GUIDE TO ECONOMIC REGULATION

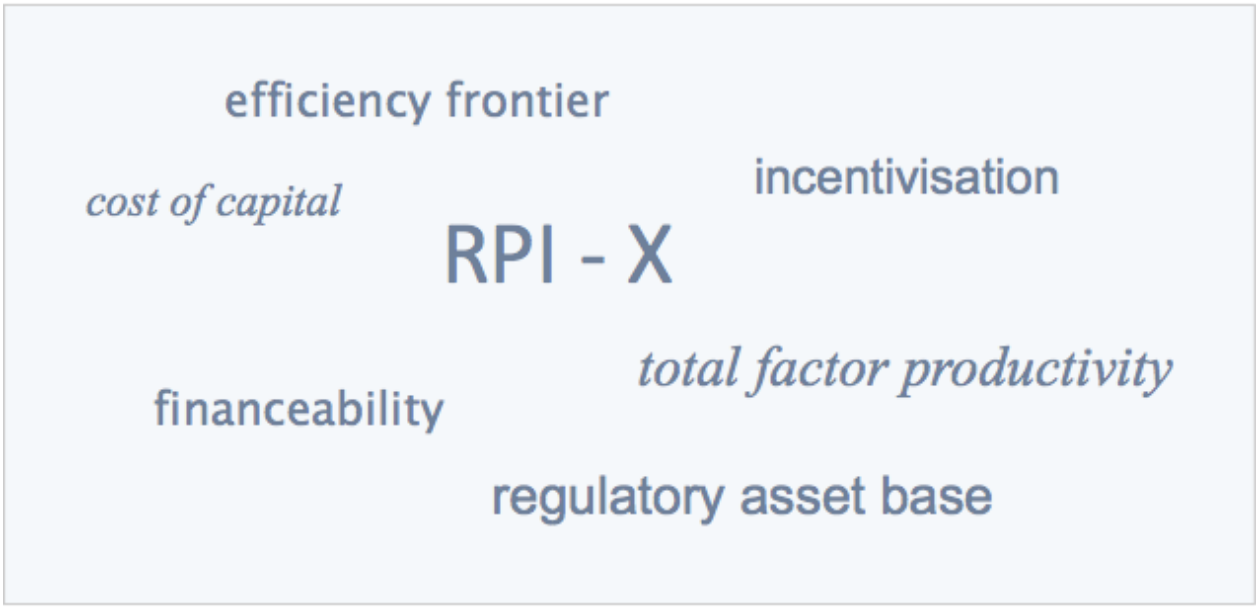

Part 5: Cost Assessment

John Earwaker

Foreword

This is Part 5 in a series of booklets which aim to provide individuals working in the regulated aviation, communications, energy, rail and water sectors with an introductory guide to the principles and practices of economic regulation.

We saw in Part 2 that regulators set regulated companies upfront allowances for operating expenditure (opex), capital expenditure (capex) and/or total expenditure (totex). We now work through some of the analysis that a regulator will assemble when sizing these £m amounts.

1. A Framework of Analysis

In order to help the reader navigate through what can often add up to multiple layers of challenge, encompassing multiple different pieces of work, we begin by defining four key terms: base costs; catch up; frontier shift; and growth and enhancement.

1.1 Base costs

"Base costs" can be thought of as the day-to-day, week-to-week, month-to-month expenditures that a regulated company incurs on a repeated basis. They will comprise most of a company's opex and, potentially, a proportion of a company's capex, wherever a firm has a regularised amount of ongoing investment activity (e.g. in the form of capital maintenance).

1.2 Catch up

The starting point in a typical cost assessment will be a challenge from the regulator to the level of efficiency that a company was exhibiting within its known base costs in its most recent complete financial year.

This starting level of costs is depicted with a cross in the chart opposite.

Figure 1

A regulator will look to determine whether it is fair and appropriate for customers to cover this starting run rate of expenditure in full. As we will see in section 2, this will typically entail the regulator benchmarking out-turn costs against comparable companies' expenditures and coming to a view about whether the firm could have spent less had it matched the cost control exhibited by the most efficient companies in the industry.

Figure 2, overleaf, depicts a situation in which a regulator finds that a company was capable of spending less than its actual recorded costs. We can label the gap between the company's actual expenditure and the efficiency frontier "catch up".

Figure 2

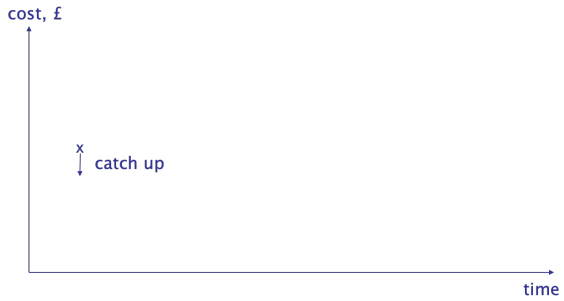

1.3 Frontier shift

The chart as drawn constitutes a snapshot efficiency assessment in a single year. The regulator's job during a price review is to size expenditure allowances prospectively for a n-year period covered by a brand new price control. This requires that the starting efficient level of base costs is rolled forward n years into the future.

Figure 3

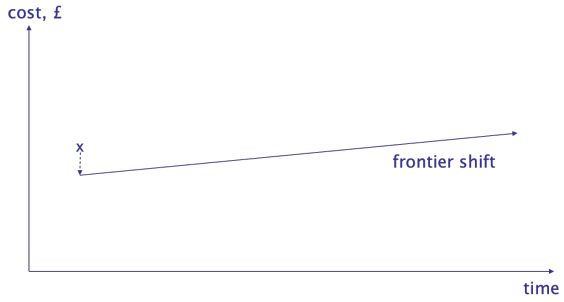

We term this roll forward of efficient costs "frontier shift". This labelling hints from the outset that the gradient of the sloped line in figure 3 should be set in accordance with the rate at which a hypothetically efficient frontier company's costs would change, rather than the actual increase or decrease in costs that the actual company might be looking at.

1.4 Growth and enhancements

Figure 3 shows the key inputs into the regulator's calculation of a suitable allowance for base costs in a stable, unchanging business.

Real-life companies are not so simple. Some regulated companies serve a growing number of customers each year, necessitating additional spending. And modern-day price review will often throw up questions around future service levels and the size of the future expenditure programme that is needed in order to meet changing customer requirements.

We must therefore also add new "growth" and "enhancement" expenditures to the preceding charts to capture the step changes that there might need to be in the scale of the firm's running costs and/or investments.

Figure 4

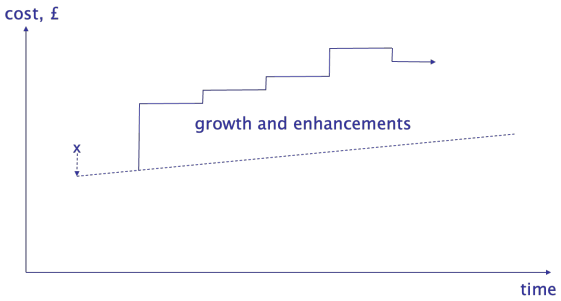

It is the final solid line in figure 4 – calibrated, as shown, by reference to the company's actual starting level of base costs, plus a catch-up efficiency challenge (where deemed appropriate), plus an allowance for frontier shift, plus allowances for growth and enhancements – that will ultimately then be input directly into a regulator's opex/capex/totex allowances.

Sections 2, 3 and 4 elaborate further on each of the steps in the calculation

2. Base Costs and Catch Up

The way in which a regulator will approach the initial sizing of a firm's base costs allowance tends to depend on the specific features of the industry that it is regulating.

In sections 2.1 and 2,2, we look first at how a regulator will ideally want to proceed, before discussing in section 2.3 what alternatives a regulator has if there are obstacles that make the first-best option practically impossible to implement.

2.1 Intra-industry benchmarking

The ideal scenario for a regulator will be when it can observe a large number of firms providing the same set of services in an industry simultaneously. This kind of set-up naturally presents an opportunity to conduct company-to-company benchmarking and for a regulator to judge the efficiency of each firm in the sector by reference to the contemporaneous performance of the firm's peers.

By way of an illustration, and to help fix ideas in the discussion that follows, figure 5 plots the out-turn base costs incurred by ten hypothetical companies from the same sector in a particular year. The chart is drawn to show each company's size on the horizontal axis and each company's base costs in £m on the vertical axis.

Figure 5

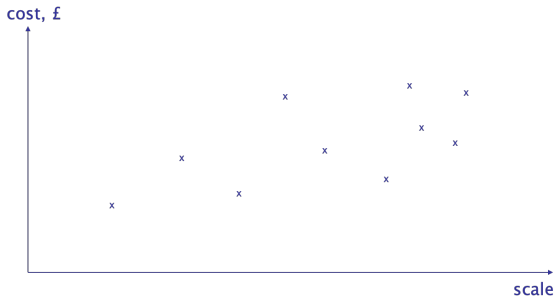

What can the regulator take from this data set? At first glance, there seems to be a discernible spread in companies' cost control, with some firms looking like higher spenders and other firms looking like they are spending less. But a simple visual inspection of the chart also suggests that there is some sort of relationship between cost and scale – i.e. smaller firms, quite naturally, seem to spend less and bigger firms seem to spend more. In order to obtain any useful information from this kind of data set, a regulator will therefore need to account for this size factor and identify only differences in expenditures that are caused by factors that are within companies' control.

Regulators can make allowance for exogenous factors by running regressions. A regression is a statistical technique that identifies an equation that best explains the way in which a dependent variable (in this

case, cost) is affected by one or more independent variables (in this case, scale).

Figure 6 plots the results of an ordinary least squares (OLS) regression.

Figure 6

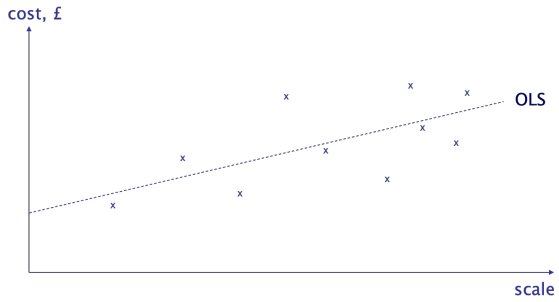

The OLS line is the unique line drawn through the scatter plot that minimises, to the greatest extent possible, the residual, unexplained variation in the data, marked by the arrows in figure 7.

Figure 7

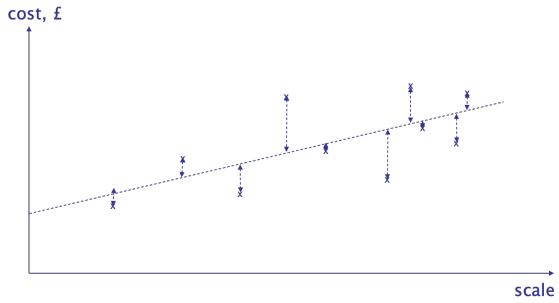

(The exact positioning of the OLS line of best fit will be determined using computer software. Strictly speaking, the algorithm is: find the equation that minimises, to the greatest extent possible, the sum of the squares of the bars shown in figure 7 – hence the name “ordinary least squares”.)

At this point, looking at figure 6, we can start to make some more informed judgments about companies’ relative efficiency. We can say, for example, that the companies circled red in figure 8 look to be relatively inefficient in that they sit above the line and seem to be spending more than we might expect for firms of their size. Conversely, the companies that sit below the regression line seem to be spending less than the modelling predicts and appear relatively efficient.

Figure 8

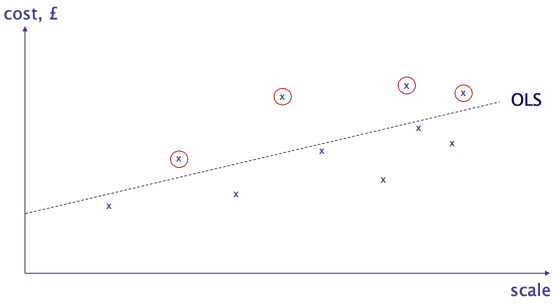

In figure 9, we reposition the regression line from the previous charts so that it runs

through what appears to be the most efficient company in the sector.

Figure 9

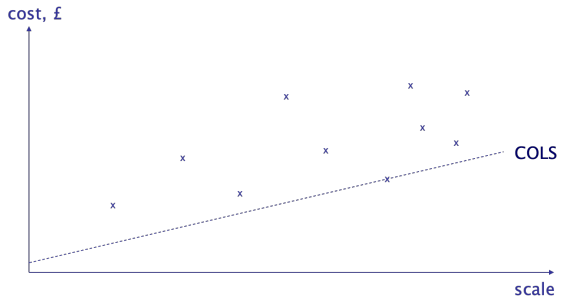

This line is labelled COLS, which stands for corrected ordinary least squares. The COLS line, as drawn, retains the information that the OLS regression revealed about the relationship between cost and scale, but in a way that exactly explains the costs achieved by the apparent frontier-defining company rather than the average company in the data set.

A regulator could at this point admissibly use the COLS line in figure 9 to size the efficiency challenge that it puts to each firm that it regulates. Take, for example, the company highlighted in figure 10. The regulator could say to this firm that, after controlling for differences in size, the company in question is not matching the cost control being exhibited by the frontier firm in the industry. As a consequence, the regulator might conclude that the firm's expenditure allowance will be built from the modelled efficient level of expenditure A in figure 10, rather than the company's actual starting costs.

Figure 10

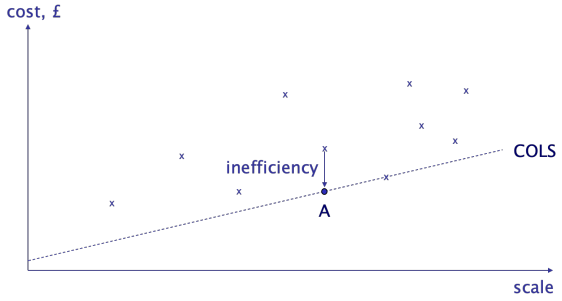

The regulator does, however, have to exercise a degree of judgment when deciding how much of a challenge it wants to table. The particular COLS line in figure 9 was fixed by reference to a single, leading company; by definition, that company is assumed to be efficient, and every other company in the sector is deemed inefficient. It is perfectly possible that a regulator may not wish to be quite that firm in its opinion. Figure 11 overleaf has a different COLS line in which the correction to the OLS regression line goes only as far as the upper quartile in the data set.

Figure 11

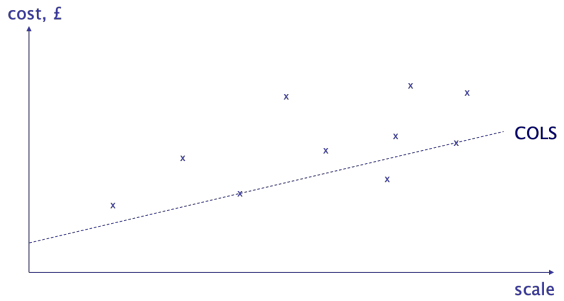

Figure 12

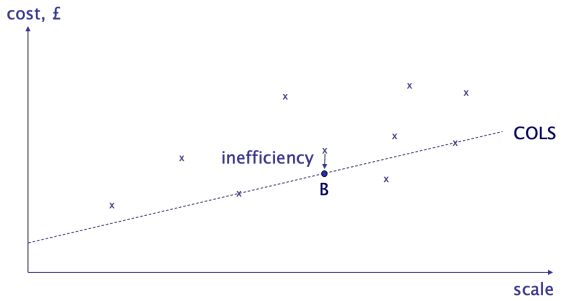

A regulator comparing figure 11 and figure 9 might feel more confident in the knowledge that the COLS line in figure 11 is defined by reference to a group of companies rather than the costs recorded by just a single firm. Among other things, the regulator may feel that the slightly more conservative upper quartile line allows for the possibility that there could be something atypical about the leading firm's expenditure in that particular year or, alternatively, something unique and special about the leading firm's performance, geography or heritage that is not being captured in the model.

When the regulator goes on to size our highlighted company's allowed expenditure, it may therefore feel that the modelled level of expenditure B in figure 12 is a more realistic and more defensible expenditure allowance than the previously identified value A.

Of course, the choice here is not simply between leading company and upper quartile. A regulator is entitled to position its COLS line at median or upper decile or upper third or at any percentile that it feels is warranted in light of the confidence that it has in it has in the robustness and the predictive power of its analysis.

The reader will note that this brings an unavoidably subjective element into the "catch up" stage of the cost assessment process. Regardless of how scientific and how exacting the numerical analysis can be made to be, there will never be an indisputably correct way of positioning the COLS line. To emphasise this point still further, a degree of judgment may similarly be required in a number of other areas, including:

- cost adjustment claims: even with the caution shown when positioning the COLS line, a regulator may want to signal a

willingness to accept submissions arguing for modest adjustments to point B on account of unique, company-specific 'special factors' that are not reflected in the modelling;

- glidepaths: a regulator is not compelled to set the highlighted company's base costs allowance exactly in line with expenditure level B at the beginning of the new price control period. It could decide instead to profile the catch up to the efficiency frontier over a period of more than one year, if it deems that the company should be afforded time to take on its high costs;

- treatment of beyond-frontier companies: figures 10 and 12 focus on the predicament of a relatively inefficient company. A regulator also has choices to make when it sets allowances for a company that sits below the chosen cost frontier. In particular, a regulator must decide whether it will initially set an allowance in line with the company's actual starting expenditure or provide a more generous modelled allowance read off from the higher COLS line.

There are no right and wrong approaches in any of these areas. As such, it is right and proper that we talk about a 'judgment' about the efficient starting level of costs rather than a categorically correct methodology for setting base costs allowances.

2.2 Extensions

The preceding charts are deliberately a simplified way of explaining how a regulator can use comparative efficiency analysis. There are several ways of extending the sophistication of the comparative analysis still further.

2.2.1 A more expansive cost function

In real life, a regulator will not be content with exploring the relationship between cost and a single measure of scale. A regulator will instead want to allow for the possibility that cost is the result of a basket of multiple explanatory factors.

Unfortunately, we cannot draw charts in a booklet like this in three dimensions, let alone show visually what a regression line looks like when we extend the cost function to include further explanatory variables. The underlying structure of the analysis will remain as set out above, however. Whereas figure 6 and all of the charts that followed thereafter depicted a simple cost function of the form:

$$\text{cost} = \alpha + \beta \,.\, \text{scale}$$

a more sophisticated benchmarking analysis might entail estimating a cost function like:

$$\text{cost} = \alpha + \beta_1 \,.\, x_1 + \beta_2 \,.\, x_2 + \beta_3 \,.\, x_3 \ldots \text{etc.}$$

where x_1, x_2, x_3 and so on are distinct, exogenous factors that the regulator thinks exert a meaningful impact on companies' costs.

A regulator that wants to allow for multiple cost drivers like this can continue to use the OLS regression technique outlined earlier to identify the line of best fit through the data. When running what is known as a 'multiple regression', the estimated coefficients labelled β_1, β_2 and β_3 in the fitted cost function will represent estimates of the amounts that each extra unit of the explanatory variables x_1, x_2 and x_3 add to costs.

2.2.2 Panel data

A regulator's ability to add more and more explanatory factors to the cost function is limited by the amount of data it has. As a matter of basic statistics, where a regulator is dealing with only a small number of firms and a small sample of observed costs, adding further x_i terms to the regression analysis can quickly exhaust the information that the data set contains and, after a point, render the results of the regression work meaningless.

This creates a dilemma: a regulator would ideally want to control for as many explanatory factors as it can in order to come to the best possible characterisation of the underlying drivers of a cost in an industry and to avoid a situation in which it wrongly ascribes differences in expenditure to efficiency and inefficiency. But it will know that using more complex regression formulae can quickly make its work less rather than more accurate.

One way of resolving this conundrum, which has found increasing favour in regulated industries in recent years, is to run regressions that contain data from more than one year. By way of an example, figure 13 goes back to a simple two-dimensional chart and depicts a data set containing three years of observations.

Figure 13

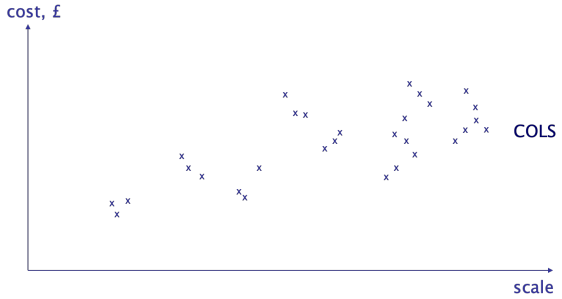

It can be seen straight away that this so-called 'panel data' has a richness that the original single-year data set lacked. As a result, it potentially permits a regulator to run regressions with a greater number of explanatory variables and so come to a more precise assessment of relative efficiency.

In practical terms, the methods that a regulator will apply to extract the greatest possible amount of information from panel data can be quite complex. A regulator may opt to use the OLS/COLS techniques that we saw in section 2.1. But a regulator can also use more sophisticated statistical techniques that recognise that the data points in figure 13 are not completely independent from one another.

One common approach is to run a different kind of regression (called generalised least squares, or GLS) that accounts for the fact that the same companies appear multiple times in the data set, as shown by the additional colouring in figure 14.

Figure 14

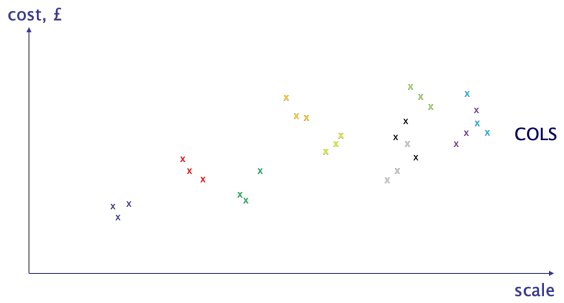

It is beyond the scope of this introductory booklet to explain the mechanics of such regressions, but the key principles outlined above continue to hold: the regulator will identify a line of best fit through the data; it will then correct this regression line to a frontier that it feels comfortable with; and it will then use this frontier to challenge companies that look to be spending less efficiently than their peers.

2.2.3 Disaggregation and triangulation

A third and final avenue that a regulator could explore involves drilling down into costs and running separate regressions for separately identifiable activities. A regulator might even assemble a compendium of top-down and disaggregated models, perhaps running two or three regression per activity type comprising different baskets of explanatory variables, in the hope of obtaining a more rounded picture of efficiency in a particular area.

The thought process here would be that a single regression can sometimes throw up rogue results. Triangulating across multiple pieces of analysis can help to minimise the risk of error and, in an ideal world, will mutually corroborate and reinforce the evidence that the regulator has of relative efficiency and relative inefficiency.

2.3 Alternative approaches

The preceding discussion hopefully shows that a regulator of multiple companies possesses the toolkit with which to make defensible judgments about efficiency. But

what about a regulator that does not have the same luxury of being able to compare multiple regional firms all doing the same job – e.g. because the regulated firm is a national monopoly or because there is a heterogeneity among service providers?

2.3.1 Other forms of regression analysis

The first question the regulator might ask is: can I generate the kind of comparative efficiency analysis outlined in sections 2.1 and 2.2 by a different means?

It might, for example, be possible to carry out some form of international benchmarking exercise. In this case, the data feeding into the regression analysis will not be from firms operating in different parts of the UK but from companies providing essentially the same services in different countries.

International benchmarking, as a general rule, tends to be somewhat more challenging than in-country benchmarking. A regulator will normally have to control, as a minimum, for the different legal and other standards that apply in different jurisdictions. It can also be surprisingly complicated to convert currencies and ensure that appropriate allowance is made for differences in country price/cost levels. However, so long as companies are carrying out broadly similar activities, my experience has been it is usually possible to obtain at least some insights into efficiency through cross-country comparisons.

In the same vein, an alternative approach might be to look at different regional units within a national company. Provided that good regional data exists, the observations feeding into the regression analysis do not necessarily need to come from different ownership groups. It is perfectly acceptable to conduct instead an internal benchmarking exercise which compares the efficiency of the different parts of an organisation.

This kind of in-company analysis obviously cannot reveal whether the company as a whole is spending efficiently or inefficiently across the board. However, it can still aid the regulatory process if the regulator is able to identify examples of good/best practice within a business and, hence, is able to challenge higher-spending regions to match the efficiency and cost control exhibited by the lowest-spending regions.

2.3.2 Other benchmarking

Regulators that, for whatever reason, cannot find suitable comparators by any of the above-mentioned methods might think next about more simplistic forms of benchmarking.

The strength of the regression method is that it allows a regulator to control for exogenous factors. But there can be categories of

expenditure which do not obviously require complex regression analysis and where a regulator can take useful information from straight comparisons of unit costs (e.g. cost per km).

2.3.3 Expert challenge

The final option that a regulator has is to subject the regulated firm to a line-by-line challenge of its expenditures.

A well-staffed regulator, employing individuals with extensive industry experience, may feel that they are capable of developing this kind of expert challenge in-house. Alternatively, the regulator may choose to call in independent, external consultancy advice. In either case, the task will be to identify in qualitative terms if and where a company could be spending more efficiently, and to size a realistic catch-up efficiency target based on the assembled analysis.

My experiences of this type of work have been quite mixed. The best expert challenges build a catalogue of compelling and objective evidence, yet stop short of telling the company, in effect, that the regulator knows how to run the business better than the company's management. The less convincing reports that tend to be based more on subjective opinions and can provoke the regulated firm into appointing its own preferred consultant, resulting in drawn out, tit-for-tat, "you say, we say" differences of views.

For this reason, most regulatory practitioners would probably put expert challenge last in the hierarchy of possible approaches to efficiency assessment. But that it is not to say that it does not have its place when circumstances rule out any of the other options detailed above.

3. Frontier shift

The comparative efficiency analysis outlined in section 2 enables the regulator to locate the efficient levels of cost within an industry immediately prior to the start of a new regulatory period. The next task is to roll these starting costs forward.

3.1 Contributors to frontier shift

No company in the economy, no matter how efficient, will see costs remain completely static from one year to the next. There are multiple reasons why the efficient level of costs in an industry tomorrow might be higher or lower than the efficient level of costs today, but two drivers of changes in expenditures stand out from the pack and merit particular attention.

The first is **input price inflation**. As a general rule, the prices of capital, labour, energy and materials all change at a regular frequencies. As a direct consequence, the cost of the basket of inputs that a company needs in order to provide its services will naturally move up or down over time.

In calibrating the gradient of the line drawn earlier in figures 3 and 4, it is important that the regulator makes a fair and reasonable allowance for the impact that the underlying, sector-wide rate of input price inflation will have on costs over the course of each new price control.

A regulator also ought to recognise that companies in most industries have historically been able to combine inputs into finished outputs more efficiently over time. Where a regulated company is likewise capable each year of reducing the quantity of workers, energy or materials that it uses to produce a given level of service, or where it ought to be able to deliver better service from a given quantity of inputs, it is right and proper that the regulator should recognise this scope for efficiency gains in its forward-looking cost allowances.

The ability to improve its input-output ratio year by year can be termed annual **productivity improvement** or **ongoing efficiency**.

(NB: for the avoidance of doubt, productivity is purely about quantities – i.e. number of workers, tonnes of materials, etc. Any changes in total labour costs, energy costs and/or materials costs driven by changes in wages, energy prices and materials prices should be accounted for in the aforementioned analysis of input prices and is separable from the concept of productivity growth.)

The rate of frontier shift in an industry may therefore be thought of, at least in simple terms, as:

Frontier shift

= Input price inflation *minus*

Ongoing productivity growth

3.2 Input price inflation

A regulator will typically look to independent forecasters to help it pinpoint the likely rate of inflation that will impact the different input types that a regulated company will be buying. Potential sources of information include the Office for Budget Responsibility's twice yearly economic forecasts, HM Treasury's monthly survey of independent economic forecasts and any one of a gaggle of sector-specific publications that cover more specialist input types (especially in the energy and materials cost categories).

In the past, all regulators would provide for a given, forecast level of input inflation out to the end of the control period. More recently, however, some of the regulators have recognised that there is also the option of indexing its allowances to actual out-turn input price data, using one of the uncertainty mechanisms outlined in section 1.2 of Part 3 of this Guide. In this way of doing things, a regulator can be fairly relaxed about its initial upfront allowance for input price inflation safe in the knowledge that a suitably calibrated indexation formula will adjust opex/capex/totex allowances up or down by an appropriate amount once the Office of National Statistics (or other chosen data providers) have published the actual rate of change in the cost of relevant input types in a given year.

At a technical level, the key requirement is that the provision for input price inflation – whether a fixed ex ante allowance or an adjustable, index-linked amount – is set equal to the real rather than nominal rate of input price inflation. This follows from the general treatment of inflation outlined in Part 2 of this Guide. If a company's price control is automatically indexing every year in line with, say, CPIH inflation, the company already has an in-built compensation mechanism that covers the general level of price inflation in the economy. The regulator therefore only needs to recognise the differential between sector-specific inflation rates and general consumer price inflation when setting allowances.

(NB: it is for this reason that allowances for input price inflation are sometimes termed "real price effects" or RPEs.)

3.3 Productivity growth

The judgment that a regulator makes about the potential for productivity savings will normally be based on historical experience.

Regulators look, in particular, at the rates of productivity growth that have historically been achieved in sectors that are bear some similarity to the UK's regulated sectors, and take the view that they are entitled to expect the firms that they regulate to match the benchmarks set by these comparator industries.

This does leave room for debate about exactly which industries in the economy and which historical time periods provide the most relevant benchmarks. At the time of writing, there is a good degree of consensus that one can ontain useful information by looking at the following comparator sectors:

- construction
- transportation and storage
- professional, scientific, technical, administrative and support services

There is less agreement on the question of time period. The UK as a whole has been suffering from very low productivity growth since the 2008 global financial crisis, and the reasons for this so-called 'productivity puzzle' are not well understood. This has made for differences of view as to how much weight regulators should give to pre-2008 comparator data (showing relatively strong rates of productivity growth) compared to post-2008 comparator data (which typically show noticeably lower rates of productivity improvement).

Notwithstanding these difference of opinion, UK regulators have more often than not tended to factor a rate of productivity improvement of around 1% per annum into recent price controls. This makes for an informal rule of thumb that can provide a starting point for the debates about productivity growth in each new price review.

3.4 Overall frontier shift

The overall rate of frontier shift will be the simple mathematical consequence of the numbers that a regulator alights on for real input price inflation and productivity growth.

There can be some years in which the overall rate of frontier shift will be positive and the line in figure 3 will slope upwards (typically when industries are encountering high rates of real input price growth) and some years in which the overall rate of frontier shift will be negative and the line in figure 3 will slope downwards (typically when industries are seeing benign or positively negative real input price inflation).

It is not, therefore, the case that the line I drew in figure 3 has a natural tilt in one direction or the other. The rate of frontier shift in an industry has to be continually assessed and reassessed in light of the economic conditions of the day.

4. Growth and Enhancement Expenditure

4.1 Overview

In contrast to the systemised way in which regulators try to approach the calibration of base costs, catch up and frontier shift, it is much less easy to describe a textbook way of setting allowances for growth and enhancements.

This reflects the bespoke and circumstance-dependent nature of such expenditures. Whereas it was possible to think of a standardised amount of base activity, the implications that growth and higher service levels have for cost can vary considerably according the specific demands that the company is being asked to accommodate.

This being said, there may be some steps that a regulator can take to streamline the cost assessment process.

It might, for example, be possible to use the regression equations from section 2 to model the additional expenditure the regulated company should expect to incur as variables like population and number of customers increase. The regressions, after all, seek to identify the relationship between scale and costs, so there is a basic logic in thinking that the models will be able to predict the change in efficient expenditure when the x_1, x_2, x_3, etc. values change.

Alternatively, it might be possible to use historical experience to discern an underlying relationship between growth and costs. This might be in the form of a simple elasticity or, conceivably, with sufficient data, it might be possible to identify a relationship between growth-related costs and multiple explanatory factors.

Some regulators have on occasion gone even further and focused their regression work from the outset on firms' projected future costs, rather than on out-turn historical costs, where they have reason to think that companies will grow at a similar speed and in a similar way.

There will come a point, however, when the scope for mechanising the sizing of allowances for new costs will be exhausted. At that point, the regulator will unavoidably need to take a more bottom-up approach focusing on two sequential considerations: need and £m amount.

4.2 Need

When confronted with plans for brand new expenditures, a regulator will need to establish first or all that there is a robust justification for the proposed new activity or scheme.

In some sectors, the regulator may be helped by an outside authority who will prescribe that specific investments must be delivered or, alternatively, that specific new quality or safety

standards must be achieved (examples of such authorities include the Health & Safety Executive, the energy sector's Future System Operator, the Environment Agency and the Drinking Water Inspectorate). In other cases, however, there may genuinely be open questions about the breadth and scale of the improvements that a regulated firm should be tasked with delivering.

The onus usually falls in the first instance on the regulated firm to justify the business cases for any new expenditures that it wishes to incur. The regulator will expect to see a business plan that sets out, as a minimum, consideration of alternative technical options, the costs and benefits of the proposed scheme, and evidence of customers' willingness to pay. The last of these components, in particular, has become a much more important part of the price review process in recent years, with companies investing considerable effort and expense to ascertain the views and preferences of their customer base.

The final decision on whether to make allowance for a new scheme will normally sit ultimately with the regulator. If unconvinced by a company's proposals, the regulator may choose to disallow a project in full or in part. But it can also leave the door open by providing for a possible in-period change to allowances if certain triggers are met or if the company is able come back with a revised plan (see the discussion in Part 3 of this Guide about regulators' use of uncertainty mechanisms).

4.3 Costing

Approved schemes, like other items of expenditure, will enter the company's opex/capex/totex allowance with a fixed costing.

At this point, it really does become impossible to give a one-size-fits-all account of the steps that the regulator will take in order to arrive at the relevant £m amounts.

Sometimes a regulator will be able to identify benchmarks it can apply to cost a particular type of scheme. Sometimes the regulator will commission an expert challenge to the regulated company's submitted costings. Sometimes the regulator will focus squarely on the effectiveness and efficiency of the company's procurement practices, and be willing to accept that the winning tender prices emerging from a well-run competitive process reveal efficient costs.

It is perhaps sufficient to conclude the discussion in this booklet by saying that the regulator will have to do the best that it can with the information that is available to it. This can make for a particularly taxing part of the price review, albeit regulator, company and customers alike can take comfort from the

knowledge that any costing error, and hence under- or over-spending, will usually be shared between shareholders in accordance with the sharing rules set out in Part 3 of the Guide.

If you have any questions about the content of this booklet, please get in contact at:

john_earwaker@first-economics.com