

Diversity and Inclusion Report

This report uses randomly generated data and does not reflect the situation in a particular company.

Dr. Philippe De Brouwer

06 March, 2021

Abstract

This report informs about how diverse the team is and tries to indicate potential areas of bias. This report cannot prove or disprove bias, it can only direct the user to potential problems. Hence this report is in the first place an invitation to think. Much like feedback from other people is an invitation to think and not a diagnosis per se.

Contents

1	Introduction	2
2	The Particular Example Behind this Demo Report	3
3	Diversity	3
3.1	Gender diversity per grade	4
3.2	Age diversity	4
3.3	Diversity in nationalities	5
4	Inclusion	7
4.1	Method	7
4.2	Legend	7
4.3	The Pay-Ratios (“Pay-Gap’’)	8
5	Conclusions – How to Use This Report	9

6 Appendix 1: The Diversity Index	10
7 Appendix 2: Further Analysis	12
Bibliography	12

1 Introduction

This report is for the manager who honestly tries to forge strong and diverse teams, and fosters an inclusive atmosphere. Nobody is free from bias and we are influenced by who we are as well as by our environment. Our brain is evolved to do pattern recognition, and just as machine learning that will pick up patterns that might be true (or true in our distorted perception of the world) on average, but forego the right of everyone to be treated as an individual. Even with the best intentions, each one of us will have certain biases: both active and passive. Active bias is where one holds an explicit or implicit bias and hence will automatically value people more based on that bias. To get you started, we refer to two possible places where you can test for your own biases:

- www.tolerance.org
- Harvard University

Secondly, there is the passive bias. Passive bias occurs where (independently of your bias), the typical behaviour of one group is different than the other. To get you started on this subject, we refer to the MBTI profiles or the – more recent and more scientific – the “theory of the big five personality traits’’. The theory identifies five factors:

- openness to experience (inventive/curious vs. consistent/cautious)
- conscientiousness (efficient/organized vs. extravagant/careless)
- extraversion (outgoing/energetic vs. solitary/reserved)
- agreeableness (friendly/compassionate vs. challenging/callous)
- neuroticism (sensitive/nervous vs. resilient/confident)

In both theories men and women are typically¹ different. For example in the

¹Typically means here “on average.” For example in the MBTI profile, we find that roughly 60% of males are “Thinking,” where 60% of females are “Sensing.” More information about the MBTI profiles is for example on Wikipedia or myersbriggs.org

Big Five one finds that women score higher on extroversion, neuroticism, and agreeableness. The combination of those two last dimension implies that men will be (on average) more confident and less likely to accept that there is no salary rise for them. Therefore men will be more likely to ask promotion and salary increase and will therefore also be more likely to obtain it.

This report aims to point out if there are possible areas where bias would have influenced the salaries in your team. We have no claims on where the bias comes from but present the analyses on such way that it becomes *actionable*.

2 The Particular Example Behind this Demo Report

The data used for this report is randomly generated with the following characteristics:

- team size: 200
- percentage of females: 0.45
- average male's salary / average female's salary: 1.05
- we do not use jobID and hence assume that all people in the same Grade do the same work (and hence can be compared)

So, while this data has a bias built in for the Gender Paygap (on average males will earn 5%) more than females; there is no bias built in for any of the other dimensions that can be studies (e.g. Citizenship)

3 Diversity

First we present how diverse the team actually is. Diversity is expressed as a number between 0 and 1. The diversity is 1 if the population in the team is exactly the same as the reference population (e.g. the gender diversity equals 1 if the team has 50% males,² because that would be the same number as the reference population of the country).

²The country has 50.2% males in the work population, but that does not seem to be the case in Krakow – so we use an equal percentage for both genders.

3.1 Gender diversity per grade

In Figure 1 you can see the diversity index for the gender distribution per grade.

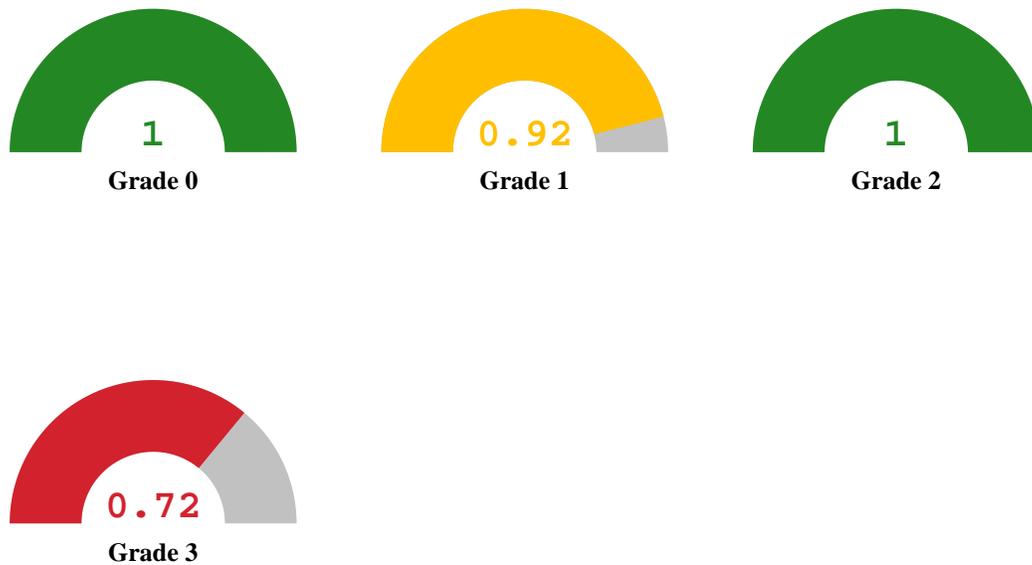


Figure 1: The diversity of the team with respect to gender per grade.

3.2 Age diversity

The diversity in function of age is represented in Figure 2. Here we do not show the distribution per grade, because one can expect a natural correlation between age, experience and grade.

Table 1: The percentages of different age groups in this team compared to the population of Poland. An equal distribution would yield a diversity of 1.

	(20,30]	(30,40]	(40,50]	(50,60]	(60,65]
team	0.49	0.43	0.07	0.01	0.01
reference	0.16	0.27	0.24	0.21	0.12

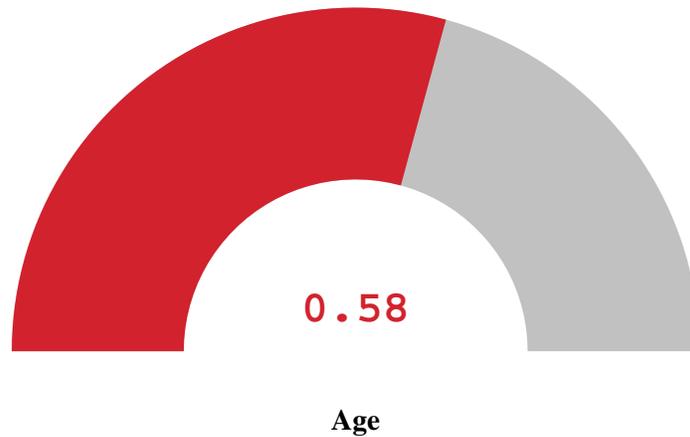


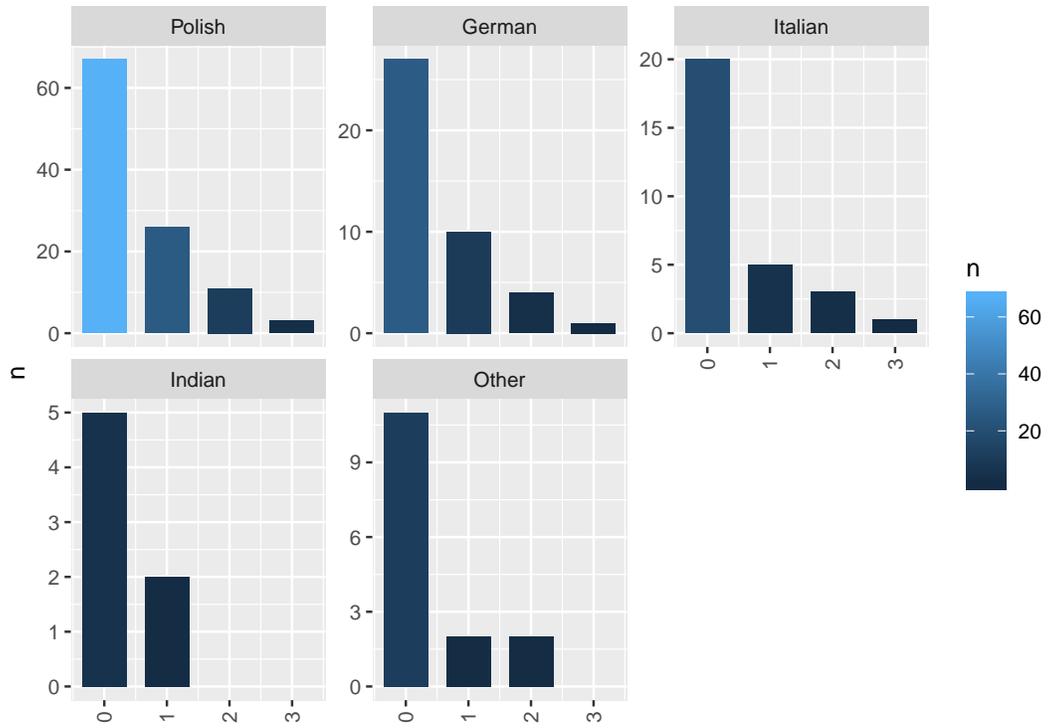
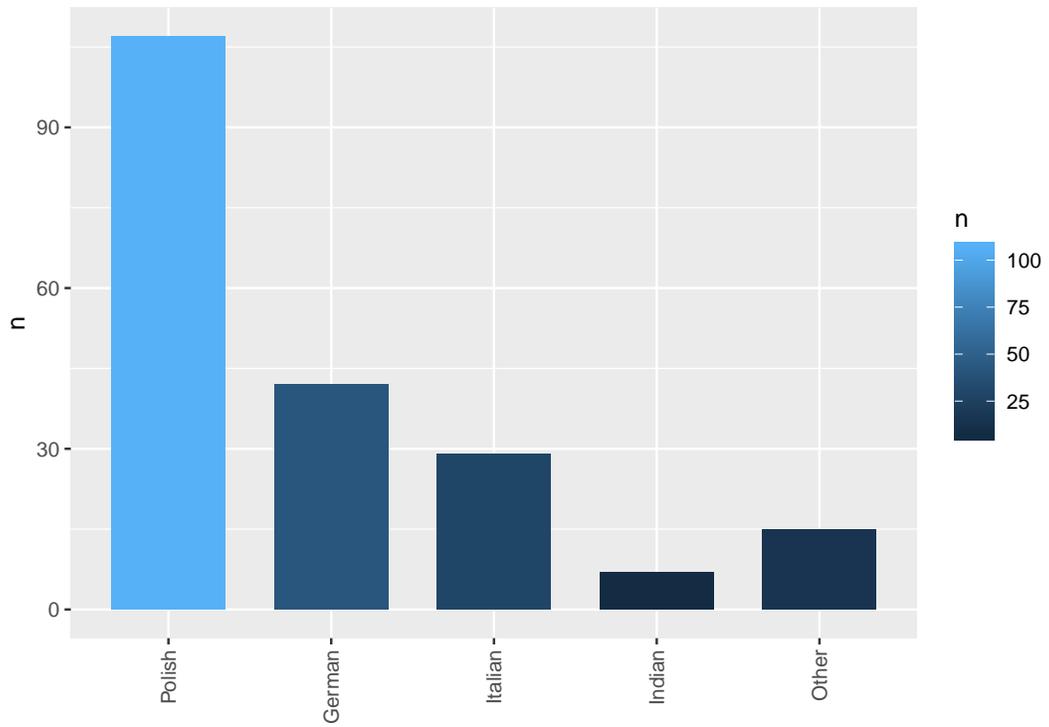
Figure 2: The diversity of the team with respect to age, assuming the age pyramid of the country as reference.

To calculate the diversity, we assume that the reference population is distributed as in Table 1. These are the numbers corresponding to Poland.

3.3 Diversity in nationalities

While ethnicity is equally valuable in diversity, nationality is measurable, available in most HR databases, and readily quantified. Few people can trace their heritage uniquely to a particular ethnic ascendance and DNA studies reveal that the concept of race is a quite blurred reality – it is rather a clustering of people that actually display a continuum instead binary racial characteristics.³

³A good introduction is here on the website of Harvard University



4 Inclusion

In previous section, we have shed some light on how diverse the team is. In this section we want to examine how inclusive the team is. We will do this via the salary and unconscious bias. Everyone has certain biases and we can investigate if certain genders or nationalities earn less in similar roles.

We can however not prove that this is the result of bias. We show only the most important control variable: experience (or rather its proxy “age”). There might be many other factors such as starting date in the firm or team, performance, productive behaviour etc. These variables are not part of this exercise, but they can be used by the manager to understand if an observed paygap is equitable or not.

4.1 Method

We realise that publishing a “paygap’” as difference between the average salary of the male and female is not very helpful as this is most likely not measuring a paygap but rather a difference in occupations.

Therefore we only present results of groups that are comparable.

4.2 Legend

- `paygap` = the ratio of median salaries of one group divided by the median of the salaries of the other group
- `NA` = numbers are too small, please look at individuals;
- `nothing` = no bias detectable;
- `.` = maybe there is some bias, but the numbers are low, check individuals;
- `*` = you should check for bias;
- `**` = bias is probably there;
- `***` = most certainly there is bias

So, there will be more stars if the probability of a bias is higher: this can be due to a higher bias and/or due to a larger sample size.

To improve readability, the column headers are abbreviated. We summarise here the columns and their content:

- `grade` = the salary grade as used in the company
- `jobID` = a unique identifier of the job category (can be abbreviated)
- `sal_F` = the median salary of the females (F), this will be different for other paygap estimations
- `sal_oth` = the median salary of the other groups (non F). The tool is open to use more than one gender. A paygap is expressed as the median of one group compared to others.
- `age_F` = the median age of the females (or `age_Po1` could be the median age of the team members with Polish citizenship)
- `age_oth` = the median age of the other groups take together (e.g. the median age of non females)
- `paygap` = the ratio of median salary earned by the selected group (e.g. females) divided by the median of the other people. So, if this number is lower than one, then median female earns less than the median non-female.
- `conf.` = the confidence level that this paygap is significant and that bias is part of the explanation (for details see in the beginning of this section).

4.3 The Pay-Ratios (“Pay-Gap”)

The “paygap” as popularly used is actually a ratio and not a gap. If the paygap for example is 0.8 in the the gender analysis and the reference group is females, then it can be understood as follows: “a woman earns 0.8 for every dollar that a man (not woman) earns.”

4.3.1 The Gender PayGap

Table 2 shows the results for the Gender paygap.

Table 2: The paygap for gender and the confidence level that this paygap is significant alongside the control variable age.

grade	jobID	sal_F	sal_oth	n_F	n_oths	age_F	age_oth	paygap	conf.
0	sales	3,971	4,274	35	37	30	30	0.93	*
0	analytic	4,121	3,968	29	29	28	30	1.04	
1	sales	8,512	8,573	21	9	31	29	0.99	

grade	jobID	sal_F	sal_oth	n_F	n_oths	age_F	age_oth	paygap	conf.
1	analytic	8,371	8,898	9	6	32	29	0.94	.
2	sales	17,684	19,300	4	5	36	36	0.92	NA
2	analytic	17,868	19,085	6	5	38	35	0.94	NA
3	sales	37,971	39,581	1	1	36	62	0.96	NA
3	analytic	NA	40,333	NA	3	NA	35	NA	NA

4.3.2 The Citizenship PayGap

Table 3 shows the results for the paygap for the main nationality compared to all other nationalities together.

Note that the data did not include a bias in function of citizenship. So any bias that appears below is due to random sampling that used to created the data. That gives you an idea about the balance between being sensitive enough to pick up a bias and not to brand random variations as biased.

Table 3: The paygap for citizenship and the confidence level that this paygap is significant alongside the control variable age.

grade	jobID	sal_oth	sal_Pol	n_oths	n_Polis	age_oth	age_Pol	paygap	conf.
0	sales	4,132	4,120	32	40	30	30	1.00	
0	analytic	4,272	3,956	31	27	29	29	0.93	.
1	sales	8,573	8,507	9	21	28	33	0.99	
1	analytic	8,645	7,904	10	5	32	27	0.91	NA
2	sales	18,945	18,058	5	4	36	37	0.95	NA
2	analytic	17,825	18,367	4	7	38	35	1.03	NA
3	sales	37,971	39,581	1	1	36	62	1.04	NA
3	analytic	40,333	40,272	1	2	51	34	1.00	NA

5 Conclusions – How to Use This Report

This report is a statistical tool to that merely points in a direction where a manager can start looking for potential bias. It cannot prove that bias exists, in all cases it is necessary to understand what is happening and why.

The reason why we cannot prove the existence of bias is that: even where we have rather large numbers and can attribute three stars, the individual performance might be so that the apparent bias is justified.

The best way to us this report is hence to use it a lead that needs investigation. The goal must not be to make sure all stars disappear, but rather that we understand and accept why certain stars are there.

6 Appendix 1: The Diversity Index

We express diversity as a number between zero and one. Our calculation is based on De Brouwer (2020) and more in particular section 36.3.1 “The Business Case: a Diversity Dashboard’’. Details can be found in the book⁴, but in summary the method has the following key aspects:

- The diversity is 0 if only one of the groups is present, and is 1 if both groups are equitably present⁵
- This calculation is similar to the established concept of entropy in physics.
- More than two categories can be used (e.g. one is not limited to two genders)
- It is possible to calibrate the probabilities so that they show maximum entropy (or diversity) for the percentages that naturally occur. Imagine for example that we are after a war and the country has 60% females. That means that we want to see a maximum diversity with a F/M ratio of 60/40. This is done in this report.

In Figure 3 we show how these properties of the diversity index evolve in function of the proportions of different classes. This example is for binary classes, because that is easier to visualise. When more classes are present, then the visualisations are to be done in a space with more dimensions.

⁴In essence, the diversity index di is defined as $di = -\sum p_i \log p_i$, where p_i are the proportions of the population in class i .

⁵What a fair and impartial distribution is must of course be carefully considered and is up the user of this report. As initial position, we used the work-population of the country Poland.

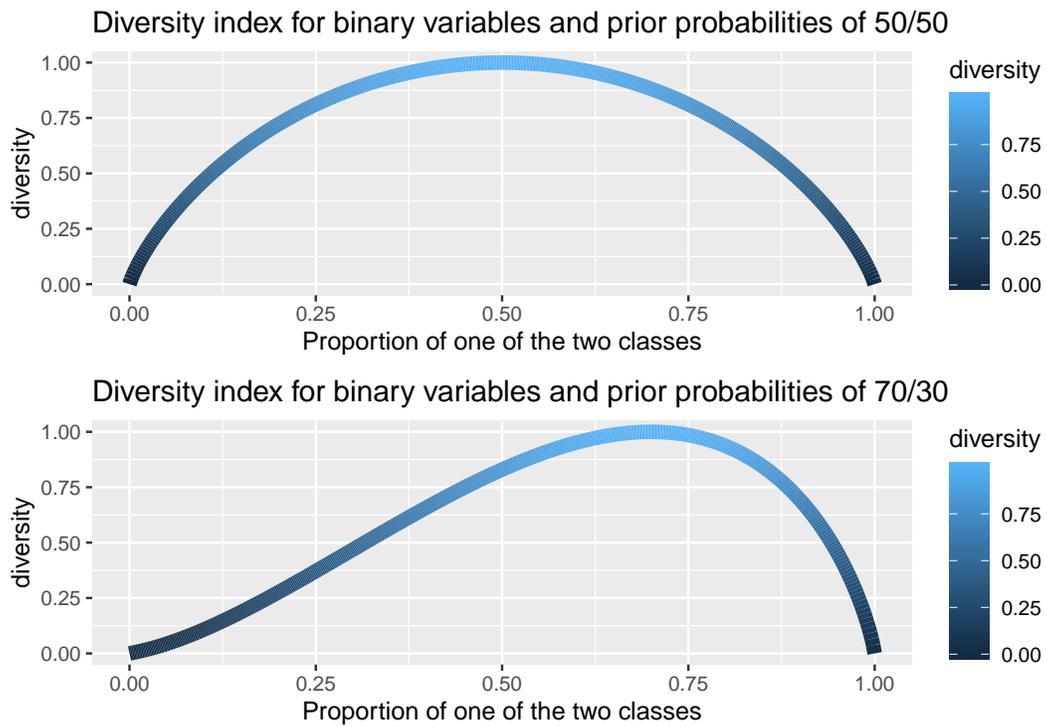
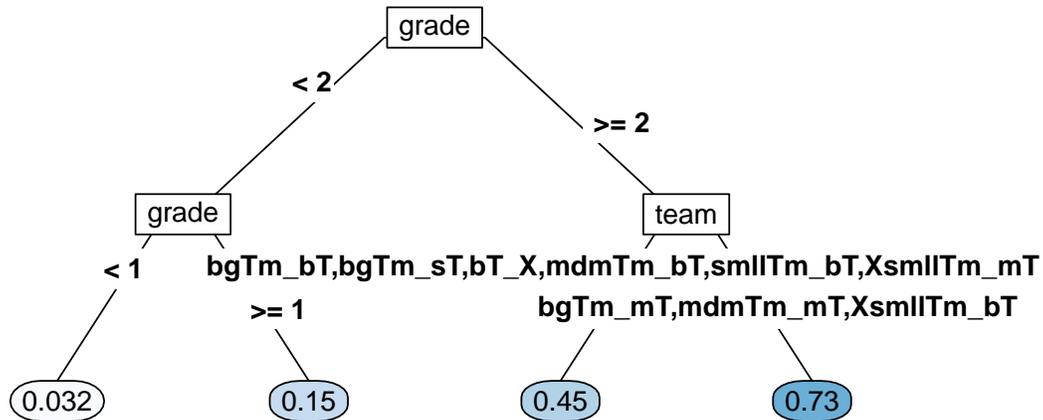


Figure 3: The diversity index illustrated for the case where there are only two possible classes (e.g. in the case of binary sex), and where the prior probabilities are respectively 50/50 (top) and 70/30 (bottom). This shows how the index reaches a maximum at a distribution equal to the prior probabilities.

7 Appendix 2: Further Analysis

What really drives your salaries? Below is a decision tree that (is over-fitted) shows the most determining variables to come to the salary (average numbers in the boxes at the bottom).



Bibliography

De Brouwer, Philippe J.S. 2020. *The Big r-Book: From Data Science to Learning Machines and Big Data*. New York: John Wiley & Sons, Ltd. <https://doi.org/https://doi.org/10.1002/9781119632757>.