WHY MOST THINGS FAIL

Paul Ormerod

Chapter 2 A Formula for Failure

It is easy to be critical about economics. PJ O'Rourke, in his entertaining book *Eat The Rich*, has a stimulating definition of the content of the subject: 'One thing that economists do know is that the study of economics is divided into two fields, "microeconomics" and "macroeconomics". Micro is the study of individual behaviour, and macro is the study of how economies behave as a whole. That is, microeconomics concerns things that economists are specifically wrong about, while macroeconomics concerns things economists are wrong about generally'.

As far as most economics textbooks are concerned, O'Rourke's definition seems perfectly reasonable. Failure is all around us, a pervasive feature of everyday life. Government policies fail, firms fail, whole economies fail and remain enmeshed in poverty. Failure is both general and specific. As Lucy Kellaway, award-winning columnist in the *Financial Times*, recently wrote, albeit slightly tongue in cheek: '83 per cent of Chief Executive Officers fail'.

The existence of failure on this scale is simply not recognised in economics. Instead, in much of economic theory, getting the right strategy, the right policy, is straightforward. It is simply a matter of following the appropriate formula laid down in the textbook. The formula itself may be expressed in varying degrees of mathematical complexity, depending upon the level of student at which the text is aimed. But it will in essence remain the same.

Economics as an academic discipline does have a very valuable strength. It trains people to think analytically. And as we will see during the course of the book, the subject has not stood still over the past thirty years. Important progress has been made in understanding some aspects of how the world behaves.

But in essence, economic theory remains rooted in a vision of the world which was derived from the physical sciences of the 19th century. The achievements of the latter are plain for all to see. The analytical techniques and the mathematical tools used by 19th century scientists enable us to understand a great deal of the world around us.

These have been much less successful, however, when applied to human social and economic systems rather than to questions in the natural sciences. The fundamental

reason is that this approach regards equilibrium - a static, changeless state of the world - as the natural order of things. The whole panoply of the differential calculus, the branch of mathematics which is by far the most widely used in economics, is focussed on finding equilibrium solutions. Solutions in which the system is at rest, is static, in which continuity and lack of change are its hallmarks.

This is simply not the case either with society or with the economy. As we saw in the opening chapter, one of the most obvious features of firms in the real world is that they fail. To remind ourselves of just one piece of evidence, over 10 per cent of all companies in America, the largest and most successful economy in the history of the world, fail every single year. An analytical approach which at heart involves the concept of equilibrium cannot really cope with this most dramatic feature of change, namely failure. One year a company exists, the next it has disappeared.
Economic theory reflects this fundamental point. There is little or no mention of failure. Instead, the emphasis is on what firms need to do not just to succeed, but to take the best possible decision from the entire range of options which is available to them.

In the textbook world, running a business in easy. Almost all the problems which exercise the effort and activity of management are swept away. Rather, they are never mentioned at all. Despite this, economists have an irritating habit of claiming the discipline is the first to explain anything worth knowing about how the business world works.

A single example will suffice to illustrate the point. A successful textbook, by highly respected authors and specifically directed at the interface between economics and business, contains an amazing opening remark. 'The economic analysis of firms,' it is declaimed, ' provides a simple decision rule that managers and entrepreneurs will find useful'. What remarkable insight is about to be vouchsafed to us? It is none other than the truly original concept that 'Any action that adds more to revenue than it adds to cost should be undertaken' and, what is more, 'actions that add more to costs than to revenues should not' . The text goes on 'In the economic theory of the firm this rule is heavily disguised'.

I think that the authors mean by this that you have to spend years acquiring at least first-year degree level maths before you can begin to understand it[1]. Alternatively, we may regard this rule as an illustration of the technique outlined by the British satirical writer Stephen Potter about how to gain the upper hand in a conversation about business. In his excellent book *One-Upmanship*, he describes his 'Economics B' technique as the 'Approach of Utter Obviousness'.

The business school gurus, it should be said as an aside, are by no means immune to the temptation to reduce highly complex problems to a set of easy formulas. Leading figures like Tom Peters offer beguiling sets of simple rules. Business-school thinking has fallen

[1] to be fair, this particular book does get considerably better as it goes along and is one of the very few makes an effort to relate to real world business problems. It is *Economics for Business and Management* by Alec Chrystal and Richard Lipsey, Oxford University Press, 1997

into the trap of standard economics: everyone can use the same rules and prosper. But most of the firms in Tom Peters' *In Search of Excellence* later failed to maintain their excellence, to say the least.

An important defect, a failure I am almost tempted to say, of the equilibrium approach in economics is that it tells us nothing about the time-scale of the process of change from one equilibrium to another, even in the textbook world in which equilibria exist by definition. In reality of course, as we have noted above, they might not exist at all.

But in the comforting world of economic theory, we can think, for example, of a single market in which the price is set so that supply and demand balance. The amount which firms are induced to supply to this market at this price is equal to the amount which consumers are willing to buy. This is one of the most basic features of any economics textbook, even at the most elementary level. As it happens, as we shall see below and again in chapter 5 , even this apparently straightforward concept gives rise to some difficult and awkward questions for conventional economics.

But the relevant point here is that the emphasis in economics is on a description of such equilibrium situations, in which supply and demand balance exactly. If the system receives a shock, a tax being imposed, say, on the product in question so that its price goes up, economic theory describes the new levels at which supply and demand will once more be equal. Both producers and consumers adjust their decisions in the light of the new price. We can then compare the two equilibria, both before and after the tax was imposed.

The theory tells us nothing at all about how long the system will take to move from the old equilibrium to the new one, following the introduction of the tax. And it says nothing at all about the path the system follows between these two equilibria. Is most of the adjustment made rapidly, for example, and the last few steps taken more slowly, or is the path one of entirely smooth progress? Nobel laureate Vernon Smith of George Mason University makes the point clearly in his 2002 Nobel lecture: 'the prediction of the competitive model is silent as to how long it takes to respond to a change..'.

In the context of a very basic market, such as cornflakes or bread, for example, these questions might not really matter in practice. But once we start thinking about more complicated problems, the process by which the system moves from one equilibrium to the next, and the time this takes, become very important issues indeed.

Economics, for example, has a theory of economic growth, growth being the single feature which most distinguishes the market-oriented economies of capitalism from all other previously and actually existing societies. No other system has ever generated such steady growth over such a long period of time. For example, during the period from 1500 to 1800, the world economy was far from stagnant. Gradually, progress was being made towards the Industrial Revolution, which by 1800 had just in fact begun in Britain, then the world's most advanced economy. But the total increase in world output (after

allowing for inflation) is now more in a single year than it was during the entire three centuries from 1500 to 1800.

As long ago as 1969, the then young Cambridge academic Tony Atkinson, now recently retired as Warden of Nuffield College, Oxford, investigated the time scales of various models in economic theory. In particular, he looked at equilibrium solutions in the standard theory of growth, and calculated how long it takes for the system to move to a new equilibrium once a change is introduced. He found that the answer was typically over 100 years. This is not a misprint. One hundred years. In other words, even within the safe confines of economic theory, a system which is in equilibrium will take more than a century to move to a new one once a shock is administered to it. So the system spends a long time in *dis*equilibrium, when things alter and no longer run completely smoothly.

Economic and social systems are essentially dynamic and not static. Even if we were to retain a faith, for that is what it is, in some Platonic idea of equilibrium towards which the economy is moving, most of the time we will experience and observe behaviour not in but out of equilibrium. Firms fail, policies alter, behaviour changes. A theory which is based on describing what the economy is like once it reaches an equilibrium, a static solution, will of necessity give us at best only a partial understanding of the world.

In practice, firms fail for a whole variety of reasons, some of which might be common to a number of firms, such as an economic recession causing a drop in sales, or some might be specific to an individual company, such as losing sales to a more effective competitor. But the proximate reason for failure, the thing which happens just before they fail, is that firms run out of money. They do not have enough revenue to cover costs, and are unable to beg, borrow or steal sufficient funds to fill the gap.

Firms raise revenue by selling their products, whether goods or services. A statement of, in Stephen Potter's phrase, the utterly obvious. Equally obvious is that revenue depends upon how much of the product is sold, and what price it is sold at.

Choosing the price of a product is a challenging and difficult decision for companies. If a firm makes a big enough mistake on this, or even persists with perhaps relatively minor mistakes for a sufficiently long period, it will fail. Many factors need to be considered. The same product need not have the same price at the same time. There might, for example, be discounts for bulk buying, or special terms for either loyal or new customers. And who, for example, are the main competitors, and how might they react to a change in the price of your product? Are there any potential competitors likely to enter your market with a rival offer or, worse, be developing one which is better than yours? How can pricing strategy be used to deal with this more shadowy but often dangerous threat of potential as opposed to actually existing competitors?

In the world of textbooks, the formula which most students of economics learn reduces these and other difficulties to a single phrase: 'set price equal to marginal cost'. The word 'price' is used in its everyday sense. But the phrase 'marginal cost' is a piece of

jargon. ‘Marginal’ is a word which permeates the whole of standard economic theory. It has an entirely specialised meaning in this context. A workaday translation is ‘additional’. So, stripped to its essentials, the phrase ‘marginal cost’ means the cost of producing an additional unit of output.

We might usefully recall here the piece of advice given by economists to decision makers in companies: 'Any action that adds more to revenue than it adds to cost should be undertaken'. If we are able to sell an additional unit of our product, the additional revenue we obtain from it will be the price at which it is sold. ‘Marginal cost’ is the cost of producing an additional unit of the product. So if we are able to sell another unit when the extra revenue is greater than the extra costs we incur in producing it, we should do so. And we should go on doing so until we reach a point where the additional revenue, the price, is equal to the additional cost of doing so. This is what the formula ‘set price equal to marginal cost’ means.

By following this formula, according to the textbook model firms will automatically succeed in making the maximum amount of profit available to them. The formula can be expressed at different levels of sophistication. It can be demonstrated graphically. It can be shown using differential calculus. It might even, for graduate students, involve a knowledge of Brouwer’s extension of Kakutani’s fixed point theorem in the proof of the existence of general equilibrium. The formula is by no means purely of academic interest. It is the principle which underlies a great deal of the activity of regulatory authorities around the world, and is the standard against which potentially anti-competitive behaviour is often judged.

But, regardless of the mathematical language in which it is couched, it is not our concern here to prove that proposition that the formula does indeed lead to a firm making not just a profit but the maximum amount possible. Any non-economist reader who wishes to understand this can readily consult any one of a vast range of textbooks. Rather, our concern is to discuss the formidable difficulties of following this rule in practice. Many factors which can and do lead to failure in reality are entirely missing from this way of thinking.

The single most important difficulty with this rule of how to set price is the amount of information, of knowledge which the firm is assumed to have. The company is presumed to understand exactly how many sales it will be able to make at each particular level of price. In the jargon of economics, the firm knows its demand schedule, or curve. From the other aspect, the firm is required to know its cost curve, or how the costs of production vary with the level of production.

These may seem at first sight to be innocuous assumptions, but the practical difficulties of obtaining this information with certainty are enormous. Consider first of all the demand curve, the information which in principle will tell the firm how much of a particular product people will buy at any given price.

Two enormous, related industries, advertising and consumer market research, have developed which attempt both to understand the complicated motivations of consumers, and to try to shape them in the interests of their client, the producer. If the process of running a business were easy, if competitors were completely obvious, if all the factors which attract consumers to a product were known, if the reaction of sales to price were understood exactly, advertising and market research would scarcely be necessary.

The very existence of these industries on this scale is testimony to the pervasive nature of uncertainty in business, even on such a basic concept as setting the price of the product. And when uncertainty exists on such a scale, the potential for getting things wrong, for failing, exists along side it.

It is worth mentioning at this point a distinction which is made in much technical work between the concepts of risk and uncertainty. In casual, everyday English, the terms might be used more or less synonymously. But in their specialised usage, particular within economics, there is a difference to be drawn. Risk refers to situations in which the outcome cannot be known with certainty, but the probability of any given outcome *is* understood perfectly. A simple example would be a toss of a fair coin. There is a 50-50 chance of it being either heads or tails. If we are gambling on the next toss being heads, there is a risk that we lose our money if it turns out to be tails. But we know precisely what the chances are. Uncertainty, in its strict sense, refers to situations in which the probability of the various outcomes is itself unknown. So, for example, we might place a bet on the Earth being visited by creatures from outer space in the next ten years. Here, no one knows for certain what is the true probability of this happening. Almost everyone would agree the chances are very small, but are they 1 in 100 million, 1 in 100 billion, or are they literally zero?

Modern economics deals extensively with the concept of risk. Risk involves quantifiable probabilities. This makes it amenable to mathematical manipulation, a practice dear to the hearts of economists. But the great economists of the past wrote much more about situations involving genuine uncertainty, regardless of their own ideological persuasion.

Keynes, for example, believed in more extensive government involvement in the economy. Equally, however, he wrote a great deal about situations which were uncertain rather than merely risky. Frank Knight, who founded the free market school of economics at the University of Chicago, and who Milton Friedman described as ‘one of the most original and influential social scientists of the twentieth century’, also thought instinctively about uncertainty rather than risk. His doctoral dissertation in 1921 was indeed the seminal piece of work which made this distinction. Interestingly, in the context of this book, Knight famously responded to his own rhetorical question ‘how far is life rational’ with the blunt answer ‘not very far’.

In practice, of course, the two concepts almost always blur into one another. We rarely face in practice situations in which are clear cut as the outcome of the toss of a fair coin. Equally, it is unusual for us to have no idea whatsoever about the likely distribution of possible outcomes. The precise mix will depend upon circumstances.

Returning to the problem of companies grappling to understand their customers, both actual and potential, an immense amount of effort and money is expended by large firms in trying to understand the shape of the demand curve. Focus groups, sophisticated consumer surveys, complicated statistical modelling of data, are all used in an attempt to reduce a fundamental uncertainty confronting any business. Namely, what happens if I change the price of my product? But despite the striving for knowledge, uncertainty persists.

Decision makers in companies are particularly uncertain about how their competitors might react to any changes in price which they might make. Reducing prices to boost sales might make sense, but not if your key competitors follow suit, and the end result is that everyone has to charge a lower price and hence make less of a profit margin on sales. Judging how competitors might react takes us to the concept of game theory, an esoteric idea made famous to the general public in the film *A Beautiful Mind*. The film tells the story of the eccentric American academic John Nash, who invented one of the fundamental propositions of game theory. But game theory is such an important idea, both in its own right and in terms of what it can and cannot tell us about failure, that it merits a prolonged discussion of its own in chapter 5.

The uncertainty which firms face in understanding the nature of the demand schedule for their product is pervasive. As I was writing the first draft of this chapter in 2003 in the agreeable winter climate of New Mexico, the same day's newspapers carried two such stories.

At what appears to be a mundane level, but all the more important for being so normal, were the stories in the American press about the Japanese vehicle manufacturers' strategy for capturing part of the lucrative US market for pick-up trucks. This is a highly profitable industry, with US sales of just over 2 million a year bringing in almost $20 billion profit. Until recently, the large American manufacturers faced no challenge in this market from the Japanese. Now, a dangerous new competitor is starting to roam across the pastures.

In early 2003, Ford launched their redesigned F-150 pick-up, the most popular vehicle in America in this market for over twenty years. The very next day, Nissan unveiled their first full-size pick-up, the Titan. The Nissan Motor CEO was quoted as saying 'We have done our homework and know what the big pickup buyer wants'. But the General Motors' Vice-President in charge of pick up trucks thought otherwise. Gary White conceded that in the 1970s Japanese cars made inroads into the US market because 'we had the wrong products and we had poor quality'. He insisted that this time it was different, a view echoed by the Chrysler Group market research director Dave Bostwick: 'We paid too little attention and let these guys in the door in the past, but this time I think everybody's ready for them'. But GM overall product chief Bob Lutz was prepared to concede that a danger existed: 'The risk for us is if consumers prefer Nissan styling and their power trains'.

So here were some of the largest industrial companies in the world, each with enormous experience of the market, each spending vast amounts on trying to discover what their individual demand schedules look like or, in plain English, how many pick-ups they would sell in 2003. Nissan and the other Japanese firms thought they would do well and gain significant market penetration. General Motors and the American firms thought not. Both of them could not be right. One or the other of these competing groups, or maybe both, had an imperfect knowledge of the demand schedule which they face. In advance, before the event, no-one knew for sure who was right and who was wrong.

It is worth repeating the quote given by Bob Lutz at General Motors: 'The risk for us is if consumers prefer Nissan styling and their power trains'. How might we expand and deconstruct this statement, possibly prepared at considerable expense through the public relations advisers? The following seems a fair attempt: '*If* it turns out to be the case that consumers like what the Japanese have to offer more than the stuff they are used to buying, then they will buy the Japanese product. But if they don't, then they won't. And I don't really know whether they will or not. I can only say that there is a risk of it happening'.

My intention is not to mock Mr Lutz's statement, but to pay him due honour for recognising so clearly and so openly the fundamental uncertainty about the demand schedule which companies face. And if General Motors, with its vast resources, does not really know, if GM cannot really tear down completely the veil of uncertainty which shrouds the future, who can?

Certainly not the airline industry, which is the feature of the second of the stories in the newspapers. The world's airlines lost $13 billion in 2002, an improvement upon the $18 billion loss in the aftermath of September 11. And despite enormous increases in the number of air travellers in the second half of the twentieth century, the combined profit made by airlines over the 1955-2000 period was less than the losses they incurred in just the two years 2001 and 2002. These losses are not surprising, once we know that the attack on the World Trade Centre took place. But this was, admittedly in dramatic form, an example of the sudden changes in the external environment which can buffet even a well-run company.

An example from the European market emphasises the point. In the last few years, the cartel of national air carriers which dominated the market for so long has been challenged aggressively by low-cost, no-frills airlines. The major airlines were eventually forced to respond, both in terms of their own pricing strategies, and in terms of setting up their own, dedicated low-cost operations. The Dutch company KLM set up Buzz, which incurred multi-million pound losses in each of its three years of existence before it was sold in January 2003 to another low cost airline, Ryanair of Ireland. The price for the company, after deducting its cash balances, was a mere £4 million. As Gert Zonnenveld, airlines analyst at WestLB Panmure said 'It sounds very easy to say "let's start up a low cost airline and make lots of money". But in reality it's not easy, and it takes quite a few years – if you are lucky'.

There have already been numerous casualties, numerous firms which have failed in this sector. Ciao Fly began operations in the summer of 2002. Described as 'a Swiss company run by a German, operating an Italian airline to Britain', it lasted just six weeks. Swedish entrepreneurs launched Goodjet in April 2002. Visitors to its website are now directed to the site of a firm of accountants, where a sad little note reads 'Goodjet was declared bankrupt on 17 January 2003 by the Gothenburg district court' Even the doyen, if the term might be used in this context, of European low cost flying, Easyjet, is currently making investors nervous as its seat occupancy falls sharply.

Suppose, however, business is as easy as the economics textbooks pretend it is. The whole process of trying to penetrate the uncertainty which surrounds identifying all your key competitors, which surrounds quantifying the impact of pricing strategies both by you and by your competitors on how much you will sell, understanding the effect of advertising and promotion, solving the problem of distribution to actually get your products in the shops for customers to buy – all this process, meshed in uncertainty, is assumed away in the textbooks. You know what your demand schedule looks like. In other words, you know how much of your product will be bought at whatever particular price you decide to charge.

Suspending for a moment the disbelief that all this information has appeared as if by magic, if we know the demand curve, we know how much extra revenue we can get if we are able to sell one unit of output in addition to our current level of output. And we know this for all possible current levels of output.

The formula 'set price equal to marginal cost' means, we recall, that we should produce exactly that level of output at which the additional revenue obtained from selling an extra unit of output is equal to the additional cost of producing it. So we also need to discover how our costs of production vary with the level of our output.

Unlike the question of finding out about how the demand for the company's products might vary as things like price are changed, this task does not seem so unreasonable. The process of production is internal to the firm. The company is not faced to the same extent by the vastness of the external world and the associated uncertainty, as it is when it thinks about how its products might be sold.

Even so, the discovery of how costs vary over different ranges of output might not be completely straightforward. For example, the British government recently brought in a scheme to assist in meeting pollution reduction targets under the Kyoto protocol, specifically for the reduction of carbon dioxide production. Companies can participate in what is known as the emissions trading scheme. Under this, the firm must undertake a binding agreement to reduce its use of energy - and hence its emissions of carbon dioxide - by a specified amount by a certain date. In return for this commitment, it receives allowances to produce emissions at the agreed level.

The real innovation, however, was that these permits can be traded. If a company finds that is proves easier to meet its target than it originally thought, it will be able to sell

some of its permits. Its emissions will be below the level which it contracted to 'produce'. Equally, if a firm miscalculates in the opposite direction and discovers that it is much harder and more expensive to cut its emissions, it can buy permits which allow it to produce emissions at a higher level than its original target.

Now, if firms knew exactly how the costs they incur in reducing the energy content of their production processes varied, from their point of view, these allowances would not be required. No firm would make a mistake. But the scheme has proved very popular in practice, with very large companies being active participants. Like Sherlock Holmes's dog that did not bark, this tells us a great deal. It tells that the companies which use large amounts of energy which participate in the trading scheme realise that they *can* make mistakes about their knowledge of how their own costs evolve.

Just as we imagined earlier that a firm knows its demand curve, suppose as well that it understands its cost curve. In other words, it knows what its costs of production are at all possible levels of output. The management can then work out, or get some accountants to do it for them thereby adding even more to the costs, how the costs per unit of output change as the scale of production varies. For it is this change which is the key to the concept of 'marginal cost'. Marginal cost means the cost of producing an additional unit of output. So, in principle, we can calculate, starting from every single possible level of production, the cost of producing the next unit of output.

Perhaps we are finally there. Perhaps we can now put into practice the formula: set price equal to marginal cost. We know the cost of producing an extra unit of output over and above any particular level of output. We know the revenue we can get from selling an extra unit of output over and above any particular level of sales. So we find the point at which the two are equal and – abracadabra! – we know not only the exact amount to produce, but we know that we will make the maximum possible amount of profit, given our demand and cost curves.

But the practical difficulties are not yet exhausted. In many situations, the higher the price which is charged for a product or service, the less of it will be bought. So, in the jargon, the demand curve slopes downwards. If we draw a chart with the level of sales on the left hand axis and the price on the bottom axis, the demand curve will slope down from left to right across the chart. At low levels of price, demand will be high and at high levels, it will be low.

We can immediately think of exceptions to this. In financial markets, for example, it is often the case that when the price of a share rises, traders buy more of it and not less, since they imagine that the price will continue to rise and so they can sell it later at a profit. This motivation was particularly strong during the stock market boom of the late 1990s. In consumer markets where high fashion plays a key role, a higher price might serve to attract rather than deter those individuals who enjoy conspicuous consumption. At the other end of the scale, a second hand car which seemed to be priced too cheaply might very well not sell at all, on the grounds that if it is so cheap there must be something wrong with it. But in general, we seem on reasonably safe grounds if we

assume that the demand curve slopes down, that the lower the price, the more will be sold.

We are almost, but not quite, there. We need to know a little bit about the shape of the marginal cost curve, how marginal cost varies at different levels of output. The standard assumption made in teaching generations of students Economics 101, the introductory course to university economics, is that as the level of output increases, marginal cost first of all falls but eventually rises.

At first sight, this might not seem unreasonable. We might imagine a car plant being built. The cost of producing just one car is enormous - all the costs of the plant are allocated to this single car. When two are produced, the costs of producing the second car are dramatically less. So the marginal cost of making the second car is much less than that of the first. And as more are produced, the marginal cost of each continues to fall. But at some point, the capacity of the plant will be reached, and producing an additional car will need a whole new plant - costs per unit of output will once more rise. Even before then, there may well be bottlenecks in capacity. The workers may require overtime payments at premium rates, for example, or a particular segment of the production line might require upgrading. Such factors will lead to an increase in the cost of making that extra car, even before a whole new plant is needed.

Focus now on the part of the marginal cost curve where it starts to slope upwards. At levels of output lower than this, price must be higher than marginal cost, otherwise we make no profit. Indeed, if marginal cost is higher than price over the entire downward sloping part of its range, there is no point in producing at all, because each additional unit of output makes a loss. So when marginal cost starts to rise, at some level of output, having been less than price, it becomes equal to price. This is the magic spot, where price equals marginal cost and by definition no profit is made on that particular unit of output.

So there we are! We can all go off and happily maximise our profits like all good capitalists. Or can we? The problem is that in substantial sections of the economy, it is not at all clear that the marginal cost curve ever gets to the point where it starts to slope upwards rather than downwards.

An obvious contemporary example is when a company offers subscription services over the world wide web. Once the process of debiting credit card payments has been automated, the material set up in an appropriate form, and the initial teething problems of running the system overcome, the marginal cost, the cost of selling to an additional subscriber, is virtually zero. The marginal cost curve never slopes upwards. On the contrary, it approaches zero. A company which charged marginal costs in these circumstances would soon go bankrupt, because it could never recoup the costs involved in setting up the service in the first place.

More generally, firms strive to avoid a situation in which the marginal cost curve starts to slope upwards. The managers may not see it in precisely the jargon of economics. But a

great deal of their time and effort goes into trying to provide their product more efficiently. Companies carry out research and development, they invest in new equipment, they look for cheaper and better sources of the supplies they need to produce their product, they evaluate their promotional and marketing material.

An important and mysterious effect which seems to occur in many industries is that the costs of producing a given level of output fall, the greater the cumulative level of production has been. In other words, simply by producing more of a product over time, the process of production becomes more efficient. As long ago as 1936, an innocuously title article 'Factors Affecting the Cost of Airplanes' provided empirical evidence that the direct labour cost of manufacturing an airframe fell by 20 per cent with every doubling of cumulative output. Many further studies have found a similar qualitative effect in a wide range of industries, although the exact gains in efficiency differ.

All these activities mean that the marginal cost curve is not the fixed, immutable line of economic theory. It is the outcome of a dynamic process, and it changes and evolves over time. The quest to produce a better product more efficiently is the driving force of all successful businesses.

Interestingly, this was a theme emphasised again and again by some great economists of the 1920s and 1930s. Allyn Young, for example, who held professorships first at Stanford, then Cornell and finally Harvard before dying tragically early at 53, and who supervised Frank Knight who we met briefly earlier in this chapter. Young wrote a classic article in 1928 in what was then the world's leading academic economics journal, showing how falling marginal cost curves undermined the traditional economic concept of equilibrium. Joseph Schumpeter, another Harvard professor who we will meet again in the final chapter emphasised that the key features of actually existing Western economies were change and discontinuity and not equilibrium.

Piero Sraffa, a Fellow of Trinity College Cambridge, published an article in 1926 arguing that in very few industries indeed did marginal cost curves ever slope upwards. He ridiculed the importance which conventional economics attached to this concept, writing that 'Business men, who regard themselves as being subject to competitive conditions, would consider absurd the assertion that the limit to their production is to be found in the internal conditions of production in their firm, which do not permit of the production of a greater quantity without an increase in cost. The chief obstacle against which they have to contend when they want gradually to increase their production does not lie in the cost of production—which, indeed, generally favours them in that direction—but in the difficulty of selling the larger quantity of goods without reducing the price, or without having to face increased marketing expenses'

.

Sraffa, like many Cambridge dons of his generation, held a lifelong political attachment to Stalinism. But he did not allow this to influence his own conduct. For example, in 1945, after the atomic bombs had been dropped on Hiroshima and Nagasaki, he invested all his money in Japanese government bonds, whose price then was obviously very low indeed. He judged that capitalist Japan would recover, despite the apparently devastating

blow it had just suffered. And indeed it did, surpassing its pre-war levels of output by 1953, and going on the another thirty years of stupendous growth. Sraffa made a fortune.

How *do* firms actually behave when it comes to deciding on the price of a product? Immense efforts have been made in economics to develop and make more precise and rigorous the whole corpus of theory associated with the phrase 'price equals marginal cost'. The theory is distilled and passed down to generations of students, at all levels of the discipline.

Yet the empirical evidence suggest that firms rely in practice much more on simple rules of thumb. The classic study was carried out by two Oxford economists, Hall and Hitch, as long ago as 1939. Hall and Hitch questioned the owners of 38 firms and found that rather than profit maximising by producing where marginal cost is equal to marginal revenue, the majority in fact used cost-plus pricing. The entrepreneurs added up their costs of production and then added what they thought was a fair profit margin. A few took into account what the market price was but none was able to calculate marginal costs and revenues.

Numerous studies since that time have confirmed the validity of the view of the world presented in this pre-World War Two research. For example, as the final version of this chapter is being prepared, yet another such study has been published in the *IBM Systems Journal*, which argues not only that 'Historical evidence suggests that cost-plus pricing has been in use at least since the end of the 18th century', but that the principle is widely used in the pricing of services as well as products. Of course, there are many complicating factors, such as the price charged by competitors, or the price which might be charged by competitors if your price is set at a particular level. But the basic finding remains valid.

The use of fairly simple rules to guide behaviour is in fact a rather sensible way to behave when confronted with a situation which is both enormously complicated and massively uncertain. The process of operating a business, far from being easy, is a very demanding activity.

The capacity of firms to deal with market situations in a cognitive sense, their capacity to process information and turn it into knowledge, is small compared to the sheer dimension of the problems which confront them. Companies can never deal completely with the complexity of the real world. The uncertainty which shrouds the future is not so much a veil but an iron curtain. In the current state of scientific knowledge, it cannot be penetrated. There is ample opportunity at any point in time for any firm, no matter how large, to fail.

So far, we have looked at firms operating in the uncertainties of the market. At both how firms have fared in practice, at the failure rates of the corporate giants of capitalism, and at what, if anything, economic theory has to say about such failures. Governments are different. In the Western world, the state has much more guaranteed continuity.

Nevertheless, failure extends from the realm of economics into the world of public policy. It is to examples of failure of government policy to which we now turn.