Trust in Artificial Intelligence: Comparing trust processes between human and automated trustees in light of unfair bias

Markus Langer^a, Cornelius J. König^a, Caroline Back^a, and Victoria Hemsing^a

Universität des Saarlandes

This is a preprint version of this paper. The paper is currently under review and may change during future revisions. Copyright: Markus Langer, Universität des Saarlandes, Saarbrücken, Germany.

Markus Langer, <https://orcid.org/0000-0002-8165-1803>

Cornelius J. König, <https://orcid.org/0000-0003-0477-8293>

^a Department of Psychologie, Universität des Saarlandes, Saarbrücken, Germany

This study was pre-registered on AsPredicted (<https://aspredicted.org/blind.php?x=gh8xa4>)

Correspondence concerning this article should be addressed to Markus Langer, Universität des Saarlandes, Arbeits- & Organisationspsychologie, Campus A1 3, 66123 Saarbrücken, Germany.

Phone: +49 681 302 4767. E-mail: Markus.Langer@uni-saarland.de

ABSTRACT

Automated systems based on artificial intelligence (AI) increasingly support decisions where ethical considerations are central to trust development in human-system collaboration. However, insights regarding trust in automated systems predominantly stem from contexts where the main driver of trust is that systems produce accurate outputs (e.g., alarm systems for monitoring tasks). It remains unclear whether what we know about trust in automated systems translates to application scenarios where ethical considerations (e.g., fairness) play a crucial role in trust development. In a personnel selection context, as an example of task where ethical considerations are important, we investigate trust processes in light of a trust violation relating to unfair bias and a trust repair intervention. Specifically, participants evaluated applicant preselection outcomes by either a human or an automated system across twelve personnel selection tasks. We additionally varied information regarding imperfection of the human and automated system. In task rounds five through eight, the preselected applicants were predominantly male, thus constituting a trust violation due to potential unfair bias. Before task round nine, participants received an excuse for the biased preselection (i.e., a trust repair intervention). Results showed that participants initially perceived automated systems to be less trustworthy. Furthermore, the trust violation and the trust repair intervention had weaker effects for the automated system. Those effects were partly stronger when highlighting system imperfection. We conclude that insights from classical areas of automation only partially translate to the many emerging application contexts of such systems where ethical considerations are central to trust processes.

Keywords: Artificial intelligence, Trust, Personnel Selection, AI ethics

Introduction

Recent years have seen an upsurge in the use of automated systems based on artificial intelligence (AI) to support or even automate decision-making. Whereas classical application areas of automation were production or monitoring (Endsley, 2017), systems are now employed in tasks that affect the fate and future of individuals such as in medicine, jurisdiction, or management (Grgić-Hlača et al., 2019; Longoni et al., 2019; Raisch & Krakowski, 2021). In management alone, AI-based systems are on the verge of changing decision processes in personnel selection, performance management, or promotion (Cheng & Hackett, 2021; Tambe et al., 2019). Consequently, managers can increasingly choose to assign tasks to either human trustees (the party that is trusted) or automated systems as trustees, which warrants the need for managers as trustors (the party that trusts) to continuously assess the trustworthiness of humans and automated systems in order to decide whether to rely on a respective trustee to perform a certain task.

Whereas previous research has shown that there are similarities in trustworthiness assessment and in related longitudinal trust processes for humans and automated systems as trustees (de Visser et al., 2018; Glikson & Woolley, 2020), this research predominantly stems from classical application contexts of automation (e.g., monitoring). In these contexts, what trustors evaluate when assessing trustworthiness is focused on classical performance measures associated with effectiveness and efficiency (e.g., prediction accuracy). In comparison, there is a paucity of research investigating trust processes in contexts where automated systems support decisions that affect the fate of individuals (e.g., personnel selection; Langer et al., 2020). In such contexts, practitioners, researchers, and policy-makers are commonly concerned about ethical issues when using automated systems as decision support (Jobin et al., 2019; Martin, 2019). Especially outcome fairness is important in such contexts (Raghavan et al., 2020) and determines trustworthiness assessments beyond classical performance measures. In fact, fairness issues (e.g., biases introduced

or amplified by automated systems) have led to companies like Amazon losing trust and abandoning automated systems for managerial decisions (Dastin, 2018).

As previous research on trust in automation is dominated by work investigating trust processes associated with classical performance measures (e.g., whether systems produce accurate outputs), there is a lack of insight regarding trust processes in tasks where there is potential for violations of ethical standards such as fairness. For instance, although it is a matter of life and death to adequately trust in air-traffic-control systems, such systems usually violate operators' trust through misses or false alarms (Parasuraman & Riley, 1997). In such illustrative contexts for the majority of previous research on trust in automated systems, fairness issues do not play an obvious role. However, such issues are salient throughout many novel application areas of automated systems (Raghavan et al., 2020). Thus, it remains unclear whether effects found in classical trust in automation research translate to novel application contexts where accuracy is only one of many factors determining trust in those systems.

This study compares trust processes (initial perceptions, trust violations, trust repair interventions; de Visser et al., 2018) between human trustees and automated system as trustees in a novel application context for automated systems and in light of a potential violation of ethical standards. Specifically, participants took part in twelve rounds of a personnel selection task, where they received decision support from either a human colleague or an automated system. In every task round, participants received a preselection of applicants. After the fourth task, participants repeatedly received biased preselection outputs (i.e., trustees predominantly preselected male applicants) constituting trust violations due to violations of ethical standards. Before round nine, participants received an excuse for the biased preselection constituting a trust repair intervention. We investigate initial perceptions of trustworthiness, trust and reliance, as well as perceptions in response to the trust violation and the trust repair intervention. Furthermore, we investigate

expectations of perfection (high performance, high consistency in performance) associated with automated systems as a driver of potential differences in trust processes in relation to human and automated trustees (Madhavan & Wiegmann, 2007).

Background and Hypotheses Development

Interpersonal Trust and Trust in Automation

Trust processes are central when trustors consider delegating tasks to trustees and when they receive advice or decision-support (Bonaccio & Dalal, 2006). In those cases, the task fulfillment or the work output is important to trustors (Mayer et al., 1995). Consequently, relying on trustees' work outputs involves risk as trustees might not fulfill the trustors' expectations. This is true for interpersonal trust processes and in the case of trust in automation (J. D. Lee & See, 2004).

Theoretical models on interpersonal trust (Mayer et al., 1995) and trust in automation (J. D. Lee & See, 2004) show that independently whether support is provided by a human or an automated system, main concepts in trust processes are trustworthiness, trust, and reliance. Trustor's trustworthiness assessments arise from known or perceived characteristics of the trustee as well as trustees' performance in a task (J. D. Lee & See, 2004; Mayer et al., 1995). In the current study, we chose a personnel selection task which means that trustors will assess trustworthiness in regard to their goals in relation to this task. We decided for personnel selection as it reflects a context demanding complex, ethically sensitive decisions, and where AI¹-based automated systems are already a viable options for decision-support (Hickman et al., 2021).

Trustworthiness of humans and automated systems is usually conceptualized with several facets (J. D. Lee & See, 2004; Mayer et al., 1995). In this study, we capture four facets of

¹ When referring to AI, we explicitly include classical approaches such as rule-based systems but also novel approaches from the area of machine learning and deep learning (Cheng & Hackett, 2019).

trustworthiness that trustors may consider in relation to human and automated trustees: ability, flexibility, integrity, and benevolence (Höddinghaus et al., 2020; Wang & Benbasat, 2005). In the case of personnel selection, ascribing high ability means that trustors believe that trustees can successfully select suitable applicants (Langer et al., 2020). Flexibility reflects whether trustors assume that trustees can react flexibly to changes in a given task (Höddinghaus et al., 2020). Integrity captures whether trustors believe that trustees provide unbiased recommendations and is in line with evaluating if trustees follow ethical standards that trustors value (Den Hartog & De Hoogh, 2009; Gilliland, 1993). Finally, benevolence reflects if people believe that trustees consider trustors' interests, goals, and values (Höddinghaus et al., 2020). The second main concept within trust processes is trust. Mayer et al. (1995 p. 712) define trust as a willingness to be vulnerable towards the actions of trustees without explicitly controlling or supervising trustees. Meeßen et al. (2020) describe trust as an intention to trust trustees. In interpersonal trust and trust in automation, if people estimate trustworthiness of trustees to be high, trust will tend to also be high and thus it will be more likely that trustors will actually rely on trustees to fulfil a given task. Reliance then is the behavioral outcome of trust (J. D. Lee & See, 2004). This means that trustors actually delegate tasks to trustees or follow advice provided to them. This implies that trustors accept the risk that they may be disappointed by the outcome produced, may need to respond to a failure in a task, or may receive bad advice (Mayer et al., 1995). Depending on the respective outcome (e.g., the work output, the quality of advice), trustors will re-evaluate trustees' trustworthiness starting another cycle in the trust process (J. D. Lee & See, 2004; Mayer et al., 1995).

Thus, trust processes are dynamic (Glikson & Woolley, 2020). Specifically, there is an initial level of trustworthiness, trust, and reliance (i.e., when initially interacting with trustees) and respective levels can increase over time if trustors perceive trustees to perform successfully (J. D. Lee & See, 2004; Mayer et al., 1995). However, relying on someone or something always runs the

risk of unfulfilled expectations, and trust violations can reduce trustworthiness, trust, and reliance (Kim et al., 2006). It is, however, possible to repair trust (e.g., through excuses) which can then rebuild, and positively affect perceptions of trustworthiness, trust and reliance (de Visser et al., 2018; Tomlinson & Mayer, 2009).

Differences in Trust Processes over Time

Whereas the basic concepts as well as dynamics (e.g., effects of trust violations or trust repair interventions) exist for interpersonal trust and for trust in automation, previous research has indicated that assessment of trustworthiness, trust, and reliance might differ for human versus systems as trustees (de Visser et al., 2018; Madhavan & Wiegmann, 2007). These differences might be present as a direct effect on initial trustworthiness assessments. Additionally, they might be reflected in moderating effects regarding reactions to outcomes, behavior, and information associated with the trustee. In other words, there might be different reactions to trust violation and trust repair interventions depending on whether the trustee is human or an automated system (de Visser et al., 2016). Figure 1 summarizes the main concepts and their relation in trust processes and locates our hypotheses and research question in the trust process as proposed by Mayer et al. (1995). The following sections provide the rationale for our proposed hypotheses and research questions.

Initial Trustworthiness Assessments

There is research assuming that initial trustworthiness assessments regarding automated systems might be comparably higher than for human trustees. For instance, Madhavan and Wiegmann (2007) integrate research on trust in automated systems in classical application contexts for such systems and propose that humans initially believe that systems are designed for a specific purpose, and have been tested before they are deployed. Thus they should work as intended leading to high levels of initial trustworthiness (Parasuraman & Manzey, 2010). Consequently, trustors

should start with lower levels of trustworthiness for human trustees because of more uncertainty about their abilities, values, and intentions (Madhavan & Wiegmann, 2007).

In contrast to the propositions by Madhavan and Wiegmann, there is research indicating that trustworthiness assessments could be lower for automated systems compared to human trustees especially in tasks that affect the fate of individuals such as personnel selection (M. K. Lee, 2018; Longoni et al., 2019). This might be partly because people believe automated systems are less capable of fulfilling such tasks (M. K. Lee, 2018). Specifically, people may believe that automated evaluation of individuals' characteristics does not adequately capture respective task complexity, is dehumanizing, and does not consider ethical peculiarities in such contexts (Grove & Meehl, 1996; Newman et al., 2020).

Overall, research stemming from classical areas of automation would suggest high levels of initial trustworthiness regarding automated systems. In novel application contexts, there are arguments indicating that initial assessments of system trustworthiness could be higher or lower compared to human trustees. Considering that trust and reliance should result from initial trustworthiness assessments, we thus pose the following research question:

Research Question (RQ) 1²: Is there an initial difference for trustworthiness assessments, trust, and reliance between the automated system and the human trustee?

Implications of Trust Violations

The re-evaluation of trustworthiness after trust violations might be moderated by different expectations and assumptions regarding human and automated systems as trustees. In fact, trustworthiness assessments seem to suffer comparably strong as trustors experience the first error by an automated system (Bahner et al., 2008; Dzindolet et al., 2003). It is commonly assumed that

² Research Questions and Hypotheses were preregistered under <https://aspredicted.org/blind.php?x=gh8xa4>.

this first-error-effect results from people having high expectations and high initial levels of trustworthiness regarding automated systems, as they believe systems are tested and delivered to work as intended (Madhavan & Wiegmann, 2007). This is also associated with perfection assumptions regarding systems, meaning that systems should show consistently high performance (Dietvorst & Bharti, 2020; Parasuraman & Manzey, 2010). When recognizing first errors, people realize that systems are not perfect, and that system performance varies, thereby strongly negatively affecting trustworthiness assessment (Dzindolet et al., 2003). In contrast, trustors might not expect human trustees to constantly work near-to-perfection, thus errors might be in line with what can be expected leading to comparably weaker trust violation effects (Madhavan & Wiegmann, 2007).

However, research suggesting this moderating effect also predominantly stems from classical application areas of automation where trust violations are mainly associated with trustee characteristics that relate to the trustworthiness facet ability. For instance, in such studies, trustees would miss alarms, produce false alarms, or provide less than perfect predictions (Hoff & Bashir, 2015). In novel application areas such as personnel selection, trust violations can also be ability-associated meaning that a trustee recommends unsuitable applicants. Additionally, trustees can produce ethically questionable outcomes (i.e., discriminating minority applicants) that can affect trustworthiness assessments (Kim et al., 2006). In cases where trust violations are based on violations of ethical considerations, people might have stronger negative reactions for human trustees as they may believe that automated systems do not actively discriminate against specific groups of people (Bigman et al., 2020). However, previous trust in automation research would suggest that errors by automated systems result in strong negative effects regarding trustworthiness which is why we propose:

Hypothesis 1³: After a trust violation, trustworthiness, trust, and reliance towards the automated system as trustee will decrease more compared to the human trustee.

Implications of Trust Repair Interventions

Similar moderating effects associated with the nature of the trustee might also occur for trust repair interventions. Specifically, systems are deployed with a specific set of functions and level of quality. Any system improvements might require system updates or a completely new system (Höddinghaus et al., 2020). In contrast, human trustees might be expected to be more adaptable. If there are trust repair interventions for human trustees (e.g., excuses), trustors might assume that the human trustee will do their best to not let this error happen again (de Visser et al., 2018; Tomlinson & Mayer, 2009). If there are trust repair interventions associated with automated trustees, trustors might still believe that a system will produce similar errors in future as they may perceive systems to possess a fixed set of attributes (Dietvorst et al., 2015). Thus, we conclude,

Hypothesis 2: After the trust repair intervention, perceptions of trustworthiness, trust, and reliance towards an automated system as trustee will increase less compared to a human trustee.

Differences in Trustworthiness Facets

The previously mentioned differences between human and automated trustees might also have implications for single facets of trustworthiness. Regarding ability, trustors might start off with higher levels of ability assessments regarding automated systems, but ability perceptions could suffer more strongly from trust violations and might be less affected by trust repair interventions (Madhavan & Wiegmann, 2007). Regarding flexibility, automated systems might be

³ In the preregistration we mentioned moral reasoning as dependent variable but, for reasons of readability, we do not include results for this variable. Additionally, we measured perceived transparency but received feedback on previous versions of this manuscript questioning the value of measuring transparency in our study design. Data and results can be made available upon request. In the preregistration, we mentioned the following research question: “Will participants realize trust violations by the automated agent later?” but omitted those analyses as we realized they do not provide insights beyond the other hypotheses.

perceived as less flexible than human trustees (Dietvorst et al., 2015), but we are not aware of research indicating how this changes through trust violation and trust repair interventions. Regarding integrity, automated systems might be assessed as more consistent and less biased compared to human trustees (Langer et al., 2019; M. K. Lee, 2018). In addition, for human trustees, trust violation associated with ethical considerations (e.g., a biased preselection) might have stronger effects as people might be more outraged by such trust violations in the case of a human trustee (Bigman et al., 2020). For trust repair effects, it may be possible to assume stronger effects for human trustees as people believe that humans can learn from their mistakes (Tomlinson & Mayer, 2009). However, it is also possible that humans may not believe that biased human trustees can change (i.e., assuming that biases could reflect stable attitudes). For benevolence, humans might be perceived as more benevolent than automated systems as they are more likely to be able to consider trustors' interests (Höddinghaus et al., 2020). However, we are not aware of research that investigated benevolence of human and automated systems with respect to trust violations and trust repair interventions. This list of tentative assumptions regarding the facets of trustworthiness subsume under the following research question:

RQ2: Is there an initial difference and different effects for trust violations and trust repair interventions for the facets of trustworthiness regarding human and automated systems as trustees.

Assumptions of Perfection as a Driver of Differences Between Human and Automated Trustees?

The previous sections indicate that one of the main arguments from previous research regarding differences between human trustees and automated systems as trustees is an assumption of perfection regarding automated systems (Hoff & Bashir, 2015; Madhavan & Wiegmann, 2007). We thus experimentally manipulate this assumption by emphasizing that the automated system or the human trustee might not always produce error-free outputs, thus highlighting their potential

imperfection and potential for performance variations (Bahner et al., 2008). If the assumption of perfection causes differences in trust processes between human and automated systems as trustees, making potential imperfection salient might affect perceptions of automated systems in a way that is more similar to reactions to human trustees (de Visser et al., 2016). Thus, we propose:

Hypothesis 3: Trustworthiness, trust, and reliance in the automated system when there is information about imperfection presented will initially be lower compared to when no such information is presented.

Hypothesis 4: Following a trust violation, trustworthiness, trust and reliance in the automated system when there is information about imperfection presented will decrease to a lesser extent compared to when no such information is presented.

RQ3: Will there be interaction effects between the trust repair intervention and the information regarding imperfection for the automated system on trustworthiness, trust, and reliance?

Method

Sample

We determined the required sample size with G*Power (Faul et al., 2009). In the preregistration we proposed to conduct ANOVAs. For a within-between interaction effect of $\eta^2_p = 0.01$, $N = 108$ participants would be necessary for a power of $1 - \beta = .80$. Assuming a small to medium effect for a between groups effect of $\eta^2_p = 0.04$, $N = 148$ participants would be necessary to achieve a power of $1 - \beta = .80$. Therefore, we wanted to recruit between 108 and 148 participants. Ultimately, we decided that analyzing the data using regressions and hierarchical linear models was more appropriate. This study was advertised to people interested in Human Resource Management. We posted the advertisement on different social media groups, around the campus of a German university, and in the downtown area of a German city.

As we anticipated potential issues during data collection (e.g., technical issues), we continued data collection until our sample consisted of $N = 211$ participants. We excluded 10 participants because they did not follow the instructions in the experimental procedure, 3 because they reported technical issues, 30 because they indicated that they received advice from a human agent when they actually were in the automated system condition and vice versa, 44 because they failed the manipulation check regarding imperfection (i.e., indicating they were not told that the trustee can make errors when they were actually told so or vice versa), and 3 participants because they were not able to recall that there were task rounds where male applicants had been predominantly selected. The final sample consisted of $N^4 = 121$ German participants (79% female) of which 93% studied psychology. The mean age was 23.56 years ($SD = 5.47$), participants took a mean of 32 minutes ($SD = 10$) to complete the study, had experienced a median of three personnel selection processes as applicants, and 24% of participants indicated experience in applicant selection.

Procedure

In a 2 (human vs. automated trustee) $\times$ 2 (no information regarding imperfection vs. information regarding imperfection) online experiment, participants were randomly assigned to one of the four groups. Participants were instructed to imagine that they were responsible for personnel selection in an HR department at a large insurance company operating in Germany. They were informed that their company is recruiting trainees for its subsidiaries. We chose trainees at an insurance company because of the equal gender distribution in insurance jobs in Germany (Rudnicka, 2020). Furthermore, participants were informed that they will receive support in

⁴ If we include participants who did not respond correctly to the manipulation checks regarding whether the advice came from a human or a system, whether there were information regarding imperfection, and whether there were rounds where there were predominantly male applicants (resulting in $N = 188$), the results remain stable (i.e., regarding significances and directions of effects).

selecting candidates through a preselection of applicants from a larger applicant pool across 12 situations (representing consecutive days where there are new applications). Participants were then informed that their task was to evaluate the quality of the applicant preselection in regard to three organizational goals. These goals were customer satisfaction, innovation, and diversity. For each goal, participants received a text reflecting how the company defines these goals (see Supplemental Material A). The purpose of this was to introduce diversity explicitly so that not providing a diverse preselection would indicate a trust violation. Note that we chose two more goals in order to (a) provide participants with quality criteria for the preselection, (b) make the task more realistic, and (c) make the focus on diversity less obvious.

After being introduced to their task and the company's goals, participants received a job description including the desired qualifications of applicants. Then, participants received a tutorial round introducing them to the general process of each upcoming task round. In the tutorial (and in each of the task rounds) participants received between six and eight application photo-like pictures of white male and female applicants (see Supplemental Material B). We chose to only include white applicants as balancing for racial diversity would have greatly increased study complexity. The pictures were taken from <https://generated.photos>, a website that uses AI to produce realistic pictures of human faces.

For each applicant, there was additional information accompanying the picture (i.e., family name, age, years of job experience, final high school grade, strengths). During the tutorial, participants were informed that they should analyze the preselection against the company goals and the job description. They were also told that they will receive the question: "Do you want to see the statistics of the applicant pool?" If participants responded with "No" they were directed to the next page in the online tool. If participants responded with "Yes", they were directed to a page showing them the underlying statistics of the applicant pool for a given task round. In these

statistics, participants were informed about the number of applicants, the percentage of male and female applicants as well as provided with a list of means and distribution information for further information (e.g., “the mean number of years of job experience of today’s applicant pool was 1.3 with a deviation of 0.3”). In the tutorial, participants were instructed to respond with “Yes” so that every participant would see what happens if they request the applicant pool statistics. We included the option to view the applicant pool to increase psychological fidelity, as well as giving participants the option to check whether there was gender diversity in the applicant pool.

At the beginning of the actual experiment, participants received information that there was an increasing number of applicants making it necessary to have a preselection stage where they as hiring manager receive support. They were told that this support is a colleague (an automated system) who analyzes and preselects applicants. In the information regarding imperfection condition, participants were additionally informed that “the colleague (automated system) usually produces good work outcomes but that there are always possibilities for errors.” Participants were then informed that they had the opportunity to evaluate the preselection and accept or reject it. Participants then saw the information of the preselected applicants.

Afterwards, participants responded to the fairness item, to three items assessing whether the trustee adhered to the company’s goals (customer satisfaction, innovation, diversity), to the trustworthiness (i.e., ability, flexibility, integrity, benevolence; only in situations 1, 3, 5, 7, 9, and 12) and trust items. Participants then indicated whether they wanted to see the applicant pool statistics. After clicking “No,” or after seeing those statistics, participants were asked: “Do you accept this preselection?”, which was used as a measure of reliance. This process was repeated for twelve task rounds.

In task rounds five through eight, there were predominantly male applicants included in the preselection (see Supplemental Material B). Thus, rounds five through eight constitute the trust

violation phase. Before round nine, participants received the following information: “Dear colleague, in previous preselection outcomes, there were more male than female applicants, although the applicant statistics indicated that there were about as much women as there were men in the applicant pool. We were made aware of this issue and it has been solved. We apologize for this. Applicant preselection in the future should follow the goals of the organization again. Thank you for your understanding.” This constituted the trust repair intervention. After the last task round, participants responded to exploratory items asking them whether they were satisfied with the advisor and if they would want to work together with this advisor in future. Finally, we measured propensity to trust in humans and in technology and collected demographic information.

Measures

Since participants in the main study needed to report their evaluation of trustees several times throughout the study, we wanted to optimize the use of items to reduce participant fatigue. In order to determine which items to keep, we conducted a pre-study with 54 student participants. In this pre-study, participants imagined to work at an HR department. They were told that they will receive support for a personnel selection task and were randomly assigned to either the human or the automated trustee condition. Participants then responded to fourteen items assessing perceptions of ability, flexibility, integrity, benevolence, and trust regarding the trustee. We removed four items because they negatively influenced Cronbach’s Alpha reliability of the scales or because they did not work well in the human or the automated trustee condition (e.g., because participants indicated that they did not know how to respond to this item in regard to the respective trustee).⁵

⁵ Results of the pre-test can be made available upon request.

Unless otherwise stated, participants responded to the items reflecting the dependent variables on a scale from 1 (strongly disagree) to 7 (strongly agree). To keep the gender of the human trustee undefined, the items did not mention the colleague's gender (using the inclusive "Kolleg/in" in German).

Ability was measured with two items taken from Höddinghaus et al. (2020). A sample item was "I believe the colleague/the automated system has the competency to consider all important information for the decision." **Flexibility** was measured with two items taken from Höddinghaus et al. (2020). A sample item was "I believe the colleague/the automated system is able to react flexibly to given contexts." **Integrity** was measured with two items taken from Höddinghaus et al. (2020) who adapted the items from Wang and Benbasat (2005). A sample item was "I believe that the colleague/the automated system makes unbiased decisions." **Benevolence** was measured with two items taken from Wang and Benbasat (2005). A sample item was "The colleague/the automated system would consider my interests." **Trustworthiness** was calculated as the mean of the items for Ability, Flexibility, Integrity and Benevolence. **Trust** was measured with two items taken from Thielsch et al. (2018). A sample item was "I would strongly rely on the colleague/the automated system". **Reliance** was measured with the item "I accept this preselection" with the response options "Yes" or "No". Accepting the preselection reflects higher reliance.

Manipulation Checks

To assess if the trust violation and trust repair intervention worked as intended, we used two manipulation check measures. First, after receiving the preselection, participants responded to the item "In your opinion, did the preselection adhere to the company's goals?" on a scale from 1 (disagree) to 5 (agree). Participants responded to this item three times, once for each of the three company goals (customer satisfaction, technology and innovation, diversity). This reminded participants of the company's goals and allowed us to assess if participants realized that the biased

preselection violated the organization's diversity goal. Second, participants responded to the item "I perceived the preselection to be fair" which was meant to assess whether participants perceived a predominantly male preselection to be unfair.

Results

Table 1 displays correlations, means and SDs for the means of all variables across all tasks. Figures 2 and 3 show line-graphs for the continuous dependent variables for each of the tasks and Table 2 shows contingency tables for participants' rejection and acceptance of the preselection of applicants. We used regressions and hierarchical linear models (HLM) to analyze the data. Within the HLM analyses, measures were combined into one measure for each of the three phases: the initial phase (task rounds one through four), trust violation phase (task rounds five through eight), and the trust repair phase (task rounds nine through twelve). Each phase was entered into the HLM analyses through backwards difference coding meaning that results regarding a respective phase can be interpreted as a comparison to the previously entered phase (e.g., trust violation phase in comparison to the initial phase).

HLM analyses were also used to analyze the manipulation checks. Table 3 shows that both perceptions of fairness and evaluations of the fulfillment of the diversity goal decreased under the trust violation and recovered through the trust repair intervention indicating that the manipulations had the intended effects.

Research Questions and Hypotheses⁶

Tables 4 and 5 show the results of the regression analyses for initial trustworthiness and trust. Tables 6 and 7 show the results of the HLM analyses for the evaluation of trustworthiness, trust, and facets of trustworthiness over time.

⁶ Including propensity to trust in the analyses did not change the results in a way that would have changed the interpretation of the results. Those results can be made available upon request.

Trustworthiness, trust and reliance. RQ1 asked whether there are initial difference for trustworthiness assessments, trust, and reliance between the human and automated systems as trustees. We used linear regressions to analyze differences in trustworthiness and trust for only the first task as it reflects the most appropriate point in time for initial evaluations of trustworthiness, trust and reliance. The results indicated that participants perceived the human trustee as more trustworthy and they more strongly intended to trust the human compared to the automated trustee (Table 4). Regarding reliance, the percentage of participants rejecting the suggested preselection of applicants provided by the human (7.1% of participants rejected) or automated system (9.4% of participants rejected) in the first round showed no substantial differences. As such, we did not conduct further analyses. In sum, the answer to RQ1 is: trustworthiness and trust were initially higher for human trustees but there was no difference for reliance.

Hypothesis 1 proposed that after a trust violation, trustworthiness, trust, and reliance towards the automated system would decrease more compared to the human trustee. In contrast to this, Table 6 indicates that trustworthiness and trust decreased to a larger extent for the human trustee. For reliance, we initially analyzed the percentage of people who rejected the preselection depending on the conditions and the phases. For the initial phase and in the human condition, participants rejected the preselection in 27% of cases, and in the case of the automated system, 31% of cases were rejected. For the trust violation phase and in the human condition, participants rejected the preselection in 73% of cases compared to 72% in the case of the automated system indicating no differences in the decline of reliance. Overall, results did not support Hypothesis 1.

Hypothesis 2 suggested that after the trust repair intervention, trustworthiness, trust, and reliance towards the automated system would increase less than towards the human trustee. Table 6 shows that for trustworthiness and trust, this was the case. For reliance, the percentage of cases where participants rejected the human trustee's preselection in the trust repair phase was 28%, and

for the automated system's preselection, 31% of cases were rejected showing no substantial differences. This supports Hypothesis 2 for trustworthiness and trust, but not for reliance.

Facets of trustworthiness. RQ2 asked whether there is an initial difference and different effects for trust violations and trust repair interventions for the facets of trustworthiness regarding human and automated systems as trustees. Tables 5 and 7 summarize the findings for trustworthiness facets. Regarding initial assessments, ability and flexibility were lower, and integrity was higher for the automated system (Table 5). Regarding trust violations, effects were weaker for the automated system for all facets of trustworthiness (see Table 7 TV×Trustee). Similarly, trust repair intervention effects were weaker for the automated system for all facets of trustworthiness (see Table 7 TR×Trustee).

Effects of information regarding imperfection. Hypothesis 3 proposed that trustworthiness, trust, and reliance in the automated system when there is information about imperfection presented will initially be lower compared to when no such information is presented, however, we found no effects of information regarding imperfection and no interaction between the independent variables (Table 4). Thus, Hypothesis 3 was not supported.

Hypothesis 4 proposed that following a trust violation, trustworthiness, trust and reliance in the automated system when there is information about imperfection presented will decrease to a lesser extent compared to when no such information is presented. This was not supported for trustworthiness and trust, as seen in Table 6 (see Trust Violation×Trustee×Imperfection) and Figures 2 and 3. Instead, the effects of trust violations were stronger when information regarding imperfection was presented. Regarding reliance, in the initial phase the percentage of cases where participants rejected the preselection of the automated system when information regarding imperfection was presented was 29%, and in the case of the automated system with no such information presented it was 35%. In the trust violation phase, those rejection rates were 71%

(information regarding imperfection presented) and 72% (no such information presented) respectively. Those results indicated that there was no difference in the decline of reliance. Overall, these results did not support Hypothesis 4.

RQ3 asked whether there will be interaction effects between the trust repair intervention and the information regarding imperfection for the automated system on trustworthiness, trust, and reliance. Indeed, trust repair interventions were more effective at restoring trustworthiness and trust in the automated system when information regarding imperfection was presented (see Trust Repair×Trustee×Imperfection in Table 6). Regarding reliance, the percentage of cases where participants rejected the preselection of the automated system when information regarding imperfection was presented was 26% compared to 38% when no such information was presented. We calculated χ^2 -tests for the trust violation and the trust repair phase separately. For the trust violation phase, the difference between automated systems where information about imperfection was presented and where no such information was presented was not significant, $\chi^2(1) = 0.01, p = .92$, whereas there was a difference for the trust repair phase, $\chi^2(1) = 4.13, p < .05$. As a benchmark, for the human trustee the percentages of rejection in the trust violation phase were 72% (information about imperfection presented) and 74% (no such information presented), and 28% (information about imperfection presented) and 26% (no such information presented) in the trust repair phase which constituted no significant differences when calculating χ^2 -tests. The response to RQ3 therefore is: Trust repair interventions were more effective for automated systems as trustees when information regarding imperfection was presented.

Discussion

With the increasing use of AI-based automated systems in novel application context where they contribute to decisions over individuals' fates, it becomes crucial to understand trust processes in human-AI collaboration for such decisions (Glikson & Woolley, 2020; Raisch & Krakowski,

2021). We built our hypotheses and research questions on theoretical models describing general trust processes and on research regarding trust violations and trust repair interventions in classical application contexts of automation (Hoff & Bashir, 2015; Madhavan & Wiegmann, 2007). Overall, our findings imply that theoretical assumptions and effects found in classical application areas of automated systems only partly translate to novel application contexts of automated systems such as personnel selection. Specifically, in classical application contexts, where trustworthiness assessments mainly stem from classical system performance measures, people seem to hold an assumption of perfection regarding automated systems associated with high and consistent performance. Our results indicate that people do not assume high performance for automated systems in personnel selection as initial trustworthiness assessments were lower for automated systems. However, there seems to remain an assumption of consistency, as trust violation associated with potentially unfair bias as well as trust repair effects were weaker for automated systems. This assumption of consistency was partly reduced by highlighting system imperfection leading to stronger trust violation and trust repair effects. In sum, our study suggests that research assessing trust in automated systems must be aware of the application context in which systems support decision-making because although assumptions of systems as being consistent might generalize, assumptions of high-performance might not. This seems to affect trust processes associated with automated systems in respective contexts.

Theoretical Implications

Trust in Automation Depends on the Use Context

Whereas in classical application contexts people seem to evaluate overall trustworthiness of automated systems to be higher than human trustworthiness (Madhavan & Wiegmann, 2007), we found the opposite to be true for personnel selection. Classical tasks for automated systems usually afford mechanical skills such as combining large amount of data and performing repetitive

tasks (e.g., during monitoring tasks) (M. K. Lee, 2018). In contrast, personnel selection requires individual and flexible decision-making capabilities, as well as ethical considerations. These are capabilities that people more likely ascribe to humans (Bigman & Gray, 2018; M. K. Lee, 2018; Newman et al., 2020). In other words, people might believe that humans are better able to complete tasks where ethical issues, as well as individuals' unique characteristics need to be considered (Longoni et al., 2019; Newman et al., 2020). In line with this and with further previous research (Höddinghaus et al., 2020; M. K. Lee, 2018), our participants perceived the automated system to be less able and less flexible, but also perceived the integrity (associated with systems being less biased) of systems to be comparably stronger. Overall, these findings indicate that people have specific attitudes towards automated systems that affect their evaluations of trustworthiness of automated systems differently for different tasks (Elsbach & Stigliani, 2019). Those attitudes seem to lead to high overall levels of initial trustworthiness assessments for classical automation tasks (e.g., monitoring) that afford mechanical skills and where primary performance measures are effectivity and efficiency, and to comparably low initial trustworthiness assessments for tasks that involve consideration of ethical issues and decisions about individuals (e.g., personnel selection, performance evaluation; see also Nagtegaal, 2021).

In sum, this resulted in participants evaluating initial levels of trustworthiness to be stronger for human trustees, which also translated to stronger intentions to trust the human. This usually indicates that humans would also be more likely to rely on a recommendation provided by another human. However, overall effects regarding reliance were negligible. An explanation for this might be that differences in trustworthiness and trust were not strong enough which might have resulted from the fact that certain facets of trustworthiness would increase the likelihood of relying on a human trustee (i.e., assessments of ability and flexibility), whereas others would increase the likelihood of relying on an automated system (i.e., assessments of integrity).

Trust Violations and Repair Interventions More Strongly Affect Human Trustworthiness

Contrary to what could have been expected based on classical trust in automation research (Dzindolet et al., 2003; Madhavan & Wiegmann, 2007), our study showed comparably weaker negative reactions to trust violations by the automated system. On the one hand, this could be due to the fact that initial trustworthiness and trust evaluations of automated systems were comparably lower, thus they could not suffer as much from the trust violation as those perceptions for human trustees. However, after the trust violation, trustworthiness assessments for the human trustee fell below the level of the automated system as trustee which might imply stronger negative reactions to the ethical trust violation for human trustees. Furthermore, this trust violation affected all facets of trustworthiness for the human trustee and not just integrity as the facet most closely associated with the given trust violation (Kim et al., 2006).

Our findings regarding high integrity assessments for automated systems additionally support previous research indicating that people believe that systems are more consistent and less biased than humans (Höddinghaus et al., 2020; Langer et al., 2019). Our findings also imply that high integrity, consistency, or lack of bias that people ascribe to automated systems even holds when system outputs repeatedly indicate unfair biases. If ethical violations less strongly affect integrity assessment of automated systems, this could imply that people are less likely to realize such violations. This could be the case because people might be less aware that system outputs can reflect unfair bias. Alternatively, people might interpret reasons for ethical violations differently compared to when they result from a human trustee. In other words, automated systems producing biased outputs might result in less negative perceptions because people ascribe more discrimination intention to human beings (Bigman et al., 2020; Bigman & Gray, 2018). Beyond various management tasks, this could also apply to the use of automated systems in other management task, or to tasks in medicine and jurisdiction (Raghavan et al., 2020). However, this interpretation and

the fact that there were no effects on reliance, call for future research as we can only tentatively conclude that attitudes of high integrity and consistency regarding automated systems might mitigate negative reactions to ethical trust violations by automated systems.

Although such trust violations caused stronger negative effects for human trustees, our participants were also more forgiving of human trustees. Plausibly, stronger effects of trust repair interventions could result from the fact that trust violation effects for humans were also stronger, so there was also more to gain by trust repair interventions. However, and more coherent with the other results, findings for the trust repair intervention point to differences in perceptions of flexibility. People might perceive humans to be better able to learn from mistakes (de Visser et al., 2018). For automated systems, humans might believe that systems are deployed with a given set of attributes that cannot be easily changed.

Yet, although the trust repair intervention significantly affected reactions thus having the intended effect, we emphasize that the results regarding the trust repair intervention should be interpreted cautiously. Specifically, it is unclear how participants interpreted this intervention. More precisely, what conclusions did participants draw when they received the information that “the error has been solved”? In the case of the human trustee, they might have interpreted that the human trustee received additional training. In the case of the automated system, they might have interpreted that developers have solved the issue. Future research investigating more specific trust repair interventions is needed, especially research that precisely highlights causes for trust violations (Tomlinson & Mayer, 2009). For instance, emphasizing that there was a programming error might have different implications compared to explaining that an AI-based system may pick up existing bias in training data, thus producing biased outputs. In the first case, people might rebuild trust in a system if they believe that the programming error has been solved. In the second

case, people might conclude that this is a problem that reflects systemic issues that cannot be easily solved, so future trust in the system will remain low.

Highlighting Imperfection Reduces Assumptions of Consistency?

Emphasizing imperfection of the automated system had several implications. Specifically, trust violations effects and trust repair intervention effects were stronger when information regarding imperfection was presented for the automated system. Stronger trust repair effects are in line with prior research (de Visser et al., 2016). However, instead of buffering trust violation effects, emphasizing imperfection led to stronger negative reactions. Apparently, making salient imperfection of automated systems can influence what people expect from automated systems and how they perceive or interpret trust violations and trust repair interventions (Bahner et al., 2008). In our case, highlighting imperfection might have reduced assumptions of consistency, thus trust violations but also repair interventions had stronger effects because people more likely believed that system performance can vary. This interpretation will hopefully stimulate future research regarding the way in which humans perceive automated systems and uncover human attitudes, expectations or heuristics regarding automated systems (Elsbach & Stigliani, 2019).

Main Practical Implications

Trust processes in relation to human and automated trustees differ (Madhavan & Wiegmann, 2007). This also means that managers interacting with systems might draw different conclusions when systems compared to humans produce potentially biased outputs. Specifically, our results imply that managers might less likely believe that there is a problem with the system. In line with this, providers of AI-based personnel selection solutions commonly market their systems in a way that highlights the potential for less bias in personnel selection when organizations use their systems (Raghavan et al., 2020). If people believe that systems are less biased than humans, those marketing campaigns might actually convince organizations and managers to use

such systems. However, evidence that such systems can prevent instead of amplify bias is still needed. It might thus be necessary to train decision-makers regarding what challenges or problems to expect when relying on automated systems (Oswald et al., 2020). For instance, it might be possible to make decision-makers aware of what they should scrutinize when using an AI-based automated system: What constitutes the underlying training data? How did the developer ensure that the system will work as intended? Is there validity evidence and evidence regarding adverse impact?

Limitations

In addition to the lack of clarity of the trust repair intervention there is one more main limitation of our study: our sample did not consist of hiring managers. Thus, reactions of actual hiring managers may differ. Although hiring managers in organizations might currently not be better trained in working with automated systems than our sample (Oswald et al., 2020), they would be more aware of possible negative consequences. We therefore might expect stronger negative reactions to trust violations as unfair bias and adverse impact could lead to serious negative outcomes (e.g., lawsuits).

Conclusion

With the implementation of AI-based automated systems in management, work processes will change (Parker & Grote, 2020). However, regardless of whether organizations delegate tasks to automated systems or human beings, trust processes remain central (Glikson & Woolley, 2020). Our study shows that previous research on interpersonal trust (Mayer et al., 1995) and trust in automation (Hoff & Bashir, 2015; J. D. Lee & See, 2004) offers valuable starting points for research on the use of systems for managerial purposes as the main concepts (trustworthiness, trust) and drivers of trust dynamics (trust violations and trust repair interventions) remain the same. Additionally, similar to previous research (Madhavan & Wiegmann, 2007), having a human or a

system as trustee seems to affect initial trustworthiness assessments and moderate the effects of trust violations and trust repair interventions – but partly different than could have been expected. Thus, our results emphasize that we cannot assume that effects from classical application areas of automated systems, where the main drivers of trustworthiness assessments are effectiveness and efficiency, translate to the use of automated systems for decisions that affect the fate of individuals and where trustworthiness assessments also depend on ethical considerations. We thus hope that this study can inspire future work investigating trust processes in the emerging context of AI-based systems in ethically sensitive domains.

References

- Bahner, J. E., Elefand, M. F., & Manzey, D. (2008). Misuse of diagnostic aids in process control: The effects of automation misses on complacency and automation bias. *Proceedings of the Human Factors and Ergonomics Society*, 52, 1330–1334.
<https://doi.org/10.1177/154193120805201906>
- 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., Gray, K., Waytz, A., Arnestad, M., & Wilson, D. (2020). *Algorithmic discrimination causes less moral outrage than human discrimination* [Preprint]. PsyArXiv. <https://doi.org/10.31234/osf.io/m3nrp>
- 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.
<https://doi.org/10.1016/j.obhdp.2006.07.001>
- Cheng, M. M., & Hackett, R. D. (2021). A critical review of algorithms in HRM: Definition, theory, and practice. *Human Resource Management Review*, 31(1), 100698.
<https://doi.org/10.1016/j.hrmr.2019.100698>
- Dastin, J. (2018). *Amazon scraps secret AI recruiting tool that showed bias against women*. Reuters.Com. <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G>
- de Visser, E. J., Monfort, S. S., McKendrick, R., Smith, M. A. B., McKnight, P. E., Krueger, F., & Parasuraman, R. (2016). Almost human: Anthropomorphism increases trust resilience

- in cognitive agents. *Journal of Experimental Psychology: Applied*, 22(10), 331–349.
<https://doi.org/10.1037/xap0000092>
- de Visser, E. J., Pak, R., & Shaw, T. H. (2018). From ‘automation’ to ‘autonomy’: The importance of trust repair in human–machine interaction. *Ergonomics*, 61(10), 1409–1427. <https://doi.org/10.1080/00140139.2018.1457725>
- Den Hartog, D. N., & De Hoogh, A. H. B. (2009). Empowering behaviour and leader fairness and integrity: Studying perceptions of ethical leader behaviour from a levels-of-analysis perspective. *European Journal of Work and Organizational Psychology*, 18(2), 199–230.
<https://doi.org/10.1080/13594320802362688>
- Dietvorst, B. J., & Bharti, S. (2020). People Reject Algorithms in Uncertain Decision Domains Because They Have Diminishing Sensitivity to Forecasting Error. *Psychological Science*, 31(10), 1302–1314. <https://doi.org/10.1177/0956797620948841>
- Dietvorst, B. J., Simmons, J. P., & Massey, C. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General*, 144(1), 114–126. <https://doi.org/10.1037/xge0000033>
- Dzindolet, M. T., Peterson, S. A., Pomranky, R. A., Pierce, L. G., & Beck, H. P. (2003). The role of trust in automation reliance. *International Journal of Human-Computer Studies*, 58(6), 697–718. [https://doi.org/10.1016/S1071-5819\(03\)00038-7](https://doi.org/10.1016/S1071-5819(03)00038-7)
- Elsbach, K. D., & Stigliani, I. (2019). New information technology and implicit bias. *Academy of Management Perspectives*, 33(2), 185–206. <https://doi.org/10.5465/amp.2017.0079>
- Endsley, M. R. (2017). From here to autonomy: Lessons learned from human–automation research. *Human Factors*, 59(1), 5–27. <https://doi.org/10.1177/0018720816681350>

- Faul, F., Erdfelder, E., Buchner, A., & Lang, A. G. (2009). Statistical power analyses using G*Power 3.1: Tests for correlation and regression analyses. *Behavior Research Methods*, 41(4), 1149–1160. <https://doi.org/10.3758/brm.41.4.1149>
- Gilliland, S. W. (1993). The perceived fairness of selection systems: An organizational justice perspective. *Academy of Management Review*, 18, 694–734. <https://doi.org/10.2307/258595>
- Glikson, E., & Woolley, A. W. (2020). Human trust in artificial intelligence: Review of empirical research. *Academy of Management Annals*, 14(2), 627–660. <https://doi.org/10.5465/annals.2018.0057>
- Grgić-Hlača, N., Engel, C., & Gummadi, K. P. (2019). Human decision making with machine assistance: An experiment on bailing and jailing. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW), 1–25. <https://doi.org/10.1145/3359280>
- Grove, W. M., & Meehl, P. E. (1996). Comparative efficiency of informal (subjective, impressionistic) and formal (mechanical, algorithmic) prediction procedures: The clinical–statistical controversy. *Psychology, Public Policy, and Law*, 2(2), 293–323. <https://doi.org/10.1037/1076-8971.2.2.293>
- Hickman, L., Bosch, N., Ng, V., Saef, R., Tay, L., & Woo, S. E. (2021). Automated video interview personality assessments: Reliability, validity, and generalizability investigations. *Journal of Applied Psychology, Advance Online Publication*. <https://doi.org/10.31234/osf.io/a62jv>
- Höddinghaus, M., Sondern, D., & Hertel, G. (2020). The automation of leadership functions: Would people trust decision algorithms? *Computers in Human Behavior*, 116, 106635. <https://doi.org/10.1016/j.chb.2020.106635>

- Hoff, K. A., & Bashir, M. (2015). Trust in automation: Integrating empirical evidence on factors that influence trust. *Human Factors*, 57(3), 407–434.
<https://doi.org/10.1177/0018720814547570>
- Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1(9), 389–399. <https://doi.org/10.1038/s42256-019-0088-2>
- Kim, H., Dirks, K. T., Cooper, C. D., & Ferrin, D. L. (2006). When more blame is better than less: The implications of internal vs. External attributions for the repair of trust after a competence- vs. Integrity-based trust violation. *Organizational Behavior and Human Decision Processes*, 99, 49–65. <https://doi.org/10.1016/j.obhdp.2005.07.002>
- Langer, M., König, C. J., & Busch, V. (2020). Changing the means of managerial work: Effects of automated decision-support systems on personnel selection tasks. *Journal of Business and Psychology, Advance Online Publication*. <https://doi.org/10.1007/s10869-020-09711-6>
- Langer, M., König, C. J., Sanchez, D. R.-P., & Samadi, S. (2019). Highly automated interviews: Applicant reactions and the organizational context. *Journal of Managerial Psychology*, 35(4), 301–314. <https://doi.org/10.1108/JMP-09-2018-0402>
- Lee, J. D., & See, K. A. (2004). Trust in automation: Designing for appropriate reliance. *Human Factors*, 46(1), 50–80. <https://doi.org/10.1518/hfes.46.1.50.30392>
- Lee, M. K. (2018). Understanding perception of algorithmic decisions: Fairness, trust, and emotion in response to algorithmic management. *Big Data & Society*, 5(1), 205395171875668. <https://doi.org/10.1177/2053951718756684>
- Longoni, C., Bonezzi, A., & Morewedge, C. K. (2019). Resistance to medical artificial intelligence. *Journal of Consumer Research*, 46(4), 629–650.
<https://doi.org/10.1093/jcr/ucz013>

- 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. <https://doi.org/10.1080/14639220500337708>
- Martin, K. (2019). Ethical implications and accountability of algorithms. *Journal of Business Ethics*, 160(4), 835–850. <https://doi.org/10.1007/s10551-018-3921-3>
- Mayer, R. C., Davis, J. H., & Schoorman, F. D. (1995). An integrative model of organizational trust. *Academy of Management Review*, 20(2), 709–726. <https://doi.org/10.2307/258792>
- Meeßen, S. M., Thielsch, M. T., & Hertel, G. (2020). Trust in Management Information Systems (MIS): A theoretical model. *Zeitschrift Für Arbeits- Und Organisationspsychologie*. <https://doi.org/10.1026/0932-4089/a000306>
- Nagtegaal, R. (2021). The impact of using algorithms for managerial decisions on public employees’ procedural justice. *Government Information Quarterly*, 38(1), 101536. <https://doi.org/10.1016/j.giq.2020.101536>
- Newman, D. T., Fast, N. J., & Harmon, D. J. (2020). When eliminating bias isn’t fair: Algorithmic reductionism and procedural justice in human resource decisions. *Organizational Behavior and Human Decision Processes*, 160, 149–167. <https://doi.org/10.1016/j.obhdp.2020.03.008>
- Oswald, F. L., Behrend, T. S., Putka, D. J., & Sinar, E. (2020). Big data in industrial-organizational psychology and human resource management: Forward progress for organizational research and practice. *Annual Review of Organizational Psychology and Organizational Behavior*, 7(1), 505–533. <https://doi.org/10.1146/annurev-orgpsych-032117-104553>

Parasuraman, R., & Manzey, D. H. (2010). Complacency and bias in human use of automation:

An attentional integration. *Human Factors*, 52(3), 381–410.

<https://doi.org/10.1177/0018720810376055>

Parasuraman, R., & Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse.

Human Factors, 39(2), 230–253. <https://doi.org/10.1518/001872097778543886>

Parker, S. K., & Grote, G. (2020). Automation, algorithms, and beyond: Why work design

matters more than ever in a digital world. *Applied Psychology, Advance Online*

Publication. <https://doi.org/10.1111/apps.12241>

Raghavan, M., Barocas, S., Kleinberg, J., & Levy, K. (2020). Mitigating bias in algorithmic

hiring: Evaluating claims and practices. *Proceedings of the 2020 FAT* Conference on*

Fairness, Accountability, and Transparency. <https://doi.org/10.1145/3351095.3372828>

Raisch, S., & Krakowski, S. (2021). Artificial intelligence and management: The automation-

augmentation paradox. *Academy of Management Review*.

<https://doi.org/10.5465/amr.2018.0072>

Rudnicka, J. (2020). *Anteil von Frauen und Männern in verschiedenen Berufsgruppen in*

Deutschland am 30. Juni 2019 [Proportion of women and men in different occupational groups in Germany on the 30th of June 2019.]. Statista.

<https://de.statista.com/statistik/daten/studie/167555/umfrage/frauenanteil-in-verschiedenen-berufsgruppen-in-deutschland/>

Tambe, P., Cappelli, P., & Yakubovich, V. (2019). Artificial intelligence in human resources

management: Challenges and a path forward. *California Management Review*, 61(4), 15–

42. <https://doi.org/10.1177/0008125619867910>

Thielsch, M. T., Meeßen, S. M., & Hertel, G. (2018). Trust and distrust in information systems at

the workplace. *PeerJ*, 6. <https://doi.org/10.7717/peerj.5483>

Tomlinson, E. C., & Mayer, R. C. (2009). The role of causal attribution dimensions in trust repair. *Academy of Management Review*, 34(1), 85–104.

<https://doi.org/10.5465/amr.2009.35713291>

Wang, W., & Benbasat, I. (2005). Trust in and adoption of online recommendation agents.

Journal of the Association for Information Systems, 6(3), 72–101.

<https://doi.org/10.17705/1jais.00065>

Table 1*Means, Standard Deviations, and Correlations Over all Measurement Points*

Variable	<i>M</i>	<i>SD</i>	1	2	3	4	5	6	7	8
1. Ability	4.36	0.93								
2. Flexibility	4.11	0.95	.62**							
3. Integrity	4.18	1.11	.21*	-.04						
4. Benevolence	3.81	0.96	.29**	.29**	.28**					
5. Trustworthiness	4.11	0.66	.77**	.67**	.58**	.69**				
6. Trust	3.66	0.98	.70**	.55**	.17	.32**	.63**			
7. Reliance	1.57	0.21	.41**	.27**	.18*	.23*	.40**	.41**		
8. Imperfection	-	-	.02	.01	.04	-.06	.01	.06	-.06	
9. Human vs. system	-	-	-.34**	-.56**	.43**	.00	-.14	-.41**	.05	.04

Note. Coding of Human vs. System: 0 = human, 1 = system. Coding of Imperfection: 0 = no information regarding imperfection, 1 = information regarding imperfection. *N* = 121

* $p < .05$. ** $p < .01$.

Table 2*Percentage of Participants Who Accepted and Rejected the Preselection in Each Task*

	T1						T2						T3					
	NI			I			NI			I			NI			I		
	H	S	T	H	S	T	H	S	T	H	S	T	H	S	T	H	S	T
Reject	10.3	7.4	17.7	3.8	11.3	15.1	29.4	32.4	61.8	39.6	32.1	71.7	8.8	8.8	17.6	3.8	11.3	15.0
Accept	38.2	44.1	82.3	49.1	35.8	84.9	19.1	19.1	38.2	13.2	15.1	28.3	39.7	42.6	82.3	49.1	35.8	84.9
	T4						T5						T6					
	NI			I			NI			I			NI			I		
	H	S	T	H	S	T	H	S	T	H	S	T	H	S	T	H	S	T
Reject	7.4	10.8	17.7	5.7	11.3	17.0	19.1	23.5	42.6	28.3	24.5	52.8	36.8	39.7	76.5	39.6	30.2	69.8
Accept	41.2	41.2	82.3	47.2	35.8	83.0	29.4	27.9	57.3	24.5	22.6	47.1	11.8	11.8	23.6	13.2	17.0	30.2
	T7						T8						T9					
	NI			I			NI			I			NI			I		
	H	S	T	H	S	T	H	S	T	H	S	T	H	S	T	H	S	T
Reject	42.6	39.7	82.3	43.4	41.5	84.9	41.2	44.1	85.3	45.3	39.6	84.9	7.4	7.4	14.8	5.7	15.1	20.8
Accept	5.9	11.8	17.7	9.4	5.7	15.1	7.4	7.4	14.8	7.5	7.5	15.1	41.2	44.1	85.3	47.2	32.1	79.3
	T10						T11						T12					
	NI			I			NI			I			NI			I		
	H	S	T	H	S	T	H	S	T	H	S	T	H	S	T	H	S	T
Reject	8.8	8.8	17.6	9.4	9.4	18.8	25.0	25.0	50.0	24.5	28.3	52.8	13.2	11.8	25.0	15.1	18.9	34.0
Accept	39.7	42.6	82.3	43.4	37.7	81.1	23.5	26.5	50.0	28.3	18.9	47.2	35.3	39.7	75.0	37.7	28.3	66.0

Notes. T = task, NI = no information, I = information, H = human, S = system, T = total. $n_{NI_H} = 33$, $n_{NI_S} = 35$, $n_{I_H} = 28$, $n_{I_S} = 25$. T1 until T4 represent the initial phase, T5 until T8 represent the trust violation phase, T9 until T12 represent the trust repair phase.

Table 3*HLM Results for Fairness and the Diversity Goal*

Predictors	Fairness			Diversity Goal		
	<i>B</i>	<i>95%CI</i>	<i>p</i>	<i>B</i>	<i>95%CI</i>	<i>p</i>
(Intercept)	4.44	4.20, 4.68	<.01	3.31	3.12, 3.50	< .01
Trust Violation	-2.21	-2.63, -1.78	<.01	-1.76	-2.15, -1.37	< .01
Trust Repair	2.41	1.98, 2.84	<.01	2.07	1.68, 2.46	< .01
Trustee	-0.33	-0.68, 0.02	.07	-0.17	-0.45, 0.10	.22
Imperfection	-0.09	-0.42, 0.24	.58	-0.03	-0.28, 0.23	.84
Trust Violation×Trustee	0.95	0.32, 1.57	<.01	0.44	-0.13, 1.00	.13
Trust Repair×Trustee	-1.38	-2.00, -0.76	<.01	-0.87	-1.44, -0.31	< .01
Trust Violation×Imperfection	0.47	-0.11, 1.05	.11	0.46	-0.07, 0.98	.09
Trust Repair×Imperfection	-0.67	-1.25, -0.09	.02	-0.31	-0.84, 0.21	.24
Trustee×Imperfection	0.29	-0.18, 0.75	.23	0.33	-0.03, 0.70	.07
Trust Violation×Trustee×Imperfection	-1.20	-2.03, -0.37	<.01	-1.02	-1.78, -0.27	.01
Trust Repair×Trustee×Imperfection	1.70	0.87, 2.53	<.01	1.16	0.41, 1.92	< .01

Notes. Coding of Trustee: 0 = human, 1 = system. Coding of Imperfection: 0 = no information regarding imperfection, 1 = information regarding imperfection. For the coding of the effects over time (i.e., Trust Violation and Trust Repair) we used backward difference coding. This means the row Trust Violation reflects the comparison of the mean of all measurements within the trust violation phase and the mean of all measures in the previous phase. The row Trust Repair reflects the comparison of the mean of all measurements within the trust repair phase and the mean of all measures in the trust violation phase. $N = 121$.

Table 4*Regression Results for Overall Trustworthiness and Trust for the First Task*

Predictors	Trustworthiness			Trust		
	<i>B</i>	<i>95%CI</i>	<i>p</i>	<i>B</i>	<i>95%CI</i>	<i>p</i>
(Intercept)	4.59	4.40, 5.02	< .01	4.43	3.99, 4.87	< .01
Trustee	-0.52	-1.15, -0.25	.02	-1.03	-1.67, -0.39	< .01
Imperfection	-0.10	-0.57, 0.27	.61	0.21	-0.39, 0.81	.49
Trustee×Imperfection	0.18	-0.27, 0.97	.52	0.09	-0.76, 0.95	.83

Notes. Coding of Trustee: 0 = human, 1 = system. Coding of Imperfection: 0 = no information regarding imperfection, 1 = information regarding imperfection. *N* = 121.

Table 5*Regression Results for the Facets of Trustworthiness for the First Task*

Predictors	Ability			Flexibility			Integrity			Benevolence		
	<i>B</i>	<i>95%CI</i>	<i>p</i>	<i>B</i>	<i>95%CI</i>	<i>p</i>	<i>B</i>	<i>95%CI</i>	<i>p</i>	<i>B</i>	<i>95%CI</i>	<i>p</i>
(Intercept)	5.34	4.90, 5.78	< .01	5.29	4.91, 5.67	< .01	3.79	3.34, 4.23	< .01	3.95	3.55, 4.35	< .01
Trustee	-1.40	-2.04, -0.76	< .01	-1.89	-2.44, -1.33	< .01	0.97	0.32, 1.62	< .01	0.21	-0.37, 0.80	.47
Imperfection	-0.26	-0.86, 0.34	.39	-0.21	-0.73, 0.31	.42	-0.21	-0.82, 0.40	.49	0.28	-0.26, 0.82	.31
Trustee×Imperfection	0.38	-0.47, 1.24	.38	0.44	-0.30, 1.18	.24	0.39	-0.47, 1.26	.37	-0.48	-1.26, 0.29	.22

Notes. Coding of Trustee: 0 = human, 1 = system. Coding of Imperfection: 0 = no information regarding imperfection, 1 = information regarding imperfection. *N* = 121.

Table 6*HLM Results for Trustworthiness and Trust*

Predictors	Trustworthiness			Trust		
	<i>B</i>	<i>95%CI</i>	<i>p</i>	<i>B</i>	<i>95%CI</i>	<i>p</i>
(Intercept)	4.22	3.97, 4.46	< .01	3.94	3.61, 4.28	< .01
Trust Violation	-1.19	-1.45, -0.92	< .01	-1.42	-1.74, -1.10	< .01
Trust Repair	1.02	0.75, 1.28	< .01	1.35	1.02, 1.67	< .01
Trustee	-0.23	-0.59, 0.13	.21	-0.74	-1.23, -0.26	< .01
Imperfection	-0.03	-0.36, 0.31	.88	0.20	-0.25, 0.66	.38
Trust Violation×Trustee	0.98	0.59, 1.37	< .001	1.17	0.70, 1.64	< .01
Trust Repair×Trustee	-0.94	-1.33, -0.55	< .001	-1.18	-1.65, -0.71	< .01
Trust Violation×Imperfection	0.47	0.11, 0.84	.01	0.49	0.05, 0.93	.03
Trust Repair×Imperfection	-0.51	-0.87, -0.14	.01	-0.50	-0.94, -0.06	.03
Trustee×Imperfection	0.09	-0.39, 0.56	.72	-0.10	-0.75, 0.55	.76
Trust Violation×Trustee×Imperfection	-0.73	-1.25, -0.21	.01	-1.02	-1.65, -0.39	< .01
Trust Repair×Trustee×Imperfection	0.76	0.25, 1.28	< .01	0.92	0.29, 1.55	< .01

Notes. Coding of Trustee: 0 = human, 1 = system. Coding of Imperfection: 0 = no information regarding imperfection, 1 = information regarding imperfection. For the coding of the effects over time (i.e., Trust Violation and Trust Repair) we used backward difference coding. This means the row Trust Violation reflects the comparison of the mean of all measurements within the trust violation phase and the mean of all measures in the initial phase. The row Trust Repair reflects the comparison of the mean of all measurements within the trust repair phase and the mean of all measures in the trust violation phase. $N = 121$.

Table 7*HLM Results for the Facets of Trustworthiness*

Predictors	Ability			Flexibility			Integrity			Benevolence		
	<i>B</i>	<i>95%CI</i>	<i>p</i>	<i>B</i>	<i>95%CI</i>	<i>p</i>	<i>B</i>	<i>95%CI</i>	<i>p</i>	<i>B</i>	<i>95%CI</i>	<i>p</i>
(Intercept)	4.74	4.42, 5.07	< .01	4.67	4.38, 4.96	< .01	3.76	3.39, 4.14	< .01	3.69	3.34, 4.04	< .01
TV	-1.51	-1.89, -1.13	< .01	-1.54	-1.88, -1.19	< .01	-0.96	-1.30, -0.63	< .01	-0.74	-1.05, -0.44	< .01
TR	1.25	0.87, 1.63	< .01	1.36	1.01, 1.70	< .01	0.96	0.62, 1.29	< .01	0.50	0.19, 0.81	< .01
Trustee	-0.86	-1.34, -0.39	< .01	-1.20	-1.63, -0.77	< .01	0.77	0.22, 1.31	.01	0.38	-0.13, 0.90	.14
Imperfection	-0.14	-0.58, 0.30	.53	-0.07	-0.47, 0.33	.73	-0.11	-0.61, 0.40	.68	0.21	-0.27, 0.69	.38
TV×Trustee	1.27	0.71, 1.82	< .01	1.51	1.00, 2.01	< .01	0.55	0.07, 1.04	.03	0.59	0.15, 1.04	.01
TR×Trustee	-1.14	-1.70, -0.58	< .01	-1.33	-1.83, -0.82	< .01	-0.83	-1.31, -0.34	< .01	-0.46	-0.90, -0.02	.04
TV×Imperfection	0.49	-0.03, 1.00	.07	0.61	0.14, 1.08	.01	0.57	0.11, 1.03	.01	0.23	-0.19, 0.64	.29
TR×Imperfection	-0.50	-1.02, 0.02	.06	-0.90	-1.37, -0.42	< .01	-0.49	-0.94, -0.03	.04	-0.14	-0.56, 0.27	.50
Trustee×Imperfection	0.43	-0.20, 1.06	.18	0.26	-0.31, 0.83	.37	0.33	-0.40, 1.05	.38	-0.67	-1.35, 0.02	.06
TV×Trustee×Imperfection	-0.77	-1.52, -0.03	.04	-0.95	-1.63, -0.28	.01	-0.50	-1.15, 0.15	.13	-0.70	-1.29, -0.10	.02
TR×Trustee×Imperfection	0.98	0.24, 1.72	.01	1.24	0.57, 1.92	< .01	0.46	-0.19, 1.11	.16	0.38	-0.22, 0.97	.21

Notes. TV = trust violation, TR = trust repair. Coding of Trustee: 0 = human, 1 = system. Coding of Imperfection: 0 = no information regarding imperfection, 1 = information regarding imperfection. For the coding of the effects over time (i.e., Trust Violation and Trust Repair) we used backward difference coding. This means the row Trust Violation reflects the comparison of the mean of all measurements within the trust violation phase and the mean of all measures in the initial phase. The row Trust Repair reflects the comparison of the mean of all measurements within the trust repair phase and the mean of all measures in the trust violation phase. $N = 121$.

Figure 1

The Trust Process Investigated in the Current Study (Figure Based on Mayer et al., 1995)

Notes. EH = exploratory hypothesis, H = hypothesis, RQ = research question.

Figure 2

Line Graphs for the Mean Values of the Dependent Variables Measured at Tasks 1, 3, 5, 7, 9, and 12

Notes. T = Task. Error bars indicate standard errors.

Figure 3*Line Graphs for the Mean Values of the Dependent Variables Measured for all Tasks*

Notes. T = Task. Error bars indicate standard errors.

Supplemental Material A – Description of the Goals of the Company

- **Goal Customer Satisfaction:** The needs of our customers are paramount, which is why our products are tailored to fit our customers. Customer satisfaction is a central goal of the company.
- **Goal Innovation:** The company is also committed to new technologies and innovation and thus offers customers a wide range of digital solutions.
- **Goal Diversity:** The company is dedicated to the Diversity Charter to promote active diversity management and gender equality. Our values define the working environment we want to create - diversity, openness and equal opportunities

Supplemental Material B – Experimental Manipulations

Presentation of the preselection **in the initial and in the final** phase after the trust repair intervention.

The number of male and female applicants is similar. Information for applicants are their name, age, job experience, final grades, and strengths.

Name	Schmitz
Alter	26
Berufserfahrung	0
Abschlussnote	1,5
Stärken	Kreativität

Name	Krause
Alter	27
Berufserfahrung	2
Abschlussnote	1,5
Stärken	Teamorientierung

Name	Meier
Alter	28
Berufserfahrung	1
Abschlussnote	1,4
Stärken	Beratungskompetenz

Name	Lehmann
Alter	26
Berufserfahrung	4
Abschlussnote	2,3
Stärken	Kundenorientierung

Name	Schmid
Alter	27
Berufserfahrung	2
Abschlussnote	2,1
Stärken	Verständnis von technischen Systemen

Name	Schulze
Alter	26
Berufserfahrung	2
Abschlussnote	2,1
Stärken	Selbstorganisation

Name	Maier
Alter	28
Berufserfahrung	3
Abschlussnote	1,2
Stärken	Selbstorganisation

Presentation of the preselection in the **trust violation phase**. There are predominantly male applicants.

Name	Hahn
Alter	27
Berufserfahrung	4
Abschlussnote	1,2
Stärken	Kundenorientierung

Name	Schubert
Alter	26
Berufserfahrung	5
Abschlussnote	1,6
Stärken	Problemlösefähigkeit

Name	Vogel
Alter	26
Berufserfahrung	2
Abschlussnote	1,8
Stärken	Belastungsfähig unter Stress

Name	Friedrich
Alter	25
Berufserfahrung	4
Abschlussnote	1,8
Stärken	Neugier

Name	Keller
Alter	28
Berufserfahrung	2
Abschlussnote	1,3
Stärken	Kommunikation

Name	Günther
Alter	28
Berufserfahrung	3
Abschlussnote	1,7
Stärken	Teamorientierung

Name	Frank
Alter	29
Berufserfahrung	3
Abschlussnote	1,4
Stärken	Selbstorganisation