CAPACITY TO CONTRIBUTE: INTRODUCTION TO ALTERNATIVE STATISTICAL SUMMARY MEASURES FOR USE IN THE DMI METHODOLOGY

Direct Measure of Income Refinement Working Group Paper

January 2021

CAPACITY TO CONTRIBUTE: INTRODUCTION TO ALTERNATIVE STATISTICAL MEASURES FOR USE IN THE DMI METHODOLOGY

Executive Summary

In 2020-21, the Department of Education, Skills and Employment (DESE) engaged the Australian Bureau of Statistics (ABS) to evaluate the fitness-for-purpose of alternative statistical summary measures to the median that could be used to summarise school community income, under the Direct Measure of Income (DMI) methodology for Capacity to Contribute (CTC).

As a condition for use in the DMI methodology, the alternative statistical summary measures are required to: (i) produce a single number, that represents the anticipated CTC of the school community; and (ii) allow schools to be ranked and compared based on that single number.

Based on these requirements the ABS has identified six alternative statistical measures for consideration in the DMI methodology: (i) first quartile; (ii) mid-hinge; (iii) tri-mean; (iv) mean; (v) trimmed mean; (vi) winsorised mean. See Table E.1 below for the definitions of these alternative summary measures as well as the median.

Usage of these statistical summary measures (including the median) in national and international income and economic statistics varies. The median and the mean, despite the latter not being robust to extreme values, are used extensively. This property of the mean causes it to often be paired with other, more robust statistics (such as the median). Alternatively, more robust versions of the mean, such as the trimmed mean and the winsorised mean are sometimes used instead.

The first quartile is similar to the mean in that it is also often used in income and economic statistics but is also paired alongside other percentiles in order to provide a fuller picture of the data. The mid-hinge and the tri-mean are not commonly used as statistical summary measures of income or economic data.

Analysis of the fitness-for-purpose of these six alternative statistical summary measures compared to the median, based on actual school community income data, will be discussed in a February 2021 Working Group paper. For this current paper the behaviours of the six alternative statistical summary measures under different distributional conditions were investigated using ten simulated school community income distributions.

There were two 'base' income distributions: (i) a small school (with 103 families); and (ii) a large school (with 819 families). Building upon these base distributions were four scenarios which were developed to illustrate how different statistical summary measures reflect a range of school community income distributions.

These four scenarios were based on altering the distribution of the two school communities relative to the base distributions through having:

1. A higher proportion of students with lower family income values;
2. A higher proportion of students with higher family income values;
3. A higher proportion of students with extreme family income values; or
4. No students with extreme family income values.

A summary of the results from these simulations can be found in Table E.2. It is important to note that the key findings regarding the relative differences in the summary statistics are applicable to schools with a variety of median incomes due to the above scenarios being distributional changes over the base case. In other words, the following results would hold regardless of the magnitude of the summary statistics provided that the shapes of the base distributions stay the same.

The next phase of this work investigating the alternative statistical summary measures (for the February Working Group paper) will be to assess the fitness-for-purpose of these measures for use in the DMI methodology by, using actual school community data, investigating:

1. the measures' impact on the DMI scores for schools, compared to using the median; and
2. the measures' impact on the protection of the confidentiality of persons present in the income data, compared to using the median.

The first investigation will involve comparing the DMI scores based on the median and the alternative statistical summary measures. This work includes analysing differences in the DMI score based on school size and type, state or region of the school, possibly family income distribution shape, robustness to large incomes and sensitivity to the use of imputed incomes.

The second investigation will look at the possible requirement for different and/or additional considerations and requirements for protecting the confidentiality of persons present in the income data. It will also look into whether there is any impact on the DMI scores that are not permitted to be released from the ABS DataLab, for confidentiality reasons.

Table E.1: Definitions of selected statistical summary measures.

Statistical Summary Measure	Definition
Median (or second quartile, Q2)	After sorting all family incomes in ascending order, the middle value or the average of the middle two values (depending on whether the dataset has an odd or even number of incomes respectively).
First quartile (Q1)	Similar to the median, but is instead the income value that lies 25% of the way through the ordered set of income values.[1]
Mid-hinge	The average of first quartile and the third quartile (Q3) [2] of the school community's income distribution; that is: $$\text{Mid-hinge} = \frac{Q1+Q3}{2}.$$
Tri-mean	The weighted average of the first quartile (Q1), the median (Q2) and the third quartile (Q3) of the school community's income distribution; that is: $$\text{Tri-mean} = \frac{Q1+2*Q2+Q3}{4}.$$
Mean[3]	The sum of all income values in the school community's income distribution divided by the number of income values in the distribution. That is, if there are n income values in the distribution, then: $$\text{Mean} = \frac{\sum_{i=1}^{n} x_i}{n}.$$
Trimmed mean[3,4]	Calculated by discarding extreme income values (either a certain number or proportion) from one or both sides of the school community income distribution and only using the remaining values in the calculation of the mean. If the number of discarded values is m, then the trimmed mean is the sum of the remaining (n-m) values divided by (n-m), not n.
Winsorised mean[3,4]	Similar to the trimmed mean but the extreme income values are dampened or brought closer to a pre-determined, less extreme value and are still used in the calculation of the mean, instead of being discarded.

[1] Note that if the 25% mark falls between two income values in the distribution, then Q1 is an appropriately weighted average of those two income values.
[2] Note that Q3 is calculated in a similar manner to Q1, but is instead the income value that lies 75% of the way through the ordered set of income values.
[3] If the mean, trimmed mean or winsorised mean is chosen, a decision will need to be made regarding the treatment of negative incomes. For the purposes of the analysis described in this report, any negative incomes identified in the data will be treated as zero in the calculation of means.
[4] Note that if the trimmed mean or winsorised mean were to be used in the DMI methodology, significant further investigations and decisions would be required to determine the method and amount of trimming or winsorisation to be applied to school community income distributions.

Table E.2.1: Orientation of statistical summary measures relative to the median – results from simulated school distributions

Statistical Summary Measure	Relative orientation to the Median	
	Small school distributions	Large school distributions
First quartile (Q1)	Always lesser	Always lesser
Mid-hinge	Always greater	Varies
Tri-mean	Always greater, but similar	Varies, but similar
Mean	Always greater	Always greater
Trimmed mean	Varies	Always lesser
Winsorised mean	Always greater	Always greater

Table E.2.2: Sensitivity of statistical summary measures to distributional changes relative to the median – results from simulated school distributions

Statistical Summary Measure	Sensitivity to distributional changes (compared to Median)[5]	
	Small school distributions	Large school distributions
First quartile (Q1)	Greater for decreases, Less for increases	Greater for decreases, Less for increases
Mid-hinge	Always greater	Always greater
Tri-mean	Always greater	Always greater
Mean	Similar	Similar
Trimmed mean[6]	Similar	Similar
Winsorised mean[7]	Similar	Similar

Table E.2.3: Robustness of statistical summary measures to extreme values relative to the median (from simulated school distributions).

Statistical Summary Measure	Robustness to extreme values (compared to Median)[8]	
	Small school distributions	Large school distributions
First quartile (Q1)	Similar	Similar
Mid-hinge	Less	Less
Tri-mean	Less	Less
Mean	Less	Less
Trimmed mean[6]	Less	Less
Winsorised mean[7]	Less	Less

[5] In terms of magnitude, and for increases in the proportion of either lower income or higher income families (Scenarios 1 and 2)
[6] The method used to calculate trimmed mean values was to remove of all family income values greater than the 95th percentile of the distributions. Different trimming methods will give rise to different results
[7] The method used to calculate these winsorised mean values was to bring back all family income values greater than the 95th percentile of the distribution to the 95th percentile value. Different winsorisation thresholds will yield different results
[8] In terms of magnitude, and for increases in proportion of, or removal of, families with extreme incomes (Scenarios 3 and 4)

Introduction

1. Capacity to contribute (CTC) is a measure of the anticipated capacity of a non-government school community to financially contribute to the operating cost of schooling. Under the Australian Education Act 2013 (Cth), the national school funding model uses CTC to inform Australian Government funding for most non-government schools[9].
2. Since 2020, a Direct Measure of Income (DMI) methodology has been used to inform the CTC assessment. Under this methodology, a DMI score is calculated for each non-government school based on the median income of parents and guardians of the students attending each school [see reference 1]. Due to data availability, the CTC score for each school in 2020 is the average of the DMI scores for that school for the most recent two years and from 2021 onwards, it will be the average of the DMI scores for that school for the most recent three years.
3. In 2020-21, the Department of Education, Skills and Employment (DESE) engaged the Australian Bureau of Statistics (ABS) to evaluate the fitness-for-purpose of alternative statistical summary measures that could be used instead of the median to summarise school community income.
4. As part of this engagement, DESE requested that the ABS consider such statistical summary measures as the first quartile, the tri-mean, the mean and the trimmed mean.

Requirements of the statistical summary measure used in the DMI methodology

5. In the DMI methodology, the statistical summary measure is required to be a single number that represents the anticipated CTC of the school community, relative to other schools.
6. The median meets this requirement as it is a single number that measures the ‘central tendency’ of a distribution; the use of the median in the DMI methodology enables schools to be ranked and compared relative to each other’s ‘central’ position on an income number scale.
7. It is important to note the common properties of income distributions when considering possible alternative statistical summary measures for the DMI methodology:
   - they are typically bounded on the left-hand side at $0; that is, there are no negative incomes[10];
   - they have long right-hand tails; that is, there is a rightward skew to the income data as typically there are a relatively small number of people with very large incomes;
   - they typically have a single peak (or modal value) or region of most common income values; and
   - it is possible that they have a large number of $0 income values.

[9] Some non-government schools, such as special schools, special assistance schools, sole provider schools, and majority Aboriginal and Torres Strait Islander schools, are subject to alternative funding arrangements. Also, the CTC assessment does not apply to some students, such as overseas and distance education students. These schools and students are excluded from the ABS’ analysis of alternative statistical measures.

[10] This is typically done for coherency purposes but it is possible for income to be negative (arising from losses incurred by a sole proprietor). ABS notes that, while it is likely to be extremely rare, it is possible for parents in the CTC population to report a negative Adjusted Taxable Income (ATI). Negative incomes are possible as a result of business or investment losses. With the ATI, negative incomes are less likely to occur, because the definition of ATI adjusts for financial investment losses and rental property investment losses. For more information, see the definition of ATI at the end of this paper. For more information about negative incomes in Australia, see the explanatory notes for Household Income and Wealth, Australia, 2017-18 (cat. no. 6553.0).

8. Given these typical features, any distribution of income will have at least skewness, range of values and the modal or 'central' value(s) of income as key descriptors of that distribution. It is difficult to encapsulate all of these key features into a single number, which is why the 'five number summary' - consisting of the minimum value, first quartile, median, third quartile and maximum value – or a set of percentiles – for example, the quartiles, quintiles or deciles – are commonly used to describe income distributions. While the mean and standard deviation are commonly known and understood statistical summary measures, they are typically more useful for normal or symmetrical distributions and are less commonly used to summarise (usually asymmetric) income distributions.

9. Various statistical summary measures can be used to summarise different features of income distributions, such as the inter-quartile range (the spread of the middle 50% of the distribution), skewness, kurtosis (a measure of how 'peaked' a distribution is) and the Gini co-efficient (used to measure income inequality), but these do not provide information about the relative location of income distributions on an income number scale, making them unsuitable for use in the DMI.

10. The statistical summary measures listed below all provide a different way to measure in a single number the relative position of school community income distributions on an income number scale, making them suitable candidates for consideration in the DMI methodology. For completeness, a wide range of summary measures were investigated and then narrowed down to the following, which include a mix of summary measures commonly encountered when discussing income distributions, as well as summary measures that are not often used, such as the mid-hinge and winsorised mean:

- first quartile;
- mid-hinge;
- tri-mean;
- mean;
- trimmed mean; and
- winsorised mean.

11. The mid-hinge is included because it has some similarities with the tri-mean, while the winsorised mean is included because it is considered as an alternative to the trimmed mean. The fitness-for-purpose of these statistical summary measures for use in the DMI methodology will be assessed.

Main properties of the statistical summary measures

12. In Table 1, below, the main properties of the median and the six alternative statistical summary measures are described and compared, in the context of the DMI methodology.

Table 2: Definition and properties of selected statistical summary measures.

Statistical Summary Measure	Definition	Properties
Median (or second quartile, Q2)	All family income values are sorted in ascending order and depending on whether the school's distribution of family income values consists of an odd or even number of values, the middle value or the average of the middle two values, respectively, represent the median.	• A measure of central tendency. • Focuses on the centre of the income values and does not provide any indication of the skewness or spread of the income distribution. • Robust to extreme high income values.
First quartile (Q1)	All family incomes are sorted in ascending order and the income value that lies 25% of the way through the ordered set of income values is Q1.[11]	• Not a measure of central tendency. • Places greater focus on the lower end of values in a school community's income distribution. • Robust to extreme high income values.
Mid-hinge	The mid-hinge is the simple average of the first quartile (Q1) and the third quartile (Q3) [12] of the school community's income distribution; that is: $$\text{Mid-hinge} = \frac{Q1+Q3}{2}.$$	• A measure of central tendency. • A less direct measure of the centre of a distribution since data further out on both sides of the school community's income distribution are used in its formula. • Relatively robust to extreme high income values. • Equal to the median when the median is equidistant from the Q1 and Q3 values. • Less than the median when the median is closer to the Q3 value than to the Q1 value – this occurs when a school community income distribution has proportionally more higher income values. • Greater than the median when the median is closer to the Q1 value than to the Q3 value – this occurs when a school community income distribution has proportionally more lower income values.
Tri-mean	The tri-mean is a weighted average of the first quartile (Q1), the median (Q2) and the third quartile (Q3) of the school community's income distribution; that is: $$\text{Tri-mean} = \frac{Q1+2*Q2+Q3}{4}.$$	• A measure of central tendency. • Equal to the average of the median and the mid-hinge. • Relatively robust to extreme high income values. • Behaves similarly to the mid-hinge under different distributional conditions but deviates less from the median due to weighting the median twice as much as the Q1 and Q3 values in its formula.

[11] Note that if the 25% mark falls between two income values in the distribution, then Q1 is an appropriately weighted average of those two income values.
[12] Note that Q3 is calculated in a similar manner to Q1, but is instead the income value that lies 75% of the way through the ordered set of income values.

Statistical Summary Measure	Definition	Properties
Mean[13]	The mean, or the arithmetic average, is the sum of all income values in the school community's income distribution divided by the number of income values in the distribution. That is, if there are n income values in the distribution, then: $$\text{Mean} = \frac{\sum_{i=1}^{n} x_i}{n}.$$	• A measure of central tendency because the sum of the deviations of each income value in the distribution from the mean equals zero. • Not robust to extreme high income values, nor to asymmetry in the school community's income distribution. • Equals the median when the distribution is symmetrical. • Less than the median when there is a leftward skew to the distribution. • Greater than the median when there is a rightward skew to the distribution (as is more common for income distributions).
Trimmed mean[13,14]	The trimmed mean is calculated by discarding extreme income values (either a certain number or proportion) from one or both sides of the school community income distribution and only using the remaining values in the calculation of the mean. If the number of discarded values is m, then the trimmed mean is the sum of the remaining (n-m) values divided by (n-m), not n.	• A measure of central tendency. • Due to the removal of extreme values, it is more robust to extreme values than the mean (depending on the extent of trimming). • Trimming extreme values introduces bias into the trimmed mean. • Behaves similarly to the mean under different distributional conditions but may deviate less from the median depending on the extent of trimming.
Winsorised mean[13,14]	The winsorised mean is similar to the trimmed mean but the extreme income values are dampened or brought closer to a pre-determined, less extreme value and are still used in the calculation of the mean, instead of being discarded.	• A measure of central tendency. • Due to the treatment applied to extreme values, it is more robust to extreme values than the mean (depending on the extent of the winsorisation). • Winsorising extreme values introduces bias into the winsorised mean. • Behaves similarly to the mean under different distributional conditions but may deviate less from the median depending on the extent of winsorisation.

[13] If the mean, trimmed mean or winsorised mean is chosen, a decision will need to be made regarding the treatment of negative incomes. For the purposes of the analysis described in this report, any negative incomes identified in the data will be treated as zero in the calculation of means.
[14] Note that if the trimmed mean or winsorised mean were to be used in the DMI methodology, significant further investigations and decisions would be required to determine the method and amount of trimming or winsorisation to be applied to school community income distributions.

Usage of the statistical summary measures

13. Examples of the usage of the median and the six alternative statistical summary measures in national and international income and economic statistics publications (if available) are provided in the Appendix. It should be noted that:

- The median is used extensively as a statistical summary measure in national and international income and economic statistics publications;
- The mean is also widely used as a statistical summary measure, even though it is not robust to extreme values. To partly counteract this the mean is often provided alongside other summary statistics such as the median;
- There are also a number of mean-based summary measures used such as the trimmed mean (which is used extensively to calculate Consumer Price Index (CPI) movements) and the winsorised mean (which is used by the ABS for selected earnings, expenditure and employment statistics);
- Similar to the mean, the first quartile is typically used alongside other percentiles (such as deciles or quartiles) as a statistical summary measure;
- No examples of the mid-hinge or tri-mean being used in national and international income and economic statistics publications were able to be found.

Behaviours of the statistical summary measures under various scenarios

14. In this section, the behaviours of the median and the six alternative statistical summary measures (compared to that of the median) under different distributional conditions are demonstrated using simulated school community income distributions.

15. Two simulated school community income distributions, referred to as the ‘base’ income distributions, were randomly generated to simulate the income distributions of a smaller and a larger school community. The small school base income distribution has 103 families and the large school base income distribution has 819 families. Four variations on the small and large school base income distributions were also randomly generated to illustrate how different statistical summary measures reflect a range of school community income distribution scenarios; namely:

- Scenario 1, where the school community consists of a higher proportion of students with lower family income values than the base income distribution;
- Scenario 2, where the school community consists of a higher proportion of students with higher family income values than the base income distribution;
- Scenario 3, where the school community consists of a higher proportion of students with extreme family income values than the base income distribution; and
- Scenario 4, where the school community consists of no students with extreme family income values, unlike the base income distribution.

16. While the simulated school community income distributions presented in this section were designed to reflect the kinds of income distributions one might expect to see in the CTC data, no income data of parents/guardians of students attending any school in Australia were used in the creation of the simulated school community income distributions.

17. ABS notes that the simulated income distributions for large schools resulted in a relatively high median income. While the magnitude of the summary statistics in this example may not be representative of a majority of schools, it is important to note that the key findings regarding the relative differences in the summary statistics are still applicable to schools with a variety of median incomes. This is because the various scenarios can be considered to introduce compositional or distributional changes over the base case and thus hold regardless of the magnitude of the summary statistics provided that the shapes of the base distributions stay the same.

18. It should be noted when viewing the results in the following sub-sections that the school sizes and scenarios were only generated for demonstrative purposes and are not intended to explore all the potential variability in the simulated values (and distributions), and thus the possible variability in behaviour for the alternative summary measures, under different situations.

Median

19. The behaviour of the median in different scenarios is presented and compared to that of the six alternative statistical summary measures in the figures provided throughout this section. In summary and as expected, compared to the base case for both the small (Figures 1-5) and large school (Figures 6-10) distributions, the median:

- is lower in the low income scenario (Scenario 1);
- is higher in the high income scenario (Scenario 2); and
- demonstrates robustness to extreme values in the extreme outlier scenario (Scenario 3), where it increases by a small amount. The robustness of the median is also demonstrated in the no outlier scenario (Scenario 4), which results in a slight decrease to the median.

First quartile

20. The behaviour of Q1 in different scenarios is presented and compared to that of the median in Figures 1-10 below using the simulated school community income distributions.

21. Compared to the base case, under the low income scenario, Q1 decreases more than the median in response to the higher proportion of lower family incomes in the distributions. For small schools, this is visible by comparing Figures 1 and 2, and for large schools, Figures 6 and 7. This is to be expected as Q1 is more affected than the median by the lower 50% of the distribution.

22. In contrast, Q1 increases less than the median in response to the higher proportion of higher family incomes in the distributions, under the high income scenario. This is shown in Figure 3 for the small school and Figure 8 for the large school. This is also to be expected as Q1 is less affected than the median by the upper 50% of the distribution.

23. The robustness of the Q1 to extreme values is demonstrated and is similar to that of the median. The presence of more extreme family income values under Scenario 3 results in only a slight increase in the Q1 (see Figure 4 for the small school and Figure 9 for the large school distribution). Similarly, the presence of no extreme family incomes under Scenario 4 results in only a slight decrease in the Q1 (see Figure 5 for the small school and Figure 10 for the large school distribution).

Figure 1: Small school, base case

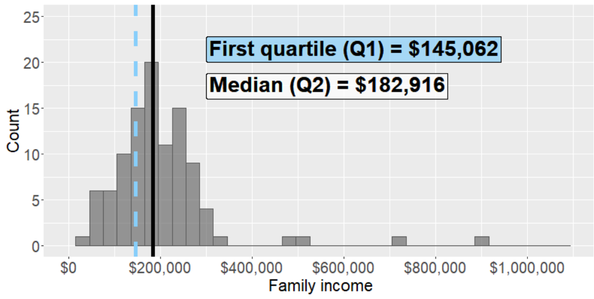

Figure 2: Small school, low income scenario (1)

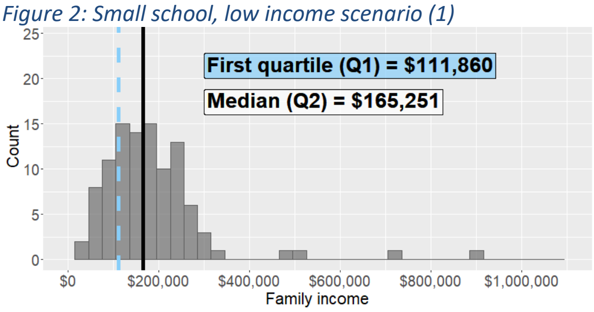

Figure 3: Small school, high income scenario (2)

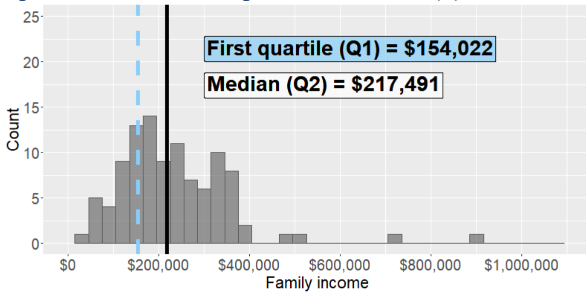

Figure 4: Small school, extreme outlier scenario (3)

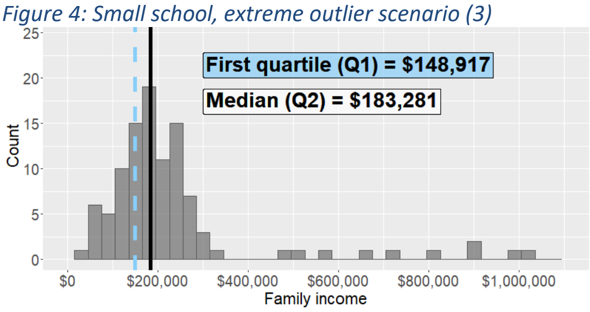

Figure 5: Small school, no outlier scenario (4)

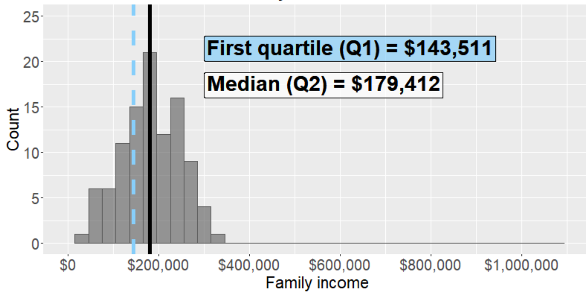

Figure 6: Large school, base case

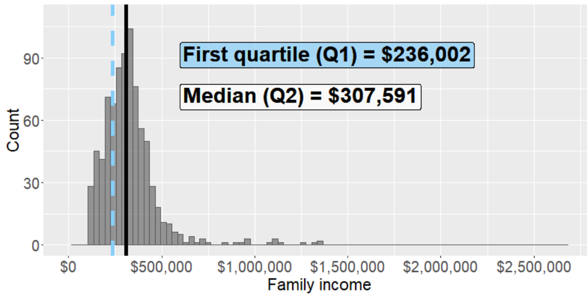

Figure 7: Large school, low income scenario (1)

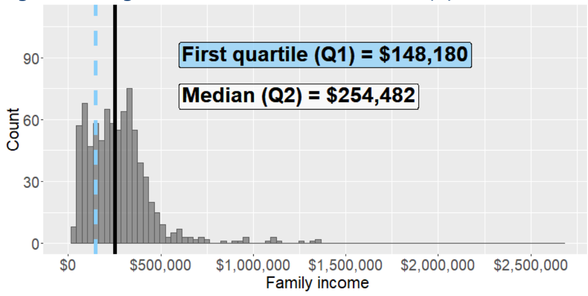

Figure 8: Large school, high income scenario (2)

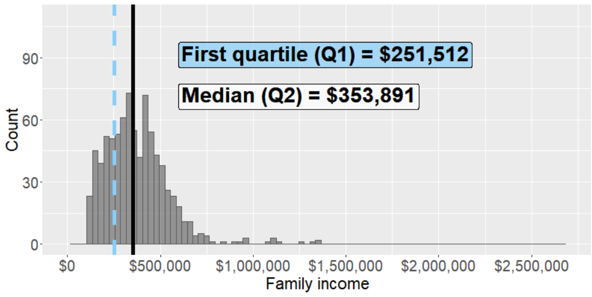

Figure 9: Large school, extreme outlier scenario (3)

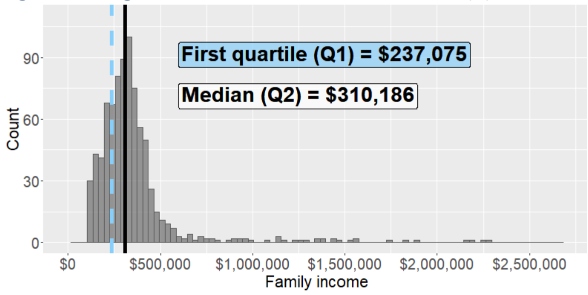

Figure 10: Large school, no outlier scenario (4)

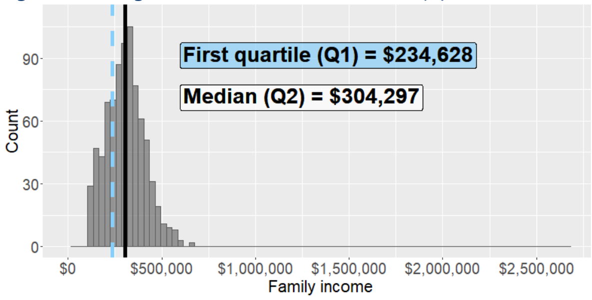

Mid-hinge

24. The behaviour of mid-hinge in different scenarios is presented and compared to that of the median in Figures 11-15 (small school distributions) and Figures 16-20 (large school distributions) below.

25. In all of the simulated income distributions for the small school, the mid-hinge and the median are similar in value but the mid-hinge is always greater than the median, which is not the case for the large school scenarios. This is the result of there being a slightly greater rightward skewness to the income distribution for the small school compared to the large school. Any change in the relationship between the mid-hinge and the median between the small school scenarios and the corresponding large school scenarios can be attributed to this difference in skewness.

26. For the large school income distribution, the median and the mid-hinge both decrease from the base case (Figure 16) in response to the higher proportion of lower family incomes under the low income scenario (Figure 17), but the magnitude of the movement in the mid-hinge value is slightly greater than that of the median – indicating that the mid-hinge is slightly more sensitive to distributional changes. Similarly, for the high income scenario (Figure 18), the mid-hinge increases by more than the median, where there are a relatively more high income families in the distribution.

27. When considering the extreme outlier scenario (Figure 14), the mid-hinge increases more than the median in response to the presence of more extreme income values in the distribution, compared to the base case (Figure 11). This indicates that the mid-hinge is less robust to extreme values than the median. This is also true of the large school distribution, as shown by comparing Figures 16 and 19.

Figure 11: Small school, base case

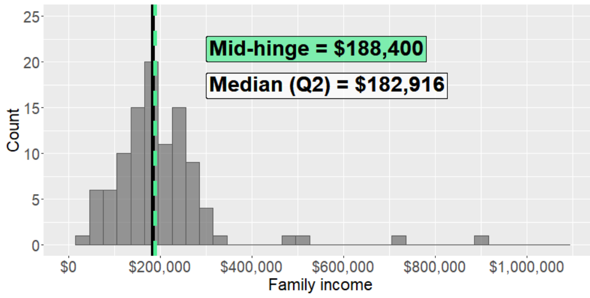

Figure 12: Small school, low income scenario (1)

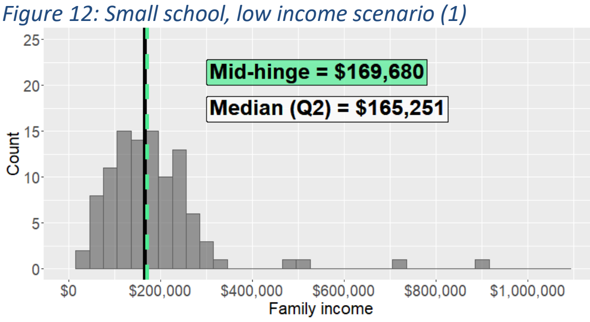

Figure 13: Small school, high income scenario (2)

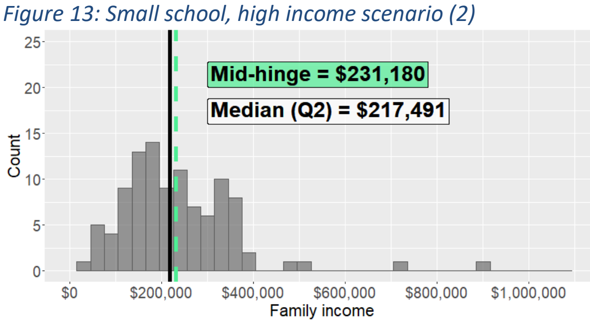

Figure 14: Small school, extreme outlier scenario (3)

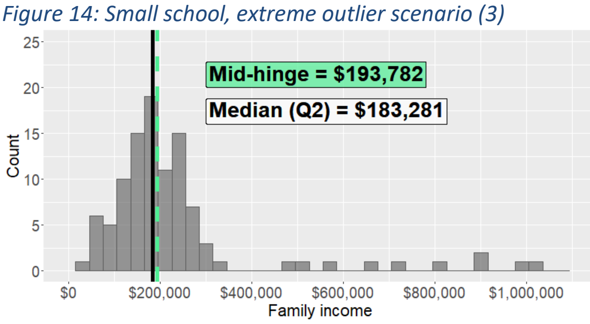

Figure 15: Small school, no outlier scenario (4)

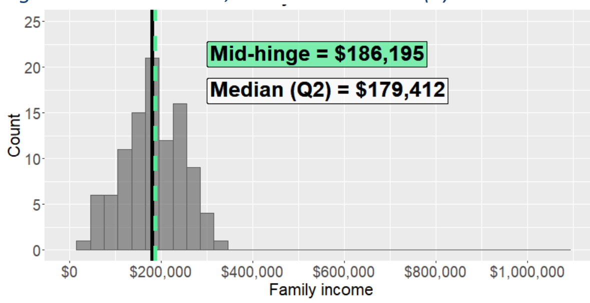

Figure 16: Large school, base case

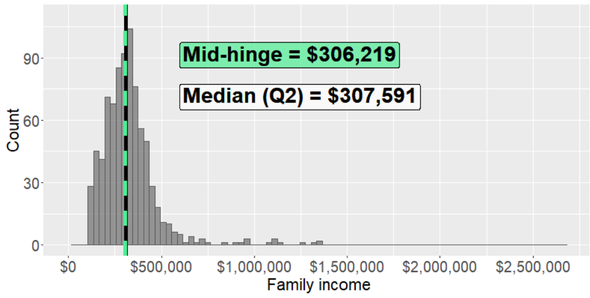

Figure 17: Large school, low income scenario (1)

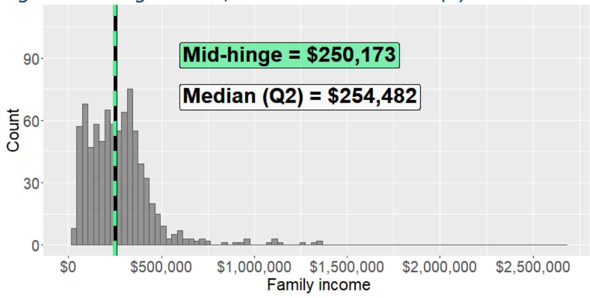

Figure 18: Large school, high income scenario (2)

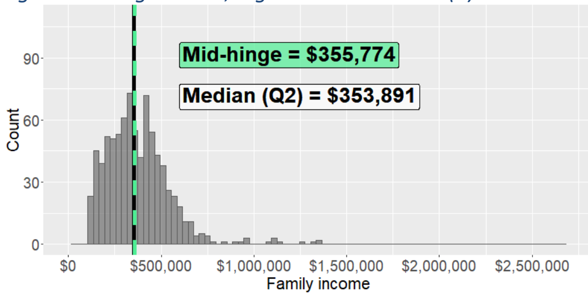

Figure 19: Large school, extreme outlier scenario (3)

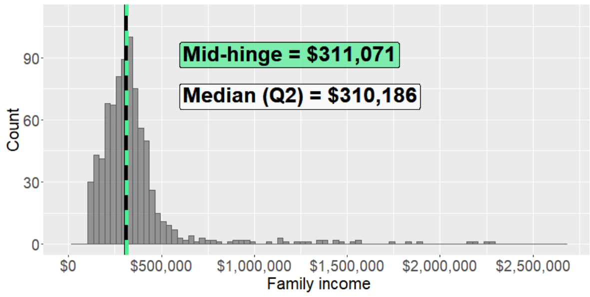

Figure 20: Large school, no outlier scenario (4)

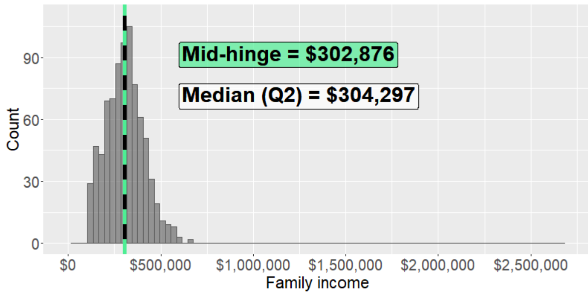

Tri-mean

28. The behaviour of the tri-mean in different scenarios is presented and compared to that of the median in Figures 21-25 for the small school, and in Figures 26-30 for the large school simulated income distributions below.

29. For the small school distributions, the tri-mean and the median are similar in value but the tri-mean is always greater than the median, indicating a slight rightward skewness to the small school income distributions (Figures 21-25).

30. For the small school distributions, under the low income scenario (Figure 22), the median and the tri-mean both decrease to reflect the higher proportion of lower family incomes, but the magnitude of the movement in the tri-mean is slightly greater than that of the median – indicating that the tri-mean is more sensitive to distributional changes. Similarly, the tri-mean also increases by more than the median does, under the high income scenario presented in Figure 23.

31. Under the extreme outlier scenario (Figure 25), the tri-mean increases by more than the median in response to the presence of more extreme family income values in the distribution. This indicates that the tri-mean is slightly less robust to extreme values than the median.

Figure 21: Small school, base case

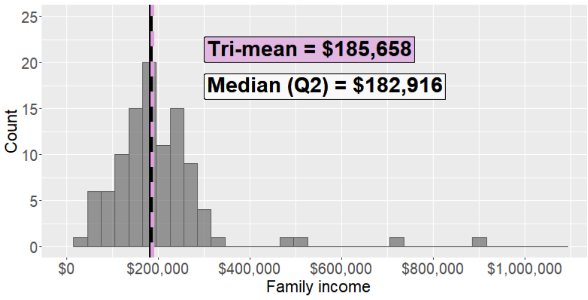

Figure 22: Small school, low income scenario (1)

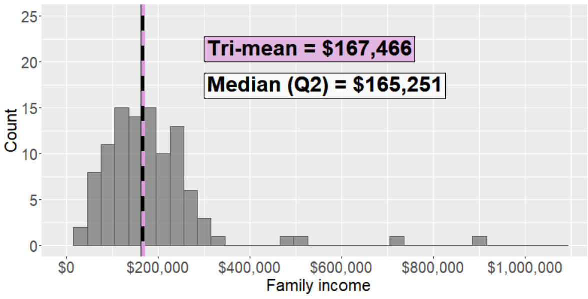

Figure 23: Small school, high income scenario (2)

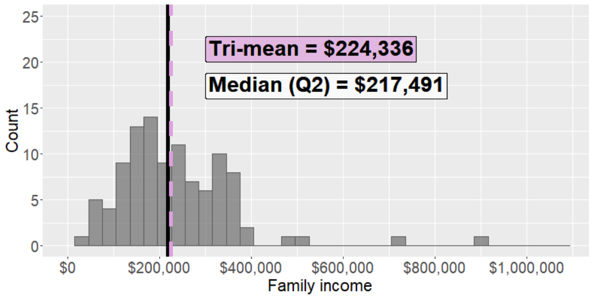

Figure 24: Small school, extreme outlier scenario (3)

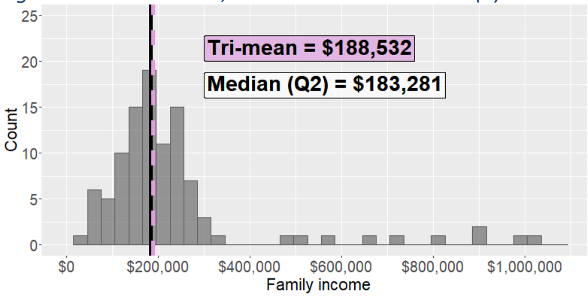

Figure 25: Small school, no outlier scenario (4)

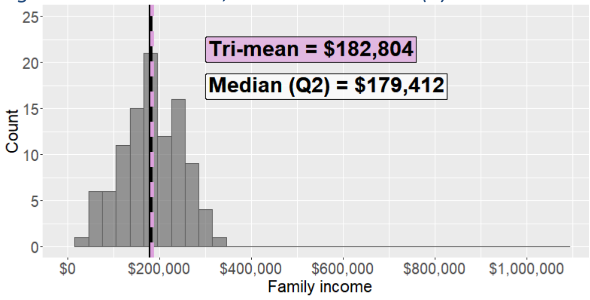

Figure 26: Large school, base case

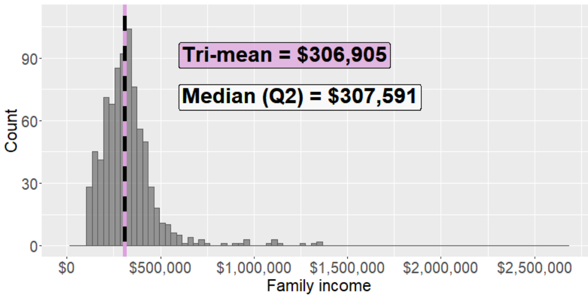

Figure 27: Large school, low income scenario (1)

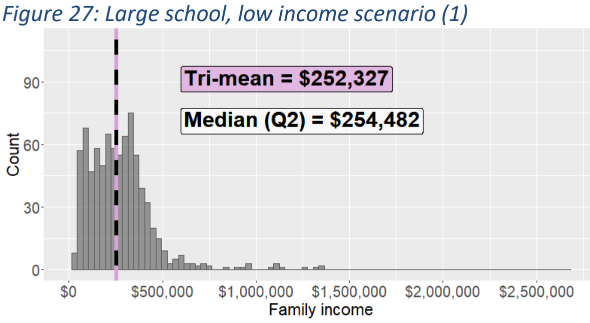

Figure 28: Large school, high income scenario (2)

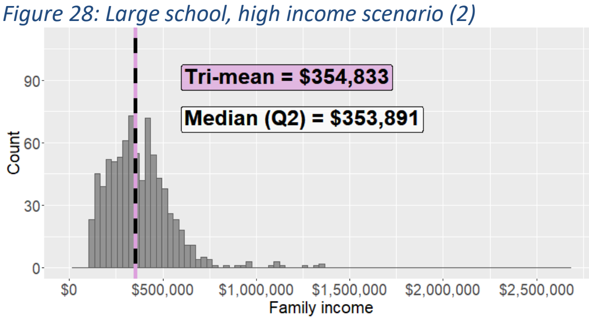

Figure 29: Large school, extreme outlier scenario (3)

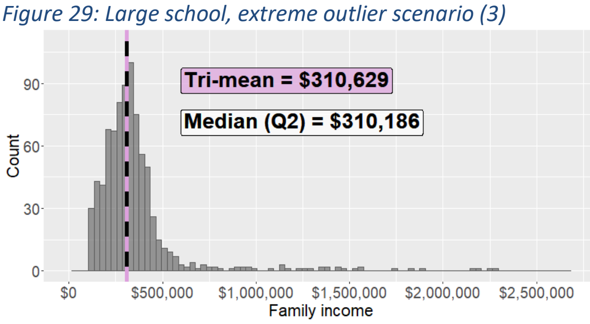

Figure 30: Large school, no outlier scenario (4)

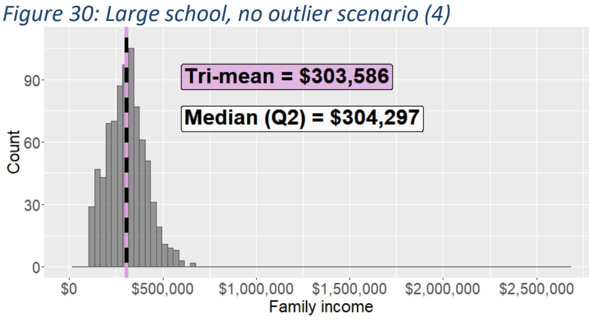

Mean

32. The behaviour of the mean in different scenarios is presented and compared to that of the median using the simulated income distributions for small schools (Figures 31-35) and large schools (Figures 36-40) below.

33. The mean and the median both decrease in the low income scenarios (Figures 32 and 37) compared to the base cases (Figures 31 and 36), and increase in the high income scenarios (Figures 33 and 38). In both scenarios, the magnitude of the change in the mean is similar to that of the median, indicating that the mean and the median share a similar sensitivity to distributional changes under these two scenarios.

34. The most marked differences in behaviour between the mean and the median can be seen for the small school under the extreme outlier scenarios (Figures 34 and 39). The presence of more extreme family income values in the distributions causes the mean to increase by approximately $37,400 for the small school distribution and $30,600 for the large school distribution while the median only increases by approximately $300 and $2,600 respectively. This demonstrates the lack of robustness of the mean to the number and magnitude of extreme values present in the distribution.

35. In the no outlier scenario (Figures 35 and 40), the absence of any extreme family income values results in an almost symmetrical income distribution, and therefore the mean and the median are almost equal in value. The presence of extreme family income values in all of the other scenarios, for both the small and large school distributions, cause the mean to be greater than the median.

Figure 31: Small school, base case

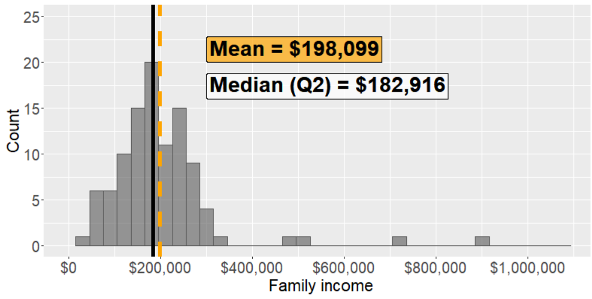

Figure 32: Small school, low income scenario (1)

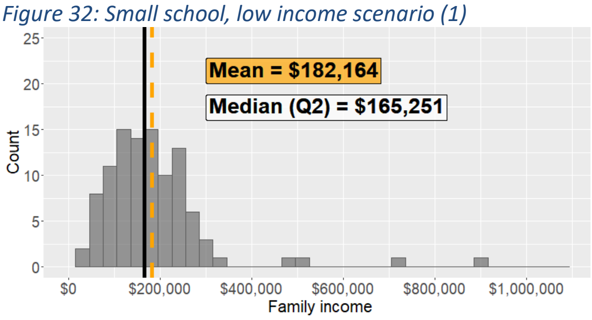

Figure 33: Small school, high income scenario (2)

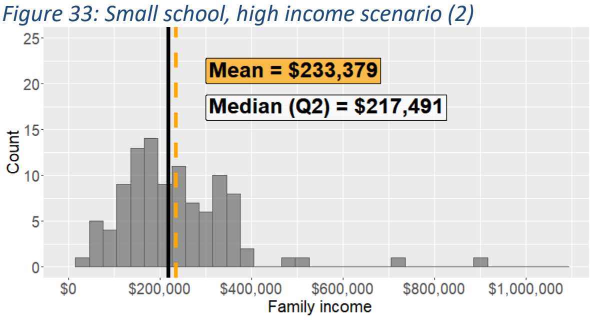

Figure 34: Small school, extreme outlier scenario (3)

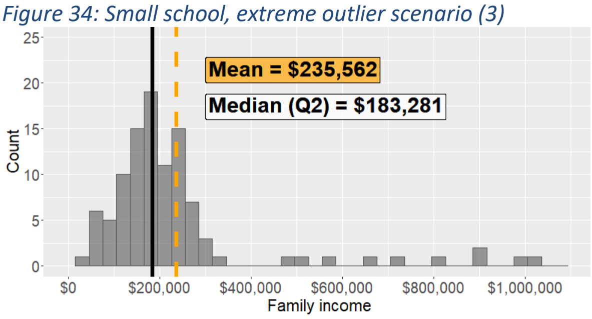

Figure 35: Small school, no outlier scenario (4)

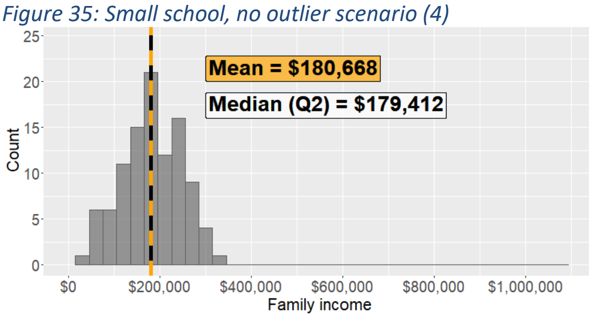

Figure 36: Large school, base case

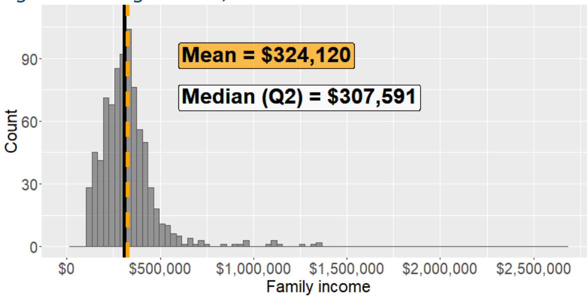

Figure 37: Large school, low income scenario (1)

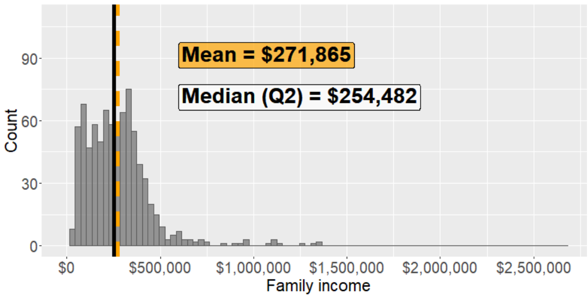

Figure 38: Large school, high income scenario (2)

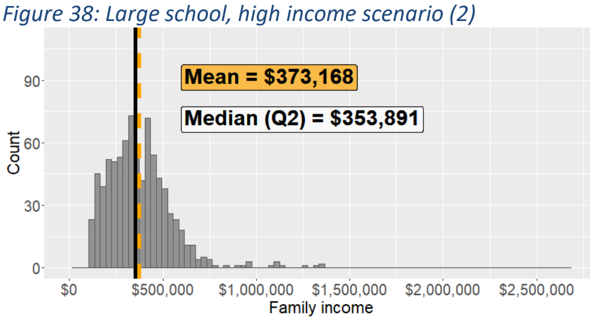

Figure 39: Large school, extreme outlier scenario (3)

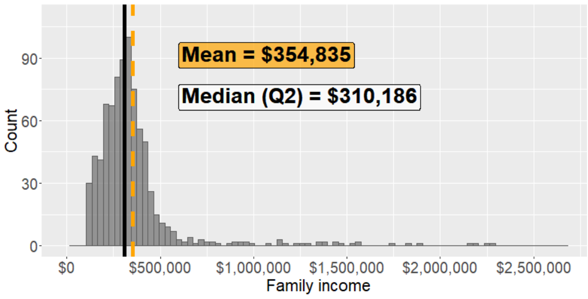

Figure 40: Large school, no outlier scenario (4)

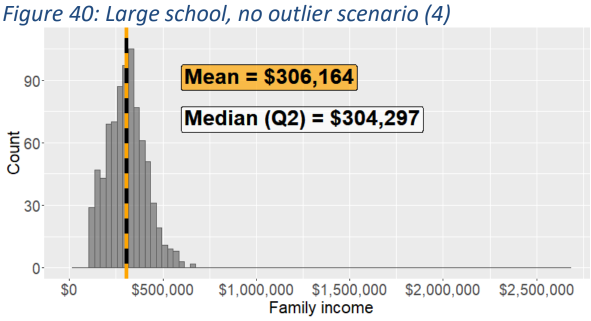

Trimmed mean

36. The behaviour of the trimmed mean under different distributional conditions is presented and compared to that of the median using the simulated income distributions for small schools (Figures 41-45) and large schools (Figures 46-50) below. The method used to calculate these trimmed mean values was to remove of all family income values greater than the 95th percentile of the distribution.

37. The method of trimming has the effect of the trimmed mean being less than the median for the large school distributions under all of the scenarios (Figures 46-50). For the small school distributions, the trimmed mean is less than the median for all except the extreme outlier scenario (Figure 44). This is simply a result of the number of the outliers added being large enough that trimming to the 95th percentile still had a larger right tail than in the base case. Adding fewer outliers might have resulted in the relationship between the trimmed mean and the median being different.

38. Under the extreme outlier scenario (Figures 44 and 49), the continued presence of more extreme family income values in the distribution, even after trimming to the 95th percentile removed several outliers, causes the trimmed mean to increase significantly more than the median – demonstrating that the trimmed mean is less robust to any remaining extreme values

than the median under this trimming method. An alternative trimming method could enable the resulting trimmed mean to be more robust to extreme values.

39. Considering the low income scenario (Figures 42 and 47), the trimmed mean and the median both decrease with a similar magnitude in response to the higher proportion of lower family incomes, compared to the base case. Similarly, under the high income scenario (Figures 43 and 48), both the trimmed mean and median increase by a similar amount when a larger proportion of high income families are present in the distribution. This indicates that the trimmed mean and the median share a similar sensitivity to these distributional differences.

40. If the trimmed mean were used in the DMI methodology, significant further investigations and decisions would be required to determine the method and amount of trimming to be applied to school community income distributions.

Figure 41: Small school, base case

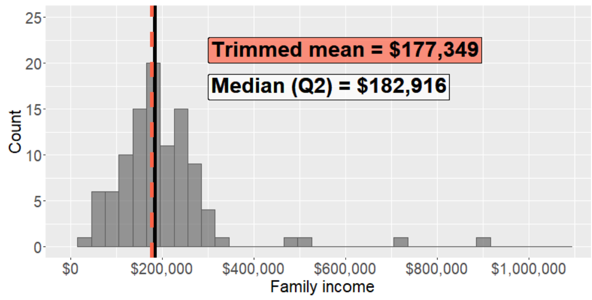

Figure 42: Small school, low income scenario (1)

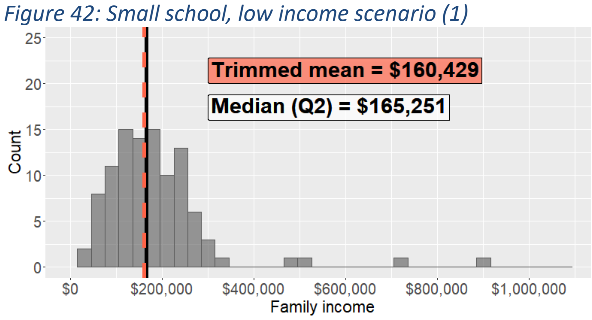

Figure 43: Small school, high income scenario (2)

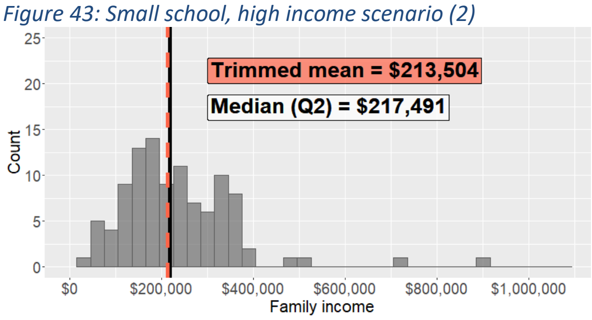

Figure 44: Small school, extreme outlier scenario (3)

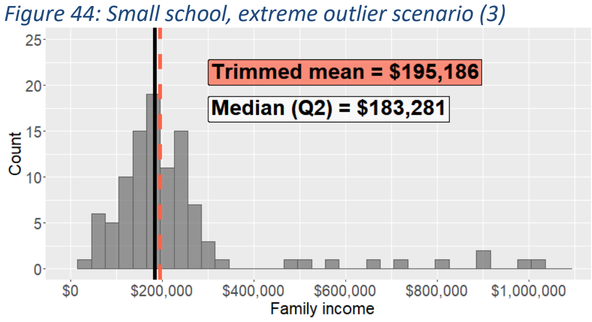

Figure 45: Small school, no outlier scenario (4)

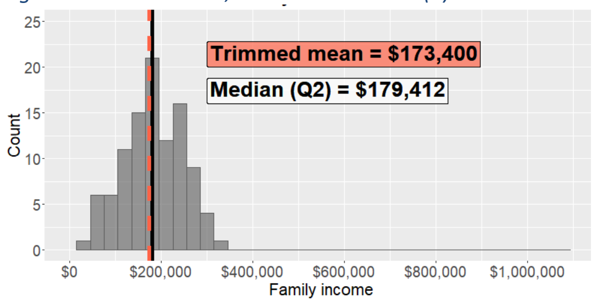

Figure 46: Large school, base case

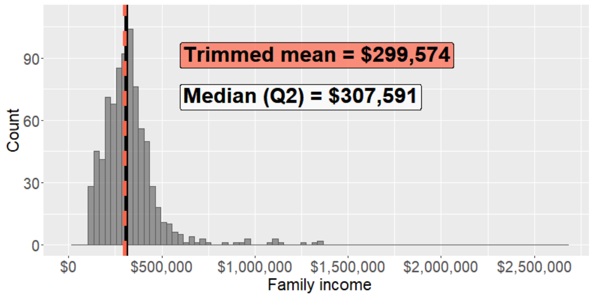

Figure 47: Large school, low income scenario (1)

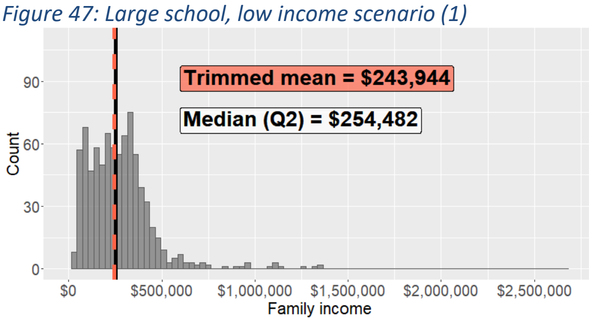

Figure 48: Large school, high income scenario (2)

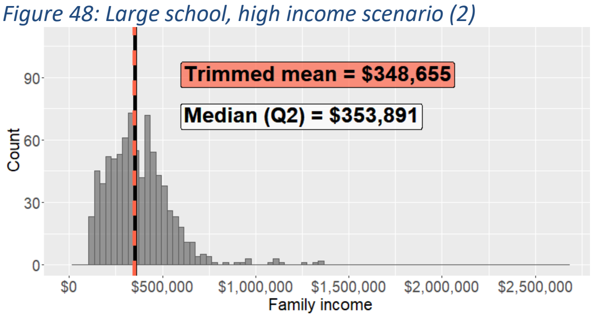

Figure 49: Large school, extreme outlier scenario (3)

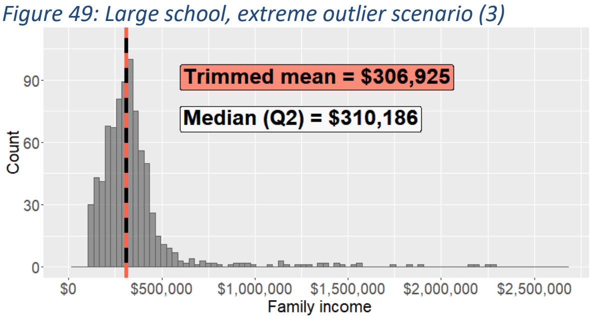

Figure 50: Large school, no outlier scenario (4)

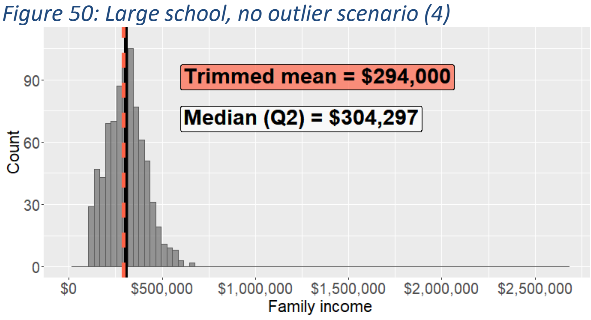

Winsorised mean

41. The behaviour of the winsorised mean under different distributional conditions is presented and compared to that of the median using the simulated income distributions for small schools (Figures 51-55) and large schools (Figures 56-60) below. The method used to calculate the winsorised mean values reduced all family income values greater than the 95th percentile down to the value of the 95th percentile.

42. The method of winsorisation has the effect of the winsorised mean being greater than the median in all of the small school income distributions.

43. The lack of extreme values in the no outlier scenario compared to the base case shows that, for the small school distributions, the method of winsorisation counteracts the slight rightward skewness present in the 'main bulk' of the distribution given that the winsorised mean and the median are essentially equal in value (Figures 51 and 55). For the large school distributions, however, the method of winsorisation results in the winsorised mean being less than the median (Figures 56 and 60).

44. Where there is a higher proportion of more extreme family income values in the distribution (Scenario 3, Figures 54 and 59), the winsorised mean increases significantly more than the median – demonstrating that the winsorised mean is less robust to extreme values than the

median under this winsorisation method. An alternative winsorisation method could enable the resulting winsorised mean to be more robust to extreme values.

45. Under the low income scenario, the winsorised mean and the median both decrease with a similar magnitude in response to the higher proportion of lower family incomes (Figures 52 and 57). Similarly, under the high income scenario, the winsorised mean and median increase by a similar amount (Figures 53 and 58). This indicates that the winsorised mean and the median share a similar sensitivity to these different income distributions.

46. If the winsorised mean were used in the DMI methodology, significant further investigations and decisions would be required to determine the method and amount of winsorisation to be applied to income distributions.

Figure 51: Small school, base case

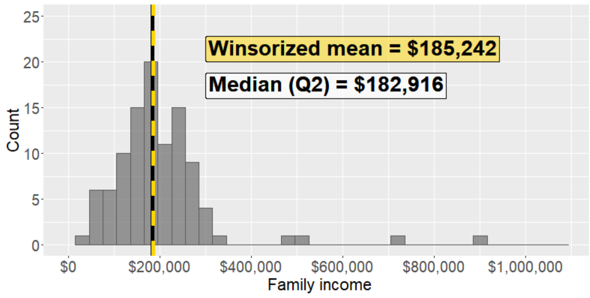

Figure 52: Small school, low income scenario (1)

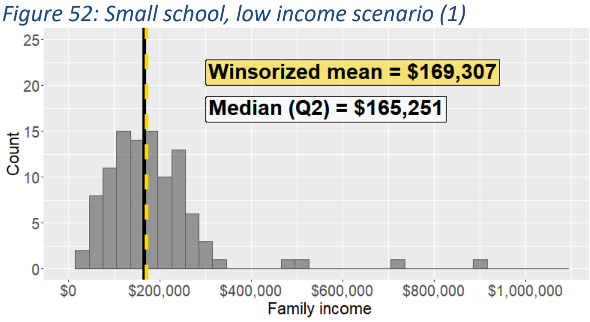

Figure 53: Small school, high income scenario (2)

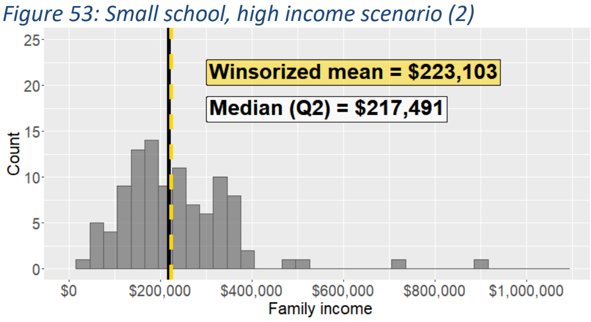

Figure 54: Small school, extreme outlier scenario (3)

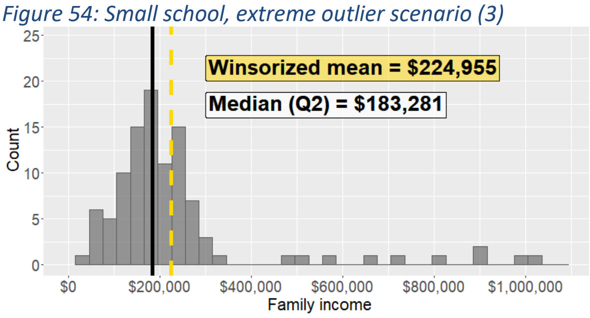

Figure 55: Small school, no outlier scenario (4)

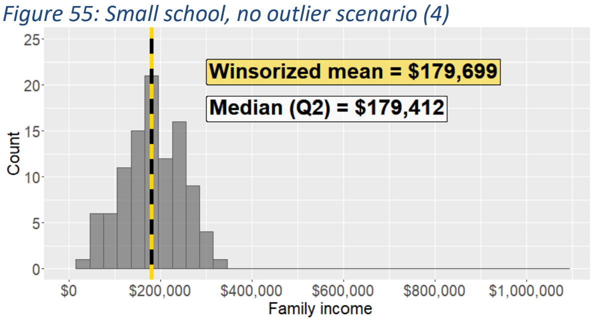

Figure 56: Large school, base case

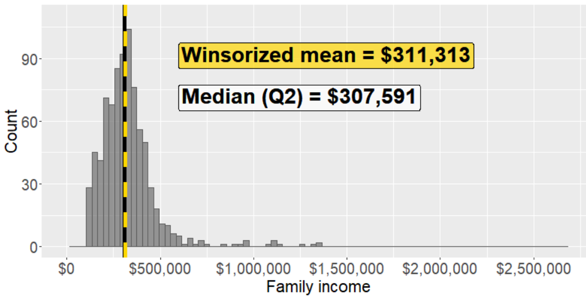

Figure 57: Large school, low income scenario (1)

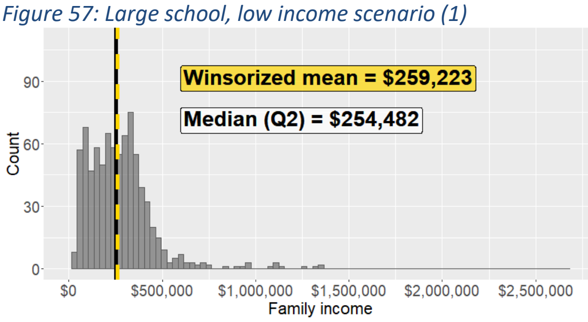

Figure 58: Large school, high income scenario (2)

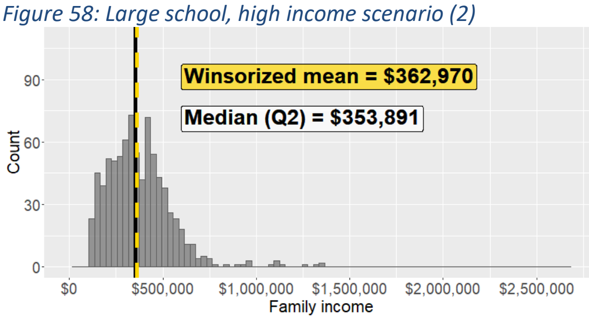

Figure 59: Large school, extreme outlier scenario (3)

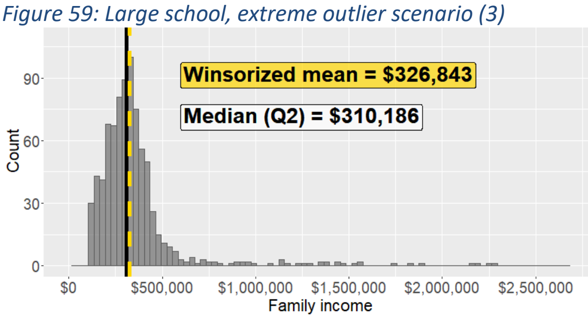

Figure 60: Large school, no outlier scenario (4)

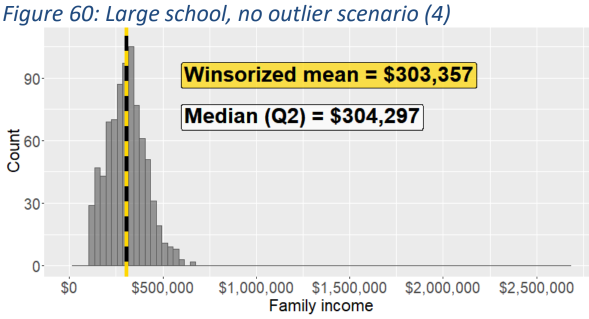

Proposed assessment of fitness-for-purpose of alternative statistical summary measures

47. To assess the fitness-for-purpose of alternative statistical summary measures for use in the DMI methodology, actual school community data will be used and the following two main factors will be investigated:

    (a) their impact on the DMI scores for schools, compared to using the median; and

    (b) their impact on the protection of the confidentiality of persons present in the income data.

48. Investigation of the impact on DMI scores of using the alternative statistical summary measures instead of the median will include:

    - analysing any differences in the impact for (i) schools of different sizes and types (that is, primary, secondary and combined), (ii) schools in different states and regions, and potentially (iii) schools with unusual family income distributions (for example, a long tail with unusually high income values);
    - assessing how sensitive each alternative statistical summary measure is to using imputed incomes values versus missing income values, compared to that of the median;

- testing how robust each alternative statistical summary measure is to unusually high income values, compared to the robustness of the median; and
- examining the effect each alternative statistical summary measure has on the volatility of the resulting DMI scores for schools over time, compared to using the median.

49. Investigation of the impact on person-level confidentiality of using the alternative statistical measures instead of the median will look at:

- whether there are different and/or additional considerations and requirements for protecting the confidentiality of persons present in the income data; and
- whether there is any differential effect on the number of DMI scores that are not permitted to be released from the ABS DataLab (as per Quality Gate 4, in reference [2]).

50. It is worth noting some important scoping points and assumptions that underlie the above proposed assessment of the fitness-for-purpose of the alternative statistical measures for use in the DMI methodology, which are as follows:

- the impact of the use of the alternative statistical measures on the DMI scores, rather than the CTC scores, for schools will be investigated;
- the process for calculating the DMI scores for schools (as specified in [1] and [2]) will be held constant, regardless of the statistical summary measure being used;
- in the event that DESE has not made a final decision on the income imputation method to be used for missing parent/guardian incomes values, a promising interim imputation model will be used to analyse the sensitivity of the alternative statistical summary measures to imputed incomes values versus missing income values.

Appendix - Usage of the statistical summary measures

1. In this section, examples of the usage of the median and the six alternative statistical summary measures in national and international income and economic statistics publications (if available) are provided.

Median

2. The median is used extensively as a statistical summary measure in national and international income and economic statistics publications. Some examples include:

   - the ABS's published statistics on (i) median weekly gross and equivalised disposable household income [3], (ii) median total personal income [4], (iii) median weekly and hourly earnings [5] and (iv) median house price at the state level [6];
   - the Melbourne Institute's published income distribution statistics from the Household, Income and Labour Dynamics in Australia (HILDA) Survey, which present income distribution information via the median, the ratio of the median to the 10th percentile and the ratio of the 90th percentile to the median [7];
   - the US Census Bureau's median household income statistics at the county level [8];
   - the UK Office of National Statistics' median real equivalised household disposable income statistics [9]; and
   - the weighted median for the Consumer Price Index (CPI) published by the ABS, Reserve Bank of Australia (RBA) and the Department of the Treasury [10].

First quartile

3. Some examples of the first quartile, being used as a statistical summary measure, usually alongside other similar percentiles, in national and international income and economic statistics publications include:

   - the ABS's published statistics from the Survey of Income and Housing [3], which use quintiles to present information about many income data items;
   - the ABS's 'National Accounts: Distribution of Household Income, Consumption and Wealth' publication [11], which also uses quintiles to present income data information;
   - the Organisation for Economic Co-operation and Development's (OECD's) 'Income Distributions in OECD countries' publication [12], which uses Q1 along with other percentiles to present income statistics; and
   - the US Department of Labor Statistics' 'Average weekly wages' publication [13], which uses a mix of quartiles and deciles to present income data.

Mid-hinge and tri-mean

4. The mid-hinge and the tri-mean date from developments in statistics during the 1970s [14]. However, they are not commonly used as statistical summary measures of income and

economics related data at the present time. A review conducted by the ABS did not yield any examples of the use of the mid-hinge or tri-mean in any national and international income and economics statistics publications.

Mean

5. Despite the weaknesses of the mean as a measure of central tendency of income distributions, its conceptual simplicity means that it is still often used [15]; albeit, it is often paired with other, more robust statistical summary measures like the median in order to provide a fuller picture of the income distribution.
6. For example, published income statistics from the ABS's Survey of Income and Housing, such as weekly equivalised disposable household income, weekly gross household income and household net worth, are presented as both means and medians [3]. Another example is the Melbourne Institute's published income distribution statistics from the Household, Income and Labour Dynamics in Australia (HILDA) Survey, which presents income distribution information via the mean, as well as the median, the ratio of the median to the 10th percentile and the ratio of the 90th percentile to the median [7].

Trimmed mean

7. The trimmed mean is commonly used in economic statistics when reporting price movements such as the Consumer Price Index (CPI), which is the weighted average price movements across the economy. The trimming allows for the CPI to be less affected by extreme (and likely unrepresentative) price movements [10]. The ABS, RBA [16, 17] and the Department of the Treasury [18] all use the trimmed mean to calculate the CPI and other price inflation measures. Similarly, the national statistical organisations of the US [19], UK [20] and the European Union [21] also use a trimmed mean when calculating their CPI measures.

Winsorised mean

8. Some examples of the winsorised mean being used as a statistical summary measure in national and international income and economic statistics publications include the ABS's published statistics for earnings, expenditure and employment from the Average Weekly Earnings Survey [22], the Quarterly Business Indicators Survey [23] and the Capital Expenditure Survey [24].

References

[1] What is the Direct Measure of Income (DESE): https://www.dese.gov.au/quality-schools-package/fact-sheets/what-direct-measure-income

[2] ABS (2020) A Data Quality Framework for the Australian Government's Direct Measure of Income for Capacity to Contribute, available at www.dese.gov.au.

[3] Household Income and Wealth data (ABS): https://www.abs.gov.au/statistics/economy/finance/household-income-and-wealth-australia/latest-release

[4] Personal Income in Australia, 2016/17 (ABS): https://www.abs.gov.au/statistics/labour/earnings-and-work-hours/personal-income-australia/2011-12-2016-17

[5] Characteristics of Employment data (ABS): https://www.abs.gov.au/statistics/labour/earnings-and-work-hours/characteristics-employment-australia/latest-release

[6] Residential Property Price Indexes: Eight Capital Cities methodology (ABS): https://www.abs.gov.au/methodologies/residential-property-price-indexes-eight-capital-cities-methodology/jun-2020

[7] The Household, Income and Labour Dynamics in Australia Survey: Selected Findings from Waves 1 to 17 (Melbourne Institute): https://melbourneinstitute.unimelb.edu.au/__data/assets/pdf_file/0011/3127664/HILDA-Statistical-Report-2019.pdf

[8] Median income data (US Census Bureau): https://www.census.gov/library/visualizations/interactive/2019-median-household-income.html

[9] Average Household Income (ONS): https://www.ons.gov.uk/peoplepopulationandcommunity/personalandhouseholdfinances/incomeandwealth/bulletins/householddisposableincomeandinequality/financialyearending2020provisional

[10] Underlying Inflation Measures: Explaining the Trimmed Mean and Weighted Median (ABS): https://www.abs.gov.au/websitedbs/D3310114.nsf/home/ABS+Chief+Economist+-+Underlying+Inflation+Measures

[11] Australian National Accounts: Distribution of Household Income, Consumption and Wealth (ABS): https://www.abs.gov.au/statistics/economy/national-accounts/australian-national-accounts-distribution-household-income-consumption-and-wealth/latest-release

[12] Income Distribution in OECD Countries (OECD): http://www.oecd.org/els/soc/Income-distribution-in-OECD-countries-1995-SPS18.pdf

[13] Average Weekly Wages (US Bureau of Labor Statistics): https://www.bls.gov/news.release/pdf/wkyeng.pdf

[14] Tukey, J. W. (1977) Exploratory Data Analysis, Addison-Wesley.

[15] Canberra Group Handbook on Household Income Statistics (UNECE): https://www.unece.org/fileadmin/DAM/stats/groups/cgh/Canbera_Handbook_2011_WEB.pdf

[16] Measures of Consumer Price Inflation (RBA): https://www.rba.gov.au/inflation/measures-cpi.html

[17] The Performance of Trimmed Mean Measures of Underlying Inflation (RBA): https://www.rba.gov.au/publications/rdp/2006/2006-10/background-on-trimmed-mean-measures-of-underlying-inflation.html

[18] Treasury's Underlying Rate of Inflation (Australian Parliament): https://www.aph.gov.au/About_Parliament/Parliamentary_Departments/Parliamentary_Library/pubs/MSB/feature/UNDERLY

[19] 16% Trimmed-Mean Consumer Price Index (Federal Reserve Bank of Cleveland): https://fred.stlouisfed.org/series/TRMMEANCPIM158SFRBCLE#

[20] Core Inflation: A Critical Guide (Bank of England): https://www.bankofengland.co.uk/-/media/boe/files/working-paper/2004/core-inflation-a-critical-guide.pdf

[21] An Evaluation of Some Core Measures of Inflation for the Euro Area (European Central Bank): https://www.ecb.europa.eu/pub/pdf/scpwps/ecbwp053.pdf

[22] Average Weekly Earnings, Australia methodology (ABS): https://www.abs.gov.au/methodologies/average-weekly-earnings-australia-methodology/may-2020

[23] Business Indicators, Australia methodology (ABS): https://www.abs.gov.au/methodologies/business-indicators-australia-methodology

[24] Private New Capital Expenditure and Expected Expenditure, Australia methodology (ABS): https://www.abs.gov.au/methodologies/private-new-capital-expenditure-and-expected-expenditure-australia-methodology

Glossary

Adjusted taxable income (ATI)

For the 2017-18 financial year, the Australian Taxation Office defined a person's adjusted taxable income (ATI) as the sum of the following amounts:

- taxable income;
- adjusted fringe benefits (total reportable fringe benefits amounts multiplied by 0.51);
- reportable employer superannuation contributions;
- deductible personal superannuation contributions;
- certain tax-free government pensions or benefits received by the person;
- target foreign income;
- net financial investment loss;
- net rental property loss;
- less any child support payments the person provided to another person.