

---

# AUDITING A DUTCH PUBLIC SECTOR RISK PROFILING ALGORITHM USING AN UNSUPERVISED BIAS DETECTION TOOL

---

A PREPRINT

**Floris Holstege\***  
University of Amsterdam  
The Netherlands  
f.g.holstege@uva.nl

**Mackenzie Jorgensen**  
King's College London  
Alan Turing Institute  
UK

**Kirtan Padh**  
TU Munich  
Helmholtz Munich  
Germany

**Jurriaan Parie**  
Algorithm Audit  
The Netherlands  
j.parie@algorithmaudit.eu

**Joel Persson**  
Algorithm Audit  
UK

**Krsto Prorokovic**  
Algorithm Audit  
The Netherlands  
info@algorithmaudit.eu

**Lukas Snoek**  
Algorithm Audit  
The Netherlands

February 5, 2025

## ABSTRACT

Algorithms are increasingly used to automate or aid human decisions, yet recent research shows that these algorithms may exhibit bias across legally protected demographic groups. However, data on these groups may be unavailable to organizations or external auditors due to privacy legislation. This paper studies bias detection using an unsupervised clustering tool when data on demographic groups are unavailable. We collaborate with the Dutch Executive Agency for Education to audit an algorithm that was used to assign risk scores to college students at the national level in the Netherlands between 2012-2023. Our audit covers more than 250,000 students from the whole country. The unsupervised clustering tool highlights known disparities between students with a non-European migration background and Dutch origin. Our contributions are three-fold: (1) we assess bias in a real-world, large-scale and high-stakes decision-making process by a governmental organization; (2) we use simulation studies to highlight potential pitfalls of using the unsupervised clustering tool to detect true bias when demographic group data are unavailable and provide recommendations for valid inferences; (3) we provide the unsupervised clustering tool in an open-source library. Our work serves as a starting point for a deliberative assessment by human experts to evaluate potential discrimination in algorithmic-supported decision-making processes.

**Keywords** fairness · auditing · bias · clustering · indirect discrimination

## 1 Introduction

The use of algorithmic support for decision-making has grown in recent years [61], influencing decisions such as who gets a loan, who makes it to the next recruitment stage for a job, and who is flagged for fraud investigation [60, 56, 64]. As algorithmic systems are adopted, their negative impacts have also become more clear [8, 67, 44, 6, 56]. These systems may replicate the bias we see in our society, leading to potentially discriminatory effects [62]. In this study, we outline a methodology to detect such biases in algorithms without data on demographic groups using an unsupervised bias detection algorithm, evaluated on a Dutch public sector risk profiling algorithm.

---

\*Corresponding author

In the period 2012-2023, a simple rule-based risk profiling algorithm of the Dutch Executive Agency for Education (DUO) contributed to indirect discrimination in a control process to check whether students were (un)duly allocated a college grant [5]. A rule-based algorithm, consisting of three apparently neutral characteristics of students (type of education, age, distance to parents), was used to assign higher risk scores to students enrolled in vocational education, younger students, and students who lived near their parental address. The seemingly impartial risk assessment led to unequal treatment of students with a non-European migration background, remaining unnoticed for over 10 years and affecting more than 350,000 students<sup>2</sup> [4]. Given this, in 2024 the Dutch Minister for Education, Culture and Science apologized on behalf of the Dutch government for indirect discrimination experienced by students with a non-European migration background and announced a EUR 61 million compensation plan [4]. This example of a new discriminatory algorithm scandal in the Netherlands [34] highlights that not only complex statistical methods, like machine learning and deep learning [15], but also simple rule-based algorithms pose significant risks of embedding discrimination in algorithm-supported decision-making processes. In a recent interview, the chair of the Dutch Data Protection Authority warned that discriminatory rule-based algorithms are “around every corner” in executive branches of the Dutch government [59].

This type of risk profiling is prohibited by the Charter of Fundamental Rights of the European Union (EU) [2], the EU General Data Protection Regulation (GDPR) [3] and the EU AI Act [28], as well as the national Public Administration Law [1]. Limited institutional enforcement capabilities, along with a lack of awareness, have led to inadequate oversight of algorithmic-supported decision-making processes [47].

To combat bias in algorithmic systems, the academic community has presented a multitude of ways to define, detect, and mitigate different notions of bias at different points along the algorithmic pipeline [46, 17, 10, 68]. Most bias mitigation methods require data on the demographic groups to prevent discriminatory bias. These demographic groups are typically protected attributes under discrimination law, such as race, gender and age<sup>3</sup> [9]. However, data on demographic groups are often not available to organizations or external algorithm auditors [7, 14]. Upon request, public sector organizations can occasionally obtain aggregated statistics on demographic groups through official institutions, such as the National Office of Statistics, which have access to statistical population registers. Private organizations cannot request these data and face ethical, legal, and practical challenges when collecting data on demographic groups customers belong to [14, 29, 28]. Unsupervised learning can aid in this challenge by detecting groups in the data for which a given notion of equal treatment is violated, avoiding the need for ex-ante specification of sensitive groups [48].

In this paper, we use an unsupervised bias detection tool to audit the risk profiling algorithm of DUO on bias. For this audit, the country of birth and country of origin were provided by Statistics Netherlands [51]. These data enable us to verify whether the clusters identified by the unsupervised bias detection tool, which represent potentially unfairly treated groups, correspond to disadvantaged demographic groups. This allows for a comparison between an unsupervised learning approach and an approach using demographic data. As a bias metric, we use whether or not a student is categorized as “high risk” by DUO’s risk profiling algorithm. We find that the cluster with the highest bias returned by the unsupervised bias detection tool is characterized by students who are more often enrolled in vocational education, live close to their parental address, and are relatively more often 15-18 or 25-50 years old. This aligns with demographic groups found to be negatively affected by the risk profiling algorithm, based on aggregated statistics regarding students’ non-European migration background.

Our paper makes three contributions:

1. **Large scale audit of indirect discrimination:** The DUO audit presents a high-quality and socially relevant dataset, providing insights into the assessment of indirect discrimination in algorithm-supported decision-making. This study thereby contributes to the advancement of public knowledge building and establishment of standards for auditing algorithmic systems on indirect discrimination.
2. **Advancing methodology for unsupervised bias detection:** We build on existing clustering algorithms and strengthen their methodology. Using a simulation study, we examine conditions under which the tool works best and highlight key factors to consider for reliable inferences when applying it.

---

<sup>2</sup>Between 2012-2023, over 350,000 students were assigned a risk score by the algorithm. Our analysis focuses on the student population in two years: 2014 and 2019, which consists of a total of 298,882 students.

<sup>3</sup>The exact protected grounds are enshrined in European non-discrimination law, such as the Treaty on the Functioning of the European Union (TFEU), the Charter of Fundamental Rights of the EU, and specific directives like the Racial Equality Directive (2000/43/EC) and the Employment Equality Directive (2000/78/EC). The main protected attributes include: race and ethnicity, gender and sex, religion or belief, disability, age, sexual orientation, nationality and marital or family status. Throughout, we will refer to these attributes as *demographic groups*. Note that some provisions do not mention any particular grounds of discrimination, such as Art. 20 of the EU Charter. Instead using openly formulated clauses such as “everyone is equal before the law”.

3. **Open-source tool:** We have published the unsupervised bias detection method as an open-source tool in the form of a Python package, as well as an online web application. This allows a non-technical audience to easily use the tool on their own datasets to detect potential bias.

Overall, our paper showcases the potential of unsupervised learning to aid in auditing algorithmic-supported decision-making processes via an application to a real-world use case. We argue that unsupervised bias detection serves as a starting point for a deliberative assessment by human experts to evaluate potential discrimination and unfairness in the algorithmic-supported decision-making process under review.

The rest of the paper is outlined as follows. We start by reviewing related literature in Section 2. We present the methodology behind the unsupervised bias detection tool in Section 3. Then, in Section 4, we describe the application of the unsupervised bias detection tool to DUO’s risk profiling algorithm. We follow this by providing simulation studies in Section 5 which help to understand potential pitfalls when using the tool. We finish with a discussion on the limitations and conclusions of our work in Sections 7 and 8.

## 2 Literature Review

We review a selection of the literature on bias detection and non-discrimination law, structuring the review around three topics: (1) challenges in detecting and establishing indirect discrimination, (2) bias detection methods with and without access to demographic data, and (3) auditing algorithmic systems on bias.

### 2.1 Challenges in Detecting and Establishing Discrimination

Under European non-discrimination law, DUO’s control process is classified as indirect discrimination as an apparently neutral practice disproportionately disadvantaged students with a non-European migration background<sup>4</sup> [53]. In algorithms, discrimination often arises from the *proxy and correlation challenge* [30]. Proxies refer to characteristics that are correlated with a demographic group, leading to disadvantaged outcomes for this group. For example, the Court of Justice of the European Union (CJEU) found that using part-time employment as a basis for differentiation in the context of occupational pension schemes constituted indirect discrimination, as over 80% of the part-time workers in Spain in 2012 were women [54]. In some cases, courts may accept that differential treatment of a demographic group is acceptable. Under EU law, indirect discrimination can be justified if it is demonstrated that a legitimate aim is being pursued, and the means of achieving that aim are appropriate and necessary [23, 24, 25, 26]. Our paper highlights how unsupervised methods can aid in the challenge of assessing indirect discrimination in algorithmic-supported decision-making processes.

### 2.2 Bias Detection With and Without Access to Demographic Groups

In the algorithmic fairness literature, identifying systematic differences in the treatment of individuals or groups compared to others is commonly referred to as *bias detection*. Most bias detection methods assume access to demographic groups such as gender or ethnicity [11, 32, 42, 39, 19]. However, in practice, these labels are often unavailable due to privacy legislation [65]. For instance, the EU’s GDPR restricts the collection of “special categories of data” (article 9) [29], which include data on ethnicity, religion, health, and sexual orientation<sup>5</sup>. The EU AI Act introduces a potential exception to the prohibition on processing special category data: “to the extent that it is strictly necessary for the purpose of ensuring bias detection and correction” [28]. Nonetheless, ongoing legal uncertainties regarding the interplay between the AI Act and the GDPR underscore the need for methods to assess bias in algorithmic systems without relying on demographic data.

There are several alternatives to using data on demographic groups. First, researchers have presented methods to predict demographic groups, which are often referred to as proxy models or attribute classifiers [9]. For example, Bayesian Improved Surname Geocoding (BISG) [22] uses Bayes Theorem and data from the United States Census Bureau, including their surname and address, as a proxy for ethnicity or race [9]. This set of methods does require to specify the demographic group of interest beforehand. A second set of methods circumvents the problem of limited availability of protected characteristics by using a Rawlsian definition of fairness [58] that maximizes utility for the most disadvantaged group [43, 33, 16]. However, the outcome of these methods is found to not always adhere to more traditional parity-based group fairness metrics [37]. The third set of alternatives are unsupervised learning methods for

<sup>4</sup>Supra note 2

<sup>5</sup>Age and gender are not considered special categories of data. See Article 9(1) GDPR. Although, a caveat could exist [66]. The other way around, “political opinions,” “trade union membership,” “genetic,” and “biometric” data are special categories of data but are not protected by EU non-discrimination directives.

fairness [49, 40]. These methods tend to modify traditional clustering algorithms to handle fairness criteria explicitly. Our work contributes to this latter category. The main benefit of this set of methods compared to proxy inferring methods is that they do not require specification of demographic groups beforehand. Thus, unsupervised learning methods are particularly helpful to examine the effect of apparently neutral provisions that might result in indirect discrimination.

### 2.3 Auditing Algorithmic Systems

Inspired by established practices in non-algorithmic disciplines, auditing has the potential to mitigate ethical and legal risks in algorithmic systems through both internal and external oversight [57]. Legislators across various jurisdictions are implementing regulations to impose audit frameworks for these systems. For instance, since 2023, New York City has required automated employment decision-making tools to undergo audits by independent third-party auditors [52]. Similarly, since 2024, the EU Digital Services Act mandates annual independent third-party audits for online platforms and search engines with over 45M annual users [27]. In the coming years, the EU AI Act places predominantly self-regulatory obligations on producers and deployers of high-risk AI systems [28].

Implementing effective auditing regimes for algorithmic systems, however, poses practical and cultural challenges. Auditors frequently encounter obstacles such as restricted data access [31], quandaries about which fairness metrics fit the context best [18] and varying opinions on what constitutes a legitimate auditor [31]. Furthermore, auditing reports are not always required to be disclosed, reducing transparency and hindering the development of public knowledge regarding auditing standards. By discussing the audit of DUO’s risk profiling algorithm, we contribute to public knowledge on best practices for auditing algorithmic systems.

## 3 Unsupervised Bias Detection Tool

This section introduces the theoretical background of the unsupervised bias detection tool. We use examples from our DUO case study to clarify the methodology. Figure 1 provides an overview of all steps in applying the tool.

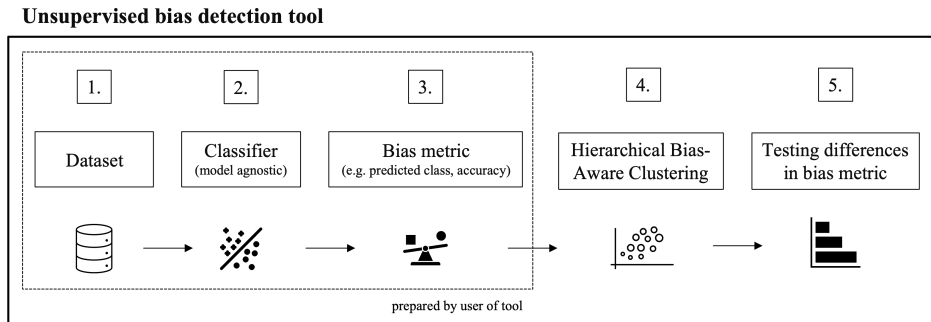


Figure 1: Schematic overview of the steps involved in applying the unsupervised bias detection tool. The required information is a dataset, classifier, and a bias metric. Part of the dataset is used to train the Hierarchical Bias-Aware Clustering algorithm [48]. Another part of the dataset is used to test whether differences in the bias metric across clusters are statistically significant.

### 3.1 Notation and Problem Description

Suppose that we have a total of  $N$  students in the dataset, labeled  $u_1, u_2, \dots, u_N$ . Each row in the dataset reflects characteristics  $x_i \in \mathbb{R}^d$  (e.g., education level, age, distance to parents) of a student  $u_i$  receiving a grant. Suppose also that we have a risk profiling algorithm  $f : \mathbb{R}^d \rightarrow \{0, 1\}$  (e.g., the risk profiling algorithm used by DUO) that takes the features of a user as input and outputs whether a user should be deemed high risk of unduly using a grant. There exists a true value of this variable of interest, denoted  $y_i$ , denoting whether the grant was, in fact, duly given (or not).

We measure the degree to which an algorithm is biased by using a bias metric  $M$ . For a user  $u_i$ , the value of the metric  $m_i = M(f(x_i), y_i)$  is some function of the prediction  $f(x_i)$  and (potentially)  $y_i$ . To give some concrete examples, this could be a measure of performance, such as accuracy (requiring  $f(x_i), y_i$ ). Alternatively, it could measure demographic parity (only requiring  $f(x_i)$ ), equality of opportunity, and so on. Without loss of generality, we assume that a larger average value for  $m_i$  for one group compared to another group indicates bias towards the former group.

For  $k \subset \{1, 2, \dots, N\}$ , let  $U_k = \{u_i\}_{i \in k}$  be a subset of all applicants  $U$ . We denote by  $\bar{M}(U_k)$  the mean of the metric  $m_i$  among all in  $U_k$ . Our goal is to test whether the difference between  $\bar{M}(U \setminus U_k)$  and  $\bar{M}(U_k)$  is statistically significant. If we had access to a demographic group for which we would like to investigate bias, such as race or gender, we would simply define the subset  $U_k$  to be different genders, race or intersections thereof. However, since the demographic group is unavailable in our setting, we leverage clustering algorithms to separate the data into subsets. Let  $\mathcal{K} = \{1, \dots, K\}$  denote the set of clusters, and  $x_{i,k}, y_{i,k}, m_{i,k}$  to denote data from a user  $u_i$  belonging to cluster  $k \in \mathcal{K}$ .

### 3.2 Clustering Algorithm

We use the Hierarchical Bias-Aware Clustering methodology from Misztal-Radecka and Indurkha [48], hereafter referred to as ‘‘HBAC.’’ This is an iterative clustering algorithm that produces a partition of the original dataset to isolate potentially unfairly treated groups into distinct clusters. We adopt this algorithm because it is a standard approach for unsupervised bias detection in the realm of recommendation systems [20].

The algorithm first considers all observations in the dataset to belong to a single cluster and then iteratively splits the observations into distinct clusters. At each iteration, the algorithm identifies the cluster with the highest standard deviation of a specified bias metric, which serves as an indicator of imbalance between clusters. This cluster is then split into two subclusters. The process repeats for a predefined number of iterations, *max\_iterations*. Pseudocode for the algorithm is provided in Algorithm 1. Compared to the original algorithm from Misztal-Radecka and Indurkha [48], we allow for clusters to be defined via  $k$ -modes in addition to  $k$ -means to allow clustering based on categorical and binary variables.

---

#### Algorithm 1: Hierarchical Bias-Aware Clustering

---

**Input:** A dataset  $\mathcal{X} = \{x_1, \dots, x_N\}$  and bias metric  $\{m_1, \dots, m_n\}$ . Set the *max\_iterations* and a minimum of samples per cluster  $n_{\min}$ .

**Output:** A partition  $\{C_1, \dots, C_k\}$

```

1 Define the partition =  $\{\mathcal{X}\}$ 
2 for  $i \leftarrow 1$  to max_iterations do
3   | Set  $C$  to be the cluster in partition with the highest standard deviation of metric  $M$  among those that have not been
   |   selected in any previous iteration.
4   | Split  $C$  into two clusters  $C'$  and  $C''$  using  $k$ -means or  $k$ -modes
5   | if  $\max(\bar{M}(C'), \bar{M}(C'')) \geq \bar{M}(C) \wedge |C'| \geq n_{\min} \wedge |C''| \geq n_{\min}$  then
6   |   | Remove  $C$  from partition
7   |   | Add  $C'$  and  $C''$  to partition
8   | end
9 end

```

---

The original algorithm from Misztal-Radecka and Indurkha [48] does not specify how to assign clusters to data points that are not part of the original training dataset. We propose using the centroids of each cluster. The datapoint is assigned to the cluster whose distance from the centroid is the smallest. In the case of  $k$ -means and  $k$ -modes, the centroids are defined as respectively the mean and mode for each of the characteristics  $x_{i,k}$  [38, 36]. The Euclidean distance ( $k$ -means) or Hamming distance ( $k$ -modes) is used to determine which centroid is closer.

### 3.3 Minimum Number of Samples per Clusters

Algorithm 1 has been claimed to not require a specification of the number of clusters [48]. Although this is true, other hyperparameters of the algorithm such as  $n_{\min}$  – the minimum number of samples per cluster – affect the number of clusters it finds. To determine the appropriate minimum number of samples per cluster, we use the Calinski-Harabasz index [12] of the bias metric. This index measures the ratio of the *between* cluster variance over the *within* cluster variance. Intuitively, we want the clusters to have little variation in the bias metric *within* each cluster, but large variation *between* the clusters. This means that the data points are well separated into distinct clusters.

### 3.4 Testing Differences in Bias Between Clusters

After conducting the clustering, our goal is to test for each of the clusters the following null versus alternative hypothesis:

$$H_0 : \bar{M}(U \setminus U_k) = \bar{M}(U_k), \quad H_A : \bar{M}(U \setminus U_k) \neq \bar{M}(U_k). \quad (1)$$

If  $\bar{M}(U \setminus U_k) \neq \bar{M}(U_k)$ , the bias is found to be higher or lower for a particular cluster found by the clustering algorithm. To test these hypotheses, we use a two-sample  $t$ -test when the bias metric is continuous, and Pearson’s  $\chi^2$ -test when the bias metric is a binary variable. For either case, the test is two-tailed, as we allow  $\bar{M}(U_k)$  to be both lower and higher than  $\bar{M}(U \setminus U_k)$ , meaning that we can detect bias both when it is greater or less for a group compared to the rest.

For valid statistical inference, the null and alternative hypothesis must be defined before observing the data. If we conduct our statistical test on data which was also used to identify clusters, this condition is violated. Since the clusters are selected to maximize the difference in the bias metric, we will therefore tend to incorrectly conclude that there is bias when there truly is none. The broader issue of testing data-driven hypotheses—often called *post-selection inference*—has been studied extensively in the context of regularized regression (for an overview, see Kuchibhotla et al. [41]). A long-standing practice in that literature is to use different partitions of the data for model estimation and inference, referred to as *sample splitting* [45, 63]. We apply this technique by fitting the HBAC algorithm on a subset of the data and subsequently inferring the clusters and their bias on a *new* subset of data. Importantly, this set cannot be used to determine our hyperparameters.

When testing for multiple  $K > 2$  clusters, an issue is that without adjustments, the likelihood of (incorrectly) detecting a statistically significant difference under the null hypothesis increases. We address this by using a Bonferroni correction [35]. In Section 5, we illustrate that the combination of the Bonferroni correction and sample splitting prevents us from (wrongly) concluding that there is a difference in the bias metric while there is none.

## 4 Detecting Bias in DUO’s Risk Profiling Algorithm

The unsupervised bias detection tool will be applied to the DUO dataset. Because this dataset includes ground truth labels for the demographic group, in the form of aggregated statistics on the migration background of students, our case study is suited to evaluate the effectiveness of the tool in detecting bias.

For the DUO audit, researchers were granted access to aggregated statistics on students’ migration by Statistics Netherlands [51]. If these data had not been available, the unsupervised bias detection tool discussed in this paper could have been used as an alternative methods for examining potential bias in the CUB process.

### 4.1 Description of the Dataset

From 2012 to 2023, DUO aimed to verify that college grants were allocated to students who genuinely lived away from their parents. Some students claimed to DUO an independent living status to qualify for a grant, even though they did not actually reside at the address they provided [5]. The process used to verify whether a student lived at the stated address is known as the *CUB process*<sup>6</sup>. As part of the CUB process, students were assigned a risk score ranging between 0 and 180 by a rule-based algorithm, of which the details are provided in Appendix A. This risk score was binned in six risk categories ranging from very high (1), high (2), average (3), low (4), very low (5) to unknown (6)<sup>7</sup>. Students in category (1) and (2) were considered to be “high risk” of misusing the subsidy. Based on the assigned risk category, human analysts of DUO followed operational instructions to subject students to an investigation. For example, students living in a care home and married students were excluded from investigation. While students from all risk categories could potentially be investigated, the majority of investigations focused on higher-risk categories [5].

After concerns were raised regarding potential bias in the CUB process affecting students with a non-European migration background, DUO requested the Dutch national office of statistics (Statistics Netherlands) to provide aggregated statistics<sup>8</sup> of student’s country of origin and country of birth<sup>9</sup> who received a college grant for the years 2014, 2017, 2019, 2021, and 2022 [51]. From these data the non-European migration background of students can be retrieved in aggregate. We do not have information for individual students about whether or not they have a non-European migration background.

This paper focuses on the years 2014 and 2019, covering in total 298,882 students who received college grants. We specifically focus on the CUB population from 2014 ( $n = 248,649$ ) as this was the last year all types of students –

<sup>6</sup>CUB stands for *Controle Uitwonendenbeurs*.

<sup>7</sup>Risk category 6 includes students with a risk score of 0. Risk categories 5, 4, 3, 2 and 1 correspond to students with risk scores in the ranges of 1-19, 20-39, 40-59, 60-79 and 80-180 respectively. See also Appendix A.

<sup>8</sup>Through protocols established by Statistics Netherlands, aggregated population statistics can be computed securely. Data were analyzed in a privacy-preserving manner within a secure environment. Groups are rounded to the nearest ten. Results pertaining to groups smaller than 10 people are not published.

<sup>9</sup>Following Statistics Netherlands, students with a non-European migration background are defined as those who were either born outside the Netherlands and Europe, or who have at least one parent born outside these regions. Countries that are considered to be European can be found here [50].

MBO 1-2 and MBO 3-4 (vocational education), HBO (higher professional education) and WO (university education) – were awarded a college grant. We do not combine data from the different years since in 2019 only MBO students were eligible to receive a college grant. Results are replicated for the CUB 2019 dataset ( $n = 50,233$  students) in Appendix B.

DUO’s risk profiling algorithm differentiated on the basis of three variables. First, the type of education the student is following, categorized as MBO 1-2, MBO 3-4 (vocational education) HBO (higher professional education) and WO (university education). Second, the age of the student, categorized as: 15-18, 19-20, 21-22, 23-24 and 25-50 years. Third, the distance between the address of the parents and the student, measured through the 6-digit postal code of the registered addresses and categorized as: 0km<sup>10</sup>, 1m-1km, 1-2km, 2-5km, 5-10km, 10-20km, 20-50km, 50-500km, unknown.

In 2014, 34,050 students were assigned the risk category 6 (unknown), e.g., because the distance to their parents was unknown. We exclude these students from the clustering analysis, since the algorithm did not make a prediction on their risk of unduly using the grant, and these cases were deprioritized in the control process [5]. Clustering will therefore be applied to remaining dataset of  $n = 214,599$  students.

## 4.2 Applying the Unsupervised Biased Detection Tool

The following choices were made in applying the unsupervised bias detection tool to the CUB 2014 dataset.

**Data preprocessing:** We transform the 3 categorical profiling characteristics to a set of 17 binary (non-overlapping) variables. We define no ordinal relationship between categories. We use the  $k$ -modes clustering within Algorithm 1.

**Bias metric:** We select whether a student is deemed “high risk” by the risk profiling algorithm or not as the bias metric. If the likelihood of being classified as “high risk” is equal for students of Dutch origin and those with a non-European migration background, it would correspond to the fairness metric— demographic parity [21]. Alternative definitions of bias often whether a college grant was duly granted ( $y_i$ ), which in this case is only available for the small set of students who have manually been controlled by DUO ( $n = 3,179$ ). As classifications by the risk profiling algorithm contributed to discrimination in the CUB process [4], we consider this an appropriate bias metric. We emphasize that the HBAC algorithm can be used for each bias metric of interest.

**Hyperparameters:** The minimum number of samples per cluster is selected by iterating over 5,000, 10,000, 20,000, 30,000 and selecting the one which maximizes the average Calinski-Harabasz index, measured via 5-fold cross-validation. We set the maximum number of iterations sufficiently high ( $max\_iterations = 1,000$ ), such that Algorithm 1 continues until the minimum number of samples per cluster is reached. The initial centroids are determined via the method proposed by Cao et al. [13].

**Testing procedure:** We split our dataset and reserve 20% of the dataset for testing the differences in the bias metric per cluster. We fit the unsupervised bias detection tool (and determine hyperparameters) on the remaining 80% of the data. Since our bias metric is binary, we use a Pearson’s  $\chi^2$ -test.

## 4.3 Analyzing Identified Clusters

The unsupervised bias detection tool returns three clusters for the CUB 2014 dataset. The characteristics of these clusters are analyzed.

**Cluster size:** Cluster 1 consists of 56,354 (26.3%) students, where cluster 2 consists of 123,203 (57.4%) students and cluster 3 of 35,042 (16.3%) students.

**Testing the difference in bias metric per cluster:** We find stark differences in the bias metric across the three clusters. In the cluster with the highest bias (cluster 3), 50% of the students are classified as “high risk” by the risk profiling algorithm, versus 3% and 19% respectively in clusters 1 and 2. These differences are statistically significant at the 0.1% level, as indicated in Table 2 in Appendix C, where we report differences between  $\bar{M}(U \setminus U_k)$  and  $\bar{M}(U_k)$  for the validation set separated for testing, including the  $p$ -values.

---

<sup>10</sup>0km indicates that student and parents share the same postal code but not the same house number.

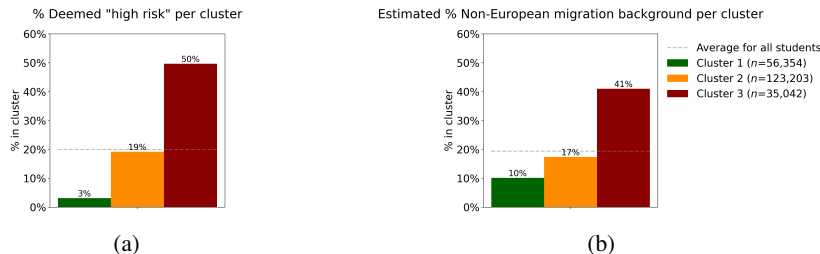


Figure 2: The percentage of students (a) deemed as “high risk” and (b) estimated as having a non-European migration background for identified clusters based on the student population of 2014, excluding students for which the risk profile was unknown ( $n = 214, 599$ ).

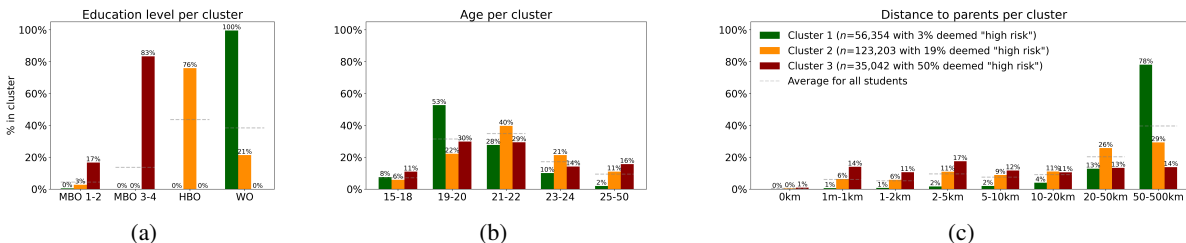


Figure 3: The percentage of students in each cluster for the three characteristics used in the risk profiling algorithm: (a) type of education, (b) age and (c) distance to parents within the student population of 2014, excluding students for which the risk profile was unknown ( $n = 214, 599$ ). For each subgroup, the average is indicated by the dashed line.

**Student characteristics per cluster:** Figure 3 illustrates the characteristics of students within each cluster. Students in cluster 3, associated with the highest average bias, are more likely to be enrolled in vocational education (MBO 1-2, MBO 3-4), be overrepresented in age groups 15-18 and 25-50, and live on average closer to their parents. In contrast, students in cluster 1, which exhibits the lowest average bias, exclusively attend university, typically live far from their parents and are more likely to be 19-20 years old. Students in cluster 2 tend to follow higher professional education (HBO), do not display distinct patterns regarding age and live slightly more often 20-50km away from their parents.

#### 4.4 Comparing Results with Aggregate Statistics on non-European Migration Background

We now examine the extent to which the identified clusters align with the non-European migration background data, using the aggregated data provided by Statistics Netherlands [51]<sup>11</sup>. We exclude cases where the distance is unknown, since the risk profiling algorithm did not apply to these students.

**Relationship between clusters and non-European migration background:** As individual-level demographic data on students’ origin are unavailable, we estimate the percentage of students with a non-European migration background per cluster using aggregate statistics. This percentage is determined using the weighted average of the proportion of students per cluster, based on the percentage of students with a non-European migration background for each combination<sup>12</sup> of the three characteristics available in [51].

There are stark differences in the proportion of students with a non-European migration background across clusters. Figure 2b shows that it is estimated that 41% of students in cluster 3 have a non-European migration background, compared to 17% in cluster 2 and 10% in cluster 1. These results indicate that clusters found by the unsupervised bias detection tool are related to the demographic group of interest.

**Proxies for non-European migration background:** In order to shed light on the relationship between clusters and non-European migration background, we investigate how characteristics of the clusters (and the risk profiling algorithm) serve as proxies for non-European migration background. This is illustrated in Figure 4. Students that are enrolled in vocational education (MBO 1-2, MBO 3-4), in the age groups 23-24 and 25-50, and living less than 10km

<sup>11</sup>The aggregated statistics are rounded to the nearest ten, hence why there are  $n = 248, 650$  reported students instead of 248, 649.

<sup>12</sup>With 4 types of education, 5 age groups and 8 distance categories, there are in total 160 possible combinations. Each possible combination is assigned to a unique cluster – e.g. if two students belong to the same combination of education level, age, and distance to parents, they belong to the same cluster.



away from their parents on average are more likely to have a non-European migration background. This corresponds to the characteristics of students in cluster 3. In contrast, university students (WO), young students, and students living far away from their parents are relatively less likely to have a non-European migration background. These characteristics correspond to students in cluster 1, and this also aligns with the logic of the risk profile: younger students and those living closer to their parents are assigned higher risk scores [5].




Figure 4: The percentage of students with a non-European migration background for the three used profiling characteristics: (a) type of education, (b) age and (c) distance to parents for students in the CUB 2014 dataset ( $n = 248,650$ ),<sup>13</sup>excluding students where the distance is unknown.

**Combinations of profiling characteristics as proxies for non-European migration background:** In addition to examining single characteristic relationships, we also investigate how specific combinations of characteristics act as a proxy for having a non-European migration background.

Figure 5a highlights that young students in university education (WO) are less likely to have a non-European migration background. This result aligns with the characteristics of cluster 1, which consists exclusively of university (WO) students and contains relatively more students aged 19-20. Relatively older students (older than 23) who follow vocational education (MBO 1-2, MBO 3-4) are also more likely to have a non-European migration background.

Figure 5b indicates that students who follow vocational education (MBO) and live relatively close to their parents are more likely to have a non-European migration background. This pattern is reflected in cluster 1, which contains MBO 1-2 and MBO 3-4 students, who tend to live closer to their parents.

Figure 5c shows a strong overrepresentation of students with a non-European migration background in the 15-18 age group who live 1m-1km from their parents. This pattern is somewhat reflected in cluster 3, which contains a higher proportion of 15-18 year old students and students living 1m-1km away from their parents.




Figure 5: The percentage of students with a non-European migration background for three (bivariate) combinations of the used profiling characteristics: (a) type of education and age, (b) type of education and distance to parents, and (c) age and distance to parents for the CUB 2014 dataset ( $n = 248,650$ ),<sup>14</sup>excluding students where the distance is unknown.

<sup>13</sup>Since the aggregated data of Statistics Netherlands is rounded to the nearest 10, we report here  $n = 248,650$  instead of the 248,649 for the CUB data.

<sup>14</sup>Supra note 10

## 5 Simulation Study to Evaluate HBAC Algorithm

This section provides simulations to support several of the design choices made when applying the HBAC algorithm, and develops empirical guidance for how to best use it. Similarly to [48], we study the performance of the HBAC algorithm on simulated data. The data-generation process is as follows. Let  $\iota_d \in \mathbb{R}^d$  be a column vector of ones, and  $I_d \in \mathbb{R}^{d \times d}$  the identity matrix. We use  $\mathcal{N}(\boldsymbol{\mu}, \Sigma)$  to refer to a (multivariate) Gaussian distribution with mean  $\boldsymbol{\mu} \in \mathbb{R}^d$  and covariance matrix  $\Sigma \in \mathbb{R}^{d \times d}$ . The features per user for cluster  $k$  are generated as:

$$\mathbf{x}_{i,k} \sim \mathcal{N}(\mu_k \iota_d, I_d), \text{ where } \mu_k \sim U(-1, 1). \quad (2)$$

This ensures the features differ on average per user. The bias metric per user is generated as:

$$m_i \sim \mathcal{N}(\eta_k, 1). \quad (3)$$

The  $\eta_k$  is defined in line with one of two scenarios:

**Scenario 1: Constant bias:** the distribution of the bias metric per user is directly generated and remains consistent across clusters, i.e.,  $\eta_k = 0$  for all  $k \in \mathcal{K}$ .

**Scenario 2: Linear increase in bias:** The expectation of the bias metric per user is directly generated as increasing from -1 to 1 (i.e.,  $\eta_k = -1 + 2 \cdot \frac{k-1}{K-1}$ , for all  $k \in \mathcal{K}$ ).

We fix the number of clusters,  $K$  and datapoints  $N$ , and keep the number of samples per cluster  $N_k$  equal per cluster.

We fit the HBAC algorithm to this synthetic data, and afterwards test if the difference between  $\bar{M}(U \setminus U_k)$  and  $\bar{M}(U_k)$  is statistically significant at a pre-specified significance level. We highlight three recommendations for this test to have the desired properties for detecting true bias, illustrated via results on the simulated dataset. We illustrate these recommendations when the simulated data have  $K = 5$  clusters,  $n = 1,000$  observations, and  $d = 2$  features.

In the simulation study, we identify three separate issues with the original implementation by Misztal-Radecka and Indurkha [48] that we argued for in Section 3. Next, we outline these issues and propose a solution.

**The necessity of sample-splitting:** Without using sample-splitting, the difference in the bias metric is measured *in-sample*, e.g. for data used to fit the HBAC algorithm. Here, there is a risk of “overfitting.” Figure 6 illustrates this. When the bias metric  $m$  is constant for each cluster, and measured on data used to fit the HBAC algorithm, there are still substantial differences between the clusters when measuring the difference between  $\bar{M}(U \setminus U_k)$  and  $\bar{M}(U_k)$  in-sample. However, these differences disappear when measuring the difference between  $\bar{M}(U \setminus U_k)$  and  $\bar{M}(U_k)$  on a sample not used to fit the HBAC algorithm, e.g., *out-of-sample*. This indicates that without sample splitting, we might falsely conclude that a classifier is biased. Figure 6 shows that when there is a difference in bias between the clusters, we are still able to detect this when measuring these differences out-of-sample (albeit slightly less likely).

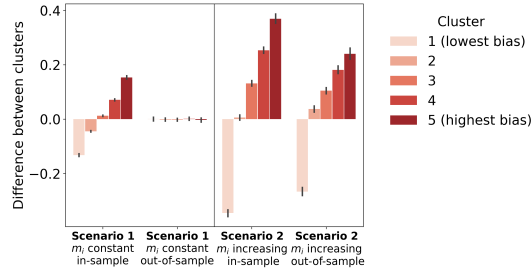



Figure 6: Average of the difference between  $\bar{M}(U \setminus U_k)$  and  $\bar{M}(U_k)$  for 1,000 simulations for the simulated dataset with  $K = 5$ ,  $n = 1,000$ , and  $d = 2$ . The error bar reflects the 95% confidence interval. In-sample refers to the data used to fit the HBAC algorithm, where out-of-sample refers to the data reserved for statistical testing via sample splitting. In Scenario 1 ( $m_i$  constant), the expectation of  $m_i$  is the same per cluster and we expect the HBAC algorithm to not find bias in the identified clusters. This is only the case when measuring the difference out-of-sample. In Scenario 2 ( $m_i$  increasing), the expectation of  $m_i$  linearly increases per cluster and we expect the HBAC algorithm to find bias both in-sample and out-of-sample.

**Testing across multiple clusters requires multiple testing correction:** In Appendix D, Figure 7a shows that without Bonferroni correction the false positive rate is much higher than the expected rate from the significance level of the test (i.e., 5% with  $\alpha = 0.05$ ), when there is no difference in bias between clusters. This shows that without Bonferroni correction, we might conclude that our classifier is biased, when it is not. With Bonferroni correction, the false positive rate does not exceed the expected rate (i.e., 5%). However, this comes at a cost: Figure 7a in the Appendix D shows that

when the bias increases per cluster, the Bonferroni correction makes it less likely that we detect differences between clusters (i.e., the true positive rate decreases).

**Different bias metrics require different null-hypotheses:** So far, our null hypothesis has been that the difference between  $\bar{M}(U \setminus U_k)$  and  $M(U)$  is zero. The validity of this null hypothesis depends on the bias metric used. For this part of the simulation study, we do not directly simulate  $m_i$ , but instead simulate the true attribute  $y$  as follows:  $y_i = \text{Bern}(p_k)$ . In the first scenario, we keep  $p_k = 0.5 \forall k \in \mathcal{K}$ , that is: the probability of  $y = 1$  is constant for all clusters. In the second scenario, we linearly increase the probability that  $y = 1$  from 0.1 to 0.9, that is:  $p_k = 0.1 + 0.8 \cdot \frac{k-1}{K-1}$ . We then fit a logistic regression to predict  $y_i$ . Here, we only consider one bias metric: the predicted value  $\hat{y}_i$  (usually used to investigate demographic parity). We consider using *accuracy* as a bias metric for this experiment in Appendix D. When using  $\hat{y}_i$  as the bias metric, we observe that even when  $y_i$  has the same distribution per cluster, there are differences in  $\hat{y}_i$  per cluster. This is because even though  $y_i$  is independent from  $x_i$ ,  $\hat{y}_i$  is not independent from  $x_i$ . Thus, if we use the null hypothesis that  $\bar{M}(U \setminus U_k)$  and  $\bar{M}(U_k)$  is zero, we are likely to conclude that our classifier is biased. In Figure 7b, we illustrate this: if  $y_i$  is constant per cluster, we incorrectly conclude there are statistically significant differences in  $\hat{y}_i$ . An alternative is to use the permutation test as suggested by Ojala and Garriga [55]. To do so, we permute the order of  $y_i$  in our data to make it independent of  $x_i$ . By repeating this  $n_{\text{perm}}$  times, we obtain samples from our null distribution, the distribution of the bias metric when  $x_i$  is independent of  $y_i$ . We can then perform a statistical test by examining the probability of observing a value of the bias metric under the null distribution. Figure 7b shows the result of using such a permutation test. The  $t$ -test frequently suggests that there is a statistically significant difference in  $\hat{y}_i$  when  $y_i$  is constant. This is not true for the permutation test.

Finally, we emphasize that the appropriate null hypothesis depends on the preference of the user. It could be that a user is not interested in differences in  $\hat{y}_i$  per cluster that are irreducible, e.g., that appear even if  $x_i$  is independent of  $y_i$ . Then, a permutation test is more appropriate than a  $t$ -test.



(a) The effect of using a Bonferroni correction

(b) The effect of using a permutation test

Figure 7: The percentage of simulations where for at least one of the clusters  $\bar{M}(U_k)$  is statistically significant from  $\bar{M}(U \setminus U_k)$  at the 5% level for 1,000 simulations for the synthetic dataset with  $K = 5, n = 1,000$  and  $d = 2$ . The error bar reflects the 95% confidence interval. The red dotted line reflects the significance level used when conducting the statistical test for the differences. We measure the difference between  $\bar{M}(U \setminus U_k)$  and  $M(U)$  out-of-sample. In Scenario 1 the expectation of (a)  $m_i$  or (b)  $y_i$  is the same per cluster. In Scenario 2 the expectation of (a)  $m_i$  or (b)  $y_i$  linearly increases per cluster. In subfigure (b), we use  $\hat{y}_i$  rather than directly simulating  $m_i$ , and we use  $n_{\text{perm}} = 1,000$  permutations.

## 6 Open Source Tool

The unsupervised bias detection tool is available on PyPI<sup>15</sup>. It implements the HBAC algorithm to detect potentially discriminated groups of users. Currently, the library supports  $k$ -means and  $k$ -modes algorithms for splitting the data into clusters. The modular design of the tool ensures that additional splitting methods can be integrated with minimal effort. To enhance accessibility, the tool is embedded within a web application, enabling people with no programming background to utilize its capabilities via an intuitive graphical interface<sup>16</sup>.

<sup>15</sup><https://pypi.org/project/unsupervised-bias-detection/>

<sup>16</sup><https://algorithmaudit.eu/technical-tools/bdt/>

## 7 Discussion

Here, we discuss the limitations of the unsupervised bias detection tool, concerning its practical application and the theoretical limitations of the simulations studies conducted.

**Limitations of clustering to audit algorithmic systems:** The unsupervised bias detection tool alone cannot establish prohibited discrimination, as human expertise is essential to evaluate bias within the context of relevant legislation. However, the tool can serve as a starting point for such deliberative assessments. For instance, the clusters it identifies could prompt further investigation into how the characteristics of individuals relate to demographic groups.

**Limitations of applying the unsupervised bias detection tool to the DUO use case:** There are several challenges associated with using aggregated data to evaluate the proportion of students with a non-European migration background in the clusters identified by the tool. For instance, this approach requires estimating the percentage of students with a non-European migration background and prevents precise analysis of the joint relationship between the bias metric, clusters and demographic groups. However, this is an unavoidable limitation due to the sensitive nature of the data<sup>17</sup>.

**Limitations relating to the simulation study conducted:** Although the simulation study is designed to be general, it relies on a relatively simple data-generation process. Future work could explore more complex scenarios, such as non-linear relationships between clusters and the bias metric, or situations where the number of data points is relatively small compared to the number of features.

## 8 Conclusion

As demonstrated through the DUO case study encompassing over 250,000 students, the unsupervised bias detection tool identifies clusters that correspond to groups of students found to be unequally treated by a risk profiling algorithm. The clustering results highlight proxy characteristics associated with a non-European migration background: vocational education (MBO) and living close to parents, as examined through aggregate statistics provided by Statistics Netherlands.

By advancing unsupervised learning methods, we enhance public knowledge on auditing algorithmic-supported decision-making processes on bias, in the absence of demographic data. The tool provides a scalable method that helps private entities, as well as internal and external auditors, to obtain inferences about seemingly neutral rules or practices that might embed unwanted biases. The tool could support further investigation into whether demographic groups face disadvantages from algorithmic-supported decision-making systems.

Through a simulation study, we highlight important considerations when applying unsupervised bias detection methods to avoid drawing incorrect conclusions. The simulations illustrate the necessity of using sample-splitting and Bonferroni corrections for valid inferences. The unsupervised bias detection method has been implemented as an open-source tool in the form of a Python package, accompanied by a graphical interface in a web application to encourage adoption by a non-technical audience. We hope that the contributions from this paper will serve the community by assisting in the detection of biases in algorithmic-supported decision-making systems.

## 9 Ethical Considerations Statement

The researchers have a contractual agreement with DUO to process CUB data. Data have been processed according to the EU's GDPR [29]. We acknowledge that if the unsupervised bias detection tool is used purely on its own without any human oversight then false conclusions could be drawn.

## References

- [1] 1994. Dutch Public Administration Law. <https://wetten.overheid.nl/BWBR0005537/2025-01-01>
- [2] 2000. Charter of Fundamental Rights of the European Union (EU) (2000/C 364/01). [https://www.europarl.europa.eu/charter/pdf/text\\_en.pdf](https://www.europarl.europa.eu/charter/pdf/text_en.pdf)
- [3] 2016. *Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the Protection of Natural Persons with regard to the Processing of Personal Data and on the Free Movement of such Data (General Data Protection Regulation)*. Official Journal of the European Union. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex:32016R0679>

---

<sup>17</sup>Supra note 7

- [4] 2024. Further research Statistic Netherlands and Algorithm Audit to CUB process. *Dutch Parliament* (2024). [https://www.tweedekamer.nl/kamerstukken/brieven\\_regering/detail?id=2024Z08699&did=2024D20430](https://www.tweedekamer.nl/kamerstukken/brieven_regering/detail?id=2024Z08699&did=2024D20430)
- [5] 2024. Internal research DUO CUB process. *Dutch Parliament* (2024). <https://www.tweedekamer.nl/kamerstukken/detail?id=2024D07565&did=2024D07565>
- [6] Muhammad Ali, Piotr Sapiezynski, Miranda Bogen, Aleksandra Korolova, Alan Mislove, and Aaron Rieke. 2019. Discrimination through Optimization: How Facebook’s Ad Delivery Can Lead to Biased Outcomes. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–30.
- [7] McKane Andrus, Elena Spitzer, Jeffrey Brown, and Alice Xiang. 2021. What We Can’t Measure, We Can’t Understand: Challenges to Demographic Data Procurement in the Pursuit of Fairness. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (Virtual Event, Canada) (FAccT ’21). Association for Computing Machinery, New York, NY, USA, 249–260. <https://doi.org/10.1145/3442188.3445888>
- [8] Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. 2016. Machine Bias. (2016).
- [9] Carolyn Ashurst and Adrian Weller. 2023. Fairness Without Demographic Data: A Survey of Approaches. In *Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*. 1–12.
- [10] Solon Barocas, Moritz Hardt, and Arvind Narayanan. 2019. *Fairness and Machine Learning: Limitations and Opportunities*. fairmlbook.org. <http://www.fairmlbook.org>.
- [11] Rachel K. E. Bellamy, Kuntal Dey, Michael Hind, Samuel C. Hoffman, Stephanie Houde, Kalapriya Kannan, Pranay Lohia, Jacquelyn Martino, Sameep Mehta, Aleksandra Mojsilovic, Seema Nagar, Karthikeyan Natesan Ramamurthy, John Richards, Diptikalyan Saha, Prasanna Sattigeri, Moninder Singh, Kush R. Varshney, and Yunfeng Zhang. 2018. "AI Fairness 360: An Extensible Toolkit for Detecting, Understanding, and Mitigating Unwanted Algorithmic Bias". <https://arxiv.org/abs/1810.01943>
- [12] Tadeusz Caliński and Jerzy Harabasz. 1974. A dendrite method for cluster analysis. *Communications in Statistics-theory and Methods* 3, 1 (1974), 1–27.
- [13] Fuyuan Cao, Jiye Liang, and Liang Bai. 2009. A new initialization method for categorical data clustering. *Expert Syst. Appl.* 36, 7 (Sept. 2009), 10223–10228. <https://doi.org/10.1016/j.eswa.2009.01.060>
- [14] Centre for Data Ethics and Innovation. 2023. *Enabling responsible access to demographic data to make AI systems fairer*. Technical Report.
- [15] Centre for Data Ethics and Innovation, Government of the United Kingdom. 2020. *Review into bias in algorithmic decision-making*. Technical Report. 151 pages.
- [16] Junyi Chai and Xiaoqian Wang. 2022. Self-supervised fair representation learning without demographics. *Advances in Neural Information Processing Systems* 35 (2022), 27100–27113.
- [17] Alexandra Chouldechova and Aaron Roth. 2020. A Snapshot of the Frontiers of Fairness in Machine Learning. *Commun. ACM* 63, 5 (2020), 82–89.
- [18] Sam Corbett-Davies, Johann D Gaebler, Hamed Nilforoshan, Ravi Shroff, and Sharad Goel. 2023. The measure and mismeasure of fairness. *The Journal of Machine Learning Research* 24, 1 (2023), 14730–14846.
- [19] Eoin Delaney, Zihao Fu, Sandra Wachter, Brent Mittelstadt, and Chris Russell. 2024. OxonFair: A Flexible Toolkit for Algorithmic Fairness. *arXiv preprint arXiv:2407.13710* (2024).
- [20] Yashar Deldjoo, Dietmar Jannach, Alejandro Bellogin, Alessandro Difonzo, and Dario Zanzonelli. 2024. Fairness in recommender systems: research landscape and future directions. *User Modeling and User-Adapted Interaction* 34, 1 (2024), 59–108.
- [21] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. 2012. Fairness through awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*. 214–226.
- [22] Marc N. Elliot, Peter A. Morrison, Allen Fremont, Daniel F. McCaffrey, Philip Pantoja, and Nicole Lurie. 2009. Using the Census Bureau’s surname list to improve estimates of race/ethnicity and associated disparities. *Health Serv Outcomes Res Method* 9, 2 (2009), 69–83.
- [23] European Parliament. 2000. Employment Equality Directive (2000/78/EC). <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex%3A32000L0078>
- [24] European Parliament. 2000. Racial Equality Directive (2000/78/EC). <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32000L0043>

- [25] European Parliament. 2004. Goods and Services Directive (2004/113/EC). <https://eur-lex.europa.eu/eli/dir/2004/113/oj/eng>
- [26] European Parliament. 2006. Recast Gender Equality Directive (2006/54/EC). <https://eur-lex.europa.eu/eli/dir/2006/54/oj/eng>
- [27] European Parliament. 2022. Digital Services Act (2022/2065). <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32022R2065>
- [28] European Parliament. 2024. AI Act (2024/1689). <https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng>
- [29] European Parliament and Council of the European Union. [n.d.]. Regulation (EU) 2016/679 of the European Parliament and of the Council. <https://data.europa.eu/eli/reg/2016/679/oj>
- [30] Janneke Gerards and Raphaële Xenidis. 2021. Algorithmic discrimination in Europe – Challenges and opportunities for gender equality and non-discrimination law. *European Commission: Directorate-General for Justice and Consumers, Publications Office* (2021). <https://doi.org/doi/10.2838/544956>
- [31] Lara Groves, Jacob Metcalf, Alayna Kennedy, Briana Vecchione, and Andrew Strait. 2024. Auditing Work: Exploring the New York City algorithmic bias audit regime. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency (FAccT '24)*. Association for Computing Machinery, New York, NY, USA, 1107–1120. <https://doi.org/10.1145/3630106.3658959>
- [32] Moritz Hardt, Eric Price, and Nati Srebro. 2016. Equality of Opportunity in Supervised Learning. *Advances in Neural Information Processing Systems* 29 (2016).
- [33] Tatsunori Hashimoto, Megha Srivastava, Hongseok Namkoong, and Percy Liang. 2018. Fairness Without Demographics in Repeated Loss Minimization. In *International Conference on Machine Learning*. PMLR, 1929–1938.
- [34] Melissa HEikkilä. 2022. Dutch scandal serves as a warning for Europe over risks of using algorithms. *Politico* (2022).
- [35] Yosef Hochberg and Ajit Tamhane. 1987. Multiple comparison procedures.
- [36] Zhexue Huang. 1998. Extensions to the k-Means Algorithm for Clustering Large Data Sets with Categorical Values. *Data Mining and Knowledge Discovery* 2, 3 (1998), 283–304. <https://doi.org/10.1023/A:1009769707641>
- [37] Rashidul Islam, Huiyuan Chen, and Yiwei Cai. 2024. Fairness without Demographics through Shared Latent Space-Based Debiasing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 38. 12717–12725.
- [38] Gareth James, Daniela Witten, Trevor Hastie, Robert Tibshirani, and Jonathan Taylor. 2023. *Unsupervised Learning*. Springer International Publishing, Cham, 503–556. [https://doi.org/10.1007/978-3-031-38747-0\\_12](https://doi.org/10.1007/978-3-031-38747-0_12)
- [39] Jon Kleinberg, Jens Ludwig, Sendhil Mullainathan, and Cass R Sunstein. 2018. Discrimination in the Age of Algorithms. *Journal of Legal Analysis* 10 (2018), 113–174.
- [40] Arvindkumar Krishnakumar, Viraj Prabhu, Sruthi Sudhakar, and Judy Hoffman. 2021. UDIS: Unsupervised Discovery of Bias in Deep Visual Recognition Models. In *British Machine Vision Conference (BMVC)*.
- [41] Arun K. Kuchibhotla, John E. Kolassa, and Todd A. Kuffner. 2022. Post-Selection Inference. *Annual Review of Statistics and Its Application* 9 (2022), 505–527.
- [42] Matt J. Kusner, Joshua Loftus, Chris Russell, and Ricardo Silva. 2017. Counterfactual Fairness. *Advances in Neural Information Processing Systems* 30 (2017).
- [43] Preethi Lahoti, Alex Beutel, Jilin Chen, Kang Lee, Flavien Prost, Nithum Thain, Xuezhi Wang, and Ed Chi. 2020. Fairness Without Demographics through Adversarially Reweighted Learning. *Advances in neural information processing systems* 33 (2020), 728–740.
- [44] Anja Lambrecht and Catherine Tucker. 2019. Algorithmic Bias? An Empirical Study of Apparent Gender-Based Discrimination in the Display of STEM Career Ads. *Management science* 65, 7 (2019), 2966–2981.
- [45] Selmer C. Larson. 1931. The shrinkage of the coefficient of multiple correlation. *Journal of Educational Psychology* 22, 1 (1931), 45.
- [46] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. 2021. A Survey on Bias and Fairness in Machine Learning. *ACM Computing Surveys (CSUR)* 54, 6 (2021), 1–35.
- [47] Anne Meuwese, Jurriaan Parie, and Ariën Voogt. 2024. How ‘algotrudence’ can contribute to responsible use of ML-algorithms. *Dutch Journal for Legal Professionals* 10 (2024).

- [48] Joanna Misztal-Radecka and Bipin Indurkha. 2021. Bias-Aware Hierarchical Clustering for detecting the discriminated groups of users in recommendation systems. *Information Processing & Management* 58, 3 (2021), 102519. <https://doi.org/10.1016/j.ipm.2021.102519>
- [49] Neda Nasiriani, Anna Squicciarini, Zara Saldanha, Sanchit Goel, and Nicola Zannone. 2019. Hierarchical Clustering for Discrimination Discovery: A Top-Down Approach. In *2019 IEEE Second International Conference on Artificial Intelligence and Knowledge Engineering (AIKE)*. 187–194. <https://doi.org/10.1109/AIKE.2019.00041>
- [50] Statistics Netherlands. 2024. Country of origin. *Website* (2024). <https://www.cbs.nl/nl-nl/maatwerk/2024/08/landenindeling-van-de-variabele-herkomstland-2022>
- [51] Statistics Netherlands. 2024. Ontvangers uitwonendenbeurs; herkomst, 2014, 2017, 2019, 2021 en 2022. *Maatwerktafel* (2024). <https://www.cbs.nl/nl-nl/maatwerk/2024/21/ontvangers-uitwonendenbeurs-herkomst-2014-2017-2019-2021-en-2022>
- [52] New York City. 2023. Automated Employment Decision Tools (AEDT). <https://www.nyc.gov/site/dca/about/automated-employment-decision-tools.page>
- [53] Court of Justice of the European Union. 1986. Case 170/84. Bilka - Kaufhaus GmbH v Karin Weber von Hartz. *EU:C:1986:204* (1986).
- [54] Court of Justice of the European Union. 2012. Case 385/11. Isabel Elbal Moreno v. Instituto Nacional de la Seguridad Social (INSS) and Tesorería General de la Seguridad Social (TGSS). *ECLI:EU:C:2012:746* (2012).
- [55] Markus Ojala and Gemma C. Garriga. 2010. Permutation Tests for Studying Classifier Performance. *Journal of Machine Learning Research* 11, 62 (2010), 1833–1863. <http://jmlr.org/papers/v11/ojala10a.html>
- [56] Manish Raghavan, Solon Barocas, Jon Kleinberg, and Karen Levy. 2020. Mitigating Bias in Algorithmic Hiring: Evaluating Claims and Practices. In *Conference on Fairness, Accountability, and Transparency*. 469–481.
- [57] Inioluwa Deborah Raji, Peggy Xu, Colleen Honigsberg, and Daniel Ho. 2022. Outsider oversight: Designing a third party audit ecosystem for ai governance. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society*. 557–571.
- [58] John Rawls. 2001. *Justice as fairness: A restatement*. Harvard University Press.
- [59] Frank Rensen. 2024. ‘Bij zo’n beetje elke tegel die we lichten, ontdekken we discriminerende algoritmen bij de overheid’, zegt de Autoriteit Persoonsgegevens. *de Volkskrant* (2024). <https://www.volkskrant.nl/tech/bij-zo-n-beetje-elke-tegel-die-we-lichten-ontdekken-we-discriminerende-algoritmen-bij-de-overheid->
- [60] Reuters. 2018. Amazon Scraps Secret AI Recruiting Tool that Showed Bias against Women. <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight-idUSKCN1MK08G>
- [61] Cynthia Rudin, Caroline Wang, and Beau Coker. 2020. The Age of Secrecy and Unfairness in Recidivism Prediction. *Harvard Data Science Review* 2, 1 (2020), 1.
- [62] Till Speicher, Muhammad Ali, Giridhari Venkatadri, Filipe Nunes Ribeiro, George Arvanitakis, Fabrício Benvenuto, Krishna P Gummadi, Patrick Loiseau, and Alan Mislove. 2018. Potential for Discrimination in Online Targeted Advertising. In *Conference on Fairness, Accountability and Transparency*. PMLR, 5–19.
- [63] Mervyn Stone. 1974. Cross-validatory choice and assessment of statistical predictions. *Journal of the royal statistical society: Series B (Methodological)* 36, 2 (1974), 111–133.
- [64] New York Times. 2017. When an Algorithm Helps Send You to Prison. <https://www.nytimes.com/2017/10/26/opinion/algorithm-compass-sentencing-bias.html>
- [65] Marvin Van Bekkum and Frederik Zuiderveen Borgesius. 2023. Using sensitive data to prevent discrimination by artificial intelligence: Does the GDPR need a new exception? *Computer Law & Security Review* 48 (2023), 105770.
- [66] Marjolein van den Brink, Peter Dunne, and Jenny Julén Votinius. 2018. *Trans and intersex equality rights in Europe—a comparative analysis*. Publications Office of the European Union.
- [67] Michael Veale and Reuben Binns. 2017. Fairer machine learning in the real world: Mitigating discrimination without collecting sensitive data. *Big Data & Society* 4, 2 (2017), 2053951717743530. <https://doi.org/10.1177/2053951717743530>
- [68] Michael Veale, Max Van Kleek, and Reuben Binns. 2018. Fairness and Accountability Design Needs for Algorithmic Support in High-Stakes Public Sector Decision-Making. In *CHI Conference on Human Factors in Computing Systems*. 1–14.

## A Risk profile

The risk score for student  $u_{1 \leq i \leq N}$  is determined as follows:

$$\text{risk score}_{u_i} = R_1 \times (R_2 + R_3).$$

$R_1$  is the risk factor assigned to different types of education (see Table 9).  $R_2$  is a risk factor determined by the combination of the distance category and the age category of the student (see Figure 10).  $R_3$  is an adjustment based on whether there were particular deviations between the age of the student known to DUO, the age known to the Municipal Basic Administration (GBA), combined with information about how long the student has been living away from home (see Table 9). Based on these risk factors a risk score is computed which are binned in 6 risk categories (see Figure 8).

Risk categories			
Coding	Description	Lower Limit	Upper Limit
1	Very High Risk	80	180
2	High Risk	60	79
3	Medium Risk	40	59
4	Low Risk	20	39
5	Very Low Risk	1	19
6	Unknown Risk	0	0

Figure 8: A tabular representation of the risk categories and the corresponding lower and upper limit of the risk score. These limits indicate when a student is deemed to be “high risk” for unduly being allocated a college grant, i.e., if the risk score for a student is determined between 80 and 180, students were allocated to risk category 1.

Type of education	Factor
MBO 1-2	1.2
MBO 3-4	1.1
HBO	1.0
WO	0.8

Figure 9: The risk factor  $R_1$  for the different types of education.

Current Age From	Current Age Up Until	Age when registered From	Age when registered Up Until	Current Age (GBA) From	Current Age (GBA) Up Until	Risk Factor $R_3$
21	22	17	18	17	18	5
21	22	17	18	19	20	0
21	22	19	20	19	20	0
23	24	17	18	17	18	15
23	24	17	18	19	20	10
23	24	17	18	21	22	0
25	65	17	18	17	18	30
25	65	17	18	19	20	25
25	65	17	18	21	22	15
25	65	17	18	23	24	0
25	65	17	18	25	65	0
25	65	19	20	19	20	25
25	65	19	20	21	22	0
25	65	19	20	23	24	0
25	65	19	20	25	65	0
25	65	21	22	21	22	15
25	65	21	22	23	24	0
25	65	23	24	23	24	0

Table 1: Values of  $R_3$  depend on the ‘Current age’ of students (referring to the current age of a student as officially registered by DUO), the ‘Age when registered’ (referring to the age known by DUO when the student was registered at an address different from their parental address) and the ‘Current age (GBA)’ (referring to the age known by the Municipal Basic Administration (GBA)).



Distance to Parent(s) (in m)	Age	Risk Score (R2)	MBO 1-2	MBO 3-4	HBO	WO
0	15-18	120	144	132	120	96
0	19-20	110	132	121	110	88
0	21-22	105	126	116	105	84
0	23-24	100	120	110	100	80
0	25-50	80	96	88	80	64
1-1000	15-18	100	120	110	100	80
1-1000	19-20	95	114	105	95	76
1-1000	21-22	85	102	94	85	68
1-1000	23-24	75	90	83	75	60
1-1000	25-50	65	78	72	65	52
1001-2000	15-18	95	114	105	95	76
1001-2000	19-20	85	102	94	85	68
1001-2000	21-22	75	90	83	75	60
1001-2000	23-24	65	78	72	65	52
1001-2000	25-50	60	72	66	60	48
2001-5000	15-18	85	102	94	85	68
2001-5000	19-20	75	90	83	75	60
2001-5000	21-22	65	78	72	65	52
2001-5000	23-24	55	66	61	55	44
2001-5000	25-50	45	54	50	45	36
5001-10000	15-18	75	90	83	75	60
5001-10000	19-20	65	78	72	65	52
5001-10000	21-22	55	66	61	55	44
5001-10000	23-24	45	54	50	45	36
5001-10000	25-50	35	42	39	35	28
10001-20000	15-18	50	60	55	50	40
10001-20000	19-20	40	48	44	40	32
10001-20000	21-22	30	36	33	30	24
10001-20000	23-24	25	30	28	25	20
10001-20000	25-50	20	24	22	20	16
20001-50000	15-18	35	42	39	35	28
20001-50000	19-20	25	30	28	25	20
20001-50000	21-22	20	24	22	20	16
20001-50000	23-24	15	18	17	15	12
20001-50000	25-50	10	12	11	10	8
50001-500000	15-18	20	24	22	20	16
50001-500000	19-20	20	24	22	20	16
50001-500000	21-22	15	18	17	15	12
50001-500000	23-24	10	12	11	10	8
50001-500000	25-50	5	6	6	5	4
Unknown	15-18	0	0	0	0	0
Unknown	19-20	0	0	0	0	0
Unknown	21-22	0	0	0	0	0
Unknown	23-24	0	0	0	0	0
Unknown	25-50	0	0	0	0	0

Figure 10: The table of  $R_2$  risk factor determined by the combination of the distance category and the age category of the student.

## B Clustering Results for CUB 2019 data

We replicate all the figures presented in Section 4 for the student population of 2019, i.e., the CUB 2019 dataset ( $n = 50,233$ ). This student population differs from the student population in 2014 in a key aspect. After 2015, students no longer received the grant for living outside of their parents home if they were following higher professional (HBO) or university (WO) education. The HBO and WO students included in the CUB 2019 dataset started receiving the college grant before 2015. Because of this difference, we separately analyze the CUB 2019 and 2014 dataset.

We follow the exact same procedure as outlined in Section 4, only changing the options for the minimum number of samples to 1,000, 2,000, 3,500, 5,000. This is because the student population of 2019 is smaller than the one in 2014.

The results are broadly in line with the results from 2014. One difference is that we find 2 instead of 3 clusters. This can be explained by the fact that there are much less HBO and WO students present in this student population.

Again, there are stark differences in the bias metric between clusters. For instance, Figure 11a shows that in cluster 2, 46% of the students are deemed "high risk", whereas only 17% are deemed "high risk" in cluster 1. This difference is statistically significant at the 0.1% level, as indicated in Table 3 in Appendix C.

We also (again) observe stark differences per cluster in the estimated percentage of students with a non-European migration background. In Figure 11b it is shown that the estimated percentage of students with a non-European migration background in cluster 2 is 36%, where it is only 14% in cluster 1.

Cluster 1 (the cluster with a relatively lower bias) consists exclusively of students following higher professional education (HBO), students that live relatively further away from their parents, and which are relatively older, i.e., exclusively consisting of students older than 21. Cluster 2 (the cluster with a relatively higher bias) consists of 96% of students following vocational education (MBO 1-2, MBO 3-4), students that live relatively closer to their parents, and are relatively younger.

Based on the aggregated data available on the origin of students, as provided by Statistic Netherlands ( $n = 50,230$ )<sup>18</sup>. We again observe how particular characteristics used in the risk profile serve as proxies for non-European migration background. Similar to 2014, students following vocational education (MBO 1-2, MBO 3-4) are relatively more likely to have a non-European migration background, and students living relatively closer to their parents, specifically less than 5km, are also relatively more likely to have a non-European migration background. This corresponds to the characteristics of students in cluster 2. Particular combinations of characteristics also observe as proxies for non-European migration background. For instance, similar to 2014, students older than 23 years who are enrolled in vocational education (MBO 1-2, MBO 3-4) are more likely to have a non-European migration background.

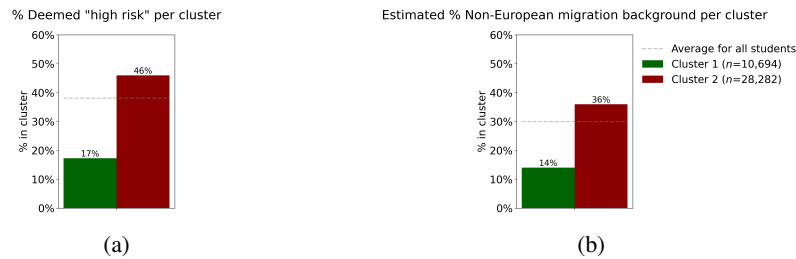


Figure 11: The percentage of students (a) deemed “high risk” and (b) non-European migration background for identified clusters based on the CUB 2019 dataset, excluding students for which the risk profile was unknown ( $n = 38,976$ ).




Figure 12: The percentage of students in each cluster for the three characteristics used in the risk profiling algorithm: (a) type of education, (b) age and (c) distance to parents within the CUB 2019 dataset, excluding students for which the risk profile was unknown ( $n = 38,976$ ). For each subgroup, the average is indicated by the dashed line.

<sup>18</sup>The aggregated statistics are rounded to the nearest ten; hence, why there are  $n = 50,230$  reported students instead of ( $n = 50,233$ ).

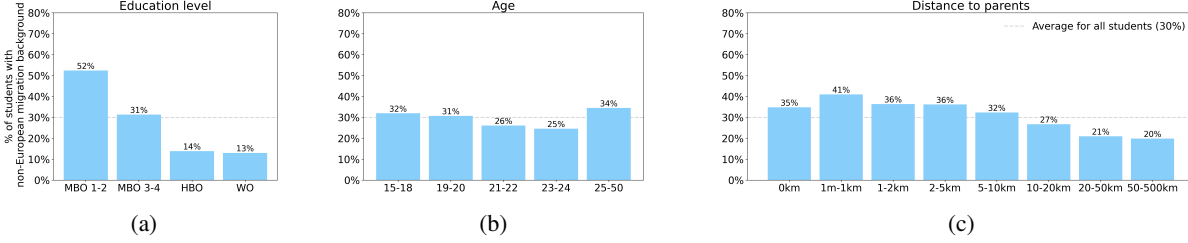


Figure 13: The percentage of students with a non-European migration background for the three used profiling characteristics: (a) type of education, (b) age and (c) distance to parents for students in the CUB 2019 dataset ( $n = 50,230$ ), excluding students where the distance is unknown.

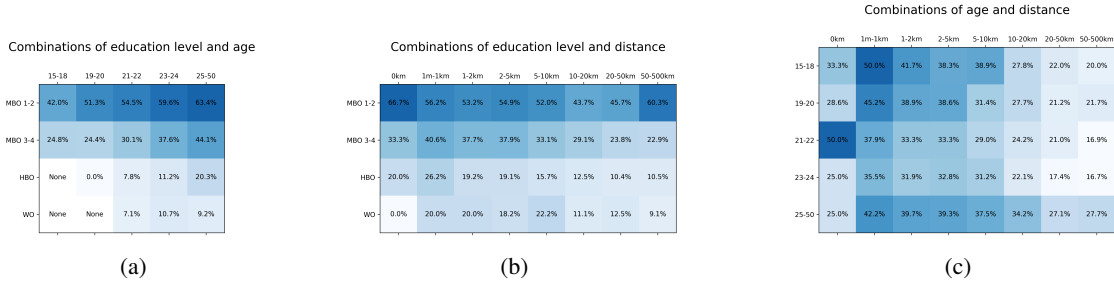


Figure 14: The percentage of students with a non-European migration background for three (bivariate) combinations of the used profiling characteristics: (a) type of education and age, (b) type of education and distance to parents, and (c) age and distance to parents for the CUB 2019 dataset ( $n = 50,230$ ), excluding students where the distance is unknown. The ‘None’ indicates there were no observations available for these categories.

### C Details on Statistical Tests

Here, we provide details for the statistical tests used to investigate whether or not differences between the bias metric per cluster are statistically significant. In Table 2 we report results for the student population in 2014, and in Table 3 for the student population in 2019.

Cluster	$n$ in cluster	High risk (%) in cluster	High risk (%) outside cluster	Difference (%)	P-value
1	11,356	3.39	25.58	-22.19	$< 10^{-16}$
2	24,650	18.88	20.83	-1.95	$1.73 \times 10^{-6}$
3	6,914	49.46	13.99	35.47	$< 10^{-16}$

Table 2: Results of testing difference in bias metrics across clusters returned by the unsupervised bias detection tool on the validation set for the year 2014 ( $n = 42,920$ ), using a Pearson’s  $\chi^2$ -test. The ‘ $< 10^{-16}$ ’ is used to indicate that the  $p$ -value was smaller than the smallest decimal that can be represented due to integer underflow.

Cluster	$n$ in cluster	High risk (%) in cluster	High risk (%) outside cluster	Difference (%)	P-value
1	2,065	17,29	46,34	-22.19	$< 10^{-16}$

Table 3: Results of testing difference in bias metrics across clusters returned by the unsupervised bias detection tool on the validation set for the year 2019 ( $n = 7,796$ ), using a Pearson’s  $\chi^2$ -test. The ‘ $< 10^{-16}$ ’ is used to indicate that the  $p$ -value was smaller than the smallest decimal that can be represented due to integer underflow. In this case, we only conduct a single statistical test, since there are only two clusters returned.

## D Additional results for Simulation Study

In this section, we report an additional result of the simulation study, complementing the results in Section 5.

In Figure 15 we report the results of using the permutation test when accuracy is used as a bias metric. Unlike the results for using predicted value  $\hat{y}_i$  as a bias metric in (see Figure 7b), we are not very likely to conclude there is a statistically significant difference in the bias metric. Figure 15 also suggests that using the permutation test in this scenario results in the correct probability of falsely rejecting the null hypothesis, given the significance level (5%).

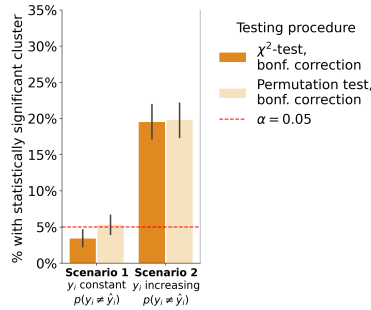


Figure 15: The percentage of simulations where for at least one of the clusters  $\bar{M}(U_k)$  is statistically significant from  $\bar{M}(U \setminus U_k)$  at the 5% level for 1,000 simulations for the synthetic dataset with  $K = 5$ ,  $n = 1,000$  and  $d = 2$ . The error bar reflects the 95% confidence interval. The red dotted line reflects the significance level used when conducting the statistical test for the differences. We measure the difference between  $\bar{M}(U \setminus U_k)$  and  $M(U)$  out-of-sample. In Scenario 1 the expectation of  $y_i$  is the same per cluster. In Scenario 2 the expectation of  $y_i$  linearly increases per cluster. As a bias metric, we use the accuracy  $p(y_i \neq \hat{y}_i)$  rather than directly simulating  $m_i$ , and we use  $n_{\text{perm}} = 1,000$  permutations.