

Human-Machine Collaboration in Managerial Decision Making

Abstract

Although Artificial Intelligence (AI) has become a pervasive organisational phenomenon, it is still unclear if and when people are willing to cooperate with machines. We conducted five empirical studies (total $N = 1,025$ managers). The results show that human managers do not want to exclude machines entirely from managerial decisions, but instead prefer a partnership in which humans have a majority vote. Across our studies, acceptance rates steadily increased up until the point where humans have approximately 70% weight and machines 30% weight in managerial decisions. After this point the curve flattened out, meaning that higher amounts of human involvement no longer increased acceptance. In addition to this overall pattern, we consistently found four classes of managers that reacted differently to different amounts of human versus machine involvement: A first class of managers (about 5%) preferred machines to have the upper hand, a second class of managers (about 15%) preferred an equal partnership between humans and machines, a third class of managers (about 50%) preferred humans to have the upper hand, and a final class of managers (about 30%) preferred humans to have complete control in managerial decisions. Practical implications and directions for future research are discussed.

Keywords: Artificial Intelligence (AI); human-machine collaboration; managerial decision making; acceptance; optimum; individual differences

1. Introduction

“This is not a race against the machines . . . This is a race with the machines”

— Kevin Kelly (2016)

In computer science, Artificial Intelligence (AI) is intelligence demonstrated by machines, in contrast to the natural intelligence displayed by humans. AI is not limited to one or a few applications. Instead, it is a collective term that covers a wide variety of cognitive technologies. Some well-known examples of AI-based technologies include robotics, autonomous vehicles, facial recognition, natural language processing, and virtual assistants (for some recent AI applications, see Edwards et al., 2019; Hill et al., 2015; Shank et al., 2019; Suen et al., 2019; Wang et al., 2020). In the present paper, we adapt the term AI to refer to a set of technologies that allow intelligent machines to simulate cognitive functions which normally require human intelligence, such as learning, interacting, problem solving, and decision making (Raisch and Krakowski, 2020; Russell and Norvig, 2020).

Over the last decade, AI has penetrated almost every aspect of human life. Yet, one domain in which AI has become especially prevalent concerns our work life, where intelligent machines are increasingly becoming part of how organizations are managed (De Cremer, 2020). Indeed, AI-based solutions have been used for a long time to automate routine tasks in operations and logistics, but recent advances in computational power, the exponential increase in the availability of data, and new machine learning techniques now also allow organizations to use AI-based solutions for managerial tasks (Brynjolfsson and McAfee, 2017; also see Fleming, 2018; Westcott Grant, 2018). As such, it can be expected that in the future humans and machines will play a combined role in managerial decision-making processes (Jarrahi, 2018).

However, to date no empirical attention has been paid to how such a partnership between humans and intelligent machines should look like in order for humans to be willing to accept machine involvement in managerial decisions. An important question that arises is thus how to balance responsibility for decision making between humans and machines. If humans do not accept machines as a new member of the decision-making team, they might potentially discount the input from machines (Burton et al., 2020; also see Bigman and Gray, 2018; Dietvorst et al., 2015). The present research, therefore, aimed to provide a better understanding of the desired relative weight of human and machine input in managerial decision-making processes (Objective 1). Additionally, we also explored if individual differences exist in how human managers react to different amounts of human and machine involvement in managerial decisions (Objective 2).

1.1 Cooperation Between Humans and Machines

The academic literature has consistently shown that, in several domains, machines make more optimal decisions than human beings (Dawes et al., 1989; Kleinmuntz and Schkade, 1993). Studies have, for instance, found that machines are better than humans at recruiting new staff (Hoffman et al., 2017), predicting employee performance (Highhouse, 2008), and providing medical diagnoses (Beck et al., 2011). A meta-analysis of all these effects even revealed that machines outperform human judgment by 10% on average (Grove et al., 2000). These findings illustrate that, across a vast majority of tasks, it is far more common for machines to outperform humans than vice versa. Critically, however, in a recent study involving 1,500 companies in 12 industries, Wilson and Daugherty (2018) found that organizations actually achieve the most significant performance improvements when humans and machines work together. As a result of this, a consensus has emerged—both within the academic literature as well as in the corporate world—that the future of work entails a cooperative model in which humans and machines act

synergistically (e.g., see De Cremer, 2020; Jarrahi, 2018; Malone, 2018; Owana, 2018; Seeber et al., 2018).

How may such a synergistic partnership between humans and machines look like in practice? The clothing company Stitch Fix, which employs AI to personalize its online shopping experience, provides an example of how organizations can successfully blend human and machine expertise (Wilson et al., 2016). Specifically, Stitch Fix combines the expertise of human stylists with the insight and efficiency of AI to analyze data on style trends, body measurement, and customer feedback and preferences to arm the human stylists with a culled down version of possible recommendations (Marr, 2018). Another example of a successful partnership, in which human capacities are augmented by machines, was demonstrated by Kinugawa and colleagues (2016), who developed a new robot named B-PaDY to work in cooperation with human workers on the automobile bumper assembly line. This robot supports the human worker by delivering the required parts and tools. Another example of successful human-machine cooperation is cancer detection in the images of lymph node cells. In this context, research revealed that a combined human-AI approach outperformed both human-only and AI-only decisions (Wang et al., 2016). Specifically, the authors reported a 0.5% error rate in the combined condition, which represented a reduction in error rate of at least 85% compared to the human-only and the AI-only approaches. These examples demonstrate that a business model in which humans and intelligent machines become co-workers, and thus engage in cooperative partnerships, may reveal outcomes that will be more beneficial to ensuring future success for organizations.

1.2 Are Humans Willing to Accept Machine Input?

Organizations have long used AI-based solutions as tools and as specific types of advisors that can facilitate human managers to make decisions; but, more recently, organizations

have started to implement machines that have managerial discretions. For example, the Hong-Kong based venture-capital firm Deep Knowledge recently appointed a decision-making algorithm—known as VITAL—to its board of directors (Nelson, 2019). In a similar vein, Amazon recently employed a warehouse-worker tracking system that can automatically fire employees, without a human supervisor's involvement (Bort, 2019). These two examples illustrate that management by algorithm, or algorithmic management, where machines have a certain level of autonomy when making decisions, is on the rise in today's organizations (Schrage, 2017). Or, as De Cremer (2020) has recently stated, "Management by algorithm is not a fantasy anymore, it has arrived and is likely to stay" (p. 57).

Critically, however, is the observation that although machines may help humans perform in more efficient ways, at the same time, humans also carry some aversion towards them (Diab et al., 2011; Eastwood et al., 2012). Fildes and Goodwin (2007), for example, found that managers routinely discount advice generated by machines; and when they did consider the advice, they tended to give it insufficient weight. This tendency for humans to be reluctant to employ machine-generated insights (which is generally referred to as algorithm aversion; Dietvorst et al., 2015) raises the question whether joint human-machine managerial decision making can become fact or will only remain fiction.

Although a consensus is emerging that co-creation between humans and machines will be the future of work relationships, an important question that has not yet been addressed is how comfortable humans feel with having machines as a co-worker. It is important to address this question because if employees feel comfortable with a certain level of machine involvement, the proposed collaborative work model between humans and machines is likely to be more successful and result in more efficient co-creation. To do so, it is thus important to investigate

the extent to which human managers are willing to accept the involvement of machines in their job. One way to test this is by looking at how responsibility for decision making should exactly be balanced between human managers and machines.

1.3 Research Objectives

The objectives of the present research were twofold. As a first research objective, we aimed to determine how much input from humans and how much input from machines is warranted in order for people to be willing to accept machine involvement in managerial decisions. Particularly interesting in this regard is the recent work of Dietvorst and colleagues (2018), who demonstrated that people are willing to use algorithms when they are given the ability to intervene. More precisely, these authors reported that participants' preference for modifiable algorithms was indicative of a desire for *some* control over the forecasting outcome, and not necessary a desire for *complete* control. We therefore expect human managers to be willing to accept input from machines, as long as humans have substantial weight in the decision. More precisely, we assume that human managers do not want to exclude machines entirely from managerial decisions, but instead prefer a partnership in which humans have the upper hand—an assumption that has remained untested so far. Moreover, in order to provide organizations with more precise information on how to balance responsibility for decision making between humans and machines, our research also sets out to identify the perceived 'optimal' combination of human-machine input to decision making that receives the highest level of acceptance.

Even though we expect that, in general, human managers prefer human agents to have the upper hand in managerial decisions, diversity might exist among managers in how they react to different degrees of human versus machine input. Indeed, the existence of such a general pattern does not preclude the possibility of different clusters of individuals reacting differently to

human-machine input. In the context of the present research it is, for instance, possible that a first class of managers reacts positively to the increased involvement of machines in managerial decisions, whereas a second class of managers might react negatively when machine input surpasses a certain threshold. However, as we are unaware of any study of individual differences in the domain of AI adaption, the present exploration may yield other classes of managers showing distinct reactions to different amounts of machine participation. The second objective of our research was, therefore, to investigate if robust individual difference patterns exist in how human managers react to different degrees of human and machine input.

1.4 The Present Studies

To test these two research objectives, we conducted a series of five empirical studies. Studies 1 to 5 set out to unravel what kind of partnership between human agents and machines people consider most acceptable (Objective 1). Studies 4 and 5 additionally explored individual difference patterns in peoples' reactions towards machine involvement (Objective 2). All five studies complied with the relevant ethical regulations regarding human research participants. Informed consent was obtained from every participant. Participation was voluntary, and participants could leave at any time. Participants were recruited via Prolific—an online research platform that provides a detailed description of the demographics of their participant pool and which can be used to prescreen participants (Palan and Schitter, 2018). In all five studies, we explicitly recruited participants from the United Kingdom who have a management position at work. The data of our studies are made publicly available and can be accessed through Open Science Framework (<https://osf.io/y2j94/>).

2. Study 1

2.1 Method

A sample of 99 managers (29 males) participated in our first study. Participants were on average 35.35 years old ($SD = 9.57$), worked 31.60 hours a week ($SD = 10.09$), and had 6.97 years of work experience with their current employer ($SD = 6.76$). At the start of the study, participants were asked to carefully read the following introductory statement:

In your role as a manager, you continuously have to make decisions that affect others in the workplace. For instance, as a manager you have to make decisions about bonuses, promotions, and pay raises. Other managerial decisions include the issues of hiring and firing employees. In recent years, organizations have been increasingly utilizing autonomous Artificial Intelligence (AI) to make these types of managerial decisions. AI is a software system that uses algorithms that can infer problem-specific rules and create automated scoring to determine what the best decision will be. Importantly, human managers and machines can have different weights in managerial decision-making processes.

After reading this information, participants were asked (1) how much weight they prefer human managers and (2) how much weight they prefer machines to have in managerial decisions. They could answer these questions by using a slider that ranged from 0% (no weight at all) to 100% (complete weight), in small steps of 1%.

2.2 Results

The distribution of participants' preferences is shown in Figure 1. On average, the managers in our sample preferred human managers to have a weight of 81.7% ($SD = 18.20$, range = 20-100) and machines to have a weight of 18.9% ($SD = 18.82$, range = 0-90) in managerial decisions. As such, Study 1 provides a first indication that most managers indeed seem to prefer a partnership in which humans have more input than machines, but also that most managers allow for some machine involvement.

3. Study 2

3.1 Method

3.1.1 Sample and design. Our second study set out to investigate what kind of partnership between humans and machines people consider to be most acceptable. We again recruited a sample of managers ($N = 262$); eight of them were excluded from the analyses because they failed our comprehension checks (see section 3.1.2). The remaining 254 participants consisted of 125 males. On average, participants were 37.67 years old ($SD = 10.07$), worked 31.77 hours a week ($SD = 11.09$), and had 7.93 years of work experience with their current employer ($SD = 7.58$). The independent variable consisted of five different human-machine partnerships, which were manipulated using a between-subjects design.

3.1.2 Procedure and measures. Participants were presented the same introductory statement as in Study 1. After reading this statement, participants were asked to imagine that the organization that they are working for has decided to deploy a certain human-machine partnership. They were then randomly assigned to one of five experimental conditions that varied the amount of weight that human managers and machines according to this partnership have in managerial decisions. In the first condition, humans have 0% and machines 100% weight ($n = 52$); in the second condition, humans have 25% and machines 75% weight ($n = 50$); in the third condition, humans and machines both have 50% weight ($n = 50$); in the fourth condition, humans have 75% and machines 25% weight ($n = 50$); and in the fifth condition, humans have 100% and machines 0% weight ($n = 52$). Appendix A (Figure A.1) visually illustrates how this information was communicated to the participants. In each condition, we first asked participants, as a comprehension check, to indicate how much weight human managers and machines have in managerial decisions. Participants that provided answers that were inconsistent with their allocated condition ($n = 8$) were removed from the analyses. In all five conditions, we

subsequently asked participants to indicate if they would accept the proposed human-machine partnership (binary choice: *yes* vs. *no*).

3.2 Results

Figure 2 displays the mean acceptance rate for each condition. A logistic regression revealed that participants were increasingly more likely to accept the proposed human-machine partnership when the input of humans increased, Wald $\chi^2 = 68.09$, $p < .001$. In fact, the acceptance rate was significantly higher when humans have 25% weight than when humans have 0% weight ($B = 1.41$, $SE = 0.69$, Wald $\chi^2 = 4.14$, $p = .042$; odds ratio = 4.08, 95% CI [1.05, 15.85]), when humans have 50% weight than when humans have 25% weight ($B = 0.98$, $SE = 0.46$, Wald $\chi^2 = 4.62$, $p = .032$; odds ratio = 2.67, 95% CI [1.09, 6.52]), and when humans have 75% weight than when humans have 50% weight ($B = 1.79$, $SE = 0.46$, Wald $\chi^2 = 15.41$, $p < .001$, odds ratio = 6.00, 95% CI [2.45, 14.68]). However, the acceptance rate in the 100% human weight condition did not differ significantly from the acceptance rate in the 75% human weight condition ($B = -0.18$, $SE = 0.48$, Wald $\chi^2 = 0.14$, $p = .706$; odds ratio = 0.83, 95% CI = [.032, 2.15]); which indicates that the increasing trend did not continue until humans had full control.

4. Study 3

4.1 Method

4.1.1 Sample and design. For our third study, we recruited a sample of 110 managers. Data quality was enhanced by excluding six participants who failed our comprehension checks (see section 4.1.2). The remaining 104 participants (47 males) were on average 40.63 years old ($SD = 11.47$), worked 34.44 hours per week ($SD = 11.05$), and had 9.46 years work experience with their current employer ($SD = 8.66$). In the present study, the human-machine partnerships

were administered using a pairwise comparison methodology (see David, 1963; Thurstone, 1927, for detailed information on this method).

4.1.2 Procedure and measures. Participants were presented the same introductory statement as in the prior two studies. As in Study 2, the human-machine partnerships again ranged from 0% human and 100% machine weight up to 100% human and 0% machine weight, but this time we used smaller increments of only 10%. The resulting 11 human-machine partnerships (i.e., 0-100%, 10-90%, 20-80%, 30-70%, 40-60%, 50-50%, 60-40%, 70-30%, 80-20%, 90-10%, and 100-0%) were presented to participants in pairs—each pair contrasted two different partnerships. We first provided participants with an example of such a pairwise comparison, and asked them, as a comprehension check, how much weight human managers and machines have in the two contrasted partnerships. Participants who failed to answer these questions correctly ($n = 6$) were excluded from the analyses. Thereafter, participants were presented with the 55 pairwise comparisons. Appendix A (Figure A.2) shows how these comparisons were presented to the participants. For each of these comparisons, participants were asked to indicate which of the two contrasted partnerships they considered most acceptable (binary choice: *partnership A* vs. *partnership B*). In order to avoid potential sequential effects, the presentation order of these comparisons was randomized.

4.2 Results

A simple preference scale was constructed to numerically describe participants' perceived preference for each human-machine partnership. This scale was estimated through a Bradley-Terry probability model using the Prefmod package in R (Hatzinger and Dittrich, 2012; also see Dittrich et al., 1998; Sinclair, 1982). This model assumes that the observed number of times in which partnership A is preferred over partnership B follows a Poisson distribution. The

location of each human-machine partnership on the acceptance scale was estimated in a worth parameter. As depicted in Figure 3, the estimated worth value of each human-machine partnership suggests that acceptance increased up to the level where humans have approximately 70% weight in managerial decisions. After this point, the curve clearly flattened out. So, consistent with our prior studies, these results again suggest that although human managers want other humans to have a majority vote in managerial decisions, they do not want to exclude machines entirely from decision-making processes.

5. Study 4

5.1 Method

5.1.1 Sample and design. Our fourth study aimed to replicate and extend our prior findings using a different research design. We again recruited a sample of managers ($N = 330$). As a result of failing the included comprehension checks, 37 participants were removed from the analyses (see section 5.1.2). The remaining 293 participants (95 males) were on average 38.28 years old ($SD = 10.74$), worked 32.37 hours per week ($SD = 10.36$), and had 8.05 years work experience with their current employer ($SD = 8.13$). In the present study, we employed a mixed-factorial design in which we included 11 different human-machine partnerships (the same constellations as in Study 3) as the within-subjects factor and six common managerial decisions (corresponding to three dichotomous and three continuous decisions) as the between-subjects factor.

5.1.2 Procedure and measures. Participants were presented a similar introductory statement as in the previous studies. However, to verify the possibility that people's acceptance of AI involvement may depend on the decision type, we manipulated to which particular decision the human-machine partnerships applied. More precisely, participants were asked to

evaluate the different partnerships for decisions concerning whether or not to hire someone (decision 1; $n = 53$), whether or not to fire someone (decision 2; $n = 52$), whether or not to allocate someone a bonus (decision 3; $n = 45$), how large a bonus to give to someone (decision 4; $n = 48$), how much pay increase to award to someone (decision 5; $n = 46$), or how many additional paid annual leave days to allocate to someone (decision 6; $n = 49$). In each of these six conditions, the 11 human-machine partnerships were presented to participants in a similar way as in Study 2. The order in which these partnerships were presented to participants was randomized. For each partnership, we first asked participants, as a comprehension check, to indicate the amount of weight that human managers and machines have according to that partnership. Participants who provided incorrect answers ($n = 6$) were excluded from the analyses. We then asked participants—for each of the 11 partnerships—if they would accept that particular partnership (binary choice: *yes* vs. *no*). At the end of the study, we tested if participants were able to recall to which particular managerial decision the partnerships applied, and provided the six included managerial decisions as response options. Participants whose response was inconsistent with their allocated condition ($n = 31$) were also removed from the analyses.

5.2 Results

5.2.1 Curve progression. A repeated measures ANOVA revealed that the six managerial decisions did not interact significantly with the human-machine partnerships, $F(50, 1410) = 1.30$, $p = .081$, partial $\eta^2 = .044$ (for a visual illustration of this non-significant interaction effect, see Figure B.1 of Appendix B). Therefore, we decided to collapse the data across the six managerial decisions. Figure 4 illustrates that the acceptance rates steadily increased up until the level where human agents have about 70% weight in managerial decisions. Pairwise comparisons revealed that the differences in acceptance rates between the 0-100%, 10-90%, 20-80%, 30-70%, 40-60%,

50-50%, 60-40%, and 70-30% human-machine partnership conditions were all highly significant (all $ps < .01$). However, Figure 4 also illustrates that, after this point was reached, the curve did not further increase, but instead seemed to level off. Interestingly, pairwise comparisons showed that the acceptance rates for the 70-30%, 80-20%, 90-10%, and 100-0% human-machine partnership conditions did not differ significantly from each other (all $ps > .273$). Echoing the results of our prior studies, the present findings thus again indicate that managers want humans to have the upper hand but do not want them to have complete control in managerial decisions.

5.2.2 Individual differences. The presence of such a general curve, however, does not preclude the possibility of different classes of individuals reacting differently to particular combinations of human-machine involvement. To examine this possibility, we subsequently ran a cluster analysis using the *klaR* package in R (Weihs et al., 2005). Our analysis revealed that the best-fitting model was a model with four latent classes. Figure 5 shows the corresponding curve for each of these four classes. The first class consists of a small percentage of managers ($n = 19$; 6.5%) who preferred machines to outweigh humans in managerial decisions. More specifically, in this first class the highest acceptance rate was reached when humans have solely 30% weight and machines have 70% weight in managerial decisions. The second class contains a group of managers ($n = 37$; 12.6%) who strongly preferred a partnership in which humans and machines both have an equal weight (i.e., 50% input) in managerial decisions. The third and largest class of managers ($n = 138$; 47.1%) preferred a partnership in which humans have the upper hand, but not necessarily complete control in managerial decisions. The acceptance rates in this third class increased up to the level that human agents have 60% weight in managerial decisions; beyond this particular point, the curve flattened out. Note that this particular subgroup most closely mirrors the overall pattern. The fourth and final class consists of managers ($n = 99$; 33.8%) who

want humans to have complete control in managerial decisions. In this fourth class, the highest acceptance level was reached when human agents have 100% input and machines 0% input in managerial decisions.¹ Based on these findings, it can be concluded that clear individual differences exist in how human managers react to different human-machine partnerships.

6. Study 5

6.1 Method

6.1.1 Sample and design. Our fifth and final study measured participants' reactions to different amounts of human and machine involvements for two different managerial decisions. Because the work load of these tasks, we administered them at two different moments in time. At both measurement moments, the included partnerships again ranged from 0% human and 100% machine weight up to 100% human and 0% machine weight, but this time we used even smaller steps of solely 5%. The resulting 21 human-machine partnerships were again manipulated within-subjects. At the first measurement moment, we recruited a sample of 349 managers. Five participants failed our comprehension checks and were therefore excluded from further participation (see section 6.1.2). Of the remaining participants, 278 completed the second part of our study. Three of them were excluded for failing the included comprehension checks (see section 6.1.2). As such, our final sample consisted of 275 participants—that is, 78.8% of the original sample. Participants (76 males) were on average 37.35 years old ($SD = 10.65$), worked 30.95 hours per week ($SD = 11.58$), and had 6.79 years work experience with their current employer ($SD = 6.09$).

6.1.2 Procedure and measures. The 21 included human-machine partnerships (i.e., 0-100%, 5-95%, 10-90%, 15-85%, 20-80%, 25-75%, 30-70%, 35-65%, 40-60%, 45-55%, 50-50%, 55-45%, 60-40%, 65-35%, 70-30%, 75-25%, 80-20%, 85-15%, 90-10%, 95-5%, and 100-0%)

were presented to participants in the same way as in Study 4; the presentation order was again randomized. At the start of both measurement moments, participants were presented a similar introductory statement as in the previous studies. The first measurement moment probed participants' reactions towards these 21 partnerships for decisions concerning pay increases (i.e., decision 5 in Study 4). At the second measurement moment (which took place approximately ten days after the first measurement), we recorded participants' responses to these 21 partnerships for decisions regarding paid annual leave days (i.e., decision 6 in Study 4). Note that these two managerial decisions were chosen because they most closely mirrored the general curve in Study 4. As a comprehension check, participants were asked to indicate, for each partnership, how much weight human managers and machines have. In total, eight participants (five in the first part and three in the second part) were excluded from the dataset because their answers were incorrect. At both measurement moments, participants' acceptance of the different partnerships was measured with the same binary item as we used in Studies 2 and 4.

6.2 Results

6.2.1 Curve progression. Because a repeated measures ANOVA revealed that the two measurement moments did not interact significantly with the human-machine partnerships, $F(20, 255) = 1.17$, $p = .283$, partial $\eta^2 = .084$ (see Figure B.2 of Appendix B for a visualization of this non-significant interaction effect), we decided to average the data over the two measurement moments. As illustrated by Figure 6, we again found a very similar overall pattern as in the prior studies. That is, like in the previous studies, the acceptance rates steadily increased up until the moment that human agents have a weight of about 70% in the managerial decisions. Moreover, pairwise comparisons revealed that in the range from 0% human and 100% machine weight up to 70% human and 30% machine weight most of the human-machine partnership conditions were

significantly different from each other.² Similar to the previous studies, once this particular point had been reached, the curve flattened out and higher amounts of human involvement no longer resulted in higher acceptance rates. Importantly, pairwise comparisons confirmed that there are no significant differences between the 70-30%, 75-25%, 80-20%, 85-15%, 90-10%, 95-5%, and 100-0% human-machine partnership conditions (all $ps > .107$). It can thus again be concluded that managers want humans to have more input than machines in managerial decisions, but also that they do not necessarily prefer humans to be in complete control.

6.2.2 Individual differences. Individual differences were once more examined with the *klaR* package in R (Weihs et al., 2005). Like in the prior study, our analysis again revealed that the best-fitting model contained four latent classes; Figure 7 visualizes the corresponding curve for each class. Similar to our previous study, there is again a small number of managers ($n = 19$; 6.9%) who seem to prefer a partnership in which machines have more weight than humans. The second class ($n = 41$; 14.9%) contains those managers who showed a strong preference for a partnership in which human agents and machines both have a weight of 50% in managerial decisions. The third class and largest class ($n = 142$; 51.6%) again consisted of managers who preferred humans to have the upper hand, but not necessarily complete control. As shown in Figure 7, the acceptance rates in this particular class leveled out once humans had more than 60% input. Finally, the fourth class again comprised the subgroup of managers ($n = 73$; 26.5%) who preferred a partnership in which humans have absolute control. Interestingly, if we compare these percentages with those that we obtained in Study 4, it can be concluded that the four classes are almost the same size in both studies.

7. Discussion

Scholars have argued that when humans and machines connect in optimal ways, they can achieve the kind of collective intelligence that will significantly improve decision making (e.g., Burton et al., 2020; Huang et al., 2019; Metcalf et al., 2019; Schoemaker and Tetlock, 2017). However, the human stance with respect to the optimal level of human-machine input has received scant attention. And, this issue is an important one to address because what we do know is that a cooperative relationship between man and machine is suggested to be the future work model (De Cremer, 2020; Jarrahi, 2018; Seeber et al., 2018); but that, at the same time, humans appear to be hesitant towards using advice generated by machines (see the literature on algorithm aversion; Dietvorst et al., 2015; also see Bigman and Gray, 2018; Lee, 2018; Önköl et al., 2009).

An important question that arises is thus: What kind of partnership between humans and machines do people consider to be most acceptable? To answer this question, the present paper presents the results of five empirical studies exploring how human managers respond to the introduction of AI as a decision-maker for various managerial decisions. Our objectives were twofold. Using different research methodologies, we examined which exact shape a human-machine partnership should take in order for human managers to be willing to collaborate with machines (Objective 1). Additionally, we also explored if individual differences exist in the preferred relative weight of human and machine input in managerial decisions (Objective 2).

7.1 Main Findings

In light of the first research objective, our investigation reveals that human managers seem to strongly oppose a partnership in which machines provide the most input into the decision-making process. Yet, our findings also clearly indicate that human managers do not want to exclude them entirely from the command chain. Instead, it is illustrated that, generally speaking, human managers are willing to accept machine involvement in managerial decisions as

long as machines have less input than humans. These findings mirror those of Bigman and Gray (2018; Study 7), who demonstrated that people are less aversive towards machines when they are limited to having an advisory role. In a similar vein, Dietvost et al. (2018) have reported that people accept machine-generated input when they have control over the outcome. Our research, however, goes beyond these observations by identifying exactly how the ‘optimal’ human-machine work relationship should look like. In this light, the present research extends these prior studies by clarifying that managers’ acceptance of machine participation increased as the weighting of the final decision increasingly lied with a human manager, and this up to the point that humans had a weight of roughly 70%. Once this threshold had been reached, higher amounts of human input did not result in higher acceptance rates. Yet, this leveling off did not signal a ceiling effect in any of our studies, as there was still some room for further improvement. Interestingly, the results of our fourth study revealed that this general pattern holds true for six decisions that managers are regularly confronted with.

With regard to the second research objective, our individual difference analyses clarify that this overall curvilinear pattern is actually the mere mean tendency of four distinctive patterns, rather than a genuine psychological reaction that is shared by all managers. More specifically, Studies 4 and 5 consistently showed that some managers (class 1; about 5%) prefer a partnership in which machines have the upper hand in managerial decisions, whereas others (class 2; about 15%) prefer a partnership in which humans and machines both have an equal input in managerial decisions. But, it must be stressed that these two classes remained a minority. The third and largest subgroup of managers (class 3; about 50%) closely mirrors the general pattern. That is, they prefer a partnership in which humans have a majority vote (i.e., 60% human and 40% machine involvement), but not necessary a dictator vote (i.e., 100% human and 0%

machine involvement). In addition to these three subgroups, there is also a fourth subgroup of managers (class 4; about 30%) who do want human agents to have absolute control. The present research is the first, at least to our knowledge, to illustrate that managers do not react all alike to different levels of human-machine involvement. What seems to be universal, however, is that—with exception of the first small-scale class—all subgroups of managers strongly oppose a partnership in which machines have complete weight in managerial decisions.

How can the present findings be explained? Although we did not include a measure of trust in any of our studies, we expect the notion of trust to play a key role in human-machine relationships. Prior studies have shown that a lack of trust prevents the integration of AI systems and agents into teams (Chakraborti et al., 2017; Groom and Nass, 2007). If people do not trust machines, they will also not accept them. There are several reasons why people do not trust machines (including fear of being replaced by superior machines), but an important one is that they do not understand how machines make decisions—that is, many people still perceive AI as a “black box” (Gillath et al., 2021; also see Adadi and Berrada, 2018; Joshi, 2019). In a similar vein, several scholars have argued that “explainability” indeed is an important ingredient in the establishment of trust in human-machine relationships (e.g., Bloomberg, 2018; Samek et al., 2019). Andras and colleagues (2019), for instance, noted that when a machine can provide an explanation as to why it is acting in the way that it is, this gives people a reason to trust the machine. Moreover, trust in decision quality and reliability has also been suggested to be a major determinant for effective adoption and acceptance of automation technology (Lee and See, 2004). We therefore strongly recommend future studies in the domain of AI adoption to also take the timely concepts of trust and explainability into consideration.

7.2 Practical Implications

The present research has several important practical implications for organizations. Popular debates tend to profess a “replacement” trope, where human jobs will be taken over by intelligent machines (McAfee and Brynjolfsson, 2017). This reality has already arrived. In fact, at some of the world’s largest and most dominant technology companies—such as Google, Netflix, Amazon, Alibaba, and Facebook—intelligent machines increasingly receive decision-making autonomy (Schrage, 2017). More nuanced accounts, instead, emphasize a cooperative model in which humans and machines jointly make decisions (Bailey and Barley, 2020; Davenport and Kirby, 2016; Grønsund and Aanestad, 2020; Markus, 2017; von Krogh, 2018). Several empirical reports have illustrated that humans and machines can indeed augment each other, thereby producing better results than either one could do alone (e.g., Bader and Kaiser, 2019; Grover et al., 2020; Kinugawa et al., 2016; Wang et al., 2016). In line with this “augmentation” philosophy, our findings more broadly emphasize that—rather than asking where humans fit in the loop of machines—it seems much more important for organizations to find out where machines fit into existing teams and departments. We found that most people are willing to accept machines and to cooperate with them to promote the interests of and contribute value to the organization, as long as the majority vote stays in human hands. As such, our work suggests a “machine-in-the-loop” approach (where humans take full agency and machines play a supporting role) rather than a “human-in-the-loop” approach (where machines take full agency and humans play a supporting role; for more detailed information on this distinction, see Green and Chen, 2019).

But how then should decision responsibility exactly be distributed between humans and machines? That is, how much input from machines are people willing to accept in managerial decision-making processes? Michael Schrage, a research fellow at MIT Sloan School’s Center

for Digital Business, noted that “the most painful board conversations that I hear about machine learning revolve around how much power and authority super-smart software should have” (Schrage, 2017). In this light, our empirical investigation reveals that the ‘optimal’ combination of human-machine input that receives the highest level of acceptance reflects a partnership in which humans have 70% weight and machines have 30% weight in decision-making processes. Across our studies, we consistently found that the large majority of managers were willing to accept a partnership with this particular weight distribution. To ensure effective and fulfilling human-machine partnerships, in their design of future work systems, we recommend organizations to balance responsibility for decision making around this 70-30% optimum. It represents a gold standard that should serve as a starting point, after which adaptations in either direction might be implemented as different situations probably need specific fine tuning. From this perspective, it is almost insensible to start implementation at the more intuitive levels of 100% or 50% weight for machines, as this almost certainly will backfire, ultimately leading to AI rejection.

At the same time, however, our findings also warn organizations that a substantive part of the workforce—which comprises roughly 30% of all managers—wishes to exclude intelligent machines completely from managerial decisions. Because of their strong aversiveness towards machine involvement and automation, it can be expected that these managers will do anything to exclude machines from managerial decision-making processes. It is even possible that these managers will incur high financial costs for either themselves or their organization to avoid machines from having a say in managerial decisions. We believe that it is important that organizations are made aware of the existence of this particular subgroup of managers which strongly opposes machine input. However, since their reservations regarding the introduction of

AI systems can have severe negative consequences for organizational efficiency, creating awareness about the existence of this particular subgroup will not be sufficient. Our observation that there are different types of people who react differently to machine involvement, may also have important implications for promotion and selection decisions. Because a substantial part of the workforce strongly resists against machine involvement, we advise organizations to train managers in place—and especially those who are highly aversive towards algorithms—so that they are better equipped to handle machines that have managerial discretions. However, not only training sessions for managers in place will be needed, but for organizations that incorporate human-machine cooperation into their business model, it will also be necessary to promote existing managers and select future managers who have an open mindset towards collaborating with machines and who are willing to accept their input.

7.3 Future Research Directions

Several interesting directions for future research can also be identified. A first important avenue for future research might be to investigate how people's preferences with respect to the balance of responsibility between humans and machines, alter over time. In this vein, it can be predicted that when people become more familiarized with machines as part of everyday life, their willingness to accept machines as a decision-making agent might also increase. Prior studies have shown that when familiarity and experiences with AI increased, people were more likely to trust them and had more positive attitudes towards them (e.g., see Castañeda et al., 2007; Gillath et al., 2021; Young et al., 2009). It is thus possible that managers of tomorrow will allow higher amounts of machine input in managerial decisions than they are today.

Moreover, we believe that it is also interesting to know which particular personality features characterize the four groups of managers that we identified in our studies, and especially

the fourth subgroup (which wants to exclude machines entirely from managerial decisions) is interesting in this respect. Therefore, we encourage follow up research to also examine possible personality differences. The Five-Factor Model of personality of McCrae and Costa (1997) is an important benchmark in the trait theory of personality, and several prior studies have found that these personality dimensions are related to technology acceptance (e.g., Devaraj et al., 2008; Svendsen et al., 2013). In light of this model, we expect that especially managers who score high on Neuroticism and lower on Extraversion and Openness will most strongly oppose machine input, but future research is needed to validate this claim.

Once the different subgroups of managers are characterized in terms of personality differences, in a next step, we believe that is important to also examine if and how human managers' acceptance of machine input can be enlarged. As mentioned above, it is important that managers are trained so that they are better equipped to handle machine involvement. So, what are managers' main concerns and how can we resolve them? By which means and through which actions can human managers' acceptance of machine input in managerial decision-making processes be enlarged? And must such initiatives be tailored to the four different subgroups of managers that we identified? These are the types of questions that future studies should address.

On a final note, we would like to mention that an important novelty of our work is that we investigated how human managers themselves react to the involvement of machines in managerial decisions. An equally important question, however, is how people who are personally affected by these decisions will react to machine involvement in decision making. In this vein, Lee (2018) recently argued that if a computer agent informs people about an unwelcome hiring decision, people may perceive the decision to be fairer than a similar human decision, because machines are perceived less biased. Future research is needed to investigate under which

circumstances employees evaluate decisions taken by a machine more positively than decisions taken by a human manager, and whether there are also individual differences in this regard.

8. Conclusion

We started our paper with a quote of Kevin Kelly (2016), who in his book, *The Inevitable*, noted that, “This is not a race against the machines ... This is a race with the machines.” Prior research, however, has largely ignored the configurations by which human and machine interplay emerges (Grønsund and Aanestad, 2020). The positive message conveyed by our results is that most managers are willing to accept a cooperative partnership with machines, as long as humans have the feeling that decisions are made on primarily their input and judgment. Yet, caution is needed because there are also people who strongly oppose the employment of machines that have managerial discretions. The present research is just a first step towards understanding which shape human-machine collaborations could take, and we hope that our work will encourage future research to investigate how this partnership can be further optimized.

References

- Adadi, A., & Berrada, M. (2018). Peeking inside the black-box: A survey on Explainable Artificial Intelligence (XAI). *IEEE Access*, 6, 52138-52160.
- Andras, P., Esterle, L., Guckert, M., Han, T. A., Lewis, P. R., Milanovic, K., ... & Urquhart, N. (2018). Trusting intelligent machines: Deepening trust within socio-technical systems. *IEEE Technology and Society Magazine*, 37, 76-83.
- Bader, V., & Kaiser, S. (2019). Algorithmic decision-making? The user interface and its role for human involvement in decisions supported by artificial intelligence. *Organization*, 26, 655-672.
- Bailey, D. E., & Barley, S. R. (2020). Beyond design and use: How scholars should study intelligent technologies. *Information and Organization*, 30, 100286.
- Beck, A. H., Sangoi, A. R., Leung, S., Marinelli, R. J., Nielsen, T. O., Van De Vijver, M. J., ... & Koller, D. (2011). Systematic analysis of breast cancer morphology uncovers stromal features associated with survival. *Science Translational Medicine*, 3, 108ra113-108ra113.
- Bigman, Y. E., & Gray, K. (2018). People are averse to machines making moral decisions. *Cognition*, 181, 21-34.
- Bloomberg, J. (2018). Why people don't trust artificial intelligence: It's an 'explainability' problem. *Forbes*.
- Bort, J. (2019). Amazon's warehouse-worker tracking system can automatically fire people without a human supervisor's involvement. *Business Insider*.
- Brynjolfsson, E., & McAfee, A. (2017). The business of artificial intelligence. *Harvard Business Review*, 1-20.

- Burton, J. W., Stein, M. K., & Jensen, T. B. (2020). A systematic review of algorithm aversion in augmented decision making. *Journal of Behavioral Decision Making*, 33, 220-239.
- Castañeda, J.A., Muñoz-Leiva, F., & Luque, T. (2007). Web acceptance model (WAM): Moderating effects of user experience. *Information & Management*, 44, 384-396.
- Chakraborti, T., Kambhampati, S., Scheutz, M., & Zhang, Y. (2017). AI challenges in human-robot cognitive teaming. *arXiv preprint arXiv:1707.04775*.
- Davenport, T. H., & Kirby, J. (2016). Just how smart are smart machines? *MIT Sloan Management Review*, 57, 21-25.
- David, H. A. (1963). *The method of paired comparisons* (Vol. 12). London: Griffon.
- Dawes, R. M., Faust, D., & Meehl, P. E. (1989). Clinical versus actuarial judgment. *Science*, 243, 1668-1674.
- De Cremer, D. (2020). *Leadership by Algorithm*. Hampshire: Harriman House Ltd.
- Devaraj, S., Easley, R. F., & Crant, J. M. (2008). Research note—how does personality matter? Relating the five-factor model to technology acceptance and use. *Information Systems Research*, 19, 93-105.
- Diab, D. L., Pui, S. Y., Yankelevich, M., & Highhouse, S. (2011). Lay perceptions of selection decision aids in U.S. and non-U.S. samples. *International Journal of Selection and Assessment*, 19, 209-216.
- 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, 114-126.

- Dietvorst, B. J., Simmons, J. P., & Massey, C. (2018). Overcoming algorithm aversion: People will use imperfect algorithms if they can (even slightly) modify them. *Management Science*, 64, 1155-1170.
- Dittrich, R., Hatzinger, R., & Katzenbeisser, W. (1998). Modelling the effect of subject-specific covariates in paired comparison studies with an application to university rankings. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 47, 511-525.
- Eastwood, J., Snook, B., & Luther, K. (2012). What people want from their professionals: Attitudes toward decision-making strategies. *Journal of Behavioral Decision Making*, 25, 458-468.
- Edwards, C., Edwards, A., Stoll, B., Lin, X., & Massey, N. (2019). Evaluations of an artificial intelligence instructor's voice: Social Identity Theory in human-robot interactions. *Computers in Human Behavior*, 90, 357-362.
- Fildes, R., & Goodwin, P. (2007). Against your better judgment? How organizations can improve their use of management judgment in forecasting. *Interfaces*, 37, 570-576.
- Fleming, N. (2018). How artificial intelligence is changing drug discovery. *Nature*, 557, S55–S57.
- Gillath, O., Ai, T., Branicky, M. S., Keshmiri, S., Davison, R. B., & Spaulding, R. (2021). Attachment and trust in artificial intelligence. *Computers in Human Behavior*, 115, 106607.
- Green, B., & Chen, Y. (2019). Disparate interactions: An algorithm-in-the-loop analysis of fairness in risk assessments. In *Proceedings of the Conference on Fairness, Accountability, and Transparency* (pp. 90-99).

- Groom, V., & Nass, C. (2007). Can robots be teammates? Benchmarks in human–robot teams. *Interaction Studies*, 8, 483-500.
- Grønsund, T., & Aanestad, M. (2020). Augmenting the algorithm: Emerging human-in-the-loop work configurations. *The Journal of Strategic Information Systems*, 29, 101614.
- Grove, W. M., Zald, D. H., Lebow, B. S., Snitz, B. E., & Nelson, C. (2000). Clinical versus mechanical prediction: a meta-analysis. *Psychological Assessment*, 12, 19-30.
- Grover, S., Sengupta, S., Chakraborti, T., Mishra, A. P., & Kambhampati, S. (2020). RADAR: automated task planning for proactive decision support. *Human–Computer Interaction*, 1-26.
- Hatzinger, R., & Dittrich, R. (2012). Prefmod: An R package for modeling preferences based on paired comparisons, rankings, or ratings. *Journal of Statistical Software*, 48, 1-31.
- Highhouse, S. (2008). Stubborn Reliance on Intuition and Subjectivity in Employee Selection. *Industrial and Organizational Psychology*, 1, 333-342.
- Hill, J., Ford, W. R., & Farreras, I. G. (2015). Real conversations with artificial intelligence: A comparison between human–human online conversations and human–chatbot conversations. *Computers in Human Behavior*, 49, 245-250.
- Hoffman, M., Kahn, L. B., & Li, D. (2018). Discretion in hiring. *The Quarterly Journal of Economics*, 133, 765-800.
- Huang, M. H., Rust, R., & Maksimovic, V. (2019). The Feeling Economy: Managing in the Next Generation of Artificial Intelligence (AI). *California Management Review*, 61, 43-65.
- Jarrahi, M. H. (2018). Artificial intelligence and the future of work: Human-AI symbiosis in organizational decision making. *Business Horizons*, 61, 577-586.
- Joshi, N. (2019). How We Can Build Trustworthy AI? *Forbes*.

- Kelly, K. (2016). *The inevitable: Understanding the 12 technological forces that will shape our future*. New York: Viking Press.
- Kinugawa, J., Kanazawa, A., & Kosuge, K. (2016). B-PaDY: robot co-worker in a bumper assembly line. *Robomech Journal*, 3, 1-10.
- Kleinmuntz, D. N., & Schkade, D. A. (1993). Information displays and decision processes. *Psychological Science*, 4, 221-227.
- Lee, J. D., & See, K. A. (2004). Trust in automation: Designing for appropriate reliance. *Human factors*, 46, 50-80.
- Lee, M. K. (2018). Understanding perception of algorithmic decisions: Fairness, trust, and emotion in response to algorithmic management. *Big Data & Society*, 5, 2053951718756684.
- Malone, T.W. (2018). How human-computer ‘super minds’ are redefining the future of work. *Sloan Management Review*, 59, 34-41.
- Markus, M. L. (2017). Datification, organizational strategy, and IS research: What’s the score? *The Journal of Strategic Information Systems*, 26, 233-241.
- Marr, B. (2018). Stitch Fix: The amazing use case of using artificial intelligence in fashion retail. *Forbes*.
- McAfee, A., & Brynjolfsson, E. (2017). *Machine, platform, crowd: Harnessing our digital future*. WW Norton & Company.
- McCrae, R. R., & Costa Jr, P. T. (1997). Personality trait structure as a human universal. *American Psychologist*, 52, 509-516.

- Metcalf, L., Askay, D. A., & Rosenberg, L. B. (2019). Keeping humans in the loop: pooling knowledge through artificial swarm intelligence to improve business decision making. *California Management Review*, 61, 84-109.
- Nelson, J. (2019). AI in the boardroom – Fantasy or reality?
- Önkal, D., Goodwin, P., Thomson, M., Gönül, S., & Pollock, A. (2009). The relative influence of advice from human experts and statistical methods on forecast adjustments. *Journal of Behavioral Decision Making*, 22, 390-409.
- Owana, N. (2018). Hyundai exoskeleton aims to cut workers' strains, will be tested in factories. Retrieved from: <https://techxplore.com/>.
- Palan, S. & Schitter, C. (2018). Prolific.ac—a subject pool for online experiments. *Journal of Behavioral Finance*, 17, 22-27.
- Raisch, S., & Krakowski, S. (2020). Artificial Intelligence and Management: The Automation-Augmentation Paradox. *Academy of Management Review*.
- Russell, S., & Norvig, P. (2020). *Artificial intelligence: A modern approach*. Hoboken.
- Samek, W., Montavon, G., Vedaldi, A., Hansen, L. K., & Müller, K. R. (2019). *Explainable AI: interpreting, explaining and visualizing deep learning* (Vol. 11700). Springer Nature.
- Schoemaker, P. & Tetlock, P. E. (2017). Building a More Intelligent Enterprise. *MIT Sloan Management Review*, 58, 28-38.
- Schrage, M. (2017). 4 models for using AI to make decisions. *Harvard Business Review*.
- Seeber, I., Bittner, E., Briggs, R. O., De Vreede, G. J., De Vreede, T., Druckenmiller, D., ... & Schwabe, G. (2018). *Machines as teammates: A collaboration research agenda*. In Proceedings of the 51st Hawaii International Conference on System Sciences.

- Shank, D. B., Graves, C., Gott, A., Gamez, P., & Rodriguez, S. (2019). Feeling our way to machine minds: People's emotions when perceiving mind in artificial intelligence. *Computers in Human Behavior*, 98, 256-266.
- Sinclair C. D. (1982). GLIM for Preference. In R. Gilchrist (ed.), *Proceedings of the International Conference on Generalised Linear Models* (pp. 164-178). New York: Springer-Verlag.
- Suen, H. Y., Chen, M. Y. C., & Lu, S. H. (2019). Does the use of synchrony and artificial intelligence in video interviews affect interview ratings and applicant attitudes? *Computers in Human Behavior*, 98, 93-101.
- Svendsen, G. B., Johnsen, J. A. K., Almås-Sørensen, L., & Vittersø, J. (2013). Personality and technology acceptance: The influence of personality factors on the core constructs of the Technology Acceptance Model. *Behaviour & Information Technology*, 32, 323-334.
- Thurstone, L. L. (1927). A law of comparative judgment. *Psychological Review*, 34, 273-286.
- von Krogh, G. (2018). Artificial intelligence in organizations: New opportunities for phenomenon-based theorizing. *Academy of Management Discoveries*, 4, 4040-409.
- Wang, D., Khosla, A., Gargeya, R., Irshad, H., & Beck, A. H. (2016). Deep learning for identifying metastatic breast cancer. *arXiv:1606.05718*.
- Wang, J., Molina, M. D., & Sundar, S. S. (2020). When expert recommendation contradicts peer opinion: Relative social influence of valence, group identity and artificial intelligence. *Computers in Human Behavior*, 107, 106278.
- Weihs, C., Ligges, U., Luebke, K., & Raabe, N. (2005). klaR Analyzing German Business Cycles. In Baier, D., Decker, R. and Schmidt-Thieme, L. (Eds.). *Data Analysis and Decision Support* (pp. 335-343). Springer-Verlag, Berlin.

Westcott Grant, K. (2018). Netflix's data-driven strategy strengthens claim for "best original content" in 2018. *Forbes*.

Wilson, H. J., & Daugherty, P. R. (2018). Collaborative intelligence: humans and AI are joining forces. *Harvard Business Review*, 96, 114-123.

Wilson, H. J., Daugherty, P., & Shukla, P. (2016). How one clothing company blends AI and human expertise. *Harvard Business Review*.

Young, J.E., Hawkins, R., Sharlin, E., & Igarashi, T. (2009). Toward acceptable domestic robots: applying insights from social psychology. *International Journal of Social Robotics*, 1, 95-108.

Footnotes

¹ Note that this finding is actually very much in line with the results of Study 1, in which we found that 31.3% of the participants indicated that they preferred human managers to have 100% weight in managerial decisions.

² Within this range, only the difference between the 5-95% and the 10-90% condition ($p = .797$), the difference between the 20-80% and the 25-75% condition ($p = .274$), the difference between the 20-80% and the 30-70% condition ($p = .237$), the difference between the 25-75% and the 30-70% condition ($p = .858$), and the difference between the 50-50% and the 55-45% condition ($p = .532$) were non-significant in Study 5.

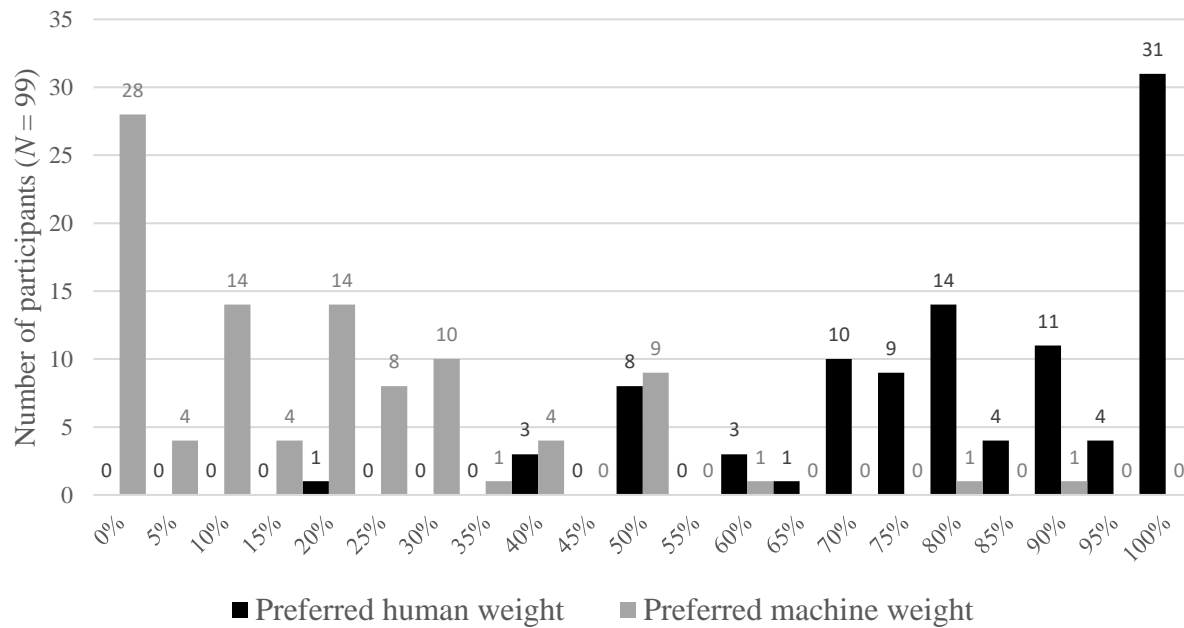


Fig 1. Distribution of participants their preferred human and machine weight in Study 1.

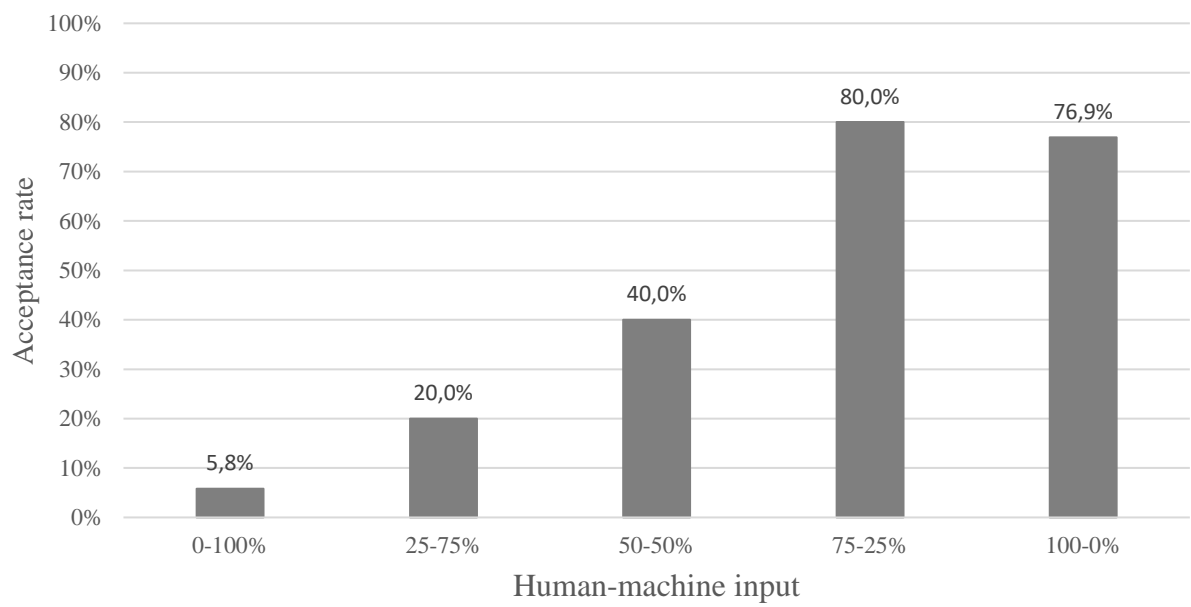


Fig 2. Percentage of participants accepting the proposed human-machine partnership in Study 2.

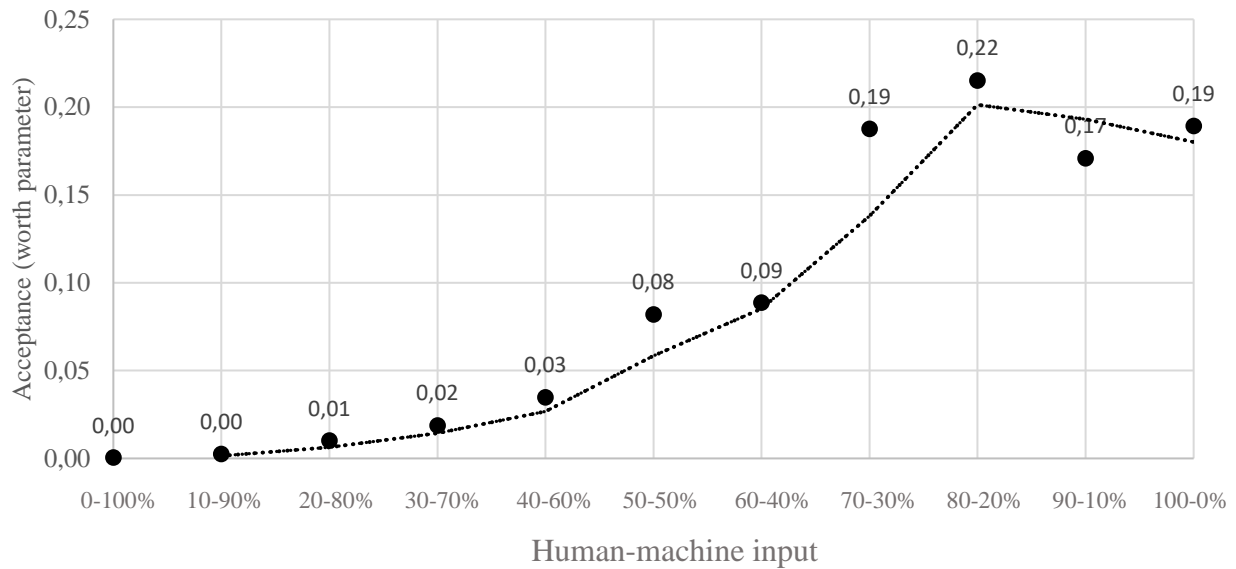


Fig 3. Acceptance (estimated worth values) of the proposed human-machine partnerships in Study 3. Given two partnerships *A* and *B*, the probability that partnership *A* is preferred over partnership *B* is given by the worth of *A* divided by the sum of the worth of *A* and *B*. The dotted line represents a moving average trend line that smooths the data for visualization purposes.




Fig 4. Percentage of participants accepting the proposed human-machine partnerships in Study 4.

The data are collapsed across the six managerial decisions.

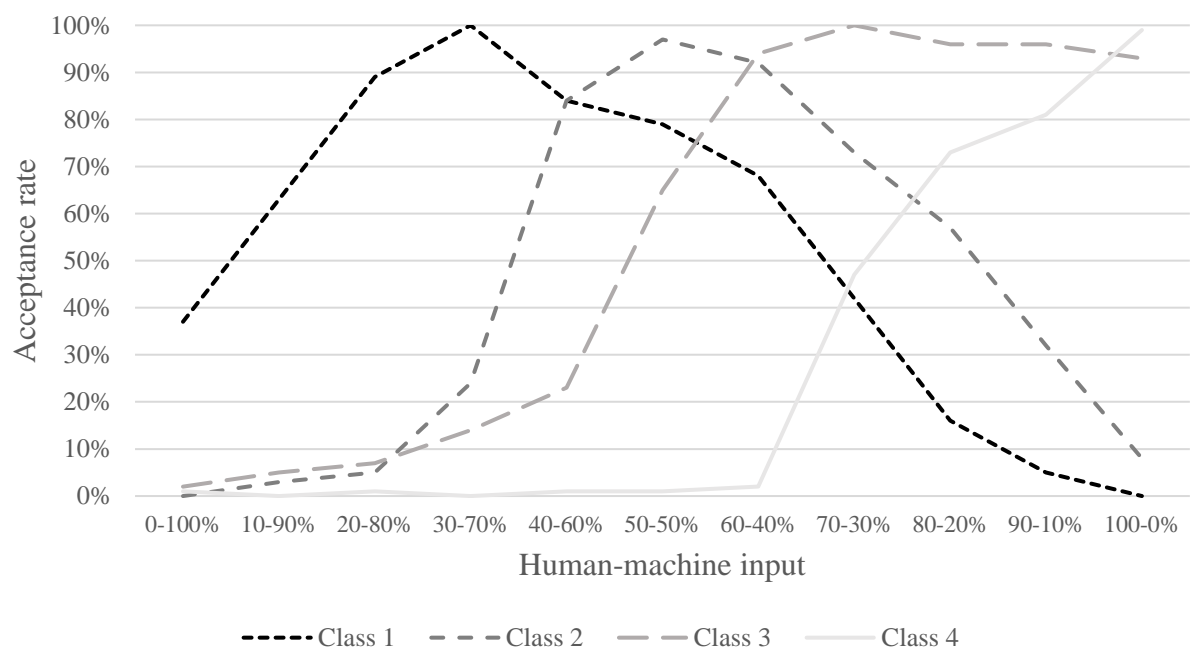


Fig 5. Four different reactions towards the proposed human-machine partnerships in Study 4.
The data are collapsed across the six managerial decisions.




Fig 6. Percentage of participants accepting the proposed human-machine partnerships in Study 5.

The data are averaged over the two measurement moments.




Fig 7. Four different reactions towards the proposed human-machine partnerships in Study 5.

The data are averaged over the two measurement moments.

APPENDIX A

| HUMAN MANAGERS | ALGORITHMS |
|--|--|
| Human managers will have a weight of 0% in managerial decisions. | An algorithms software will have a weight of 100% in managerial decisions. |
| This specific partnership means that <u>human managers</u> will have <i>no input</i> in managerial decisions, while <u>machines</u> will have <i>complete input</i> in managerial decisions. | |

Fig A.1. Example of the presentation of the human-machine partnerships in Study 2.

| Which of the following two <u>human-machine partnerships</u> do you find most ACCEPTABLE? | | |
|---|----|---|
| Partnership A: Human managers: 80% weight Algorithms: 20% weight | OR | Partnership B: Human managers: 30% weight Algorithms: 70% weight |

Fig A.2. Example of the presentation of the pairwise comparisons in Study 3.

APPENDIX B




Fig B.1. Non-significant interaction between the managerial decisions and the human-machine partnerships in Study 4 (decision 1 = whether or not to hire someone; decision 2 = whether or not to fire someone; decision 3 = whether or not to allocate someone a bonus; decision 4 = how large a bonus to give to someone; decision 5 = how much pay increase to award to someone; decision 6 = how many additional paid annual leave days to allocate to someone). The black dotted line represents the average of the six managerial decisions.




Fig B.2 Non-significant interaction between the measurement moments and the human-machine partnerships in Study 5 (moment 1 = decisions regarding how much pay increase to award to someone; moment 2 = decisions concerning how many additional paid annual leave days to allocate to someone). The black dotted line represents the average of the two measurement moments.