HSC Data Analysis Tutorial Transcript

This document is a transcript of the narrative provided for each of the 17 modules in the online HSC Data Analysis Tutorial.

Contents

MODULE 1: COMPARATIVE LEARNING GAIN:

The basic idea of the HSC Data Analysis is called "Comparative Learning Gain". You can think of this as the answer to the question, "How have our students gone in this course, compared to similar students in other schools?" Importantly, this gives you a like-students rather than a much weaker like-schools analysis.

Any achievement result for a student can be considered in terms of the "Results Equation". This says that any result comes from two buckets of factors: what the students had at the starting point, and what they learned along the way. In Modern History, for example, the student begins with some knowledge and capabilities, and gains more learning along the way. It's important to see here that the "starting point" is not just knowledge of Modern History – it includes the student's literacy, their research skills and their ability to express themselves. At a deeper level, it includes their engagement, forward focus, resilience, wellbeing, self-concept and many other underpinnings of learning. And each of these can be – along with specific Modern History knowledge and skills – factors which are learned along the way. "Starting point" and "Learning" are buckets of factors, not just single factors. The same is true of any other achievement result such as NAPLAN Year 5 Numeracy – the result depends on what the student begins with and what they learn along the way. It follows therefore that there are only two ways to improve results. You either select students with a higher starting point, or you get the students you have to learn more.

You can consider the Results Equation the other way around. The Learning – or more precisely, the "Learning Gain" – is the difference between the Result and the Starting point. - This is the difference between where you finished and where you began. Note that the term "learning gain" is a better description than the expression that is sometimes used - "value-added". "Value-added" comes from a manufacturing analogy, where you take in raw materials and add value by transforming say wood into a table. We are not in the manufacturing business. In education, our work is learning.

It is easy enough to conceptualise Learning Gain in NAPLAN, because the four NAPLAN tests are all marked on the same scale. The learning gain measured for each student is the difference between the previous and the current score, so the learning gain for the class is just the average of these individual gain measures.

This leads into the idea of comparative learning gain. We are interested in learning gain because it is a measure closely tied to quality pedagogy. However, unlike NAPLAN, in many cases we do not have starting and finishing achievement results on the same scale. We solve this by asking where our students have finished compared to students who were in a similar place to them previously. This gives us a measure of comparative learning gain as the comparison between our achieved result and the "typical" result.

It's important to get the distinction between Learning Gain and Comparative Learning Gain clear. You have a starting point and a finishing point and the difference between them is the Learning Gain.

However, we don't have measures of the starting points on the HSC scale. If students with the same starting characteristics as our students typically finished at this point in the HSC, and your students finished here the difference between your achieved result and the typical is the measure of comparative learning gain. In other words, the typical learning gain is less that our achieved learning gain. It's important to see that if your achievement is less than typical, this does not represent "negative learning gain". The blue arrow still goes to the right, not just as far as the black arrow. It is true to say that the comparative learning gain is negative, but to say that the overall learning gain is negative is just wrong.

So that leaves us with the question, "How do you work out the "Typical" result? Each year, for each board-determined course where there are more than about 100 students across the NSW Catholic sector, take the results and for each student in that course, find their NAPLAN results, their gender, what school they attend, the Farish SES index for that school and the student's postcode-average SES value, and the student's ATSI status, then relate the overall result to these underlying data using a multilevel statistical model. While results vary

from course to course, overall this process accounts for about 56% of the variance in the HSC each year, a remarkably high figure. We then remove the "school" factor from the model, and so derive the typical result.

There are many ways in which the Analysis presents comparative learning gain. The first, the "Primary Analysis" – which is discussed in detail in Module 2 – shows the difference between typical and achieved such that if they are equal they lie on the 45 degree line but if, as in the example given earlier, the Achieved value is higher than Typical, the distance from the centre of the ellipse to the diagonal line represents a positive comparative learning gain.

A second way in which comparative learning gain is represented, discussed in Module 8, is the individual student contribution to the overall CLG. It is important to not over-interpret individual student results. As discussed in Module 8, this representation is about inferring patterns in the data.

A third way in which CLG is represented in the Analysis is in the 4th graph of a Trends Graph – discussed further in Module 7 – and in the Overall School Result – discussed further in Module 12. These graphs show the difference between Achieved and Typical, plotted over time. The CLG is represented as the vertical distance, corresponding to distance between centre of the ellipse in the Primary Analysis and the 45 degree line, with the diameter of the ellipse shown as the length of the whiskers.

The value of comparative learning gain is that it gives us a measure, derived from the student achievement data in the HSC, which is closely related to the pedagogy and can be represented straightforwardly. The question to be investigated is, "Why is it like this?'

MODULE 2: THE PRIMARY ANALYSIS

The Primary Analysis is called "primary" because it gives the simplest and most direct depiction of the Comparative Learning Gain in a course. It is the comparison of what the students have achieved in their HSC results with the typical HSC results achieved by similar students in other Catholic schools.

The primary analysis thus answers our main question, "How have our students gone in the HSC, compared with similar students in other Catholic schools?'

We will see that this leads to three possible follow-up questions. If we are using data well – and this applies to HSC CLG data as well as any other data – then we are using the data to frame good questions. There is always a temptation to rush to judgment with data. This isn't particularly useful, as it leads usually to either blame or self-satisfaction, without giving a clear way forward. It is only in digging into the data and framing good question around the data that we can make progress.

1. How have our students gone in the HSC, compared with similar students in other Catholic schools?
2. What – if anything – does this ask us about how teaching and learning are going in our school?
3. What are we going to do about it?
4. What do we need to learn?

The Primary Analysis graph is a plot of Achieved results versus Typical results.

When people start considering HSC results, they often come to the task with a concept of "what would be expected for this group". Often the assumption is that the group should be about State Average. This may or may not be a valid assumption. But if they rush into a judgment - "This wasn't a very good group", or "Haven't we done well", they are implying a statement of expectation. The consideration of the results is then one of what has been achieved compared with what might have been expected.

We used to use the idea of "Expected" to describe the bottom axis of the Primary Analysis Graph, but it has been changed to "Typical" on the argument that "I expect my students to be better than similar, read typical, students in other catholic schools…" You can expect more than what is Typical, but it's a bit illogical to expect more than Expected.

The primary analysis takes this achieved/typical concept and makes it explicit in a graph.

To begin the analysis, we use a multilevel model incorporating NAPLAN, gender, school, socioeconomic status measures and ATSI status to develop a common scale for all courses. There is a separate model developed for each course, because the relationship between HSC achievement and the underlying factors varies from course to course.

We then extract the Typical results just for one course. For the sake of the example, let's say we are considering Biology. The brown line represents the distribution of Typical scores for students doing Biology this year.

From the group of all students in the analysis doing Biology, extract just those doing Biology in this particular school. For the sake of making a pointed example, the green band here suggests that in this particular school the Biology students had much lower Typical scores that for students this year across the state. The middle of the green band is considerably to the left of the middle of the brown band.

Next, find the average of Typical scores for the students in this school. Remember that the calculation of the average is never an absolute process – such a calculation always includes some assumptions about the nature of the measurement.

Therefore, calculate the uncertainty of measurement of the average, to give a range within which it is 95% probable that the average will lie.

Then, consider the HSC results for all students in the analysis in Biology which we start with as one-unit marks.

Re-scale the Biology marks to the same mean and standard deviation as the Typical Biology results for the same students. The two brown lines now have the same mean and standard deviation; the HSC marks now lie on the same scale as the Typical results.

Extract from the re-scaled HSC Biology marks, the marks for the students in this particular school doing Biology. For the sake of making the point, in this example the students taking Biology in this school are shown as well above average in their HSC results - the middle of the green line is considerably above the middle of the brown line.

Find the average of the HSC Biology re-scaled scores.

You can then also find the uncertainty of measurement of this average, so giving a range within which it is 95% probable that the average lies.

If the achieved result was exactly what was typical, the average for achieved would be exactly equal to the average for typical, and the point representing both would be exactly on the line.

So in this example, we have a very high level of achievement with a relatively low typical result giving a Comparative Learning Gain that is very high However, we know that the average values have uncertainties around them, so we incorporate the confidence limits

We do this by replacing the dot with an ellipse, showing the confidence limits.

We therefore say that this represents Achievement significantly above Typical. Any time the ellipse is fully above the diagonal line, it represents achievement that is significantly above what is typical of similar students in other Catholic schools.

Whenever the ellipse crosses the line, this represents achievement in the typical range – although it is clearly better to have the centre of the ellipse above rather than below the diagonal line.

If the ellipse is fully below the diagonal line, this represents Achievement significantly below Typical. Note that in the particular example given here, the typical value lies well to the right indicating that this is a highly capable group of students – typically, they would have achieved very high results. In this case, they have not reached those high results.

There are two key variables in the Analysis which can be read from the Primary Analysis. The first of these is "Diff" – the difference between Achieved and Typical, which is seen in the Primary Analysis as the vertical distance between the centre of the ellipse and the diagonal line. The second is "Delta", the distance from the edge of the ellipse to the line. If delta is positive, as in this case, that represents achievement significantly above typical. Delta is negative when the ellipse is fully below the line indicating achievement significantly below typical. Delta is zero in any case where the ellipse cuts the line and the achievement is in the typical range.

Users of the Analysis can make the mistake of trying to over-interpret the size and shape of ellipses. There's not much educational information in these two factors; the user's interest should focus on how close the ellipse is to the diagonal line. However, it's not hard to understand the three factors that influence the size and shape of ellipses. Smaller cohorts have bigger ellipses, and vice versa, because the larger the group the more confident we can be in the mean. The spread of the results can affect the height and width of the ellipse differently causing it to be circular in some cases and stretched in one or other dimension in other cases. In some cases, a weaker

statistical model will cause the uncertainty in the typical values to increase the width of the ellipse. But to repeat, the user's focus should not be on these three factors, it should be on the distance of the ellipse from the diagonal line.

We can summarise the information in this module by looking at a typical set of results for an English KLA report. Note first that the English Extension ellipse is much larger than either of the others, probably indicating that it was only a small group of students, compared to the other two courses. Note also that the achievement in English Advanced is significantly above typical, with a delta value that is positive.

Thirdly, note that although the Diff value for English Extension 1 is larger than that for English Advanced – the centre of the Extension ellipse is further from the diagonal line than the centre of the Advanced ellipse is – Extension is still described as "in the typical range" because of its much larger ellipse.

Finally, note that each of these courses (and every other course analysed) lie on the same background Typical scale, so the fact that the Advanced cohort is to the right of the Standard cohort tells us that they had a higher starting point – which is not particularly surprising. The Advanced course generally has a school's more-able English students in it.

We are interested in the Primary Analysis because it gives us a way of depicting Comparative Learning Gain. And we are interested in Comparative Learning Gain because it is arguably the strongest indicator, derived from student achievement data, about something we are critically concerned with: the quality of the pedagogy. Subsequent modules will give us further ways of exploring this.

MODULE 3: THE SECONDARY ANALYSIS

This module gives an overview of the Secondary Analysis in the HSC Analysis of Results. While the primary analysis is a comparison of the results in each course with what is typical of comparable students in other Catholic schools, the secondary analysis is a comparison of the HSC results in a course with the average of the results for the same student in each of his/her other courses. The second part of the secondary analysis is a comparison with the state average. The secondary analysis is done using no NAPLAN information. The secondary analysis is done on the basis of marks re-scaled to TES (ATAR) parameters. In this way, all marks lie on a common scale.

First, take all of the students in this course, and get their NESA one-unit mark.

Re-scale the marks for each course in each school to get the TES (ATAR) value of those marks. The mapping points to do this have been published by UAC each year on their website and in the Technical Committee for Scaling's report.

For each school, find the variation of the school from the all-of-state average. Note that this won't be exactly the same as the "School/ State variation" published by the board of studies, as this analysis has put all subjects onto a common scale, the TES scale.

For this subject, keep the difference as a TES-based school/state difference.

Now consider how each course compares with all others in the school. The problem we have with NESA marks is that they lie on different scales. I.e., a mark of 70 in Physics may be more difficult to obtain than a mark of 70 in Senior Science. The TES scaling process is intended to place all courses on a common scale, so it is these marks that we use.

Take one subject at a time. Say Biology, and re-scale it to the TES scale. Consider each student who takes biology in turn. This first student we have found has a mark in Biology a bit above the average. Now for that student, consider his mark in…… English Standard. He hasn't done as well in English, the mark is a bit less. The green course is Maths, and he has done quite poorly there. And in Business Studies – the yellow line – he has done better than Maths and English, but not as well as Biology. Studies of Religion, for this student, is a bit less again The whole school did pretty terribly compared to the State in the purple course, say PDHPE, but the student whose marks we are looking at was close to the top of that class, and he has done just a bit better than in Biology.

This demonstrates that for this student, Biology is his second-best subject. Now for this student, measure the difference between his mark in Biology and his mark in each other subject.

Considering just these differences, find their average for this one Biology student. Repeat this process to find the average "all-other-courses" difference for every Biology student. Because we are looking only at the Biology students' marks in their other courses, and doing it on a common (TES) scale, it is a fair way to compare courses. This sort of comparison is called a "same-student" analysis.

Then we find the average difference across all students. This gives a measure of how Biology has compared with all of the other courses in the school taken by students of biology - hence it is a comparison on a common (TES) scale, across common students.

This value, with the value we worked out before for the comparison with state, Gives the point on the graph where we put A little dot representing this course compared to school and state averages. Note that this dot doesn't have "uncertainties of measurement" around it. These are shown in the Trends Analysis, described later, where these data are represented over time.

MODULE 4: UNDERSTANDING "+H" AND "+L"

There are several places in the Analysis where you might see either +H or +L written beside a course's analysis for the year. This is a clue to understanding something about the pedagogy in that course. We call +H and +L second-order effects in the analysis. The first-order effect is the comparative learning gain, while the second order effect enables us to dig a little deeper into the pedagogy that resulted in that comparative learning gain.

Think of the comparison between Achieved and Typical that is given by the Primary Analysis. In this graph, each triangle on this graph represents an individual student's result. Only a few students are shown for simplicity. You actually need several hundred students and at least 20 schools to get a multilevel model to converge.

Take a look at just one school's students. These are marked here in red and show each student's comparison of Achieved and Typical.

The average of the individual students' results is represented in the Primary Analysis as an ellipse – here shown in red.

This gives the usual representation of the Primary Analysis, in this case showing Achievement above Typical.

However, there is other information here, besides the average of the results. It is possible to draw a line of best fit through the individual

Think of another school, shown here in green. The line of best fit through these points is much steeper than typical: the Comparative Learning Gain for the higher end of the student distribution (those to the right of the graph) has been greater than that for those at the lower end. We call this a +H effect.

A plus-H effect is neither particularly good news or bad news. It just tells you that the line of best fit is steep. In the case shown here, achievement is less than typical, and the +H effect tells us that this is likely to be the case because the lower end of the lass distribution has become disengaged.

A plus-L tells us the opposite. The line of best fit is a shallower slope than is typical, showing that the comparative learning gain for those at the lower end of the distribution is greater than that for those at the higher end of the distribution. This is a common pattern we see in courses where achievement is above typical.

A plus-L effect with a significant achievement below typical as is shown here is not common.

If the line of best fit is more-or-less typical, then the course that year has neither a +H nor a +L.

We don't use the terms minus-L nor minus-H, as they would mean just the same as +H or +L. The most important effect to look for is the comparative learning gain shown in the primary analysis. A later module will look at the individual student scatterplot of results to see how we can examine these for patterns.

The point of looking at second-order effects, the +H and +L, is to frame questions around the pedagogy. A string of +L results –across years, or across different courses – asks whether we have a particular focus on working with struggling students, and therefore if we need to look more closely at how we might stretch able-students. Conversely, a string of +H results with negative CLG asks what we need to do to help struggling or disengaged students.

It's valid to adopt a see-saw principle: if I succeeded in doing well with struggling students last year, what can I do with able students this year? Or vice versa?

MODULE 5: TRENDS GRAPHS (A): GRAPHS 1 & 2

This is the first of three modules which will deal with understanding the Course Trends Report. This module deals with graphs 1 and 2 from a Trends Report. The next trends module shows the links between the Secondary Analysis and trends graphs 1 and 3. The third trends module shows the links between the Primary Analysis and the trends graphs 2 and 4.

Individual Course Trends graphs are found in the Analysis portal under the Course Reports tab.

Note that in these three modules, we are considering just the individual Course Trends graphs, not the Overall School Result graph. While it looks much the same, the Overall School Result graph has some important differences from the individual Course Trends graphs.

A Trends Report has 4 separate graphs. Graphs 1 and 2, which we will be concentrating on in this module, are comparisons of results with state average. The horizontal line of the graph represents state average. Graph 3 is the comparison with the school all-other-courses value found in the Secondary Analysis. Graph 4 is the Comparative Learning Gain from the Primary Analysis.

Note that in each of these 4 graphs, each annual point on the line has whiskers around it. These are the so-called confidence limits within which the value being reported lies. They are a similar concept to the ellipse used in the Primary Analysis. Their use will be discussed in more detail in Module 6.

It may help you while doing these three Trends modules to have a copy of a Trends Report from your own school in front of you.

Trends Reports cover the current and the previous five years, as shown at the top of the report. If there are too few students doing the course in any given year, it is not possible to produce a Primary Analysis. It is necessary to have both a Primary and a Secondary Analysis to get a Trends Report.

Under each year there are two rows of student numbers.

The first is TotNos – Total Numbers. It shows the total number of students doing this course in this school this year. This is always the same as the number of students whose results contribute to Graph 1.

The second is NoIncl – Number Included. This shows the total number of students doing this course in this school this year, for whom we have sufficient NAPLAN, SES and gender data to be able to include them in the Primary Analysis. This is the same as the number of students whose results contribute to Graph 4. The Number Included can never be greater than the Total Number.

Graphs 1 and 2 both show a comparison with state average. What's the point of doing this?

The point of doing this is to have a valid result to compare Graph 4 with. So, in Graph 2, we extract those students included in the Graph 4 Comparative Learning Gain graph and get their comparison with state average. The same students are included in Graph 2 as in Graph 4, so the number of students in Graph 2 is the Number Included.

Generally, you will find that Graphs 1 and 2 look much the same. What that tells us is that any students who didn't have sufficient background data – NAPLAN etc – to be included in Graph 4 are nevertheless pretty typical of the school cohort. The important conclusion from this is that their absence from Graph 4 doesn't offer any explanation for features of Graph 4.

Just occasionally, you will note a sharp difference between Graph 1 and Graph 2. That is true in this example for 2018. There were 7 students who did the course – and hence are included in the Total Number – but are not included in the Graph 4 CLG graph. And we note that in 2018 Graph 1 has spiked much higher than Graph 2.

This tells us that these 7 students have contributed more than the rest of the course cohort to the above average results, and hence might have produced a significantly different Graph 4, if background data had been available for them. It is not certain that they would have, but it is possible.

The final thing to note in the headers on a Trends Report is that you are shown the second-order effect, +H or +L or blank for each year, as described in module 4. We will see in module 7 how this information can be used to start to infer patterns in the data.

The next module will look at the comparison between Graph 1 and Graph 3, and see how this gives a representation of the Secondary Analysis over time.

MODULE 6: TRENDS GRAPHS (B): THE SECONDARY ANALYSIS AND WHISKERS

This is the second of the three modules which will deal with understanding Course Trend Graphs. This module deals with the links between the Secondary Analysis and the Trends graph, and explains how the confidence limits – the whiskers – on the trend graphs are used.

Individual Course Trends graphs are found in the Analysis portal under the Course Reports tab.

Graphs 1 and 3, which we will be concentrating on in this module, allow us to see how the Secondary Analysis has changed over time. Graph 1 is the comparison with state average. Graph 3 is the comparison with the school all-other-courses value found in the Secondary Analysis.

In module 3, we looked at the Secondary Analysis, where the vertical distance represents the comparison of the course with state average, and horizontal distance represents the comparison of the course with the school all-other-courses average.

The same things can be seen in the trends report. Graph 1 of a trends report shows the vertical distance, for that year, on the Secondary Analysis.

Similarly, trends graph 3 shows the horizontal distance in the secondary analysis.

Satisfy yourself that you can see that trends graphs 1 and 3 show you what is happening to the Secondary Analysis over time. Select a Trends Report of interest, then go to the KLA reports and get the corresponding Secondary Analysis. Make sure that you can see that the dimensions shown on the previous slide are true for your course. Then, go back to the previous year, get the Secondary Analysis for that year, and do the same. Doing this, you will see that the Trends Report gives you, through graphs 1 and 3, a view of the Secondary Analysis over time.

Trends graphs 1 and 3 give you more than just the pattern of the Secondary Analysis. They also provide confidence limits or whiskers which help you to decide whether differences are notable or not.

On any of the 4 graphs in a Trends Report, the whiskers represent what are known as confidence limits. Our measurement of the average, the mean, is shown by the dot. The whiskers show how confident we can be that the mean lies exactly where the dot is. Importantly, the whiskers do not represent the spread of results, which will always be much bigger than the whiskers. The remainder of this module will deal with what the whiskers show, and how they can be used.

Think about the individual student results which make up, for example, a primary or secondary analysis. The mean of these results can be represented by the red dot. But how sure can we be that this dot exactly represents this group? There are many ways to think about this question, but one way is to think about what happens when you take out just one result. If you remove just one student, the mean shifts a little. If you remove another, it shifts again.

We use whiskers to show us where we can be confident the mean of the group of values is. The size of the whiskers can be affected by three factors.

As the group gets smaller, we can be less confident of the location of the mean, so the confidence limits get bigger.

As the data points become more bunched, we can be more confident of the mean, so the whiskers get smaller. But it is important to note that the size of the whiskers is always smaller than the spread of the results.

So, what do you use the whiskers for in reading a trend graph? The first use is to tell if the mean is significantly different from the average represented by the horizontal line – say the state average, or the school all other average or the typical value. In the example given, the mean is significantly higher in the first five years, but in the range in the sixth.

The second use of the whiskers is to tell if a change between two years is likely to be notable or not. This is where we use the half-whisker rule: if the change is more than about half a whisker, it is more likely than not to be significant. Under this rule, the first change indicated is not likely to be significant, but the second is.

Sometimes, when you have very large groups you have very small whiskers. In this example, even small changes are significant because they are more than half a whisker.

There are two important things to take from this module: Firstly, we can use graphs 1 and 3 from a Trends Report to show the Secondary Analysis over time, and secondly we can use the whiskers on the graphs to interpret significance of what we are looking at.

MODULE 7: TRENDS GRAPHS (C): GRAPHS 2 & 4 WITH THE RESULTS EQUATION

In this third module concerning trend graph interpretation, we concentrate on graphs 2 and 4 and see how they give us key insights to the factors underlying the HSC results.

Remember that in a Trends Report, graph 4 shows the comparative learning gain, over time, for the students for whom we had a full set of NAPLAN, SES, gender, school and ATSI data. This is the number included at the top of the report. Graph 2 shows the average of these included students' results in this course, compared with state average. So, in all cases graphs 2 and 4 represent the same students – graph 2 showing the results and graph 4 showing the comparative learning gain.

Graph 4 is the heart of the analysis, showing comparative learning gain over time. If you compare graph 4 with the primary analysis showing the same course, you see that the vertical distance on graph 4 is just the same as the distance on the primary analysis from the centre of the ellipse to the diagonal line. As the ellipse gets smaller on the primary analysis, the whiskers on graph 4 get smaller. As the ellipse gets bigger, the whiskers get bigger.

It is useful at this point to pause this module and check one of your own course trends graphs. Open the KLA Primary Analysis including that course for each of the last few years and see that the way in which graph 4 changes reflects the changes in the primary analyses.

The Results Equation, covered in Module 1, reminds us that there are only two buckets of factors influencing the final HSC result – where the students started from, and how much they have learned along the way. We can see the result simply in the comparison of the group's average score with state average. The Comparative Learning Gain gives us our best available measure of learning gain, and so tells us about pedagogy. But we have no direct way of measuring starting point. This module shows how to infer this.

In the academic literature there are two broad ways in which the term "pedagogy" is used. It either means the art and science of teaching or it means the teaching/learning interaction. Because the same teaching can produce different results with two different groups having different learning needs, we use the second of these definitions in this module.

On a Trends Report, graphs 2 and 4 give you the Results Equation analysis, over a 6-year period, for the students shown as number included at the top of the report. Note that you also see the second-order effect, as discussed in Module 4. Graph 2 shows you the Result and Graph 4 shows you the comparative learning gain, so gives you an insight to the pedagogy.

You can look at the results equation diagrammatically. Graph 2 on a trends graph shows the result. Graph 4 shows the CLG, our best indicator of the pedagogy. An average cohort getting typical pedagogy will achieve an about state average result. Both Graphs 2 and 4 will be on the horizontal axis. If – as is the case in this diagram – the students have received above typical pedagogy, you would see an average cohort achieve an above state average result.

We now get into the business of explaining our results. The results are what you are looking at in graph 2, and are shown as a difference from state average. They can be explained by the quality of the cohort, the quality of the pedagogy, or a mixture. The rule is simple: if you see the same thing happening in Graph 4 as Graph 2, the explanation of the results is likely to be a pedagogy effect – it is pedagogy, rather than a cohort effect, that is producing the result. Otherwise, it is a cohort effect, or sometimes a mixture of both cohort and pedagogy – for example an above average cohort with above typical pedagogy.

It is important to be specific on what you are explaining: it is always in graph 2, but there are many features in graph 2. In this Graph 2, you could be explaining this year's result. Alternatively, you might be explaining the improvement in results over the last 4 years.

Alternatively, you might be explaining why the results have been consistently above state average for six years. Be specific about what you are going to explain. Say you have chosen to explain this year's result. The rule is that if you see the same thing in graph 4 as in graph 2, the explanation is likely to be pedagogy. That's the case here: graph 4 is way above the horizontal axis, just as graph 2 is. So, we can safely conclude that this looks like an about average cohort which has received above typical pedagogy - see Graph 4 - and so has achieved above average results - see Graph 2.

Consider this set of results. You would probably choose from Graph 2 to explain the fact that for 6 years the results have been consistently below state average. But you see a very different thing in graph 4: the CLG line is consistently above, showing that these students are getting better results that similar students in other schools. So, you must conclude that this is a cohort effect. The reason the results are below state average is that these students started secondary with very little prior schooling.

How would you interpret this set of results? It is a consistently above average set of results. However, you see a very different thing in graph 4. The students are achieving consistently lower than similar students in other schools. The explanation of the above state average results in graph 2 has to be that they are a very able cohort – even though graph 4 shows they are not reaching their potential.

Try providing an explanation of this graph before going on to the next page.

Whether you choose to explain the six year trend of below state average results, or just this year's result, you get the same explanation: you see the same thing in graph 4, so the explanation is less than typical pedagogy.

Try explaining this Trend Report before going on to the next slide.

The most obvious feature of this graph 2 is the way the results go up and down every second year. You see the same thing in graph 4, indicating a pedagogy effect. Do you have alternating teachers? Does one teacher need to learn something from the other?

This is another example of a pedagogy effect. Graph 2 is slowly rising, and you see a similar thing in graph 4. Notice that the second order effect at the top of the report indicates a +L for the last 3 years. Does this point to better engagement?

Considering results in this way leads to the conclusion that your objective is to get graph 4 to go up. The results are driven by both cohort and pedagogy. The changeable factor in the teacher's hands is pedagogy, which is indicated by graph 4.

It would be useful at this stage to open some course Trends Reports of interest to you. Apply the principles covered in this module to them, until you are confident you can see where the questions in the data lie.

MODULE 8: SCATTERPLOTS (A): ACHIEVED VS TYPICAL

In this module we look at the use of the first of the two scatterplots that are available for each course. These are useful, but easily mis-used.

When you open the Analysis portal, there is a button for Scatterplot of Students Achieved versus Typical results. In doing this module, it will be helpful to have a scatterplot of a course that is of interest to you in front of you.

The scatterplot shows each student's Achieved HSC mark compared to his/her Typical mark, calculated as described in Module 1. In this course, the Primary Analysis ellipse, as described in Module 2, would be something like this. Note that in this scatterplot, the values being plotted are actual marks, which are also shown to you in the Student Results report as described in the next two slides.

The Student Results report can be gained for any course as the fourth report listed under each course.

The Student Results report gives you both the Achieved HSC results, rounded to a whole number, and the Typical values – as well as the TES equivalent 1-unit mark, and the All-others TES value used in the Secondary Analysis described in Module 3.

What's the point of this scatterplot? Firstly, it helps you to put faces on the data shown in the Primary Analysis. Secondly, and even more importantly, it enables you to look for patterns in the data. Is there something about what Adam and Rebecca, and Alfonso and Victoria did that really helped them to achieve just good learning gain?

This exercise is not one of finding excuses – these students have left school! – but looking for patterns that may give a direction for improving the teaching/learning process.

The third thing that the scatterplots are useful for is in seeing differences between different classes of students within a course or different categories of students. The class at the top of the report is by default set as A. In Module 9, we will look at how to split students within a course into different classes, and thereby see differences.

The fourth use of the scatterplot is to translate the Comparative Learning Gain into HSC marks. The Class Mean Difference at the top of the report is the value of the average difference between Achieved and Typical for students in this class. It is directly a measure of Graph 4 on a Trends Report. Importantly, once you have divided a course into different classes, it will tell you the different Class Mean Difference or Comparative Learning Gain for each class.

MODULE 9: LOADING CLASS DATA

In module 8, we looked at the scatterplot graphs of students' Achieved and Typical results. In that module, we saw that the scatterplot, by default, places all students for a course in the same graph. This is okay if you only have a small group in the course, but presents difficulty for courses with large numbers of students. It is often of interest to know how the students in different classes of a course have gone. This module gives you a way of exploring this.

On the top right-hand corner of the Analysis Portal is a link, "Edit classes and factors".

3. If you click on the "Edit classes and factors" link, you are taken to a page with a drop-down menu enabling you to select any course for the year, and underneath that an alphabetic list of all students who sat the HSC in that course. There are two fields beside each student's name. On the left-hand side is "Factor", which we will discuss in Module 10. The "Class" field allows you to enter any code you like for each student. The easiest thing to do here is to just use "A", "B", "C" etc for classes, but you can use anything you like – teacher names, class timetable codes or whatever makes sense. The important thing to note is your codes need to be identical for all members of the same class. Capital-A will be treated as a different class to lower-case-a. And A following by a space will be treated differently to an A without a space after it. So be careful to not leave spaces! Once you have entered the students' classes, make sure you press one of the "Save All Changes" buttons before navigating away from this page. Otherwise, all your work disappears! Note that when you press "save all changes" your changes are saved to the server where the Analysis resides, not to your own computer. Hence, your changes will be visible to any user of your school's analysis.

So, what's the point of entering the code for different students' classes? You will remember that in Module 8 we saw that at the top of each scatterplot of Achieved versus Typical, you see both the Class for the students, and the "Mean Difference" for the group. If the whole group is in the default class "A", you see them all on the one scatterplot. However, if you enter different classes for the students, you get as many scatterplots as the different codes you have entered.

Say you have entered "A" and "B" as the class code for each student. You get a scatterplot for those students you have left in class A, and a second scatterplot for class B. Notice that the Class Mean Difference has changed for both classes.

"Class" doesn't have to only refer to what teacher taught the student. You might want to investigate categorising students according to which elective they studied. Or whether they took a particular course in Year 10. Or whether they are part of the football team! The Class Mean Difference figure gives you a way of expressing Comparative Learning Gain in HSC marks, to make sense of whether there is a difference between any of these categories.

MODULE 10: LOADING FACTOR DATA

In module 9, we saw how to enter students' classes within a course, or any other category that you are interested in. But sometimes, the things we are interested in investigating are numerical rather than categorical. For example, you might be interested in the effect of submitting practice essays. Or attending after-school seminars. Or absences from class. This module gives you a way of exploring this.

You go to the same place on the top right-hand corner of the Analysis Portal and open the link "Edit classes and factors".

Use the drop-down menu to select any course for the year. We are now interested in the left-hand side column, labelled "Factor". Here you can enter any numerical data of interest to you concerning the student. You might have records of the number of student absences. You might have records of student homework submission. What the Factor column enables you to do is to enter these data. Again, once you have entered the students' classes, make sure you press one of the "Save All Changes" buttons before navigating away from this page.

In Module 11 we will look at how you can use the second scatterplot graph to see if the factor data you have entered is making a difference to the observed Comparative Learning Gain in the course.

MODULE 11: SCATTERPLOTS (B): DIFFERENCE VS FACTOR

In this module we look at the use of the second of the two scatterplots that are available for each course. The second scatterplot is useful only if you have loaded some factor data. It plots the Difference between the Achieved and the Typical result against the Factor that you have loaded.

When you open the Analysis portal, the second scatterplot is called "Scatterplot of Students' Difference vs Factor".

You may have seen a Difference vs Factor graph, looking something like this. It doesn't tell you much at this stage, because every student has the same default factor value of zero. The problem we are trying to solve here is this: we have three different things we are interested in, the Achieved HSC score, the Typical score and the Factor. It's hard to plot a graph with three axes, so what we do is plot the Difference between Achieved and Typical against the Factor. Remember that the Difference gives us our best measure of comparative learning gain. So, on the graph shown on this page, we have the differences plotted, but no factors.

Maybe you have recorded how often students attended your after-school or holiday seminars, and want to know if they worked. If you have entered the number of seminar attendances for each student as a factor (as explained in Module 10), you might get a Diff versus Factor graph like this. A graph like this would be powerful evidence that attending your seminars is related strongly to better learning gain, and hence better results for students. Alfonso and Rebecca have attended all seminars and got the best CLG figure; Marco and Abe attended none and got the lowest.

Of course, it is quite possible that whatever you have entered as a Factor is not related to the student comparative learning gain. A graph like this would indicate not much relationship.

It is quite possible to have a negative relationship between Difference and Factor. Say the factor you put in was absences from school. A graph like this would indicate pretty strongly that there is a connection between absences from school and poorer learning gain.

The Diff versus Factor graph is useful for giving you an idea of whether the factor you are investigating has a relationship with the HSC results by looking at the relationship with student-level comparative learning gain. If you want to do any more detailed analysis, you will need to download the dataset, as explained in Module 16.

MODULE 12: THE "OVERALL SCHOOL RESULT" REPORT

If your concern is how the school as a whole has performed in the HSC, your first place to look each year is at the Overall School Result (OSR) report. While it uses the same software as individual course Trends reports, it has a number of significant differences from these which are outlined in this module.

The OSR report is found in the Analysis portal close to the top of the page.

Graphs 1, 2 and 4 for an OSR are weighted averages of the individual Course Trends reports. To say that they are "weighted" averages means that courses with more students have a bigger impact on the average than small courses. Graph 3 however is different. You will remember from Module 6 that in an individual Course Trends report, Graph 3 shows this course compared with the average of all other courses done by these students. So if we did a weighted average across all courses of these differences, they would all cancel each other out, and you would get a flat line along the graph telling you nothing. So the OSR Graph 3 is used to show you something completely different. It takes, for all students who have 10 units of Board Determined Courses (BDCs), the value of their Achieved Tertiary Entrance Score (TES) and uses multilevel modelling to compare it with their Typical TES value, as outlined in Module 1 for individual courses.

So in an OSR Graph 3 gives you CLG across just the 10 best BDC units, but Graph 4 is the average CLG across all units. This means that the comparison between the two graphs gives you interesting information. Mostly, as in this case, you will see that both graphs are similar. Where they are noticeably different, the differences are caused by differences in units included. Graph 4 will show units which are not in the 10 best units for students who take more than 10 units of BDCs. Graph 4 will also include the BDC units undertaken by any student who has not yet accumulated 10 BDC units – these could be accelerated students, or students on a non-ATAR pathway. Graph 3 will include units studied in previous years by any accelerated student who now has 10 units, and because these units were not studied this year they will not show in Graph 4. The interpretation of any differences between Graphs 3 and 4 needs to take account of which of these factors apply in your school.

The same comparison between Graphs 2 and 4 applies for an OSR as we saw in Module 7 for an individual course Trends report. As for individual course Trends reports, the Graphs 2 and 4 in the OSR report on the same students. Before you click to go on to the next slide, check out how well you remember module 7. What is the likeliest explanation of (a) this year's result; and (b) the change from last year to this year in this OSR?

In looking at this year's result in Graph 2, we can see it is well above the line indicating average results across all courses above state average. Graph 4 is just above the line, indicating somewhat better-than-typical learning gain. So we would describe this as a mixture of pedagogical and cohort effects – good pedagogy (Graph 4 is above the line) for a capable group (Graph 2 is even more above the line). Looking secondly at the change from last year to this year, we can see that Graph 2 is pretty steady over the two years. However, Graph 4 has dropped somewhat. This suggests that this year's cohort was somewhat more capable, getting the same results (shown in Graph 2) off slightly lesser CLG.

The final difference between the OSR and individual course Trends Reports is in the numbers of students reported at the top of the page. The "Total Numbers" reported is the number of students who had at least one Board Determined Course result included in Graph 1. The "Number Included" for an OSR is the number of students included in Graph 3. Hence, number included is the total number of students who had accumulated 10 units of BDCs.

As for a course Trends report, the objective for the whole school is simply stated and hard to achieve: Make Graph 4 of the OSR go up.

MODULE 13: THE SCHOOL SUMMARY

When you have looked at the Overall School Result report – as described in Module 12 – the question that arises is, "How do different courses contribute to this Overall School Result". The School Summary report gives you a starting point to investigate this question.

The School summary can be selected from the portal just below the Overall School Result, near the top of the page.

There are six important things to know about the School Summary in order to be able to use it effectively. While this report contains lots of detail, it is very easy to over-interpret it in ways that are not helpful. Keep these six things in mind.

In this report, the course at the top has the highest comparative learning gain within the school, and the course at the bottom of the list has the least. Hence, when the course at the top has a Trends Report, its Graph #4 will be the highest above the line of any course within the school.

To produce a list ordered on comparative learning gain, we calculate the difference between the average Achieved score and the average Typical score within the class and call this "Diff". This value is the fourth numerical column in the report. Note that in this report we are using standard deviations rather than HSC marks, so a Diff value of 0.810 as is seen here for Design and Technology is quite large. Diff values which are positive indicate a course Trends graph with Graph #4 above the line, and those which are negative indicate a Graph #4 below the line.

For most courses, the achievement level is noted as a comparison with Typical. As noted in Module 5, if there are too few students in the school undertaking the course it is not possible to produce a Trends report for that course because we cannot confidently compare the average of Achieved scores with the average of Typical. However, in the School Summary these courses appear where they would have if it were possible to do the comparison.

It is the courses with more students that have the larger impact on Graph #4 of the Overall School Result. In deciding which courses to focus on, the courses with more students are a better place to start.

The point of the School Summary is to give you a basis for investigating individual course Trends reports. Has Graph #4 for this course stayed consistent, risen or fallen? Why? What is the pedagogical story?

It is possible to use any ranked list as a league table. It is also usually unproductive. The useful way to use these data is to know where to focus attention in using the techniques outlined in Module 7 for understanding whether cohort or pedagogy is making the larger impact on results.

Getting Graph #4 of the Overall School Result to go up depends on getting individual course Graph #4 to go up. The School Summary report is a way into considering where the productive points of investigation lie.

MODULE 14: PARTICIPATION RATES

It is often of interest to consider how many of your students undertake different courses. This module outlines the two different ways in which you can use the Analysis to gain a perspective on this.

The first way we can look at participation is through the Participation report on the Analysis portal.

This report gives a "Cohort Participation Rate" for both the State as a whole, and your school. What percentage of the state – or of your school – took this course? The ordering of the report is important. The first course on this report – in this case, Music 1 - is the one with the least participation in your school (in this case 1.4%) compared to the state as a whole- 7.9%. The school participation rate is less than a quarter of the state participation rate. Towards the middle of the list, the state and school participation rates are about equal.. Towards the bottom of the list, the school participation rate is much greater than the state participation rate. It is not surprising that a Catholic school would have a higher participation rate in Studies of Religion than the state as a whole. But are there other surprises in this report, for your school, for you?

What is the point of this report? It allows you to see where your school is different from the state as a whole, and gives the basis for asking if this is an issue. As one example from this report, consider English participation. In English Standard, school participation is much larger than state participation. But for English Advanced, the opposite is true. Should we be challenging more of our students to undertake the higher English course? Module 17 will consider some of the research that addresses the questions on how we advise students in course preferences, using data such as are shown in this report.

There is a second way we can consider participation. Comparing the Course Participation Rate for the state with that of the Catholic sector as a whole gives rise to further interesting questions. In the Analysis Portal there is a "Resources" tab. Clicking on this gives the user access to the report prepared each year on the Analysis for the preceding year. Appendix 1 in this report gives Course participation rates for the Catholic sector.

A Course Participation Rate is the percentage of the candidature of a course that comes from a particular group – in this case, Catholic schools across NSW. We know that students in Catholic schools generally comprise about 25% of the HSC candidature each year. (The exact value is given at the top of Appendix 1 each year.) Therefore you might expect that Catholic school students would comprise about 25% of each course. That is the case for many of the courses in the right-hand column of this table. In very small courses, it is not surprising that the figures vary widely. But can you explain Mathematics Extension 2, 12% in this year? Again, Module 17 will address issues like this.

Participation rates can show up some surprising trends in course enrolment. The point of considering them is to consider whether our practice in student course preference advice is the best it could be.

MODULE 15: THE TERTIARY ENTRANCE SCORE (TES) REPORTS

It is useful to have a summary indicator of each student's performance in the HSC. This module outlines the two TES reports available in the Analysis which give one way of summarising a student's performance.

The TES Summary report gives an alphabetic listing of students, showing the value calculated for TES for each student. This value can then be used to locate the student in the second report, TES Details. This second TES report is listed in descending order of TES.

The TES is the sum of the student's ten best units, which forms the basis for the ATAR. The value calculated for the Analysis is close to but not exactly the same as that used for the official ATAR. It is a close enough approximation to serve well as a summary of student performance in the HSC.

The TES Summary report lists students alphabetically. To the right-hand side of the report is the calculated value of the TES. Students who have fewer than 10 BDC units do not get a value for TES. This report also gives the NAPLAN scores for the student, from Years 9 and 7.

The second TES report lists students in order of their calculated TES values with each of their BDCs undertaken this year. To the right-hand side of the report is the calculated value of the TES-scaled one-unit mark for each course. Note that in this report the listing of each student's courses is in descending order of TES-scaled mark, and so indicates which course counts most, and which counts least, towards the student's ATAR.

The two TES reports give a useful summary of student performance. However it is only a summary. The work of using the analysis lies in the investigation of the detail described in earlier modules.

MODULE 16: DOWNLOADING YOUR DATA

Sometimes, you may want to do more with analysing your data than is presented in the various reports provided by the portal. For this reason, the portal gives you the opportunity for downloading three different datasets which may be of use.

Near the top of the main portal page is a tab, "Download data". When you click on this, the three options show up.

The first option is a simple listing of the school's students' Typical and Achieved marks in each course for the year. The second option, described in more detail later, gives a fuller picture of the data held on the server for each student. The third download is the set of mapping points published each year by the Universities Admissions Centre, which are used to produce TES-scaled marks, as described in Module 15. In all three cases, the data downloads as a Microsoft Excel spreadsheet.

The first download file is simple. It lists each course in alphabetic order, then within that course, each student who took the course, again alphabetically. The Typical value listed is given to the first decimal place, but the Achieved value is rounded to a whole number. Note that Class and Factor data are downloaded too. By default, these have values of "A" and zero, but if you or anyone else have entered class or factor data as described in Modules 9 or 10, whatever you have entered will be downloaded here.

The second download file is packed with data. It has all of the data of the first file, and adds more. The Band Awarded is given beside the achieved mark. The one-unit value of the Moderated Assessment Mark is given beside the one-unit value of the Scaled Examination Mark. The next column provides the one-unit value of the achieved mark, re-scaled as described in Module 15 to the TES scale. The next column shows the "All Others" value for this student in this course. This is the value described in Module 3 of the average, on the TES scale, of all the other BDC courses besides this one that the student has attempted. Finally, the file includes all of the NAPLAN data available for this student. Note that the Literacy values for Years 9 and 7 are calculated as the average of Reading, Writing, and the mean of Spelling and Grammar/Punctuation.

The third download file gives whole-of-state data, rather than that for an individual school. It shows data published by UAC, and gives the current year's mapping points for each course which are used – as described in Module 15 – to calculate the "TES". All data in this file give one-unit values for each course. The first three columns are a summary of HSC marks received from NESA (referred to in this file by its old acronym, "BOS") Note that some of these data differ slightly from those published by NESA. The next six columns give the one-unit HSC score for firstly, the highest score in the course, and then for a student on the 99th, 90th, 75th, 50th and 25th percentiles. These are matched in the last 6 columns by the corresponding data on the TES scale. These two sets of data then provide the mapping points which can be used to get a reasonably accurate estimation of the TES value of any mark on the HSC scale provided by NESA.

The point of having download files available from the Analysis is to enable the development of further insights beyond those already available through the portal. Some of these are discussed in the final module, number 17, which looks at the use of analysis data in advising students. However, it is important to ensure that any use of the data from the analysis is both educationally and statistically sound.

MODULE 17: USING THE ANALYSIS TO ADVISE STUDENTS ON CHOICES

Teacher advice to students can have a considerable impact on the courses students choose. This module will briefly look at the whole field of advice to students, before highlighting some aspects of the Analysis which are particularly useful in shedding light that may help both students and teachers.

There are three critical factors to balance in advising students on which courses they should do. The first of these is Interest. Look for real interest among the available options, rather than just attempts to game the ATAR system. The second factor is Challenge. Students should be stretched by courses they attempt. High expectations – from both the student and teachers – can effectively drive achievement. This does need some balance with Realism. Students have to know something of what different courses will involve, and their capacity to meet these requirements.

In the Resources section of the portal, there are two useful papers which are intended to help teachers in advising students. The first of these is Professor George Gooney Cooney & Dr John DeCourcy's 2016 paper, Choosing Wisely, Choosing Well. The intended audience for this is teachers, not students. Don't give it out to Year 10, use it as a discussion paper with staff. The second resource is Prof Cooney's paper on Catholic sector selection of more challenging courses. This paper demonstrates that as a sector, Catholic schools have fewer students taking high-level courses such as Extension Mathematics than would be expected by consideration of the students enrolled. What does your school's participation in high-level courses look like?

How well we are challenging our students can be considered in terms of participation, as was outlined in Module 14. Is our advice too heavily biased to failure-avoidance, and not enough to challenge?

There is a report in the Analysis, called "Ranked All" which is useful for shedding some light on common ATAR myths. This report has no student names, and presents every HSC result from the school in descending order of the one-unit value of that result on the TES scale. Hence, it shows which courses count more for an ATAR, and by how much. The two factors to take account of in reading this report are Equivalence and Frequency. Two results which are close in TES 1-unit value are more-or-less equivalent in their impact on the ATAR. However, you need to look also at the Frequency with which any given course gets results in a particular range of TES values.

Another way of looking at ATAR impact was outlined in Module 15, looking at the TES Details report. Because for each student these results are ranked according to their TES one-unit value, you can go through and highlight any course in this list, thereby giving a clear impression of the course's impact on ATAR.

Finally, it is possible to put some realism into students' and teachers' consideration of courses by looking at past performance. Module 16 outlines how to download data from the Analysis. This includes both NAPLAN and HSC data. If you download several year's data, you can build a scatterplot relating any HSC course to any NAPLAN result – or any other data you have available about those students. This can then be used in a discussion with a student about what has actually happened in the past, as distinct from trying to make predictions for the future.

There is much else beside the Analysis which should come to bear on student course decisions. The Analysis can assist by providing accurate data on past performance.