

Title: Guide to using the ‘Sex, area background and ethnic group’ reports

Date: 17 December 2020

Status: Final

Security marking: Public

Guide to using the 'Applications by sex, area background, and ethnic group' reports

Quick guide to using key statistics to look at the equality of representation and offer-making

What is the purpose of these statistics?

To provide data on individual providers about any differences in representation and offer-making by sex, ethnic group, and area background.

Our national-level reporting, such as the End of Cycle Report, details the entry rates for different groups to higher education. These show some differences in representation, for instance young women are more likely to enter higher education than young men.

The national-level reporting has also examined a key stage of the admissions process for young people applying with predicted grades, in whether they are made an offer or not. This analysis shows there are large differences in the offer rates to young people from different groups (for example, different ethnic groups) but when the combination of predicted grades held and courses applied to are considered, these differences are very much smaller.

This new set of provider statistics uses developments of these methodologies to give the same kind of detail about individual providers.

What type of questions do the statistics help answer?

The new data directly answers questions such as how many people from a particular group apply to a provider, how many get offers, and how many are placed at the end of the cycle. All the statistics are calculated on the same basis so that questions about how many more applicants apply to one provider compared to another can also be answered.

Some additional statistics are reported for 18 year old applicants, such as entry rates and comparisons between offer rates and average offer rates. These statistics are designed to more directly answer questions about the equality of representation, or the equality of process, under certain assumptions.

How do I look at how equal the representation of different groups is?

The representation of different groups is covered by the tables about applicants and placed applicants. The best statistic for looking at how equal representation is between groups is called 'All placed applicants per 10,000 population'.

This statistic reports how many, from a group of 10,000 18 year olds in the UK population, are placed at that provider. This is a form of entry rate – 100 placed applicants per 10,000 population is equivalent to an entry rate of 1 per cent. This statistic has the advantage that differing population sizes are taken into account.

For example, there are many more young people in the 'White' ethnic group than in other ethnic groups, so it is quite often the case that placed applicants from the 'White' ethnic group can be the most numerous at a provider, but this can represent fewer placed applicants per 10,000 of population.

Typically, larger providers will have higher values of these statistics than smaller providers, simply because they have more placed applicants and therefore a larger share of the population.

The value of these statistics comes in comparing the values for different groups within a provider. For instance, if the number of placed applicants per 10,000 of the population is 12.0 for men and 14.5 for women, then it means that men and women do not attend the provider at the same rate. In this case, an extra 2.5 women per 10,000 population are placed at the provider compared to men (14.5 - 12.0) which means that women are around 21% more likely to enter this provider than men (14.5/12.0).

It is better to compare ratios between providers rather than the absolute values, which are driven in part by the size of the provider. For example, if another provider had placed 3.0 applicants per 10,000 population for men, and 4.0 for women, it would be best to calculate how much more likely women are to enter than men as a proportion ($4.0/3.0 = 33\%$ more likely) and compare that to the equivalent proportion at another provider.

How do I interpret different levels of representation?

Where different groups have different levels of placed applicants per 10,000 of the population, it is indicating a difference in this outcome, not what is causing it.

For instance, at the national level, differences in entry rates between groups are similar to the differences in the level and type of qualifications obtained by young people who apply. Some of the differences in representation at particular providers will relate to differences in the proportion of each group that obtain qualifications the provider judges necessary for the course, or differences in whether people apply to UCAS or to that particular provider.

National statistics also show there are large differences in the subject choices of applicants between different groups. For example, some subjects are chosen much more frequently by women than men, and other subjects are chosen more frequently by men than women. If a particular provider offers mostly one of these types of subject, this is likely to be reflected in the statistics about representation.

The placed applicants per 10,000 of the population statistics relate only to the 18 year old UK population. For some groups (for example, men and women), the relative population shares do not vary much across the UK. But where the distribution of the population varies across the UK (for example, ethnic group or area

background), and a provider has predominantly local students, then the placed applicants per 10,000 of the population will not necessarily reflect the relative chances across groups of local populations in entering that university.

Why is offer-making important?

Most applicants are placed at a university or college by making an application, receiving an offer, choosing that offer, then satisfying any terms of that offer (for example, achieving particular grades). All these stages play a part in whether an applicant is placed at a provider. The offer-making stage is important when looking at the different groups at a provider, because it is the one stage of the process that is in the direct control of the provider itself.

Why are the detailed offer-making statistics only reported for 18 year olds?

Using 18 year olds gives the best comparison because these applicants are typically similar, in that they will typically only have predicted grades at the point when the offer is made, and they are relatively large in number so that more accurate comparison statistics can be created.

How do I look at how equal the outcomes of the offer-making process are?

Offers are made in response to applications, so it is the tables containing statistics on applications and offers which help answer these questions.

The best statistics for looking at offer-making to different groups are those that use the offer rate (the proportion of applications that receive an offer) to 18 year olds, and report it as a difference against the average offer rate.

These statistics show whether the offer rate to that group is different from the offer rate to all applicants who apply to the same courses at that provider, with the same level of predicted grades. If the statistic is negative, for example -1.8, it means that the group is less likely to receive an offer, by (in this case) 1.8 percentage points, compared to all applicants with the same predicted grades and course choices. Positive figures mean the group is more likely to receive an offer.

Statistics are also provided on the actual offer rates to each group. These statistics describe the actual experience of applicants in different groups, in terms of getting an offer in response to their applications. Typically, there will be larger differences between these groups than between the rate for any particular group and the average offer rate. This is because applicants from different groups may be applying to courses with varying difficulty of securing an offer, and doing so with differing levels of predicted grades.

How do I interpret offer rates that are higher or lower than the average offer rate?

The first thing to assess is whether the difference is likely to be a random fluctuation, rather than a real difference in the underlying offer rate. In many cases, such as some ethnic groups, or for POLAR4 quintiles 1 and 2, the number of applications that the offer rate is based on can be small, especially at smaller providers.

In these cases (where the contribution of a group to the average offer rate is also likely to be small) the reference lookup table of the expected range of random variation for the average offer rate can be helpful. If the expected range of random variation is similar or larger than the difference, then it is possible that the difference may have arisen simply by chance.

If this is the case, it might be because the number of applications is particularly small. The calculation can be repeated for all six years aggregated together to test this.

In general, a larger provider might have 500 to 1,000 applications per year from POLAR4 quintile 1, or from the Black ethnic group. Typically, this would mean that a difference in any particular year would need to be around two percentage points or more to be confident it had not occurred by chance. Alternatively, a difference of around one percentage point or more each year would indicate a difference may be statistically significant over the entire period.

If the difference is statistically significant, it indicates that the difference is not accounted for by the applicants' predicted qualifications and courses applied to alone.

This does not mean that exactly equivalent applications in terms of predicted qualifications have been treated differently. In particular, the measure of predicted grades is a summary. For example, it does not reflect grades in particular subjects that may be considered important for admission to that course. For qualifications other than A levels, where there are typically fewer applicants to establish an average offer rate, there is necessarily more aggregation of grades.

This means that, for example, if one group had predicted grades in different subjects from another group, this might be a reason for why the offer rate is different from the average offer rate (since the subjects of the predicted grades are not used in the calculation of the average offer rate).

Detailed guide to the reports and how they are defined

Overview

The reports cover UK domiciled applicants for the 2015 to 2020 cycles, and all numbers reflect the position at the end of each cycle. A separate report is available for each provider. In addition, a csv data file is available which contains all the numbers in the reports, plus numbers for the 2010-2014 cycles.

Statistics are reported for three characteristics: sex, ethnic group, and area background (POLAR4). For providers in Scotland, SIMD 2016 quintile is reported as an additional area background measure. For each characteristic, a common set of statistics are reported. These statistics are split over two tables according to whether they are at applicant level or application level. Applicant-level tables report numbers of applicants and placed applicants, while application-level tables look at numbers of applications and offers.

Each characteristic has tables that cover all UK domiciles, and tables that cover only 18 year old UK domiciles. For the tables that cover only 18 year olds, a series of context statistics are included to aid interpretation of the raw numbers. In total, each report contains 16 tables (20 for providers in Scotland).

The reports also have a title page which gives summary information, and they come with a comprehensive set of technical notes and definitions which explain each of the statistics reported in the tables in detail.

June deadline applicants, June deadline placed applicants, and all placed applicants per 10,000 population

The June deadline applicants, June deadline placed applicants, and all placed applicants per 10,000 population, are context statistics that reference the number of applicants and placed applicants to the underlying 18 year old UK population. They are the number of applicants and placed applicants per 10,000 people in the population from that group. By referencing the underlying population, these statistics show how the number of applicants and placed applicants are changing in relation to the available pool of potential 18 year old applicants, and so give the chances that somebody from the group will apply to, or be placed at, a provider. The more applicants or placed applicants a provider has from a group, the bigger these statistics will be. Providers for which the chances of applying or being placed are similar for different groups, will have values for those groups that are broadly equal. Even if a provider has the same number of applicants or placed applicants in different years, the value of these statistics could change if the size of the population changes.

In 2018 there was a change in the methodology used to calculate the value of the 18 year old UK population by ethnic group. See “Notes to the report” for more details.

Average offer rate

The average offer rate places the offer rate statistic into context, by estimating what the offer rate would be if academic strength (using predicted grades) and the subject choices of those who applied to the provider were the only factors that influenced whether an applicant was made an offer. Put another way, the average offer rate represents the offer rate to all applicants applying to the same courses with the same grades as the group being looked at.

Often the offer rates at a provider for different groups can vary, but a substantial proportion of this difference is usually accounted for by variation in the predicted grades and subject choices across these groups, meaning the difference between the offer rate and the average offer rate for a group is typically smaller than the variation in offer rates across different groups.

The average offer rate does not control for other factors that may play a part in the decision to make an offer, such as performance at interview, subject of predicted qualifications and relevance to the course applied to, the exact profile of the predicted grades, personal statement, teacher references, or any other criteria (such as work experience or portfolios) that may be part of the admissions decision.

When calculating the average offer rate, we consider the predicted grades of applicants holding A levels (using their best three predicted grades), BTECs, International Baccalaureates, and SQA Highers and Advanced Highers, including combinations of these types of qualifications. Subjects are accounted for by using the 26 subject groups in the HESA JACS 3.0 classification. Further information on how the average offer rate is calculated can be found in the technical notes and definitions section of the reports.

Percentage point difference between offer rate and average offer rate

This context statistic shows how the offer rate differs from the offer rate that would be expected, as given by the average offer rate statistic. Positive values indicate where the offer rate is higher than the average offer rate, and negative values indicate where it is lower.

Even if the average offer rate is taken as a fixed value, there will be a range of random variability around the offer rate, with greater variability when the number of applications from a group is smaller. The average offer rate lookup table (supplied alongside the reports) indicates the likely range of uncertainty around the offer rate, given the number of applications it is based on, for the purposes of comparing it to an average offer rate.

This range can then be compared with the percentage point difference. When the difference is smaller than the range of uncertainty, the interpretation is that it is possible to obtain such a reading, by random variation, even if the underlying propensity for the group to receive an offer is the same as the average offer rate.

Where the difference is greater than the range of uncertainty, it means there is higher confidence the difference represents a real difference between the offer rate and what might be expected after controlling for predicted grades and subject choice. In these cases, it may be that the difference can be explained by other factors not accounted for by the average offer rate, for example performance at interview, subject of predicted qualifications and relevance to the course applied to, the exact profile of the predicted grades, personal statement, teacher references, or any other criteria (such as work experience or portfolios) that may be part of the admissions decision.

Even if the percentage point difference lies within the range of uncertainty for all years, it might be that when the offer rate and average offer rate are calculated across all years, the difference falls outside the range of uncertainty. This is most likely to be the case when the difference is consistently above or below zero.

Graphs of the percentage point difference between offer rate and average offer rate statistic for each of the characteristics covered are included in the reports.

Contribution of group to the average offer rate

The average offer rate for a group is always based, in part, on the number of applications from that group and the number of offers they receive. In some instances, this contribution can be large. This is most likely to happen when one group of applicants to a provider is much larger than the others (for many providers this will be the 'White' ethnic group, while for some providers it might be POLAR4 quintile 5). In these circumstances, the average offer rate will tend to reflect the observed patterns of offer-making to the group, resulting in a small value for the difference statistic. In these cases, the estimated range of uncertainty from random variation will be overstated.

The contribution of a group to the average offer rate context statistic helps users identify the size of the contribution a group makes to its own average offer rate. The statistic can take values between 0 and 1, with groups that make larger contributions to their own average offer rates having higher values. High values, for example 0.8 or more, mean the group is largely being compared to itself – limiting the usefulness of the statistic.

Country specific reports and reports for UCAS high, medium, and low tariff providers

Four equivalent reports covering all providers in each of the four UK countries (England, Northern Ireland, Scotland, and Wales) are also available to enable comparisons to be drawn between providers and the national picture. A further three reports covering UCAS higher, medium, and lower tariff provider groups are also available to enable comparisons to be drawn between a provider and the tariff group as a whole to which they belong.

Changes in provider codes during the reporting period

Reports are generated as per current UCAS provider codes. Where a provider has undergone a change of UCAS provider code during or since the reporting period, for example as part of a merger or demerger, the numbers reported will only pertain to those cycles where the UCAS provider code is the same as it is at the time of publication. For cycles where the code is different, values will be reported as zero.