Executive Summary

The objective of this project is to examine whether or not an instructor discriminates between his students in the classroom based on their gender and/or clothing type. Based on video footage of one class period, two observers rate each instance of a student-teacher interaction as either positive, negative or neutral/ambiguous. Data is available for 231 students, on the students' gender (male or female), clothing type (unisex, standard, or other), and the number of positive and negative interactions with the instructor. Ambiguous/neutral interactions are recorded as zero values.

Due to the presence of several zero values in the data provided by the client, several suitable statistical models are tested, such as the zero-inflated Poisson, zero-inflated negative binomial, negative binomial, Tobit, binomial and Multinomial logit models. In light of the limited information provided by the data and in the interest of parsimony and simplicity, the binomial and multinomial models are chosen for statistical analysis.

Treating each student-teacher interaction as a Bernoulli experiment, the binomial regression model finds that conditional on an interaction occurring, it is positive in nature with a probability of 92.4% for students in unisex clothing and 84.8% for students in non-unisex clothing. Interestingly, gender is not a significant factor in predicting discrimination. However, the model suffers from low explanatory power.

The multinomial model posits that for each student, there is a probability that, during the course of a class period, the student will have only positive, or only negative or some of both types of interactions. In contrast to the binomial model, the multinomial model finds gender (and not clothing) to be the significant factor in predicting the probability of falling in one of three aforementioned groups. Furthermore, both male and female students are more than 50% likely to have only positive interactions, and less than 50% likely to have only negative interactions with the instructor. Male students, however, are 20% more likely than female students to have both types of interactions, while female students are 20% more likely than male students to have only positive interactions with the instructor.

Ultimately, the data may explain associations or correlation between the variables provided; however, evidence for discrimination by the instructor is hard to establish. Finally, some recommendations about a better design and data collection are made that would extend the study, and potentially find statistically and practically significant evidence on discrimination in the classroom.

Introduction and Data Exploration

The client sought the consultants' expertise to help answer the following question: Does an instructor discriminate among his students based on their gender and/or clothing type? Data, provided by the client, is collected by two observers, who watch the video recording of one class period. The observers unanimously rate each interaction between individual students and teachers as either positive or negative. Ambiguous interactions are treated as 'no interaction' and recorded as either zero or as a missing value.

Upon the client's recommendation, all missing values are treated as zeroes; thereafter, the modified data set is explored in more detail. In total, there are 231 students in the class, out of whom 26 have had no interaction at all with the instructor. There are two categorical predictors, namely gender and clothing. Table **1** describes the distribution of students in the class by gender and type of clothing. Clearly, male and female students are evenly distributed, with each group accounting for half the population. As far as clothing type is concerned, nearly half of the students wear 'other' type of clothing, with 'unisex' and 'standard' making up the other half. Due to this uneven distribution of clothing types, it is suspected that 'clothing' as it is in the data, may not provide an appropriate classification category for students.

Table 1. Numbers of students in each category

Female	111	48.1%
Male	120	51.9%
Total	231	100%
Unisex	54	23.4%
Standard	72	31.2%
Other	105	45.4%
Total	231	100%

The students are further classified into six categories, each category representing a gender-clothing combination. For example, a male student wearing 'unisex' clothing is a category different from a male student wearing 'standard' clothing. The top panel of Figure 1 shows the distribution of the positive responses, while the bottom panel shows the distribution of the negative responses. Overall, there are more positive responses

than negative responses. Furthermore, male students wearing unisex clothing have the most number of positive responses. However, due to sparse data for negative responses, it is hard to say much about negative responses, making it necessary to explore the behavior of positive responses further.

Figure 1. The distribution of different responses vs. Clothing-Gender combinations

Figure 2 shows the mean proportion[1] of positive responses for the six combinations of gender and clothing. It is immediately clear that both male and female students wearing 'unisex' clothing have a higher rate of positive interactions with the instructor. Even more, the six combination categories are distinct, in that they differ from one another in the rate of positive interactions with the instructor. However, this difference is not substantial, which can be seen by taking a closer look at the scale of the Y-axis: its range is from 0.88 to 0.96, which indicates that the actual difference between groups is not obvious. This may indicate that the conclusions of the statistical analyses, which follow, should be interpreted cautiously.

[1] Proportion of positive responses = (Number of positive responses) / (Total number of responses of both kinds). The mean proportion of positive responses was calculated by taking the sample average value for the proportion of positive responses for each category.

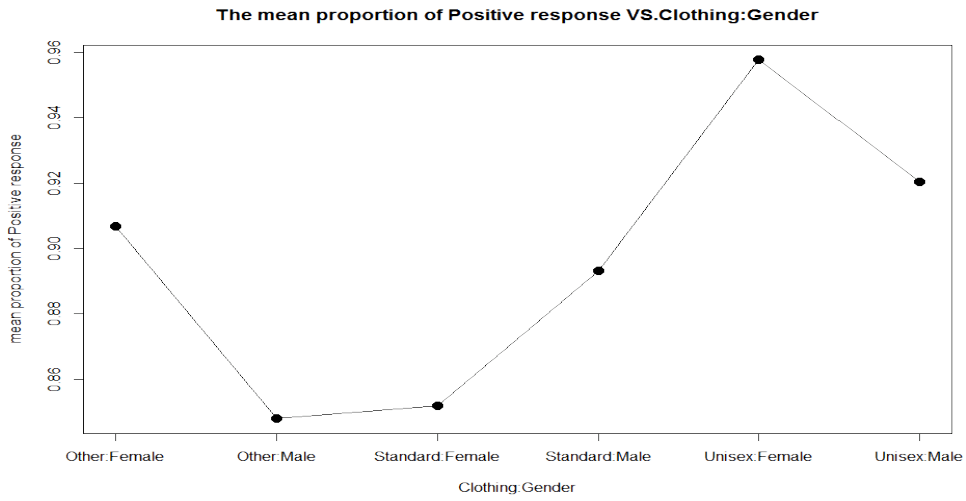

Figure 2. The mean proportion of Positive vs. Clothing-Gender combinations

Objective: With this backdrop, the task is to understand if certain gender and clothing type, or a combination of the two, lead to differential treatment by the instructor. In this sense, each instructor-student interaction may be treated as a unit of observation; then statistical modeling methods may be used to see whether certain gender-clothing combinations affect the likelihood of getting either type of interaction.

Methodology and Results

1. *Binomial Model:* A Binomial model can be used to estimate the probability of the occurrence of an event that has only two possible outcomes. With regard to the present scenario, instructor-student interactions are the 'event', and the two possible outcomes are 'positive interaction' and 'negative interaction'. Consequently, the question can be restated in the following way: given a student's gender and clothing, what is the probability that an interaction with the instructor, when it occurs, is positive rather than negative?

   Results: Based on the analysis, 'unisex' clothing is statistically significant in determining the probability of positive and negative interactions. In other words, the difference in estimated probability of getting a positive/negative interaction between 'unisex' and 'non-unisex' clothing is significant enough for the model to detect. More specifically, students wearing 'unisex' clothing have a 92.7% probability of having a positive interaction compared to students wearing non-unisex clothing, who have an 84.4% probability for the same. Conversely, students wearing 'unisex' type of clothing have 7.3% probability of having negative interaction with the instructor,

compared to students wearing non-unisex clothing, who have 15.2% probability for the same.

The other variable of interest, 'gender', was not statistically significant, implying that being a male or a female does not affect the probability of having either type of interaction.

2. *Multinomial Logit Model:* Alternatively, suppose a student could describe the class experience as good, bad or mixed. These three categories ('somePos', 'someNeg' and 'Both') roughly correspond to having only positive interactions, only negative interactions and both positive and negative interactions. If the probability of falling into one of these three categories differs greatly among students with different genders and clothing styles, that would be evidence of potential discrimination by the instructor. The statistical model that describes this situation is the *multinomial* model.[2] It assumes that the probabilities of each category can be described by the equation below. In this case, the "base" group is set to be the 'Both' group.

$$log\left(\frac{p_{Alternative\ Category}}{p_{Base\ Category}}\right) = \beta_0 + \beta_1 x$$

Results: First, it is found that the gender variable is sufficient to predict the probability of falling into one of the three interaction categories. The output of this model shows that being male decreases the probability of having only positive interactions (statistically significant) and decreases the probability of having only negative interactions (not statistically significant).

The estimated probabilities are shown in figure 3. The three categories are ordered (from left to right) as 'Both', 'someNeg' and 'somePos'. The blue dots represent male students and the red dots represent female students. The dashed line represents a 50% chance of a falling into a certain category. Within each gender, the estimated probabilities must add up to one.

Clearly, for both male and female students, there is a greater than 50% chance of having only positive interactions. Also, both genders have a miniscule chance of having only negative interactions. Both genders have a less than 50% chance of having positive and negative interactions. However, there is a roughly 20% gap

[2] The model was fit using the 'multinom' function in the R package 'nnet'.

between males and females. Males are 20% more likely to be in the 'Both' group, while females are 20% more likely to be in the 'somePos' group.

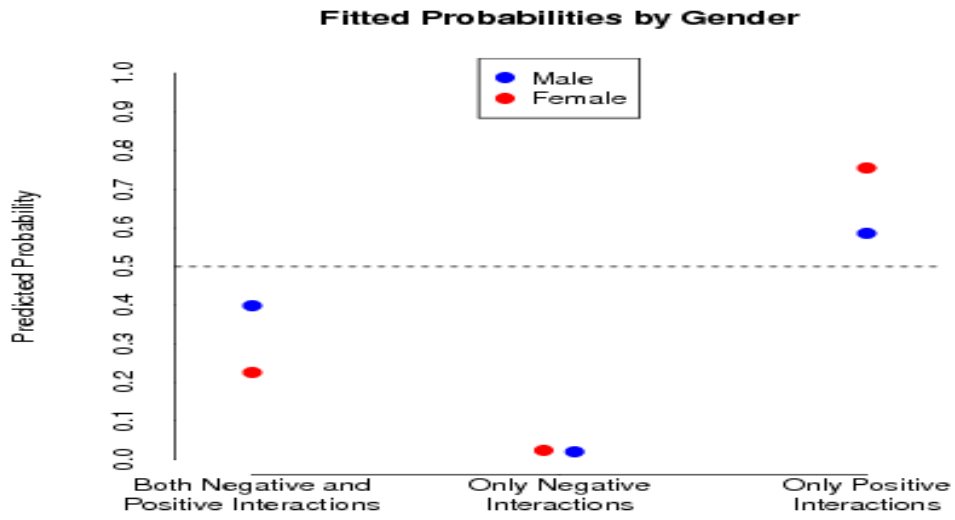

Figure 3. Fitted Probability of having either all positive, all negative or both types of interaction

Conclusion

There are two perspectives from which the question of discrimination can be approached. If students' class experience is the main concern, female students are 20% more likely than male students to have positive interactions with the instructor. On the other hand, if the instructor's behavior is of concern, he has a 7.9% higher chance of giving positive feedback to students wearing unisex type of clothing than to those wearing non-unisex type of clothing.

Although results from the two analytical approaches are significant in the statistical sense, it is hard to interpret these findings as evidence of discrimination (rather than correlation). The main reason for this is the nature of the data. Not only is information to distinguish between students available on only two variables; it is also sparse in the sense of several zero or missing observations. As a result, the scope of the analysis is substantially limited.

Furthermore, the data contains only gender and clothing information to explain the amount of positive and negative feedbacks in the classroom. However, it is likely that there exist several important covariates to explain the nature of these interactions: for example, students' academic performance or major are bound to affect the quality and nature of students' interaction with the instructor. Additionally, even within the available data, the categories are somewhat ambiguous: female standard clothing may

be significantly different from male standard clothing; this is also true for 'other' type of clothing, which comprises of nearly half the population. As a result, it is hard to justify the use of type of clothing as a variable to predict instructor's behaviour towards students. Finally, an instructor's attitude towards students is a product of a multiplicity of important factors, such as the students' behaviour in the classroom, student's academic performance (as measured by GPA), among several other factors. Moreover, to evaluate the effect of instructor's prejudice toward gender or clothing type on his interaction with the students, it is imperative that data be collected in the first week of the term, when a relationship between instructor and students has not developed yet.

In sum, there isn't enough evidence to reject the hypothesis of fairness to believe that the instructor discriminates among his students based on their gender and/or clothing type. Until more information is available, it can be concluded that, in all likelihood, the instructor is fair in his behaviour towards students in his class. Hence, in order to establish with certitude whether or not the instructor displays discriminatory behavior, a follow-up study is recommended, which, it is recommended, should include several other descriptive and objective variables, including those that are suggested in this report.