www.f1f9.com

Technical notes **on escalation**

At one level, financial modelling boils down to just two challenges: first, placing values correctly in time, and secondly, adjusting these values correctly for 'the passage of time'.

Monetary inflation is the most common, but not the sole, driver of such 'passage of time' adjustments. For example, a buyer and seller of gas may agree that the gas price increase by a certain, contractually agreed amount independent of developments in inflation. For this reason, we use the more generic term 'escalation' to describe the process of adjusting the value of a modelling parameter over time.

While values generally increase over time, it is of course also conceivable that certain values will decrease over time. One such example would be a devaluing currency exchange rate. Thus and somewhat confusingly, 'escalation' modelling also covers 'de-escalation' scenarios.

Experience will prove that escalation is an area of modelling where it pays to be particularly careful for two reasons.

First, escalation is inevitably applied to calculations throughout a financial model, many of them of substantial materiality. As one example of possible error, incorrectly lagging (or leading) the calculations by a single model period, even with a low escalation rate, can often affect the results more than is acceptable.

Second, escalation involves various issues that arise in translating the continuous reality of time to the discrete 'points in time' approximation required in the column-by-column structure of spreadsheets. Being aware of these issues, and then making appropriate and explicit (read: intentional) approximations, is vital to good modelling.

Hence, it's essential that we first develop a good understanding of what is predicted to happen in reality. From this understanding we can develop various methods for approximating this reality within the financial model.

Gathering the facts

Understanding reality begins by developing the list of assumptions that will drive the escalation forecast, which generally result from answers to the following three questions:

1. What is the structural form of the escalation, 'stepped' or 'continuous'?
2. What is the expected rate of escalation from period to period?
3. What is the reference date, the so-called basis date, for which the value(s) to be escalated apply?

Escalation form

There are two general forms of escalation that are important to identify, as the modelling approximations that apply to each are different.

The first, most often referred to as 'stepped' escalation, applies to a value that has a **known** date (or dates) upon which it will escalate. The second, generally referred to as 'continuous' escalation, applies to values that are presumed to generally escalate from one modelling period to another, but the dates upon which these changes occur are either too frequent or too uncertain to be modelled reasonably as stepped[1]. For example, a gas price may be agreed contractually to increase, or 'step-up', twice per year, say on 1 Feb and 1 Aug, the first such escalation date in the forecast being, for example, 1 Aug 2000. Hence this circumstance would best be modelled as a 'stepped' escalation form.

Alternatively, the gas price in question could be a 'spot' price, whereby price movements are determined purely by market forces. While it may be reasonable to expect the spot gas price to increase over the long term, in this case it would be nonsense to project (and therefore model) the exact date (or dates) on which these price increases will occur. Consequently, the modelling approximation commonly made for such 'un-contracted' parameters is that they escalate smoothly, or continuously, from one modelling period to the next. In this circumstance, the gas price would be best modelled as a 'continuous' escalation form.

[1] While presented in an inflation context, the modelling principles discussed in this note apply generally to modelling a change in a value over time (e.g. production volumes, traffic forecasts, etc).

Escalation rules

Having determined the escalation form (and - if stepped - the date or dates upon which step changes occur), the next step is to determine the appropriate rate of escalation (i.e. the amount of increase or decrease per modelling period).

Here there are two issues to consider. First, does the escalation **rate** remain constant (e.g. 3% p.a. for all modelling periods) or does it vary from one modelling period to the next (e.g. 4% for 2000 and 2% for 2001)? Second, does the time period of the escalation rate (e.g. annual) match the time period of the model (e.g. semi-annual)? If not, we need to adjust the escalation rate to fit the modelling period. In making this adjustment with escalation (as distinct from interest rates), note that a quoted annual inflation of, say, 4.00% invariably means that a loaf of bread costing £1.00 today will cost **exactly** £1.04 in a year's time. In a semi-annual model, this is **not** properly modelled by simply dividing the annual rate by 2 (i.e. 2.00% inflation per period). Such an approach would yield a price of £1.02 after six months and then £1.02 * (1 + 2%) after a year, or £1.0404. While this variation may seem trivial, and in fact is modest for short periods of time with low inflation, the error compounds over time and can be particularly significant in higher inflation environments. In any event, it is important to be aware of what is precisely correct, and then to **consciously** approximate. Hence, in a semi-annual model, a 4.00% p.a. escalation rate becomes 'de-compounded' to 1.98% per modelling period[2].

Escalation base date

Finally, for any given value (or series of values), we need to ascertain the date at which these values apply[3].

For the stepped escalation form, this is usually fairly straightforward, as any value will remain constant over some stretch of time, and therefore what is most operative is to ensure the first forecast date upon which a step-up occurs is properly understood, e.g. in our contracted gas price example, 18.00 p/therm on 1 Feb 2000 first stepping up on 1 Aug 2000. For the continuous escalation form, the escalation base date needs to reflect that a given price may be effectively an average value over a period (say, in the spot gas price example, 19.45 p/therm for the 6 month period from 1 Jan 2000 to 30 Jun 2000). In this sense the price in reality is not as at a specific date, but over a given period. Therefore, as we will discuss further below, specifying a particular date, which is necessary for modelling purposes, needs a little more thought before using it directly as a modelling assumption.

[2] Derived, in this case, from the expression, rannual = (1 + rsemi) ^ 2 - 1 or rsemi = (1 + rannual) ^ (1 / 2) – 1.

[3] For costs, this is often referred to as the 'cost estimate date'.

www.f1f9.com

Approach to modelling

In the following, and before attempting to model the two escalation forms, we first draw a picture of what the respective gas price examples might look like in reality.

In order to make our examples as clear as possible, we focus on a two-year period only (1 Jul 1999 to 30 Jun 2001). From a modelling perspective, we will focus only on the two semi-annual periods ending 31 Dec 2000 and 30 Jun 2001.

STEPPED ESCALATION

DRAWING STEPPED REALITY

Figure 1 illustrates a stepped gas price profile based on an initial gas price of 18.00 p/therm as at 1 Feb 2000, which we will assume will escalate by 1% on 1 Aug 2000 and by 3% on 1 Feb 2001.

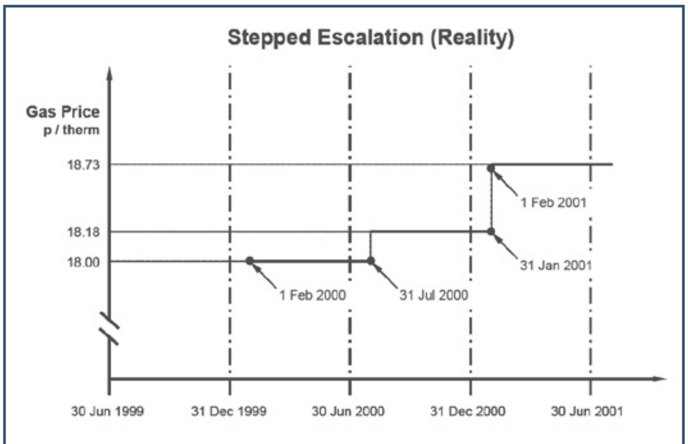

Figure 1
Figure 1 plots the 'reality' of the contracted gas price as being 18.00 p/therm from 1 Feb 2000 to 31 Jul 2000 (end of the day). On 1 Aug 2000 (beginning of the day) the price increases by 1% to 18.18 p/therm and remains at this level until 31 Jan 2001. On 1 Feb 2001, it increases by 3% to 18.73 p/therm.

MODELLING STEPPED ESCALATION

Bearing in mind the discrepancy between real time and modelling time, let us consider how in a semi-annual model we might best approximate the escalation scenario in Figure 1 above.

For the sake of simplicity, we assume that we have a two-period model only (i.e. two Excel columns), covering 'Period A' from 1 Jul 2000 to 31 Dec 2000 (Excel Column J) and 'Period B' from 1 Jan 2001 to 30 Jun 2001 (Excel Column K). In modelling terms, this means we have only two calculation opportunities, one for Column J and one for Column K. In other words, there can be only **one** model price (or number) for Column J and only **one** model number for Column K. Thus, the question becomes: What is the **best single number** to use in each of these Excel columns?

For stepped escalation, the approach commonly taken is to calculate the time-weighted average (in this example, the time-weighted average contract gas prices for Period A and Period B). For Period A in Figure 2 below, this equates to 1 month at 18.00 p/therm and 5 months at 18.18 p/therm, or 18.15 p/therm. Using the same approach, the modelling price for Period B is 18.64 p/therm (1 month at 18.18 p/therm and 5 months at 18.73 p/therm).

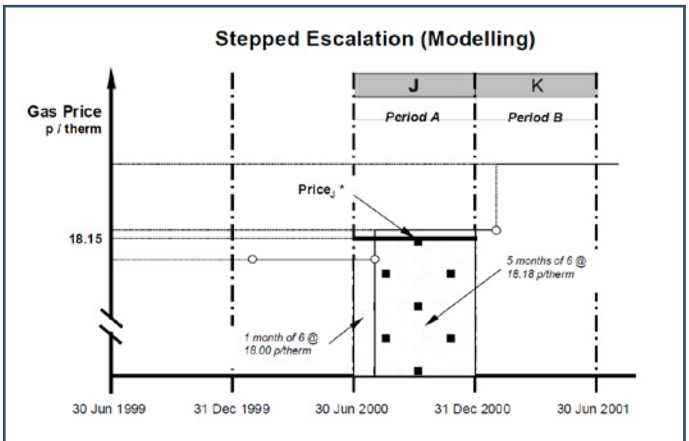

Figure 2

THE CONSTANT MULTIPLIER CAVEAT

How valid the above time-weighted average modelling approximation is depends on the broader modelling context. In our gas price example, let us assume that we will be multiplying the gas price by a gas volume (in therms) in order to arrive at a gas cost (in £s). **Only** if the rate of gas consumption is constant (say 100,000 therms per month) throughout Period A (giving a total of 600,000 therms for the modelling period) will our time-weighted average price approach produce the correct gas cost (i.e. 18.15 p/therm x 600,000 therms = £108,900). If, however, the Period A gas volume is **unevenly** distributed over the modelling period, the time-weighted average modelling approach will lead us to either over or underestimate the correct gas cost. Depending on the degree of bias in this distribution, this inaccuracy can be substantial (e.g. modelling retail toy sales, where 50% of annual sales may occur in month of December).

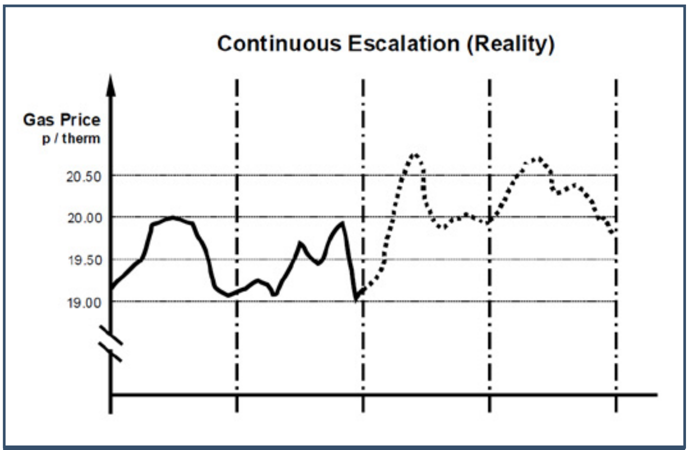

Figure 3

CODING IMPLEMENTATION

The exact coding implementation of the above modelling will vary by modeller and model. However, the preceding analysis of the problem provides a good general framework for a well-structured modelling approach. In this case, our analysis suggests that we might first model the pre and post-escalation date prices (in our example 18.00 and 18.18 p/therm for Column J, and 18.18 and 18.73 p/therm for Column K). Next, we would model the number of months (or more generally the fraction of the modelling period) applicable to each pre and post-escalation date price. Finally, we put the above two elements together to produce the time-weighted average price.

CONTINUOUS ESCALATION

DRAWING 'CONTINUOUS' REALITY

Drawing a picture of reality for our continuous escalation form example (spot gas price) is much more difficult. Being purely market force driven, the price is unpredictable and likely to show substantial volatility. Nevertheless and for the sake of argument, let us consider Figure 3 as one possible rendition of reality. (Drawing is generally a good place to start, even if the picture itself is guesswork.)

IMPOSING AN ESCALATION LINE

Now let us suppose it is 30 Jun 2000 and we have data (say, daily spot gas prices) available for the past six months. From this data, we calculate the average price for the six-month **period** ending 30 Jun 2000 to be 19.45 p/therm. From this point (30 Jun 2000) forward, trying to predict the daily price movements for the next year would be more than a little unrealistic (the dotted line 'projection' in Figure 3 is simply a **possible** reality, not a projection with any confidence). However, a forecaster might hazard a guess at future **average** price levels. In this spirit, let us suppose that the forecaster projects a 4% p.a. increase in the average price level over the next two semi-annual periods (i.e. 1.98% per period), averaging 19.83 p/therm and 20.23 p/ therm in the next two succeeding periods respectively. We can impose an 'escalation line' on our continuous escalation graphic by joining up the average price points as shown in Figure 4.

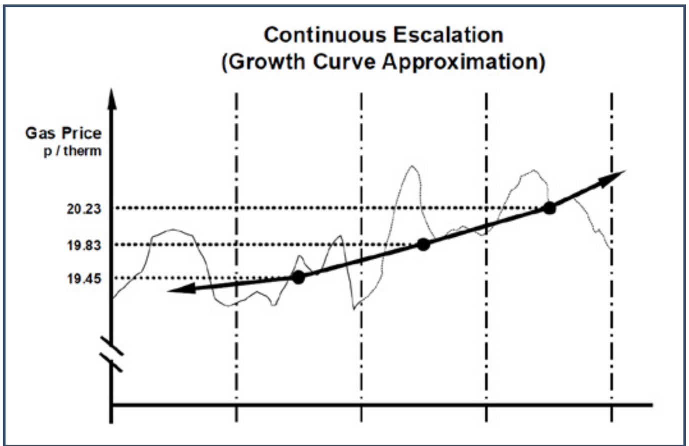

Figure 4

[3] For costs, this is often referred to as the 'cost estimate date'.

MODELLING CONTINUOUS ESCALATION

Once again, let us assume that we have a two-period model only (i.e. two Excel columns), covering 'Period A' from 1 Jul 2000 to 31 Dec 2000 (say Excel Column J) and 'Period B' from 1 Jan 2001 to 30 Jun 2001 (in Excel Column K). Once again, the issue boils down to determining the **best single number** to use in each of these Excel columns.

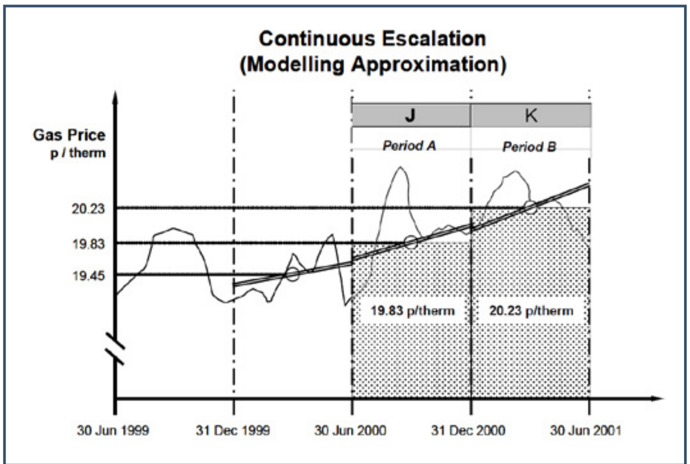

Figure 5

For continuous escalation, as with stepped escalation, the approach commonly taken requires the time-weighted average predicted for the period in question (in this example, the time-weighted average spot gas prices for Period A and Period B). The difference in approach, however, is that for continuous escalation we must estimate the time-weighted average of an artificially imposed 'escalation curve' and not, as is the case for stepped escalation, on reality.

For low escalation rates, a continuous escalation curve can be fairly closely approximated by a series of straight line segments (see double lines in Figure 5). In evaluating the fit of these line segments to an (exponential growth) curve it works out that the segments drawn through the value at the middle of the period will be a little below the true value at the start of the period and a bit above at the end of the period[4]. Hence the segments will not link up (this is not an unintentional glitch in the graph).

However, if one accepts this construction as a reasonable approximation of the curve, then the time-weighted average price for a period is equal to the average of the beginning and ending prices for the period. Hence, the final analytical step is to realise that, for the continuous escalation form, an even simpler way to calculate the time-weighted average price is to calculate the price as at the mid-period point (i.e.30 Sep 2000 for Period A and 31 Mar 2001 for Period B).

MID-PERIOD APPROXIMATION REVISITED

To elaborate on the mid-period approach, rather than being given an **average** price of 19.45 p/therm (for the six-month period 1 Jan 2000 through 30 Jun 2000), let us assume that we were given a price of 19.45 p/therm **as at**, say, 31 Jan 2000. Assuming a continuous 4.0% p.a. increase in the average price level, would we still model our 'best guess' as 19.83 p/therm for the period 1 Jul 2000 through 31 Dec 2000 (i.e. Period A)? In short, no. Intuitively, we would want a higher price (on average) for Period A given the earlier base date of 31 Jan 2000 (and hence a longer stretch of time over which escalation occurs to get to the mid-point of Period A). Applying the mid-period approach on this revised assumption results in a modelling price of 19.97 p/therm (i.e. five months from 31 Jan 2000 to 30 Jun 2000, and a further 3 months to the middle of Period A) for Period A[5].

[4] As the rate of escalation increases, e.g. for high inflation environments, the magnitude of this discrepancy increases.

[5] 19.45 p/therm * (1 + 4.0%) ^ (8 / 12) = 19.97 p/therm.

WHEN RATES CHANGE OVER TIME

The preceding example presumed that the prediction of a 4.0% escalation rate stayed constant over the forecasting period. But what if the economic forecast predicts an accelerating (or decelerating) rate of escalation, say 16.0% p.a. for the period 1 April 2000 to 30 Sep 2000 and then 8.0% p.a. for the next six-months? For Period A, we would then be predicting that the first three months had prices escalating at a 16.0% **annual** rate, and the last three months escalating at 8.0% (numbers deliberately picked large with substantial change between periods to unearth level of any variations that may arise). To apply the mid-period approximation most directly, we need to derive a rate for the period whose value at mid-period is closest to the time-weighted average price level for the period.

This suggests we should first be confident we know what we're aiming for. For these purposes we will presume a £1.00 price for a loaf of bread at 30 Jun 2000 and ignore the effects of compounding (i.e. in our earlier example the true 1.98% period rate still being close enough to ½ the annual rate, or 2.0%). It might be natural to leap to a position that 3 months at 16.0% and 3 months at 8% averages to 6 months at an annual rate of 12%, or approximately 6% for the period in question. This would predict a price in six months' time of approximately £1.06, and an average price, taken at the period's mid-point, of £1.03.

However, is this quite right? After 3 months at 16% inflation, the price will rise to £1.04. From **that price** at mid-period (and here's the rub), the price will rise a further (3/12) * 8.0% or about £0.02, hence the price at the end of Period A will rise to £1.06. The final price at least looks consistent with taking a simple time-weighted average of the rates operative for the given period. But what is the true average price level?

Well for the first half of Period A, the average price was £1.02, and for the second half £1.05, hence the average price for the period is £1.035. Hence we are ½ a pence off (or a 0.5% accuracy) on the time-weighted price level. So what should we do?

One temptation is simply to work out some means to get to the 'right answer' (in this case developing the arithmetic that would yield an effective 14% annual rate for the period, such that 3 months of inflation (to the middle of the six-month Period A) would yield an average price level of about £1.035 for the period[6].

Another option is to say 'hang it', and simply pick one of the rates to apply to the entire period, i.e. either 16% (getting a price of £1.04) or 8% (getting a price of £1.02)[7]. Given that escalation rates are often fairly modest, and that changes from period to period are rarely (if ever) forecast with drastic changes, this is not an unreasonable position. Particularly when put in context of the key presumption of 'constant volume' (see above), and recognising that the escalation forecast itself may be virtual guesswork, this may be quite reasonable.

Or finally, we can simply choose to calculate a time-weighted rate for the period, which at least adjusts for the duration of time each rate applies (e.g. if we had only 1 month in Period A at 16.0% p.a. and 5 months at 8.0%, we might be less happy with the preceding approximation of simply presuming 16.0% p.a. for the period).

Take your pick. The point of this discussion is not to provide 'an answer', but to reinforce good analytical methodology:

- First understand what is actually going on in reality (at least on the basis that assumptions prove correct)
- Then think carefully as to whether some approximation is warranted in the interests of improving modelling simplicity and in the context of other inaccuracies that may be embedded in the modelling methodology.

Your decision will be based on circumstances of the business being modelled, and an approach reasonable in one circumstance may be less valid in another.

[6] Left as a modelling challenge for anyone interested. Does it matter (much) whether escalation is coming down or going up over time, i.e. would we get about the same answer if we assumed 8.0% p.a. escalation rising to 16.0%?

[7] It appears to be the higher rate that gets closer to correct answer, but this of course in a situation where both rates were operative for the same fraction of time during Period A.

Model implementation specifics

ESCALATION FACTORS

To this point in the discussion, we have worked directly with average price levels, mostly in the interests of illustrating with tangible values. However, as a matter of good modelling practice (i.e. breaking things up), the exercise of escalation is best separated from the amounts being escalated into a series of non-dimensional values that can be multiplied by the 'base' value in question to derive the final answer. These intermediate calculations are commonly called escalation factors. At the most basic level, nothing more need be done than to replace the 'starting price' with 1.000 and use the techniques of averaging (for either stepped or continuous) discussed above. Presuming that the base values in question are relevant for the start of the forecast period, this is simple and fairly transparent (another key objective of good modelling practice).

WORKING WITH INDEX VALUES

However in practice, the most direct link with the 'real world' is the use of indices. These may be directly quoted (e.g. RPI in the UK) or constructed from a weighted average of various indices (e.g. a blend of commodity price and labour cost indices). The value in forecasting indices becomes clear when one considers that we have to this point simplified the modelling requirement by noting that the base value starts at the beginning of the forecast horizon. Hence we have simply to forecast from this point forward, i.e. the basis date was conveniently presumed to be at the start of our modelling exercise.

But as is commonly the case, what if this basis date is some ways prior to the start of our modelling? To handle this situation the best modelling solution starts from a recognition that two indices are required (because two dates are important). The first index is what one might term the 'last actual' index, i.e. the one most recently quoted and hence closest to the start point of the forecast horizon (e.g. a recent CPI quoted in the Financial Times). We would simply use our escalation techniques to forecast a stream of indices from this 'last actual' figure forward in time. However, in deriving an escalation factor, we then need to be conscious of the index at the basis date, e.g. a cleaning maintenance contract negotiated at £4.2 million per year at a point in time when the labour cost index (upon which this contract will be annually indexed) was 210.6, i.e. the 'basis date' index. We would then simply forecast from the 'last actual' index available and divide by the basis index to derive the relevant escalation factor, which we would then apply to the £4.2 million 'base' contract price to derive costs in any given period[8].

Different time resolutions (modelling 201 only)

In the above examples, we have presumed a semi-annual time resolution, which is typical of modelling operating periods where the repayment of debt and/or dividends are scheduled twice annually. Six months is a long enough time period where dealing with intra-period effects can be modestly important to accuracy. Hence calculating time-weighted average levels for stepped escalation and / or mid-period approximation for continuous escalation may be relevant. In models that contain higher resolution periods, say monthly resolution modelling of a construction period, matters can become quite a bit simpler. Even presuming step-ups may not occur on the 1st of a month (and they often do anyway), it would be rare to consider intra-month effects as material to the overall forecast. Hence deriving monthly escalation factors becomes simply a matter of forecasting an index for a given month (inevitably using a flag to note the months when a rate may change, otherwise taking the preceding value) and dividing back by the basis index. The only additional twist that can arise in such dual-period models is the objective to seamlessly pass this monthly forecast into the operating period forecast (presuming this period occurs chronologically prior to the lower resolution modelling). This can be achieved by various means. If the (semi-annual) operating period simply picks up from the end of construction, then one can simply 'pick' the index derived during the monthly period for the month that begins the (semi-annual) operating period, say the last day of construction or conversion date of the construction loans. This index then is passed into the semi-annual escalation factor calculations as what heretofore has been considered the 'last actual' index and all proceeds as described above.

[8] In practice, modelling contract escalation is often further complicated by a need to lag the indices used for a given period by some months, as inevitably data is only available some months after the period in question. We will leave this to another day, but it is another area where you are advised to be careful.

