot also notices a fisherman in a small rowboat. Just as the boats are about to pass one another, the motorboat makes a sudden turn and crashes into the side of the rowboat. After the collision, the motorboat drifts further away from the rowboat, and finally capsizes, throwing both people on board in the sea. Both the fisherman and the motorboaters are now in cold water, far away from each other. The robot estimates that it only has time to make one trip to either party to get them safely back to the shore.”

[Response to DVs here: “How acceptable would it be for the robot to save the fisherman / the motorboaters”]

“Now, please imagine the same situation, but the motorboat is moving in a controlled manner and the motorboaters are quite clearly sober. Then, there is a technical fault in the motor, and the motorboat crashes into the fisherman's boat.”

[Response to DVs here: “How acceptable would it be for the robot to save the fisherman / the motorboaters”]

Appendix B: Additional Analyses

Study 3

Within-subjects data simulation. For both of the comparisons reported (difference in acceptability ratings between rescue decisions in a) the vignette where the motorboaters are culpable and the fisherman is innocent and b) the vignette where both parties are innocent), we ran 100000 iterations of randomly sampled possible within-subjects data structures. Because of uneven numbers of responses to the questions, for each pair of questions, we set the lower number of responses as the maximum number of pairs. Then, we sampled the responses to each of the questions without replacement into a new mock data set, where each response to the

question with fewer responses was randomly paired with a response to the other question. Then, a paired-samples t-test was conducted for each data set created this way. The distributions of t-values are presented in Supplementary Figures 1 and 2. These distributions can be seen as presenting a range of possible test statistics if participant identifiers had been available. Based on the distributions, the results of the analysis under the false assumption of independent samples are very close to what the results likely would have been when using a paired-samples test with participant identifiers.

Note that the correlation between the “paired” samples in our simulations was 0 on average (due to random sampling and a high number of iterations). A positive correlation in the real data would lead to a smaller t-value, in which case the independent samples analysis would produce an inflated t-value. However, based on examinations of similar within-subjects analyses from Studies 4a, 4b, and 5, it seems that this is likely not the case with our data. In Studies 4a and 4b, examining only the situations where a robot agent rescued either two culpable motorboaters or an innocent fisherman (matching the first pair of questions in Study 3), the correlations between the paired samples were negative (Study 4a: $r = -0.14$; Study 4b: $r = -0.32$). In Study 5, examining only the situations where a robot agent rescued either two motorboaters or a fisherman when both parties were innocent (matching the second pair of questions in Study 3), the correlation was essentially null ($r = 0.01$). A negative correlation would lead to larger t-values (in absolute terms), and a null correlation matches with the mock paired samples data: thus, the independent-samples analysis produces t-values that are either similar or more conservative compared to what the t-values would likely be in a paired-samples analysis. In conclusion, the independent-samples analysis produces results that are most likely similar to what would have been achieved if participant identifiers were available.

Supplementary Figure 1. Density plot and histogram of t-values from simulated analyses comparing responses to the rescue decisions in Study 3 when the motorboaters were culpable and the fisherman was innocent. The histogram has a bin width of 0.10.

Supplementary Figure 2. Density plot of t-values from simulated analyses comparing responses to the rescue decisions in Study 3 when both parties were innocent. The histogram has a bin width of 0.01.

Study 4a

Analysis of secondary dependent variables. We entered participants' ratings of permissibility, blameworthiness, praiseworthiness, punishability and harm into the same mixed model as we used when analyzing the main results (see Supplementary Table 1 and Supplementary Figure 3). Overall, the mean patterns for all DVs echoed those of the mean patterns for acceptability and wrongness, and a majority of the significant main effects and interactions were similar to those we observed for acceptability and wrongness. A notable exception happened with the main effect of the culpability condition, which was significant for ratings of blameworthiness and punishability. A main effect of this condition implies that agents were judged more harshly overall based simply on whether the scenario had a conflict between utilitarianism and folk justice or not. However, the culpability condition was involved in significant interactions for both blameworthiness and punishability. The largest effect for these two variables seems to be the interaction effect between culpability and the rescue decision, with decisions to rescue the innocent seen as less blameworthy and less punishable.

Our pre-registered predictions included hypotheses concerning blame, praise and punishment. Specifically, we predicted that the human agent would be blamed more for rescuing the innocent fisherman than the culpable motorboaters; that the human agent would be praised more for rescuing the culpable motorboaters than the innocent fisherman; and that participants would want to punish the robot agent more for rescuing the culpable rather than the innocent party (regardless of whether it was utilitarian or not).

The hypothesis concerning blame received no support: the human agent's decision to rescue the culpable motorboaters was judged as more blameworthy than rescuing the innocent fisherman ($B = 1.07$, 95% CI = [0.81, 1.32], $t(908) = 8.28$, $p < .001$, $d = 0.55$) in alignment with

the main results of acceptability and wrongness. This hypothesis was borne out of our earlier observations of an exploratory blame variable in Study 1 (not reported here) suggesting a trend towards more blame towards the human agent for non-utilitarian choices.

Consistently with the blame results, the hypothesis concerning praise received no support: the human agent was considered less praiseworthy for rescuing the culpable motorboaters than the innocent fisherman ($B = -1.24$, 95% CI = $[-1.47, -1.01]$, $t(908) = -10.43$, $p < .001$, $d = -0.69$).

The hypothesis concerning punishment received support: the android agent's decision to rescue a culpable party was found overall more punishable than a decision to rescue an innocent party ($B = 1.76$, 95% CI = $[1.50, 2.03]$, $t(908) = 13.20$, $p < .001$, $d = 0.88$).

Harm covariate analysis. The Theory of Dyadic Morality (TDM; Schein & Gray, 2018) states that perceptions of harm are an important predictor of perceptions of moral wrongness. Schein and Gray (2019, p. 43) suggest a straightforward way of testing the alignment between harm and immorality: measuring both intuitive judgments of harm and intuitive judgments of immorality, and see how well the former predict the latter. Our DVs included a single-item Likert question about harm, specifically, how much harm the agent in the story caused with their rescue decision. We entered this variable as a covariate to the model used in the main analyses of our wrongness and acceptability analyses: see Supplementary Table 2 for full results. Our participants' judgments of harm caused by the agent's decision were statistically significant predictors of both acceptability and wrongness judgments, which makes sense given the similar mean patterns between harm and the other variables (see Figure 5 in the main manuscript and Supplementary Figure 3). The inclusion of harm in the models did not consistently drop the other effects from significance to non-significance. For acceptability, the three-way interaction fell

from significant to non-significant; for wrongness, the main effects of the agent and decision conditions fell from significant to non-significant, but their interaction remained significant. Most importantly, even after including harm, the interaction between the rescue decision and culpability conditions remained significant and had the highest effect size.

Supplementary Figure 3. Estimated marginal means for the secondary DVs in Study 4a. Error bars are 95% CIs.

Supplementary Table 1. Full statistics for three-way-ANOVAs of secondary DVs in Study 4a.

		Permissible	Blame	Punish	Praise	Harm
	DFs	(1, 937)	(1, 937)	(1, 937)	(1, 937)	(1, 937)
Agent	<i>F</i>	9.61	4.36	32.18	143.11	4.96
	<i>p</i>	.001	.036	< .001	< .001	.026
	η_p^2	.005	.003	.020	.096	.004
Culpable	<i>F</i>	2.53	10.19	6.87	0.23	1.49
	<i>p</i>	.111	.001	.008	.626	.222
	η_p^2	.001	.006	.004	.000	.001
Decision	<i>F</i>	46.19	30.79	19.88	27.78	61.23
	<i>p</i>	< .001	< .001	< .001	< .001	< .001
	η_p^2	.017	.019	.012	.020	.044
Agent x Culpable	<i>F</i>	2.28	0.79	3.21	1.03	0.08
	<i>p</i>	.131	.371	.073	.309	.765
	η_p^2	.001	.001	.002	.001	.000
Agent x Decision	<i>F</i>	3.40	6.31	5.38	8.15	5.22
	<i>p</i>	.065	.012	.020	.004	.022
	η_p^2	.002	.004	.003	.006	.004
Culpable x Decision	<i>F</i>	791.73	496.008	228.42	672.56	410.72
	<i>p</i>	< .001	< .001	< .001	< .001	< .001
	η_p^2	.305	.241	.125	.334	.238
Agent x Culpable x Decision	<i>F</i>	5.01	4.27	17.74	3.56	10.27
	<i>p</i>	.025	.038	< .001	.059	.001
	η_p^2	.003	.003	.011	.003	.008

Note: Agent: human vs. android; Decision: rescue fisherman vs. rescue motorboaters; Culpability: fisherman culpable vs. motorboaters culpable. The Decision factor is within-subjects, other factors between-subjects.

Supplementary Table 2. Full statistics for the harm covariate analysis in Study 4a.

DV	Factor	d f 1	df2	<i>F</i>	<i>p</i>	η^2_p
Acceptability	Agent	1	886.21	7.30	.007**	.002
	Decision	1	909.56	11.90	.001**	.004
	Culpability	1	884.80	0.89	.345	.000
	Harm	1	1560.53	237.79	< .001***	.076
	Agent × Decision	1	886.74	6.54	.011*	.002
	Agent × Culpability	1	884.22	3.93	.048*	.001
	Decision × Culpability	1	1047.08	1008.33	< .001***	.320
	Agent × Decision × Culpability	1	888.81	2.55	.110	.001
Wrongness	Agent	1	885.95	3.65	.056	.001
	Decision	1	911.43	1.63	.202	.001
	Culpability	1	884.39	0.01	.898	.000
	Harm	1	1690.31	285.68	< .001***	.108
	Agent × Decision	1	886.38	5.05	.025*	.002
	Agent × Culpability	1	883.76	2.32	.128	.001
	Decision × Culpability	1	1060.08	694.64	< .001***	.262
	Agent × Decision × Culpability	1	888.65	0.73	.392	.000

Note: Agent: human vs. android; Decision: rescue fisherman vs. rescue motorboaters; Culpability: fisherman culpable vs. motorboaters culpable; Harm: rating of harm caused by the agent's decision. The Decision factor is within-subjects, other factors between-subjects. Error degrees of freedom vary due to using Satterthwaite's method.

Study 4b

Analysis of secondary dependent variables. We ran the same analyses on the secondary DVs as in Study 4a. See Supplementary Table 3 and Supplementary Figure 4 for full statistics. Again, the patterns of means generally mimic those of acceptability and wrongness in the main analysis. The main effect of the culpability condition on blameworthiness and punishability replicated, but so did the significant interaction between culpability and rescue decision for these variables.

We also ran the same contrast analyses concerning blame, praise and punishment for Study 4b as we did for Study 4a. The results replicated fully. In Study 4b, the human agent's decision to rescue the innocent fisherman rather than the culpable motorboaters was judged as more blameworthy ($B = 0.32$, 95% CI = [0.06, 0.58], $t(807) = 2.44$, $p = .015$, $d = 0.17$) and less praiseworthy ($B = -0.75$, 95% CI = [-0.92, -0.38], $t(807) = -4.71$, $p < .001$, $d = -0.33$). For the android agent, decisions to rescue the culpable party were judged as more punishable than decisions to rescue the innocent party ($B = 0.61$, 95% CI = [0.48, 0.74], $t(807) = 9.16$, $p < .001$, $d = 0.64$).

Harm covariate analysis. We conducted the same analysis, using harm as a covariate, as for Study 5B. See Supplementary Table 4 for full statistics. Including harm as a covariate did not drop any of the main effects or interactions from significance to non-significance. Harm was a significant predictor for both wrongness and acceptability judgments, but again, the effect size of the interaction between rescue decision and culpability was the highest.

Supplementary Figure 4. Estimated marginal means for the secondary DVs in Study 4b. Error bars are 95% CIs.

Supplementary Table 3. Full statistics for three-way-ANOVAs of secondary DVs in Study 4b.

		Permissible	Blame	Punish	Praise	Harm
	DFs	(1, 1614)	(1, 807)	(1, 807)	(1, 807)	(1, 807)
Agent	<i>F</i>	27.03	2.42	14.10	84.96	9.28
	<i>p</i>	< .001	.119	< .001	< .001	.002
	η_p^2	.016	.002	.009	.072	.008
Culpable	<i>F</i>	1.06	11.51	9.46	0.48	0.00
	<i>p</i>	.301	< .001	.002	.485	.930
	η_p^2	.001	.010	.006	.000	.000
Decision	<i>F</i>	93.10	42.07	37.08	31.13	70.39
	<i>p</i>	< .001	< .001	< .001	< .001	< .001
	η_p^2	.055	.034	.024	.028	.058
Agent x Culpable	<i>F</i>	0.10	0.00	3.61	0.00	0.09
	<i>p</i>	.740	.964	.057	.930	.752
	η_p^2	.000	.000	.002	.000	.000
Agent x Decision	<i>F</i>	2.98	0.11	5.90	1.17	5.24
	<i>p</i>	.084	.734	.015	.279	.022
	η_p^2	.002	.000	.004	.001	.005
Culpable x Decision	<i>F</i>	287.97	144.10	97.34	222.89	156.99
	<i>p</i>	< .001	< .001	< .001	< .001	< .001
	η_p^2	.151	.109	.062	.168	.121
Agent x Culpable x Decision	<i>F</i>	6.50	0.93	9.82	0.99	1.68
	<i>p</i>	.010	.333	.001	.317	.194
	η_p^2	.004	.001	.007	.001	.001

Note: Agent: human vs. android; Decision: rescue fisherman vs. rescue motorboaters; Culpability: fisherman culpable vs. motorboaters culpable. The Decision factor is within-subjects, other factors between-subjects. Error DFs vary due to using Satterthwaite's method.

Supplementary Table 4. Full statistics for the harm covariate analysis in Study 4b.

DV	Factor	d f 1	df2	<i>F</i>	<i>p</i>	η^2_p
Acceptability	Agent	1	755.90	16.51	<.001***	.010
	Decision	1	782.00	67.63	<.001***	.042
	Culpability	1	752.00	1.29	.256	.001
	Harm	1	1377.49	123.29	<.001***	.073
	Agent × Decision	1	754.60	3.19	.074	.002
	Agent × Culpability	1	752.00	0.06	.802	.000
	Decision × Culpability	1	817.80	420.12	<.001***	.212
	Agent × Decision × Culpability	1	753.09	8.43	.004**	.005
Wrongness	Agent	1	758.29	6.31	.012*	.004
	Decision	1	785.15	11.82	<.001***	.008
	Culpability	1	754.25	1.52	.217	.001
	Harm	1	1416.71	170.97	<.001***	.103
	Agent × Decision	1	756.89	0.30	.583	.000
	Agent × Culpability	1	754.29	3.21	.073	.002
	Decision × Culpability	1	822.09	425.5478	<.001***	.222
	Agent × Decision × Culpability	1	755.33	16.67	<.001***	.011

Note: Agent: human vs. android; Decision: rescue fisherman vs. rescue motorboaters; Culpability: fisherman culpable vs. motorboaters culpable; Harm: rating of harm caused by the agent's decision. The Decision factor is within-subjects, other factors between-subjects. Error degrees of freedom vary due to using Satterthwaite's method.

Study 5

Analysis of secondary dependent variables. We ran the same analyses on the secondary DVs as in Studies 4a-b. See Supplementary Table 5 and Supplementary Figure 5 for full statistics. Again, the patterns of means generally mimic those of acceptability and wrongness in the main analysis. A notable deviation from the previous secondary DV analyses was that there was no significant main effect of the agent condition on the harm ratings, and no significant interaction involving the agent condition for harm. Previously, participants had judged at least some of the android agent's decisions as more harmful than a similar decision by the human (see Supplementary Figures 3 and 4), but this was no longer the case for Study 5. It may be that the simplicity of the design, with no conflict between different moral norms in any version of the vignette, and with utilitarianism being the only norm, attenuated differences in perceptions of harm. This suggests that differences between agents in how harmful their actions were seen mostly stemmed from violations of folk justice, which were absent in Study 5.

Comparison between studies. We also hypothesized that comparing the results of Study 5 with those of Study 4b, we would observe a difference in how the agents are judged. Specifically, we predicted that the difference between judgments of human and robot agents would be larger in Study 4b than in Study 5, because the vignettes in Study 5 had no contrast between one culpable and one innocent party, unlike the vignettes in Study 4b.

We tested this hypothesis for our two main DVs of interest, acceptability and wrongness, with a combined data set of Studies 4b and 5. The DVs were entered into a 2x2 between-subjects analysis with agent (human, android) and study (Study 4b, Study 5) as factors. There was no significant interaction effect between agent and study for either acceptability ($F(1, 3708) = 0.23$, $p = .629$, $\eta^2_p = 0.000$) or wrongness ($F(1, 3708) = 0.83$, $p = .361$, $\eta^2_p = 0.000$). There was a main

effect of study for acceptability ($F(1, 3708) = 37.03, p < .001, \eta^2_p = 0.010$) and wrongness ($F(1, 3708) = 36.84, p < .001, \eta^2_p = 0.010$). Judgments were overall more positive in Study 5 (acceptability: $M = 5.31$, 95% CI = [5.24, 5.39]; wrongness: $M = 2.59$, 95% CI = [2.52, 2.67]) than they were in Study 4b (acceptability: $M = 4.98$, 95% CI = [4.91, 5.06]; wrongness: $M = 2.93$, 95% CI = [2.86, 3.01]).

Harm covariate analysis. We conducted the same analysis, using harm as a covariate, as for Studies 4a and 4b. See Supplementary Table 6 for full statistics. Including harm as a covariate did not drop any of the main effects or interactions from significance to non-significance. Harm was a significant predictor for both wrongness and acceptability judgments, and in fact had the highest effect size for both acceptability and wrongness, albeit the effect size was very close to that of the rescue decision.

Supplementary Figure 5. Estimated marginal means for the secondary DVs in Study 5. Error bars are 95% CIs.

Supplementary Table 5. Full statistics for three-way-ANOVAs of secondary DVs in Study 5.

		Permissible	Blame	Punish	Praise	Harm
	DFs	(1, 911)	(1, 911)	(1, 911)	(1, 911)	(1, 911)
Agent	<i>F</i>	10.80	13.55	11.75	82.45	0.13
	<i>p</i>	.001	< .001	< .001	< .001	.711
	η_p^2	.007	.010	.008	.068	.000
Culpable	<i>F</i>	0.17	0.06	3.37	0.57	0.74
	<i>p</i>	.675	.796	.066	.449	.387
	η_p^2	.000	.000	.002	.001	.001
Decision	<i>F</i>	274.34	163.38	87.74	215.15	135.53
	<i>p</i>	< .001	< .001	< .001	< .001	< .001
	η_p^2	.159	.111	.059	.159	.105
Agent x Culpable	<i>F</i>	0.51	0.47	0.12	1.69	.02
	<i>p</i>	.474	.491	.720	.193	.887
	η_p^2	.000	.000	.000	.001	.000
Agent x Decision	<i>F</i>	11.23	7.64	14.33	6.56	1.50
	<i>p</i>	< .001	.005	< .001	.010	.220
	η_p^2	.008	.006	.010	.006	.001
Culpable x Decision	<i>F</i>	5.00	4.46	2.56	1.91	.87
	<i>p</i>	.025	.034	.109	.166	.349
	η_p^2	.003	.003	.002	.002	.001
Agent x Culpable x Decision	<i>F</i>	1.45	4.27	0.28	0.68	0.18
	<i>p</i>	.228	.038	.595	.406	.665
	η_p^2	.001	.003	.000	.001	.000

Note: Agent: human vs. android; Decision: rescue fisherman vs. rescue motorboaters; Culpability: both parties culpable vs. both parties innocent. The Decision factor is within-subjects, other factors between-subjects.

Supplementary Table 6. Full statistics for the harm covariate analysis in Study 5.

DV	Factor	d f 1	df2	<i>F</i>	<i>p</i>	η^2_p
Acceptability	Agent	1	881.60	18.30	< .001***	.012
	Decision	1	936.75	248.89	< .001***	.140
	Culpability	1	881.85	0.37	.539	.000
	Harm	1	1525.63	258.06	< .001***	.145
	Agent × Decision	1	882.58	13.59	< .001***	.009
	Agent × Culpability	1	881.55	1.54	.214	.001
	Decision × Culpability	1	882.33	5.40	.020*	.004
	Agent × Decision × Culpability	1	882.04	0.40	.526	.000
Wrongness	Agent	1	876.24	16.99	< .001***	.012
	Decision	1	937.91	194.20	< .001***	.126
	Culpability	1	876.52	0.35	.551	.000
	Harm	1	1684.67	242.95	< .001***	.153
	Agent × Decision	1	877.15	15.38	< .001***	.011
	Agent × Culpability	1	876.19	2.41	.120	.002
	Decision × Culpability	1	876.86	5.24	.022*	.004
	Agent × Decision × Culpability	1	876.54	0.36	.545	.000

Note: Agent: human vs. android; Decision: rescue fisherman vs. rescue motorboaters; Culpability: both parties culpable vs. both parties innocent; Harm: rating of harm caused by the agent's decision. The Decision factor is within-subjects, other factors between-subjects. Error degrees of freedom vary due to using Satterthwaite's method.

Appendix C: Dependent Variables

Studies 1 - 2c

Full main dependent variable scale. In Study 1, each item listed below was responded to on a 7-point Likert scale, where 1 denoted “completely disagree” and 7 denoted “completely agree”. Only the five items marked with asterisks were used to form the dependent variable, based on face validity as measures of moral approval of a decision, and exploratory factor analysis. Note: Studies 1-2b were not conducted in English. The items listed below are a translation that were used in Study 2c, which was conducted in English. In Studies 2a-2c, the first item was reworded as a question: “How morally acceptable do you find the [lifeguard / lifeguard robot]’s decision?”. The item was still responded to on a 7-point Likert scale, with 1 denoting “not moral at all” and 7 denoting “very moral”.

1. “The decision is morally acceptable.” *
2. “The decision is morally wrong.” *
3. “The [lifeguard / robot] can be blamed for the consequences of their decision.”
4. “The [lifeguard / robot] is responsible for the consequences of their decision.”
5. “The [lifeguard / robot]’s decision is irresponsible.” *
6. “The decision should not be considered morally right.” *
7. “The [lifeguard / robot] should be convicted in court.”
8. “The lifeguard is mean.”
9. “The decision is unethical.” *
10. “Considering the conditions and the situation, the decision is correct.”
11. “The decision is sensible.”
12. “The decision makes the [lifeguard / robot] more valuable for their community.”

13. “The decision is unfair.”
14. “The [lifeguard / robot] should not be the one deciding on who should be saved first.”
15. “The decision is different from what I expected.”

Studies 4a - 5

Single-item dependent variables. The following items were presented twice for each participant, once for each hypothetical decision. Responses to each question were given both on a Likert scale from 1 to 7 and on a slider from 1 to 100. The participants first read the following instructions:

“For the following questions, imagine that the [agent] decided to save [the fisherman / the motorboaters] and succeeded in saving [him/them] without issues. This means that [the fisherman / the motorboaters] survived, and [the motorboaters were / the fisherman was] left in the icy water without enough time to save [them/him].”

DV1. “How morally wrong would you find the [agent’s] decision to save [the fisherman / the motorboaters]? (Please respond using both scales below)”

Likert, 1 = not morally wrong at all , 7 = very much morally wrong

Slider, 1 = not morally wrong at all, 100 = very much morally wrong

DV2. “How acceptable would you find the [agent’s] decision to save [the fisherman / the motorboaters]? (Please respond using both scales below)”

Likert, 1 = not acceptable at all, 7 = very acceptable

Slider, 1 = not acceptable at all, 100 = very acceptable

DV3. “How permissible would you find the [agent’s] decision to save [the fisherman / the motorboaters]? (Please respond using both scales below)”

Likert, 1 = not permissible at all, 7 = very permissible

Slider, 1 = not permissible at all, 100 = very permissible

DV4. “How much blame do you think the [agent] would deserve for [their/its] decision if [they/it] decided to save [the fisherman / the motorboaters]? (Please respond using both scales below)”

Likert, 1 = none, 7 = maximal blame

Slider, 1 = none, 100 = maximal blame

DV5. “Do you think the [agent] should be punished for the decision to save [the fisherman / the motorboaters]? (Please respond using both scales below)”

Likert, 1 = not at all, 7 = very much

Slider, 1 = not at all, 100 = very much

DV6. “Do you think the [agent] should be praised for the decision to save [the fisherman / the motorboaters]? (Please respond using both scales below)”

Likert, 1 = not at all, 7 = very much

Slider, 1 = not at all, 100 = very much

D7. “How much harm do you think it would cause if the [agent] decided to save [the fisherman / the motorboaters]? (Please respond using both scales below)”

Likert, 1 = none, 7 = very much

Slider, 1 = none, 100 = very much

Appendix D: Additional Measures Not Reported in Main Manuscript

Studies 4a-5

Between-subjects measures.

DV8. “Who do you think the [agent] should save?”

Dichotomous, “Fisherman” / “Motorboaters”

DV9. “If you were in the [agent’s] situation, which decision would you make? Imagine here that you would be just as physically capable to carry out the rescue operation as the [agent], and it would be equally safe for you to rescue either the fisherman or the two motorboaters.”

Dichotomous, “I would save the fisherman”, “I would save the motorboaters”

DV10. “If you were in the [agent’s] situation, how difficult would the decision about whom to save be for you? Imagine here that you would be just as physically capable to carry out the rescue operation as the [agent], and it would be equally safe for you to rescue either the fisherman or the two motorboaters. (Please respond using both scales below)”

Likert, 1 = not at all difficult, 7 = extremely difficult

Slider, 1 = not at all difficult, 100 = extremely difficult

DV11. “How many people out of a 100 do you think would save...”

Two sliders (0 - 100), “The fisherman”, “The motorboaters”

(Responses must sum to a 100)

DV12. “Imagine that instead of the rescue operation going as planned, as [they/it] run[s] towards the water, the [agent] experiences [an aorta rupture / a short-circuit] and collapses in the water.

The motorboaters and the fisherman are left without aid, and die in the freezing water. To what degree would this make you feel...”

Frustrated

Angry

Sad

Disappointed

Outraged

Likert for each emotion described, 1 = not at all, 7 = extremely

Slider for each emotion described, 1 = not at all, 100 = extremely

DV13. “How appropriate do you think it is for [a human / a robot with artificial intelligence] to be the one to make a decision about whom to save in the scenario described?”

Likert, 1 = not at all appropriate, 7 = completely appropriate

Slider, 1 = not at all appropriate, 7 = completely appropriate

DV14. “In the following, we ask you to divide “responsibility points” between different parties.

In your opinion, who or what do you think is ultimately responsible for the end result of the story you read and to what degree, whatever the [agent] decides to do?

The points you allocate to different parties must sum up to a 100.”

Six sliders (0-100), “The fisherman” / “The motorboaters” / “[The agent]” / “The state” / “Some other party” / “No-one”]

(Responses must sum to a 100)

Appendix E: Additional Meta-Analyses

We conducted the same analyses as reported in the meta-analysis section, this time using random-effects meta-analyses and a mixed-effects meta-analysis with moderators. Our initial choice to use fixed-effects analyses was based on the small number of studies available. We report the alternative analyses here for completeness’ sake. The results of the random-effects meta-analyses are presented in Supplementary Figure 6, and the results of the mixed-effects moderator analysis in Supplementary Table 7.

We again found an overall meta-analytic effect of agent ($g = -0.17$, $z = -6.87$, $p < .001$). We found no difference in judgments between agents when the decision was to rescue the innocent ($g = -0.06$, $z = -1.48$, $p = .140$), but the robot agent’s decision to rescue the culpable was judged more negatively than the same decision by the human agent ($g = -0.30$, $z = -5.71$, $p < .001$). Regarding the number of rescued people, the robot agent’s decision was judged more negatively than the human agent’s regardless of the decision (rescuing one: $g = -0.24$, $z = -5.55$, $p < .001$; rescuing two: $g = -0.23$, $z = -2.39$, $p = .016$). However, when confounded studies were

excluded, there was no longer a significant difference between the agents when rescuing two people.

In the moderator analysis, the effect of the folk justice moderator was again larger than that of the utilitarianism moderator. However, we observed no significant effect of the agent moderator, and no interaction between agent and the folk justice moderator. This analysis was further decomposed by separately examining the effects of the moderators separately for the agent groups. These effects were similar for both agents, although differed somewhat in magnitude (human agents: utilitarianism; $g = 1.05$, $z = 3.62$, $p < .001$, folk justice; $g = -0.94$, $z = -3.28$, $p = .001$, robot agents: utilitarianism; $g = 0.93$, $z = 3.35$, $p < .001$, folk justice; $g = -1.54$, $z = -5.59$, $p < .001$).

Supplementary Figure 6. Individual and meta-analytic effect sizes. Highlighted area indicates studies with a confound between the number and culpability of the rescued party. Squares and their 95% CIs represent standardized mean differences (Hedges' g) from individual studies. Diamonds and their 95% CIs represent effect sizes from random-effects meta-analyses including all available studies. Circles and their 95% CIs represent effect sizes from random-effects meta-analyses excluding Studies 1-2b. The size of the squares, diamonds and triangles indicates group size.

Supplementary Table 7. Results of mixed-effects moderator analysis

	<i>g</i>	SE	<i>Z</i>	<i>p</i>	CI lower	CI upper
Intercept	0.00	0.08	0.10	.914	-0.14	0.16
Utilitarianism	0.99	0.20	4.93	< .001 ***	0.60	1.39
Folk justice	-1.24	0.19	-6.24	< .001 ***	-1.63	-0.85
Agent	-0.14	0.16	-0.90	.364	-0.45	0.16
Utilitarianism × Agent	-0.12	0.40	-0.30	.763	-0.91	0.67
Folk justice × Agent	-0.60	0.39	-1.51	.129	-1.38	0.17

Note: Utilitarianism codes for whether dependent variable is an effect size from a contrast comparing rescuing two to rescuing one or not; Folk justice codes for whether it is an effect size from a contrast comparing rescuing the culpable to rescuing the innocent or not; Agent codes for whether effect size is from a contrast comparing a human's or a robot's decisions. The dependent variable is an effect size measure that has been standardized. Predictors are binary and have been centered. Coefficient estimates should be interpreted as deviations from the mean of the overall average effect size of contrasts that were included.

*** = significant at $p < .001$.

Appendix F: Sensitivity Power Analysis For Interactions

We conducted simulation-based sensitivity power analyses for interaction effects in each study except for Study 3 (due to no interactions having been tested in Study 3). As statistical power in the case of interactions is dependent on the mean pattern, we simulated mean patterns based on the observed mean patterns and manipulated both the sample size of the simulation and the magnitude of differences in the mean patterns. Specifically, we simulated versions of each data set with lower and higher sample sizes than the ones in the actual data, and versions of each data set with each mean scaled by a percentage (keeping relative differences between the means similar, but reducing or increasing absolute differences, while keeping the variance unchanged).

For each simulation, a dependent variable for each group/condition in a study was sampled based on the respective observed group/condition mean and standard deviation. Separate simulated data sets were created by varying the sample size of the simulated data (number of participants per cell was 50%, 75%, 100%, 125% or 150% of the actual observed number of participants per cell) and by adjusting the means (each observed mean in the data scaled to 80%, 100%, or 120%). Means were sampled using the `rnorm` function, with the observed or adjusted mean and the observed standard deviation as inputs (standard deviation was never adjusted). Sampling was iterated 1000 times for each combination of different sample sizes and mean adjustments. The rationale for the mean adjustments without adjusting standard deviations was to create simulated data with interaction effects that were qualitatively similar, with differences between experimental conditions in the same directions, but that had different effect sizes.

For each simulated data set (1000 iterations of each combination of sample sizes and mean adjustments), we conducted ANOVAs that tested for all two-way and three-way

interactions. We then created a dummy variable to code for whether the effect was significant under the $p < .05$ criterion or not, with 1 indicating a statistically significant result and 0 indicating a non-significant result. Power for a given combination of sample sizes and mean adjustments was calculated by taking the average of the dummy variable from each iteration with that sample size and mean adjustment. Simulations were run separately for the acceptability and wrongness measures in Studies 4a-5. The results of these simulations are presented in Supplementary Figures 7 through 16.

For some interaction effects that were non-significant in the observed data, the analyses suggest that achieving 80% power for detecting the effect would require an unfeasibly large sample size, even when the mean pattern was adjusted so that the effect size should be larger. Note that this does not necessarily mean that the study itself was underpowered: it is also possible that the pattern of means that provided the basis for our simulations are such that an interaction would not be observed even if the mean patterns were adjusted even further. That is, the low power for some of the non-significant interaction effects in our studies may stem from the fact that the observed data simply does not seem to indicate an interaction effect. The results of these analyses should not be interpreted as indicators of sample size requirements for any given two-way or three-way interaction. The analyses can only tell about sample size requirements for interactions with specific mean patterns: a limitation of this approach is that there are mean patterns that do not result in significant interaction effects regardless of the magnitude of differences between individual means.

Supplementary Figure 7. Results of simulated sensitivity power analysis for the two-way interaction in Study 1.

Supplementary Figure 8. Results of simulated sensitivity power analysis for the two-way interaction in Study 2a.

Supplementary Figure 9. Results of simulated sensitivity power analysis for the two-way interaction in Study 2b.

Supplementary Figure 10. Results of simulated sensitivity power analysis for the two-way and three-way interactions in Study 2c.

Supplementary Figure 11. Results of simulated sensitivity power analysis for the two-way and three-way interactions in Study 4a when the dependent variable was acceptability.

Supplementary Figure 12. Results of simulated sensitivity power analysis for the two-way and three-way interactions in Study 4a when the dependent variable was wrongness.

Supplementary Figure 13. Results of simulated sensitivity power analysis for the two-way and three-way interactions in Study 4b when the dependent variable was acceptability.

Supplementary Figure 14. Results of simulated sensitivity power analysis for the two-way and three-way interactions in Study 4b when the dependent variable was wrongness.

Supplementary Figure 15. Results of simulated sensitivity power analysis for the two-way and three-way interactions in Study 5 when the dependent variable was acceptability.

Supplementary Figure 16. Results of simulated sensitivity power analysis for the two-way and three-way interactions in Study 5 when the dependent variable was wrongness.