Good Advice is Beyond All Price, but What if it Comes from a Machine?

Nicholas Hertz & Eva Wiese

George Mason University

ABSTRACT

As nonhuman agents become integrated into the workforce, the question becomes whether humans are willing to consider their advice, and to what extent advice-seeking depends on the perceived agent-task fit. To examine this, participants performed social and analytical tasks and received advice from human, robot, and computer agents in two conditions: in the *Agent First* condition, participants were first asked to choose advisors and were then informed which task to perform; in the *Task First* condition, they were first informed about the task and then asked to choose advisors. In the *Agent First* condition, we expected participants to prefer human to non-human advisors, and to subsequently trust their advice more if they were assigned the social as opposed to the analytical task. In the *Task First* condition, we expected advisor choices to be guided by stereotypical assumptions regarding the agents' expertise for the tasks, accompanied by higher trust in their suggestions. The findings indicate that in the *Agent First* condition, the human was chosen significantly more often than the machines, while in the *Task First* condition advisor choices were calibrated based on perceived agent-task fit. Trust was higher in the social task, but only showed variations with the human partner.

Keywords: Human-Machine Interaction, Trust, Decision-Making, Advice, Compliance

INTRODUCTION

Humans are inherently social creatures, and the level of social learning in human groups is unparalleled in the animal kingdom (Biele, Rieskamp, Krugel, & Heekeren, 2011). To achieve this sophisticated level of social learning, humans rely on a number of learning mechanisms that are based on observation and advice taking from others (Biele et al., 2011). Advice taking has been shown to be especially influential in situations where reliable social information based on others' expertise is readily available and erroneous decisions would be associated with a high cost (Henrich & McElreath, 2003). For example, when searching for something to eat in environments that contain a lot of poisonous food, individuals who are part of a functioning social group are able to potentially spare themselves the cost of false decisions by asking for advice from others (Henrich & McElreath, 2003). These days, we are not confined to advice from a small number of trusted friends who live in the same environment as we do but have access to the opinions of a multitude of different people from all over the world with different skill sets via Internet platforms like *yelp* or *wikihow.com*. Not only are these advice givers strangers to us, but also are they sometimes not even human, as in the case of virtual avatars or *chat bots*. The advice we receive from nonhuman advisors ranges from navigation instructions provided by a GPS to recommendations when placing an order for Mexican food with the help of *TacoBot* to potentially life-saving advice for TSA employees when searching personal luggage for weapons or explosives. Given this increased exposure to machine operators in modern societies, it is essential to understand to what extent human operators trust the advice of non-human agents and whether there are situations where machine advisors might even be preferred to human advisors.

In human-human interactions, a number of factors have been found to influence the receptivity of advice, including *personal trust* (i.e., the extent to which a user is confident in and

willing to act on the basis of the recommendations, actions, and decisions of an advisor; Mayer, Davis, Schoorman, 1995; Madhavan & Wiegmann, 2007), *source expertise* (i.e., extent to which advisor's capabilities are in line with task demands; Van Swol & Snizek, 2005; Madhavan & Wiegmann, 2007), and *source reliability* (i.e., extent to which advisor responds consistently and predictably; Bonaccio & Dalal, 2006). Furthermore, a model developed by Rempel et al. (1985) predicts the acceptance of advice from a human advisor to depend on its *predictability* (i.e., consistency and stability of the advisor's performance over time), *dependability* (i.e., confidence in the advisor as a person), and *faith* (i.e., expectation regarding the advisor's accuracy and behavior in the future). Van Swol and Snizek (2005) also showed that human operators are more willing to follow advice when it is given by an expert versus novice human advisor.

With regard to advice taking in human-machine interactions, Dzindolet and colleagues (2002) have argued that decision makers may respond similarly to advice from automation or technology as they do to human advice, based on the assumption that people believe that automated aids are essentially perfect (i.e., *perfect automation schema*). Neuroscientific studies also show that surprisingly similar brain areas involved in mentalizing are activated both in human-human and human-machine interaction (Goodyear, Parasuraman, Chernyak, Madhavan, Deshpande et al., 2016). Other researchers argue trust in advice from robotic or computer agents to be rather different from trust in human advice (Madhavan & Wiegmann, 2007). For example, Jien and colleagues (2000) found that automated agents were more likely to be labeled as "distrusted" than human agents, which indicates that automated agents may be trusted less by default. Further examples of trust differences between human-human and human-machine interaction include the observation that people tend to view expert computer decision support systems as more objective and rational than humans (Dijkstra, Liebrand, & Timminga, 1998), react more negatively to er-

rors made by computer agents (Dzindolet, Pierce, Beck, Dawe, & Anderson, 2001), and show a more immediate decline in trust after errors made by machines (Wiegmann, Rich, & Zhang, 2001).

Differences between human and machine advisors are also seen in the degree to which their advice is trusted. It was found that trust in a machine advisor is modulated by general factors like its perceived expertise (Muir & Morray, 1996; Muir & Morray, 1994; Liu, 2010), but also factors that relate to its similarity to humans such as its human-like appearance (Smith, Allaham, & Wiese, 2016; Madhavan & Wiegmann, 2007) or sensitivity to the advisee's emotional states (Picard, 1997; Goetz et al., 2003). In particular, it was shown that agents with a human-like physique receive higher trust ratings than mechanistic looking agents (Smith et al., 2016; Madhavan & Wiegmann, 2007), and that participants comply more with instructions given by human-like agents as opposed to machine-like agents (Goetz et al., 2003). Interestingly, however, when different agent features, such as the capacity to feel sensations such as hunger (i.e., experience) and the capacity to plan (i.e., agency), communicated inconsistent levels of humanness (e.g., agent seemed able to plan, but unable to feel), participants had negative attitudes towards the advisor and showed low levels of trust (although at least one of the features was high in humanness; Gray & Wegner, 2012). Similarly, Goetz and colleagues (2003) showed that trust was impaired when agent appearance and perceived abilities did not match the task, suggesting that the expertise a nonhuman advisor was expected to have, based on displayed features, determined how strongly its advice was trusted. Trust in a nonhuman agent's ability to support a given task was also shown to depend on the extent to which it was deemed appropriate to handle the situation (Liu, 2010), and to support the user to achieve his goals (Lee & Morray, 1992). In addition, trust in non-human agents is higher when they are perceived as specialized expert systems as op-

posed to general-purpose novice systems (Shaefer et al., 2014; Koh & Sundar, 2010; Madhavan & Wiegmann, 2007).

While differences in advice *taking* from human versus non-human advisors and the effects thereof on trust are well researched, general preferences in advice *seeking* (i.e., which agent to choose for advice) and potential effects on trust (i.e., to what extent is the advice followed) are less well examined. Madhavan & Wiegmann (2007) propose that human and nonhuman advisors are associated with different strength and weaknesses, which might in turn modulate agent preferences during advice seeking. According to their model, human advisors are perceived as being adaptable and are thought to have the ability to change their behavior more easily to suit the demands of the task, while nonhuman agents are characterized by performance invariability across situations. Human advisors are also judged more based on knowledge-linked information (i.e., effort and expertise), while machine advisors are mainly judged based on performance-linked information (i.e., accuracy and reliability). This indicates that information about the task-to-be-performed might critically determine whether individuals seek advice from human or nonhuman advisors, with human advisors being preferred if the task is unknown since they are seen as more adaptable to changing circumstances and a wider range of tasks. Human advisors should also be preferred when nonhuman agents are considered as being incompetent for tasks that require core human skills, such as judging emotions or making friends. Nonhuman agents, on the other hand, should be preferred as advisors on tasks that humans are considered as having difficulties with, since they are highly technical, such as mathematical tasks or tasks that require a high level of precision and consistency, such as manufacturing. To the best of our knowledge this hypothesis has not been empirically investigated yet.

Aim of Study

The goal of the current experiment is to explore whether humans distrust machine advisors in general or whether they sometimes even prefer machine advisors when they seem more competent for the task at hand. We would expect individuals to always prefer human to machine advisors if the task is unknown, but to calibrate their advice seeking based on the advisors' perceived expertise if the task is known. Furthermore, we would expect higher levels of trust in the advice given, if the advisor was selected based on its perceived fit for the task and not just based on a general preference for human-like agents. To test these hypotheses, we ask participants to perform a social (i.e., identifying emotional states) and an analytical task (i.e., standard arithmetic calculations), and to consider advice from a human, a robot, or a computer agent, which they can subsequently accept or reject depending on whether they trust it. To be able to distinguish between general preferences for certain advisors and the extent to which their perceived expertise for a given task plays a role in advice seeking, we implement two experimental conditions: In the *Agent First* condition, participants are asked to first choose the agent they want to get advice from and then get either the social or analytical task assigned (on each trial). Which agent they choose in this condition will inform us about their preferences when calibration according to the agents' expertise is not possible (i.e., human agent for social task, machine agents for analytical task), and the rate at which they follow the given advice will inform us about their trust in the chosen agent's advice. We hypothesize that participants prefer the human agent to the non-human agents since calibration according to expertise is not possible, and will trust human advice more than machine advice when they subsequently get the social versus the analytical task assigned (*Hypothesis 1*).

In the *Task First* condition, participants first receive their task assignment (on a trial-by-trial basis) and then get to choose their preferred advisor, which allows them to calibrate advice seeking according to the agents' perceived expertise (i.e., human agent for social task, machine agents for analytical task). We expect participants to choose the human agent significantly more often for the social task assignment than for the analytical task assignment (with inverted pattern for the nonhuman agents), accompanied by high post-decision trust since agent expertise should have already been figured into the advisor choice (*Hypothesis 2*).

METHOD & MATERIALS

Participants

65 undergraduate students (53 women; $M_{\text{age}} = 20.47$ years, age range: 18-58 years) at George Mason University participated in the experiment. They were recruited via SONA (an online study pool recruitment platform), and compensated for participation through course credit. The decision time point for agent selection was manipulated as a between-participants factor: in condition 1 (Agent First), participants had to choose the agent first before they were told which task to perform, while in condition 2 (Task First), the task type was first stated and then an agent had to be chosen. Participants were randomly assigned to either the agent-first condition (32 participants, 29 women, $M_{\text{age}} = 21.16$) or the task-first condition (33 participants, 24 women, $M_{\text{age}} = 19.79$) at the beginning of the experiment.

Task

Two tasks were used during the experiment: a social task and an analytical task. As social task, we used the *Reading the Mind in Films* task, which assesses recognition of complex emo-

tion and mental states from complex film scenes (Golan, Baron-Cohen, Hill, & Golan, 2006).

The task measures similar social competencies as the *Reading the Mind in the Eyes* task (Baron-Cohen et al., 2001) but is thought to be ecologically more valid since it is arguably more challenging and closer to real social situations (Golan et al., 2006). The task consists of 22 film clips (5-30 seconds long, $M=14.8$, $SD = 9.2$) that come from four British TV films (e.g. *Lost for Words*; Golan et al., 2006). The analytical task was adapted from Smith, Allaham & Wiese (2016) and required participants to solve standard addition and subtraction operations.

The social task and the analytical task were presented simultaneously: The film stimulus was shown in the center of the screen, automatically playing at the start of a new trial, while a series of white dots appeared above and below the stimulus. To complete the social task, participants had to determine what emotion or mental state a specific character in the film clip was exhibiting (e.g. “At the end of the scene, how is the woman feeling?”). To complete the analytical task, participants had to complete an addition or subtraction problem using the dots surrounding the film clip (e.g. “Top Dots + Bottom Dots =”). In both tasks the participants would respond by clicking on of four multiple-choice options; see **Figure 1**.

Agents

Participants had to perform the tasks together with other agents: a human, a robot or a computer (adapted from de Visser et al., 2012; see **Figure 2**). Participants were introduced to all agents at the beginning of the experiment via images and short cover stories: the computer was introduced as a standard desktop that was equipped with algorithms designed to learn and answer complex cognitive tasks. The robot was described as a humanoid robot (“Meka”) that has similar cognitive abilities as the computer but possesses a human-like appearance. The human was introduced as Matt O’Shea, an English major at an American university. On each trial, participants

were asked to pick the one agent they wanted to work with on this particular trial: in the Agent-First condition, participants first made their agent choice and were then presented with either the social or the analytical task; in the Task-First condition, participants were first assigned either the social or the analytical task and then had to make their agent choice.

Procedure

The stimuli were presented on a computer screen using the software Inquisit 5 (Millisecond Software). At the beginning of the experiment, participants were randomly assigned to the Agent-First or the Task-First condition. Afterwards, participants gave informed consent and completed a short demographic questionnaire. Participants were then asked if they had seen any of the films that the *Reading the Mind in Films* task utilizes by presenting them with thumbnails and titles of the movies. The participants were then introduced to the three agents (human, robot, computer) via pictures and short descriptions. If participants did not have questions, the test phase began.

Participants were presented with 22 trials during which the *Reading the Mind in Film* task and the arithmetic task were presented simultaneously (i.e., both the dots and the film clip were visible at the same time). In the Task-First condition participants would be presented with either the social (“At the end of the scene, how is the woman feeling?”, for example) or the analytical question (“Top Dots + Bottom Dots =”, for example) first, and then asked to answer the question. After they had given their answer, they were asked to select the one agent that they would like to ask for advice on the given task and the response of the respective agent for the given question was displayed. Following the agent answer, participants were given the opportunity to change their answer or to stick to their original choice. In the Agent-First condition, the agent

choice had to be made after the stimuli were presented, but before the task was assigned to them. Otherwise, the task proceeds identically to the Task-First condition (see **Figure 1**). In both conditions, the analytical and the social task appeared with equal frequency (11 social, 11 analytical) and the order in which they were presented was randomized across participants; the agent's responses to both the social and analytical task had an accuracy rate of approximately 64% (7 of the 11 were answered correctly).

RESULTS

Agent Choice Data

A 2 x 3 x 2 mixed ANOVA was conducted on frequencies of agent choices with *Condition* (Agent-First vs. Task-First) as between-participants factor and *Agent* (Human vs. Robot vs. Computer) and *Task* (Social vs. Analytical) as within-participants factors. Whenever Mauchly's test indicated that the assumption of sphericity was violated, degrees of freedom for the respective factor were Greenhouse-Geisser corrected. The ANOVA showed a main effect for *Agent* ($F(0.771, 49.344) = 17.93, p < .001, \eta^2 = .014$), indicating a significant difference in selection across agent types, as well as significant interactions of *Agent* x *Task* ($F(1.542, 97.146) = 13.34, p < .001, \eta^2 = .008$), indicating that selection of agents depended on task type, and *Condition* x *Agent* x *Task* ($F(1.542, 97.146) = 14.01, p < .001, \eta^2 = .008$), indicating that participant's selections differed depending on task type and condition (see **Figure 3**). All other effects were not significant (all $F_s < .001, p_s > .05$).

A follow-up 2 (*Task*: Social vs. Analytical) x 3 (*Agent*: Human vs. Robot vs. Computer) mixed ANOVA (one per *Condition*) was ran. The Agent-First condition showed a significant main effect of *Agent* ($F(1.526, 48.832) = 9.96, p < .001$), with the human agent being chosen significantly more often than the computer ($t(32) = -4.752, p < .001$) and the robot agent being

chosen significantly more often than the computer ($t(32) = -2.567, p < .05$). The main effect of *Task* ($F(1, 32) < .001, p > .05$) and the interaction of *Agent* x *Task* ($F(2, 64) = .265, p = .768$) were not significant, indicating no difference in selection between social and analytical tasks, as well as showing that selecting did not differ across agent and task type in the Agent-First condition. For the Task-First condition, the analysis also showed a significant main effect of *Agent* ($F(1.556, 48.236) = 8.79, p < .05$), and no significant main effect for *Task* ($F(1, 31) < .001, p > .05$). In contrast to the Agent-First condition, however, the interaction of *Agent* x *Task* was significant ($F(1.61, 49.91) = 18.77, p < .001$). Post-hoc paired t-tests for each of the agents showed that the computer ($t(61.938) = -4.444, p < .001$) and the robot ($t(58.132) = -3.358, p = .004$) were chosen significantly more often in the analytical than the social task, while the human was significantly more often chosen in the social than the analytical task ($t(51.914) = 5.809, p < .001$); t-tests were Bonferroni corrected for multiple comparisons.

Trust Data

Trust was measured in terms of compliance, which was calculated as the percentage of time participants changed their initial response to match the advice given by the chosen agent (Fox, 1996; Lee & Morray, 1992; Smith, Allaham, & Wiese, 2016). Due to the nature of the experiment, participants did not choose agents at equivalent rates. The human was selected significantly more often than the computer (43.9% vs. 23.3%; $t(129) = -5.75, p < .001$) and the robot (43.9% vs. 32.8%; $t(129) = 2.78, p < .01$), and the robot was chosen significantly more often than the computer (32.8% vs. 23.3%; $t(129) = -3.857, p < .001$). To accommodate this inequivalence in selection rates we used linear mixed-effects models with restricted maximum likelihood for the analysis of the compliance data. The model used the variables *Agent*, *Task*, and *Condition*

to predict compliance rates. This test revealed a significant main effect of *Task* ($b = .009$, $SE = .002$, $t(284.6) = 6.128$, $p < .001$), with higher compliance rates in the social versus the analytical task. The three-way interaction of *Human* $\times$ *Condition* $\times$ *Task* was also significant ($b = .002$, $SE = .002$, $t(282.8) = 2.011$, $p = .045$). Two follow up models separated by *Condition* were conducted to better understand the structure of the three-way interaction: the Task-First model indicated a main effect of *Task* ($b = .099$, $SE = .019$, $t(136.11) = 5.146$, $p < .001$), with significantly more compliance for the social task than the analytical task, and a *Task* $\times$ *Human* interaction ($b = -.069$, $SE = .027$, $t(134.67) = -2.185$, $p = .031$), indicating that participants complied more with the human agent on the social task and less on the analytical task compared to the other agents. The model for the Agent-First condition indicated only a significant main effect for *Task* ($b = .084$, $SE = .029$, $t(147.58) = 3.73$, $p < .001$), once again depicting greater compliance for the social task than the analytical task (see **Figure 4**). All other effects were not significant ($p < .05$).

DISCUSSION

The goal of the current study was to investigate whether humans prefer human advisors in general or whether their advice seeking is calibrated to the advisor's perceived expertise on a given task, and how this would affect subsequent trust behaviors. Our first hypothesis was that participants would prefer the human agent to the nonhuman agents if they were not aware of the task at the time of their decision (*Agent First*), and would subsequently trust the advice more on the social than the analytical task. This is based on the assumption that since calibration based on task expertise was not possible for the advisor choice, it would take effect on post-decision trust. Our second hypothesis was that this pattern would change if participants knew the task before choosing an advisor (*Task First*), since expertise-based calibration of their advisor choice would now be possible. In particular, we expected participants to choose the human agent significantly

more often for social task assignments than analytical task assignments, with an inverted pattern for the nonhuman agents, and that this would be accompanied by high levels of post-decision trust (since the presumably most competent advisor was already selected).

The advisor choice and post-decision trust data indicate that when the task is not known (*Agent First*), participants choose the human advisor significantly more often than the robot and computer advisor, but do not adjust post-decision trust according to the advisor's presumed task expertise (in partial support of Hypothesis 1). When the task is known before the advisor had been chosen (*Task First*), participants adjust advisor choices according to the agents' presumed task expertise and pick the human advisor significantly more often for social task assignments, and the robot and computer advisors significantly more often for analytical task assignments. Post-decision trust is also higher for the human advisor on the social than the analytical task, with no effect of agent type on post-decision trust for the machine agents (in partial support of Hypothesis 2). These results indicate that participants prefer human advisors to nonhuman advisors when the task is unknown and a more informed decision cannot be made. However, when the task is known, advisor choice is adjusted based on the presumably best agent-task fit (in line with our hypotheses). In contrast, post-decision trust was not strongly affected by agent-task fit, but mostly depended on task type only (i.e., higher compliance with given advice for the social, 33.93%, as opposed to the analytical task, 17.47%).

Our results are in line with Madhavan and Wiegmann (2007), who proposed a theoretical model that predicted that humans might be preferred as advisors if the task is unknown since they are seen as more adaptable to changing circumstances and different task types. The observation that advisors are picked based on their perceived task expertise is also in line with Dzindolet and colleagues (2002) who showed that participants have preconceived notions regarding the

utility of automated and human agents for different tasks. They found that participants expected automated aids to make fewer mistakes, but were less likely to disuse information from human aids (Dzindolet, Pierce, Beck, & Dawe, 2002). The idea that advisors are chosen based on their perceived abilities is also supported by findings from Momen and colleagues (2016) who showed that the usefulness of virtual agents on a feedback learning task depended on judgments of agent abilities as well as the relevance of the task for the participant.

While the effect of agent type on advice seeking is easy to interpret, the interpretation of the post-decision trust data is a little more puzzling. In both the *Agent First* and *Task First* conditions, trust rates were higher for the social and the analytical task, which seemed to be harder to perform in general based on participants' accuracy scores (Analytical Task Accuracy = 53.68%; Social Task Accuracy = 44.78%). The observation that compliance with given advice increases as task difficulty increases has been reported in previous studies (e.g., Gino & Moore, 2007), and is likely linked to an increased need for social information seeking under uncertainty (Snizek & Van Swol, 2001). In line with our predictions though, trust was not modulated further by advisor type in general in the *Task First* condition, presumably since considerations regarding agent-task fit had already been figured into participants' advisor choices. Interestingly, however, there were selective effects of human agent type on post-decision trust, such that the human advisor was not only picked significantly more often for the social than the analytical task, but also trusted more on the social than the analytical task. This might be best explained in light of findings by Ramberg (1996) who showed that participant's knowledge in a task domain together with their existing attitudes and experiences with the advisor has a strong impact on trust behaviors. Thus, the participants may have viewed themselves and other humans as experts in the social task and novices in the analytical task, which influenced their trust on both tasks (Shanteau, 1992). In contrast

to our predictions, we did not find a modulation of trust based on advisor type in the *Agent First* condition, showing that agent-task fit did not have a calibrating effect on post-decision trust. This is in line with Madhavan & Wiegmann (2007) who showed that although participants perceived differences in an agent's utility, similar to our agent choice results, subsequent use of and trust in their advice did not necessarily correlate with their perceived utility. Similarly, Ramberg (1996) suggests that under increased uncertainty, advice given by a computer may be trusted more than that of a human, which may help explain the trending decrease in trust as humanness increased. Additionally, having participants select advisors before actually receiving their recommendations could have produced *cognitive anchoring* (i.e. the tendency to overly rely on initial information) which has been shown to reduce trust in received advice during decision-making and might explain why advisor type did not modulate post-decision trust (Goodyear et al., 2016; Madhavan & Wiegmann, 2005). In line with this hypothesis, a previous study conducted by Smith et al. (2016) involving the same agents showed that when human, robot, and computer advisors were randomly assigned to either the social or analytical task, trust in the agents' advice showed the expected advisor by task type interaction, with stronger compliance with the human advisor on the social task and the machine agents on the analytical task. This indicates that allowing individuals to select their preferred advisor before receiving specific advice from them produces a different trust outcome than when agents are assigned without participant control.

There are some open questions and limitations associated with the current experiment. First, although it has been shown that people seek more advice on tasks that are considered to be important to them (Harvey & Fischer, 1997), task importance was not manipulated in the current experiment, which could potentially have an effect on decisions regarding agent choice and post-decision trust. Future research should manipulate task importance with outcomes contingent up-

on accuracy to further clarify this relationship. Second, participants received advice only after they had already decided, which agent they wanted to receive advice from (which was necessary given the current research question). In line with previous observations and as mentioned earlier, this might have caused cognitive anchoring and could have lead to an underestimation of compliance rates in the current experiment, as well as muted differences in trust between agent types. Future studies should examine whether compliance to agent advice is modulated by the order in which advisor choice and advice reception occur. Third, the assumption that human advisors are associated with high social skills, while machine agents are associated with high analytical skills is based on theoretical assumptions (Gray et al., 2007; Chidambaram, Chiang, & Mutlu, 2012; Andrist, Spannan, & Mutlu, 2013), as well as empirical findings that human agents are trusted more on social tasks and machine agents are trusted more on analytical tasks (Smith et al., 2016). However, actual beliefs about the social and analytical capabilities of different advisor types or general attitudes towards machine agents were not assessed. The reason why we did not obtain these measures is that we were concerned that inquiring about the agents' skill sets would potentially lead to biased data: measuring beliefs before participants perform the task might make pre-existing beliefs unnecessarily explicit or even reveal too much about the experimental manipulation, while assessing beliefs after task completion would be biased by advisor choices and interactions with the agents throughout the task and not accurately reflect pre-existing beliefs. Future research should address this limitation by explicitly manipulating beliefs regarding agent expertise via instructions and by assessing a-priori beliefs about the social and analytical abilities of computer outside the experimental setting (similar to Gray et al., 2007).

In summary, the presented results show that agent choice and subsequent trust behaviors are influenced by the to-be-performed-task. The findings expand on the existing body of litera-

ture by demonstrating that nonhuman agents are not simply ill favored or distrusted in general, but that the agent-task fit seems to be taken into account when choosing and trusting human and nonhuman agents. A better understanding of behavioral tendencies when receiving advice from nonhuman agents will allow for the design of safer and more effective automated systems. The use of automated agents (i.e., computers or robotic agents) in military, medical, or domestic applications may be better informed based on the contextual factors that influence trust with those agents (in line with Hancock et al., 2011). For example, consider investment firms that intend to create programs using automated advisors to help customers make decisions about what stocks to buy or sell, or a military pilot using automated loyal wingmen that are giving him advice regarding potential targets. Knowing that human and non-human advisors are perceived differently and that task demands may alter the use of advice from these agents should inform the final designs of these systems better. In both cases, it is important to understand when advice from non-human agents is used, misused, or disused and designers of these systems should apply these findings to create systems that avoid dangerous and unwanted outcomes and result in optimized human performance.

REFERENCES

- Biele, G., Rieskamp, J., Krugel, L. K., & Heekeren, H. R. (2011). The neural basis of following advice. *PLoS biology*, 9(6), e1001089.
- Bonaccio, S., & Dalal, R. S. (2006). Advice taking and decision-making: An integrative literature review, and implications for the organizational sciences. *Organizational Behavior and Human Decision Processes*, 101(2), 127-151.
- Buckner, R. L., Andrews-Hanna, J. R., & Schacter, D. L. (2008). The brain's default network. *Annals of the New York Academy of Sciences*, 1124(1), 1-38.
- Connors, M. M. (1990). Crew system dynamics: Combining humans and automation. In *19th Intersociety Conference on Environmental Systems, San Diego, CA*. SAE Technical Paper 891530.
- de Visser, E. J., Krueger, F., McKnight, P., Scheid, S., Smith, M., Chalk, S., & Parasuraman, R. (2012, September). The world is not enough: Trust in cognitive agents. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 56(1), 263-267.
- Dijkstra, J. J., Liebrand, W. B. G., & Timminga, E. (1998). Persuasiveness of expert systems. *Behavior & Information Technology*, 17(3), 155-163.
- Dzindolet, M. T., Peterson, S., & Pomranky, R. (2003). The role of trust in automation reliance. *International Journal of Human-Computer Studies*, 58(6), 697-718.
- Dzindolet, M. T., Pierce, L. G., Beck, H. P., Dawe, L. A., & Anderson, B. W. (2001). Predicting misuse and disuse of combat identification systems. *Military Psychology*, 13(3), 147-164.

- Dzindolet, M. T., Pierce, L. G., Beck, H. P., & Dawe, L. A. (2002). The perceived utility of human and automated aids in a visual detection task. *Human Factors*, 44, 79–94. doi: 10.1518/0018720024494856
- Engelmann, J. B., Capra, C. M., Noussair, C., & Berns, G. S. (2009). Expert financial advice neurobiologically “offloads” financial decision-making under risk. *PLoS one*, 4(3), e4957
- Fox, J. E. (1996, April). The effects of information accuracy on user trust and compliance. In *Conference Companion on Human Factors in Computing Systems* (pp. 35-36). ACM.
- Gino, F., & Moore, D. A. (2007). Effects of task difficulty on use of advice. *Journal of Behavioral Decision Making*, 20, 21-35.
- Goetz, J., Kiesler, S., & Powers, A. (2003, October). Matching robot appearance and behavior to tasks to improve human-robot cooperation. In *Robot and Human Interactive Communication, 2003. Proceedings. ROMAN 2003. The 12th IEEE International Workshop on* (pp. 55-60). IEEE.
- Golan, O., Baron-Cohen, S., Hill, J. J., & Golan, Y. (2006). The 'reading the mind in films' task: complex emotion recognition in adults with and without autism spectrum conditions. *Social Neuroscience*, 1(2), 111-123.
- Goodyear, K., Parasuraman, R., Chernyak, S., Madhavan, P., Deshpande, G., & Krueger, F. (2016). Advice taking from humans and machines: An fMRI and effective connectivity study. *Frontiers in Human Neuroscience*, 10, 542.
- Goodyear, K., Parasuraman, R., Chernyak, S., de Visser, E., Madhavan, P., Deshpande, G., & Krueger, F. (2017). An fMRI and effective connectivity study investigating miss errors dur-

- ing advice utilization from human and machine agents. *Social neuroscience*, 12(5), 570-581.
- Gray, K., & Wegner, D. M. (2012). Feeling robots and human zombies: Mind perception and the uncanny valley. *Cognition*, 125(1), 125-130.
- 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.
- Harvey, N. & Fischer, I. (1997). Taking advice: Accepting help, improving judgement, and sharing responsibility. *Organizational Behavior and Human Decision Processes*, 70(2), 117-133.
- Henrich, J. & McElreath, R. (2003). The evolution of cultural evolution. *Evolutionary Anthropology*, 12, 123-135.
- Hertz, U., Palminteri, S., Brunetti, S., Olesen, C., Frith, C., & Bahrami, B. (2017). Neural computation underpinning the strategic management of influence in advice giving. *Nature Communications*, 8. doi: 10.1038/s41467-017-02314-5
- Jian, J. Y., Bisantz, A. M., & Drury, C. G. (2000). Foundations for an empirically determined scale of trust in automated systems. *International Journal Cognitive Ergonomics*, 4(1), 53-71.
- Koh, Y. J., & Sundar, S. S. (2010). Effects of specialization in computer, web sites, and web agents on e-commerce trust. *International Journal of Human-Computer Studies*, 68(12), 899-912.

- Krach, S., Hegel, F., Wrede, B., Sagerer, G., Binkofski, F., & Kircher, T. (2008). Can machines think? Interaction and perspective taking with robots investigated via fMRI. *PloS one*, 3(7), e2597.
- Lee, J., & Moray, N. (1992). Trust, control strategies, and allocation of functions in human-machine systems. *Ergonomics*, 35, 1243–1270. doi:10.1080/00140139208967392
- Lerch, F. J., Prietula, M. J., & Kulik, C. T. (1997). The Turing effect: The nature of trust in expert system advice. In Feltovich, P. J., & Fords, K.M. (Eds.), *Expertise in Context: Human and Machine* (pp. 417–448). Cambridge, MA: The MIT Press.
- Liu, C. (2010). Human-machine trust interaction: A technical overview. *International Journal of Dependable and Trustworthy Information Systems*, 1(4), 61-74.
- Madhavan, P. & Wiegmann, D. A. (2005). Cognitive anchoring on self-generated decisions reduces operator reliance on automated diagnostic aids. *Human Factors*, 47, 332–341. doi: 10.1518/0018720054679489
- Madhavan, P., & Wiegmann, D. A. (2007). Similarities and differences between human–human and human–automation trust: An integrative review. *Theoretical Issues in Ergonomics Science*, 8(4), 277-301.
- Mayer, R. C., Davis, J. H., & Schoorman, F. D. (1995). An integrative model of organizational trust. *Academy of management review*, 20(3), 709-734.
- Momen, A., Sebrechts, M. M., & Allaham, M. M. (2016, September). Virtual agents as a support for feedback-based learning. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* (Vol. 60, No. 1, pp. 1780-1784). Sage CA: Los Angeles, CA: SAGE Publications.

- Muir, B. (1994). Trust in automation: Part 1. Theoretical issues in the study and human intervention in automated systems. *Ergonomics*, 37, 1905–1923. doi:10.1080/00140139408964957
- Muir, B., & Moray, N. (1996). Trust in automation. Part II. Experimental studies of trust and human intervention in a process control simulation. *Ergonomics*, 39(3), 429–460. doi:10.1080/00140139608964474
- Picard, R. W. (1997). *Affective computing*. Cambridge, MA: MIT Press.
- Ramberg, R. (1996). People's perception of human and computer advice. *Computers in Human Behavior*, 12(1), 17-27.
- Rempel, J. K., Holmes, J. G. and Zanna, M. P. (1985). Trust in close relationships. *Journal of Personality and Social Psychology*, 49, 95–112.
- Shanteau, J. (1992). Competence in experts: The role of task characteristics. *Organizational Behavior and Human Decision Processes*, 53, 252-266.
- Shaefer, K. E., Billings, D. R., Szalma, J. L., Adams, J. K., Sanders, T. L., Chen, J. Y. C., & Hancock, P. A. (2014). A meta-analysis of factors influencing the development of trust in automation: Implications for human-robot interaction. *Army Research Laboratory* (No. ARL-TR-6984).
- Smith, M. A., Allaham, M. M., & Wiese, E. (2016, September). Trust in automated agents is modulated by the combined influence of agent and task type. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* (Vol. 60, No. 1, pp. 206-210). Sage CA: Los Angeles, CA: SAGE Publications.
- Snizek, J. A., & Van Swol, L. M. (2001). Trust, confidence, and expertise in a judge-advisor system. *Organizational Behavior and Human Decision Processes*, 84(2), 288-307.

Transportation Safety Administration (2014). *Transportation Safety Administration*. Available

at: <https://www.tsa.gov/travel/security-screening>

Uddin, L. Q. (2017). *Salience network of the human brain*. London: Academic Press.

Van Swol L. M., & Sniezek J. A. (2005). Factors affecting the acceptance of expert ad-

vice. *British Journal of Social Psychology*, 44(3) 443–461. doi:

10.1348/014466604X17092

Wiegmann, D. A., Rich, A., & Zhang, H. (2001). The effects of aid reliability on users' trust and reliance. *Theoretical Issues in Ergonomics Science*, 2(4), 352-367.

Ziefle M., & Jakobs, E.M. (2010). New challenges in human computer interaction: Strategic directions and interdisciplinary trends. In *4th International Conference on Competitive Manufacturing Technologies*, 389–398.

FIGURES

Figure 1. *Task Sequence:* task first condition with social task (top) and agent first condition with analytical task (bottom). First, the combined *Reading the Mind in Film* and arithmetic stimuli would be shown. In the task first condition, participants would then be asked to EITHER determine which emotion a specific character was displaying OR solve a math problem with the dots. Then the participant would choose what agent they would like to see answer the question, followed by that agent's answer and a chance to finalize their answer. In the agent first condition, participants are first asked to choose which agent they would like to see answer, before they know what question they will be presented with. Film clips ranged from 5 to 30 seconds and immediately continued upon completion. There were no time constraints on the decision-making components of the task.

Figure 2. *Agent Representations:* Agents used in the experiment. Human and computer agents were adapted from deVisser et al. (2012).

Figure 3. Agent Choice Results: The mean frequency percentage each agent was selected for both tasks separated by condition, agent first (left) and task first (right). Error bars represent standard error and points are offset horizontally so that error bars are visible. A $2 \times 3 \times 2$ mixed ANOVA showed a main effect for *Agent*, indicating a significant difference in selection across agent types, as well as significant interactions of *Agent* $\times$ *Task*, indicating that selection of agents depended on task type, and *Condition* $\times$ *Agent* $\times$ *Task*, indicating that participant's selections differed depending on task type and condition.

Figure 4. The mean compliance percentage for each agent by task-type, separated by condition; agent first condition (left) and task first condition (right). Error bars represent standard error and points are offset so that error bars are visible. Compliance was calculated as the percentage of times participants changed their initial response to match the advice given by the chosen agent. A linear mixed-effects model with restricted maximum likelihood revealed a significant main effect of *Task*, with higher compliance rates in the social versus the analytical task, and a three-way interaction of *Human* x *Condition* x *Task* was also significant.