


GUIDE TO ECONOMIC REGULATION



Part 3: Incentives

John Earwaker

Foreword

This is Part 3 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.

The focus in this booklet is on incentive design. In Part 2 of the series we explained how a regulator calculates a firm's revenue requirement and how this revenue requirement can be turned into a price control. We now move on to see how the regulator will try to do more than limit prices. In particular, we consider how a carefully crafted set of conventions and rules can incentivise a regulated firm to make cost savings and improve the quality of the service provided to customers.

1. Efficiency incentives

We begin by looking at the incentives that a regulated firm has to improve its efficiency. There are two parts to the story: we must first highlight the importance of having a fixed interval between price reviews; then we can look at some of the mechanisms that regulators have devised to further target and strengthen incentives in specific areas.

1.1 The regulatory lag

At first sight, the convention that regulators conduct price reviews at pre-determined intervals may seem like a way of containing administrative burden, given that the task of costing up a firm's or an industry's revenue requirement(s) using the methodology detailed in Part 2 of this guide requires all of the parties involved to expend considerable effort. What may not be obvious, however, is that a fixed "regulatory lag" is intended to provide companies with a very strong incentive to go out and minimise their costs.

In order to explain why this is the case, it is helpful to go back in history to the 1980s when the nation's policymakers were considering how to regulate prices in the UK's newly privatised industries. The template for economic regulation at that point in time could be found in the US, where privately owned utility companies had been supervised by regulators over a period of many decades.

Many of the ideas that we introduced in Part 2 of this guide were also evident in the US approach to regulation (such as the principle that customers pay in instalments for investment, and the inclusion of an allowed return in the calculation of companies' revenue entitlements). However, crucially, the rule was that price caps, once agreed, could be revisited and reset at the request of either the regulator or the regulated company whenever one side was of the view that circumstances had changed and the company's actual costs were drifting too far away from the assumptions that the regulator had made in its most recent price review.

The consensus was that this form of regulation did not work particularly well. In particular, the perception was that the US system resulted in unnecessarily high costs and unnecessarily high prices for consumers. This was for at least three reasons:

First, whenever a company was confronted with a problem – say a contractor that was asking for higher payments, or assets that were costing more to maintain – the company's first reaction, rather than address the problem at source, would be to turn to the regulator for an offsetting increase in the prices that the company could charge to its customers. The regulator, of course, would push back on a company's claim for higher prices, but would often be outgunned by the company's better access to information and

data (a problem that is sometimes labelled “information asymmetry”).

Second, the regulated firm also knew that its regulator would very likely file for a price reduction if it saw that the firm had made cost savings. The company would therefore quite understandably feel that it had very little reason to make efficiencies, particularly if those efficiencies required effort and expense to achieve.

Third, in the absence of any clear imperative to minimise costs, and because the profits that a company was able to earn were calibrated by reference to a cost of capital x asset base calculation, companies had realised that the easiest way to grow the returns they generated for shareholders was to increase the size of their asset bases by adding more and more investment. This was seen as leading to over-investment or unnecessary “gold-plating” of companies’ networks, contrary to the interests of customers.

A solution to these problems was offered by an academic, Professor Stephen Littlechild, in a report produced for the Department of Industry in 1983. Littlechild realised that the way to avoid the problems that had afflicted the US was for the regulator to commit to a fixed profile of prices for a fixed period of time – i.e. whereas in America reviews could sometimes happen as frequently as every one or two years depending on circumstances, in

the UK the convention was to be that reviews would take place strictly every n years.

We can show the power of this seemingly very simple idea using the following charts.

Figure 1


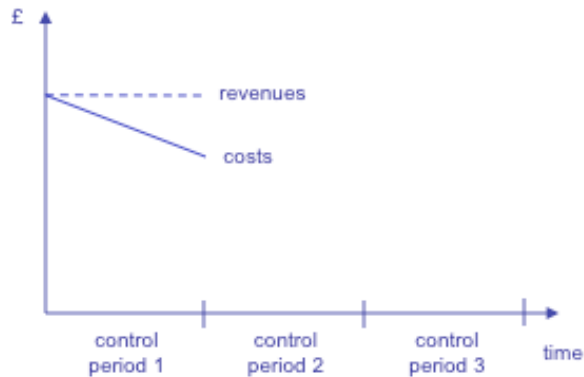


Figure 1 depicts the outcome of a regulator’s price review as an entitlement to collect a fixed amount of revenue over a period of, say, five years. To keep things simple, we draw this as a flat line, but it could just as easily be that the annual revenue entitlement changes from year to year in line with changes in the values of the building blocks that we identified in Part 2. What is crucial here for the purpose of the discussion that follows is that the regulated company knows with certainty how much revenue it will receive in each of the next five years.

How should a profit-maximising company respond to a fixed revenue entitlement? By design, the company cannot increase its profits by growing its revenues because its revenues have been capped by the regulator. However, it might be able to increase its profits if it is able to reduce its costs below the level assumed by the regulator in its price control calculations.

The solid line in figure 2 depicts a company that succeeds in finding new, unforeseen efficiencies. In practical terms, this could mean spending less on opex, spending less on capex, reducing the cost of capital and/or paying less in tax.


Figure 2



It can be seen that the regulated company in this illustration will benefit from additional profit

equal to the difference between revenues and out-turn costs, as highlighted in red in figure 3.

Figure 3

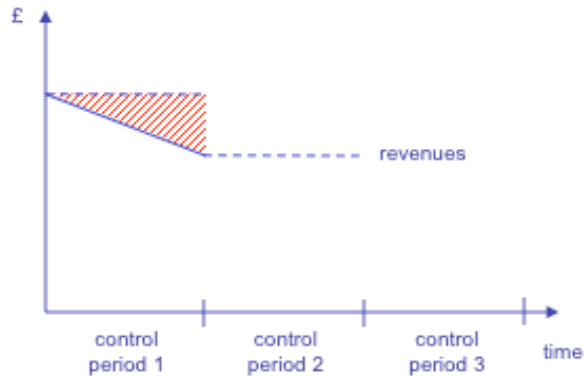


The chart as drawn therefore holds the promise of a very happy outcome for shareholders. A company that finds ways to out-perform will be able to pass the benefit of that out-performance to the equity owners. We would expect, therefore, that shareholders would be pushing the firm's managers to go out and find cost savings wherever possible so as to maximise the size of the red triangle.

What about customers? Well, initially customers are still required to pay the revenues that the regulator determined at the start of the control period. Figure 4, however, shows that the short-term win for shareholders is then transformed into a longer term benefit for customers when the five years comes to

an end and the regulator resets the company's revenues at the start of the next control period.

Figure 4




The regulator when reviewing the company's revenue requirement at the scheduled reset of prices will see that the company has made efficiencies and is now spending less than was the case in the past. The regulator will at this point recognise the lower opex or the unspent capex or the lower cost of capital or the lower tax bill in its allowed revenue calculation for the next five-year period, and capture those efficiencies for the benefit of customers in the form of lower prices.

What we now have here is a win-win situation: if a company is able to make efficiencies, figure 3 shows that shareholders will likely benefit for a period of several years depending

on when in a period a saving is made. But figure 4 shows that customers also soon benefit from lower prices, potentially for many years afterwards.

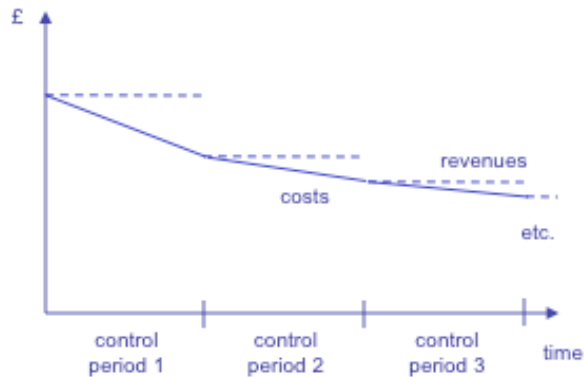
Moving forward, at the start of the new regulatory period the company will again be confronted with a situation in which its revenues are fixed for a period of five years. Once again, the company and its owners will see that there are profits to be made if the company is able to reduce its costs below the level assumed by the regulator in its latest review. And once again, if the company is able to succeed in making efficiencies, the regulator will be able to capture that benefit for customers in the form of lower prices at the next review.

Figure 5



In an ideal world, the virtuous cycle shown in the above charts would continue in perpetuity, with each new control period seeing the company unlock new, previously unforeseen efficiencies and each new price review capturing those efficiencies for customers in the form of lower prices.

Figure 6



If, however, for whatever reason, the wind blows in the other direction, and a company for whatever reason finds it cannot out-perform its regulator's price control assumptions, the incentives that were present in the preceding charts operate in the other direction and act as a deterrent against overspending.

Figure 7 shows what happens if costs were to go up rather than down (e.g. because a firm spends more on opex, spends more on capex,


faces a higher cost of capital, and/or pays more in tax). In this case, the red highlighting is a loss of profit for the firm and consequent lower returns for shareholders.

Figure 7



A company that is under-performing, or which is looking at the prospect of under-performing, ought to realise that it is in its shareholders' best interests for it to minimise its additional costs as much as possible. This in part because of the short-term pain it will suffer but also in part because it cannot assume that it will get relief from inefficiency when its allowed revenues are reset at the end of the period - i.e. while we can be reasonably certain that the regulator will capture the benefit of lower costs, the company will be uncertain exactly how the regulator will look at higher costs at the next price review.

Figure 8



In all scenarios, therefore, there is a strong incentive for a company that wants to maximise the profits it makes for its shareholders to spend as efficiently as possible.

The incentive arises, to be clear, because the regulated company knows that it has a fixed revenue entitlement for a fixed period of time. If this were not the case, say if the regulator were to reserve the right to step in at its discretion to reset revenues when it saw the company spending less than it had forecast, or say if the company were entitled to a price adjustment in circumstances where it needed to spend more, the incentive to minimise expenditure would be severely weakened, in line with the story we told at the start of this section about the old-style US system of regulation.

1.2 Extensions to the basic framework

All of the UK's regulators upon creation took up the suggestion of a fixed regulatory lag, and the vast majority found exactly the kind of virtuous cycle that we depicted in figure 6. As is often the way, several of the regulators also felt with the benefit of experience that there was room to improve on the basic structure that Littlechild had identified. This means that the framework of regulation that we now have in some of the UK's regulated industries is a slightly modified version of the original design, albeit one that remains by and large true to the underlying principles of fixed-period, fixed price regulation.

We next consider some of the extensions to the framework.

1.2.1 Sharing rules

The first modification relates to the way in which efficiencies are nowadays shared between companies and customers.

After a few price review cycles, regulators noticed that there is one significant problem with the framework that we laid out in section 1.1. In figure 9 we can examine the benefit that the company obtains from an opex saving and in figure 10 we look at the benefit that arises from a capex saving of the same £m amount. In both cases, the building blocks on the left-hand of the chart are the regulator's

price control allowances and the amounts on the right-hand chart are the company's actual costs.

Figure 10





Figure 11



In the case of the opex saving, the rule (see Part 2 of this guide) that opex is matched pound-for-pound with revenue each year

means that every £100 of opex allowance not spent equates to £100 of additional profit for the company. But this is not the case with a capex saving. Here, the convention that customers pay for capex in instalments means that the benefit to the company of not spending £100 of capex allowance will be the benefit of keeping one year's instalment on £100 of unused allowance plus one year's of associated allowed return. These two items summed together will almost always be only a fraction of the £100 saving.

Illustration: Suppose, by way of an example that a company's assets are depreciated over 40 years and that the cost of capital is 3%. The annual profit that a company can make by spending £100 less on capex is:

$$\text{Depreciation} = £100 \div 40 = £2.50$$

$$\text{Return} = £100 \times 3\% = £3.00$$

$$\text{Total} = £5.50$$

This compares to the annual profit that a company makes from a £100 opex saving of £100.

It follows that, even if a company gets to keep the revenues associated with unspent capex for a period of several years prior the next regulatory reset, there is less profit in making a capex saving than there is in making an opex saving.

We could also draw the same pictures in reverse if a company is looking at the prospect of over-spending by £100. Again, £100 of additional opex translates directly into a £100 loss, while the downside of incurring £100 more on capex is felt via the absence of one year of depreciation and one year of return on the £100 spent, which sums to only a fraction of £100.

On seeing the above arithmetic, a rational firm might reasonably conclude that it is worth looking carefully at the mix of its expenditures. Suppose that it is facing a problem on its network and could potentially address that problem with either £1m of opex or £2m of capex. The right thing to do for customers would be to choose the opex solution. But the maths we just identified probably points in the opposite direction in that, counter-intuitively, the firm's loss if it spends £1m of additional opex is much greater than the loss it suffers if it spends £2m of additional capex.

After finding evidence that such considerations were distorting company's decision-making – in particular by creating a “capex bias” – some of the UK regulators concluded that they need to take steps to ‘equalise’ incentives across expenditure types. Several different approaches have been applied over the years, but the preferred solution at the time of writing entails announcing that companies will be allowed to retain a fixed percentage share of any under-

or over-spending against expenditure allowances, irrespective of whether the costs involved are in opex or in capex.

A regulator might, for example, set a sharing rate of 50%. If a company is able to under-spend a cost allowance by £100, the rule would be that, irrespective of whether the £100 saving is classed as opex or classed as capex or even a mix of both:

- the company will be able to keep £50 of the unused allowance, increasing the company's profit by £50; and
- the remaining £50 will be passed to customers in the form of lower prices.

Similarly, if a company over-spends by £100:

- the company will have to pay for £50 of extra cost; and
- the remaining £50 will be passed to customers at the next regulatory reset via a price increase.

In practical terms, this kind of sharing rule requires the regulator to implement a ‘true-up’ – i.e. a top-up or a mark-down – at the end of each control period. The regulator will look back at the company's actual costs in each year of the five-year period and calculate the amounts of any under- and over-spending. The regulator will then run a fairly complicated reconciliation spreadsheet calculation to establish how much benefit or cost the

company has already suffered as a result of the natural out- and under-performance depicted in the earlier charts, and will calculate the residual price reduction or price increase that is required in order to achieve the desired percentage allocation of under- and over-spending between company and customers.

1.2.2 Exceptional items

The second important change that regulators have made to the basic five-year, fixed-price system of regulation involves singling out certain types of cost for special treatment.

While the general description of regulation that we provided in section 1.1 may have accurately described how regulation worked 20-30 years ago, nowadays most regulators tend to be quite reticent to give fixed allowances for costs that are particularly volatile or unpredictable, especially if changes in costs are not wholly within a company's control. They do, however, still recognise the importance of providing companies with incentives to minimise their expenditure and the damage that a straight cost pass-through arrangement can cause. This leads to the thought that the best way of setting a five-year control might be to provide for a set formula or set formulae that will be used to fix a company's revenue entitlement rather than a completely fixed £m revenue amount.

A modern-day price cap may therefore take the form:

$$\text{Revenue cap} = \text{£}100\text{m} + A_t + B_t + C_t + D_t$$

where £100m (say) is a base amount of revenue, and the A_t , B_t , C_t and D_t terms provide for monies to be added to or subtracted from that £100m depending on particular events that may happen or on movements in specified published data.

To illustrate how this might work, suppose, for example, that a regulator is concerned about its ability to forecast future interest rates and, hence, the possible error that could creep into its calculation of the 'allowed return' building block. A price control might specify that the company's revenue entitlement need not be a fixed amount of £x, but will be £x plus/minus an adjustment term that will add revenue or subtract an appropriately sized amount of revenue depending on whether interest rates turn out to be higher or lower than the regulator's assumptions.

Provided that the adjustment term references a published market benchmark, and not the company's actual out-turn cost of debt, all of the incentive properties that we described in section 1.1 remain intact. The company may no longer have a known, fixed £m revenue entitlement, but the direction and the scale of any in-period adjustments lie wholly outside of its control. The company therefore ought to

see that its best strategy is still to go out and minimise its costs to the greatest extent possible. The only difference from the charts that we drew in section 1.1 is that the red triangle of out-performance will be against a dotted blue line whose value is determined formulaically in period rather than a completely fixed £m revenue entitlement.

Other obvious applications for the kind of indexation mechanisms we have just described might include:

- changes in economy-wide wage growth;
- changes in other relevant input price indices;
- changes in corporation tax rates.

A regulator can also devise adjustment mechanisms to capture changes in external drivers of activity that a company undertakes – say, the amount of new load connecting to a network or the volume of asset replacement mandated by the health and safety authority.

When a regulator cannot pre-programme a formula using any sort of exogenously determined and easily observable reference value, the regulator may still wish to consider some form of in-period adjustment. In this case, the last-resort option will be a price control reopener, in which, say, the term D_t in the above revenue cap formula will be an amount whose value will be set by the

regulator as part of a narrowly defined in-period review.

What this all means in practical terms is that today's price controls are often significantly more complex than the price controls of the 1980s and 1990s. However, the thinking that Littlechild first tabled almost forty years ago very definitely lives on in that the addition of indexation mechanisms, volumes drivers and price control reopeners is implemented in a way that does not cut across the principle that the regulated company must know in advance that it will make additional profit for its shareholders if it finds ways of reducing its costs and suffer a loss if it under-performs against its regulatory allowances.

2. Service quality incentives

A regulatory regime which focuses incentives exclusively on price and costs is an incomplete regulatory regime. It leaves open a big loophole in that a very simple way for a company to reduce its costs and profit from out-performance is to skimp on service quality. We must turn next, therefore, to the measures that a regulator puts in place to protect the interests of customers as regards the outcomes that a regulated company delivers.

2.1 Minimum service standards

The simplest form of intervention that a regulator can make in this area is to impose a set of minimum service standards. The regulated company will know that, even in the face of an incentive to reduce costs, it must expend sufficient resources to meet these standards. It will also know that the regulatory toolkit (see Part 1 of this guide) gives the regulator ample powers to take enforcement action and fine the company in the event that the required performance levels are not achieved.

2.2 Financial incentives

The question that regulators have asked themselves over the years, though, is: is this enough? Yes, minimum service standards put a floor under service quality. Yes, a regulator can adjust required service levels at a price

review if it wants to push a company to deliver better outcomes. But is it enough for a regulated company just to deliver to a set of regulator-designed specifications each year? Shouldn't a regulated company be much more responsive to the needs of its customers in real time and without waiting for direction from the regulator?

This thought has led to some innovative developments in the scope and design of financial incentives. As a result, most current-day price controls provide for incentive schemes that either allow a company to earn monetary bonuses or require a company to pay monetary penalties depending on the service quality levels it achieves.

The basic principles that underpin these schemes can be outlined using the following framework.

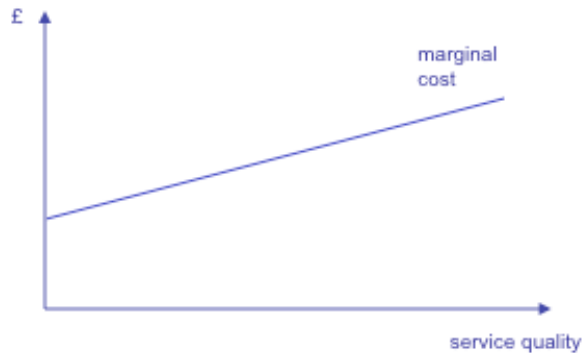
Figure 12



The above chart has service quality on the horizontal axis and pound notes on the vertical axis. Better service quality in this thought experiment could mean customers getting more of something that they want (e.g. faster response times or higher overall satisfaction) or less of something they dislike (e.g. less delay or fewer supply interruptions). In either case, customers are better off as we move from left to right in the chart.

We can draw first of all the perspective that the regulated firm has of the costs it will need to incur if it seeks to deliver better service quality.

Figure 13



This cost line is labelled 'marginal cost' because it represents the extra cost, or the incremental cost, that the company will incur in achieving each successive extra unit of

service quality. The curve is deliberately drawn to slope upwards on the assumption that a company will find it relatively straightforward to make a few quick service improvements, but much more costly to keep on pushing service standards ever higher over time.

(Think here of the resources that a train company would have to deploy if were to be asked to drive performance closer and closer to 100% on-time arrivals, or the effort that a water company would have to go to in order to ensure that its pipes almost never leak.)

Next we can layer on the perspective that customers have.

Figure 14



The 'marginal benefit' line captures the incremental valuation that customers attach to

service improvement. This curve is drawn in such a way as to indicate that each unit of service quality improvement is valuable and worth something to customers. In practice, the line might not be completely flat as drawn in the chart – i.e. it could conceivably slope upwards or downwards – but for the purposes of the discussion we do not worry too much about the exact shape of the curve. What is important is that customers attach a non-zero value to better service quality.

The sweet spot in the above chart is the point of intersection.

Figure 15



At service quality level A, the marginal cost that the firm incurs in achieving the last step up in quality exactly matches the marginal benefit to customers. Anywhere to the left of the point of intersection, we can see that a

firm could spend more to improve service quality and the value that the firm's customers would place on that improvement more than justifies the expenditure involved (i.e. the marginal benefit line lies above the marginal cost line). A service quality outcome of less than A is, therefore, distinctly sub-optimal – i.e. it is in everyone's interests that the firm goes out and does what it is needed to get service quality up to level A.

As we move to the right of point A, however, we can see that even though a firm could conceivably continue to spend money to push service quality even higher, the benefit that customers take from better performance is no longer sufficient to justify the resources that are required (i.e. the marginal benefit line lies below the marginal cost line). Moving to the right of A cannot therefore be optimal – i.e. customers prefer not to have to pay for a higher level of service quality than A, notwithstanding that this means that the service they receive falls short of perfect.

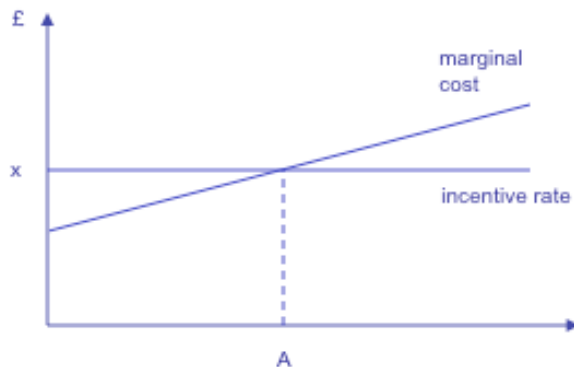
The job that company and regulator have during a price review is to determine where this sweet spot sits, make appropriate allowance for the cost to achieve in the price control calculation, and formalise the specific numerical value of the company's performance obligation.

How then to persuade a company to do more than just the minimum required to meet a regulatory target?

Let's take another look at the marginal benefit line in figure 15. Suppose the regulator can see that customers value each unit of service quality at $\pounds x$ per unit. Imagine also that a regulator were to say to the company that the company can earn a bonus from customers in the form of a small addition to prices/revenues set equal to $\pounds x$ for every unit of service quality that it delivers beyond point A, but will pay a penalty to customers in the form of a small mark-down to charges equal to $\pounds x$ for every unit of service quality deterioration below point A.

The firm would then face the financial incentive scheme shown below.


Figure 16



We know from the preceding discussion that the sweet spot for the firm is initially point A. We can see this by again comparing the marginal cost and marginal incentive payment lines in the chart below: any move to the right of point A is not financially viable since the marginal cost to achieve would exceed the bonus that the firm can collect, while any drift to the left of point A is also sub-optimal because the costs that the company would save by letting service quality slip would fail to cover the penalty the firm would have pay.


But what if a company comes up with some clever new ideas after the price review has concluded? Picture, in particular, a situation in which a company is not content just to deliver a plan that it devised several years previously but keeps putting its mind to the possible ways that there are of improving service quality for customers. In figure 17 we show what happens when these efforts bear fruit and the company comes up with a new way of working or a new investment or a new technology that enables it to improve service more cheaply than it previously thought was possible.

Figure 17



By looking once again for the point of intersection, we see that the sweet spot has shifted to the right from service level A up to service level B – i.e. it is now worth the firm's while to spend a little bit of extra money knowing that payments from the regulator's service quality incentive scheme will more than pay for the required expenditure.

Figure 18



What we are looking at in figure 18 is a business case for a new project. Looking at the numbers as presented, a rational company will choose voluntarily to spend additional money knowing that the sum of the incentive payments it can claim via an increase in customers' bills will more than pay for its costs. This in turn benefits customers who are better off when they receive a better quality of service, both in the short term and in later years.

It follows that a regulator can foster much more dynamic, customer-focused behaviours by putting financial incentives around service quality measures. Indeed, while we have worked through the consequences of devising a single reward scheme for a single performance metric, the same principles could potentially be applied simultaneously to multiple metrics covering multiple facets of the outcomes that a firm is responsible for.

In the end, this could ultimately all come together into a sort of rate card that attaches specified amounts of money to a range of different performance measures based on the marginal benefits that consumers attach to each outcome.

An outcome delivery incentive rate card

Metric 1

Benchmark = A_1 units

Payment rate = \pm £10.00 per unit

Metric 2

Benchmark = A_2 units

Payment rate = \pm £8.50 per unit

Metric 3

Benchmark = A_3 units

Payment rate = \pm £100 per unit

·
·
·


Metric 10

Benchmark = A_{10} units

Payment rate = \pm £0.50 per unit

Such a scheme would also have one final benefit for customers. Suppose that, despite all of our optimism, a company ends up falling short of its performance obligation and only delivers performance at the level C shown in the chart below.

Figure 19



We said at the beginning of this section that a regulator is able to fine the company for failing to meeting its obligations. In the past, regulators have found that this is not completely ideal because by law any fines that a regulated company pays go to HM Treasury rather than to customers. A financial incentive scheme solves this problem because the $\text{£}x$ per unit penalty that the company would pay in the situation depicted above would go directly to customers in the form of a deduction from bills and, hence, would act as a form of automatic redress for the sub-standard quality of service received.

Financial service quality incentive schemes – or outcome delivery incentives (ODIs) – therefore provide not only a reason for companies to push forward on service improvement, but also an in-built incentive not to let standards slip.

3. Reputational incentives

In sections 1 and 2 we focused on the monetary incentives that regulators can put in front of regulated companies. We now conclude this Part 3 of this guide to economic regulation by looking at the use that regulators can make of non-monetary, reputational incentives.

To fix ideas, imagine a regulator like Ofgem or Ofwat which regulates multiple companies, or a regulator like ORR which has a national company comprised of multiple regional business units. Think also of a situation in which the regulator can compile a simple league table which ranks companies or business units according to a certain measure that is important to customers – say, companies’ quality of service or the quality of companies’ business plans.

League table

1. Company H
2. Company A
3. Company C
4. Company G
5. Company D
6. Company J
7. Company B
8. Company I
9. Company F
10. Company E

Might it be that companies would care quite a lot about where they sit in such a table, even if there is no monetary prizes at stake? Would a company want to be at the top of the ranking, or at least want to be near the top? And, conversely, would a company want to avoid coming last or appearing near the bottom of the list?

The answer to these questions, as observed in several different industries over a period of many years, seems to be “yes”.

There are different explanations for why this should be the case. One possibility is that companies near the top of a league table are more attractive to investors than companies near the bottom of a league table. Another view is that a high ranking has intrinsic worth insofar as that a company is likely to get an easier ride from the regulator at its next periodic review. It could also be that there is a direct correlation here between the outcomes captured in the table, which do not confer direct financial reward and penalty, and other outcomes such as efficiency, which do have a direct impact on a firm’s bottom line.

Arguably the most persuasive explanation of all, though, is that a set of company rankings matter because the management of a company care at a personal level where their company is ranked, and by extension how they are perceived, compared to their peers. Looked at in this way, so-called reputational

incentives play first and foremost into an individual's professional pride in his or her work, or perhaps even a person's ego.

It follows that a regulator may not always need to design clever financial incentives of the kind that we described in the previous sections in order to persuade companies to take efficiency and service quality seriously. If the conditions are right, it could be that non-monetary reputational incentives are just as powerful a motivator for individuals.

What are these conditions? The litmus test is whether boards will be willing to expend effort and deploy resources to attain the best possible position in the regulator's relative rankings even in the absence of a direct monetary reward. This entails first of all that everyone believes the competition is broadly fair to all the participants – i.e. that the regulator has not rigged its rules and assessment criteria to the benefit of already favoured companies and against out-of-favour companies. It also requires that the results of the competition confer genuine status and 'matter' to a wide group of stakeholders. This could require that the regulator will go out of its way to get publicity for its rankings through a variety of channels (e.g. published documents, industry comms, local and national press).

Ideally, people that matter will know at any given point in time who are the leaders and who are the laggards. Thereafter, the hope will be that the leaders will strive to remain leaders and the laggards will take the steps that are necessary to improve their standing.

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

john_earwaker@first-economics.com