

When eliminating bias isn't fair: Algorithmic reductionism and procedural justice in human resource decisions

David T. Newman^{a,*}, Nathanael J. Fast^a, Derek J. Harmon^b

^a University of Southern California, Marshall School of Business, United States

^b University of Michigan, Ross School of Business, United States

ARTICLE INFO

Keywords:

Procedural justice
Fairness
Algorithms
People analytics

ABSTRACT

The perceived fairness of decision-making procedures is a key concern for organizations, particularly when evaluating employees and determining personnel outcomes. Algorithms have created opportunities for increasing fairness by overcoming biases commonly displayed by human decision makers. However, while HR algorithms may remove human bias in decision making, we argue that those being evaluated may perceive the process as reductionistic, leading them to think that certain qualitative information or contextualization is not being taken into account. We argue that this can undermine their beliefs about the procedural fairness of using HR algorithms to evaluate performance by promoting the assumption that decisions made by algorithms are based on less accurate information than identical decisions made by humans. Results from four laboratory experiments ($N = 798$) and a large-scale randomized experiment in an organizational setting ($N = 1654$) confirm this hypothesis. Theoretical and practical implications for organizations using algorithms and data analytics are discussed.

1. Introduction

A recent study of the tech industry found that a perceived lack of fairness was the single largest driver of employee turnover, costing the industry \$16 billion a year (Scott, Klein, & Onovakpuri, 2017). Indeed, scholars have long recognized the importance of understanding and improving employee perceptions of fairness, particularly the perceived fairness of decision-making procedures (Colquitt, Conlon, Wesson, Porter, & Ng, 2001). This work shows that people perceive decision-making procedures as fairer when they are (1) consistent, (2) based on accurate information, and (3) free of influence from the personal biases of decision makers (Brockner, 2006; Leventhal, 1980). These factors are especially important when decisions concern human resource (HR) considerations (e.g., hiring, firing, promotions, etc.), where the subject of a decision is another person. Organizations, therefore, are strongly motivated to find ways not only to make better HR decisions but also to ensure that those affected by such decisions view the decision-making procedures as fair (Weaver & Trevino, 2001).

One approach quickly gaining favor among HR experts is to use algorithms to help make these types of decisions (Christin, 2017). There is a growing awareness that the emergence of big data (George, Haas, & Pentland, 2014) and the use of algorithms to harness this big data to

make decisions (Aral, Brynjolfsson, & Wu, 2012) presents an enormous opportunity for organizations. In part, this is because algorithmic decision making is highly efficient (Wilson, Alter, & Shukla, 2016) and has been shown in some cases to outperform human decision makers at selecting high-performing job candidates (Cowgill, 2017; Kuncel, Klieger, Connelly, & Ones, 2013). Equally important to consider, however, is whether employees view algorithmic decisions as fairer than human decisions. Indeed, if humans and their biases are removed from the decision-making process, this could improve employees' perceptions of the fairness of the decision-making procedures (Colquitt, 2001) and bolster overall procedural justice (Leventhal, 1980; Lind & Tyler, 1988; Thibaut & Walker, 1975).

While this line of reasoning seems possible, it also seems inconsistent with how engineers at Google reacted when their People Analytics department proposed using algorithms to make hiring and promotion decisions. Even though management could demonstrate increased efficiency and equally high levels of accuracy, the engineers rejected it, leading the Vice President of People Analytics to conclude that they "should let people make people decisions" (Setty, 2014). One explanation for this reaction likely has to do with the preferences of the decision makers themselves. Research has shown that decision makers prefer to rely more on their own judgment than on linear models and

* Corresponding author at: Management and Organization, USC Marshall School of Business, 701 Exposition Blvd, HOH 431, Los Angeles, CA 90089-1424, United States.

E-mail address: dtnewman@usc.edu (D.T. Newman).

<https://doi.org/10.1016/j.obhdp.2020.03.008>

Received 25 April 2018; Received in revised form 19 March 2020; Accepted 24 March 2020

0749-5978/ © 2020 Elsevier Inc. All rights reserved.

algorithmic decision aids (e.g., Dana, Dawes, & Peterson, 2013; Dietvorst, Simmons, & Massey, 2014, 2018; Highhouse, 2008), perhaps because they are motivated to justify their own value and maintain control over their decision making. While this line of research helps us understand the reactions of those given the option to use algorithms when making predictions or decisions, it offers little insight into the potential reactions of those *affected* by such decisions. In fact, it could be that the latter would be more open to the use of algorithms, particularly given their potential to increase consistency while eliminating the decision maker's susceptibility to bias.

In this paper, we shed light on this issue by proposing that those affected by decisions made by HR algorithms will perceive those decisions as less fair than decisions made by humans, even when the outcomes are identical, because they will perceive the algorithmic decision-making process as fundamentally reductionistic. Perceptions of algorithmic reductionism can emerge in two ways. Specifically, we suggest that people affected by algorithmic decisions will perceive that the decision-making process reduces the qualitative aspects of their performance to quantifiable metrics (i.e., what we call “quantification”) and consequently fails to adequately consider performance in a broader context (i.e., what we call “decontextualization”). These perceptions of reductionism, we argue, lead people to assume that decisions made by algorithms are based on less accurate (i.e., incomplete) information and thus are less fair than those made by humans (e.g., Leventhal, 1980; Colquitt, 2001). In this sense, the very characteristics that could make algorithms an attractive solution for organizational justice problems—the removal of humans and their (often biased) contextualized decision making—may be precisely what leads people affected by such decisions to perceive algorithms as unfair.

Our theoretical account of algorithms provides three key contributions to the literature. First, our perspective extends current theories of procedural fairness by highlighting the importance of reductionism as a key factor that influences how people interpret decision processes. In particular, we claim that when making decisions with human implications, processes that quantify and decontextualize human performance are judged as less fair than procedures that emphasize qualitative attributes, as a result of the perceived reduction in accuracy of the information considered. Second, the present research provides further theoretical rigor to scholars' efforts to understand the implications of using algorithms to make decisions that affect employees. Indeed, our predictions suggest that people will resist algorithmic evaluations and, instead, have a deep need for qualitative considerations and holistic contextualization of their performance. Third, our findings have implications for research on organizational commitment—particularly affective commitment, which is the desire (rather than need) to remain with an organization (Allen & Meyer, 1990). By eroding perceptions of fairness, reliance on algorithms is likely to have negative downstream organizational consequences.

1.1. Determinants of procedural fairness

Justice and fairness have typically been used interchangeably in the literature, though scholars have recently sought to distinguish the two (Colquitt & Rodell, 2015). In this paper, we treat them as interchangeable and will primarily use the term fairness. Broadly speaking, perceived fairness in organizations refers to employees' global perception that decisions and procedures adhere to agreed-upon rules about equitable treatment. Importantly, different sets of rules influence perceptions of fairness (Colquitt & Zipay, 2015). For example, distributive justice is associated with how people view the fairness of allocation decisions and their outcomes (Adams, 1965; Leventhal, 1976), whereas procedural justice emphasizes how they view the fairness of processes or rules used throughout the decision-making process (Leventhal, 1980; Thibaut & Walker, 1975). Although there are many determinants of

fairness,¹ this paper is concerned with the emerging procedural shift from human- to algorithm-driven decisions, in contexts where the outcomes of such decisions are often unknown. Therefore, we are interested primarily in the procedural determinants of fairness, which require that procedures be consistent, free of bias, and based on accurate information (Barrett-Howard & Tyler, 1986; Brockner, 2002; Colquitt, 2001; Leventhal, 1980; Thibaut & Walker, 1975).

The foregoing criteria of procedural fairness are by no means exhaustive; for example, Tyler and Blader (2003) have highlighted the importance of procedures that emphasize respect, particularly under conditions of strong group identity. We have chosen to focus our theorizing around accuracy, however, because we are interested in determinants of procedural fairness in various organizational contexts, where group identity may be either strong or weak. With that said, it is likely that, in contexts implicating people's livelihood (e.g., human resources), respectful treatment can be linked to the perceived accuracy of the decision. Indeed, a procedure that denies an employee a promotion without taking into account all the relevant facts or considering the broader context will likely feel inherently disrespectful.

The perceived fairness of decision procedures is especially crucial in the design and implementation of a firm's processes affecting personnel (Conlon, Porter, & Parks, 2004). For instance, the perceived fairness of the procedures used to make decisions, such as layoffs, can affect the behavioral intentions (i.e., to engage with the organization) of both customers and potential employees (Skarlicki, Ellard, & Kelln, 1998). Moreover, when employees view decisions that affect them as unfair, a variety of negative consequences can result. For example, perceptions of unfair layoff procedures have been shown to decrease managerial self-esteem, leading to less effective managerial behaviors and, consequently, more negative subordinate perceptions of the work environment (Wiesenfeld, Brockner, & Thibaut, 2000). Moreover, such perceptions of unfairness predict employees' willingness to recruit for former employers, as well as their desire for the regulation of layoffs (Konovsky & Folger, 2006). These organizational outcomes all point to a broader consequence stemming from perceptions of unfairness with respect to decision procedures, namely the damage to different forms of organizational commitment (McFarlin & Sweeney, 1992; Sweeney & McFarlin, 1993). In particular, affective commitment is associated with how much employees *want* to remain with an organization, as distinct from their *need* to do so (Allen & Meyer, 1990), and may therefore be especially salient as an outcome of procedural fairness judgments.

1.2. Data analytics and algorithms

One potential solution to some of these procedural fairness concerns may be the increasing use of algorithms to help make decisions. Algorithms are defined as “computational procedures drawing on some type of digital data that provide some kind of quantitative output through a software program” (Christin, 2017, p. 2). While algorithmic decision making has been applied in many areas, the present work is focused on the use of such computer-based algorithms to make decisions about people, an increasingly prevalent practice referred to as people analytics (Davenport, Harris, & Shapiro, 2010; Waber, 2013). Indeed, leaders and practitioners across a wide variety of industries have suggested that people analytics provides a simple, cost-effective, and efficient way to improve organizational decision making (Fecheyr-Lippens, Schaninger, & Tanner, 2015).

¹ Scholars have also articulated two other dimensions of justice that can affect perceptions of fairness: interpersonal justice, which emphasizes whether decision makers' *conduct* during enactment of procedures treated people with dignity and respect (Boswell et al., 1986; Tyler & Bies, 1990), and informational justice, which emphasizes the candor, thoroughness, and timeliness of authorities' explanations of procedures (Shapiro, Buttner, & Barry, 1994).

While the emergence of certain types of algorithms is new, it is notable that scholars have long been interested in the use of data analytics to improve decision making. Paul Meehl (1954), for example, showed that simple statistical models outperformed human experts at predicting numerical variables of psychological interest. Following Meehl, researchers have shown that statistical linear models (e.g., actuarial, formal, mechanical, and algorithmic) outperform human (e.g., clinical, informal, subjective, and impressionistic) forecasting in such areas as psychiatric diagnosis, academic performance, parole violations, hiring decisions, and other domains (e.g., Arkes, Dawes, & Christensen, 1986; Dawes, 1979; Dawes, Faust, & Meehl, 1989; Grove & Meehl, 1996; Highhouse, 2008). Yet despite the superior performance of statistical models, researchers have found that decision makers themselves prefer to rely on their own judgment (Dana et al., 2013; Dawes et al., 1989; Grove & Meehl, 1996; Hastie & Dawes, 2010; Highhouse, 2008), likely to assuage ethical concerns or maintain personal control in the decision-making process (e.g., Dana et al., 2013; Grove & Meehl, 1996; Highhouse, 2008; see also Highhouse, Nye, & Zhang, 2019).

Importantly, this research has focused on *decision makers'* use of algorithms (see also Legg, Minson, & Moore, 2019). While this work may be relevant to the psychology of those considering the use of an algorithm to make a prediction or decision, it offers little insight into the potential reactions of those *affected* by the decision, which is our concern here. Indeed, because employees are the *subjects* of algorithmic decisions and do not have the same drive as managers to maintain control over these procedures, it is likely that those affected by HR decisions are operating with a different set of psychological needs. For this reason, it is important to examine their reactions to the use of algorithmic decision making.

1.2.1. Algorithms and the suppression of human bias

One possibility is that people affected by algorithmic decision making will view such decisions as fairer than decisions made by humans. Indeed, to view decisions as procedurally fair, one must perceive that they are free of bias (Leventhal, 1980). Empirical research has shown that the removal of bias is particularly important to procedural fairness in formal, business-like situations (Barrett-Howard & Tyler, 1986). Moreover, suppressing bias in decision making has been shown to increase employees' perceptions of fairness even when they receive lower performance evaluations (Taylor, Tracy, Renard, Harrison, & Carroll, 1995), and Sheppard and Lewicki (1987) study of organizational executives found that bias suppression was one of the most frequently cited characteristics of procedural fairness.

Building on this line of reasoning, computer-based algorithms have been publicly hailed as the next frontier in eliminating bias (Loehr, 2015). Whereas human decision makers are prone to judgment errors due to biases derived from intuition and other heuristics (e.g., Gilovich, Griffin, & Kahneman, 2002), algorithms can reduce, or even eliminate, such biases by relying on mathematical logic that converts various considerations (both quantitative and qualitative) into numerical factors. In other words, algorithms seem to remove bias and improve decision making by removing human subjectivity from the process of judging and comparing individuals (e.g., Cowgill, 2017). Thus, algorithms enable organizations to relieve decision making from subjectivity and other limitations, which has led to some organizations deploying algorithms to improve the fairness of personnel decision rules (O'Connor, 2016).²

² To be clear, algorithms are not a silver bullet. If not carefully implemented, algorithms can still contain biases stemming from flawed assumptions or incomplete data (Dana et al., 2013; Jackson, 2017). Human decision makers might also unintentionally import their biases (or those of preexisting datasets) into the algorithmic code itself. With these limitations acknowledged, it nevertheless remains a possibility that algorithm-driven procedures have the potential to reduce bias by reducing human involvement and increasing consistency.

1.2.2. Algorithms and perceived reductionism

However, there are reasons to believe that removing human decision makers in favor of algorithms may also generate concerns. In fact, while such quantification processes may make organizational decisions objectively more efficient, less biased, and even fairer in one respect, organizations that quickly adopt algorithms as a solution to biases in human decision making might do so at the cost of appearing to neglect (1) the qualitative characteristics of human nature and, by extension, (2) the contextualized circumstances in which they occur. More specifically, we argue that this may lead people to perceive algorithms as fundamentally reductionistic. Conceptually, reductionism is an analytic, mechanistic process associated with quantifying information about the world (Von Bertalanffy, 1972). As a result, evaluative procedures will be perceived as reductionistic when they take various inputs commonly considered to be qualitatively rich and either delete them from the calculus in favor of quantifiable variables or reduce them to numerical representations.

This has led some scholars to associate reductionism with decontextualization. In contrast to holism, which considers each element of a situation in light of all the others, reductionism decontextualizes individual characteristics and considers them in isolation of the broader context (Choi, Koo, & Choi, 2007; Nisbett, Peng, Choi, & Norenzayan, 2001). The apparent limitations of this practice have motivated work urging the development of holistic thinking in management (Jackson, 2003). It is important to note that algorithms may today, or in the future, be sophisticated enough to handle holistic analyses. However, regardless of their objective efficacy at this point in time, our focus in the present research is to examine how those affected by algorithmic decision making *subjectively* perceive that process. In that regard, we contend that perceptions of decontextualization are a natural consequence of the view that algorithms reduce information through quantification. Indeed, the CEO of an HR consulting firm recently remarked that he believed algorithms “miss the value of what [an] individual brings to the table in terms of personality, connectivity, and those intangible pieces” (Associated Press, 2015).

While quantification and decontextualization are two separate aspects of reductionism, we theorize that the former leads to the latter. First, people will think that algorithms quantify information about people and thereby fail to accurately measure certain important qualitative characteristics. Second, this apparent failure leads to perceptions of decontextualization, or the inability to accurately weigh and combine disparate pieces of information, precisely because important information has been rendered unavailable via quantification. Taken together, we call these two steps—quantification leading to decontextualization—the process of reductionism.

Given that fair procedures require reliance on accurate information (e.g., Leventhal, 1980), these perceptions of reductionism may prove problematic. Indeed, extant research in HR management has shown that job candidates respond negatively to structured recruitment practices to the extent that such practices cause them to feel like their personal characteristics have been reduced to “a number” (Boswell, Roehling, LePine, & Moynihan, 2003, p. 33). Because factors such as personality, intentionality, and potential are not *prima facie* reducible to numbers, or at least it is challenging for the average person to understand how they could be reducible to numbers, an algorithmic decision-making process that engages in a numerical calculus to evaluate employees will be perceived to quantify and thereby decontextualize the human qualities that make up a whole person. This, in turn, will lead people affected by algorithmic decisions to assume that they rely on less accurate information and are therefore less fair than those made by humans.

Hypothesis 1. Individuals affected by personnel decisions will perceive algorithm-driven decisions as less fair than the same decisions made by humans.

Hypothesis 2. The perception that algorithms are reductionistic (i.e., that they reduce the use of accurate information via quantification and

decontextualization) will mediate the effect of algorithm-driven decisions on perceived fairness.

Justice theory suggests that organizational commitment is a downstream consequence of fairness perceptions, particularly as respondent to procedural justice (Colquitt et al., 2001; Lavelle et al., 2009). For example, when employees see personnel decision-making procedures as fair, they show high commitment to the organization, even in the face of dissatisfying personal outcomes (McFarlin & Sweeney, 1992; Sweeney & McFarlin, 1993). Hence the fairness of an organization's decision procedures bears a direct relationship to its employees' willingness to commit. To the extent that procedures are seen as reductionistic and thus unfair, they may undermine commitment from those whom they affect and lead to a desire to leave. As alternative work arrangements increasingly disrupt the formerly linear nature of careers (Spreitzer, Cameron, & Garrett, 2017), organizations wishing to attract and retain human capital should take seriously the relationship between procedural justice and employee commitment. Based on this reasoning combined with the predictions above, we contend that the use of HR algorithms will have the effect of reducing the desire to remain in one's organization.

Hypothesis 3. Individuals affected by personnel decisions will voice lower levels of organizational commitment when subjected to algorithm-driven decisions compared to the same decisions made by humans.

Hypothesis 4. The perception of fairness will mediate the effect of algorithm-driven decisions on organizational commitment.

Finally, while it may be possible to entirely automate algorithmic decisions (i.e., with no human involvement), many decisions are made by combining human and algorithmic judgment. In line with this possibility, there is some evidence that decision makers prefer to rely on algorithms if and when they have the opportunity to adjust the algorithm's decision (Dietvorst, Simmons, & Massey, 2018). Indeed, decision makers were more comfortable using algorithms when they were given the ability to slightly tweak the algorithm's output (by 2%, 5%, or 10%). Dietvorst et al. (2018, p. 1156) concluded: "If allowing people to adjust an imperfect algorithm by only a small amount dramatically increases their willingness to use it, then people's judgments will be much more reliant on the algorithm, and much more accurate as a result."

It is important to note, however, that in the present study we are concerned not with reliance on algorithms by people vested with decision-making authority, but with perceptions of procedural fairness by those affected by such decisions and who lack opportunities for control. Because such fairness perceptions hinge on the *apparent* consideration of accurate information, we argue that reactions to conjoint human-algorithmic procedures will depend on whether human or algorithmic processes predominate. This line of reasoning is broadly consistent with prototype theories of cognitive representations (e.g., Rosch, 1975) in that the central features of an object determine its fundamental categorization. Specifically, we contend that when algorithmic decision processes are the default and a human can only tweak an algorithm-driven decision, then this will be seen as similarly unfair as a decision made purely by an algorithm. In contrast, when human decision processes are the default and the human decision maker is given the option to incorporate algorithmic input into his or her own decision, then this will be perceived as similarly fair as a decision made purely by a human.

Hypothesis 5. The negative effect of algorithm-driven decisions on perceived fairness will be mitigated by an algorithm-human partnership, but only when a human (rather than an algorithm) is the default decision maker.

1.3. Overview of studies

We conducted five studies to test our hypotheses. Study 1 is a laboratory experiment where we manipulated the decision process and then measured participants' perceptions of reductionism and fairness (H1, H2). Study 2 replicated these findings in a large-scale field experiment that used random assignment methodology to introduce employees of a large organization to different kinds of centralized HR decision-making processes and assess the effects of these treatments on fairness perceptions and organizational commitment (H3, H4). Study 3 extended our findings by experimentally varying the degree of algorithmic or human involvement in the decision process (H5). Studies 4 and 5 explored boundary conditions and finer-grained analyses of our mediating mechanisms to better understand the robustness of our results. Whereas Studies 1–3 examined quantification, Studies 4–5 brought in measurements of decontextualization. Study 4 examined our predictions under conditions of greater transparency about the inputs into the decision-making process and tested a serial mediation model that included both of our mechanisms (i.e., quantification and decontextualization) of perceived reductionism. Study 5 tested our predictions in a behavioral laboratory experiment with significant stakes and high external validity, leveraging a research design that mirrors a real-world context where participants' job interview videos are evaluated. Studies 4 and 5 were preregistered on AsPredicted.org. We designed all studies to hold other decision-making factors constant so as to isolate the effects of algorithms on perceived fairness.

2. Study 1: Algorithms and fairness perceptions

In this study, we tested our predictions in a sample of working adults who had a variety of work experiences. We assessed people's fairness perceptions associated with using algorithms (versus humans) to make personnel decisions, predicting that algorithms would be perceived as less fair. Additionally, we examined whether differences in fairness perceptions were driven by perceptions of reductionism (via quantification). To measure the robustness of the effects, we included experimental conditions that described layoffs as well as promotions, to see whether one or the other had a greater consequence. As promotions are likely to be perceived as a fairer outcome than layoffs (Brockner & Wiesenfeld, 1996), this enabled us to examine whether differences in outcome favorability influenced perceptions of algorithms.³ Finally, we assessed other possible mechanisms, such as the number of overall factors thought to be considered in the decision (i.e., thoroughness) as well as the perceived typicality of using algorithms to make personnel decisions. Because some researchers have found other moderators of people's preference for algorithms, such as personal control over model output (Dietvorst et al., 2018), we held constant the basis for the decision-making process.

2.1. Sample and procedure

A sample of 199 participants (41.2% female, $M_{age} = 32.9$, $SD = 10.3$, 76.4% Caucasian, 6% African American, 10.1% Asian/Asian American, 4.5% Hispanic, 0.5% Native American, 2.5% Other) from Amazon's Mechanical Turk worker pool took part in a 2×2 survey design that asked them to evaluate one of two types of personnel judgments (layoffs versus promotions) made by one of two types of decision-making processes (an algorithm versus managers). We examined whether using algorithms would have a differential effect on fairness perceptions, and also whether the outcome valence for employees (positive versus negative) would make a difference.

³ While outcome valence is not identical to varying levels of distributive fairness, research suggests that they overlap considerably enough to treat as convergent constructs (Brockner, 1996).

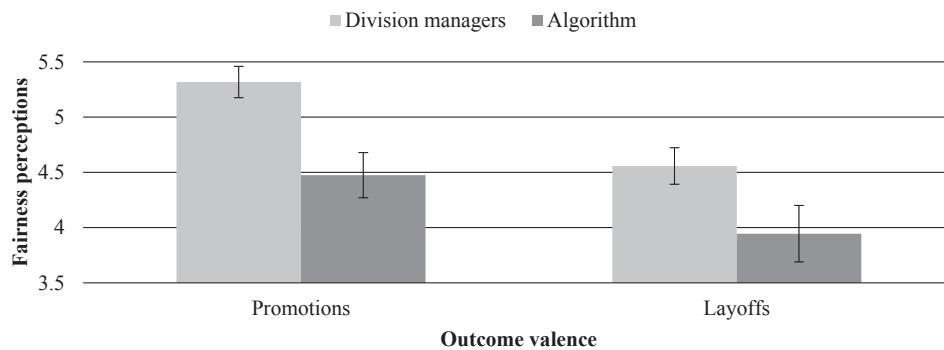


Fig. 1. Fairness of algorithms vs. managers.

Participants were randomly assigned to read one of the following four scenarios: “Company X is [promoting] [laying off] 50% of its workforce. The workers [promoted] [laid off] are determined by [an algorithm (i.e., a computerized decision-making tool) that] [the company’s division managers, who] take[s] into account a variety of factors.” This reflects an emerging use of algorithms by companies like Amazon to automate employee termination (Lecher, 2019).

2.2. Measures

2.2.1. Perceived fairness

Adapting the organizational justice scale from Conlon et al. (2004), we developed a 4-item measure using the items that were most relevant to fairness. On a 1–7 scale (Strongly Disagree to Strongly Agree), participants indicated their agreement with four statements (presented in random order): “The way this company determined which workers were [promoted] [laid off] seems fair,” “The company’s process for deciding which workers were [promoted] [laid off] was fair,” “The decision regarding which workers were [promoted] [laid off] was fair,” and “The outcome of this decision was fair.”

2.2.2. Quantitative and qualitative factors

Participants indicated the degree to which they thought each of six considerations (presented in a randomized order) was part of the decision on a 1–7 scale (“This consideration was not part of the decision” to “This consideration was an essential part of the decision”). Three items assessed quantitative considerations (i.e., profitability to the company, number of hours worked, and client billing), and three assessed qualitative considerations (i.e., leadership skills, attitude toward clients, and potential for productivity). These measures were developed based on the notion that quantification is a form of reductionistic thinking, which involves reducing the types (or forms) of factors considered to those that can be represented by numbers. Conversely, qualitative considerations, such as leadership, attitude, and potential, are potentially richer in nature, and are often necessary to consider if the objective is to make holistic, contextualized decisions.

2.2.3. Typicality

One factor that could influence participants’ perceptions of fairness is the difference in how typical, or normative, it is to make decisions with algorithms versus humans. In order to address this concern, we asked participants to rate the typicality of the company’s decision-making process on a slider from 0 to 100 (“Not typical at all” to “Completely typical”).

2.2.4. Number of factors thought to be considered

We have hypothesized that differences in the type of factors (qualitative versus quantitative) should drive participants’ views of fairness. However, another possible variable that could influence perceptions of fairness is the number of factors thought to be considered. Perhaps participants will perceive that humans consider more factors than

algorithms and, as a result, view them as fairer. To control for this, we asked participants to use a slider from 1 to 20 to indicate how many different factors they thought were considered in determining which workers were affected.

2.3. Results

The results support Hypothesis 1 in that decisions made by algorithms ($M = 4.21$, $SD = 1.65$) were perceived as significantly less fair than those made by human managers ($M = 4.92$, $SD = 1.16$), ($F(1, 197) = 12.61$, $d = 0.50$, $p < .001$).⁴ Additionally, promotions ($M = 4.89$, $SD = 1.30$) were perceived as significantly fairer than layoffs ($M = 4.26$, $SD = 1.55$), ($F(1, 197) = 9.78$, $d = 0.44$, $p = .002$). The interaction of the manipulations was not significant. See Fig. 1.

We also found support for Hypothesis 2, which predicted that algorithms would be perceived as less fair because of their perceived reduction of accurate information. In terms of considering quantitative factors, algorithm-driven decisions ($M = 5.07$, $SD = 1.24$) were perceived to be no different from manager-driven decisions ($M = 5.00$, $SD = 1.03$), ($F(1, 197) = 0.157$, $d = 0.06$, $p = .692$). However, algorithmic decisions were perceived to give significantly less consideration to qualitative factors ($M = 4.53$, $SD = 1.57$) compared to managerial decisions ($M = 5.43$, $SD = 1.15$), ($F(1, 197) = 21.57$, $d = 0.66$, $p < .001$). See Fig. 2.

While the perceived consideration of quantitative factors was associated with perceived fairness ($r = 0.186$, $p = .008$), the perceived consideration of qualitative factors was more strongly correlated with fairness perceptions ($r = 0.505$, $p < .001$). In a combined regression model using both types of factors as predictors, the link between qualitative factors and fairness persisted ($\beta = 0.50$, $p < .001$), while the link between quantitative factors and fairness disappeared ($\beta = 0.03$, $p = .596$). Hence although consideration of each type of factor was, in general, predictive of perceived fairness, consideration of qualitative factors was the stronger predictor.

We then assessed whether the main effect of algorithms on perceived fairness was mediated by perceived reductionism (i.e., perceived failure to consider qualitative factors). We used the Preacher and Hayes (2008) procedure for testing for multiple mediation and found support for this prediction, determining that perceived consideration of qualitative factors mediated the effect of the algorithm condition on perceived fairness, even controlling for ratings of typicality (the bias-corrected confidence interval ranged from -0.48 to -0.04).

One possible alternative reason for the fairness advantage of human decision makers is that they could be perceived as more typical than algorithms. Participants gave the algorithmic process an average typicality rating of 38.04 ($SD = 28.03$) and the managerial process an average typicality rating of 66.76 ($SD = 24.04$); this difference was

⁴ This effect was not moderated by gender or ethnicity (either Caucasian vs. non-Caucasian or underrepresented minority (URM) vs. non-URM).

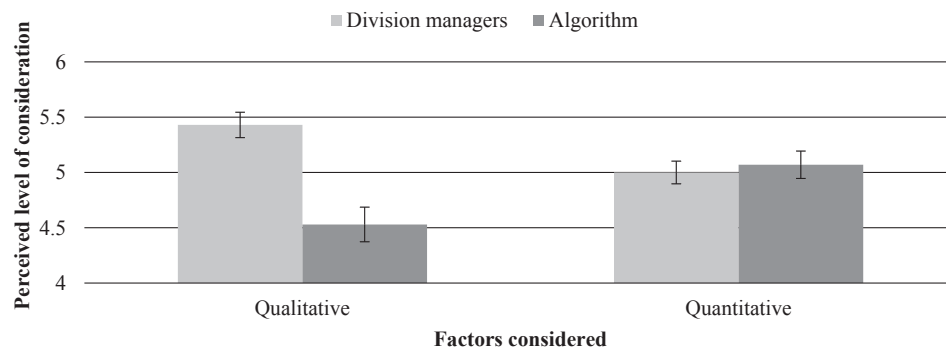


Fig. 2. Perceived consideration of factors, algorithms vs. managers.

statistically significant ($F(1, 197) = 60.24, d = 1.1, p < .001$). Participants' typicality ratings were significantly correlated with ratings of perceived fairness ($r = 0.424, p < .001$). Hence our mediation analysis (above) controlled for ratings of typicality to rule it out as an alternative mechanism.

Another possible alternative reason for the fairness advantage of human decision makers is that they could be seen to consider more factors than algorithms. Indeed, participants' ratings of the number of factors considered by the process were significantly correlated with ratings of fairness ($r = 0.174, p = .014$). However, the data do not support this interpretation. Participants on average rated algorithms as considering 7.46 factors ($SD = 3.00$) and managers as considering only 5.98 factors ($SD = 4.49$). This difference was statistically significant ($F(1, 197) = 7.55, d = 0.39, p = .007$). Thus, humans were not perceived to be fairer as a result of considering more factors.

2.4. Discussion

These results offer initial support for Hypotheses 1 and 2. Compared to human decision makers, algorithm-driven processes are perceived to render less fair decisions, and this effect on fairness perceptions is mediated by the belief that algorithms reduce information about employees by disregarding their qualitative attributes. Importantly, algorithms are viewed as no different from managers in their consideration of quantitative inputs (e.g., client billing), but they are thought to be far apart in their consideration of qualitative factors (e.g., leadership skills). Although people believe that algorithms consider numerically more factors overall than do managers, the positive effect of the number of factors considered on fairness is more than offset by reduced consideration of qualitative factors (i.e., quantification).

3. Study 2: Algorithms, fairness, and organizational commitment

Study 1 provided support for our first two predictions. In Study 2, we moved to an organizational setting to further test our prediction that algorithms would be seen as less fair than humans when making identical HR decisions. Moreover, we assessed our prediction that this decreases organizational commitment, and that this effect is mediated by a decrease in perceived fairness. To test these ideas, we randomly assigned employees of a large private university to read about new policies that would make personnel decisions with either humans or algorithms and either in a quantifying or nonquantifying manner. Aiding the believability of our experiment, the staff was aware that future changes would be made in HR practices but was not yet aware of what the changes would be, so we were able to take advantage of the timing to amplify the personal resonance and external validity of our scenarios.

In addition to testing our predictions about the downstream consequences of algorithmic decisions, we examined the persistence of people's beliefs about algorithms. To do so, we deliberately

manipulated both the reductionistic and algorithmic nature of the HR process. Thus, we created four procedures varying the balance between reductionism and human involvement: (1) humans using qualitative factors, (2) humans using quantitative factors, (3) algorithms using qualitative factors, and (4) algorithms using quantitative factors. One possible outcome would be an interaction effect (e.g., where algorithms decrease perceived fairness, but only relative to human decision makers). A second possible outcome would be a stepwise effect with the strongest effects for perceived unfairness (fairness) occurring when both (neither) variables are manipulated. In other words, it is possible that algorithms will be viewed as reductionistic (relative to human decision makers) even when it is claimed that they measure qualitative factors. If another assumption about algorithms is that they decontextualize information, it may be almost impossible to completely remove their reductionistic (and thus unfair) aspects, even by explicitly stating that they take into account specific qualitative factors.

3.1. Sample and procedure

A sample of 1654 employees (66.8% female, $M_{age} = 38.6, SD = 11.4$, 42.1% Caucasian, 6.5% African American, 20.6% Asian/Asian American, 22.3% Hispanic, 0.4% Native American, 8.1% Other) from a large private university took part in a 2×2 survey design that asked them to evaluate a new HR decision-making process described by one of two organizational charts (involving an HR algorithm versus an HR team), which emphasized specifically one of two types of factors (quantitative versus qualitative). Using this method, we experimentally examined the causal impact of reductionism (via quantification) enacted by both algorithmic and human decision makers. While we did not, in fact, implement any new HR policies, we introduced (and assessed reactions to) the idea of a new policy, which commonly occurs in organizations.

We began by obtaining permission from the university's human resources department to conduct a full-scale field experiment sampling the entire corpus of over 13,000 employees. Using a list of email addresses drawn from the payroll system, we contacted our sample population via an email asking employees to participate in a survey about a new HR decision-making process. Employees were incentivized to respond to the survey with entry into a lottery to win a \$300, \$200, or \$100 Amazon.com gift card. A total of 2151 people participated in the study, yielding a response rate of approximately 16.3%. This number was reduced by the removal (before analysis) of 497 participants who did not fully complete the survey materials.

Participants were randomly assigned to read one of four scenarios, each accompanied by its own organizational chart, which described a new employee evaluation process under consideration at many large private universities. According to the new process, decisions with respect to promotions, layoffs, raises, and pay cuts are determined by either a "Human Resources Algorithm" or a "Human Resources Team," which takes into account "a variety of factors, such as" either "(1)

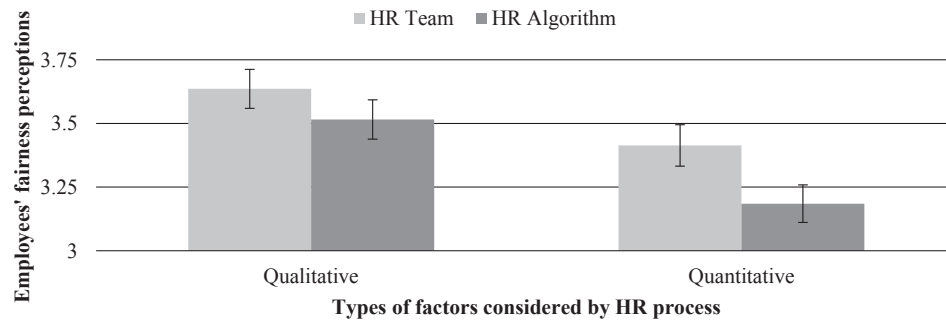


Fig. 3. Fairness of HR processes.

profitability to the university, (2) number of hours worked, and (3) project completion rate” (quantitative) or “leadership skills, attitudes toward others, and potential for growth” (qualitative). We refrained from numbering the qualitative factors to avoid priming participants with the concept of numbers. The complete organizational charts used to operationalize our experimental conditions may be found in Appendix A.

In addition to manipulating the algorithmic nature of the decision process, we directly manipulated reductionism (via explicit emphasis on quantitative factors) to see if its effect would vary between human- and algorithm-driven procedures.

3.2. Measures

3.2.1. Perceived fairness

On a 1–7 scale (Strongly Disagree to Strongly Agree), participants indicated their agreement with four (randomized) statements adapted from Conlon et al. (2004): (1) the way this process determines which workers receive raises and pay cuts seems fair, (2) this process for making personnel decisions is fair, (3) this process regarding which workers are promoted and laid off is fair, and (4) the outcomes of this process are fair.

3.2.2. Organizational commitment

Employees were asked how they would feel if their organization adopted the previously described decision-making process. On a 1–7 scale (Strongly Disagree to Strongly Agree), participants indicated their agreement with eight statements (presented in random order) from the organizational commitment scale (affective commitment subscale) developed by Allen and Meyer (1990). Sample items include “I would be very happy to spend the rest of my career with this organization” and “I would not feel a strong sense of belonging to my organization.”

3.2.3. Employment characteristics

Employees answered questions regarding faculty status, number of years with the organization, seniority within their department, power over others resulting from their role, and whether they supervised any other employees. Including these factors as controls did not change our findings.

3.3. Results

Analysis of variance revealed that the HR algorithm ($M = 3.35$, $SD = 1.54$) was perceived as significantly less fair than the HR team ($M = 3.52$, $SD = 1.61$), ($F(1, 1652) = 5.30$, $d = 0.11$, $p = .021$), and explicit emphasis on quantitative factors ($M = 3.30$, $SD = 1.59$) was perceived as significantly less fair than explicit emphasis on qualitative factors ($M = 3.58$, $SD = 1.55$), ($F(1, 1652) = 12.98$, $d = 0.18$, $p < .001$).

Planned linear contrast effects ($-3, 1, 1, 1$) revealed that quantitatively driven algorithms were perceived as less fair than all other HR

processes ($t(1650) = 3.81$, $d = 0.22$, $p < .001$), and qualitatively driven teams were perceived as more fair than all other HR processes ($t(1650) = 2.97$, $d = 0.17$, $p = .003$). There was no interaction of the manipulations, suggesting that the effects of these conditions on perceived fairness layered additively. See Fig. 3.

Analysis of variance revealed that the HR algorithm ($M = 3.63$, $SD = 1.10$) led to significantly less organizational commitment than did the HR team ($M = 3.80$, $SD = 1.06$), ($F(1, 1650) = 10.98$, $d = 0.16$, $p < .001$), and the explicit emphasis on quantitative factors ($M = 3.60$, $SD = 1.11$) led to significantly less organizational commitment than did the explicit emphasis on qualitative factors ($M = 3.83$, $SD = 1.05$), ($F(1, 1650) = 17.21$, $d = 0.20$, $p < .001$).

Planned linear contrast effects ($-3, 1, 1, 1$) revealed that quantitatively driven algorithms led to less organizational commitment than did all other HR processes ($t(1648) = 4.56$, $d = 0.26$, $p < .001$), while qualitatively driven teams led to more organizational commitment than did all other HR processes ($t(1648) = 4.02$, $d = 0.23$, $p < .001$). Once again, there was no interaction between the manipulations, suggesting that the effects of these conditions on organizational commitment layered additively.⁵ See Fig. 4.

Mediational analysis revealed that perceived fairness mediated the effect of the HR algorithm on organizational commitment (bias-corrected confidence interval: -0.13 to -0.01), and that perceived fairness also mediated the effect of quantification on organizational commitment (bias-corrected confidence interval: -0.17 to -0.05).

3.4. Discussion

In addition to replicating Study 1s support for $H1$ (algorithms are perceived as less fair than humans), the results of Study 2 provide causal evidence for the notion that processes emphasizing quantitative factors are considered less fair than processes emphasizing qualitative factors, regardless of their algorithmic or human nature. Hence the requirement that fair procedures be based on accurate information is particularly linked to qualitative consideration. Additionally, Study 2 offers support for $H3$ and $H4$: an organization’s use of reductionistic decision-making processes (whether implemented by algorithms or by humans) to evaluate employees will reduce organizational commitment, an effect mediated by the perception that fairness has been violated.

The results of our analyses of variance with planned linear contrast effects demonstrated that quantitatively driven algorithms were perceived least favorably compared to all other conditions, while qualitatively driven teams were perceived most favorably compared to all other conditions. Our results show additive main effects for algorithms and quantitative reductionism, implying that the belief in the unfairness of using algorithms to evaluate employees is persistent, even when algorithms are explicitly described as tools that consider qualitative

⁵ None of the effects on fairness or commitment showed moderation by gender or ethnicity (either Caucasian vs. non-Caucasian or URM vs. non-URM).

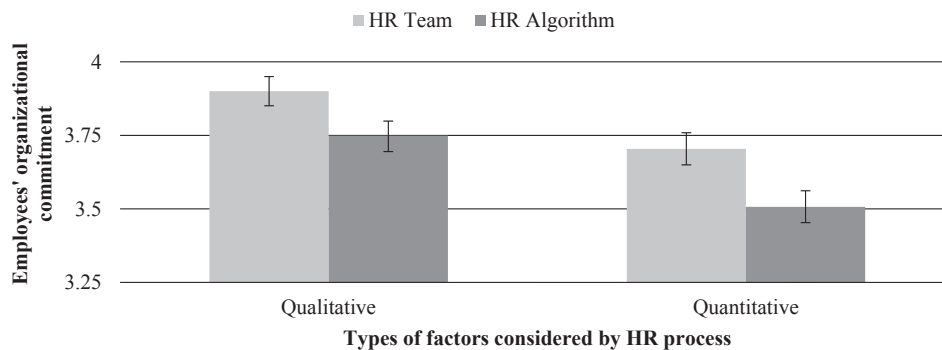


Fig. 4. Organizational commitment induced by HR processes.

factors, likely as a result of how those factors are thought to be decontextualized (which we examine in Studies 4 and 5). Similarly, human decision makers may be seen as better at contextualizing information even when emphasizing quantitative factors.

A potential limitation of this study is that there are no conditions where either algorithms or humans specifically consider both quantitative and qualitative measures of performance. This was a deliberate choice so as to set up situations in which there was no room for people to assume differences in the type of information under consideration. This limitation is addressed, to some extent, by Studies 4 and 5, which included explicit descriptions of a wider variety of factors that were considered.

4. Study 3: Examining algorithm-human combinations

Studies 1 and 2 supported our predictions that algorithmic processes would be seen as less fair than human-driven processes as a result of perceptions of reductionism. In this study, we further tested our predictions by varying the degree of algorithmic versus human involvement in the decision process. Furthermore, we examined the possibility that including human involvement at any point in the decision process might serve to enhance perceptions of fairness.

On the one hand, human involvement, even if minimal, might serve to overcome the negative consequences of algorithms on perceived fairness observed in Studies 1 and 2. For instance, recent work has shown that, in some forecasting contexts, people show a greater preference for relying on algorithms to aid their own decisions if they have the chance to tweak their output by a margin of 2%, 5%, or 10% (Dietvorst et al., 2018). However, we suggest that these findings may not carry over to the current domain. First, Dietvorst et al. (2018) examined the preference of the decision maker him- or herself to choose an algorithm to aid in decision making, whereas we are interested in perceptions of fairness from those affected by the decisions. As participation in decision making is a contributor to fairness judgments (e.g., Lind, Kanfer, & Earley, 1990), we propose that while a tweak in control may be personally meaningful to people using algorithms as prediction tools (Dietvorst et al., 2018), giving the decision maker the ability to make a small tweak may not move the dial for those subjected to the algorithmic decisions. Second, prior research examined reliance on algorithms, whereas we are interested in perceived fairness, and, as we have seen, the two concepts are distinct. From a justice perspective, we predict that the opportunity for a human decision maker to make a minor adjustment to an algorithm's output will not substantially diminish perceptions of the algorithm as an unfair reductionistic process because the algorithm will be perceived as the primary driving force in the decision procedure. On the other hand, we predict that the option to include algorithmic determinations as just one additional tool to include for consideration at the end of a human-driven decision process will be seen as comparatively less reductionistic and therefore fairer. This is because the human would have already arrived at an initial position on the decision before the algorithm comes into play.

4.1. Sample and procedure

A sample of 189 participants (45% female, $M_{age} = 35.3$, $SD = 10.5$, 79.4% Caucasian, 7.4% African American, 9.5% Asian/Asian American, 2.6% Hispanic, 0.5% Native American, 0.5% Other) from Amazon's Mechanical Turk worker pool took part in a 4-condition experimental design that asked them to evaluate a company's process for allocating end-of-the-year bonuses, which were determined with varying degrees of algorithmic versus human involvement. We examined whether the point at which algorithms were introduced to the decision-making process would have a differential effect on fairness perceptions. In particular, when the algorithm is introduced first, the ability of a human to tweak it is not enough to overcome the perception of unfairness, and vice versa.

Participants were randomly assigned to read one of four scenarios. Each scenario began by stating that "Company X just went through the process of making its end of the year bonus point, the scenarios diverged:

Algorithm-only: "In order to determine the size of the bonus for each employee, Company X relied on an algorithm (a computerized decision-making tool) that took into account a variety of factors. After the algorithm made a series of computations, it determined how employee bonuses should be allocated."

Algorithm with human adjustment: "In order to determine the size of the bonus for each employee, Company X relied on an algorithm (a computerized decision-making tool) that took into account a variety of factors. After the algorithm made a series of computations, it determined how employee bonuses should be allocated. At this point, the company's human resources team was able to adjust the algorithm's determinations up or down within a margin of 10%."

Human with algorithmic option: "In order to determine the size of the bonus for each employee, Company X relied on its human resources team, which took into account a variety of factors. After the human resources team made a series of deliberations, it determined how employee bonuses should be allocated. The human resources team had the option to consider the recommendations of an algorithm (a computerized decision-making tool) to assist in making its determinations."

Human-only: "In order to determine the size of the bonus for each employee, Company X relied on its human resources team, which took into account a variety of factors. After the human resources team made a series of deliberations, it determined how employee bonuses should be allocated."

4.2. Measures

4.2.1. Perceived fairness

On a 1–7 scale (Strongly Disagree to Strongly Agree), participants indicated agreement with four (randomized) statements adapted from Conlon et al. (2004): (1) the way this company determined how to allocate employee bonuses seems fair, (2) this company's process for

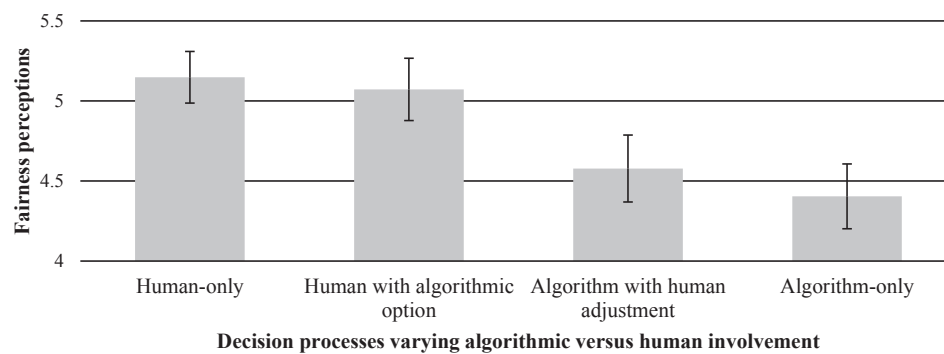


Fig 5. Fairness of various decision processes.

deciding which employees receive bonuses was fair, (3) this decision regarding which employees received bonuses was fair, and (4) the outcome of this decision was fair.

4.2.2. Quantitative and qualitative factors

Participants indicated the degree to which they thought each of six considerations was part of the decision process on a 1–7 scale (“This consideration was not part of the process” to “This consideration was an essential part of the process”). Three of the factors were numbers-based and more easily identifiable as quantitative: profitability to the company, number of hours worked, and client billing. The other three factors were more qualitative in nature and not easy to quantify with numbers: leadership skills, attitude toward clients, and potential for productivity. All factors were presented in randomized order.

4.3. Results

Analysis of variance revealed significant differences in perceived fairness across all four decision processes ($F(3, 185) = 3.62, d = 0.55, p = .014$). Further probing with planned linear contrast effects ($-3, 1, 1, 1$) revealed that the algorithm-only process was seen as less fair than all other processes ($t(185) = 2.38, d = 0.40, p = .019$), while the human-only process was seen as more fair than all other processes ($t(185) = 2.11, d = 0.35, p = .036$). Consistent with Hypothesis 5, planned linear contrast effects ($-1, -1, 1, 1$) revealed that processes driven primarily by algorithms were seen as less fair than processes driven primarily by human agents ($t(185) = 3.22, d = 0.47, p = .002$).⁶ See Fig. 5.

As in Study 1, we assessed whether these effects on perceived fairness were mediated by perceived quantification (i.e., perceived failure to consider qualitative factors). We used the Preacher and Hayes (2008) test for multiple mediation and found support for this prediction, determining that perceived consideration of qualitative factors mediated (1) the difference between algorithm-only and all other conditions (bias-corrected confidence interval ranged from -0.43 to -0.01), (2) the difference between human-only and all other conditions (bias-corrected confidence interval ranged from 0.11 to 0.50), and (3) the difference between primarily algorithm-driven and primarily human-driven processes (bias-corrected confidence interval ranged from -0.44 to -0.11).

4.4. Discussion

These results replicate our findings supporting Hypotheses 1 and 2. Additionally, they provide support for Hypothesis 5, predicting that the perceived fairness of hybrid procedures depends on the degree of human control over algorithms. If human decision makers have the opportunity to adjust the output of an algorithmic process, that process

is nonetheless still viewed as reductionistic and therefore unfair. On the other hand, if the algorithm is just one factor human decision makers have the option to consider, that process is seen as relatively less reductionistic and therefore fairer. This finding indicates that overcoming the consequences associated with an algorithmic decision may be difficult, but it also suggests that there are opportunities to add algorithm-based assessment as one factor in decisions without risking perceptions of unfairness.

5. Study 4: Algorithmic performance reviews, reductionism, and fairness

In this study, we tested our predictions in a sample of undergraduate students. We assessed people’s fairness perceptions associated with the emerging practice of using algorithms to make performance review decisions (Greenfield, 2018), predicting that algorithms would be perceived as less fair than human evaluators. Additionally, we examined whether differences in fairness perceptions were driven by a serial mediation model of perceived quantification leading to perceived decontextualization. We randomized whether participants saw the quantification or decontextualization items first. Furthermore, because one possible explanation of the fairness effects in our previous studies could be a lack of transparency about the process (an algorithmic “black box” effect) (cf. Yeomans, Shah, Mullainathan, & Kleinberg, 2019), we gave participants a high degree of detail about the specific factors considered.

5.1. Sample and procedure

A sample of 197 undergraduate students (45.7% female, $M_{age} = 20.9, SD = 1.8$, 41.6% Caucasian, 1% African American, 45.2% Asian/Asian American, 9.6% Hispanic, 2.5% Other) from a large private university took part in a randomized 2-cell survey design (manager vs. algorithm) that asked them to read about a new performance review procedure:

Many sales departments within companies are adopting a new performance review practice in which they evaluate individual performance data over the previous year as well as provide an opportunity for employees to articulate what they saw as their biggest challenges and contributions throughout the year. For the latter, employees record brief videos of themselves explaining what they achieved, what challenges they overcame, and what they learned during the previous year. This gives employees the opportunity to convey their thoughts and, if needed, put their performance data in context.

Employees’ performance reviews are conducted by [their manager] [an algorithm]. The [manager] [algorithm] evaluates the following data: the employee’s total sales, customer experience survey results, duration of employment, amount of overtime worked, contributions to coworkers’ sales, and potential for strong performance in the

⁶ None of these effects showed moderation by gender or ethnicity (either Caucasian vs. non-Caucasian or URM vs. non-URM).

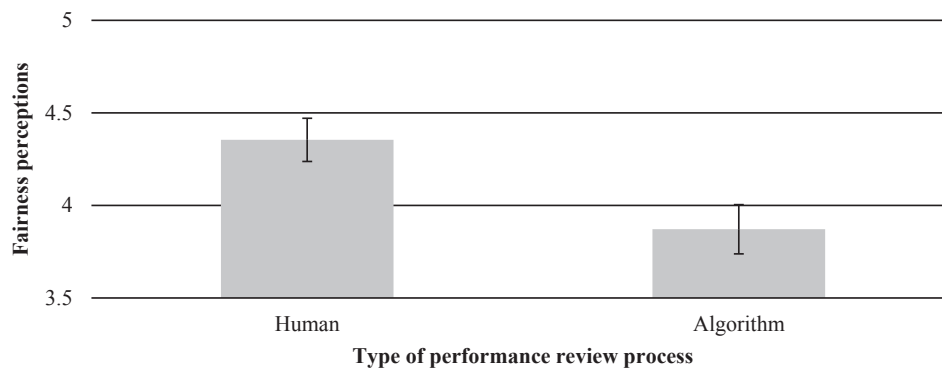


Fig. 6. Fairness of algorithmic vs. human performance review.

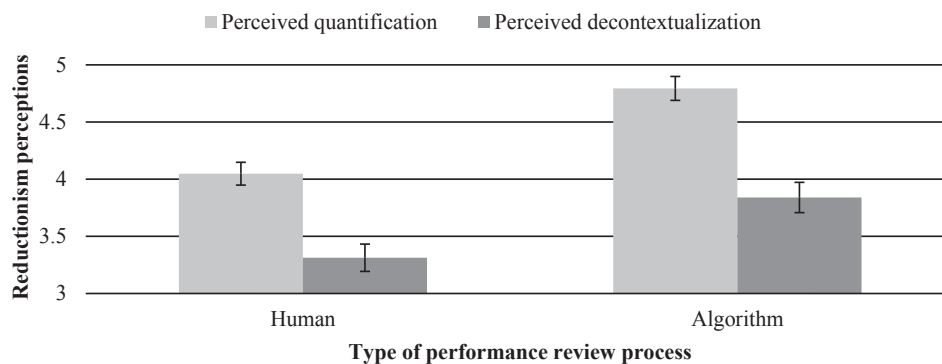


Fig. 7. Quantification and decontextualization, algorithm vs. human.

future. In addition, the [manager] [algorithm] is trained to evaluate the employee-recorded videos, paying specific attention to the employees' stated contributions, the lessons they say they have learned, and their nonverbal communication (i.e., facial expressions, mood, positivity, persuasiveness, and honesty).

After evaluating the data and the video recording, the [manager] [algorithm] comes to a final decision on the performance review, which could affect the employee's salary, bonus, eligibility for promotion, and, in some cases, dismissal.

5.2. Measures

5.2.1. Perceived fairness

On a 1–7 scale (Strongly Disagree to Strongly Agree), participants indicated their agreement with four (randomized) statements adapted from Conlon et al. (2004): (1) the way this process determines which workers receive raises and pay cuts seems fair, (2) this process for making personnel decisions is fair, (3) this process regarding which workers are promoted and laid off is fair, and (4) the outcomes of this process are fair.

5.2.2. Quantification

Participants were asked to imagine themselves as an employee going through the review process and to indicate their agreement or disagreement with three statements (presented in a randomized order) on a 1–7 scale ("Strongly disagree" to "Strongly agree"): (1) I feel like this evaluation process would just reduce me to a number, (2) This evaluation process would adequately recognize my qualitative attributes, abilities, and performance, and (3) I think some information about my performance would be lost in this evaluation process.

5.2.3. Decontextualization

Participants were asked to imagine themselves as an employee going through the review process and to indicate their agreement or

disagreement with three statements (presented in a randomized order) on a 1–7 scale ("Strongly disagree" to "Strongly agree"): (1) This evaluation process would be able to put these factors into context when evaluating me, (2) This evaluation process appropriately takes into account how all the factors fit together, and (3) This evaluation process would look at each factor in light of the others to paint a complete picture of my performance.

5.3. Results

Results supported the prediction that algorithms would be considered less fair than managers (Hypothesis 1). Analysis of variance revealed that performance reviews conducted by algorithms ($M = 3.87$, $SD = 1.34$) were perceived as significantly less fair than those conducted by human managers ($M = 4.35$, $SD = 1.14$), ($F(1, 195) = 7.39$, $d = 0.39$, $p = .007$).⁷ See Fig. 6.

We next tested Hypothesis 2, which predicted that algorithms would be perceived as less fair because of their perceived reduction of accurate information. This prediction was supported. Analysis of variance revealed (1) that algorithms ($M = 4.80$, $SD = 1.06$) were perceived to quantify employees significantly more than human managers ($M = 4.03$, $SD = 0.98$), ($F(1, 195) = 27.53$, $d = 0.75$, $p < .001$), and (2) that algorithms ($M = 3.85$, $SD = 1.34$) were perceived to decontextualize information significantly more than human managers ($M = 3.30$, $SD = 1.17$), ($F(1, 195) = 9.75$, $d = 0.44$, $p = .002$). See Fig. 7.

To dig deeper into the relationship between our causal mechanisms, we conducted our preregistered analysis to assess whether the main effect of algorithms on perceived fairness was serially mediated by perceived quantification and perceived decontextualization. We used the Hayes (2017) procedure for testing serial mediation (Model 6) and

⁷ This effect was not moderated by gender or ethnicity (either Caucasian vs. non-Caucasian or URM vs. non-URM).

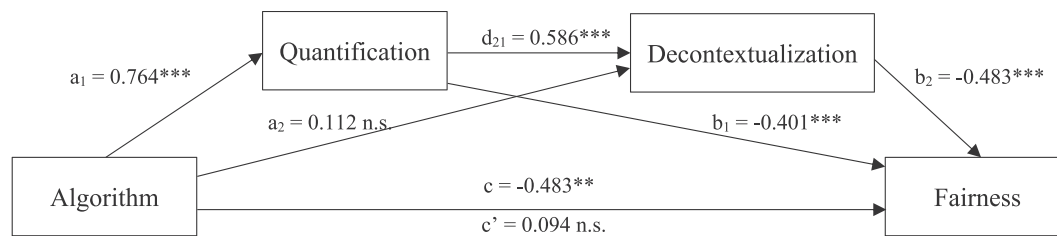


Fig. 8. Algorithmic reductionism and fairness, serial mediation model.

found support for this prediction. The indirect effect from Algorithm to quantification to decontextualization was significant (confidence interval: -0.49 to -0.16), the indirect effect from Algorithm to quantification to decontextualization to Fairness was significant (confidence interval: -0.35 to -0.11), and the indirect effect from Algorithm to decontextualization to Fairness was nonsignificant (confidence interval: -0.22 to 0.10), indicating that the (more direct) effect from Algorithm to decontextualization to Fairness is not significant unless the model accounts for the mediating (more indirect) effect of quantification. See Fig. 8.

5.4. Discussion

These results provide further support for Hypotheses 1 and 2. Additionally, they provide a fuller picture of how reductionism works, by showing that the effect of algorithms on the perceived fairness of decision processes is serially mediated by perceived quantification causing perceived decontextualization. This offers nuance to our theory of reductionism, demonstrating how apparent constraints both on qualitative information and on how information is put together work to influence perceptions of fairness. Furthermore, because this study gave participants greater specificity regarding the factors entering into the decision-making process, we can be more confident that the effect of algorithms on fairness is not merely the consequence of insufficient transparency.

6. Study 5: Algorithmic interview evaluations, transparency, and fairness

In this study, we tested our predictions in a sample of undergraduate students, using the context of video interview evaluations, based on the emergence of companies like HireVue, which have developed algorithms to assess recordings of job candidates for clients including Goldman Sachs and JP Morgan (Institute, 2018). We assessed people's fairness perceptions associated with using algorithms (versus humans) to evaluate video interviews, predicting that algorithms would be perceived as less fair. Additionally, we varied the degree of transparency surrounding the process by manipulating the level of detail, so as to directly test the possibility that algorithms might be seen as fair if only they were not construed as "black boxes." In particular, we created a scenario (i.e., being evaluated via a brief video recording) that was likely to be perceived as particularly unfair when done under conditions of low transparency. By including low- and high-transparency conditions, we were able to assess participants' beliefs about the capacity of both humans and algorithms to engage in holistic decision making. We predicted that increasing transparency about evaluation factors would increase perceived fairness, but only when the evaluator was human. Such a finding would suggest that perceived unfairness of algorithms is not simply due to people's assumptions that they will not consider the right information, but also due to the assumption that algorithms cannot even make use of the right information when they have it.

6.1. Sample and procedure

A sample of 213 undergraduate students (41.3% female, $M_{age} = 19.9$, $SD = 2.1$, 39.9% Caucasian, 5.6% African American,

37.6% Asian/Asian American, 9.4% Hispanic, 0.5% Native American, 7.0% Other) from a large private university took part in a 2×2 laboratory experiment (algorithm vs. human crossed with low vs. high transparency) in which they were incentivized to prepare for 5-minute video interviews:

Many companies these days are switching from a traditional interview process for employee hiring to a video interview process in which job candidates submit recordings of themselves speaking about their qualifications for the position.

We are creating a list of high performing video interviewers⁸ to recommend for a series of interview and job opportunities in the future.

After filling out the remainder of this questionnaire, you will be taken to another part of the lab to make a 5-minute recording of yourself answering some of the most common video interview questions.

On the next page, we varied the algorithmic or human nature of the video interview evaluation process. Additionally, participants in the high transparency condition saw an additional paragraph of text presented below in italics:

We are partnering with a university-affiliated startup that uses [a proprietary algorithm] [HR professionals] to evaluate video interviews. This process has been developed in conjunction with some of the leading industrial-organizational psychologists, working to better the hiring practices of various firms.

[The algorithm] [someone] will evaluate your recorded responses to a few questions and give you a score. Top scores will be set aside on a short list eligible for future interviews and opportunities.

The [algorithm] [evaluator] is trained to analyze the content of your responses (specifically: quality of past accomplishments, insightful thoughts and ideas, professional vocabulary) and your non-verbal cues (specifically: facial expressions, posture, tone of voice). Content and non-verbals will be weighted equally.

Click the button to reveal the interview questions.

Participants then read that they had 8 minutes to take notes on scratch paper to prepare for recording themselves answering the following video interview questions: (1) What can you tell us about yourself? (2) What are two accomplishments of which you are particularly proud? (3) What is one area where you are in need of improvement?

After the preparation time elapsed, participants answered some questions about the video interview evaluation process (perceived decontextualization and perceived fairness) before being taken to the video recording lab, at which point they were told that technical difficulties had caused a delay in recording the previous group, and they were dismissed from the study. No video interviews were in fact recorded.

⁸ "Interviewers" here was a typographical error (instead of "interviewees"). However, in light of the subsequent description of scoring, the use of "interviewees" in a fairness item, and the lack of any indication from participants of confusion about their roles, we believe that participants accurately understood the context.

6.2. Measures

6.2.1. Perceived fairness

On a 1–7 scale (Strongly Disagree to Strongly Agree), participants indicated their agreement with four (randomized) statements adapted from Conlon et al. (2004): (1) the way this process determines which candidates receive job opportunities seems fair, (2) this process for making interview performance decisions is fair, (3) this process regarding how interviewees are evaluated is fair, and (4) the outcomes of this process are fair.

6.2.2. Decontextualization

Participants were asked to consider the video interview evaluation process and to indicate their agreement or disagreement with three statements (presented in a randomized order) on a 1–7 scale (“Strongly disagree” to “Strongly agree”): (1) This evaluation process will be able to put my responses in context when evaluating me, (2) This evaluation process appropriately takes into account how all my responses fit together, and (3) This evaluation process will look at each response in light of the others to paint a complete picture of my interview performance.

6.3. Results

Analysis of variance revealed significant differences across all conditions both for perceived fairness ($F(3, 209) = 2.72, d = 0.49, p = .046$) and for perceived decontextualization ($F(3, 209) = 3.37, d = 0.53, p = .019$). These effects were driven by the human, high-transparency condition differing from the other three conditions for both fairness ($t(209) = 2.63, d = 0.42, p = .009$) and decontextualization ($t(209) = -3.16, d = 0.50, p = .002$). See Figs. 9 and 10.

Although no main effects were observed for either the algorithmic or transparency manipulation, two-step regression models found significant interactions between the algorithmic and transparency manipulations for both perceived fairness ($t = -2.41, \beta = -0.29, p = .017$) and perceived decontextualization ($t = 2.13, \beta = 0.25, p = .034$). Specifically, under conditions of high transparency, the algorithm condition yielded perceptions of significantly less fairness ($t(209) = -2.57, d = 0.50, p = .011$) and more decontextualization ($t(209) = 2.76, d = 0.53, p = .006$). Similarly, for human procedures, the high-transparency condition yielded perceptions of significantly more fairness ($t(209) = 2.33, d = 0.46, p = .021$) and less decontextualization ($t(209) = -2.58, d = 0.51, p = .01$). Additionally, the Hayes (2017) procedure for testing moderated mediation (model 7) revealed that transparency moderated the mediating effect of decontextualization between algorithms and perceived fairness (confidence interval for index of moderated mediation: -0.79 to -0.04).

6.4. Discussion

Because the student participants were incentivized with the possibility of future interview and job opportunities, this study created a

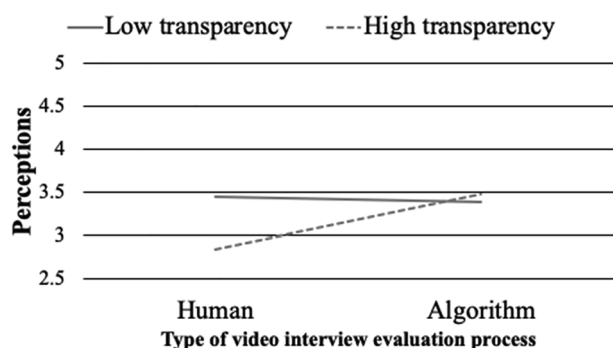


Fig. 9. Decontextualization.

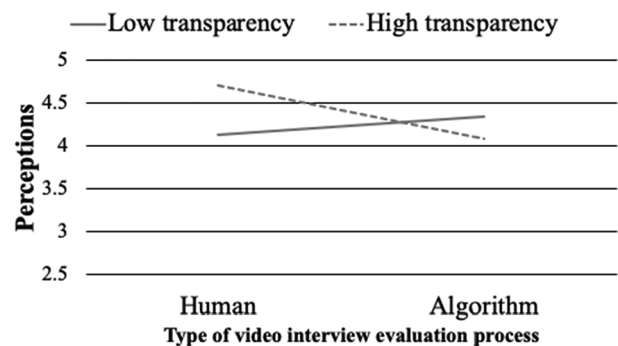


Fig. 10. Fairness.

situation involving greater stakes than our vignette-based experiments (Studies 1, 3, and 4). Under such stakes, combined with the fact that they would be assessed in a somewhat novel context (via video footage), increasing transparency in the human evaluation condition was seen as fairer than a more opaque human evaluation process. In contrast, however, increasing transparency failed to increase perceived fairness in the algorithmic condition. This result reinforces that of Study 2 (in which actual employees judged a possible change to HR policies specifying somewhat transparent factors), while also showing the moderating impact of transparency on human but not algorithmic procedures. These findings provide further support for our predictions that algorithms are seen as reductionistic and, therefore, less fair than humans.

7. General discussion

The results of five studies indicate that personnel decisions driven by algorithms are perceived to be less fair than identical decisions featuring more human involvement. Although it is possible (either now or in the future) that decision-making algorithms might objectively enhance procedural justice by eliminating biases to which human judgment is famously prone, they are subjectively perceived to violate procedural justice by reducing considerations to easily quantifiable performance data and failing to consider performance holistically.

Study 1 revealed that people tend to view algorithms as less fair than human decision makers in rendering decisions about both layoffs and promotions. The driving mechanism behind this effect was the perception that algorithms fail to adequately account for and contextualize qualitative attributes, such as leadership, attitude, and potential, which comprise part of each human being's unique personal character. Reductionism, whether real or imagined, is perceived to limit the use of accurate information on which fair procedures must rely. To be perceived as fair, then, treatment seems to require more than consistency and/or the elimination of bias—it necessitates holistic consideration of human characteristics, which calls for a less mechanistic and more human-oriented approach.

Study 2 further demonstrated these effects through a field experiment conducted in a large organization. In addition to corroborating the findings of Study 1, the results in Study 2 provided evidence that the use of algorithms to make people-related decisions may introduce attitudes and behaviors that ultimately have negative consequences for organizational performance. In particular, in our study it caused employees to express lower levels of organizational commitment, an effect that was driven by perceived unfairness. This indicates that the perceptions of unfairness associated with the use of algorithms to make personnel decisions matter in that they foster attitudes that could hinder performance and organizational loyalty.

Study 3 replicated and elaborated on our findings in a design that compared purely algorithmic and purely human decision processes to procedures combining algorithms with human judgment. Critically, human decision makers with the option to rely on an algorithm were

seen as comparably fair to the purely human process, whereas a default algorithmic process with the opportunity for a 10% human adjustment was perceived much like a pure algorithm. Mere opportunity for human agents to adjust algorithmic output was inadequate to dispel perceptions of reductionism and concomitant unfairness; fairness requires a high level of human discretion. These results highlight the central role of human involvement in fostering a climate of organizational justice.

Studies 4 and 5 demonstrated that increasing transparency about the factors being considered does not eliminate perceptions of unfairness surrounding algorithms. Additionally, they provided greater insight into the mechanism of algorithmic reductionism. In Study 4, a serial mediation model demonstrated that the effect of algorithms on perceived fairness was mediated by perceived quantification which, in turn, caused perceived decontextualization. Study 5 demonstrated that increasing transparency of factors considered leads to greater perceived fairness when humans are the decision makers, but not algorithms. This indicates that people do not believe that algorithms are as capable as humans at handling and contextualizing a range of qualitative and quantitative factors.

Crucially, our results indicate that algorithm-driven decision-making processes are seen to be unfair across a range of scenarios, including when outcome valence is positive or negative (Study 1), when qualitative factors are explicitly considered (Studies 2, 4, and 5), when humans have the ability to tweak the outcomes (Study 3), and when transparency is high (Studies 4 and 5). Across all studies, our findings indicate that people are not comfortable with an organizational paradigm that picks winners and losers through an automated software system. Although human decision makers may fall victim to bias, they are not viewed as treating employees reductionistically and are therefore preferable to resigning one's professional fate to an algorithm.

7.1. Theoretical contributions

The present findings offer a number of important theoretical contributions. First, our work extends procedural justice theory by examining new and emerging decision-making tools that are poised to be increasingly used in organizations. Specifically, we investigated employees' reactions to algorithms used to make personnel decisions. In so doing, we contribute to existing theory by showing the importance of human involvement, which is perceived as improving the use of accurate information. Consistent with this idea, across multiple studies we found that participants viewed human involvement as increasing fairness by affording greater consideration of qualitative information as well as using this information for increased contextualization.

Second, we add illumination to the fundamental psychology of algorithms by focusing on the reactions of people both observing and affected by algorithm-driven procedures. Whereas prior work has concentrated on decision makers' tendency to accept or reject algorithms as a basis for decision making, or on the strategic relationship between algorithms and organizational performance outcomes, we address the need to understand the impact of algorithms on individual attitudes and behavior. Our theoretical lens and the empirical evidence have enabled us to make a conceptually important observation about the human psychology of algorithms. Specifically, we contend that people simply do not believe that algorithms apply to qualitative factors. People are prepared to accept that a competent human decision maker can consider both quantitative and qualitative factors. However, because they see algorithms as inescapably associated with reductionism, they resist the notion that algorithms are capable of fairly accounting for qualitative human attributes.

Third, our findings show the theoretical importance to procedural justice of a distinction between quantitative versus qualitative factors. To be perceived as fair, procedures must give due consideration to all aspects of the individual, including those not easily quantified, such as leadership potential, intentions, and psychological commitment, to name a few. In fact, these kinds of factors may be seen as fundamentally more associated with accurate information about the individual as a whole. Importantly, this insight extends beyond automated algorithms to human managers and practices that reduce people to numerical representations via overreliance on analytics in general. Broadly, analytics is the enterprise of organizing and utilizing data in more sophisticated ways (Angrave, Charwood, Kirkpatrick, Lawrence, & Stuart, 2016). In our context, analytics involves "the extensive use of data, statistical and quantitative analysis, explanatory and predictive models, and fact-based management to drive decisions and actions" (Davenport & Harris, 2007, p. 7). As we found, HR teams relying on pure quantification are also likely to face charges of unfairness, possibly despite other efforts to maximize procedural and distributive justice.

Fourth, our findings have theoretical implications for research investigating the psychological and organizational determinants of organizational commitment. It is not surprising that perceptions of fairness are associated with commitment. However, our findings in Study 2 indicate that the mere possibility of introducing algorithms into an organization can lead to negative downstream consequences for organizational commitment. Given the escalating reliance on people analytics in organizations combined with the important consequences of perceived fairness, researchers should seek to better understand how to integrate analytic decision making while minimizing damage to morale and organizational commitment.

7.2. Limitations and future directions

Although the present findings make clear contributions to theory and practice, there are several limitations and remaining questions that should be addressed by future research. For example, we have focused on accuracy, through reductionism, as the mechanism by which algorithms affect the perceived fairness of decision procedures. However, there are other possible factors that could be in play, such as respect. In particular, Tyler and Blader (2003) argue that respect is likely to play a central role in fairness considerations when group identity is strong. It is worth noting that our findings persisted under conditions of both strong and weak group identity. For example, group identity is (presumably) high in Study 2, which involved employees of a single organization, but it is low in the remainder of our studies, in which participants took on the role of interviewees (Study 5) or third-party observers (Studies 1, 3, and 4) and were unlikely to feel strong social connection to the organizations in question. As such, in the majority of our studies, participants' concerns are more plausibly related to the apparent inaccuracy of the procedures than to feelings of disrespect. Nevertheless, respect is a factor worth investigating in future work on algorithmic decision making, along with other elements of procedural fairness, including perceptions of voice, bias, and consistency, in order to provide a more complete picture of the relationship between the use of algorithms and perceived fairness.

Future work on procedural justice should also further examine distinctions between quantitative and qualitative factors, including both the consequences of considering one or the other as well as determinants that lead employees to view a factor as quantitative or qualitative in the first place. While we found that people perceived human decision makers to consider both quantitative and qualitative factors (and perceived algorithms to

consider primarily quantitative factors), additional work needs to be done to assess the malleability of these perceptions as well as how procedures might balance the weighting of quantitative and qualitative factors to achieve optimal fairness.

Future work might also illuminate the boundary conditions of our effects, looking to different industries (which may vary in terms of their preference for quantification) or different objects of evaluation (such as plant and equipment, the disposition of which by analytics-based algorithms might well be seen as fair). Relatedly, future research should investigate effects associated with the author of the algorithm. Given our findings, it appears that people psychologically remove humans from the equation when considering algorithms, but further work manipulating awareness of who is writing the analytics software that compiles and implements the data might intensify or ameliorate the effects we observed. Additional research might investigate reactions to analytics when already institutionalized within an organizational or industrial context. Finally, we think it would be interesting for future work to investigate explicit messaging and information regarding consideration of demographic factors such as race and gender. If, for instance, an algorithm is explicitly framed as ignoring race (or, alternatively, assessing race and assigning different weights to particular categories), this may moderate perceptions of the fairness of the algorithm.

Finally, future work could examine the effect of having subjective or objective control over the use of analytics on perceived fairness. For example, although having the opportunity to exert a measure of control over algorithmic decisions may make decision makers more willing to rely on them (Dietvorst et al., 2018), it remains unclear whether such control has a similar effect on decision makers' judgments of fairness. It is possible that subordinates and supervisors alike might share an instinctive reaction to people analytics as essentially unjust compared to managers and HR teams making the exact same decisions, but this question should be answered by further research.

7.3. Managerial implications

The present findings also introduce several practical contributions. First, they indicate that employees are likely to reject people analytics as unfair. Yet, at the same time, there are a number of reasons companies may wish to use analytics-based decision procedures. It may be possible for managers to roll out algorithms in a manner allowing for positive responses from employees. Framing techniques could highlight how relying on analytics is objectively more, rather than less, fair than relying on human decision makers. Organizations could also pair algorithms with a human decision maker who assesses qualitative factors; this combination of processes might be perceived as the fairest of them all.

Second, our findings have practical implications for organizational commitment. In particular, drawing on the theory that justice has important downstream consequences for organizationally relevant attitudes and behaviors (Lind, 2001), we demonstrate that the adoption of algorithms has a negative impact on key organizational outcomes. Additionally, we take view of the long term, in which reliance on analytics may become common in organizational settings. As

algorithms become institutionalized within organizations, employee reactions to them may change, with natural implications for managerial decision making and implementation.

Our findings highlight the potentially unfavorable downstream implications of people analytics for organizational life. Processes perceived to be reductionistic, and consequently unfair, dampen employees' commitment to the organization, with potentially devastating consequences for retention, recruitment, productivity, and organizational citizenship behaviors. It is possible that increasing reliance on data analytics undermines a very important organizational function: leaders getting together to determine organizational values and priorities in light of discussions about whom to promote, hire, and let go. Hence organizations in search of efficiency gains from computerizing personnel decisions should carefully consider whether the cure might be worse than the disease. On the flip side, our findings indicate that organizations wishing to tap the positive potential of analytics ethically must begin by using algorithms as but one tool within a people-driven decision process.

8. Conclusion

Perceived unfairness notwithstanding, algorithms continue to gain increasing influence in human affairs, not only in organizational settings but throughout our social and personal lives. How this influence plays out against our sense of fairness remains to be seen but should undoubtedly be of central interest to justice scholars in the years ahead. Will the compilers of analytics and writers of algorithms adapt their methods to comport with intuitive notions of morality? Or will our understanding of fairness adjust to the changing times, becoming inured to dehumanization in an ever more impersonal world? Questions such as these will be asked more and more frequently as technology reshapes modes of interaction and organization that have held sway for generations. We have sought to contribute answers to these questions, and we hope that our work will encourage others to continue studying these and related topics.

CRedit authorship contribution statement

David T. Newman: Conceptualization, Methodology, Validation, Formal analysis, Investigation, Resources, Data curation, Writing - original draft, Writing - review & editing, Visualization, Project administration. **Nathanael J. Fast:** Conceptualization, Methodology, Resources, Writing - review & editing, Supervision. **Derek J. Harmon:** Conceptualization, Methodology, Writing - review & editing.

Acknowledgement

The second author was supported in part by the Air Force Office of Scientific Research, under grant FA9550-18-1-0182. The content does not necessarily reflect the position or the policy of any Government, and no official endorsement should be inferred.

Appendix A. Organizational charts used as materials in Study 2

See [Charts 1–4](#).

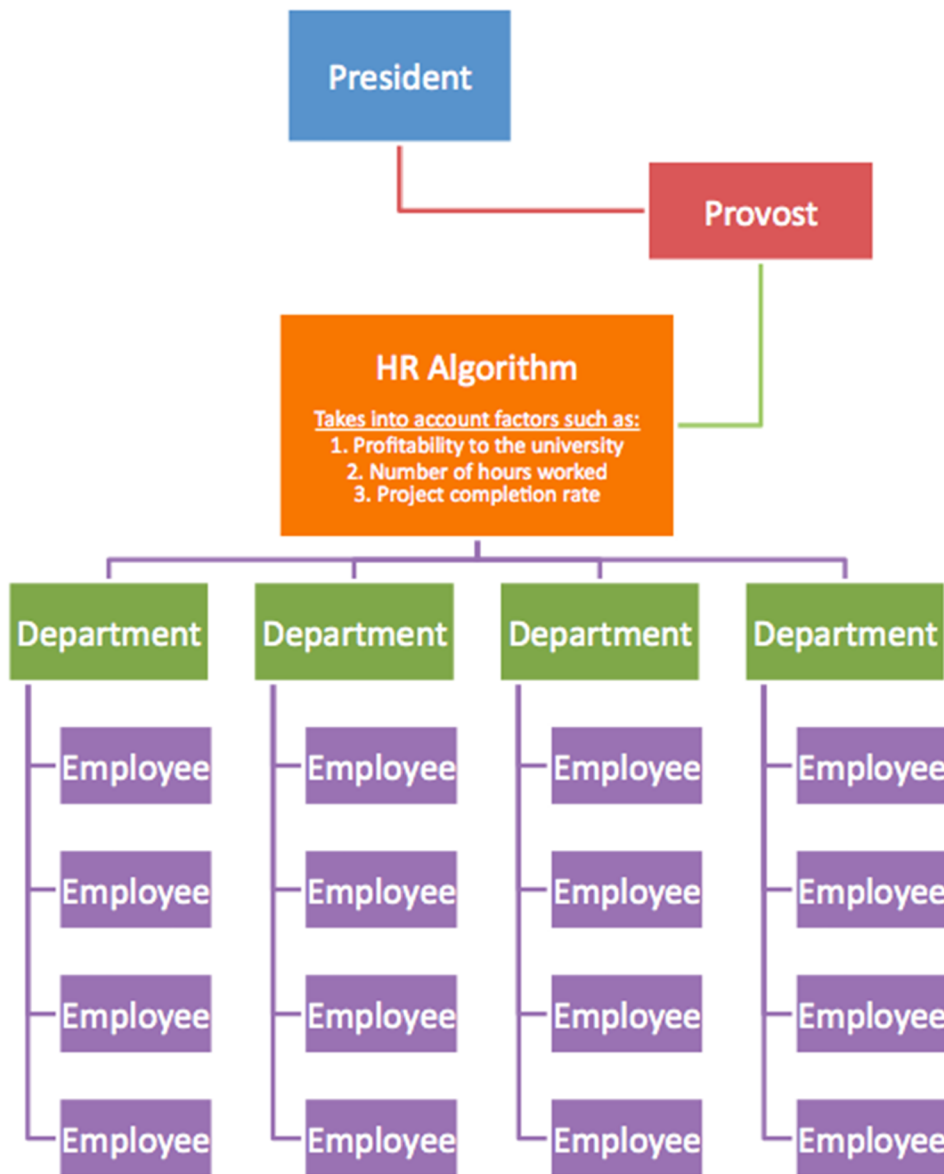


Chart 1. Algorithmic process emphasizing quantitative factors.

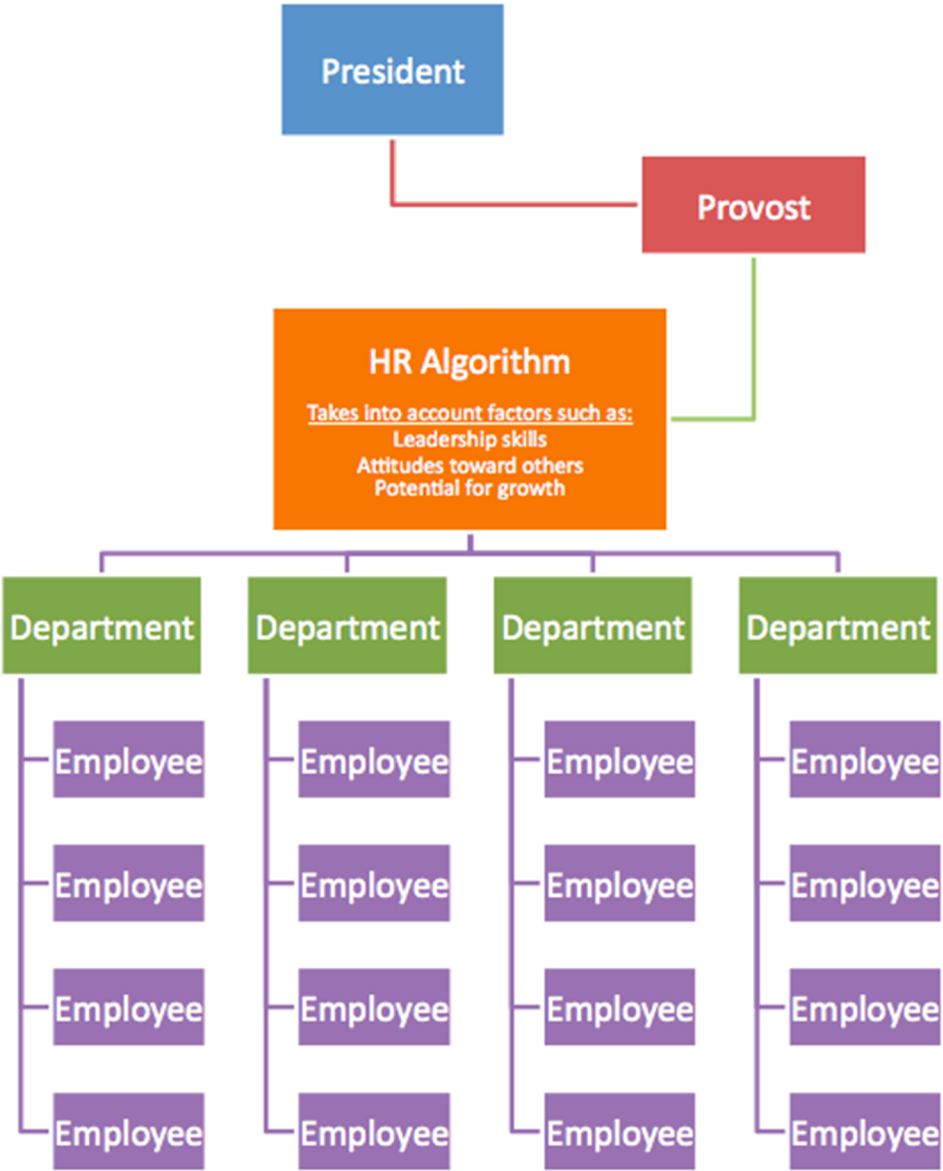


Chart 2. Algorithmic process emphasizing qualitative factors.

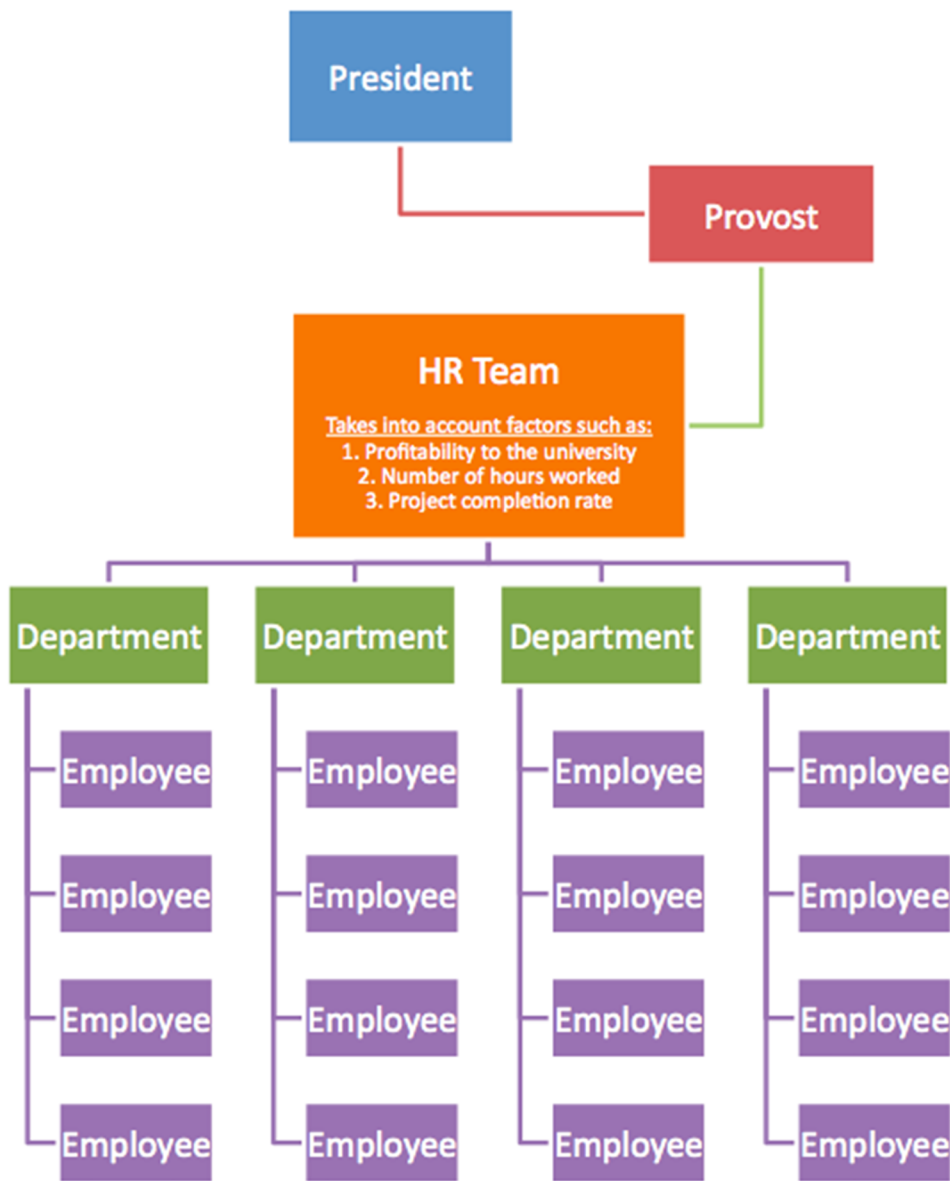


Chart 3. Human process emphasizing quantitative factors.

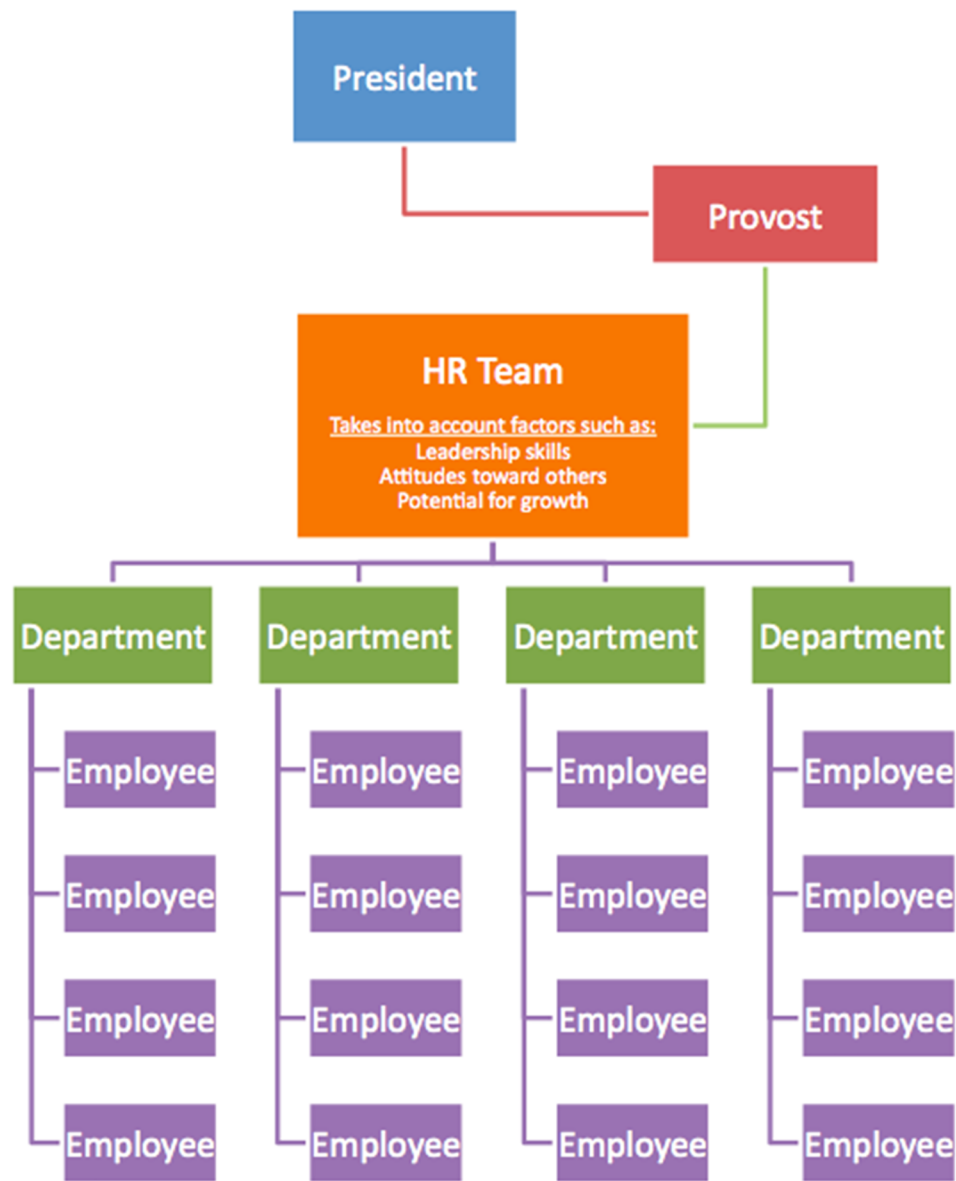


Chart 4. Human process emphasizing qualitative factors.

Appendix B. Supplementary material

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.obhdp.2020.03.008>.

References

- Adams, J. S. (1965). Inequity in social exchange. In L. Berkowitz (Vol. Ed.), *Advances in experimental social psychology: Vol. 2*, (pp. 267–299). New York: Academic Press.
- Allen, N. J., & Meyer, J. P. (1990). The measurement and antecedents of affective, continuance and normative commitment to the organization. *Journal of Occupational Psychology*, 63, 1–18.
- Angrave, D., Charlwood, A., Kirkpatrick, I., Lawrence, M., & Stuart, M. (2016). HR and analytics: Why HR is set to fail the big data challenge. *Human Resource Management Journal*, 26, 1–11.
- Aral, S., Brynjolfsson, E., & Wu, L. (2012). Three-way complementarities: Performance pay, human resource analytics, and information technology. *Management Science*, 58, 913–931.
- Arkes, H. R., Dawes, R. M., & Christensen, C. (1986). Factors influencing the use of a decision rule in a probabilistic task. *Organizational Behavior and Human Decision Processes*, 37, 93–110.
- Yeomans, M., Shah, A., Mullainathan, S., & Kleinberg, J. (2019). Making sense of recommendations. *Journal of Behavioral Decision Making*, 32, 403–414.
- Associated Press (2015). Like Amazon, more firms use data on workers to drive productivity. August 18. Retrieved from <http://www.nbcnews.com>.
- Barrett-Howard, E., & Tyler, T. R. (1986). Procedural justice as a criterion in allocation decisions. *Journal of Personality and Social Psychology*, 50, 296–304.
- Bies, R. J., & Moag, J. F. (1986). Interactional justice: Communication criteria of fairness. In R. J. Lewicki, B. H. Sheppard, & M. H. Bazerman (Vol. Eds.), *Research on negotiations in organizations: Vol. 1*, (pp. 43–55). Greenwich, CT: JAI Press.
- Boswell, W. R., Roehling, M. V., LePine, M. A., & Moynihan, L. M. (2003). Individual job-choice decisions and the impact of job attributes and recruitment practices: A longitudinal field study. *Human Resource Management*, 42, 23–37.
- Brockner, J. (2002). Making sense of procedural fairness: How high procedural fairness can reduce or heighten the influence of outcome favorability. *Academy of Management Review*, 27, 58–76.
- Brockner, J. (2006). Why it's so hard to be fair. *Harvard Business Review*, 84, 122–129.
- Brockner, J., & Wiesenfeld, B. M. (1996). An integrative framework for explaining reactions to decisions: Interactive effects of outcomes and procedures. *Psychological Bulletin*, 120, 189–208.
- Choi, I., Koo, M., & Choi, J. A. (2007). Individual differences in analytic versus holistic thinking. *Personality and Social Psychology Bulletin*, 33, 691–705.
- Christin, A. (2017). Algorithms in practice: Comparing web journalism and criminal justice. *Big Data & Society*, 4, 1–14.

- Colquitt, J. A. (2001). On the dimensionality of organizational justice: A construct validation of a measure. *Journal of Applied Psychology*, 86, 386–400.
- Colquitt, J. A., Conlon, D. E., Wesson, M. J., Porter, C. O. L. H., & Ng, K. Y. (2001). Justice at the millennium: A meta-analytic review of 25 years of organizational justice research. *Journal of Applied Psychology*, 86, 425–445.
- Colquitt, J. A., & Rodell, J. B. (2015). Measuring justice and fairness. In R. Cropanzano, & M. L. Ambrose (Vol. Eds.), *The Oxford handbook of justice in the workplace: Vol. 1*, (pp. 187–202). New York: Oxford University Press.
- Colquitt, J. A., & Zipay, K. P. (2015). Justice, fairness, and employee reactions. *Annual Review of Organizational Psychology and Organizational Behavior*, 2, 75–99.
- Conlon, D. E., Porter, C. O. L. H., & Parks, J. M. (2004). The fairness of decision rules. *Journal of Management*, 30, 329–349.
- Corporate Finance Institute (2018). HireVue Interview Guide. Corporate Finance Institute. Retrieved from <https://corporatefinanceinstitute.com/resources/careers/interviews/about-hirevue-interview/> on June 30, 2019.
- Cowgill, B. (2017). Automating judgement and decision-making: Theory and evidence from résumé screening. Columbia University, 2015 Empirical Management Conference.
- Crawford, K. (2013). The hidden biases of big data. *Harvard Business Review*. April 1. Retrieved from <https://blogs.hbr.org>.
- Dana, J., Dawes, R., & Peterson, N. (2013). Belief in the unstructured interview: The persistence of an illusion. *Judgment and Decision Making*, 8, 512–520.
- Davenport, T. H., & Harris, J. G. (2007). *Competing on analytics: The new science of winning*. Boston, MA: Harvard Business Press.
- Davenport, T. H., Harris, J. G., & Shapiro, J. (2010). Competing on talent analytics. *Harvard Business Review*, 88, 52–58.
- Dawes, R. M. (1979). The robust beauty of improper linear models in decision making. *American Psychologist*, 34, 571–582.
- Dawes, R. M., Faust, D., & Meehl, P. E. (1989). Clinical versus actuarial judgment. *Science*, 243, 1668–1774.
- Dietvorst, B. J., Simmons, J. P., & Massey, C. (2014). 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*.
- Fecheyr-Lippens, B., Schaninger, B., & Tanner, K. (2015). Power to the new people analytics. *McKinsey Quarterly*, 51, 61–63.
- George, G., Haas, M. R., & Pentland, A. (2014). Big data and management. *Academy of Management Journal*, 57, 321–326.
- Gilovich, T., Griffin, D., & Kahneman, D. (Eds.). (2002). *Heuristics and biases: The psychology of intuitive judgment*. Cambridge, UK: Cambridge University Press.
- Greenfield, R. (2018). Your raise is now based on next year's performance. *Bloomberg*. July 9. Retrieved from <https://www.bloomberg.com>.
- 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, 293–323.
- Hastie, R., & Dawes, R. M. (2010). *Rational choice in an uncertain world: The psychology of judgment and decision making*. Thousand Oaks, CA: SAGE Publications.
- Hayes, A. F. (2017). *Introduction to mediation, moderation, and conditional process analysis: A regression-based approach*. New York: The Guilford Press.
- Highhouse, S. (2008). Stubborn reliance on intuition and subjectivity in employee selection. *Industrial and Organizational Psychology*, 1, 333–342.
- Highhouse, S., Nye, C. D., & Zhang, D. C. (2019). Dark motives and elective use of brainteaser interview questions. *Applied Psychology*, 68, 311–340.
- Isaac, W., & Dixon, A. (2017). Why big-data analysis of police activity is inherently biased. *The Conversation*. May 9. Retrieved from <https://theconversation.com>.
- Jackson, M. C. (2003). *Systems thinking: Creative holism for managers*. Chichester, UK: John Wiley & Sons.
- Konovsky, M. A., & Folger, R. (2006). The effects of procedures, social accounts, and benefits level on victims' layoff reactions. *Journal of Applied Social Psychology*, 21, 630–650.
- Kuncel, N. R., Klieger, D. M., Connelly, B. S., & Ones, D. S. (2013). Mechanical versus clinical data combination in selection and admissions decisions: A meta-analysis. *Journal of Applied Psychology*, 98, 1060–1072.
- Lavelle, J. J., Brockner, J., Konovsky, M. A., Price, K. H., Henley, A. B., Taneja, A., & Vinekar, V. (2009). Commitment, procedural fairness, and organizational citizenship behavior: A multifoci analysis. *Journal of Organizational Behavior*, 30, 337–357.
- Lecher, C. (2019). How Amazon automatically tracks and fires warehouse workers for 'productivity.' *The Verge*. April 25. Retrieved from <https://www.theverge.com>.
- Leventhal, G. S. (1976). The distribution of rewards and resources in groups and organizations. In L. Berkowitz, & W. Walster (Vol. Eds.), *Advances in experimental social psychology: Vol. 9*, (pp. 91–131). New York: Academic Press.
- Leventhal, G. S. (1980). What should be done with equity theory? New approaches to the study of fairness in social relationships. In K. Gergen, M. Greenberg, & R. Willis (Vol. Eds.), *Social exchange: Advances in theory and research: Vol. 1*, (pp. 27–55). New York: Plenum Press.
- Lind, E. A. (2001). Fairness heuristic theory: Justice judgments as pivotal cognitions in organizational relations. In J. Greenberg, & R. Cropanzano (Vol. Eds.), *Advances in Organizational Justice: Vol 1*, (pp. 56–88). Stanford, CA: Stanford University Press.
- Lind, E. A., Kanfer, R., & Earley, P. C. (1990). Voice, control, and procedural justice: Instrumental and noninstrumental concerns in fairness judgments. *Journal of Personality and Social Psychology*, 59, 952–959.
- Lind, E. A., & Tyler, T. R. (1988). *The social psychology of procedural justice*. New York: Plenum Press.
- Loehr, A. (2015). Big data for HR: Can predictive analytics help decrease discrimination in the workplace? *The Huffington Post*. March 23. Retrieved from <https://www.huffingtonpost.com>.
- Logg, J. M., Minson, J. A., & Moore, D. A. (2019). Algorithm appreciation: People prefer algorithmic to human judgment. *Organizational Behavior and Human Decision Processes*, 151, 90–103.
- McFarlin, D. B., & Sweeney, P. D. (1992). Distributive and procedural justice as predictors of satisfaction with personal and organizational outcomes. *Academy of Management Journal*, 35, 626–637.
- Meehl, P. E. (1954). *Clinical versus statistical prediction: A theoretical analysis and a review of the evidence*. Minneapolis: University of Minnesota.
- Nisbett, R. E., Peng, K., Choi, I., & Norenzayan, A. (2001). Culture and systems of thought: Holistic versus analytic cognition. *Psychological Review*, 108, 291–310.
- O'Connor, S. (2016). When your boss is an algorithm. *Financial Times*. September 7. Retrieved from <https://www.ft.com>.
- Preacher, K. J., & Hayes, A. F. (2008). Asymptotic and resampling strategies for assessing and comparing indirect effects in multiple mediator models. *Behavior Research Methods*, 40, 879–891.
- Rosch, E. (1975). Cognitive representations of semantic categories. *Journal of Experimental Psychology: General*, 104, 192–233.
- Scott, A., Klein, F. K., & Onovakpuri, U. (2017). *Tech leavers study: A first-of-its-kind analysis of why people voluntarily left jobs in tech*. Kapor Center for Social Impact.
- Setty, P. (2014). HR meets science at Google with Prasad Setty. re: Work with Google. November 10. Retrieved from <https://www.youtube.com/watch?v=KY8v-O5Buyc&feature=youtu.be&t=2m58s>.
- Shapiro, D. L., Buttner, E. H., & Barry, B. (1994). Explanations: What factors enhance their perceived adequacy? *Organizational Behavior and Human Decision Processes*, 58, 346–368.
- Sheppard, B. H., & Lewicki, R. J. (1987). Toward general principles of managerial fairness. *Social Justice Research*, 1, 161–176.
- Skarlicki, D. P., Ellard, J. H., & Kelln, B. R. C. (1998). Third-party perceptions of a layoff: Procedural, derogation, and retributive aspects of justice. *Journal of Applied Psychology*, 83, 119–127.
- Spreitzer, G. M., Cameron, L., & Garrett, L. (2017). Alternative work arrangements: Two images of the new world of work. *Annual Review of Organizational Psychology and Organizational Behavior*, 4, 473–499.
- Sweeney, P. D., & McFarlin, D. B. (1993). Workers' evaluation of the "ends" and the "means": An examination of four models of distributive and procedural justice. *Organizational Behavior and Human Decision Processes*, 55, 23–40.
- Taylor, M. S., Tracy, K. B., Renard, M. K., Harrison, J. K., & Carroll, S. J. (1995). Due process in performance appraisal: A quasi-experiment in procedural justice. *Administrative Science Quarterly*, 40, 495–523.
- Thibaut, J., & Walker, L. (1975). *Procedural justice: A psychological analysis*. Hillsdale, NJ: Erlbaum Press.
- Tyler, T. R., & Bies, R. J. (1990). Beyond formal procedures: The interpersonal context of procedural justice. In J. S. Carroll (Vol. Ed.), *Applied social psychology and organizational settings: Vol. 6*, (pp. 77–98).
- Tyler, T. R., & Blader, S. L. (2003). The group engagement model: Procedural justice, social identity, and cooperative behavior. *Personality and Social Psychology Review*, 7, 349–361.
- Von Bertalanffy, L. (1972). The history and status of general systems theory. *Academy of Management Journal*, 15, 407–426.
- Waber, B. (2013). *People analytics: How social sensing technology will transform business and what it tells us about the future of work*. Upper Saddle River, NJ: FT Press.
- Weaver, G. R., & Trevino, L. K. (2001). The role of human resources in ethics/compliance management: A fairness perspective. *Human Resource Management Review*, 11, 113–134.
- Wiesenfeld, B. M., Brockner, J., & Thibault, V. (2000). Procedural fairness, managers' self-esteem, and managerial behaviors following a layoff. *Organizational Behavior and Human Decision Processes*, 83, 1–32.
- Wilson, H. J., Alter, A., & Shukla, P. (2016). Companies are reimagining business processes with algorithms. *Harvard Business Review*. February 8. Retrieved from <https://hbr.org>.