1 Pareto Curve Generation

To show trade-offs between disparity and prediction errors, we reduced it to a sequence of cost-sensitive classification (CSC) problems, and the solutions to it yield a randomized classifier Q with the lowest (empirical) error subject to the desired constraints, see equation (2).

$$\begin{aligned} L(Q, \lambda_0, \lambda_1) = \; & err(Q) \\ & + \lambda_0(FP(a_1) - FP(a_0)) \\ & + \lambda_1(FN(a_1) - FN(a_0)). \end{aligned} \tag{1}$$

The input to such a CSC algorithm is a data set of training examples $\{(X_i, C_i^0, C_i^1)\}_{i=1}^n$, where C_i^0 and C_i^1 denotes the costs of predicting label 1 and 0 (the defendant to reoffend or not), respectively, for X_i:

$$\operatorname*{argmin}_{h \in Q} \sum_{i=1}^{n} h(x_i)C_i^1 + (1 - h(x_i))C_i^0. \tag{2}$$

Thus, we can incorporate fairness constraints by specifying different costs for different training examples. In particular, equation (2) is equivalent to a weighted classification problem, where the input consists of labeled examples $\{(X_i, Y_i, W_i)\}_{i=1}^n$ with $Y_i \in \{0, 1\}$ and $W_i \geq 0$, and the goal is to minimize the weighted classification errors:

$$\operatorname*{argmin}_{h \in Q} \sum_{i=1}^{n} W_i \mathbf{1}\{h(X_i) \neq Y_i\}. \tag{3}$$

This is equivalent to equation (2) if we set $W_i = |C_i^0 - C_i^1|$ and $Y_i = \mathbf{1}\{C_i^0 \geq C_i^1\}$, where the costs, C_i^0 and C_i^1, can be written as the following:

$$\begin{aligned} C_i^0 &= 1[Y_i = 1] + \lambda_1(1[Y_i = 1, a_i = 1] - 1[Y_i = 1, a_i = 0]), \\ C_i^1 &= 1[Y_i = 0] + \lambda_0(1[Y_i = 0, a_i = 1] - 1[Y_i = 0, a_i = 0]). \end{aligned} \tag{4}$$

Since we constrained the problem to only consider binary protected attributes (African and White defendants, male and female defendants), it is suffices to conduct grid search on values of λ_1 and λ_2 to calculate the best response for each pair of values, and select values with the desired trade-offs between prediction errors and disparity by generating the Pareto frontier.

To solve the binary classification in the cost-sensitive classification problem, we used logistic regression for each $h \in Q$, because logistic regression can provide the probability of each label on each example which allows for thresholding to display trade-offs between the two types of prediction errors.

VizEval - launch 08142019

Start of Block: Welcome & Consent

Q1.1

Welcome

Purpose and Procedure Thank you for taking part in this survey. We are studying how to help users understand intelligent algorithms that support people in making important decisions, such as whether a student should be admitted to college, whether a prisoner should be paroled, and whether a loan applicant should be given a loan. The survey should take approximately 20 to 30 minutes to complete.

Responses will be Confidential All records from this study will be kept private. In addition, in any sort of report we might publish, we will not include any information that makes it possible to identify you. Research records will be stored securely, and only researchers will have access to the records.   Contacts and Questions This survey is being conducted by researchers from the University of Minnesota. If you have any questions about this study, please feel free to contact the researchers at bowen-yu@umn.edu.   Statement of Consent   By clicking "I agree", you consent to participate in this study.

◯ I agree (1)

◯ I do not agree (2)

Skip To: End of Survey If Welcome   Purpose and Procedure Thank you for taking part in this survey. We are studying how to... = I do not agree

End of Block: Welcome & Consent

Start of Block: Background Explanation

Q2.1 In this study, we will focus on an algorithm that makes **predictions about the likelihood of a criminal defendant's *recidivism***. Recidivism refers to a criminal defendant who eventually commits another crime. The algorithm predicts whether a defendant will relapse into criminal behavior, and this prediction can be used by a judge in determining the defendant's sentence.

In the study, we will refer to a dataset containing actual data about people who were accused of a crime. The dataset contains information about 3000 people. 1063 of the people were charged with another crime within two years, and 1937 were not charged with another crime within two years.

The algorithm has parameters that influence its performance and can be tuned to prioritize different outcomes. **In this study you will answer some questions about the algorithm and its trade-offs.**

Q2.2 To be sure you understand the study context, please choose the **correct** option for the following question.

- ◯ The algorithm is a set of rules that judges, probation and parole officers follow to manually decide if each defendant will relapse into criminal behavior. (1)
- ◯ The algorithm is a computer program that predicts whether a defendant will relapse into criminal behavior and whose performance can be tuned by setting its parameters. (2)
- ◯ The algorithm is a computer program that randomly generates a number. (3)
- ◯ The algorithm uses financial information to predict the credit score for defendants. (4)

End of Block: Background Explanation

Start of Block: Algorithm Explanation

Q3.1
The algorithm learns from **historical criminal recidivism data** to predict the likelihood that a defendant will reoffend. The basic logic is that the algorithm **considers a defendant likely to reoffend if his/her profile is similar to profiles of other defendants who have reoffended.**

The algorithm uses fourteen attributes of defendants, such as their age, gender, race, previous criminal records, etc, to evaluate defendants' recidivism risk. The algorithm may weigh some attributes more heavily than others in making its prediction.

Q3.2 To make sure you understand the information, please choose the **incorrect** option for the following question.

- ○ The algorithm randomly predicts whether a defendant will reoffend. (1)
- ○ The algorithm evaluates the likelihood that a defendant will reoffend based their attributes, e.g., age, gender, race, etc. (2)
- ○ The algorithm learns from historical data. (3)
- ○ The attributes considered by the algorithm may have be weighted differently. (4)

Q3.3 I trust the algorithm to make predictions about defendants' recidivism.

- ○ Strong Disagree (1)
- ○ Disagree (2)
- ○ Somehow Disagree (3)
- ○ Neither Disagree or Agree (4)
- ○ Somehow Agree (5)
- ○ Agree (6)
- ○ Strong Agree (7)

End of Block: Algorithm Explanation

Start of Block: Concept Explanation

Q4.1
Scholars have developed various metrics to evaluate the performance of the algorithm. Here are the definitions of metrics we will use in our context. Please spend some time to understand them (we will show them again). **True positive**: A defendant the algorithm predicts to reoffend **does** reoffend (the prediction is CORRECT) **False positive**: A defendant the algorithm predicts to reoffend actually **doesn't** reoffend (the prediction is WRONG) **False negative**: A defendant the algorithm predicts **NOT** to reoffend actually **does** reoffend (the prediction is WRONG) **True negative**: A defendant the algorithm predicts **NOT** to reoffend actually **doesn't** reoffend (the prediction is CORRECT) **Prediction errors**: the

number of defendants that are incorrectly predicted, i.e., the sum of false positives and false negatives **Disparity**: the largest difference between the two types of prediction errors (i.e., false positive and false negative) between the two groups.

Q4.2 To make sure you understand the information correctly, please choose All the Correct options for the following question.

- [ ] False positives and False negatives are the cases of incorrect predictions (1)
- [ ] True positives and True negatives are the cases of correct predictions (2)
- [ ] True positives and False positives are the cases the algorithm predicts the defendants to reoffend (3)
- [ ] True negatives and False negatives are the cases the algorithm predicts the defendants not to reoffend (4)

End of Block: Concept Explanation

Start of Block: Interface Explanation (data view)

Q5.1 For this study, you will take the perspective of a judge who is using a software tool that visualizes the performance of the prediction algorithm, and you will adjust the parameters to control trade-offs between different desired outcomes. Click on the link below to familiarize yourself with the tool. After you have done this, you will be asked questions about your experience with and understanding of the tool. Here are some tips for you: Keep in mind that you can always look at the interface when answering provided questions. For the concepts you are not familiar with on the interface, please move your mouse over to see the hovered explanation. Each dot on the interface represents a defendant.
https://value-sensitive-viz.herokuapp.com/explore_1_data.html

Q5.2 After you have explored the tool, select an option below and click the arrow to continue.

◯ I have explored the tool and am ready to proceed to the next step of the survey (1)

◯ I have not explored the tool (2)

End of Block: Interface Explanation (data view)

Start of Block: Payment Explanation

Q6.1 Bonus payments! For any question marked with a dollar sign ($), you can earn extra payment by answering correctly.   For each question you answer correctly, you get $0.20 bonus payment. If you choose "I don't know", you will get $0.05 bonus payment.
     If you answer 12 questions correctly, you get $2.40 extra;   If you answer 6 questions correctly, you get $1.20 extra;   Answering all the bonus questions correctly will increase the total payment for this HIT to $6.40.

-

End of Block: Payment Explanation

Start of Block: Interface Explanation (scenario view)

Q7.1
For this study, you will take the perspective of a judge who is using a software tool that visualizes the performance of the prediction algorithm, and you will adjust the parameters to control trade-offs between different desired outcomes.   Click on the link below to familiarize yourself with the tool. After you have done this, you will be asked questions about your experience with and understanding of the tool. Here are some tips for you: Keep in mind that you can always look at the interface when answering provided questions.   For the concepts you are not familiar with on the interface, please move your mouse over to see the hovered explanation.
https://value-sensitive-viz.herokuapp.com/explore_1_scenario.html

Q7.2 After you have explored the tool, select an option below and click the arrow to continue.

○ I have explored the tool and am ready to proceed to the next step of the survey (1)

○ I have not explored the tool (2)

End of Block: Interface Explanation (scenario view)

Start of Block: Interface Explanation (baseline)

Q8.1 For this study, you will take the perspective of a judge who is trying to understand the performance of the prediction algorithm.

---

*

Q8.2 Select an option and click the arrow to continue.

○ I am ready to proceed to the next step of the survey. (1)

○ I am not ready. (2)

End of Block: Interface Explanation (baseline)

Start of Block: Payment Explanation (baseline)

Q9.1 Bonus payments! For any question marked with a dollar sign ($), you can earn extra payment by answering correctly. For each question you answer correctly, you get $0.20 bonus payment. If you choose "I don't know", you will get $0.05 bonus payment. If you answer 6 questions correctly, you get $1.20 extra; Answering all the bonus questions correctly will increase the total payment for this HIT to $5.20.

-

End of Block: Payment Explanation (baseline)

Start of Block: Trade-off Related

Q10.1 Remember that in this study we are considering a dataset containing actual data about people who were accused of a crime. The dataset contains information about 3000 people. 1063 of the people were charged with another crime within two years, and 1937 were not charged with another crime within two years. As explained earlier, the algorithm includes parameters that can be adjusted to control its operation, specifically to make it more or less aggressive in identifying any defendants who might reoffend.   To help you answer the questions, we repeat the definitions of some key concepts: True positive: A defendant the algorithm predicts to reoffend does reoffend (the prediction is CORRECT)   False positive: A defendant the algorithm predicts to reoffend actually doesn't reoffend (the prediction is WRONG)   False negative: A defendant the algorithm predicts NOT to reoffend actually does reoffend (the prediction is WRONG) True negative: A defendant the algorithm predicts NOT to reoffend actually doesn't reoffend (the prediction is CORRECT)   Prediction errors: the number of defendants that are incorrectly predicted, i.e., the sum of false positives and false negatives

---

Q10.2 ($) Your goal is to tune the algorithm to **maximize true positives**. How would you adjust the bar?

○ Adjust the algorithm to be more aggressive in identifying defendants who might reoffend (i.e., identify more defendants who might reoffend) (1)

○ Adjust the algorithm to be less aggressive in identifying defendants who might reoffend (i.e., identify fewer defendants who might reoffend) (2)

○ Nothing you can do to maximize true positive (3)

○ I don't know (4)

Q10.3 ($) Your goal is to tune the algorithm to **maximize true negatives**. How would you adjust the bar?

○ Adjust the algorithm to be more aggressive in identifying defendants who might reoffend (i.e., identify more defendants who might reoffend) (1)

○ Adjust the algorithm to be less aggressive in identifying defendants who might reoffend (i.e., identify fewer defendants who might reoffend) (2)

○ Nothing you can do to maximize true negative (3)

○ I don’t know (4)

Q10.4 ($) Your goal is to tune the algorithm to **minimize false positives**. How would you adjust the bar?

- Adjust the algorithm to be more aggressive in identifying defendants who might reoffend (i.e.,identify more who might reoffend) (1)
- Adjust the algorithm to be less aggressive in identifying defendants who might reoffend (i.e., identify fewer who might reoffend) (2)
- Nothing you can do to minimize false positives (3)
- I don’t know (4)

Q10.5 ($) Your goal is to tune the algorithm to **minimize false negatives**. How would you adjust the bar?

- Adjust the algorithm to be more aggressive in identifying defendants who might reoffend (i.e., identify more defendants who might reoffend) (1)
- Adjust the algorithm to be less aggressive in identifying defendants who might reoffend (i.e., identify fewer defendants who might reoffend) (2)
- Nothing you can do to minimize false negatives (3)
- I don’t know (4)

Q10.6 ($) Suppose you tune the algorithm to reduce one type of error cases - false negatives (i.e., the number of defendants it predicts not to reoffend but actually reoffend). How will this affect the other type of error cases - false positives (i.e., the number of defendants it predicts to reoffend but actually do not reoffend)?

- The false positive will drop. (1)
- The false positive will increase. (2)
- There will be no change on the false positives. (3)
- I don't know (4)

Q10.7 **The three questions below do not have correct or incorrect answers. You will receive bonus payment as long as your responses are reasonable.**

Q10.8 ($) What do you value the most regarding recidivism prediction? For example, some people might believe it is most important to identify and detain all defendants who might potentially reoffend; others might believe it is important not to detain anyone who will not reoffend; and others might believe it is important to somehow balance these two goals.

____

____

____

____

____

Q10.9 ($) Adjust the prediction-aggressiveness bar until you think that the algorithm is producing the best outcomes. Now enter the number of the model you chose:

○ Model Number (1) ____

Q10.10 ($) Why do you think that the model you selected produces the best outcomes?

____

____

____

____

____

Q10.11 The visualization tool helps me create a model that represents what I value in prediction outcomes.

- Strongly disagree (1)
- Disagree (2)
- Somewhat disagree (3)
- Neither agree nor disagree (4)
- Somewhat agree (5)
- Agree (6)
- Strongly agree (7)

End of Block: Trade-off Related

Start of Block: Trade-off Related (baseline)

Q11.1 Remember that in this study we are considering a dataset containing actual data about people who were accused of a crime. The dataset contains information about 3000 people. 1063 of the people were charged with another crime within two years, and 1937 were not charged with another crime within two years. As explained earlier, the algorithm includes parameters that can be adjusted to control its operation, specifically to make it more or less aggressive in identifying any defendants who might reoffend. To help you answer the questions, we repeat the definitions of some key concepts: True positive: A defendant the algorithm predicts to reoffend does reoffend (the prediction is CORRECT) False positive: A defendant the algorithm predicts to reoffend actually doesn't reoffend (the prediction is WRONG) False negative: A defendant the algorithm predicts NOT to reoffend actually does reoffend (the prediction is WRONG) True negative: A defendant the algorithm predicts NOT to reoffend actually doesn't reoffend (the prediction is CORRECT) Prediction errors: the number of defendants that are incorrectly predicted, i.e., the sum of false positives and false negatives **Prediction error rate:** the percentage of defendants that are incorrectly predicted.

Q11.2 ($) Your goal is to **maximize true positives**. How would you adjust the algorithm?

- Adjust the algorithm to be more aggressive in identifying defendants who will reoffend (i.e., identify more defendants who might reoffend) (1)
- Adjust the algorithm to be less aggressive in identifying defendants who will reoffend (i.e., identify fewer who might reoffend) (2)
- Nothing you can do to maximize true positive (3)
- I don't know (4)

Q11.3 ($) Your goal is to **maximize true negatives**. How would you adjust the algorithm?

- Adjust the algorithm to be more aggressive in identifying defendants who might reoffend (i.e., identify more defendants who might reoffend) (1)
- Adjust the algorithm to be less aggressive in identifying defendants who might reoffend (i.e., identify fewer defendants who might reoffend) (2)
- Nothing you can do to maximize true negative (3)
- I don't know (4)

Q11.4 ($) Your goal is to **minimize false positives**. How would you adjust the algorithm?

- Adjust the algorithm to be more aggressive in identifying defendants who might reoffend (i.e., identify more defendants who might reoffend) (1)
- Adjust the algorithm to be less aggressive in identifying defendants who might reoffend (i.e., identify fewer defendants who might reoffend) (2)
- Nothing you can do to minimize false positive (3)
- I don't know (4)

Q11.5 ($) Your goal is to **minimize false negatives**. How would you adjust the algorithm?

◯ Adjust the algorithm to be more aggressive in identifying defendants who might reoffend (i.e., identify more defendants who might reoffend) (1)

◯ Adjust the algorithm to be less aggressive in identifying defendants who might reoffend (i.e., identify fewer defendants who might reoffend) (2)

◯ Nothing you can do to minimize false negative (3)

◯ I don't know (4)

Q11.6 ($) Suppose you tune the algorithm to reduce one type of error cases - false negatives (i.e., the number of defendants it predicts not to reoffend but actually reoffend). How would this affect the other type of error cases - false positives (i.e., the number of defendants it predicts to reoffend but actually do not reoffend)?

◯ The false positive will drop. (1)

◯ The false positive will increase. (2)

◯ There will be no change on the false positives. (3)

◯ I don't know (4)

End of Block: Trade-off Related (baseline)

Start of Block: Interface reminder (data view)

Q12.1 In case you close the interface link accidentally, you can reopen it by clicking the link below:

https://value-sensitive-viz.herokuapp.com/explore_1_data.html

End of Block: Interface reminder (data view)

Start of Block: Interface reminder (scenario view)

Q13.1
In case you close the interface link accidentally, you can reopen it by clicking the link below:
https://value-sensitive-viz.herokuapp.com/explore_1_scenario.html

End of Block: Interface reminder (scenario view)

Start of Block: Fairness Related

Q14.1 For the next questions, make sure you select race to explore the group effect, and you will see the error-disparity bar will be displayed to visualize trade-offs between errors and group disparities.
Assume that the prediction-aggressiveness bar and the error-disparity bar both are initially set in the middle of their range. Next, answer the following questions, adjusting the error-disparity bars as the questions specify (Note that adjusting prediction-aggressiveness bar will not affect answering questions in this section).
To help you answer the questions, we repeat the definitions of some key concepts:
Prediction errors: the number of **defendants that are incorrectly predicted, i.e., the sum of false positives and false negatives** **Disparity: the largest difference between the two types of prediction errors (i.e., false positive and false negative) between the two groups.**

Q14.2 ($) Be sure you have selected race to explore the group effect, and be sure that the error-disparity bar is set in the middle of their range. Now, adjust the error-disparity bar to minimize disparity between African-American and White defendants. To do this, how would this affect the algorithm's prediction error?

○ The number of incorrect predictions will go up. (1)

○ The number of incorrect predictions will go down. (2)

○ The number of incorrect predictions will stay the same. (3)

○ I don't know. (4)

Q14.3 **The three questions below do not have correct or incorrect answers. You will receive bonus payment as long as your responses are reasonable.**

Q14.4 ($) What do you value the most regarding recidivism prediction when the disparity of predicting the recidivism of African American and White American defendants can be considered? For example, some people might believe it is most important to treat the two groups of defendants without bias; some people might believe it is important to make fewer mistakes when making predictions; and some might believe it is important to somehow balance the two cases.

____

____

____

____

____

Q14.5 ($) Adjust the error-disparity bar (and maybe prediction-aggressiveness bar) until you think that the algorithm is producing the best outcomes. Now enter the number of the model you chose:

○ Model Number: (1) ____

Q14.6 ($) Why do you think that the model you selected produces the best outcomes?

____

____

____

____

____

Q14.7 The visualization tool helps me create a model that represents what I value in prediction outcomes.

- Strongly disagree (1)
- Disagree (2)
- Somewhat disagree (3)
- Neither agree nor disagree (4)
- Somewhat agree (5)
- Agree (6)
- Strongly agree (7)

End of Block: Fairness Related

Start of Block: Self Evaluation

Q15.1 I am confident in my responses.

- Strongly disagree (1)
- Disagree (2)
- Somewhat disagree (3)
- Neither agree nor disagree (4)
- Somewhat agree (5)
- Agree (6)
- Strongly agree (7)

Q15.2 The visualization tool helps me understand trade-offs in the recidivism prediction algorithm.

- Strongly disagree (1)
- Disagree (2)
- Somewhat disagree (3)
- Neither agree nor disagree (4)
- Somewhat agree (5)
- Agree (6)
- Strongly agree (7)

Q15.3 Help us verify that you are paying attention! Please choose the third option from the left.

- Strongly disagree (1)
- Disagree (2)
- Somewhat disagree (3)
- Neither agree nor disagree (4)
- Somewhat agree (5)
- Agree (6)
- Strongly agree (7)

Q15.4 I became familiar with the visualization tool quickly.

- Strongly disagree (1)
- Disagree (2)
- Somewhat disagree (3)
- Neither agree nor disagree (4)
- Somewhat agree (5)
- Agree (6)
- Strongly agree (7)

Q15.5 The visualization tool is easy to use.

- Strongly disagree (1)
- Disagree (2)
- Somewhat disagree (3)
- Neither agree nor disagree (4)
- Somewhat agree (5)
- Agree (6)
- Strongly agree (7)

Q15.6 I trust defendants' recidivism prediction results produced by the algorithm.

- Strongly disagree (1)
- Disagree (2)
- Somewhat disagree (3)
- Neither agree nor disagree (4)
- Somewhat agree (5)
- Agree (6)
- Strongly agree (7)

Q15.7 If your perception about the trust to the algorithm's recidivism prediction has changed, what caused that?

____

____

____

____

____

End of Block: Self Evaluation

Start of Block: Fairness Related (baseline)

Q16.1 To remind you of the context of this study, it refers to a dataset containing information about 3000 convicted criminals who were considered for parole and who actually were paroled. 1063 of the parolees reoffended within two years, and 1937 did not reoffend. The next question ask you to make some judgements based on this background information. To help you answer the questions, we repeat the definitions of some key concepts: **Prediction errors:** the number of **defendants that are incorrectly predicted, i.e., the sum of false positives and false negatives Disparity: the largest difference between the two types of prediction errors (i.e., false positive and false negative) between the two groups.**

Q16.2 ($) If you want to minimize disparity between African-American and White defendants, how would this affect algorithm’s prediction error?

- The number of incorrect predictions will go up. (1)
- The number of incorrect predictions will go down. (2)
- The number of incorrect predictions will stay the same. (3)
- I don't know. (4)

End of Block: Fairness Related (baseline)

Start of Block: Self Evaluation (baseline)

Q17.1 I am confident in my responses.

- Strongly disagree (1)
- Disagree (2)
- Somewhat disagree (3)
- Neither agree nor disagree (4)
- Somewhat agree (5)
- Agree (6)
- Strongly agree (7)

Q17.2 I can understand trade-offs in the recidivism prediction algorithm.

- Strongly disagree (1)
- Disagree (2)
- Somewhat disagree (3)
- Neither agree nor disagree (4)
- Somewhat agree (5)
- Agree (6)
- Strongly agree (7)

Q17.3 Help us verify that you are paying attention! Please choose the third option from the left.

- Strongly disagree (1)
- Disagree (2)
- Somewhat disagree (3)
- Neither agree nor disagree (4)
- Somewhat agree (5)
- Agree (6)
- Strongly agree (7)

Q17.4 I trust defendants' recidivism prediction results produced by the algorithm.

- Strongly disagree (1)
- Disagree (2)
- Somewhat disagree (3)
- Neither agree nor disagree (4)
- Somewhat agree (5)
- Agree (6)
- Strongly agree (7)

Q17.5 If your perception about the trust to the algorithm’s recidivism prediction has changed, what caused that?

____

____

____

____

____

End of Block: Self Evaluation (baseline)

Start of Block: Demographic Info

Q18.1 What's your age?

- ○ Under 18 (1)
- ○ 18 - 24 (2)
- ○ 25 - 34 (3)
- ○ 35 - 44 (4)
- ○ 45 - 54 (5)
- ○ 55 - 64 (6)
- ○ 65 - 74 (7)
- ○ 75 - 84 (8)
- ○ 85 or older (9)

Q18.2 What's your gender?

- ○ Male (1)
- ○ Female (2)

Q18.3 What is the highest degree or level of schooling you have completed? If currently enrolled, the highest degree you are pursuing.

- ○ Less than high school (1)
- ○ High school (2)
- ○ College degree (3)
- ○ Master degree or professional program (4)
- ○ Doctorate (5)

Q18.4 What's your race?

- White American (1)
- African American (2)
- Hispanic (3)
- Asian (4)
- Others (5)

Q18.5 Your familiarity with the judicial system.

- Not familiar at all (1)
- Not familiar (2)
- Moderately not familiar (3)
- Neither familiar nor not familiar (4)
- Moderately familiar (5)
- Familiar (7)
- Extremely familiar (8)

Q18.6 Your familiarity of the use of AI-powered systems (e.g., email spam filter, amazon recommendations, etc)

- Not familiar at all (1)
- Not familiar (2)
- Moderately not familiar (3)
- Neither familiar nor not familiar (4)
- Moderately familiear (5)
- Familiar (6)
- Extremely familiar (7)

Q18.7 Please let us know if you have any feedback or comments.

____

____

____

____

____

End of Block: Demographic Info