Sum Product

Still grafting on graphs?

We continue our preoccupation with charts this month as we look at the humble pie chart, the new charts of Excel 2016 (albeit in 2017) and how to create forecast charts with the click of a mouse.

There's the usual stuff as well. Our Power Query Pointers and Power Query / Power BI updates continue unabated, as do our regular features Keyboard Shortcuts and the A to Z of Excel Functions (we have found our **CALL**-ing this month)...

Until next month.

Liam Bastick, Managing Director, SumProduct

Nice as Pie

A lot of data visualisation critics have come down hard on the poor little pie chart. Edward Tufte says, "the only thing worse than a pie chart is several of them". Stephen Few says, "save the pies for dessert". My favourite is Cole Nussbaumer says, "Death to pie charts". There's some sympathy that the following is one of the world's more informative pie charts...

Percentage of Chart Which Resembles Pac-man

- Resembles Pac-man
- Does not resemble Pac-man
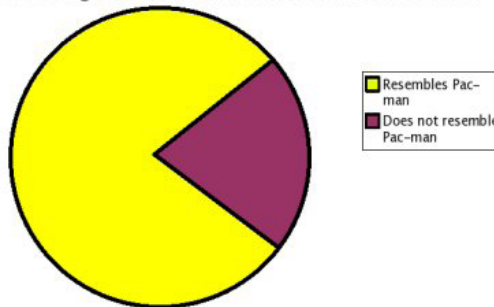
The humble pie charts may not be good at expressing accurate comparisons, but not every data exercise is about precision. It can be good for identifying largest contributors quickly and easily.

Pie charts are particularly useful when you have small number of categories and a great way of visualising data trends simply and quickly. It clearly illustrates how proportions have changed in the data more concisely then comparing bar charts side by side.

Excel doesn't provide a dual pie chart though. Solutions suggest a copy of the initial pie chart is made, place them side by side and group the two items together. However, there's a neat little trick that allows one to have their pie and eat it too. With a little work, the "pie of pie" chart in Excel can be utilised to generate the "Dual Pie chart".

It's not the most elegant, but one always needs to work hard to aspire to pie in the sky.
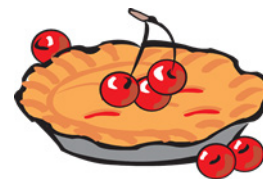
Let's at first look at the data set. In honour of the recently released Australian 2016 Census data, this article will use some data from the Australian Bureau of Statistics and compare how the age groups in the population have changed since the 1901 census to the 2016 census. We bet you're riveted:

Group	1901	2016
Age 0-19	858,539	5,786,217
Age 20-39	644,518	6,496,913
Age 40-59	334,100	6,142,602
Age 60-79	122,935	4,028,765
Age 80+	8,233	947,391
Total	**1,968,325**	**23,401,888**

People tend to read information left to right, so let's put Census 1901 on the left with the Census 2016 on the right.

First off, lets select the colours each group is to be represented by. This will become important later.

Group	Colour
Age 0-19	
Age 20-39	
Age 40-59	
Age 60-79	
Age 80+	

The trick works by using the pie of pie chart and **the smallest slice becomes the second pie chart**. The right-hand pie chart is actually made up with the proportions of that series however the total value must equal the smallest slice in the original series.

Thus, the proportion of the 2016 Census data needs to be calculated and then new numbers generated by using the *smallest category* of 1901 pro-rated. The Age 80+ group is the *smallest category* in the 1901 Census *viz.*

Group	1901	2016	2016 Proportions	2016 New Numbers
Age 0-19	858,539	5,786,217	24.73%	2,036
Age 20-39	644,518	6,496,913	27.76%	2,286
Age 40-59	334,100	6,142,602	26.25%	2,161
Age 60-79	122,935	4,028,765	17.22%	1,417
Age 80+	8,233	947,391	4.05%	333
Total	**1,968,325**	**23,401,888**	**100.00%**	**8,233**

To arrange the data for Excel to process through the chart, this is how it will be organised:

| Age Group	Value	Data
Age 0-19	858,539	1901
Age 20-39	644,518	1901
Age 40-59	334,100	1901
Age 60-79	122,935	1901
Age 0-19	2,036	2016 New Numbers
Age 20-39	2,286	2016 New Numbers
Age 40-59	2,161	2016 New Numbers
Age 60-79	1,417	2016 New Numbers
Age 80+	333	2016 New Numbers
Total	**1,968,325**	

Now it's time to generate the chart. Select the "Pie of Pie" Chart under the Recommended Charts or on the Ribbon from the pie chart drop down:
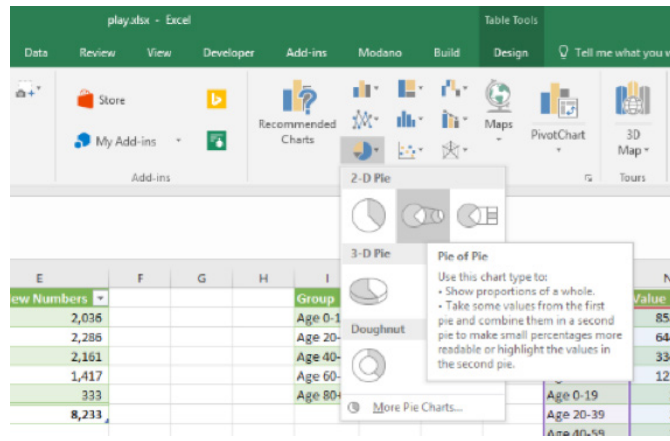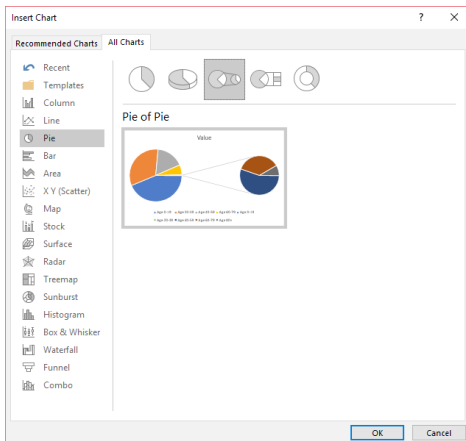
The automatically generated Excel Chart will look like this:
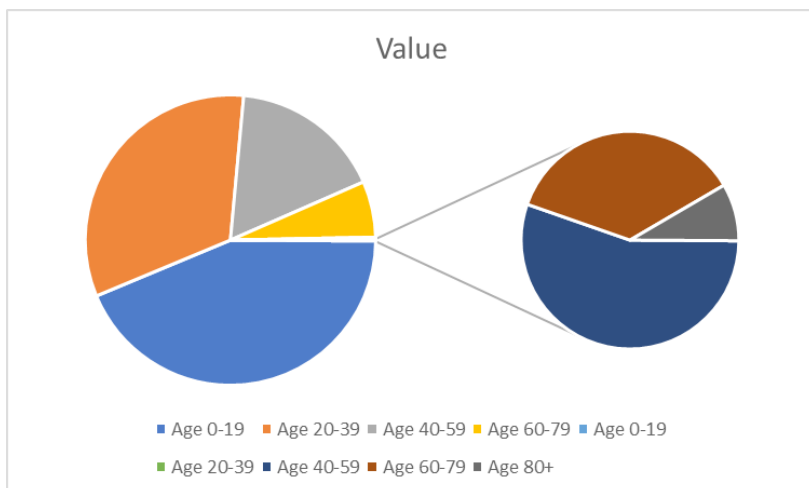
Now it's time to do some formatting magic. Right-click on either of the pies and the *"Format Data Series"* dialog should appear. Then, select the 'Series Options' menu. Change 'Split Series By' to 'Position' and 'Values in second plot' should be set to the number of categories that the data has, in this case five. The 'Second plot size' indicates the size of the second pie in relation to the first, so enter 100% to get identical sizing.
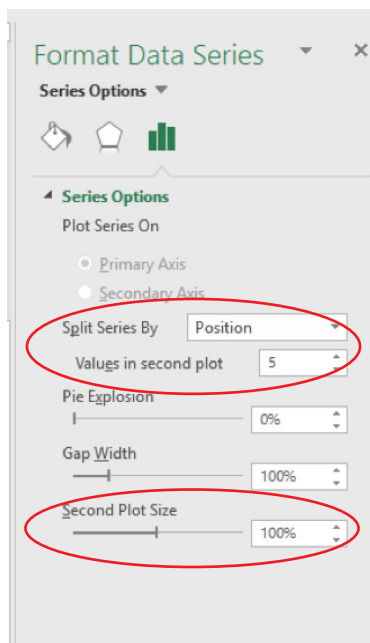

The chart should look like this:
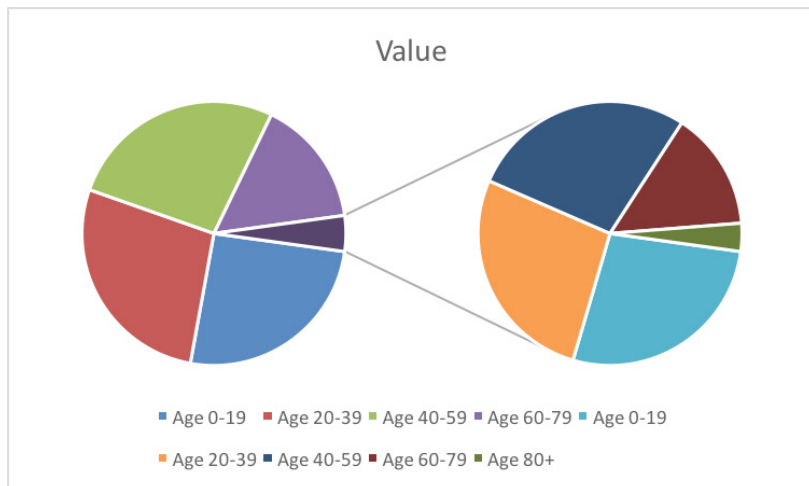
Click on the lines that connect the two pies and press delete. Then, rather arduously, go through each data point to recolour it to the colour section. Do note that in the first pie, the "Other" series (which is supposed to represent the smaller pie) will be the Age 80+ colour. The recoloured pies will look like this:
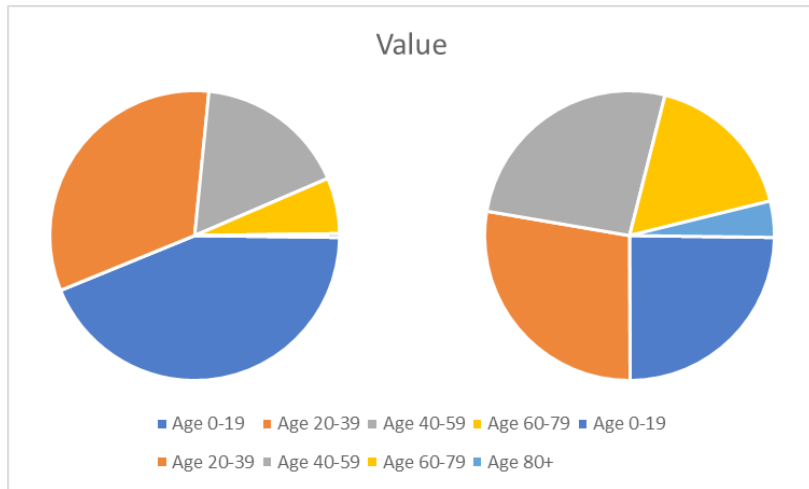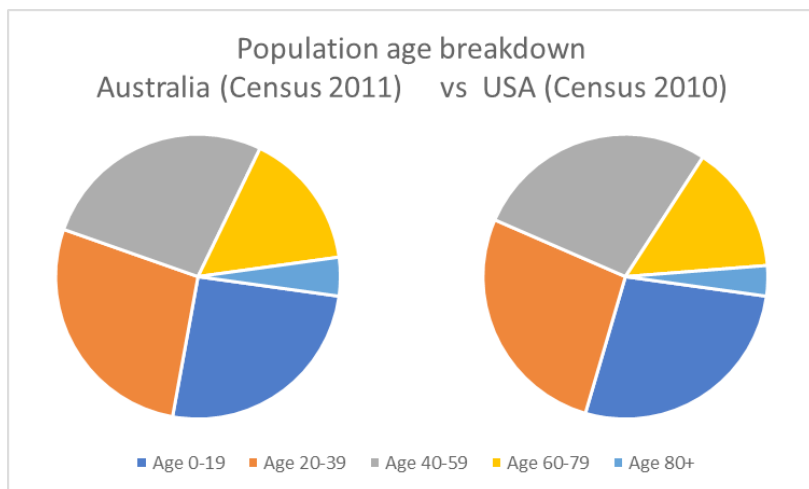

All you have to do now is fix the title and the legend which is rather straightforward (just delete the duplicate legend entries):
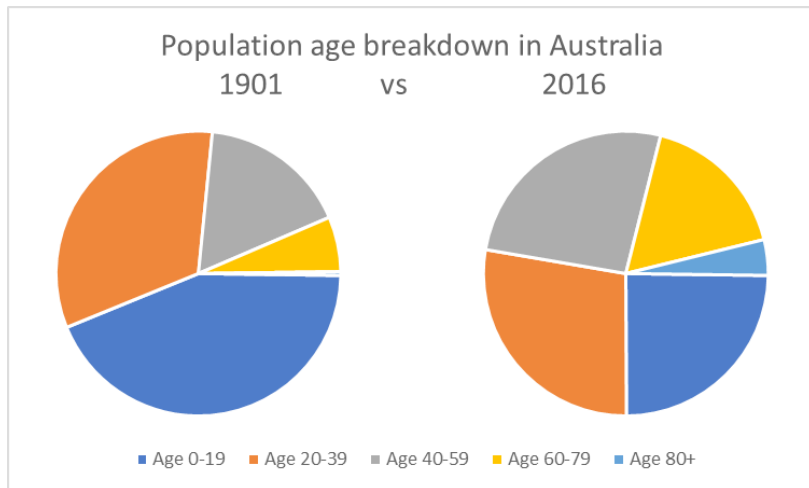
As you can see, the population proportion in 1901 is substantially younger than in 2016.

But just to have some American Pie, let's see Australia vs the USA:

Not so different after all – and easy as pie to construct!

New Charts in Excel 2016

Breaking news! "Excel 2016 comes with thrilling new charts options!". This is what we would have said if we published this article in 2016. Well we did to an extent, but given we have a focus on charts presently we thought we'd look at them in a little more detail – especially as so many people still don't use Excel 2016!
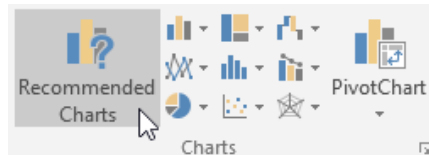
Waterfall Charts

Waterfall charts are a great medium to use to illustrate the impact of elements of your data. More specifically, they can be useful for analysing financial data in income statements or data over time, to see where you started and how you got to where you are right now.
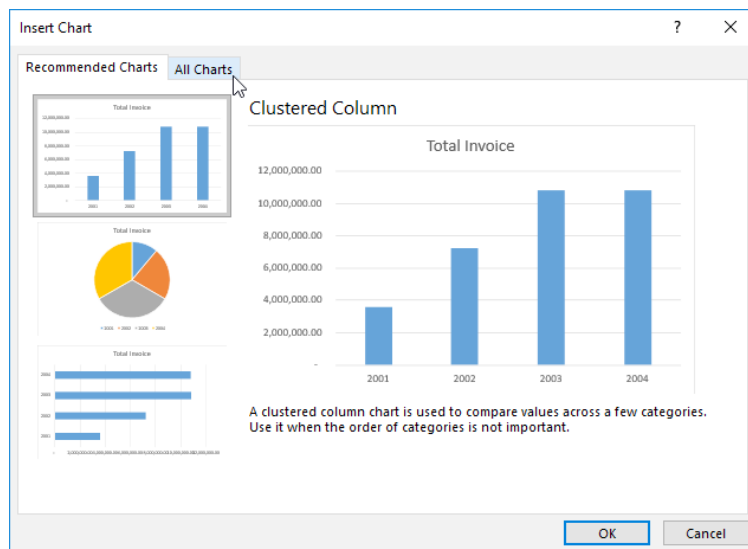
The first step to create a waterfall chart (or any chart for that matter) is to select your data and go to the 'Insert' tab on the Ribbon:
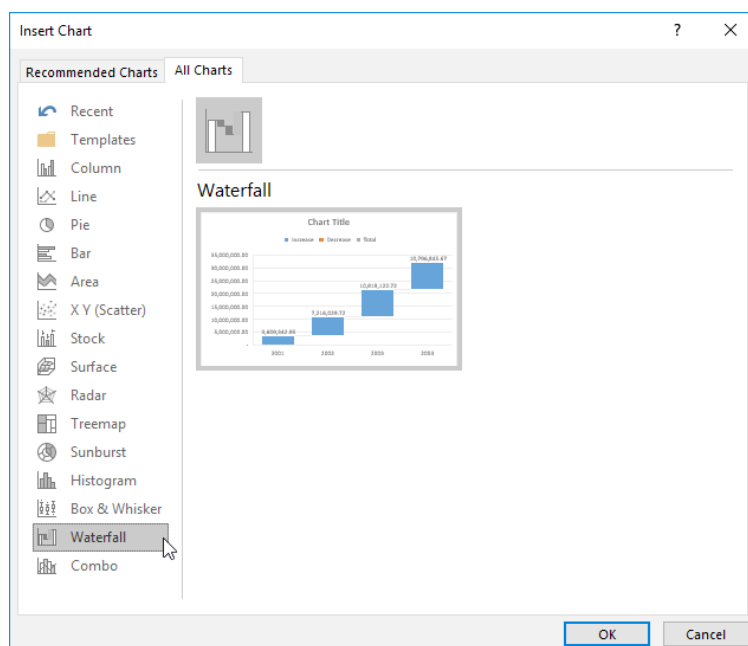

 and select the Recommended Charts option:

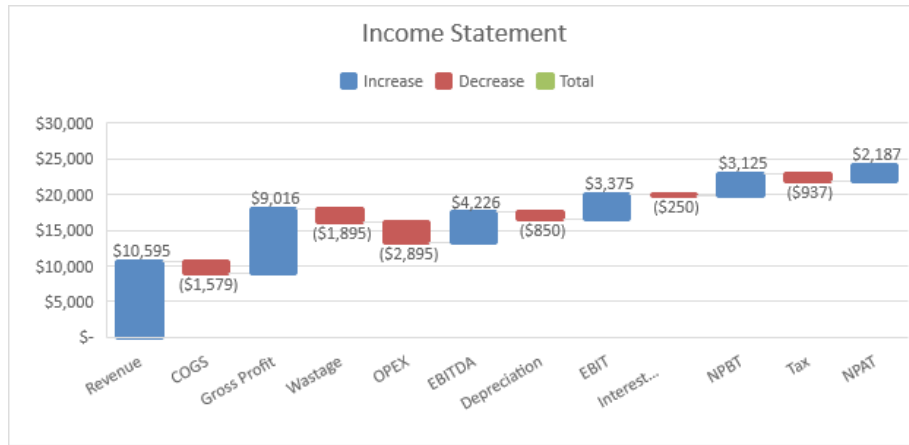Then select the 'All Charts' tab on the 'Insert Chart' dialog box:

Then select the 'Waterfall' chart option:

The resulting waterfall chart from our sample data is pictured below:

| Years	AUD (1,000's)
Revenue	10,595
COGS	(1,579)
Gross Profit	9,016
Wastage	(1,895)
OPEX	(2,895)
EBITDA	4,226
Depreciation	(850)
EBIT	3,375
Interest Expense	(250)
NPBT	3,125
Tax	(937)
NPAT	2,187

Excel automatically recognises negative values and colours them differently to positive values. However, Excel does not automatically recognise values that are meant to be totals in the waterfall chart. To change this right-click on each column you wish to change to a total and set it as a total.
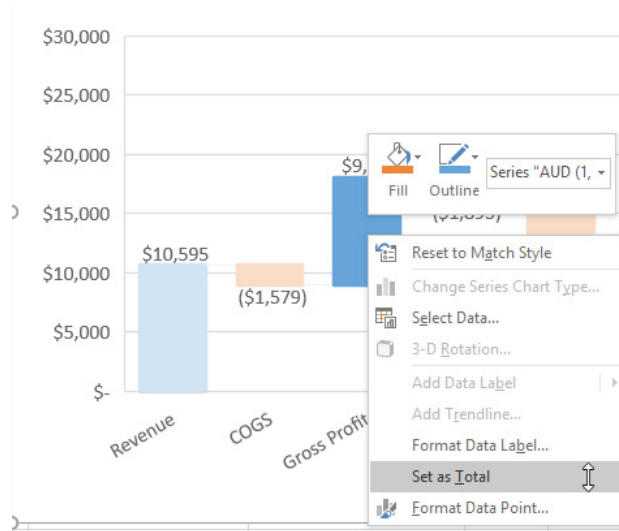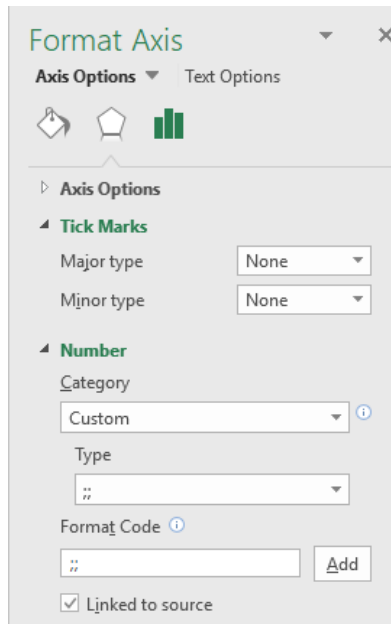

After setting all of the appropriate column as totals you will have the following chart:
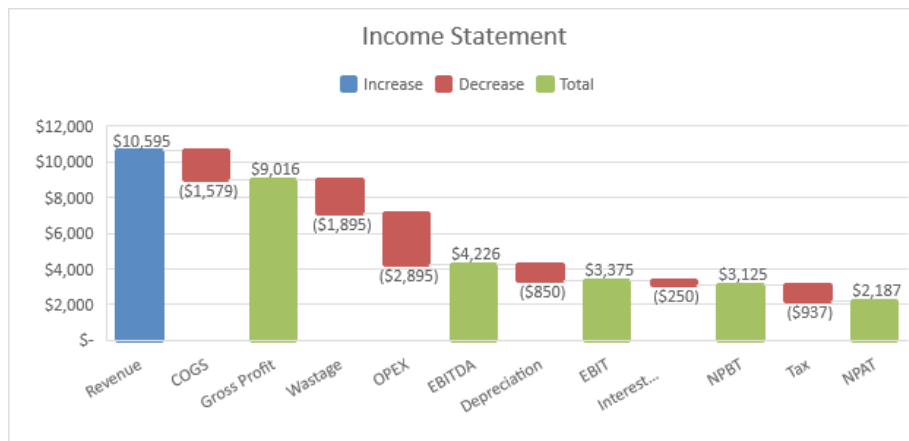
These colours and various other aspects of the Waterfall chart can be changed in the 'Design' tab on the Ribbon.
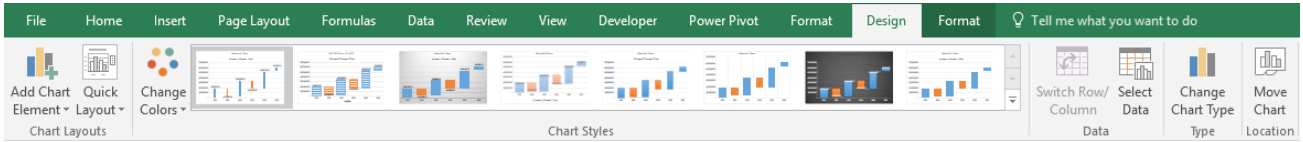
You can further customise the axis in the right-side panel, toggle tick marks on or off, or change the number type of the axis labels:


Histogram Chart

Histogram charts are best employed to show the distribution and range of data. These should come in handy for illustrating the distribution of sales in a line of products.

For instance, we have sample data of raw sales containing more than 300 rows of data. We create a Histogram chart with this data with the goal of determining the distribution of sales depending on sale amount.

| Total Sales |
| --- |
| 69 |
| 95 |
| 47 |
| 15 |
| 32 |
| 30 |
| 12 |
| 55 |
| 47 |
| 44 |
| 55 |

The initial chart that Excel gives us will look like this:
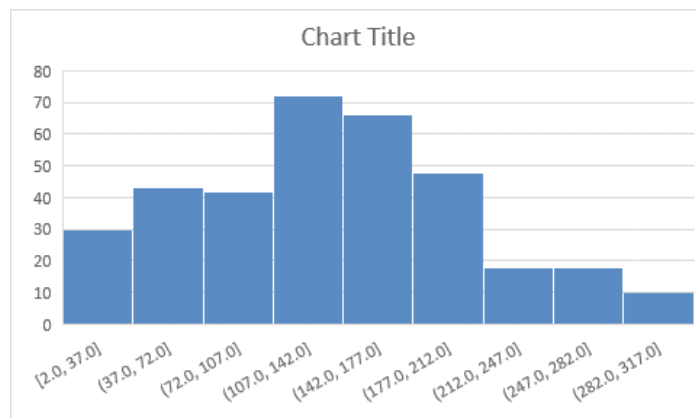
To add more 'bins' and customise other features, click on the horizontal axis then navigate to the right-side bar and click on the 'Axis Options' icon:
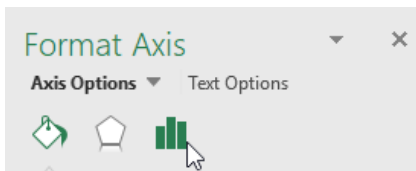
From there you can vary the number of 'bins' you want or the Bin width, and Excel will automatically determine the number of bins depending on the range of data. You can use the 'Overflow' or 'Underflow' bins to create a bin for all the values above or below the value in the box adjacent to the respective option.
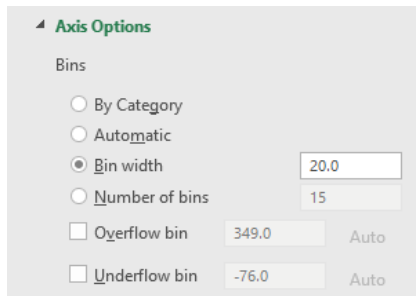
After selecting 20 bins the histogram chart will look like this:
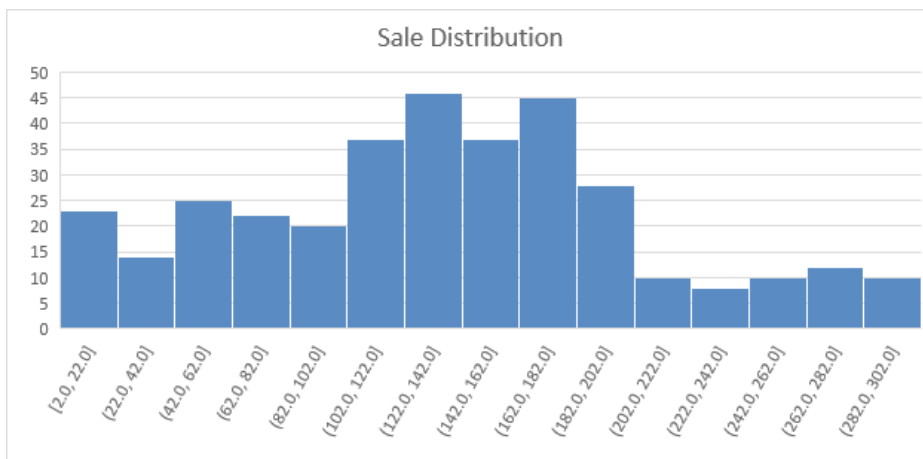
The histogram now shows us that the most popular products are the ones priced between $102 to $202 mark and the least popular products are priced between $222 dollars and $302.
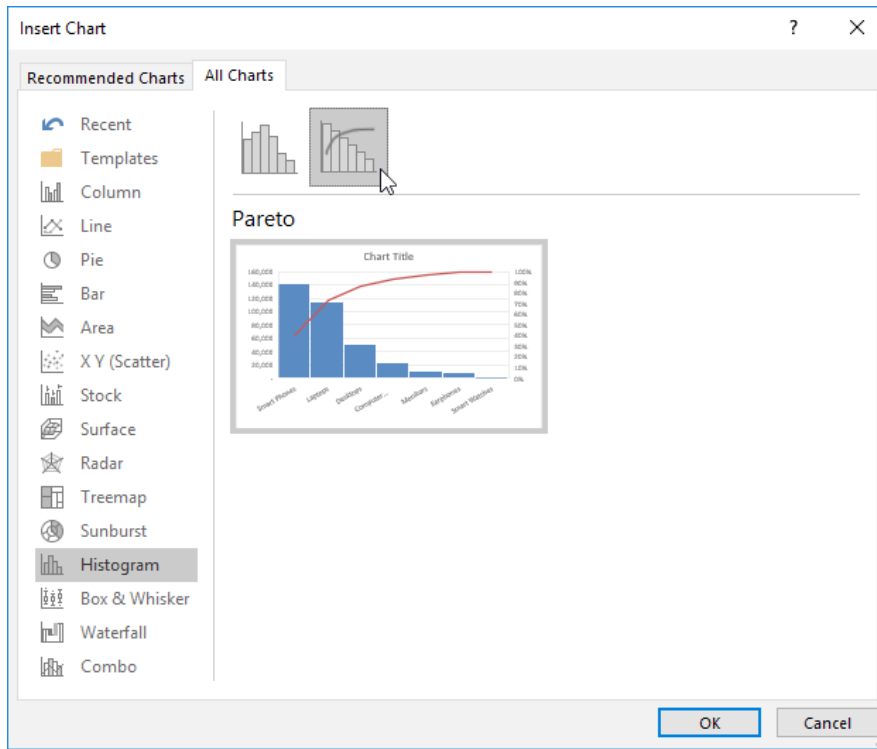
Pareto Chart

The Pareto chart combines the column chart and the line chart. The bars on the chart represent the data values in descending order, whilst the line on the same chart represents the cumulative percentage of the total amount.

For example, let's imagine that we are looking at the quarterly sales across different product categories for a small computer store:
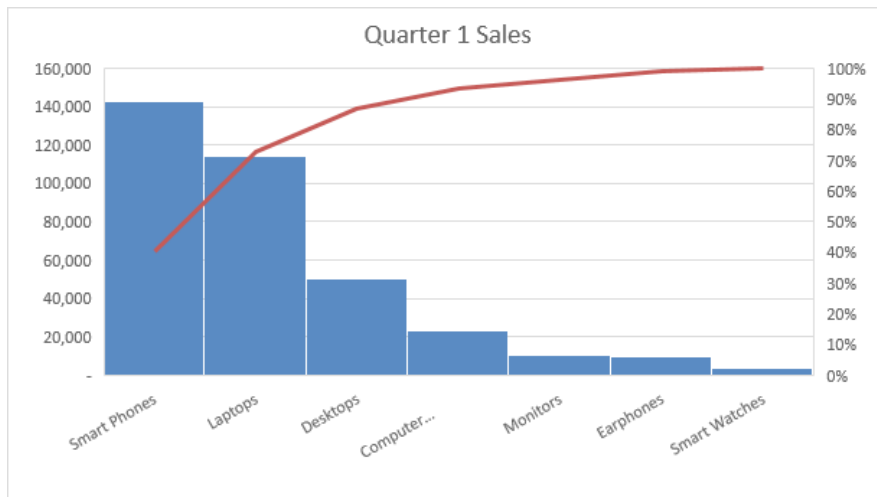
Product	Total Sales Q1
Monitors	10,596
Desktops	50,160
Laptops	114,360
Smart Phones	142,480
Smart Watches	3,400
Earphones	9,800
Computer Accessories	23,400

Select the data, then Insert a Pareto chart; note that the Pareto chart option is hidden within the Histogram chart option *viz*.


After giving our chart a name it should look something like this:

The Pareto chart makes it easy to see that Smart Phones, Laptops, and Desktops make up nearly 90% of all sales from the store.

Just like Histograms, you can customise the bin settings in the right-side panel under 'Axis Options'.
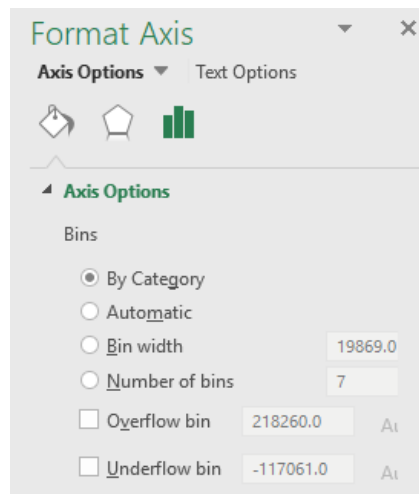
The line and column colours can be changed in the 'Format' tab on the Ribbon. The chart style can be changed in the 'Design' tab.
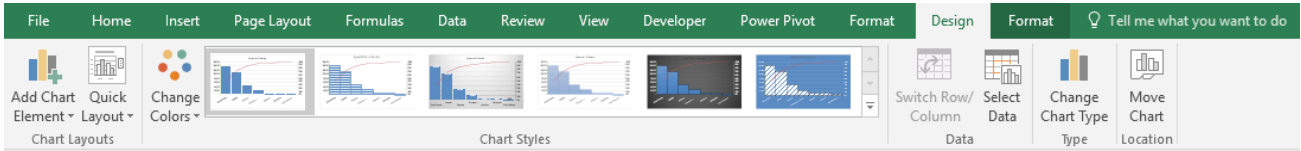

Box & Whisker Chart

This is one we have started to use quite a lot, realising what information it easily conveys. The Box & Whisker chart is commonly used in statistics to show the distribution of data in quartiles. It also shows the mean of the data and displays any outliers as separate points elsewhere in the chart.

To create the box and whisker chart, highlight your data then insert a chart, then select the 'Box & Whisker' option.
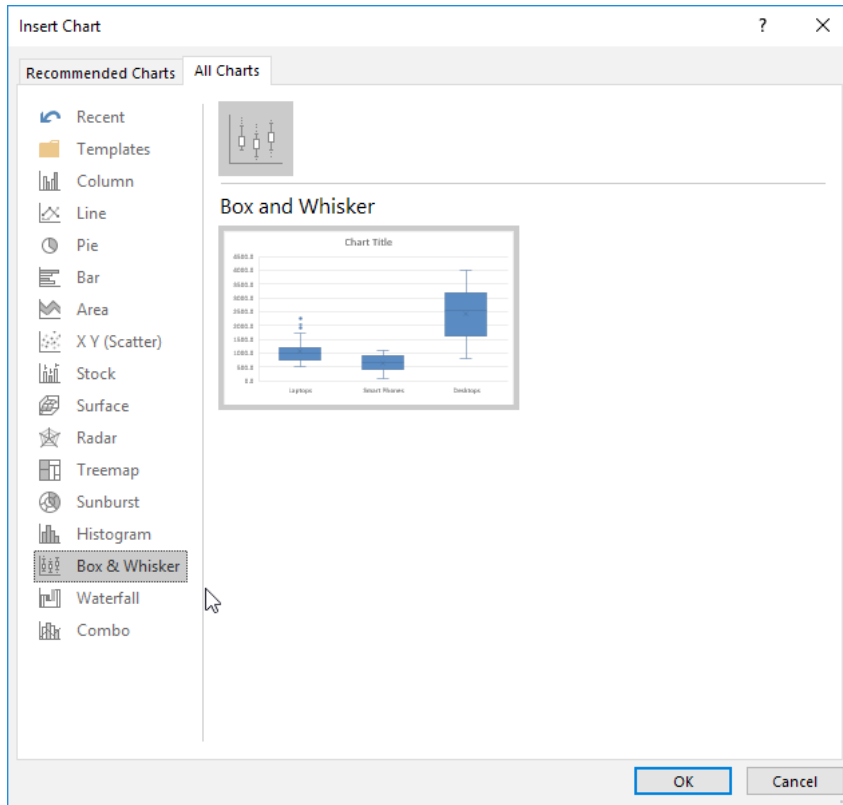
In our example the Box & Whisker chart allows us to visualise the variation of price for laptops, smart phones, and desktops, as well as other key metrics such as the mean and any outliers:
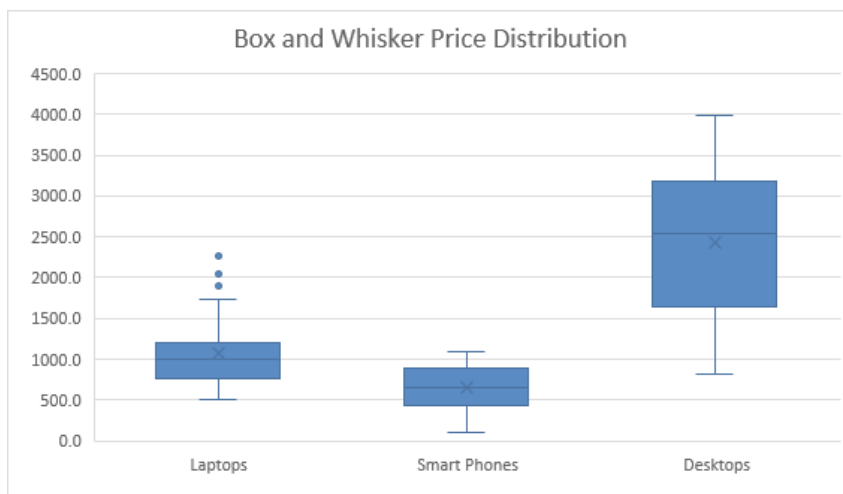
The first quartile is the lower line of the horizontal line of the lower whisker to the lower line of the box. The second quartile is the lower line of the box to the middle line of the box. The third quartile is the middle line of the box to the upper line of the box. Finally, the fourth quartile is the upper line of the box to the line of the upper whisker in the chart. The mean is represented by the 'X' symbol, usually appearing within the box. Any outliers are represented by dots, as prominently seen above the laptops' Box & Whisker.

The chart illustrates that the products with the largest amount of price variation are Desktops, but that the products with the largest skew are laptops.

You can apply further modifications to the chart from the 'Design' and 'Format' tabs on the Ribbon.
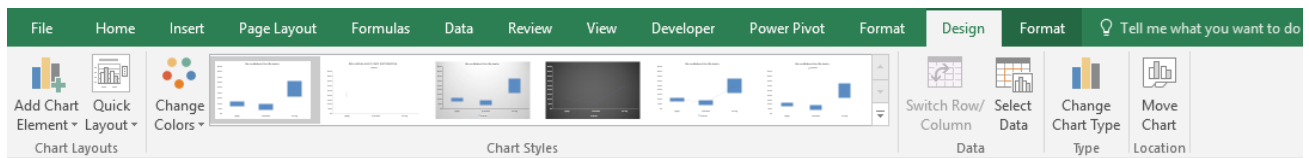
TreeMap

The TreeMap chart visually represents data in hierarchical fashion, which allows easy recognition of patterns. The TreeMap chart works by using colour coded rectangles that vary in size based on the assigned value or quantity.

For our example, we will use the following data table to produce our TreeMap chart:

Product	Subcategory	Total Sales
Baked Goods	White Bread	718
	Rye Bread	2,869
	Wholemeal Bread	8,989
	Grains Bread	5,972
	Muffins	9,658
	Cupcakes	3,319
	Cakes	302
Fruits	Apples	6,974
	Oranges	466
	Pineapples	1,602
	Grapes	3,574
	Bananas	2,950
	Pears	7,719
Meat	Beef	3,863
	Pork	9,139
	Chicken	3,959
	Kangaroo	9,220
	Monkey	8,560
	Aligator	3,585
Seafood	Squid	5,576
	Fish	5,402
	Octopus	5,847
	Shrimp	1,655
	Crab	268
	Clams	9,464
	Jelly Fish	7,024

After highlight the data, we may then insert the TreeMap chart:
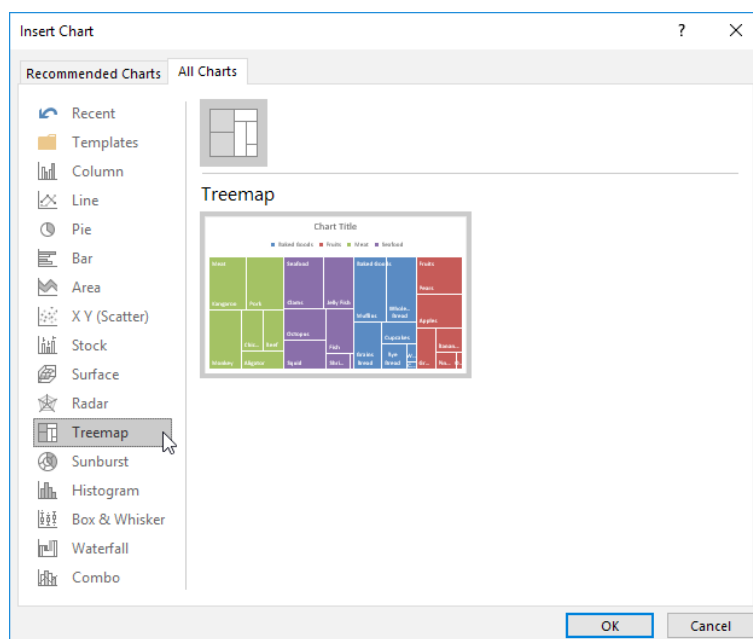
The resulting TreeMap chart clearly distinguishes between each product type with a different colour, and organizes the products based on size from right to left. It also highlights how much each subcategory sale makes up of the total product sale based on their individually sized boxes within each product. This can be a very useful way to give a summary of data 'at a glance':
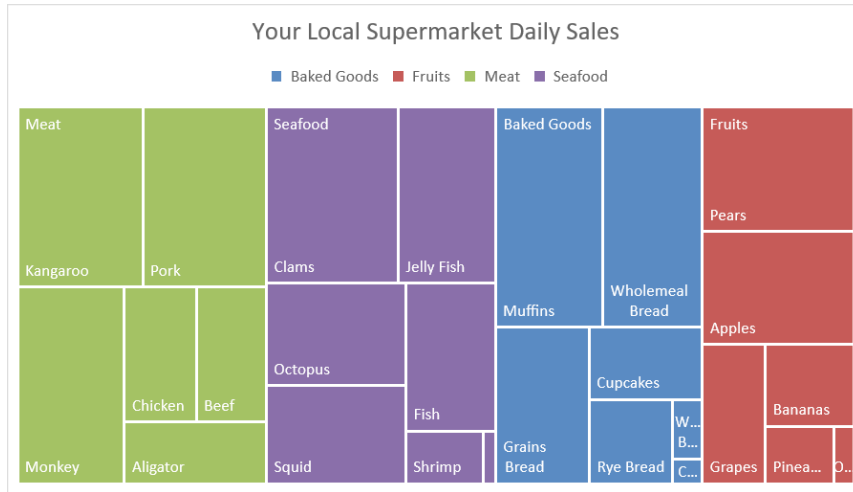

You can use the 'Design' and 'Format' tabs on the Ribbon to apply further customisations to the TreeMap chart, such as changing the colour scheme or apply borders around the chart.

Sunburst Chart

The Sunburst chart provides a hierarchical visualisation of data by representing data in the form of concentric circles (it looks like a seating chart at a concert or a sporting event). The data is further organised into multiple levels of categories, *e.g.* the inner circle in the Sunburst chart below are the main product types. The subsequent outer circle represents the sub-category within each product type, the proportion of how much each sub-category makes up is illustrated by size. The third outer circle represents another sub-sub-category for the sub-categories.

For instance, kangaroo meat makes up roughly a quarter of all meat sold, and roughly half of all kangaroo meat sold is grounded.
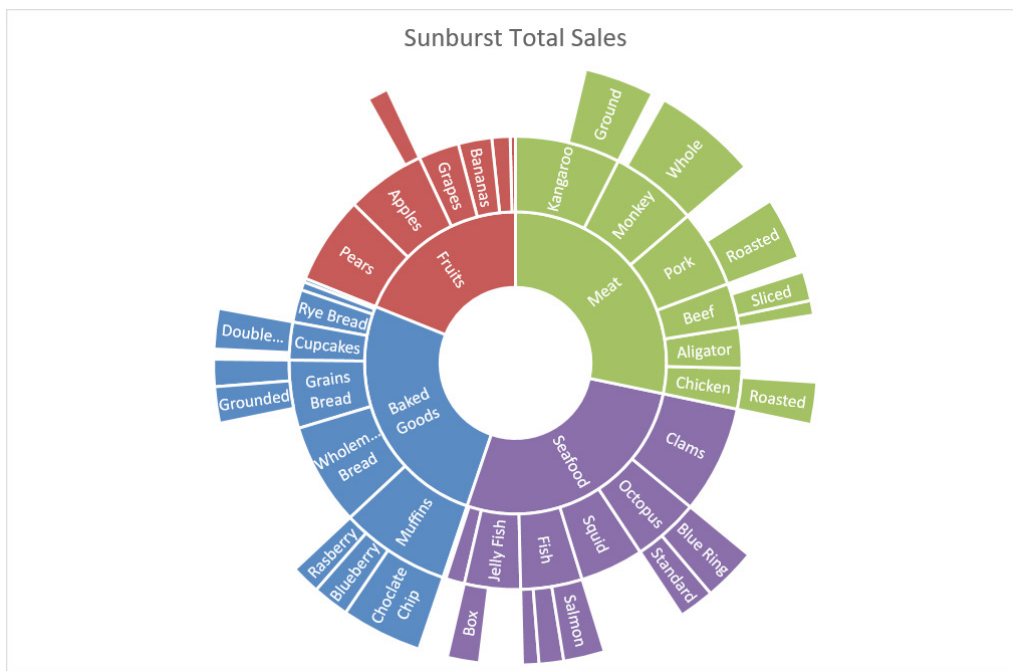

The Sunburst chart is able to display incomplete data; we do not have data on the other 50% of Kangaroo meat sold, so that region remains blank.

To create this Sunburst chart, we used similar data from the grocery store scenario, except it had to be organised in this fashion:

Product	Subcategory	Sub-Sub Category	Total Sales
Baked Goods	White Bread		718
	Rye Bread		2,869
	Wholemeal Bread		8,989
	Grains Bread		1,792
		Grounded	2,389
		Chopped Grains	1,792
	Muffins	Choclate Chip	5,312
		Blueberry	2,415
		Rasberry	1,932
	Cupcakes		664
		Double Choccolate Chip	2,655
	Cakes	Chocolate	30
		Fruit	60
		Vanilla	91
		Mix	151
Fruits	Apples		5,579
		Fuji	1,395
	Oranges		466
	Pineapples		1,602
	Grapes		3,574
	Bananas		2,950
	Pears		7,719

The next step is to highlight all of the data, click on the 'Insert' tab on the Ribbon then select the 'Recommended Charts' option. Then, scroll down to the Sunburst chart option and click 'OK':
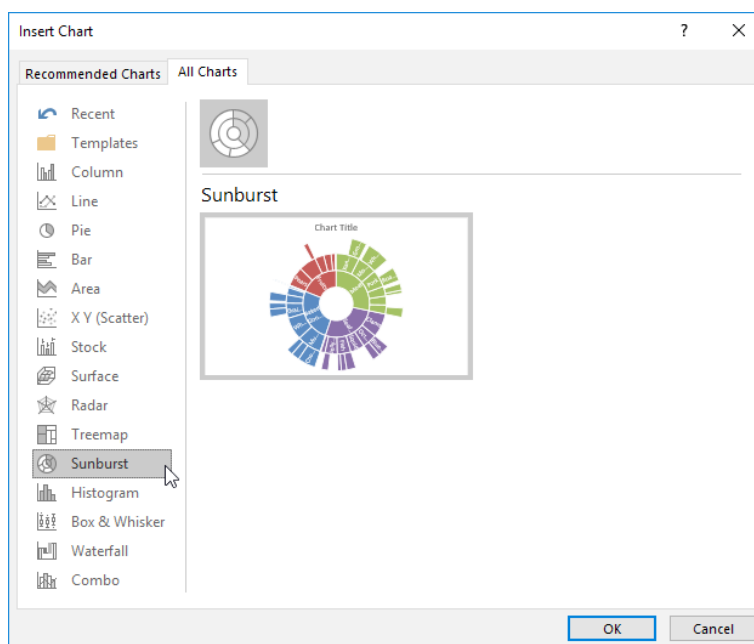
That concludes our take on Excel's 'relatively' new chart options! We hope you've found something useful in this article, if you have a query or comment for this section, please feel free to drop us a line at contact@sumproduct.com or visit the website, www.sumproduct.com.

Seasonal Forecasting

Let's face an ugly truth in the world of finance / accounting. It's the operational managers at the coalface who have the best understanding of likely demand and expected costs, yet many do not have the necessary tools or financial skills to forecast to the level that senior management requires. Similarly, most analysts do have these skills but may not be close enough to the front line to be able to understand operations sufficiently.

Operational managers and analysts need to work together to improve forecasting. That goes without saying and is never a problem in reality as the two sides work beautifully together, skipping off into the sunset hand in hand after a productive day in the office, understanding each other's needs and issues.

Yeah, right. If you believe that you're probably a Board Director…

Meanwhile, back on Planet Earth, often we find managers and analysts in a state of "co-operational flux" (a new term I have just invented). Managers sometimes feel accountants / analysts request forecasts from scratch: do they have either the time or skills to prepare a zero-based budget, whereas the analysts feel if they spend time preparing the base budget it is then torn apart by their operational counterpart. The art of war / budgeting: don't you just love it?
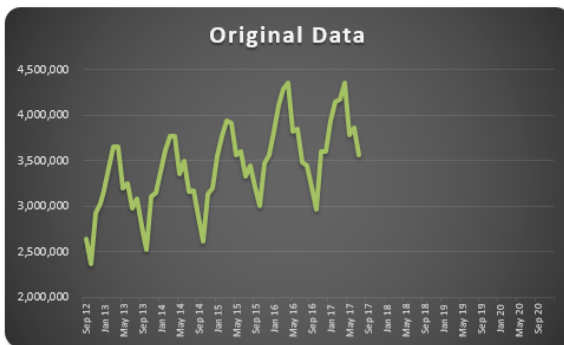
There is a need for *objective* forecasting. By this, I mean something that can be constructed simply such that if anyone follows the same process they will get the *same* figures. This needs to be a mechanical, objective process. This way, analysts may prepare this data in moments without feeling emotionally attached to the outputs. Furthermore, operational managers can review the trend and state where future numbers are wrong and all they need to do is explain the variation, *i.e.*

undertake incremental budgeting. There is no need to disagreements or confrontation: both parties may work together as a team. Simple!

Wow, is that the longest preamble ever..?

I wanted to set the scene for having a simple mechanical approach for budgeting. Let's take a look at ways Excel can do this for you. Imagine I have some historical data:

Original Data

Now let's be honest, anyone who has historical data looking this perfect should be referred to the auditors, but hey, this is for illustration purposes. I have data from September 2012 to July 2017. I want to extrapolate it until the end of 2020 (I want 2020 foresight!).
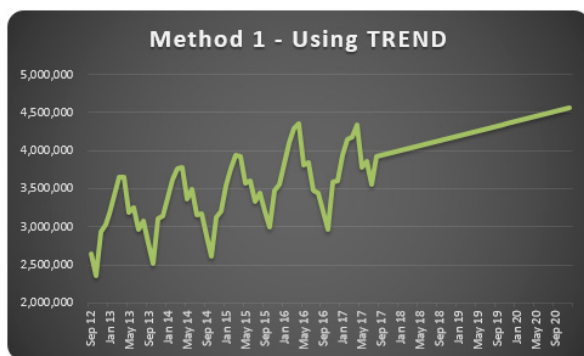
There are several functions that can help us here, with one of the simplest being **TREND**. **TREND(known_y's,known_x's,new_x's,[constant])** projects assuming that there is a relationship between two sets of variables **x** (independent variable – here, the dates) and **y** (dependent variable – the sales), through a formula **y = βx + c**, *i.e.* the equation of a

straight line (**β** is the gradient of the line and **c** is the y-intercept).

Before you disregard linear regression, bear in mind many non-linear relationships can become linear ones by taking logarithms of the variables, for example. However, this won't always work.

Here, time is our independent variable (**x**) and sales is our dependent variable (**y**). We only specify the constant if we want to force **c** in the equation (not common – it will usually be left blank). Similarly, it's preferable to leave **constant** blank in the **TREND** function. For example:

Here, we can extrapolate the data using the **TREND** function *viz*.

Method 1 - Using TREND

Ladies and gentleman, you may have heard of hockeystick projections; well, let me now present you with the swordfish. You extrapolate linearly, you get a *straight line*. Now who would like to present that projection to their senior management team?

This isn't good enough. We need to identify the cyclicality of the data. It appears to go through a cycle once every 12 months. This might not always be the case, but the concept remains the same even if the periodicity is not annual.

I want to calculate a periodic growth rate objectively. There's various ways I can do this. You might argue with me. That's fine. Feel free to write a brief note and send it to someone who cares. That's the problem here – it's *subjective* until your organisation defines how it is to be measured. Then, everyone follows that process and it becomes objective.

In my example, I am going to compare the sum of the sales over the 12 months ending 31 July 2017 with the forecast sales as calculated using **TREND** over the 12 months ending 31 July 2018:

Sum of last actual cycle ending 31 Jul 17	44,617,023
Sum of first forecast cycle ending 31 Jul 18	48,104,786
Growth implicit (CAGR)	7.82%

It is this percentage I will use to grow the forecasts. I then grow each period's value by its corresponding value in the previous period by this percentage (7.82% here). This gives me a much better chart:
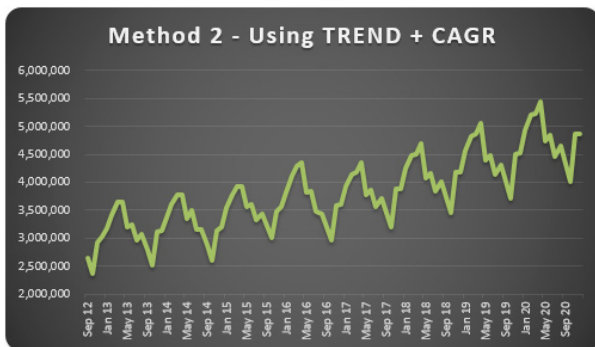
That looks *much* better. With practice, this approach doesn't take that long to prepare. Numbers may be varied from this forecast with the operational manager only having to explain these deviations. It makes life easier all round.

Once the method of assessing inferred growth rates based upon the **TREND** function have been agreed and what normalisations to historical data should be input, the process becomes much more straightforward. Of course, this method should be used for all forecast inputs separately and not just on their aggregation, otherwise confusion occurs due to sales mix changes, new products, cut-off periods, *etc.*

But there is an even faster way – if you happen to have Excel 2016…

**Exponential Triple Smoothing (ETS)** sounds like a dairy process, but it actually uses the weighted mean of past values for forecasting. It's popular in statistics as it adjusts for seasonal variations in data, like in the example above. For those who really need to know, Excel uses a variation of the Holt Winters ETS algorithm, although to be honest, I think you should get out more.

In Excel 2016, ETS has gone "native", *i.e.* it is a standard feature. This includes both a set of new functions such as **FORECAST.ETS** and other supporting functions for additional statistics. Your dataset does not need to be perfect, as the functions will accommodate up to 30% missing data.

But don't worry about using these functions. Simply highlight the actual data and click on the 'Forecast Sheet' button in the 'Forecast' group of the 'Data' tab of the Ribbon (**ALT A + FC**):
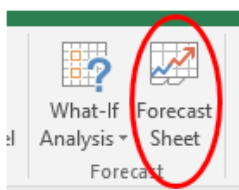
All you need to do is specify the final forecast period at the prompt and that's it. It produces a raw data sheet, together with confidence intervals (to demonstrate potential spread in the forecast error), which looks something like this:
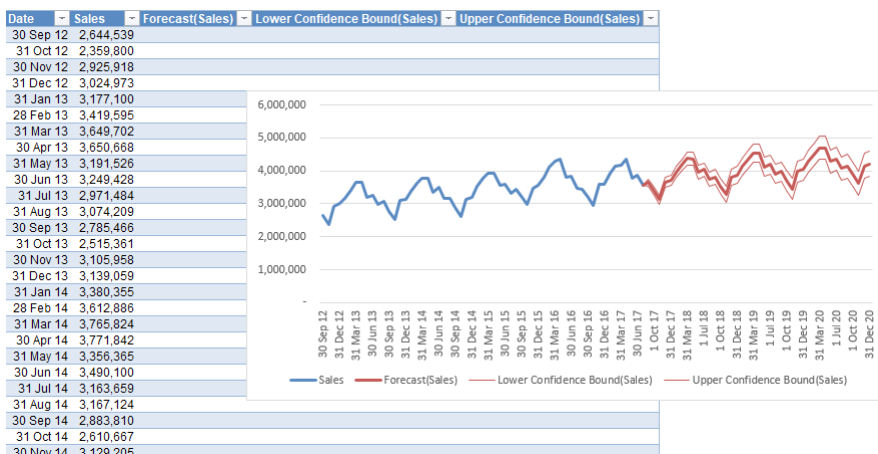
Manual versus automatic?

- The "manual" method, using **TREND**, assumes some linear relationship (possibly at a derivative level) and requires some initial subjectivity regarding the normalisations of actual data and how to determine what growth rate to use over what duration. Once it has been agreed, it becomes a simple process both to understand and maintain.

- The "automatic" method using the Forecast Sheet works it all out at the press of a button, after normalisations to historical data have been made. However, it's not all peaches and cream: it's only available in Excel 2016, it's quite "black box" to many and could you really explain to your line manager what Exponential Triple Smoothing is?

The jury remains out: whatever you decide, for your own sanity, I recommend an objective forecasting approach.

Power Query Pointers

Each month we'll reproduce one of our articles on Power Query (Excel 2010 and 2013) / Get & Transform (Excel 2016) from www.sumproduct.com/blog. *If you wish to read more in the meantime, simply check out our Blog section each Wednesday. This month we look at how queries may be merged.*

Power Query allows us to merge tables without either the need to be an expert on database structure or without having to learn formulae. You only require two (or more) existing queries.

So from previously, we have a query to load data from **ACCT_Order_Charges,** created in Power Query – (Data) Model Building, and let's merge data from another table so that we can include a description of the type of items that the charges apply to.

As with that first query creation, let's create a new query from an Access database; on the 'Power Query' tab in the 'Get External Data' section, then select 'From Database' and then choose the Access database where the tables reside:
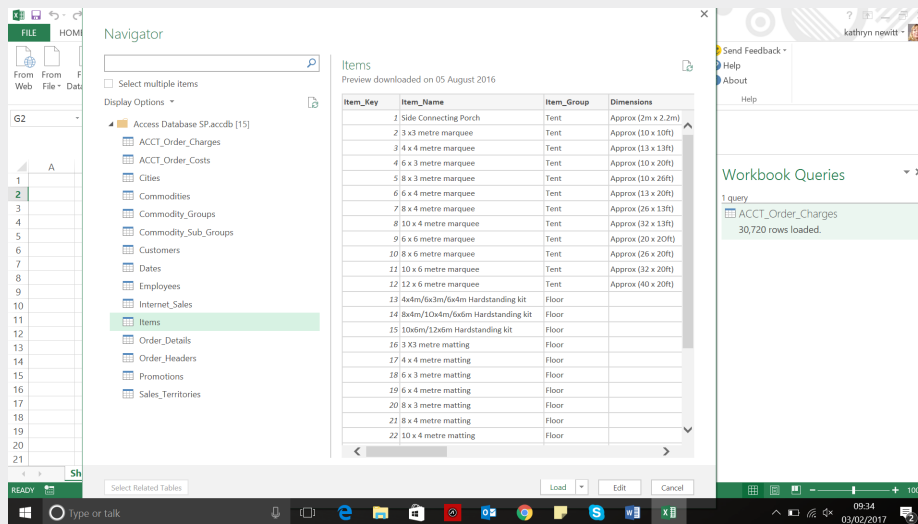

In the 'Navigator' window *(shown above)*, you can preview the table contents and choose the table that contains the data you're looking for, in this case the Items table. Note that the previous query is visible in the 'Workbook Queries' pane to the right of the screen. If the 'Workbook Queries' pane is not displayed, there is an option on the 'Power Query' tab in the 'Workbook Queries' section to 'Show Pane' accordingly.

Now, let's select the 'Edit' option and choose to extract the **Item_Group** and **Item_Name** columns from Items. As the reason for this extraction is to merge the data with the query on **ACCT_Order_Charges**, also choose to extract the **Item_Key** column to allow you to link the data together. You may select these three columns whilst holding down the **SHIFT** key and by right clicking, then choosing to delete the other columns.
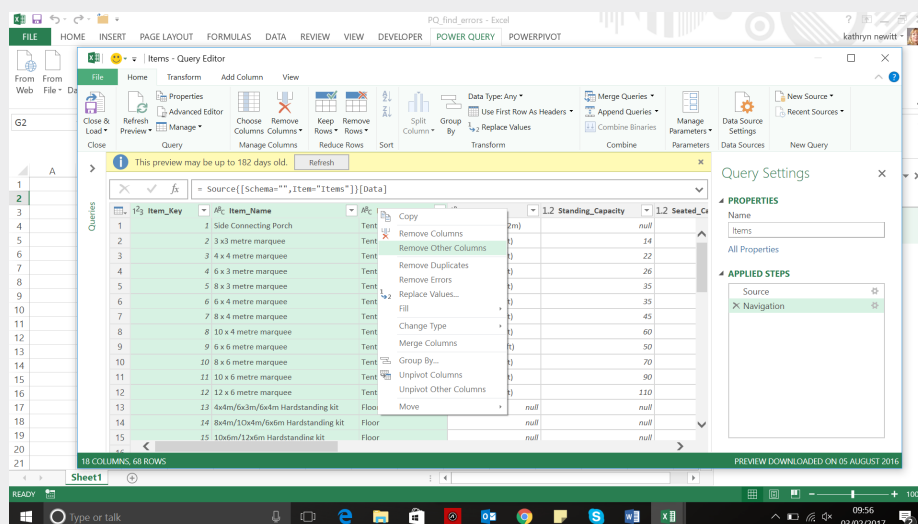
We don't wish to make further changes, so from the Home tab simply select 'Close and Load to Connection Only'.
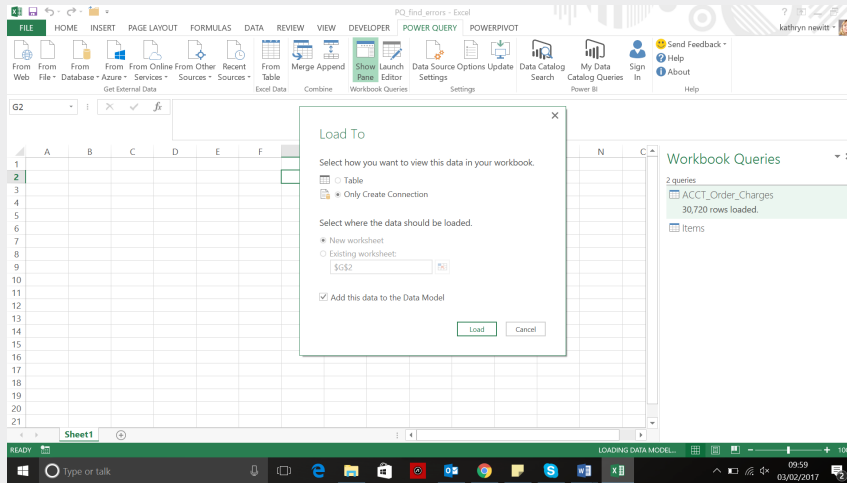
Note that the Items query appears in the 'Workbook Queries' pane, but is not complete yet. On the 'Load To' screen, the option to add the data to the Excel model can also be selected if required. Having made my selections as shown above, choose to 'Load'.

Now to join the queries – on the 'Power Query' tab there is a 'Combine' section where you may choose to 'Merge'. The following window appears where we can choose our two queries:
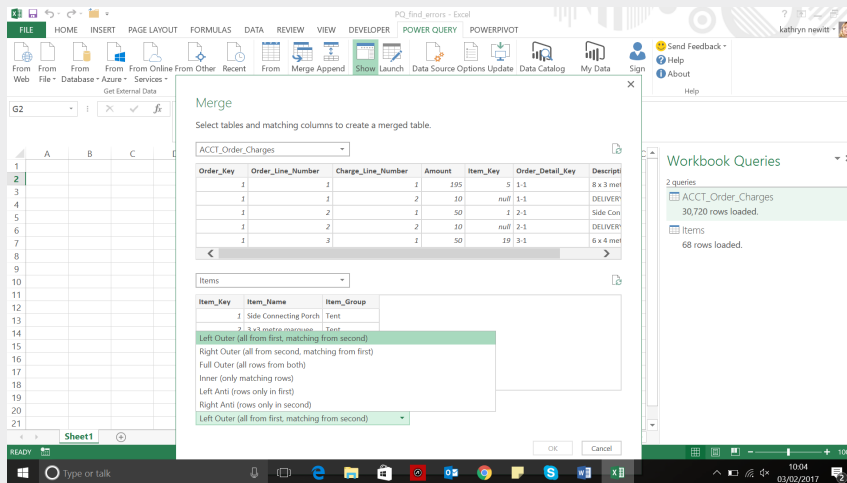

The dropdown at the bottom of the screen allows us to select how to link the tables; in this case choose 'Left Outer', as we want all of the *ACCT_Order_Charges* table and matching data from Items (if you had input Items first then we would have to pick 'Right Outer'). This is the point where some understanding of the data is required – picking 'Full Outer' can lead to duplicates.

The 'OK' button is not enabled yet – that is because we need to pick the columns to join. In this case, let's join *Item_Key* in both

tables, so select it on each table. As long as you select the same number of columns in each table, 'OK' will be enabled, which you can then click on.

The 'Query Editor' below shows all the columns from *ACCT_Order_Charges*, and a new column helpfully called *NewColumn*. Note the icon next to it with two arrows pointing away from each other – this will expand to show the new columns available:
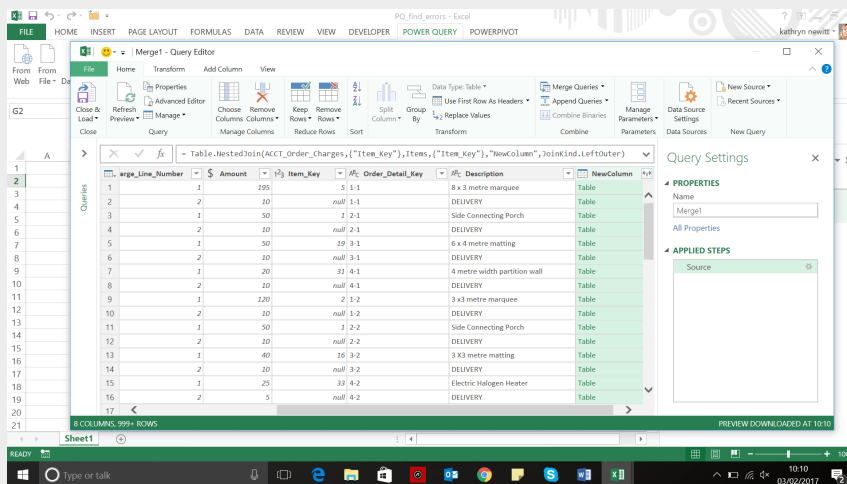
Now all the columns can be viewed, there is the option to uncheck columns that are already in the first table to prevent duplication. In this case description is very similar to *Item_Name* so uncheck *Item_Name* and *Item_Key* (since we only included the key to enable us to link the table data).
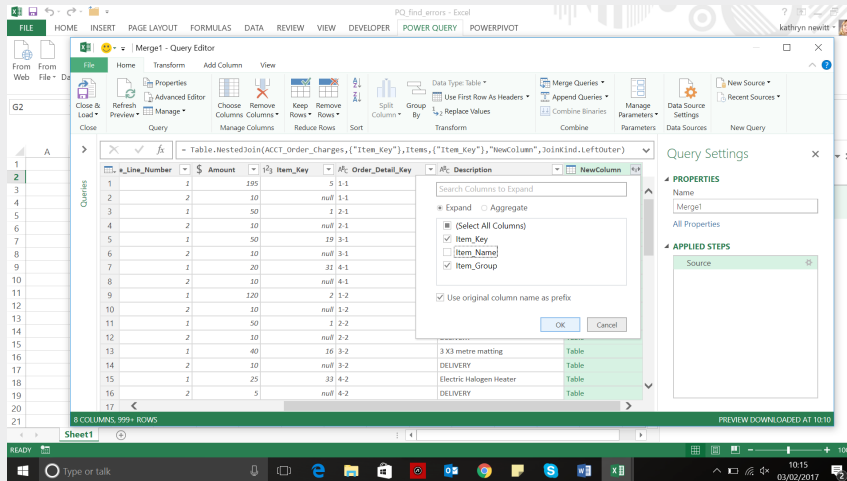
Then, click 'OK' and the data from *Items* appears with a slightly clunky column name (see below) which you can change to *Item_Group*.
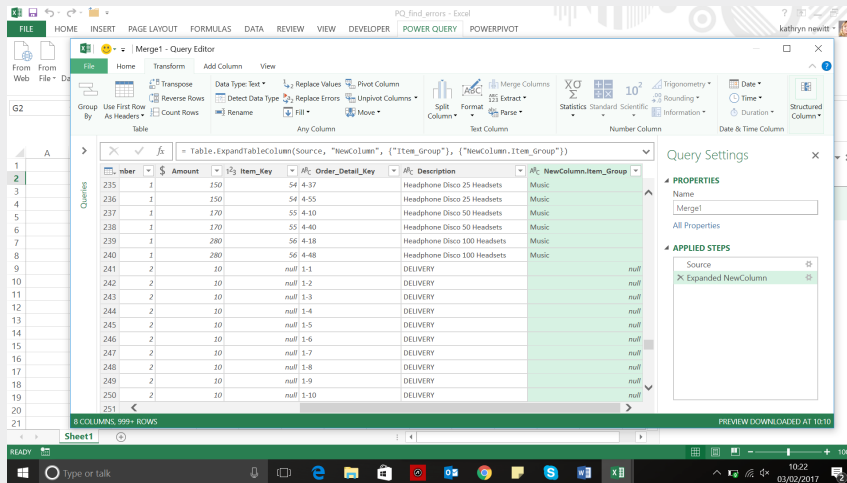
Note that the order of the rows has changed because the non-linked rows which had no item key have moved to the bottom. We can fix this by sorting on *Order_Key* and *Order_Line_Number* – there is a downward arrow icon next to each column heading which you can click and then choose to sort in ascending order – notice that Power Query combines these sorts into one step. Also choose to rename the table to something more meaningful than *Merge1*.
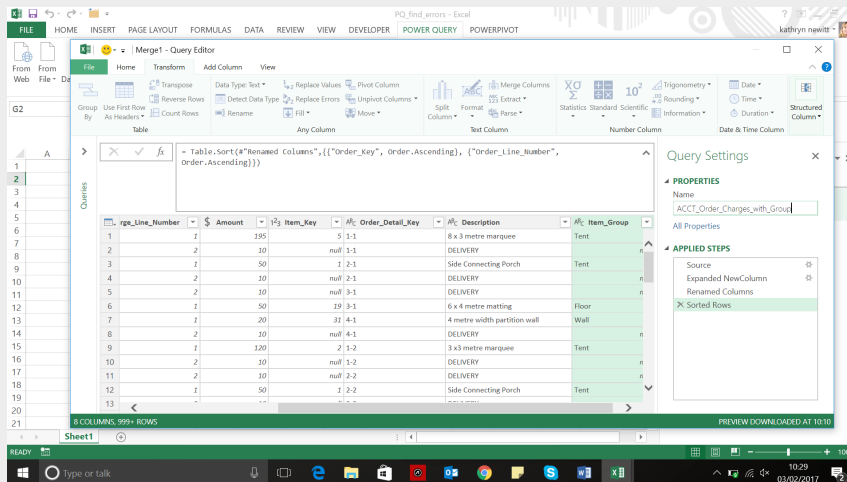
We can now 'Close & Load' to create a new merged table, which can be used in Excel workbooks and the Excel data model. When merging queries do keep in mind previous warnings in this blog series: if the queries are for use by Power Pivot only, then consider whether it would be more useful to load one merged table or separate tables that can be managed in Power Pivot. For creating Excel workbooks however, merging is a useful way of pulling in data from many tables in order to view the data in a single table. Data can of course come from other external sources.

More next time!

August 2017 Updates for Power BI Desktop

The Power BI Desktop juggernaut continued to speed forward at a rate of knots. Probably the key one this month (certainly many users have been excited about this since it was revealed at the Data Insights Summit) is the ability to interact with controls on a report in order to change a variable in the data – welcome to "What if?" analysis!

In summary, the following updates have been released:

Reporting:

- Show values on rows for matrix
- Colour scales on font colours for table and matrix
- Custom subtotal settings per level of matrix
- Line styles and legend options
- Scatter chart performance improvements
- New custom visuals:
    o Dot Plot by MAQ Software
    o Power KPI
    o Funnel plot
    o Beyondsoft Calendar

Analytics:

- What if parameters
- New scatter chart analytics features:
    o Symmetry shading
    o Ratio line
- New quick measure:
    o Weighted average

Data Connectivity:

- Live connect to the Power BI service is generally available
- Google BigQuery connector (beta).

Let's take a look at each of these in turn.

Show values on rows for matrix

After July's announcement that the new table and matrix visuals are generally available, Microsoft has added new features to these visuals. The first is the ability to show the values as row headers instead of column headers. For example, if you have multiple values in the Values field well bucket, they will by default show as column headers in the matrix:

If you would prefer for them to show as row headers instead, there's now a formatting option, 'Show on rows', that when it's switched to 'On' will show each of the values on the rows of your matrix:
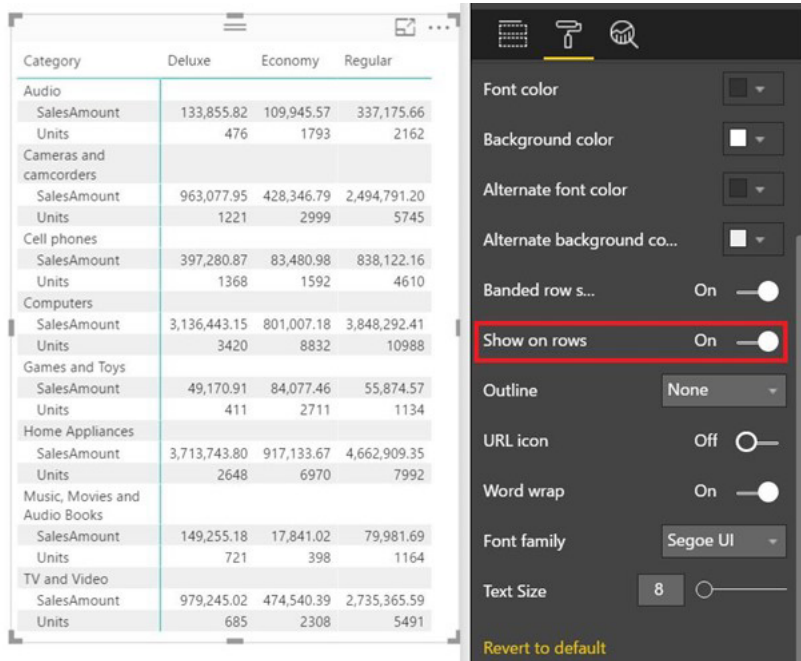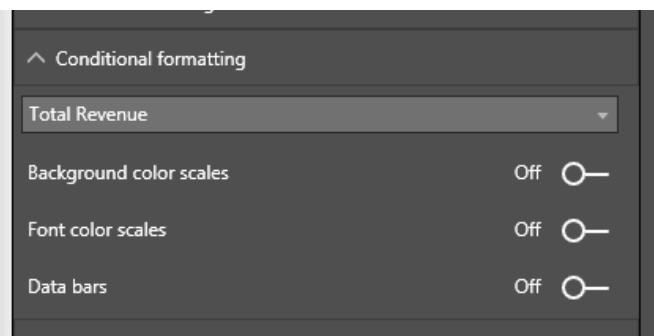

Colour scales on font colours for table and matrix
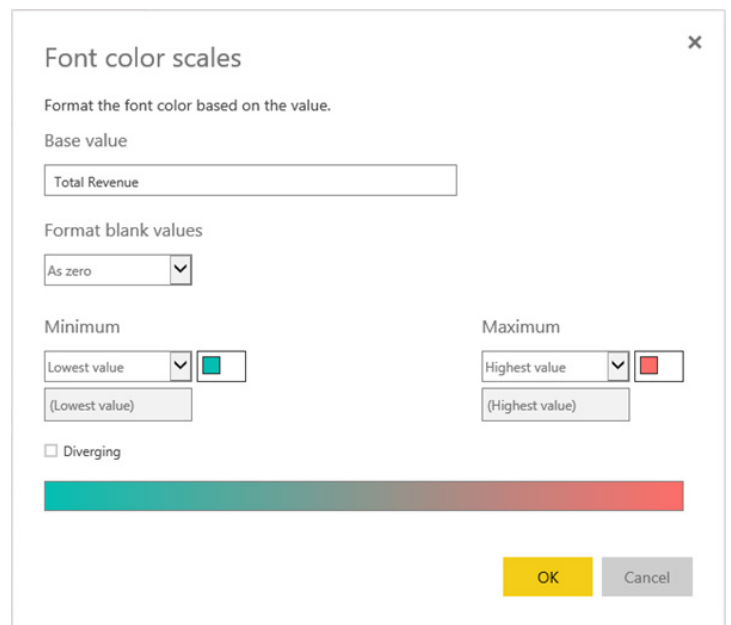
You can now apply conditional formatting to the column font colour instead of the background of cells. This option is situated on the value's menu:
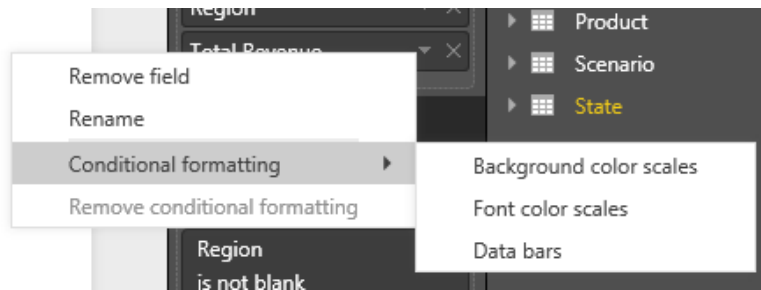
or in the formatting pane:

Once you select 'Font color scales', you'll see a dialog similar to the background colour scales that lets you customise both the colours you use and the range specified.

Once you have the style you want and close the dialog box, the formatting specified will be applied to that value's text:

Custom subtotal settings per level of matrix

Another new feature out in August for the matrix visual is the ability to customise whether subtotals are displayed for each level of the matrix hierarchy for both rows and columns. This option is located in the 'Subtotals' card of the formatting pane:

That's not all. Another new subtotal customisation is an option to move the row subtotals to the bottom of the group instead of the top, just like with PivotTables. You can also find the 'Row subtotal position' option in the 'Subtotals' card of the formatting pane too:

Line styles and legend options

Line charts have had more formatting options added as well. You may now set the style for line charts to be dotted and / or dashed in addition to the standard solid line style. Yes, we know this has been in Excel for years, but Power BI Desktop is on the catch-up! This will now mean it's easier to distinguish lines without having to rely on colour.

This month Microsoft has also introduced the ability to customise how your line chart's legend is styled. You can pick between markers only, the line style only or a combination of line and marker. This allows you to choose the legend option that best displays the chart key:
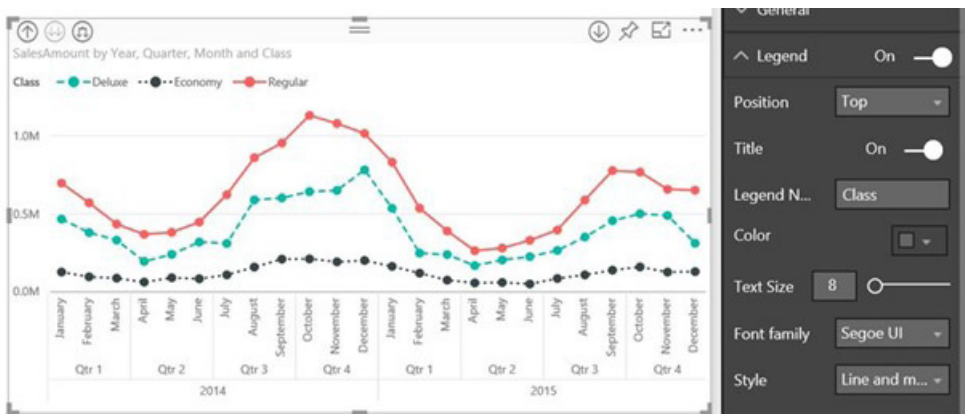
If you pick the 'marker only' option for the legend, you'll also be able to pick whether the legend colours match the line or the marker. Further, you may select whether a line or circle is used for the default icon when no marker is applied to a specific line.
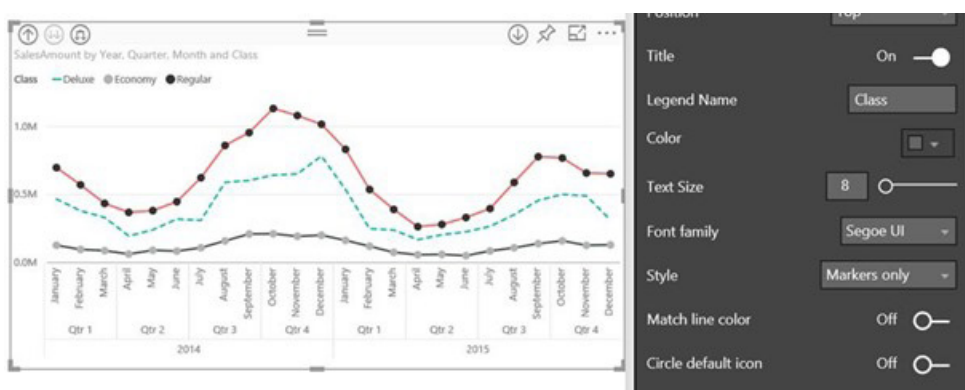
Scatter chart performance improvements

There have been changes made regarding how the scatter chart is rendered. For regular users, you should see up to a 30% performance improvement to a scatter chart's load time. If you're showing a large number of data points in your scatter charts, that's quite a significant performance improvement.

New custom visuals: Dot Plot by MAQ Software

The dot plot by MAQ Software visual plots individual data points on a simple scale, and is often considered one of the most useful chart types by many data visualisation experts, including Steven Few. *(Would it be fair to say then than Few like it? – Ed.)*

This visual is useful in many scenarios where you want to see the distributions of individual values, such as student scores, customer revenues, employee salaries, city population, movie ratings etc. This custom visual also allows you to view your data across several different categories, enabling effective comparisons.

For example, the dot plot below shows annual sales data by state (each dot) and its region (the x-axis). The legend is used to quickly tell which states perform better (blue) and which ones need attention (red) as well as the size of the dot, which shows profit margin:
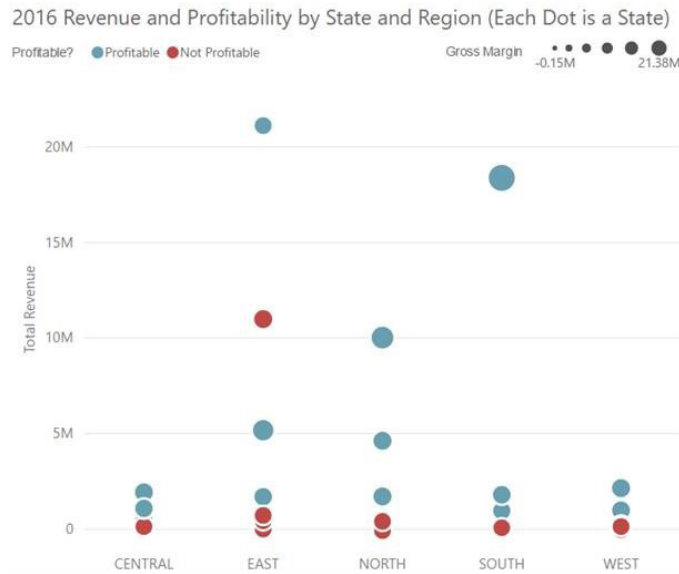
You may have multiple categories on the x-axis, allowing you to explore the data split by an additional category. In our example, I also want to explore the seasonality of the data, do I have added the quarter field to the chart, *viz.*
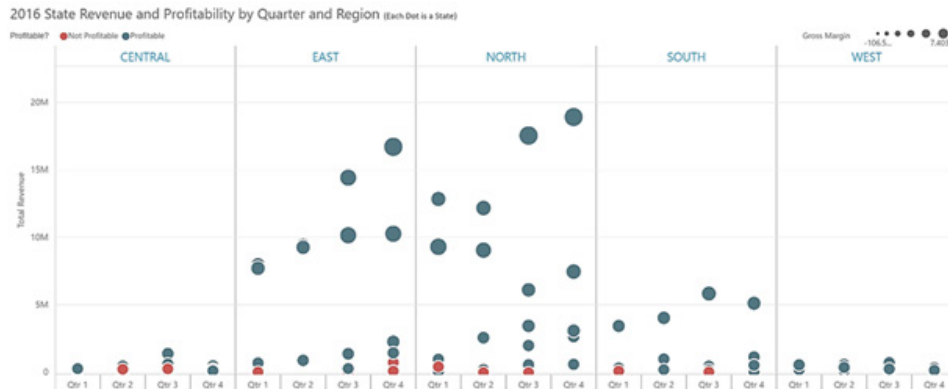
This visual also has several features in the formatting pane, including chart orientations, label controls, background shading, gridlines thickness / colours, and bubble colours / thickness.

You can test out this visual by downloading it from the custom visual store.
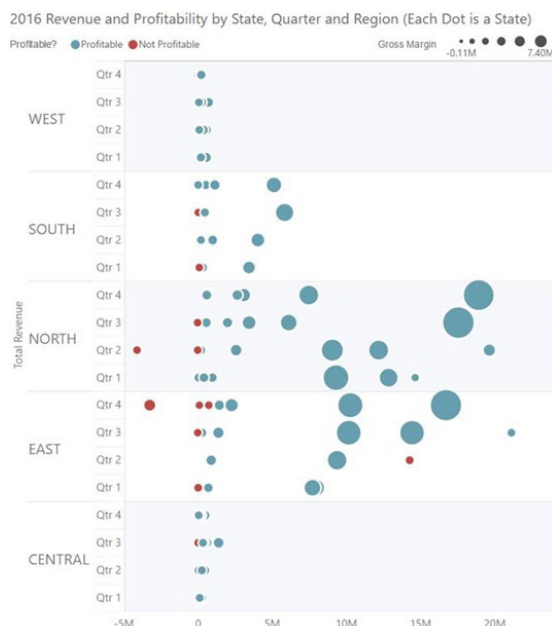
New custom visuals: Power KPI

The Power KPI visual is useful for presenting Key Performance Indicators (KPIs). The visual can include:

- title and subtitle
- KPI status and date
- a line chart with the KPI trend, forecast line, and other comparison lines.

This visual has several formatting options, such as options for KPI symbols and control for the line chart in the KPI, similar to the standard line chart visual in Power BI.

This is another chart you may download from the custom visual store.
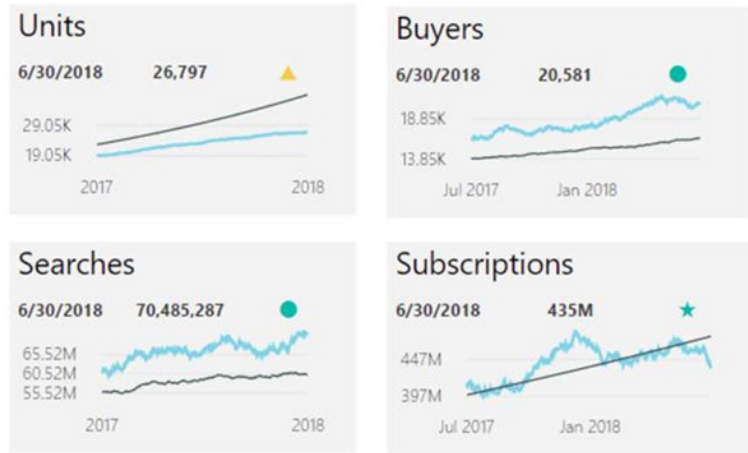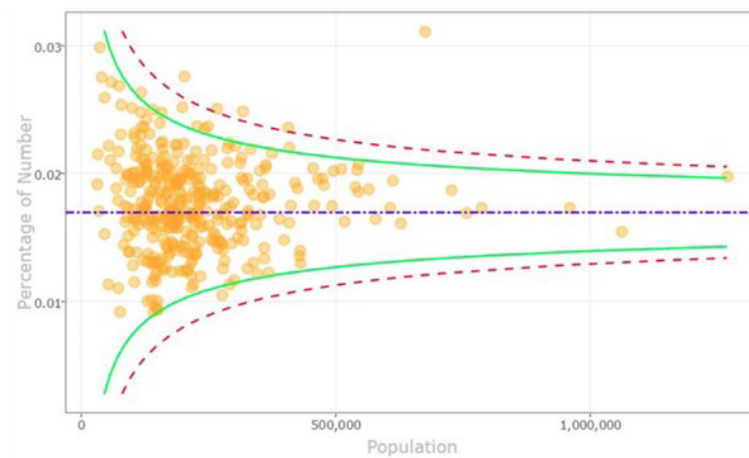
New custom visuals: Funnel plot

The funnel plot is an **R** custom visual that helps you not get distracted by statistical noise. It helps you compare samples and find true outliers among the measurements with varying precision.

To use this visual you start by defining two required fields to be analysed in the chart, such as occurrences and population. The y-axis can show either a percentage or ratio, and a "funnel" is formed by confidence limits which depict the amount of expected variation. The dots outside the funnel are therefore outliers.

Another one to consider from the custom visual store!

New custom visuals: Beyondsoft Calendar

Now this one we really like. We have had to build things like this many times in Excel for our clients. The *Beyondsoft* Calendar visual gives you a heatmap-like view of a value across all the days of a month. You can give the visual a date field and the numerical value you want to visual. From there, you can format the colours used and the layout of the calendar. This is very helpful if you have a measure you are tracking across time.

Already downloaded from the custom visual store!
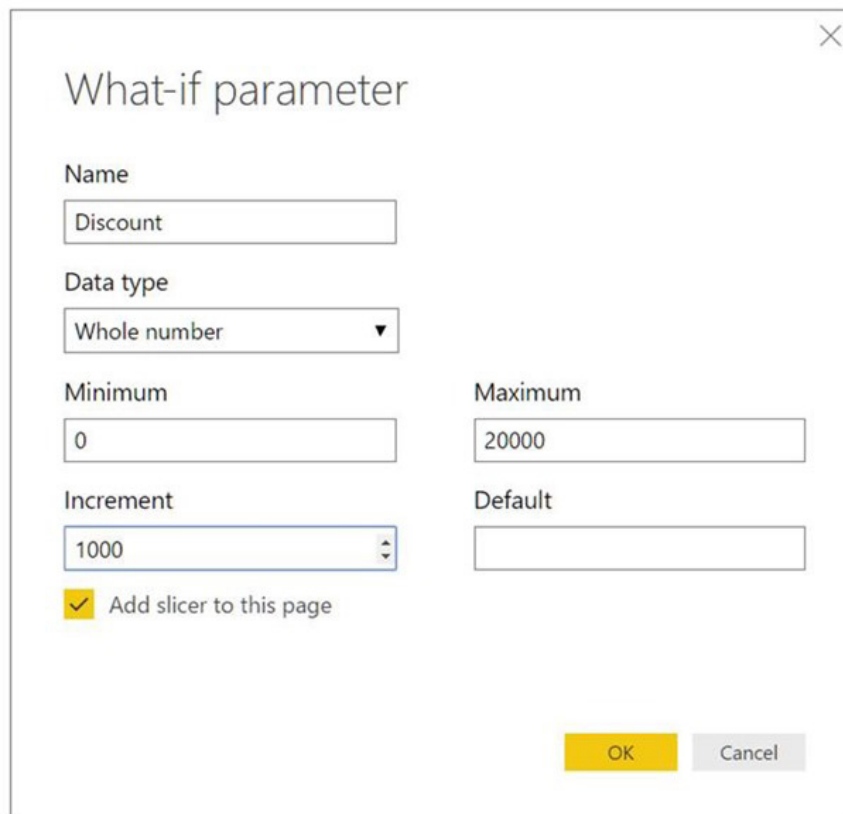
What if parameters

It's no secret that most users would love to interact with controls on a report to change a variable in their data, so-called sensitivity analysis. To demonstrate the new feature, Microsoft has created a demo scenario in the guise of a mortgage calculator report, where users would like to change the lending rate, borrowed amount and debt term in order to see the projected repayments.

To allow users to create these type of reports, Microsoft has added a new section to the 'Modeling' tab, called 'What If' *(sic)*. Using this new feature, users may create new parameters that you be used in your DAX expressions.
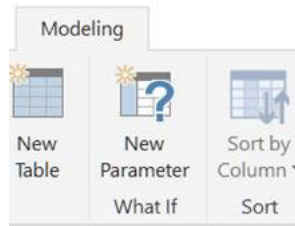
When you select 'New Parameter', you will open a dialog where you may fill in all of the details about the parameter, such as:

- Name
- Data type
- Minimum and maximum value
- Increment to increase by
- Default value

Once you finish filling out the dialog box and click 'OK', the information will be used to generate a new table with two fields, a little like a Data Table in Excel. The first field is a column with all the possible values for the parameter. It's very useful for putting in a slider for viewers to use when using the report. The second field is a measure that will also be equal to the currently selected value of the parameter column. This is what you'll want to reference in other calculations.
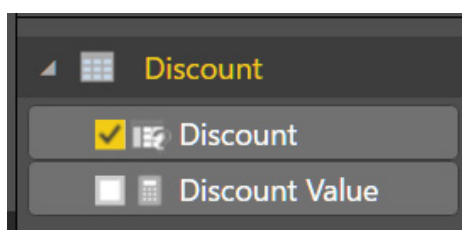
Once you use the parameter measure in your calculations, you can use those measures in your charts and have them update as you move the slider with your parameter column:
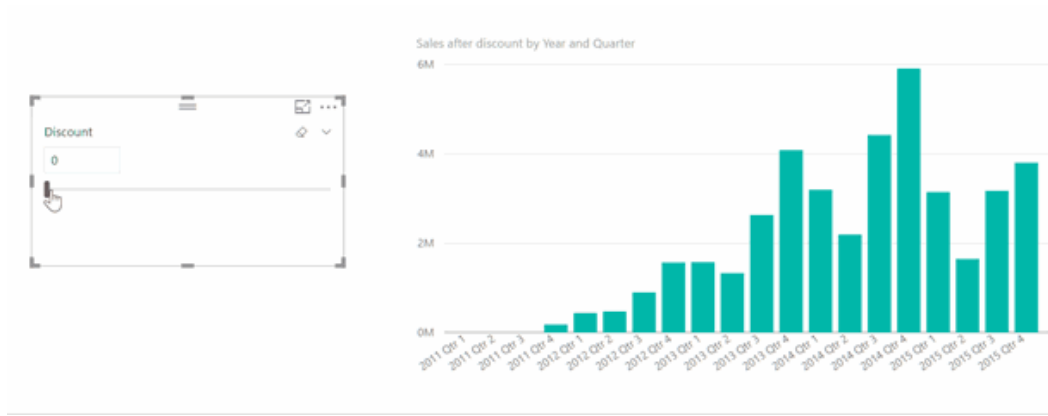
New scatter chart analytics features: symmetry shading

For scatter charts, Microsoft has added two new options in the Analytics pane. The first, symmetry shading, is a new analytics feature that allows you to easily see which points have a higher value of the x-axis measure compared to the y-axis measure and vice versa. When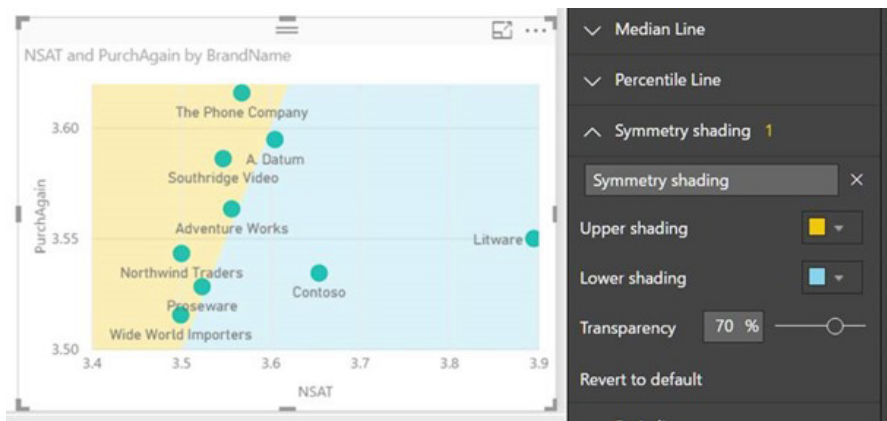 you turn symmetry shading on in the Analytics pane, Power BI Desktop will shade the background of the scatter chart symmetrically based on the current axis bounds. This is a very quick way to identify which measure a data point favours, especially when you have a different axis range for the x- and y-axes.

It doesn't half look like a coastline though: no Wonder World Wide Importers have a base there...

New scatter chart analytics features: ratio line

The second option, the ratio line feature, plots the ratio of the subtotals of x-axis and y-axis measures to let you see overall which measure has a higher value. When comparing data points to this ratio line, it allows you to see how much any given point deviates from the overall picture.

Points to the left of the ratio line favour the y-axis measure more than overall, whereas points to the right favour the x-axis more than overall.

Using these two features together may provide some very interesting insights, such as in the example (above). You can see from the shading that A. Datum in general has a higher net satisfaction score than its likelihood to be purchased again. However, when you compare it to the ratio line, you can see that it doesn't favour NSAT compared to overall.

New quick measure: weighted average

Another quick measure has been added this month: weighted average per category. This measure calculates a weighted average of the base value for each category. This measure will multiply the value by the weight for each category, sum the total, and then divide by the sum of the total of the weight. You'll find it under the aggregate by category section in the quick measures dialog:

Live connect to the Power BI service is generally available

The Power BI service connector is now "generally available". Now all users may start using this feature to connect to any of their datasets in the Power BI service. This makes it very easy to create a model once and reuse it in many reports.
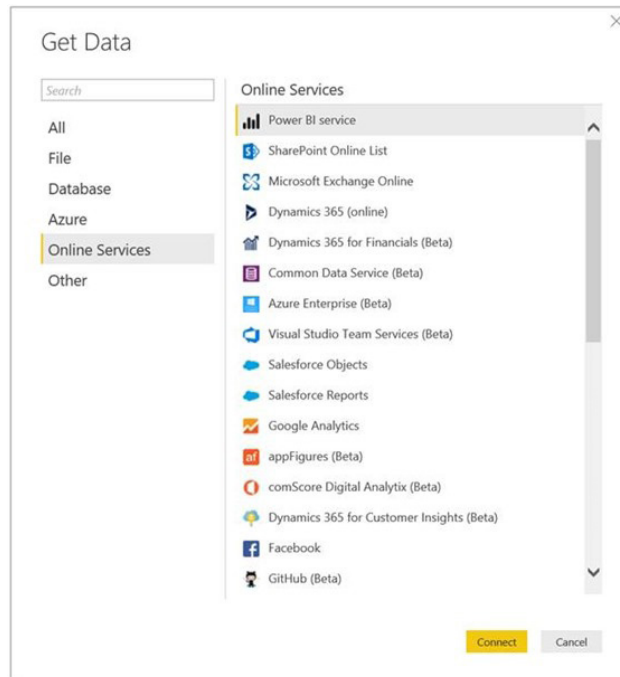
Along with making the connector generally available, there's two improvements this month as well:

- In addition to giving you a status indication in the bottom right section of Power BI Desktop, Power BI Desktop will also let you know the dataset and workspace you are currently connected to:

Connected live to the Power BI dataset: Power BI Desktop Demo in TexMex
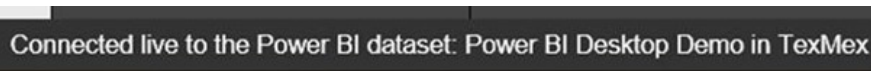
- The software will also let you update what dataset you are connecting to. This is useful when you want to use a staging or development environment when building your reports and then switch it to point to the production dataset right before publishing and sharing. This way, users may switch out what dataset they are connecting to using the 'Data source settings' option under 'Edit Queries' in the 'Home' tab on the Ribbon.
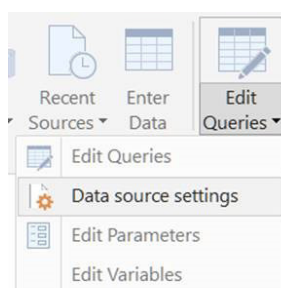
Google BigQuery connector (beta)

Finally, for August, there's a new beta connector for Google BigQuery. This connector allows you to easily create reports on top of Google BigQuery databases, either by using Import or DirectQuery mode. The new Google BigQuery connector can be found under the Database category within the Get Data dialog:

With this connector, you can select data from the bigquery-public-data project and any other project associated with your Google BigQuery account.

More next month no doubt!

August 2017 Updates for Get & Transform / Power Query Add-In

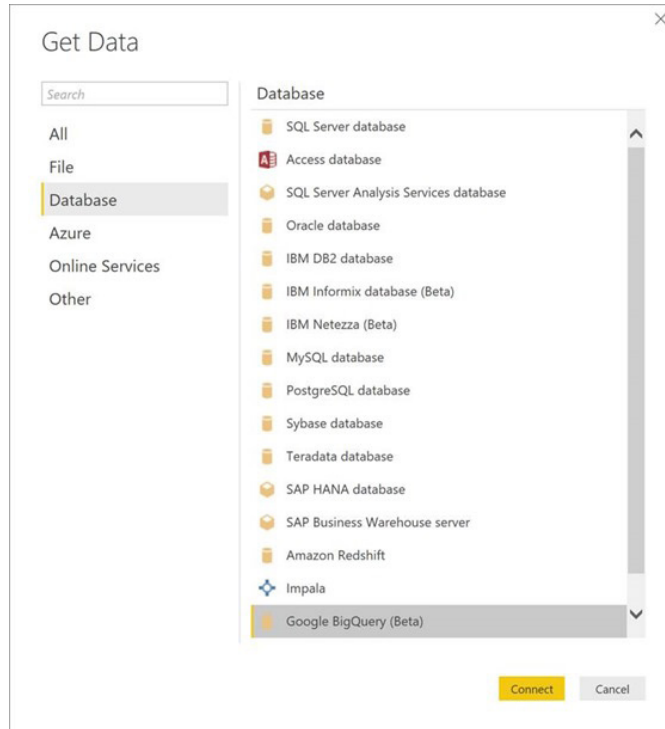We are suspecting these are beginning to taper off now, as the updates are becoming progressively more difficult to find and they are not being summarised in such detail (if the laws of grammar allow you to do such a thing!).

Anyway, Microsoft has just released another set of transformation updates for Get & Transform. The pick of this month's bunch has to be **Suggestions**, which is the latest enhancement to the recently released 'Add Column From Examples' feature, which can be accessed via the 'Add Column' tab on the Query Editor Ribbon:
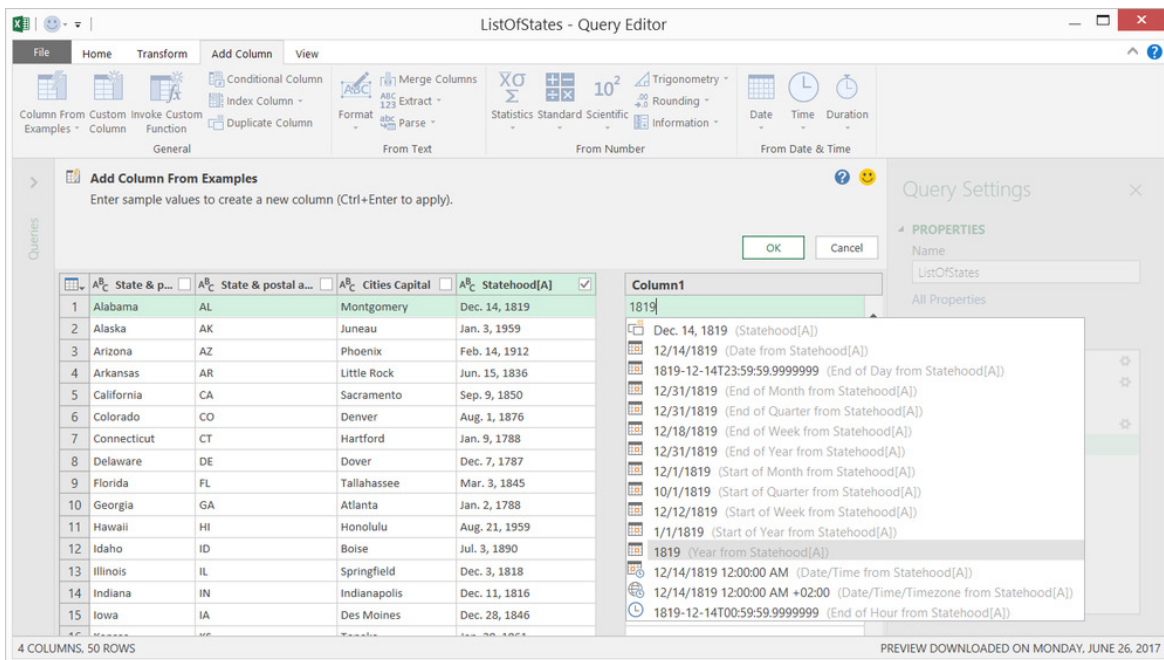
It's really quite useful for the inexperienced. As users go into 'Add Column From Examples', the software will show suggestions for output values based on the contents of other columns and partial contents typed by the user. In particular, if readers read through the suggestions, they will easily be able to distinguish between multiple transformations where the output value might be the same for a given row – and improve their skills at the same time!

There's also a raft of new transformations, including Number, Date, Time and Timezone operations. If you are sitting comfortably, here's the list *(deep breath!)*:

- Reference to a specific column, including trim, clean and case transformations
- Text transformations:
  - Combine (supports combination of literal strings and entire column values)
  - Replace
  - Length
  - Extract:
    - First Characters
    - Last Characters
    - Range
    - Text before Delimiter
    - Text after Delimiter
    - Text between Delimiters
  - Length
- Date transformations:
  - Day
  - Day of Week
  - Day of Week Name
  - Day of Year
  - Month
  - Month Name
  - Quarter of Year
  - Week of Month
  - Week of Year
  - Year
  - Age
  - Start of Year
  - End of Year
  - Start of Month
  - End of Month
  - Start of Quarter
  - Days in Month
  - End of Quarter
  - Start of Week

- o End of Week
- o Day of Month
- o Start of Day
- o End of Day
- Time transformations:
  - o Hour
  - o Minute
  - o Second
- Date / Time / Timezone transformations:
  - o To Local Time
- Number transformations:
  - o Absolute Value
  - o Arccosine
  - o Arcsine
  - o Arctangent
  - o Convert to Number
  - o Cosine
  - o Cube
  - o Divide
  - o Exponent
  - o Factorial
  - o Integer Divide
  - o Is Even
  - o Is Odd
  - o Ln
  - o Base-10 Logarithm
  - o Modulo
  - o Multiply
  - o Round Down
  - o Round Up
  - o Sign
  - o Sin
  - o Square Root
  - o Square
  - o Subtract
  - o Sum
  - o Tangent

These updates are available as part of an Office 365 subscription. If you are an Office 365 subscriber, you can find out here how to get these latest updates here. If you have Excel 2010 or Excel 2013, you can also take advantage of these updates by downloading the latest Power Query for Excel add-in.

STOP PRESS: Load Data from External Data Sources to the Common Data Service using Power Query

Just before our printing deadline, Microsoft has announced the "initial Preview availability" of a new feature in Power Query that allows users to create custom entities based on data from external data sources.

The *Common Data Service* allows users to securely store and manage data within a set of standard and custom entities. That's great if you know what an "entity" is. To explain, an entity is a set of fields used to store data similarly to a table within a database. Further, entities within the Common Data Service, both standard and custom, allow a secure and cloud-based storage option for your data. These allow you to create a business-focused definition of your data for use within your apps. The benefits include:

- **Easy to manage:** Both the metadata and data are stored in the Cloud
- **Easy to share:** You can easily share data with your colleagues because Microsoft's PowerApps manage the permissions
- **Easy to secure:** Data is securely stored so that users can see it only if you grant them access. Role-based security allows you to control access to entities for different users within your organisation *(whoops, nearly said "entity"!)*
- **Rich metadata:** Data types and relationships are leveraged directly within PowerApps. For example, defining a field type 'URL' will present your data as a hyperlink within your application
- **Productivity tools:** Entities are available within the add-ins for Microsoft Excel and Outlook to increase productivity, and ensure that your data is accessible
- **Picklists:** Include picklists from a set of standard picklists to provide quick drop downs within your entities and apps.

Business Apps Developers using Microsoft's PowerApps can now use the Power Query feature to create new entities and then load data into the Common Data Service based on data from a wide range of external data sources, including on-premises data sources such as relational databases (SQL Server, IBM DB2, etc.), Excel, Access or Text files, Online Services such as Salesforce, Azure SQL Database and Data Warehouse, Web APIs, OData feeds and many more sources. In addition to a wide range of data sources, Power Query also provides market-leading capabilities to filter, transform and combine data before loading it to the Common Data Service.
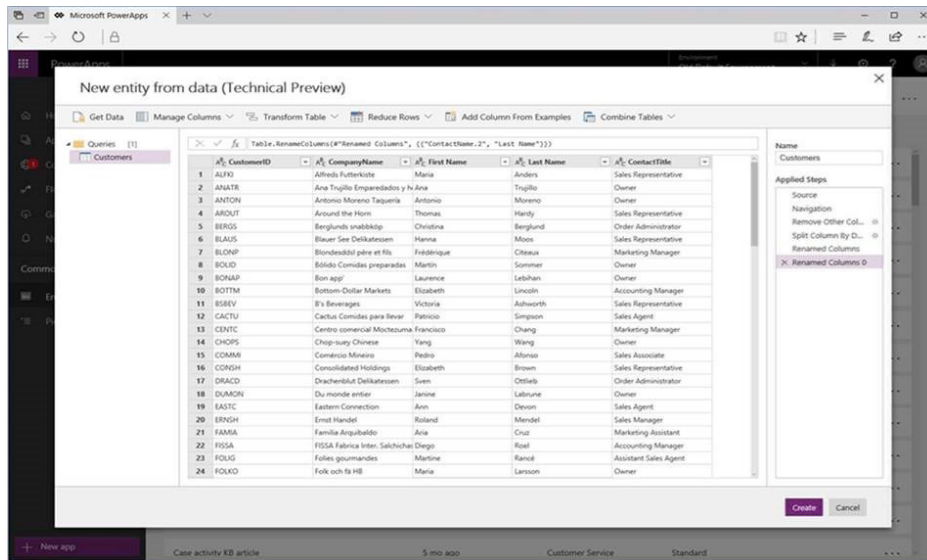
Once in the PowerApps portal, you can browse to the '**Common Data Service -> Entities**' tab in the left navigation pane. From the Entities list, you will see a 'New Entity' dropdown menu, under which you will find the 'New Entity from Data (Technical Preview)' entry:
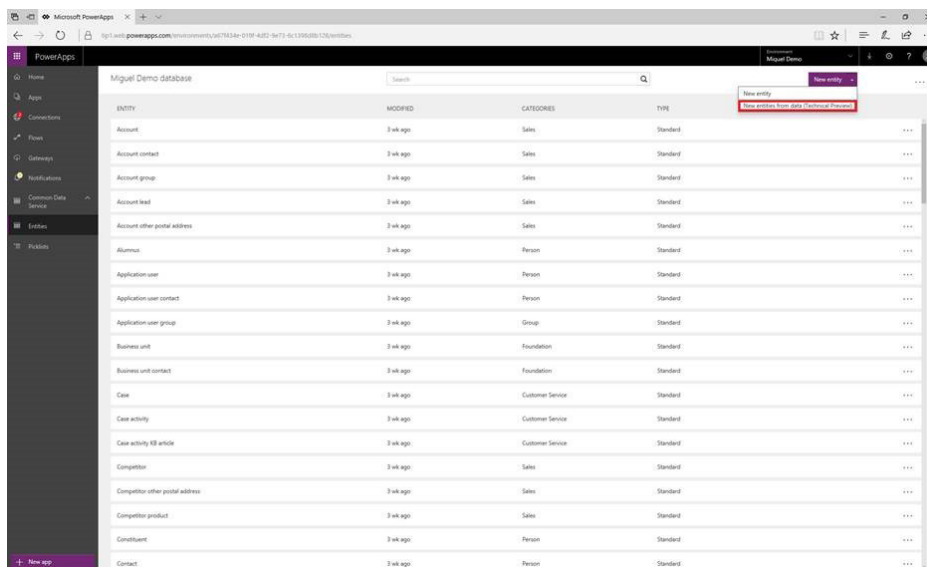
After clicking this button, you will see a dialog displaying the list of available connectors in this initial integration. Do note though that you will need to trigger this flow from a database that you created in order to be able to create new custom entities.
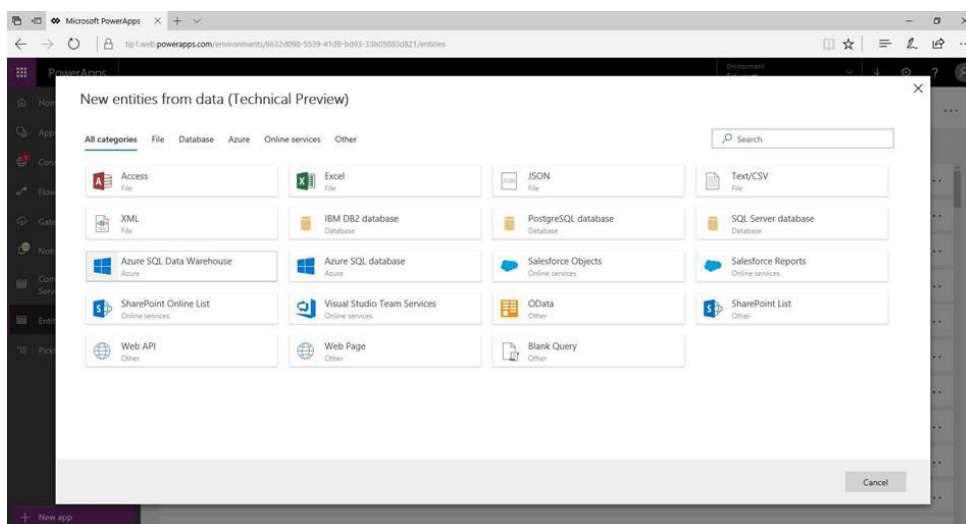
Once you have selected the desired connector to use, you may continue to the next steps to specify data source connection details and credentials, select which tables to import, etc. You may also access the 'Query Editor' (via the 'Edit' button in the 'Navigator' view) in order to apply filters or data transformation steps before importing data into CDS.
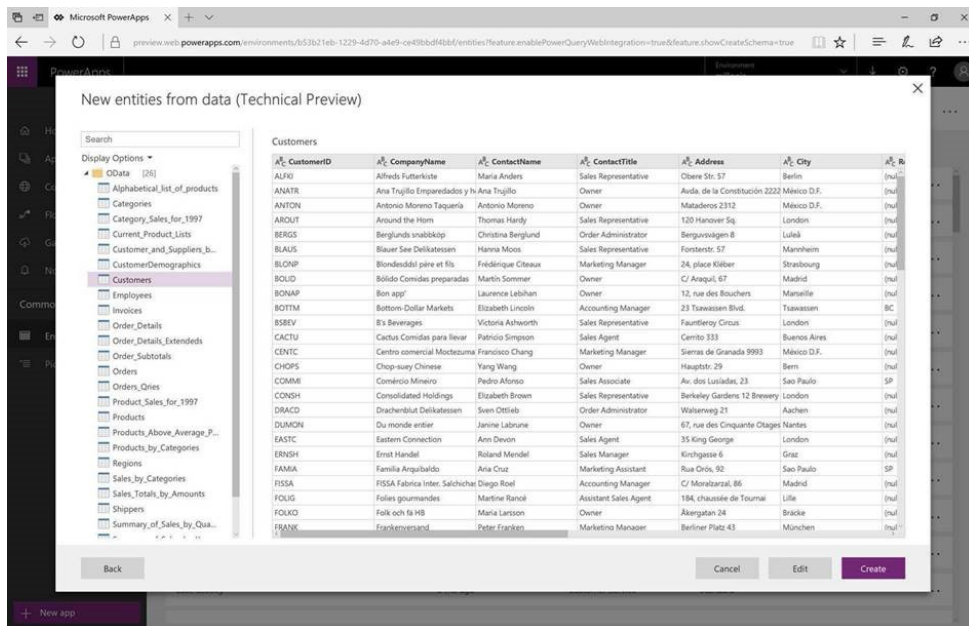
Once you are done preparing your data for loading into the Common Data Service (CDS), you will be able to adjust additional 'Load Settings', including 'Refresh' behaviour and entity-specific settings (*e.g.* Display Name, Primary Keys).

After completing these steps, you will have created new custom entities in CDS. Finally, do also note that you can 'Edit Queries' after the initial load by accessing the 'Entity' menu for a specific entity.

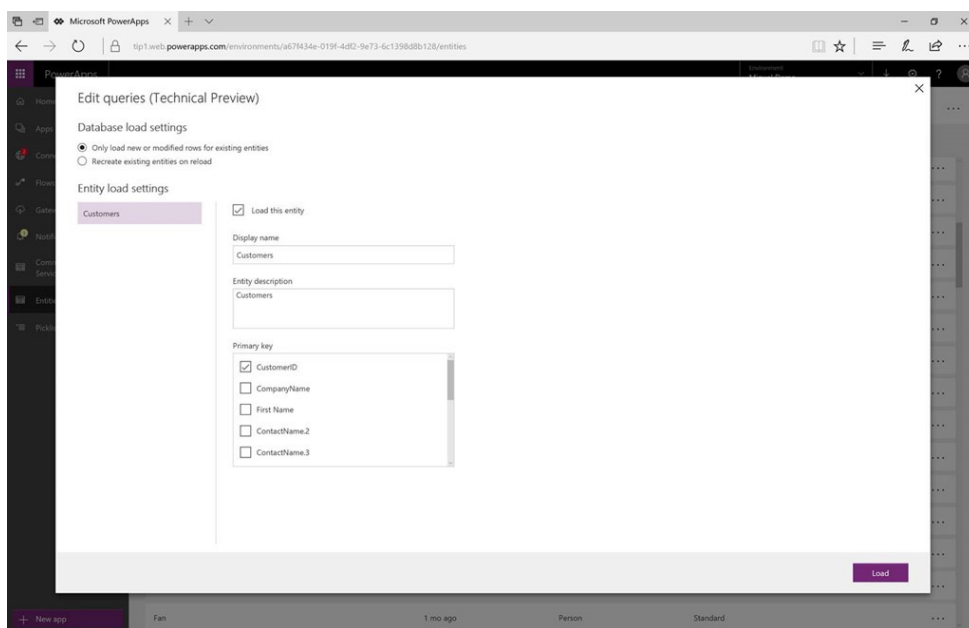Going forward, Microsoft plan to continue adding new capabilities (data connectors, data transformations, etc.) incrementally based on feedback and usage. If you are so inclined, have a play!

The A to Z of Excel Functions: CALL

Just the one this month, sorry about that. But look, before we go *anywhere* here, this function is for **advanced users only**. If you use the **CALL** function incorrectly, you might be looking for a new computer – well at best, you'll have to restart your PC. This function is only available from an Excel macro sheet in any case. Be careful with this function!

This function calls a procedure in a dynamic link library or code resource. There are two syntax forms of this function. Use the first syntax only with a previously registered code resource, which uses arguments from the REGISTER function. Use the second syntax to simultaneously **register** and call a code resource.

So, from above, the **CALL** function employs two different types of syntax operate.

Syntax 1:

Used with **REGISTER**:

CALL(register_id, [argument1],...)

where **register_id** is the value returned by a previously executed **REGISTER** or **REGISTER.ID** function.

Syntax 2:

Used alone (in Microsoft Excel for Windows):

CALL(module_text, procedure, type_text, [argument1],...])

The **CALL** function has the following arguments:

- **module_text:** this is required and is represented as quoted text specifying the name of the dynamic link library (DLL) that contains the procedure in Microsoft Excel for Windows
- **procedure:** also required. This is text specifying the name of the function in the DLL in Microsoft Excel for Windows. You can also use the ordinal value of the function from the EXPORTS statement in the module-definition file (.DEF). The ordinal value must not be in the form of text
- **type_text:** again, required. This is text specifying the data type of the return value and the data types of all arguments to the DLL or code resource. The first letter of **type_text** specifies the return value. The codes you use for **type_text** are described in detail in Using the CALL and REGISTER functions. For stand-alone DLLs or code resources (XLLs), you can omit this argument
- **argument1,...:** optional. These are the arguments to be passed to the procedure.

The following macro formula registers the *GetTickCount* function from 32-bit Microsoft Windows. *GetTickCount* returns the number of milliseconds that have elapsed since Microsoft Windows was started.

```
REGISTER("Kernel32", "GetTickCount", "J")
```

Assuming that this **REGISTER** function is in cell A5, after your macro registers GetTickCount, you can use the **CALL** function to return the number of milliseconds that have elapsed:

```
CALL(A5)
```

Maybe not for the faint hearted.

Upcoming SumProduct Training Courses

Location	Course	Date	Duration
Melbourne	Excel Tips & Tricks	4th Sep 2017	1 day
Melbourne	Financial Modelling	5th - 6th Sep 2017	2 days
Melbourne	Power Pivot, Power Query and Power BI	12th - 14th Sep 2017	3 days
Sydney	Excel Tips & Tricks	25th Sep 2017	1 day
Sydney	Financial Modelling	26th - 27th Sep 2017	2 days
Sydney	Power Pivot, Power Query and Power BI	16th - 18th Oct 2017	3 days
Melbourne	Excel Tips & Tricks	23rd Oct 2017	1 day

Melbourne	Financial Modelling	24th - 25th Oct 2017	2 days
Sydney	Visual Basic for Applications	6th Nov 2017	1 day
Sydney	Excel Tips & Tricks	13th Nov 2017	1 day
Sydney	Financial Modelling	14th - 15th Nov 2017	2 days
Melbourne	Power Pivot, Power Query and Power BI	4th - 6th Dec 2017	3 days
Melbourne	Excel Tips & Tricks	11th Dec 2017	1 day
Melbourne	Financial Modelling	12th - 13th Dec 2017	2 days

Referral Programme

Have you heard about the SumProduct Referral programme? If you successfully introduce us to a client we will pay you 10% of the net profits of the engagement won upon receipt of client payment. Sounds like a pretty good deal to us!

Terms and conditions apply. Referrer must be over 18 years of age, able to issue an invoice and not be an employee, director, agent or other related party of the client engaged. For further information, drop us a line at contact@sumproduct.com. We hope to hear from you!

Key Strokes

Each newsletter, we'd like to introduce you to useful keystrokes you may or may not be aware of. This month we thought we'd start on the numbers – it's as easy as (0,) 1, 2, 3:

Keystroke	What it does
CTRL + 0	Hide column
CTRL + SHIFT + 0	Show column
CTRL + 1	Format cells
CTRL + SHIFT + 1	Fixed decimal and comma format
CTRL + 2	Bold (toggle)
CTRL + SHIFT + 2	Time (AM / PM) format
CTRL + 3	Italic (toggle)
CTRL + SHIFT + 3	Date format

There are over 540 keyboard shortcuts in Excel. For a comprehensive list, please download our Excel file a www.sumproduct.com/thought/keyboard-shortcuts. Also, check out our new daily **Excel Tip of the Day** feature on the www.sumproduct.com homepage.

Our Services

We have undertaken a vast array of assignments over the years, including:

· **Business planning**
· **Building three-way integrated financial statement projections**
· **Independent expert reviews**
· **Key driver analysis**
· **Model reviews / audits for internal and external purposes**
· **M&A work**
· **Model scoping**
· **Power BI, Power Query & Power Pivot**
· **Project finance**
· **Real options analysis**
· **Refinancing / restructuring**
· **Strategic modelling**
· **Valuations**
· **Working capital management**

If you require modelling assistance of any kind, please do not hesitate to contact us at contact@sumproduct.com.

Link to Others

These newsletters are not intended to be closely guarded secrets. Please feel free to forward this newsletter to anyone you think might be interested in converting to "the SumProduct way".

If you have received a forwarded newsletter and would like to receive future editions automatically, please subscribe by completing our newsletter registration process found at the foot of any www.sumproduct.com web page.

Any Questions?

If you have any tips, comments or queries for future newsletters, we'd be delighted to hear from you. Please drop us a line at newsletter@sumproduct.com.

Training

SumProduct offers a wide range of training courses, aimed at finance professionals and budding Excel experts. Courses include Excel Tricks & Tips, Financial Modelling 101, Introduction to Forecasting and M&A Modelling.

Check out our more popular courses in our training brochure:

Drop us a line at training@sumproduct.com for a copy of the brochure or download it directly from http://www.sumproduct.com/training.