Can the Use of Data Analytics Tools Lead to Discriminatory Decisions?

Sepideh Ebrahimi
School of Administrative Studies
York University, Toronto, ON, Canada
sepidehe@yorku.ca

Khaled Hassanein
DeGroote School of Business
McMaster University, Hamilton, ON, Canada
hassank@mcmaster.ca

Abstract

Data Analytics (DA) has been criticized for contributing to discriminatory decisions in organizations. To date, several studies have investigated reasons for the generation of discriminatory recommendations by DA tools and how to ameliorate the issue. Nonetheless, recent studies by researchers, practitioners, and government agencies show that despite the progress made, the issue has not been eliminated. As a result, it is crucial for DA users to be vigilant about the danger of discriminatory recommendations generated by DA tools. This study represents an effort to provide empirical evidence about whether and to what extent decision makers will readily accept a discriminatory DA recommendation and about the cognition and attitudes that are associated with this behavior. The results obtained from an empirical study confirms that a majority of users readily accepted a discriminatory recommendation and sheds light on what factors influence this acceptance.

1. Introduction

The last decade has witnessed a widespread adoption of computers, smartphones, and in general Internet-connected devices by organizations and consumers. As a result, an ever-increasing amount of data is being generated. Organizations are increasingly adopting data analytics (DA) tools to derive insights from analyzing the data collected to discover patterns in support of their decision-making [42] and to make data-driven decisions [16]. Data analytics is often a combination of a number of processes and tools, including SQL queries, statistical analysis, data mining, fact clustering, and data visualization and is a way to discover customer segments, associate similar and related products, etc. [40]. Various benefits have been ascribed to using DA for making better decisions leading to favorable outcomes such as higher financial and strategic performance in organizations [7]. Nonetheless, use of such tools to support managers' decision-making has raised some major social and ethical concerns including issues associated with privacy, control, and discrimination [13, 31].

Discrimination takes place when member(s) of a socially defined group due to their membership of that group are treated differently (especially unfairly) [22]. Contrary to most of the studies on discriminatory decision-making, discriminatory decisions arising from DA recommendations are not necessarily made due to the prejudicial beliefs of the decision maker. Even fair decision makers can make a discriminatory decision drawing on a discriminatory recommendation generated by the DA tools they use to support their decision-making.

It is noteworthy that recommendations that treat a demographic class less favorably than other class(es) are considered *potentially* discriminatory. However, according to civil rights legislations (e.g., Title VII of the Civil Rights Act in the United States), these recommendations are *actually* discriminatory only when there is no legitimate business need to explain the discrepancy [47]. For instance, when making a recommendation about hiring salespeople, if a DA tool only/mainly puts forth names of male applicants, the recommendation is considered as potentially discriminatory. It is, then, required to investigate whether or not the recommendation is indeed discriminatory. If it is found out that there are legitimate business necessities behind such a situation (e.g., when driving a truck is a requirement for the job and fewer females have a license to drive trucks), then it can be concluded that the recommendation is not discriminatory against females. Similarly, throughout this paper, the notion of discriminatory recommendation refers to a recommendation that is *potentially discriminatory* and needs to be investigated further.

Discriminatory recommendations can be generated by DA tools due to biased or non-representative data and/or inadvertent modeling procedures in the DA tools [48]. Calders and Žliobaitė [5] suggest that there are three main reasons for the generation of discriminatory recommendations by unbiased algorithms. First, relations between non-sensitive and sensitive attributes in data that lead to non-sensitive attributes acting as proxies for sensitive variables. For instance,

HICSS

when zip codes are related to race, making a recommendation that relies on zip codes can generate outcomes that are different for people from different races. The second reason is data labeling. Historical data is used to build and train DA models and therefore, discriminatory data can bring about discriminatory models [9]. The third reason is flaws in the data collection process that lead to some groups of individuals being over- or under-represented in the data set such as when there are higher instances of offenders from a certain ethnic minority in the police database. If such a data set is used to train a model, it is likely to learn that a strong correlation exists between ethnicity and crime [33].

To date some technical methods have been suggested to discover and remove discrimination in data mining procedures [e.g., 11, 21, 36]. However developing computational methods that eliminate such discrimination altogether is the subject of an ongoing endeavor [48]. As such, managers who bear the responsibility for the decisions made in organizations need to be vigilant about whether discrimination exists in a recommendation put forth by a DA tool. Unfortunately it has been suggested that mostly in making decisions understanding the causes and consequences of particular patterns are neglected and finding significant connections is considered as sufficient [30]. Therefore, as the first research question in this study, we investigate *whether and to what extent do data analytics users readily accept a discriminatory recommendation generated by a data analytics tool?*

On the one hand, readily acceptance of recommendations put forth by DA tools, as Newell and Marabelli [31] suggest is due to the fact that in an organization few individuals actually understand the algorithms, what has been included in them and why [31]. Consequently, it is difficult to investigate and discern whether a DA recommendation includes a discrimination against a protected group (e.g., females) or not. Furthermore, due to the sheer volume of data being processed, decision makers are reliant on those tools to analyze the data and to support their decision-making. On the other hand, in the ethics literature it has been repeatedly suggested and shown that recognizing that there is a moral aspect to the issue at hand is required for an individual to commit an ethical behavior [e.g., 14, 37]. Therefore, it seems important to investigate whether recognizing that there is an ethical issue at hand makes a difference in terms of users' acceptance of a DA discriminatory recommendation. Therefore, the second research question in this study looks into *whether and to what extent do DA users' recognition of the moral aspect of the issue at hand is different between individuals who accept a discriminatory recommendation and individuals who reject a discriminatory recommendation?*

A number of characteristics of a moral issue increase the likelihood that an individual will recognize its moral aspect. Jones [20] suggests that "the extent of issue-related moral imperative in a situation" are determined by six elements: magnitude of consequences, social consensus, probability of effect, temporal immediacy, proximity, and concentration of effect. He suggests that these elements determine the moral intensity of a situation and therefore, impact an individual's recognition of the moral issue. In line with Jones' suggestion and following several studies, which confirmed the impact of moral intensity of an issue on individual's moral recognition [27, 43], this study also seeks to investigate *whether and to what extent do elements of moral intensity impact DA users' recognition of the moral issue when they are presented with a discriminatory DA recommendation?*

Next, we discuss the theoretical background of this study and then turn to the theoretical underpinning of the proposed hypotheses. Research methodology in support of data collection and pertinent analyses is presented in the fourth section and the results of data analyses are provided in section 5. Subsequently discussions of the results and contributions to theory and practice are discussed in sections 6 and 7 respectively.

2. Theoretical Background

To respond to the research questions outlined above, this study draws upon the four-component model of ethical decision-making and the literature on moral intensity, which will be discussed next.

2.1. The four-component model of ethical decision-making

Rest's [37] four-component model of ethical decision making is undoubtedly one of the most prevalent models in the ethics and business ethics literatures. Rest argues that, during the course of making a decision involving an ethical dimension, individuals move through a series of four sequentially ordered steps, namely, recognition of the moral issue, making a moral judgment, establishing the intent to act morally, and engaging in a moral behavior. The first step, recognition of the moral issue, also known as moral awareness, is an interpretive process in which the individual recognizes that a moral problem exists in a situation or that a moral principle is relevant to the existing set of circumstances [38].

Recognition of the moral issue, then, prompts the decision maker to make a judgment of what potential action is most moral. Moral intent is prioritizing moral values over other values and finally moral behavior is the application of the moral intent to the situation [8]. It is important to note that recognition of a moral issue plays a pivotal role in the process of making an ethical decision as without recognizing the moral issue the process might not be triggered at all as a person who fails to recognize the moral aspect of an issue will fail to employ ethical decision-making schemata and will make the decision based on other schemata such as economic factors, etc. [20].

2.2. Moral intensity

The literature on moral intensity has shown the positive impact of the moral intensity of an issue on its recognition by individuals [e.g., 4, 23, 43]. The moral intensity of an issue is comprised of six factors [20]. First, *magnitude of consequences* suggests that the issue will be more serious if its sum of harms done to the victims is higher. Second, *social consensus* states that a higher level of social agreement that a proposed behavior is unethical makes the issue more intense. Third, *probability of effect* suggests that an issue will be more intense if it has a higher likelihood to occur and to cause the anticipated harm. Fourth, *temporal immediacy* suggests that an issue with a shorter interval between when the decision is made and when the consequences occur is perceived as being more intense. Fifth, *proximity* states that the feeling of closeness that the decision-maker has for victims makes the issue more intense. Finally, *concentration of effect* suggests that an issue is perceived as being more intense if the consequences affects fewer individuals as opposed to the same consequences being more broadly distributed [20].

3. Hypotheses development

To respond to the above research questions, this study draws upon the above theoretical foundations, to propose the eight hypotheses, detailed below:

3.1. Accepting discriminatory Data Analytics' recommendations

Algorithms (e.g., Google search algorithm) are often multi-component systems built by teams and therefore, include some level of opacity that even the programmers who are insiders to the algorithms' development must deal with [41]. In the case of machine learning algorithms, the opacity is even

higher since the internal decision logic of the algorithm is altered as it learns on training data [3]. As a result, in organizations, few individuals actually understand the algorithms included in data analytics tools [31]. Therefore, often finding a strong predictive association by an algorithm is seen as sufficient and finding out the reasons for those associations in the data from different sources are neglected [31]. The fact that barely anyone knows how data analytics recommendations are generated can potentially lead to the readily acceptance of the recommendations put forth by these tools even when such recommendations are discriminatory. Therefore:

H1: When presented with a discriminatory DA recommendation, the proportion of DA users accepting it will be higher than the proportion rejecting it.

3.2. Recognition of the moral issue

Ethical reasoning has been described as a systematic framework that involves making principled assessment in questionable situations [14, 37]. Individuals engage in ethical behavior after they realize the situation at hand has an ethical aspect to it. Recognition of the moral issue is specifically important as not all moral issues are obvious [45]. Many studies have found significant relationships between recognition of a moral issue and engaging in an ethical behavior [For a review, see 8, 26, 34]. Therefore, it is reasonable to expect that users who become morally aware of a potentially discriminatory recommendation of a DA system are more likely to not accept that recommendation compared to users who are not aware. Hence,

H2: Participants who accept a discriminatory DA recommendation will exhibit a lower level of recognition of the moral issue than those who reject the recommendation.

3.3. Perceived moral intensity

As discussed previously, the moral intensity of an issue is comprised of six issue-contingent factors.

3.3.1. Magnitude of Consequences. This dimension of moral intensity suggests that the higher an individual perceives the sum of the resulting harms of an unethical behavior, the higher will be their perception of the moral intensity of the issue. For instance, an action that leads to death of one person is of higher magnitude in terms of consequences compared to an action that causes a minor injury to one individual [20]. Similarly, in the context of a DA

discriminatory recommendation, a healthcare-related decision is more likely than a marketing-related decision to be perceived as having a higher magnitude of consequences and consequently to induce in DA users recognition of the ethical aspect of the issue. The positive relationship between perceived magnitude of consequences and moral attitudes has been shown in previous studies [e.g., 27]. Therefore,

H3: In the context of a discriminatory DA recommendation, the greater the perceived magnitude of the negative consequences, the more likely that a DA user will recognize the moral issue at hand.

3.3.2. Social Consensus. An individual may not know the right behavior in a situation. Social consensus can help reduce the ambiguity in such circumstances [20]. In the context of DA use in organizations for decision-making purposes, users need to have both analytical skills as well as domain knowledge about the company to be able to make better use of DA tools [16]. If a user does not hold enough knowledge and skills to scrutinize a DA recommendation, they are less likely to realize whether or not it is discriminatory. It has been suggested that few users actually understands the logic and processes of algorithms included in DA tools. Therefore, they tend to mainly rely on the outputs generated by these systems without actively scrutinizing them [30, 31]. Such an issue can be more pronounced in organizations with weaker ethical cultures [26], where there is low social consensus about issues like discriminatory decisions and their ensuing harm. In such environments, a DA user is less likely to recognize that there can be an ethical aspect to the issue of accepting a discriminatory recommendation proposed by a DA tool. Similarly, Moberg and Caldwell [29] empirically show individuals' exposure to an organizational ethical culture to be strongly associated with their level of moral imagination (i.e., a process of considering the ethical elements of a decision thoroughly). As such we posit that,

H4: In the context of a discriminatory DA recommendation, the greater the perceived social consensus that accepting the recommendation is harmful, the more likely that a DA user will recognize the moral issue at hand.

3.3.3. Probability of effect. This dimension of moral intensity is "a joint function of the probability that the act in question will actually take place and the act in question will actually cause the harm predicted" [20]. In the context of a discriminatory DA recommendation, the probability of effect refers to a user's perception of the likelihood that the recommendation includes discrimination and that putting the recommendation into effect would actually cause some harm. A few reasons can lead to a DA user perceiving a low likelihood for the recommendation being discriminatory. For instance, previous studies have shown that having a high level of trust in a decision support system can lead to having high levels of trust in the system's advice [10]. High levels of trust in the system's advice can in turn bring about a lower perceived likelihood of it being discriminatory. This is further exacerbated if the user reasons that the likelihood of the negative consequences associated with accepting a potentially discriminatory recommendation is low. Since the intensity of a moral situation would be discounted in such circumstances [43], the user is less likely to recognize the moral aspect of the issue at hand. As such,

H5: In the context of a discriminatory DA recommendation, the greater the perceived probability of negative consequences, the more likely that a DA user will recognize the moral issue at hand.

3.3.4. Temporal immediacy. Temporal immediacy of an issue, defined as the elapsed time between the present and the time when the consequences of a moral act in question will take place is an important dimension in determining the intensity of a moral issue [20]. This is due to the fact that people tend to discount the probability and the impact of events that happen in the future [25]. In the context of decision making using DA tools, the results of accepting a discriminatory DA recommendation is less likely to happen in the immediate future or even shortly after the decision is made. Consequently, drawing on the literature on moral intensity, it is not very likely that the ethical aspect of the issue at hand will be apparent to the DA user. In light of the above discussion, we hypothesize that,

H6: In the context of a discriminatory DA recommendation, the greater the perceived temporal immediacy of negative consequences, the more likely that a DA user will recognize the moral issue at hand.

3.3.5. Proximity. The level of proximity that a decision maker feels toward the victims of a harmful decision positively impacts their perceived moral intensity of the issue and consequently the likelihood of their recognition of the moral aspect of the issue at hand. The notion of proximity is especially important in the context of using data analytics tools for

decision-making purposes due to two main reasons. First, DA tools tend to treat individuals as a set of records, each with a number of attributes and, therefore, dehumanizes them [12]. Second, computers in general and more specifically DA tools distance the decision maker from the subjects of their decisions. The further the perceived distance of the victims of the action to the decision maker, the less intense will be the moral issue in her/his mind. In such a case, the decision maker is less likely to recognize that there is an ethical aspect to the issue they are dealing with. Therefore,

H7: In the context of a discriminatory DA recommendation, the greater the perceived proximity toward the subjects of the recommendation, the more likely that a DA user will recognize the moral issue at hand.

3.3.6. Concentration of effect. Concentration of effect is defined as "an inverse function of the number of people affected by an act of given magnitude" [20]. The concept of concentration of effect is in line with the philosophy of ethical utilitarianism, which holds that "an act is right only if it produces for all people a greater balance of good consequences over bad consequences than other available alternatives (i.e., 'the greatest good for the greatest number')" [19]. Therefore, an act that is "bad" for a few people has a higher concentration of negative effect and consequently has a higher moral intensity than another act that is "bad" for a large number of people [43]. The notion of concentration of effect is even more important in the context of DA tools that are mainly used to deal with complexities associated with analyzing "big data". As in such cases, it can be expected that the number of individuals whose data are analyzed and those who could consequently be victims of discriminatory recommendations of these tools are quite high. In such circumstances, the concentration of effect tends to be low and as a result, the recognition of the moral aspect of the issue at hand is likely to also be low. Thus, we posit that:

H8: In the context of a discriminatory DA recommendation, the greater the perceived concentration of negative effects, the more likely that a DA user will recognize the moral issue at hand.

4. Research methodology

The hypotheses proposed in the present study were tested through an empirical study, where participants used a fictitious data analytics tool, designed for this research, which provided them with a discriminatory recommendation.

4.1. Generating discriminatory DA recommendations

A fictitious experimental DA tool was developed that included 200 records of individuals who work in the sales department of an organization. The aim of the analysis was to generate a list of 20 individuals to be sent to a training program on effective leadership in a sales organization. To generate the list, the system drew on various objective (education level and years of working experience at the company), and subjective factors (average of performance evaluation over the last 3 years and potential of the employee). Participants were told that the subjective factors had been provided by employees' previous/current managers.

The recommended sample of employees to be sent to the training program included discrimination against women (the proportion of female individuals in the recommended sample was considerably reduced compared to its level in the full data set; 15% versus 44%). The discriminatory recommendation was generated following the literature that suggests that when labeling the data (e.g., defining a good employee), if one or several of the defining variables are subjective, they might bring in the personal prejudice of previous/current managers into the analysis process [2, 5]. In our experiment, we simulated that the prejudice of previous/current managers toward females had led to their receiving lower performance evaluations as well as lower evaluation of potential of the employee compared to their male counterparts. Since the recommendation of our DA tool took into account these two variables, the recommendation included discrimination against females[1].

4.2. Experimental procedures

Participants for this study were told that an organization wanted to send 20 of its employees to a one-week training program on Effective Leadership in a Sales Organization. To help participants realize the importance of the training, they were told that "attending the training will bring about great experience for the selected employees. In addition, they will be more likely to receive promotions in the future". The participants' task was to use the system and evaluate the tool's recommendations. After reading a document about instructions to do the

---

[1] It is important to note that the specific cause of the discrimination is not important for the purposes of our experiment. What is important is that the output provided by the DA tool in this case is potentially discriminatory.

study, participants were presented with a dataset of the 200 sales employees that included their ID, name, and the four objective and subjective variables described above. It is noteworthy that name of each employee indicated their sex. In addition, all names were selected from one race (i.e., white) to avoid any confounding effects.

In the next step, participants were provided with a description about each of the variables that were included in the data set. Subsequently, participants received the list of the variables that in the past have been most closely associated with success in a training program similar to the one in question. Next, participants clicked on a button that based on those variables, ran a pre-designed predictive model and generated recommendations. The recommendation that participants received was the same list of all 200 potential employees for this training with one additional attribute for each employee signifying the tool's recommendation as to whether the employee should be sent to this training (1) or not (0). In the next step, participants were asked to indicate their decision on whether or not they accept the recommendations of the DA tool. Next, participants filled out a questionnaire (discussed below). Finally, they were debriefed.

4.3. Measures

To ensure content validity, all measurement instruments were adapted from existing and validated scales. Recognition of the moral issue was measured using Reynolds' [38] 3-item scale. Moral intensity factors except for proximity was measured using the Singhapakdi, et al. [43] instrument. The measure for perceived proximity in Singhapakdi, et al. [43] is not relevant to the context of our study as they measured proximity with whether the action in question would be wrong if the decision maker was a personal friend of the victim. Thus in this study, perceived proximity was measured using an alternate 3-item scale from Barnett [1]. The impacts of gender, age, and impression management (defined as the propensity of respondents to "consciously over-report their performance of a wide variety of desirable behaviors and under-report undesirable behaviors" [35]) were also controlled for in this study.

5. Data analysis results

To test the proposed hypotheses, two main methodologies were employed: analysis of variance (ANOVA) and Multiple Regression, the results of which are presented below.

5.1. Participant background information

The sample for this study, recruited by a market research firm, consisted of 73 middle managers who had more than one employee reporting to them. This sample size would assure a sufficient statistical power of 0.80 to detect a medium to large effect size [39].

Subjects were recruited from various industries including but not limited to education, government, healthcare, real estate, information services and data processing, and finance. The subjects were employed in various departments such as human resources, research and development, accounting, sales, customer service, and IT. Among the subjects, 34 (47.2%) were female and 38 (52.8%) were male. The average age of the participants was 45.

5.2. Measurement characteristics

The reliability of both multi-item constructs, measured by Cronbach's α, were well above the cited minimum value of 0.7 ($\alpha_{recognition\ of\ the\ moral\ issue}=0.93$ and $\alpha_{proximity}=0.89$), indicating that each set of measurement items was consistent in what it intended to measure [32]. In addition, discriminant validity of our two factors was assessed using SmartPLS 3.0 by examining the item loadings as recommended in the literature [18, 28, 44]. In addition, we made sure that multicollinearity was not an issue in our analysis by examining inter-factor correlations as depicted in Table 1 as well as examining Variance Inflation Factors (VIFs) that were all below 2 [6].

Furthermore, to address the concern about common method bias, two techniques of Herman's one-factor test and unmeasured latent method construct [24] were employed and the results indicated a lack of a common method bias.

Table 1. Factors' correlations							
	Recog	MoC	SC	PoE	TI	Pr	CoE
Recog	-						
MoC	0.47	-					
SC	0.56	0.54	-				
PoE	0.59	0.47	0.53	-			
TI	0.35	0.36	0.27	0.48	-		
Pr	0.29	0.19	0.15	0.18	0.12	-	
CoE	0.49	0.63	0.47	0.60	0.60	0.10	-
- Recog: Recognition of the moral issue			- PoE: Probability of Effect				
- MoC: Magnitude of consequences			- TI: Temporal Immediacy				
- SC: Social Consensus			- Pr: Proximity				
			- CoE: Concentration of Effect				

5.3. Results of hypotheses testing

The first hypothesis states that the proportion of DA users accepting a DA discriminatory recommendation will be higher than the proportion rejecting it. Results showed that 73% of the participants accepted and 27% of them rejected the discriminatory recommendation, we can conclude that enough statistical evidence has been provided to support H1 (χ^2 (1) =14.9, ρ < 0.001).

The second hypothesis suggests that participants who accept a discriminatory DA recommendation have on average a lower level of recognition of the moral issue. An ANOVA test was performed employing IBM SPSS Statistics 24 to test this hypothesis and the results, reported in Table 2, provide significant support for H2.

Hypotheses 3 to 8 are about the impact of the dimensions of moral intensity on recognition of the moral issue. To test these hypotheses, multiple regression using IBM SPSS Statistics 24 was employed. The results of the multiple regression analysis depicted in Figure 1 indicate that three out of the six dimensions of moral intensity significantly impacted recognition of the moral issue. More specifically, social consensus (β=0.28; ρ<0.05), probability of effect (β=0.32; ρ<0.05), and proximity (β=0.21; ρ<0.05) positively influenced recognition of the moral issue supporting H4, H5, and H7. However, the impacts of magnitude of consequences (β=0.08; ρ>0.05), temporal immediacy (β=0.004; ρ>0.05), and concentration of effect (β=0.11; ρ>0.05) on recognition of the moral issue turned out to be insignificant. Thus H3, H6, and H8 are not supported.

Table 2. ANOVA summary table for recognition of the moral issue					
Dependent Variable	Sum of Squares	df	Mean Square	F	Sig.
Recognition of the moral issue	45.70	1	45.70	23.55	0.00

6. Discussion

Data Analytics tools are increasingly being used to make data-driven decisions. However, there have been societal concerns raised about the use of such tools as it is possible for these tools to generate discriminatory recommendations in certain circumstances [2, 31]. To date several researchers have investigated the technical aspects of this problem and suggested methods that can help decrease the likelihood of generating such recommendations. However, such endeavors have not been able to completely eliminate the issue of discriminatory recommendations being generated by DA tools [48]. As such and since in organizations it is the managers' responsibility to make sure that their decisions are free of discrimination, it is important to look at the human aspects of decision-making while using data analytics tools too.

Figure 1. Multiple Regression Results
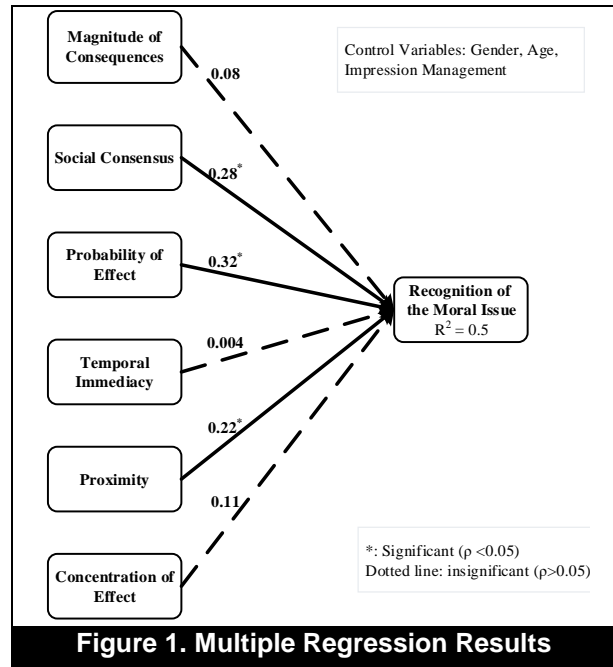
It has been suggested that due to complexity of DA tools hardly anyone in organizations understands what is included in algorithms and how they work [31]. As such, DA users tend to rely heavily on the outcomes generated by these tools without much understanding of the analyses performed to generate the results [30]. As the above discussion suggests, it is likely that DA users readily accept a discriminatory recommendation generated by a DA tools. This study provides empirical support for this argument as more than 70% of the participants in this study approved the discriminatory recommendation that was provided to them by the fictitious DA tool. In addition, in line with the business ethics literature [e.g., 15, 17] and our expectation, participants who accepted the recommendation had on average significantly lower levels of recognition of the moral issue at hand than those who rejected it.

This study investigated the impacts of the six dimensions of moral intensity on recognition of the moral issue in the context of a discriminatory DA recommendation. The results show that only 3 of these dimensions significantly increase DA users' recognition of the moral issue at hand.

The hypothesis about the positive impact of magnitude of consequences on recognition of the

moral issue was not supported. We believe that the reason might be the fact that the task that the participants were given was about selecting a number of individuals to be sent to a training program. While attending such a program can have positive impacts on one's career, the harms of not being sent to the program are not huge. Therefore, we believe that had we chosen a different task (e.g., hiring an employee) the results could have been different.

Similarly, we believe that the hypothesis about the relationship between temporal immediacy and recognition of the moral issue was not supported due to the task that the participants were given. Although in the descriptions of the task, it was written that attending the training will increase employees' chance of receiving a promotion in the future, it was not exactly specified when the training is going to take place and when the promotion decisions are going to be made. Thus, has such detail been provided, the results may have varied.

In addition, we believe that the insignificant outcome of the relationship between concentration of the effect and recognition of the moral issue might stem from the fact that participants were from different organizations with various sizes. As a result, participants' perceptions of concentration of effect might have been formed in relation to their current or past organizational sizes.

7. Contributions and limitations

This study stands to make several contributions to theory. First, to the best of our knowledge, this is the first empirical study that examines the issue of discriminatory decision making using DA tools and shows that a discriminatory DA recommendation is likely to be readily accepted by users. Previous research has only conceptually discussed this problem. Therefore, this study advances the literature on ethics of data analytics use by empirically showing the high incidence of this problem.

Second, this study shows that there is a statistically significant difference in terms of recognition of the moral aspect of the issue at hand between those users who accept the recommendation and those who reject it. Based on this outcome, researchers can further investigate tools and methods to increase the level of user's recognition of the moral issues in the case of discriminatory DA recommendations. Such tools and methods, if successful, can in turn reduce the likelihood of users' acceptance of a discriminatory DA recommendation.

Last but not least, drawing upon the literature on issue-contingent moral intensity, this study contributes to the literature on ethics of data analytics use by studying the antecedents of recognition of the moral issue in the context of discriminatory recommendations of these tools. The squared multiple correlation coefficient (R^2) of 0.5 indicates that dimensions of moral intensity explain fully half the variance in recognition of the moral issue. In addition, this study by identifying three strong antecedents of recognition of the moral issue, enables researchers to study the impact of methods that can enhance these particular antecedents. For instance, Watley and May [46] argue that by providing personal information about those impacted by a decision, perceived proximity and subsequently ethical behavioral intent would increase.

The results of this study have significant practical implications for organizations as well. This study empirically shows that DA users tend to readily rely on the recommendations provided by such tools. As such organizations should put in place various informational and training programs to alert users about the pitfalls of doing so. In addition, this study shows that DA users are less likely to accept a discriminatory recommendation if they recognize the moral issue relevant to the situation. Combining this result with arguments in the literature about the low level of users' technical skills related to DA tools provides practitioners with solid evidence that they need to invest in their employees' DA training if they are asked to make decisions with those tools.

In addition, this study sheds light on a few important antecedents of recognizing the moral aspects of the issue in the context of discriminatory recommendations of DA tools. Drawing on these results, organizations can employ methods to increase such recognition. For instance, social consensus about the harmful outcomes of approving a discriminatory DA recommendation can be increased by nurturing the culture of accountability as well as an overall ethical culture in the organization.

A number of limitations exist for this study that provide avenues for future research. First, participants for this study were selected from North American middle managers. Given the potential impacts of culture on users' attitude toward IT use as well as moral behaviors, caution should be exercised in generalizing the results of this study to DA users in other geographic regions. Second, the majority of moral intensity dimensions were measured with single-item instruments. Although these are the main instruments used in the business ethics literature for this purpose, it is fruitful to develop instruments with multiple items that are specifically geared toward the context of using IT to make organizational decisions. Third, this study only focuses on discrimination against one demographic class (i.e., females). Future

research is warranted to examine if the results of this study are generalizable to other demographic categories (e.g., age, race). Finally, it should be noted that the present study represents an early attempt to investigate the issue of users' readily accepting discriminatory recommendations generated by DA tools. As such, drawing on the literatures on DA and business ethics, the study aimed at identifying the main reasons that contribute to that issue. Future research is warranted to further investigate more concrete characteristics of tools, users, and organizations that influence recognition of the moral issue or its antecedents as identified in this work.

8. References

[1] T. Barnett, "Dimensions of moral intensity and ethical decision making: An empirical study," Journal of Applied Social Psychology, vol. 31, pp. 1038-1057, 2001.

[2] S. Barocas and A. D. Selbst, "Big data's disparate impact," California Law Review, vol. 104, pp. 671-732, 2016.

[3] J. Burrell, "How the machine 'thinks': Understanding opacity in machine learning algorithms," Big Data & Society, vol. 3, pp. 1-12, 2016.

[4] K. D. Butterfield, L. K. Trevin, and G. R. Weaver, "Moral awareness in business organizations: Influences of issue-related and social context factors," Human Relations, vol. 53, pp. 981-1018, 2000.

[5] T. Calders and I. Žliobaitė, "Why Unbiased Computational Processes Can Lead to Discriminative Decision Procedures," in Discrimination and Privacy in the Information Society: Data Mining and Profiling in Large Databases, B. Custers, T. Calders, B. Schermer, and T. Zarsky, Eds., ed Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 43-57.

[6] J. Cohen, P. Cohen, S. G. West, and L. S. Aiken, Applied multiple regression/correlation analysis for the behavioral sciences, 3rd ed. Hillsdale, NJ: Lawrence Erlbaum, 2003.

[7] N. Côrte-Real, T. Oliveira, and P. Ruivo, "Assessing business value of Big Data Analytics in European firms," Journal of Business Research, vol. 70, pp. 379-390, 2017.

[8] J. L. Craft, "A review of the empirical ethical decision-making literature: 2004–2011," Journal of Business Ethics, vol. 117, pp. 221-259, 2013.

[9] B. Custers, "Data dilemmas in the information society: Introduction and overview," in Discrimination and Privacy in the Information Society, B. Custers, T. Calders, B. Schermer, and T. Zarsky, Eds., ed Berlin Heidelberg: Springer-Verlag, 2013, pp. 3-26.

[10] P. A. Dabholkar and X. Sheng, "Consumer participation and the trust transference process in using online recommendation agents," Journal of Consumer Satisfaction, Dissatisfaction and Complaining Behavior, vol. 25, pp. 96-117, 2012.

[11] C. Dwork, M. Hardt, T. Pitassi, O. Reingold, and R. Zemel, "Fairness through awareness," in Proceedings of the 3rd Innovations in Theoretical Computer Science Conference, Cambridge, Massachusetts, 2012, pp. 214-226.

[12] S. Ebrahimi, M. Ghasemaghaei, and K. Hassanein, "Understanding the Role of Data Analytics in Driving Discriminatory Managerial Decisions," in 37th International Conference on Information Systems, Dublin, Ireland, 2016.

[13] S. Ebrahimi and K. Hassanein, "Demographic Transparency to Combat Data Analytics Discriminatory Recommendations," in 16th Annual Workshop on HCI Research in MIS, Seoul, South Korea, 2017.

[14] O. C. Ferrell, L. G. Gresham, and J. Fraedrich, "A synthesis of ethical decision models for marketing," Journal of Macromarketing, vol. 9, pp. 55-64, 1989.

[15] G. Fleischman and S. Valentine, "Professionals' tax liability assessments and ethical evaluations in an equitable relief innocent spouse case," Journal of Business Ethics, vol. 42, pp. 27-44, 2003.

[16] M. Ghasemaghaei, S. Ebrahimi, and K. Hassanein, "Data analytics competency for improving firm decision making performance," The Journal of Strategic Information Systems, vol. 27, pp. 101-113, 2018.

[17] R. Haines, M. D. Street, and D. Haines, "The influence of perceived importance of an ethical issue on moral judgment, moral obligation, and moral intent," Journal of Business Ethics, vol. 81, pp. 387-399, 2008.

[18] J. F. Hair, W. C. Black, B. J. Babin, and R. E. Anderson, Multivariate data analysis, 7th ed.: Prentice hall Upper Saddle River, NJ, 2010.

[19] S. D. Hunt and S. Vitell, "A general theory of marketing ethics," Journal of macromarketing, vol. 6, pp. 5-16, 1986.

[20] T. M. Jones, "Ethical decision making by individuals in organizations: An issue-contingent model," Academy of management review, vol. 16, pp. 366-395, 1991.

[21] F. Kamiran, T. Calders, and M. Pechenizkiy, "Discrimination aware decision tree learning," in 10th International Conference on Data Mining, Sydney, NSW, 2010, pp. 869-874.

[22] K. Kraiger, J. K. Ford, and E. Salas, "Application of cognitive, skill-based, and affective theories of learning outcomes to new methods of training evaluation," Journal of applied psychology, vol. 78, pp. 311-328, 1993.

[23] D. L. Leitsch, "Differences in the perceptions of moral intensity in the moral decision process: An empirical examination of accounting students," Journal of Business Ethics, vol. 53, pp. 313-323, 2004.

[24] H. Liang, N. Saraf, Q. Hu, and Y. Xue, "Assimilation of enterprise systems: the effect of institutional pressures and the mediating role of top management," MIS quarterly, vol. 31, pp. 59-87, 2007.

[25] N. Liberman and Y. Trope, "The role of feasibility and desirability considerations in near and distant future decisions: A test of temporal construal theory," Journal of personality and social psychology, vol. 75, pp. 5-18, 1998.

[26] T. W. Loe, L. Ferrell, and P. Mansfield, "A review of empirical studies assessing ethical decision making in business," Journal of business ethics, vol. 25, pp. 185-204, 2000.

[27] D. R. May and K. P. Pauli, "The role of moral intensity in ethical decision making a review and investigation of moral recognition, evaluation, and intention," Business & Society, vol. 41, pp. 84-117, 2002.

[28] L. S. Meyers, G. Gamst, and A. J. Guarino, Applied multivariate research: Design and interpretation. Thousand Oaks, CA: Sage publications, 2006.

[29] D. Moberg and D. F. Caldwell, "An exploratory investigation of the effect of ethical culture in activating moral imagination," Journal of Business Ethics, vol. 73, pp. 193-204, 2007.

[30] S. Newell and M. Marabelli, "The crowd and sensors era: opportunities and challenges for individuals, organizations, society, and researchers," in 35th International Conference on Information Systems, Auckland, NZ, 2014.

[31] S. Newell and M. Marabelli, "Strategic opportunities (and challenges) of algorithmic decision-making: A call for action on the long-term societal effects of 'datification'," The Journal of Strategic Information Systems, vol. 24, pp. 3-14, 2015.

[32] J. C. Nunnally and I. H. Bernstein, Psychometric theory. Ney York: McGraw-Hill, 1994.

[33] C. O'Neil, Weapons of math destruction: How big data increases inequality and threatens democracy. New York, NY: Crown Publishing Group, 2016.

[34] M. J. O'Fallon and K. D. Butterfield, "A review of the empirical ethical decision-making literature: 1996–2003," Journal of business ethics, vol. 59, pp. 375-413, 2005.

[35] D. L. Paulhus, "Measurement and control of response bias," in Measures of personality and social psychological attitudes, J. P. Robinson, P. R. Shaver, and L. S. Wrightsman, Eds., ed New York: Academic Press, 1991, pp. 17-59.

[36] D. Pedreshi, S. Ruggieri, and F. Turini, "Discrimination-aware data mining," in Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining, 2008, pp. 560-568.

[37] J. R. Rest, Moral development: Advances in research and theory. New York: Praeger, 1986.

[38] S. J. Reynolds, "Moral awareness and ethical predispositions: investigating the role of individual differences in the recognition of moral issues," Journal of Applied Psychology, vol. 91, pp. 233-243, 2006.

[39] J. L. Roldán and M. J. Sánchez-Franco, "Variance-based structural equation modeling: guidelines for using partial least squares," in Research Methodologies, Innovations and Philosophies in Software Systems Engineering and Information Systems, M. Mora, O. Gelman, A. Steenkamp, and M. Raisinghani, Eds., ed Hershey, PA: Information Science Reference, 2012, pp. 193-221.

[40] P. Russom. (2011, Big data analytics. TDWI Best Practices Report, Fourth Quarter. Available: http://techadvocategroup.com/wp-content/uploads/2012/12/TAG-BigDataAnalytics.pdf

[41] C. Sandvig, K. Hamilton, K. Karahalios, and C. Langbort, "Auditing algorithms: Research methods for detecting discrimination on internet platforms," in Data and discrimination: converting critical concerns into productive inquiry, Seattle, WA, 2014, pp. 1-23.

[42] W. Shang, Z. M. Jiang, H. Hemmati, B. Adams, A. E. Hassan, and P. Martin, "Assisting developers of big data analytics applications when deploying on hadoop clouds," in Proceedings of the 2013 International Conference on Software Engineering, San Francisco, CA, 2013, pp. 402-411.

[43] A. Singhapakdi, S. J. Vitell, and K. L. Kraft, "Moral intensity and ethical decision-making of marketing professionals," Journal of Business Research, vol. 36, pp. 245-255, 1996.

[44] D. Straub, M.-C. Boudreau, and D. Gefen, "Validation guidelines for IS positivist research," The Communications of the Association for Information Systems, vol. 13, pp. 380-427, 2004.

[45] L. K. Trevino and M. E. Brown, "Managing to be ethical: Debunking five business ethics myths," The Academy of Management Executive, vol. 18, pp. 69-83, 2004.

[46] L. D. Watley and D. R. May, "Enhancing moral intensity: The roles of personal and consequential information in ethical decision-making," Journal of Business Ethics, vol. 50, pp. 105-126, 2004.

[47] J. Yinger, "Evidence on discrimination in consumer markets," Journal of Economic perspectives, vol. 12, pp. 23-40, 1998.

[48] I. Žliobaitė and B. Custers, "Using sensitive personal data may be necessary for avoiding discrimination in data-driven decision models," Artificial Intelligence and Law, vol. 24, pp. 183-201, 2016.