10
Strategic innovation

John Beath
University of Bristol

Yannis Katsoulacos
University of Liverpool

David Ulph
University of Bristol

Chapter 10. *Mathematical Models in Economics*. MOL Bacharach, MAH Dempster, and JL Enos, eds. University of Oxford (1990).

10

Strategic innovation

John Beath
University of Bristol
Yannis Katsoulacos
University of Liverpool
David Ulph
University of Bristol

Abstract

A central question in the theory of strategic innovation is whether technological competition between initially asymmetric firms favours the incumbent (so producing *persistent dominance*) or its rivals (in which case we have *action/reaction*). The existing literature contains a variety of models, some generating the first outcome and others the second. These models are built on such different assumptions that it is often difficult to see which features of the models it is that determine the conclusion reached. In this chapter we develop a framework within which we can both set out much of the existing literature and then extend it, getting, in the process, a clearer idea of the factors that matter in determining each of the two outcomes.

In §2 we review the standard model of a single innovation under certainty. Here the only force that affects a firm's decision is the *competitive threat* that it faces from its rival. This threat is measured by the fall in future profits that it will suffer if one of its rivals innovates ahead of it. We show how the outcome of the competition to innovate depends on the size of the innovation and the nature of competition and also that such single innovation models can be very misleading when we think of firms facing a continual pressure to innovate. We model this through a *sequence* of innovations. In a sequential framework, the distinction between product and process innovation becomes vital and the conclusions reached under process innovation are almost precisely reversed under product innovation.

A major limitation of models with certainty is that firms that know that they will be losers commit no resources to R&D. This is particularly important in a sequential framework, since R&D costs incurred in future races enter the calculations of the profitability of current races. In §3 we introduce a model of a single innovation in which there is uncertainty over the date of discovery. This brings into play a second incentive to innovate—the *profit incentive*—which is reflected in the difference between a firm's current profits and those it will earn if it innovates. We show that the outcome of the race can be determined by knowing the magnitude of these two incentives for each firm, and that the effect of the assumptions employed in most of the existing models is to artifically constrain the relative magnitudes of one or other of these incentives.

In §4 we examine sequential innovation under uncertainty, and argue that another factor that now has to be taken into account is reflected in the combined effect of information dissemination and learning-by-doing, since this determines whether followers can *leapfrog* the leader, or can only gradually *catch up*. We show how this, in conjuction with most of the factors discussed in §2, determines the nature of the outcome. While the possibility of leapfrogging generally leads to action/reaction, models with a catch-up structure generate a variety of outcomes depending on the nature of competition and the rate of technical change.

1. Introduction

Some of the major advances in mathematical economics over the last twenty years have occurred in game theory, and, as a result, game theoretic models are being used in many areas of economics such as macroeconomics where strategic behaviour had not previously been analysed. The predominant area of application, however, remains the theory of imperfect competition and the analysis of industrial structure.

A particularly important type of non-price competition between firms is *technological competition*. This is where firms compete to introduce new processes or products that will give them some competitive edge over their rivals (including potential entrants), or, more defensively, will at least prevent their rivals from getting too great an advantage over them. There have been considerable advances in the analysis of this particular kind of strategic competition—especially over the last ten years—and in this chapter we wish to review and synthesize some of these developments.

A central question in this area is whether this dynamic process of competition is one in which incumbents maintain their position of supremacy (or indeed pull further ahead of their rivals) or whether they are overtaken by some rival whose incumbency is itself only short-lived. The former outcome we call *persistent dominance*, while the latter we describe as *action/reaction*, though it is perhaps more familiarly known by the Schumpeterian terminology of 'creative destruction.'

There are a number of reasons for being interested in this question. The first is that there is a great deal of concern expressed by governments and others about the relative position of firms in 'their' country—particularly in key 'high-tech' industries—and an understanding of the process of dynamic competition provides a useful framework within which the basis for this concern can be analysed and possible policy options examined. The second is that such an understanding will clearly give us a framework in which the evolution of industry structure is endogenous to the process of dynamic competition: for whether dynamic competition produces persistent dominance or action/reaction will have a major bearing on how competitive industries are, while we would expect that in turn the outcome of the process of dynamic competition will be affected by the existing market structure. As we will see, when industry structure is determined in this dynamic fashion, many of the prescriptions of static competitive theory can be overturned.

There are many factors that influence the absolute and relative amounts of resources that firms devote to R&D. One important factor is the technological opportunity that firms face, which is, in turn, related to the scientific base of the industry (Rosenberg 1974). This factor is particularly important in understanding inter-industry levels of R&D expenditure. There is also some evidence that due to the existence of *learning-by-doing* in R&D, successful firms find it easier to make further advances than less successful ones (Phillips 1971). These two factors clearly influence the *ability* of firms to innovate. Of course in analysing strategic innovation the major focus is clearly going to be on the *incentives* of firms to innovate, and here we expect that the current and anticipated structure of the market and the nature of competition between firms will be important in determining both the absolute and the relative incentives of firms to innovate. It is important to appreciate, however, that those factors which affect the ability of firms to innovate will typically also influence their incentives to do so. Thus while learning-by-doing effects clearly enhance the ability of incumbents to innovate, they will also give them incentives to maintain their leadership and so retain these benefits, while simultaneously giving rivals incentives to stop incumbents getting too far ahead and exploiting these advantages. So one should not conclude that the presence of these effects leads in any automatic way to persistent dominance being the outcome of strategic innovation.

Given the variety of factors at work in determining the outcome of strategic dynamic competition amongst firms, it is perhaps not surprising that the existing literature on this topic is rather confusing, comprising numerous different models yielding conflicting predictions about the outcome of the innovative process. Since these models are typically built on very different assumptions, it is often difficult to see just what factors are responsible for producing these differing predictions. The major purpose of this chapter is to provide a unifying framework in which most of these models can be set, and to go some way towards providing a more general model encompassing many of these more specific models.

While we will try to cover much of the literature on strategic innovation, there will be some important areas we do not examine. For example, throughout the chapter we will focus exclusively on models where the number of firms engaged in R&D competition is fixed in advance at two. Consequently we will have nothing to say about the role of innovation in generating entry barriers, so making the number of firms endogenous. We ignore this issue because the work of Dasgupta and Stiglitz (1980*a*,*b*, 1981) has provided a good understanding of it, and the material is well covered in the excellent review by Dasgupta (1986).

The structure of the chapter is as follows. In the next section we will examine a widely used class of models of innovation under conditions of certainty. We will start by considering the simple case of a single innovation and then show that while this gives some essential insights, single-innovation models can be very misleading, and that a proper understanding of the issues we are concerned with can only come from models that analyse sequences of innovations. The second part of this section will be devoted to such sequential models.

The models of certainty considered in §2 capture one important incentive to innovate—what we will call the competitive threat—but there is another important incentive that affects innovation decisions—the profit incentive—and this arises when innovation is modelled as a game of timing. While this kind of game can arise under both certainty and uncertainty, and while games with uncertainty need not be games of timing, much of the literature analysing innovation under uncertainty treats the uncertainty as arising over the timing of the innovation. Section 3 examines a general model of such games for the case of a single innovation, and shows how an understanding of these two incentives goes a long way to explaining the outcome of races, and to unifying various models in the literature. Finally, §4 examines sequential models under uncertainty—models which combine the insights from the various incomplete models considered so far.

2. Certainty

2.1 A single innovation

To fix ideas we will start with the following model. There are two firms producing an identical product under constant returns to scale technologies. As a result of earlier innovation successes, the firms' unit costs are different, with $c_1 > c_2$, where c_i is the unit cost of firm i. A new technology is discovered which lowers costs to $c_3 < c_2$, and an infinitely lived and completely effective patent for this technology is put up for auction. The patent is awarded to the firm which values it most highly, this firm having to bid an amount equal to the value placed on it by the other firm. The losing firm pays nothing. If the current low-cost firm wins the auction we will say that we have persistent dominance, while if the current high-cost firm wins we will have action/reaction. Which of these two outcomes will occur?

To answer this question we have to specify the nature of competition in the product market. For the moment suppose it is Cournot, and let $\pi(\alpha,\beta)$ be the (present value of) profits in a Cournot equilibrium of a firm whose unit costs are α while those of its rival are β. Then if the current incumbent wins the patent its profits will be $\pi(c_3, c_1)$, while, if the rival firm wins, the incumbent's profits will be $\pi(c_2, c_3)$. Thus the value that the incumbent places on the patent, and hence the maximum bid it is willing to make, B^i, is

$$B^i = \pi(c_3, c_1) - \pi(c_2, c_3). \tag{2.1}$$

Similarly the maximum bid the non-incumbent (or challenger or follower) is willing to make, B^f, is

$$B^f = \pi(c_3, c_2) - \pi(c_1, c_3). \tag{2.2}$$

These maximum bids reflect one of the major incentives driving firms to innovate—an incentive that we will call the *competitive threat*—since they measure the fall in its future profits that each firm will suffer if its rival were to successfully innovate ahead of it.

If we now let $\sigma(\alpha,\beta)$ be the industry profits made in the Cournot equilibrium, when the unit costs of the low-cost firm are α, those of the high-cost firm β, then, ignoring boundary cases where $B^i = B^f$, it follows from (2.1) and (2.2) that we will have persistent dominance iff

$$\sigma(c_3, c_1) > \sigma(c_3, c_2). \tag{2.3}$$

To understand when (2.3) is satisfied, we need to know how industry profits vary with the costs of the high-cost firm. There are two forces at work. First, holding quantities, and hence the price fixed, the lower are the costs of the high-cost firm the greater are its, and hence the industry's, profits. But of course quantities and price will not remain fixed as the costs of the high-cost firm vary. For we know that in a Cournot equilibrium the price depends on the average of the costs of the two firms, and so as the costs of the high-cost firm rise so too does the price. This increases the profits of the low-cost firm until ultimately the cost differential is so large that the low-cost firm can earn monopoly profits, which dominate the industry profits in any of the other Cournot equilibria.

Intuitively we expect that when the costs are equal, a small increase in the costs of one of the firms always lowers industry profits, since there is little scope for the low-cost firm to exploit this cost differential, and so the first of the two effects mentioned above will dominate. However, once the cost differential is large, the output of the high-cost firm will be so low that a further increase in its costs will have a negligible impact on industry profits through the first effect, while the low-cost firm will have more scope to exploit these higher costs and move industry profits closer to those obtained under monopoly. If we therefore hold the costs of the low-cost firm constant, and plot industry profits as a function of the cost difference between the two firms, we expect industry profits to be initially falling when the gap is small, then rising until the gap becomes large enough for the low-cost firm to choose the monopoly level of output.

This can readily be confirmed by examples. For example, when the industry demand is iso-elastic, then if ε, $0 < \varepsilon < 1$, is the inverse elasticity of demand, and

$$g = (c_h/c_3) - 1 \tag{2.4}$$

is the percentage gap between c_3 and the unit cost of a high-cost firm with costs c_h, it can be shown that if we plot σ as a function of g alone we obtain a relationship like that shown in Fig. 2.1, where $\overline{g} = \varepsilon/(1-\varepsilon)$ is the critical gap at which the low-cost firm can earn monopoly profits π^m, and $\underline{g}$, which depends on ε, is the value of g at which industry profits start to increase.

If we let g_i be the gap corresponding to the costs c_i, then, since $g_1 > g_2$, we have the following results.

(i) $g_2 \geq \overline{g}$ (Drastic Innovation): Here the innovation is so great that whichever firm wins will be able to act as a monopolist. In this case industry profits are the same for both firms and the auction is indeterminate.

(ii) $g_1 \geq \overline{g} > g_2$: This case would arise if in the initial position the cost difference between the two firms was sufficiently large that the incumbent could act as a monopolist. The follower should therefore

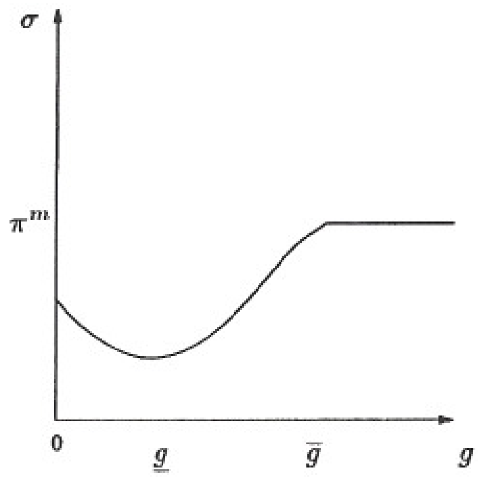

Figure 2.1

be thought of as an entrant. It is clear that in this case (2.3) is satisfied and we will have persistent dominance. Notice that the essential condition for this case to arise takes the form of a restriction on the initial position of the firms, and that, as long as it is not drastic, the size of the innovation is irrelevant.

(iii) $g_2 \geq \underline{g}$: Here the innovation is sufficiently large that the current incumbent will be able to exercise more market power than the follower, whatever the size of the initial gap between them. This will therefore generate persistent dominance.

(iv) $g_1 \leq \underline{g}$: For this to arise both the initial gap between the firms' costs and the size of the innovation must be small. From Fig. 2.1 it follows that the inequality in (2.3) will be reversed, and so in this case we will have *action/reaction*.

This list of cases does not exhaust all the possibilities, but in those we have omitted the outcome will depend on the precise magnitudes of g_1 and g_2, so not much can be said at this level of generality. We also want to note that a great deal of the discussion in the literature has focused on case (ii), where persistence is particularly important since it implies the maintenance of monopoly power.

It is also interesting to note that if we replace the assumption of Cournot competition with that of Bertrand, then industry profits are clearly going to be strictly increasing in g, $\forall g < \overline{g}$, thus generating persistent dominance in all cases. Thus the more competitive static behaviour can give rise to the less competitive dynamic outcome, a point noted by Vickers (1986).

The model we have been considering is obviously highly stylized. In particular, the auction feature of the model has two crucial limitations.

(i) Because the date of the auction is fixed, all that matters is a comparison of the future profits from winning, with those obtained from losing—the competitive threat. However, if R&D expenditure could affect the *date* of innovation, then, as explained more fully in the next section, a second force operates—the profit incentive—and this involves a comparison of future profits if successful with those being currently earned.

(ii) The unsuccessful bidder commits no real resources to R&D.[1]

Katz and Shapiro (1987) have a model which overcomes the first limitation, but suffers, as must any model of uncertainty, from the second. Rather than explore their model here, we will postpone discussion of models of timing to the next section where we also introduce uncertainty. In such a case both firms will typically commit resoureces to R&D. All the essential features of the Katz and Shapiro model appear in this more general framework.

Despite these limitations, the model we have been considering so far captures in a very direct way the competitive threats that will be an essential part of any model of strategic innovation, and has the important property that the amount of R&D done by the successful firm will be strongly affected by the valuation placed on the innovation by the unsuccessful firm. Moreover, it yields crucial insights into the effects that factors like the size of the innovation, the initial difference between the firms, and the nature of market competition will have on the outcome of the innovation process[2]

In the remaining sections of the chapter we will consider more satisfactory models which do not rely on the auction story, and in which, as a consequence, both firms will in general commit resources to R&D. In the remainder of this section, however, we will focus on another major limitation of the model—the restriction to a single innovation.

[1] This is the equilibrium that would arise in a two-stage auction in which the players had to pay what they bid. and in which the firm that valued the patent most highly had to bid first. Clearly if this firm could choose whether to bid first or second it would choose to bid second and then win the patent for virtually nothing. If instead we allowed the current incumbent to have a dominant move advantage and be able to choose whether to bid first or second. it would always choose to bid second, in which case it obtains the patent virtually free when it values it more highly, but the challenger has to pay the incumbent's valuation in order to win. This seems to give too much of an advantage to the incumbent. particularly in a sequential setting.

[2] So far, however, we have said nothing about product innovation. We will deal with this later in this section when we have examined sequential models.

2.2 Sequences of innovations

There are two directions in which we can go to remove this assumption. In the first we can introduce the idea that at any one time an incumbent might face the threat of innovation from a number of different new technologies/new firms. The case of multiple technologies, for example, is analysed by Dasgupta (1986) and, as one would intuitively have expected, the possibility of dominance is weakened, for even if the incumbent starts from an initial position of monopoly power, one has to impose some relatively strong assumptions about increasing returns to operating multiple innovations to guarantee the incumbent will win all the innovations.

However, for our purposes the more interesting direction to pursue is that of there being a sequence of single innovations. This is important because if, as we have seen, the current position in which firms find themselves is a key determinant of innovation success, then an important reason for winning the current race will be the strategic advantages conferred in future races. From this point of view there are three crucial weaknesses in single-innovation models. The first is that because whoever wins the race will, by assumption, be the leader forever, firms cannot contemplate the possibility that a lead lost today may be regained in the future. If this were a possibility then there may be less to be gained by incumbents maintaining their lead, and action/reaction becomes more likely simply because it could now be a correctly anticipated future outcome. The second factor is that in a single innovation the lead gained by the winner is fixed, and so firms cannot contemplate the possibility that success now leads to even greater success in the future. This latter fact could give current incumbents even greater incentive to win, but equally strengthen the desire of the follower to prevent incumbents opening up too great a gap. Finally, in a sequential framework firms have to contemplate not just the future profits that flow from gaining particular strategic positions, but the costs of maintaining those positions. If winning the current race makes rival firms even more determined to stop the incumbent winning future races, then, from the incumbent's point of view, winning may be more expensive than losing. However, if winning the current race so discourages rivals that they reduce their effort in the future, then winning will mean that the future bid costs will be lower if the incumbent wins than if he loses.

Because of these drawbacks, we feel that the whole issue of persistent dominance versus action/reaction is better posed in the context of sequential rather than single innovations, and there is now something of a literature on such sequential models.

The most immediate extension of the model considered so far is in Vickers (1986), where now there is a sequence of T auctions, taking place at times $\tau = 1, \ldots, T$, each introducing a technology of successively lower cost

$$c_{\tau+2} < c_{\tau+1}, \quad \tau = 1, \ldots, T.$$

To determine the outcome of this sequence of innovations, notice that when the auction is being held at the start of period τ, the firm that was successful in the previous period will enter the auction holding the patent on the technology with costs $c_{\tau+1}$, while the lowest cost technology on which the rival firm holds a patent will be some c_k, $k = 1, \ldots, \tau$, which particular one depending on the outcome of earlier auctions. To determine the outcome of the auction at τ, let $V(c_{\tau+2}, c_j)$ be the present value at τ of all current and future profits minus any successful bids which it is (correctly) predicted will be made in future auctions by the firm that has just won the τ^{th} auction and so has least-cost technology $c_{\tau+2}$, while the lowest cost technology on which the rival firm holds a patent is c_j, $j = 1, \ldots, t+1$. $V(c_j, c_{\tau+2})$ defines the analogous value function for the other firm. Notice that, because there are no future auctions,

$$V(c_{T+2}, c_j) = \pi(c_{T+2}, c_j); \quad V(c_j, c_{T+2}) = \tau(c_j, c_{T+2}).$$

Then using the same reasoning as in the single-period case, the maximum bid made by the incumbent at τ is

$$B_\tau^i = V(c_{\tau+2}, c_k) - V(c_{\tau+1}, c_{\tau+2}), \tag{2.5}$$

while the maximum bid of the follower is

$$B_\tau^f = V(c_{\tau+2}, c_{\tau+1}) - V(c_k, c_{\tau+2}). \tag{2.6}$$

With these definitions we can now in principle

(i) determine which firm wins the T^{th} auction in every conceivable situation in which this might take place

(ii) use this to define the value functions for the $(T-1)^{\text{th}}$ auction using what we now know are the correctly predicted outcomes of any subsequent auction

(iii) determine the outcome of all possible $T-1$ auctions and so, by repeating this backward recursion argument, completely determine the outcome of the sequence.

Not surprisingly, general results are harder to come by than in the case of single innovations. Vickers proves two results, the first of which is

Result 1. If, for all $\tau = 1, \ldots, T$, and for all $k = 1, \ldots, \tau$,

$$\sigma(c_{\tau+2}, c_{\tau+1}) > \sigma(c_{\tau+2}, c_k), \tag{2.7}$$

then the outcome in every auction will be action/reaction.

To understand this result, notice first of all that (2.7) can be (somewhat loosely) re-stated as requiring that industry profits be strictly decreasing in the costs of the high-cost firm throughout the entire range of costs spanned by the entire T auctions. This condition was seen to be sufficient to guarantee action/reaction in the case of a single innovation, and so it certainly guarantees it in the T^{th} auction. Suppose now we are entering the τ^{th} auction, and we know that all subsequent auctions will result in action/reaction. Then this implies that the situation the industry finds itself in after the next auction is complete will be that of one firm having costs $c_{\tau+2}$, and the other having costs $c_{\tau+3}$, and that this will be true whatever the outcome of the current auction in τ. But this means that when we look at the two value functions that appear in the bids B^i_τ and B^f_τ, those components that are attributable to future profits and future bid costs will be the same for i as for f. Thus all that can matter in determining whether or not $B^i_\tau > B^f_\tau$ are the components attributable to current profits, and then (2.7) again guarantees that the outcome of the τ^{th} auction will be action/reaction. By induction, action/reaction occurs in every auction.

This is a completely general argument establishing the result that, provided conditions sufficient to establish action/reaction in a single auction hold throughout the range of circumstances spanned by T auctions, we will get action/reaction in every auction. However, (2.7) is a much more stringent condition than the corresponding one-period one, for it requires the inequality in industry profits to hold however wide a gap opens up between the firms over the course of the sequence, and this will require much slower rates of technical change than would be implied by the single innovation.

Although (2.7) is only sufficient for action/reaction we can see intuitively how failure to satisfy it could generate dominance. For suppose that (2.7) held in the last period if the gaps were small, but not if they were large. Then there would be action/reaction in the last auction if firms entered it close together but not if they were far apart. But then the outcome of the auction in the previous period could affect the kind of auction that will take place in the last period, in which case the incumbent (at $T-1$) will enter the auction knowing that a win at $T-1$ guarantees a win at T, while the follower knows a win at $T-1$ guarantees a loss at T. This does not automatically imply that the incumbent at $T-1$ will necessarily win, because though his guaranteed win at T will bring it more profits, the firm may have to bid more in T to get those.

The second result that Vickers proves is

Result 2. If the high-cost firm's profits are always zero, and industry profits are strictly increasing in the costs of the high-cost firm, then the outcome is always persistent dominance.

These of course are precisely the conditions that prevail under Bertrand competition, so this again is just an extension of an earlier result for single innovations. The intuition is as follows. Given Bertrand competition, the incumbent will always win the last auction, and given that the loser's profits are zero, the bid he will have to make to win the patent is $\pi(c_{T+2}, c_{T+1})$, which is independent of which firm wins the second last auction. But then since the only thing that can matter in the second last auction are the profit streams of the incumbent and follower, and these are higher for the incumbent, the incumbent will win the second last auction. Moreover, the high-cost firm in $T-1$ makes no profits in that period, none in the last and has no bid costs, so the value of being the high cost firm is zero. But then the bid the incumbent needs to make to win the patent in $T-1$ is just $V(c_{T+1}, c_T)$, which is independent of the outcome of the third last auction, and so on.

The zero-profit condition is therefore quite important because it guarantees that anticipated future bid costs do not depend on the outcome of earlier auctions. What both these results therefore show is that it is only some of the strongest conditions for dominance or action/reaction in the single innovation models that carry over to sequences, and that, in general, conditions that are sufficient for dominance or action/reaction in single-innovation models can produce the opposite outcomes in at least some of the auctions in a sequence.

So far the discussion has focused exclusively on process innovation. In much of the innovation literature no great distinction is drawn between process and product innovation—they are both just ways in which firms can increase their profits. For many purposes this is perfectly correct. However, in the context of strategic innovation, it turns out that the distinction is vital.

2.3 Product innovation

In the sequential process innovation model just considered, as firms move through the sequence of auctions, they will acquire a portfolio of infinitely lived patents on various technologies. But at any time they will only employ the least-cost technology on which they have a patent. This is why the state of the game could be described by specifying just these two cost levels.

Suppose now we consider a model in which there is pure vertical product differentiation along the lines specified in, for example, Shaked and Sutton (1986). In common with most of the literature on product differentiation, the equilibrium concept employed is that of Bertrand competition.

We consider a sequence of innovations producing goods of successively higher quality, $q_1, \ldots, q_{T+2}$. Now, as the firms acquire a portfolio of patents they will in general want to produce more than one of the goods on which they hold a patent, since they will then be better able to target different consumer groups with their range of products. This makes the analysis more difficult, since in principle we now have to keep track of the entire patent portfolio of both firms.

In Beath *et al.* (1987), the problem is simplified by assuming that there are diseconomies of scope that limit each firm to producing only one good. Moreover, each firm can choose the quality of this good subject to it not exceeding the highest quality good on which they hold a patent. In this way the model is now formally very similar to that of Vickers (1986) that we have just discussed. Despite this, the striking result is that the conclusions are now almost precisely reversed. Very slow technical progress is now associated with persistent dominance, while very rapid technical change produces action/reaction.

To see why the results are reversed it is sufficient to examine how, holding fixed the quality of the high-quality good, industry profits vary with the quality gap between the two firms. This is shown in Fig. 2.2. To understand the shape of this curve, note that when the gap is zero the firms are producing identical products, and Bertrand competition eliminates profits. As long as there is some gap this enables profits to be made, so industry profits must initially increase with the gap. Recall that in these models low income consumers always have the option of not buying any of the goods. As the gap between the two goods widens, this option looks more attractive, so the low-quality good has to cut its price substantially to survive, and the high-quality firm has to do the same to compete. So profits start to fall as the gap gets too wide.

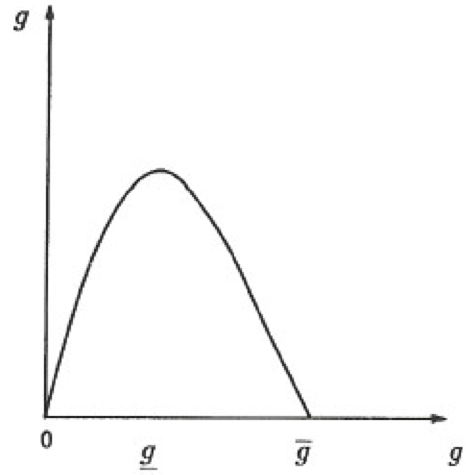

Figure 2.2

Since the logic of the previous arguments depended largely on whether industry profits rose or fell with the gap between the firms, the fact that Fig. 2.2 has precisely the reverse shape from Fig. 2.1 explains why the conclusions are reversed.

In Beath *et al.* (1989), an alternative model is proposed in which the assumption of infinitely lived patents is dropped, and patents now have only a finite life. It is also assumed that while there are only two firms engaged in the R&D race, there are many firms capable of producing goods whose patents have lapsed. Their price is driven to zero, so only the highest quality of these non-protected goods is ever bought. This keeps the number of goods in production fixed as we move through the sequence. However, there is a more important implication of having finitely-lived patents. For now if the outcome of the innovation process is persistent dominance, one firm will own all the extant patents and will be able to act as a monopolist restricted only by the presence of the highest-quality non-protected good available at a zero price. If, however, the outcome of the race is action/reaction, firms will end up holding patents on more than one good, but each of these goods will be separated from the other goods they own by goods held by the rival firm. The outcome of Bertrand competition between the two firms will be the same as if all the goods were owned by separate firms—and hence will be very competitive. In this way the nature of competition between products is entirely endogenous to the model, and is driven by the outcome of the innovation race, which will, in turn, be determined by the anticipated effects on market structure that winning or losing will produce.

Beath *et al.* show that if patents last for only two periods the outcome is always persistent dominance. The reason is the same as that behind Result 1 of Vickers. With 2-period patents the nature of future patent races is independent of the outcome of the current race. Hence all that can matter in determining this are the current levels of industry profits. But since incumbents get monopoly profits if they win which exceed profits in any duopoly, we must have persistent dominance. However, if patent life is extended to three periods, the outcome of the race can be action/reaction. For while incumbents get more profits if they win, followers are willing to pay a lot to stop them winning, so an incumbent's profits net of bid costs can be lower than those of a follower.

This is another example of the possibility that market situations which would be regarded as more competitive on the basis of static theory, can produce a less competitive outcome under dynamic competition.

2.4 Summary

The models we have been considering are undoubtedly excessively stylized—in particular, as we will show, the omission of uncertainty is crucial. Yet they have provided some valuable insights:

(i) single-period models can give very misleading results

(ii) there are crucial differences between product and process innovation

(iii) policies which, on the basis of static arguments, may be thought to produce more competitive outcomes can, in dynamic models, generate less competition.

A major difficulty with these sequential models, however, is that because the outcome depends to some extent on the anticipated future bid costs, then the rather artificial auction construction in which the loser effectively commits no resources to R&D, while the leader could put in a lot, may excessively bias the results towards action/reaction. An important feature of R&D is that resources can often be lost if a rival manages to develop a new idea first. What we need is a model in which each firm has to commit resources to R&D if it is to be have any chance of innovating, but where those resources are wasted if the firm loses. In the next section we examine a model of this kind.

3. Single innovations under uncertainty

3.1 Introduction

In the models of certainty considered in the previous section, each firm knows that by spending a little more than their rival, they can capture the innovation for sure, and the only factor which is relevant in making this decision is therefore the difference between their future profits if they, rather than their rival, innovate. This is what the competitive threat captures.

What is missing from the analysis is the standard investment calculation in which increased spending on R&D is balanced against the increase in profits over those *currently* being earned that this spending will bring about. If, however, we treat investment in R&D as having only a limited probability of success at any one time, though this probability of success can be influenced by the amount spent, then firms will face a more conventional investment decision in which they balance the gains from bringing forward the likely date at which they earn higher profits than they are currently getting, against the costs of doing so. This will introduce into the analysis an incentive to innovate that we will call the *profit incentive*. Notice that this incentive will be present even if the firm were not in a race against rivals. Indeed, as we will see, the amount of R&D expenditure generated by this force is precisely determined by calculating the amount of investment the firm would do if it were undertaking R&D in isolation. (Katz and Shapiro (1987) call this the 'stand-alone incentive.')

However, provided the risks in R&D are imperfectly correlated across firms, each firm still faces the possibility that its rival can innovate before it, and so still also faces a competitive threat. As we shall show, by understanding the relative magnitude of these two forces, the 'stick' of the competitive threat and the 'carrot' of the profit incentive, a great deal can be learned about a firm's R&D strategy, and hence about the likely outcome of strategic races in which it may be involved.

While it may seem obvious that these two forces will need to be taken account of in any R&D race, most of the models of strategic innovation in the literature turn out to be special cases of the general framework that is developed below. The reason is that they involve assumptions that effectively limit the role that one or other of these forces plays. Thus the profit incentive is the force on which Arrow (1962) focuses in reaching his conclusion that a competitive market structure is more conducive to innovation than a monopolistic one. It similarly underpins Reinganum's (1985) dynamic generalization of this result where she demonstrates that the process of dynamic competition is characterized by Schumpeterian 'creative destruction.' In contrast, the work of Gilbert and Newbery (1982) and, more recently, of Harris and Vickers (1987) is concerned with situations where only the competitive threat has any significant role to play, and this leads to the diametrically opposite conclusion that only incumbents win R&D races.

In the next section we develop a simple, yet general, model of a strategic R&D race between firms to be the first to introduce some new technology or product and then, in §3.3, show how it can be used to organize the existing literature on one-shot R&D games. It also provides the building blocks for the discussion in §4 of a sequence of strategic races under uncertainty.

3.2 The model

The model is a natural extension of that introduced by Lee and Wilde (1980). The costs of R&D take the form of a flow of expenditures that have to be incurred until the race is over (a 'subscription') rather than, as in Loury (1979), being a fixed cost at the outset (an 'entry fee').

The model is one in which two firms are engaged in a single one-stage race to be the first to introduce some new product or technology. As we shall show later, there is no great difficulty in extending the analysis to the situation where where is a (finite) sequence of races, or where there are many stages in a single race, but most of the understanding of such sequences comes from understanding the outcome of such a single one-stage race.

In general we want to allow for the possibility that firms are also engaged in production which is generating current profits, and that, possibly as the outcome of some previous race, one of the firms currently has a competitive edge over its rival enabling it to make greater profits than it. Similarly, there is no reason to suppose that the profits each firm makes conditional on winning or losing the race will be the same. However, we assume that current profits and the levels of profits made by firms contingent on winning or losing are known with certainty.

Each firm has to decide how much R&D to do at each instant of time. We make the assumption, common to other models, that the probability of a firm discovering the new product or process in the time interval $(t, t + dt)$ conditional on no one having discovered by t, depends solely on the flow rate of R&D expenditure at t undertaken by that firm, and *not* on the accumulated amounts of R&D. Thus there is no 'learning-by-doing' *within the current state*. In addition, it is assumed that the relationship between probability of discovery and R&D expenditure is time independent and exponential.

The *hazard rate* is defined as the instantaneous probability of discovery at τ, conditional on not having discovered before τ. Therefore, if one firm alone were innovating, and had chosen to spend a constant amount on R&D each period, thus generating a constant hazard rate, x, say, the probability of discovery by τ is

$$F(\tau; x) \equiv 1 - e^{-x\tau}.$$

Thus, the instantaneous probability of discovery at τ is

$$f(\tau; x) \equiv \partial F/\partial\tau \equiv xe^{-x\tau},$$

and so, the probability of discovery at τ, conditional on not having discovered before τ, the *hazard rate*, is just

$$H(\tau; x) \equiv f(\tau; x)/[1 - F(\tau; x)] \equiv x,$$

as we know it must be.

These assumptions are strong. The absence of learning-by-doing may appear serious but, as we shall see, is relatively easily accommodated by moving to a sequential framework. The restriction to a time-independent and exponential structure is motivated solely by tractability. Undertaking strategic analysis in any other framework becomes extremely difficult.

The great advantage of this framework is that if one firm chooses a constant hazard rate, then the other's best response is also to choose a constant hazard rate—since nothing is changed in the basic structure of the game as we move through time.

We assume that all firms face the same innovation technology, whereby the instantaneous resource cost of achieving the hazard rate x is given by $\gamma(x)$ and that

(i) $\gamma(0) = \gamma'(0) = 0$

(ii) $\forall x \geq 0, \;\; \gamma'(x) \geq 0, \;\; \gamma''(x) > 0$

(iii) $\gamma'(x) \to +\infty \text{ as } x \to +\infty.$

Thus there are everywhere decreasing returns. This restriction on the cost function is not too essential—though one cannot allow for too general a pattern of costs or there may be non-unique equilibria (see Harris and Vickers 1987).

Let s be the current profits of the incumbent, t those of the follower. If the incumbent innovates first, the present value, at the time of innovation, of the future profits it receives is A, while the follower's present value of future profits will be E. However, if the follower innovates first the present value of the incumbent's profits will be B, while the follower's will be D. If we let r be the common rate of interest, then we assume that $A > B$; $Ar > s$; $D > E$; $Dr > t$; where Ar and Dr are the flow rate of profits with present values A and D respectively. Thus for both firms winning is more profitable than losing, and is more profitable than their current position.

Let x denote the hazard rate chosen by the incumbent, y that chosen by the follower. Let r be the common rate of interest, V the expected present value of profits of the incumbent, W that of the follower. Then it is straightforward to show that

$$V(x,y) \equiv \frac{Ax + By + s - \gamma(x)}{x+y+r} \tag{3.1}$$

$$W(x,y) \equiv \frac{Dy + Ex + t - \gamma(y)}{x+y+r}. \tag{3.2}$$

We want to find the Nash equilibrium in hazard rates. To do that we will focus on characterizing the *reaction function* for the incumbent showing its profit-maximizing choice of x in response to any given choice of y by the follower. The reaction function of the challenger follows by analogy.

Since it will turn out that given our assumptions both firms will choose positive hazard rates, the incumbent's reaction function is defined by the equation $V_x(x,y) = 0$, where the subscript denotes partial differentiation with respect to a variable. Now $V_x(x,y) = 0$ if and only if $\varphi(x,y) = 0$, where

$$\varphi(x,y) \equiv (A-B)y + (Ar - s) + \gamma(x) - (x+y+r)\gamma'(x). \tag{3.3}$$

Given our assumptions on $\gamma(\cdot)$, $\varphi(\cdot)$ satisfies the following conditions:

(i) $\forall y \geq 0,\ \varphi(0,y) > 0$

(ii) $\forall x \geq 0,\ y \geq 0,\ \varphi_x(x,y) < 0$

(iii) $\forall y \geq 0,\ \varphi(x,y) \to -\infty$ as $x \to +\infty$.

These conditions guarantee that, for all $y \geq 0$, there exists a unique, strictly positive and finite value for x for which $V_x(x,y) = 0$, and hence that firm 1 does indeed have a well-defined reaction function. Denote this by $\rho(y)$. Then $\rho(\cdot)$ satisfies the condition

$$\varphi[\rho(y), y] \equiv 0,\ \ \forall y \geq 0. \tag{3.4}$$

Now, define x_0 by the condition

$$(Ar - s) + \gamma(x_0) - (x_0 + r) \cdot \gamma'(x_0) = 0. \tag{3.5}$$

Clearly, $\rho(0) = x_0$ and $x_0 > 0$.

We now define $\overline{x}$ by

$$\gamma'(\overline{x}) = (A - B) \tag{3.6}$$

and notice that

$$\begin{aligned} \text{sign}\{\rho'(y)\} &= \text{sign}\{\varphi_y[\rho(y), y]\} \\ &= \text{sign}\{(A - B) - c'[\rho(y)]\} \\ &= \text{sign}\{c'(\overline{x}) - c'[\rho(y)]\}. \end{aligned}$$

Moreover, if for some y, $0 \leq y < \infty$, $\rho(y) = \overline{x}$, then not only do we have $\rho'(y) = 0$ but, on substituting back into (3.4), we find that $\rho(y)$ also satisfies (3.5) and hence $x_0 = \overline{x}$. We therefore have the following results:

(i) If $x_0 < \overline{x}$, then $\forall y \geq 0$, $x_0 \leq \rho(y) < \overline{x}$, and $\rho'(y) > 0$

(ii) If $x_0 = \overline{x}$, then $\forall y \geq 0$, $\rho(y) = \overline{x}$,

(iii) If $x_0 > \overline{x}$, then $\forall y \geq 0$, $x_0 \geq \rho(y) > \overline{x}$, and $\rho'(y) < 0$.

Now, x_0 is simply the hazard rate chosen by firm 1 if its rival were doing no R&D. In this case the only motivation for doing R&D is to increase the flow rate of profits by $(Ar - s)$. Thus the incumbent's hazard rate will be chosen by balancing the gains from bringing forward the likely date of innovation against the additional costs of doing so. This is just what condition (3.5) involves. We will refer to x_0 as the profit incentive facing firm 1.

On the other hand, what we have just shown is that $\overline{x}$ is the asymptotic hazard rate chosen by firm 1 as $y \to \infty$. Notice that as $y \to \infty$, the challenger is almost certainly immediately about to innovate. If, however, the incumbent were to instantaneously successfully innovate, the marginal benefit to it would be $(A - B)$, while its marginal cost would be $\gamma'(x)$. This then explains the definition of $\overline{x}$ in (3.6). One might also note that current profits are irrelevant in determining $\bar{x}$ because, as firm 2 is almost certainly going to innovate immediately, these profits are almost certainly going to instantly disappear—whoever wins. We will refer to the hazard rate $\bar{x}$ as the *competitive threat* facing firm 1.

At this level of generality the relative magnitude of these two forces can go in either direction. However, a major determinant of their relative magnitude is clearly going to be the *ease of imitation*. Suppose first of all that that imitation is impossible, so that the new product or new technique is protected by an infinitely-lived and highly effective patent. Then, in a wide class of cases, it will be natural to assume $Ar > s > Br$ since, if you lose the race, your rival will have an even better product/process with which to compete against you, and your profits could be lower than at present. In this case it is clear we will have $x_0 < \bar{x}$. On the other hand if imitation were very easy then we could have $Ar \approx Br > s$, and we could then have $x_0 > \bar{x}$.

Having understood the nature of the two hazard rates x_0 and $\bar{x}$, the results we obtained above about the nature of the reaction function in the various cases should now be intuitively clear. Let us therefore briefly consider cases (i) and (iii) above.

(i) $x_0 < \bar{x}$

Here the incentive to do R&D in order to prevent the challenger from winning exceeds the incentive to undertake R&D to gain a greater profit stream. In such a case, if the follower were to increase its R&D effort the natural competitive response for the incumbent is to increase its own effort in response.

It is straightforward to show that moves along the reaction function to higher values of y and x reduce the profits that the incumbent makes. The reason is that such moves take the firm's hazard rate further and further above the level (x_0) that is optimal given the profits it will make if it is successful.

This case is illustrated in Fig. 3.1.

(ii) $x_0 > \bar{x}$

This case would arise if imitation were very easy and the successful innovator could be easily copied. In such a situation the incumbent's profits would be almost the same whether it won or not. Clearly, were this to be so, the competitive threat would virtually disappear. Nevertheless, the innovation could still enable the firm to make substantially greater profits than at present, which is why the profit incentive exceeds the competitive threat.

In such a situation there is a strong externality in each firm's R&D: one's rival's R&D is a very good substitute for one's own. Hence

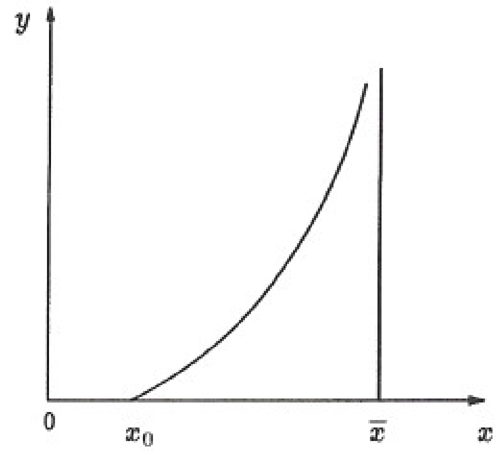

Figure 3.1

suppose that the follower were to increase its R&D effort. The incumbent's optimal response is to do less since it is going to get more or less the same profits irrespective of who wins and the follower is helping to bring forward the likely date of innovation. In this case then, as y increases and x falls along the reaction function, the profits of the incumbent *rise* because although it appears to be moving away from the hazard rate that is profit maximizing, it is really only substituting its rival's R&D for its own, keeping the overall date of innovation more or less fixed, but saving in current R&D costs. There is thus a 'free-rider' phenomenon.

This case is illustrated in Fig. 3.2.

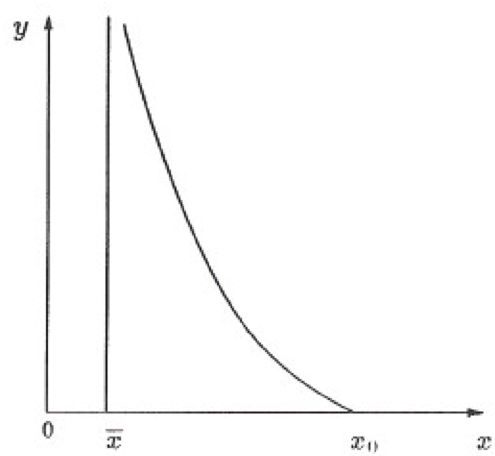

Figure 3.2

We can undertake a similar analysis for the challenger, and derive its reaction function, which again can be either increasing or decreasing depending on the relative magnitudes of the competitive threat and profit incentive it faces. Moreover, the fact that, say, the competitive threat for the incumbent is greater than its profit incentive tells us nothing in general about the relative magnitudes of these forces for the follower. So there will be four possible types of model, depending on the relative magnitudes of the forces for each of the two firms.

Under sufficiently strong assumptions about the cost function (for example that it is quadratic) then the reaction functions will have the convexity/concavity properties shown in Figs. 3.1 and 3.2 and it is easily checked that for each model type the Nash equilibrium is unique.

The model has straightforward comparative static properties. Thus an increase in the competitive threat or profit incentive facing, say, the incumbent will increase the hazard rate, x, it chooses for every positive value of y. This will increase the equilibrium value of x. Whether the equilibrium value of y is higher or lower will depend on the relative strengths of the competitive threat and profit incentive facing the challenger.

We can now establish a straightforward result about the likely winner in the race. If we let $x^\star$ and $y^\star$ be the equilibrium hazard rates, then:

$$\textit{if} \quad (A - B) \geq (D - E) \quad \textit{and} \quad (Ar - s) \geq (Dr - t),$$

$$\textit{then} \qquad x^\star \geq y^\star.$$

If one of these two conditions is satisfied as a strict inequality, then $x^\star > y^\star$. Similarly, if the inequalities governing the competitive threats and profit incentives are reversed, so too is the conclusion about the hazard rates.

While these results are very obvious, a great deal of the literature is covered by them.

3.3 Sorting out the literature

The first set of results we can look at are those contained in the papers by Loury (1979), Dasgupta and Stiglitz (1980) and Lee and Wilde (1980). These authors consider the case of identical firms competing to *enter* the industry with some innovation. Thus the payoffs to success or failure are identical, as are current profits. Hence we have

$$A = D > 0$$

$$B = E = 0$$

$$s = t = 0.$$

From this it follows that the competitive threats and profit incentives are identical, that the former exceed the latter, and so the expected outcome of the race is indeterminate.

The paper by Reinganum (1985) focuses on drastic process innovations. In this case the prize structure is the following:

$$A = D > 0$$

$$B = E = 0$$

$$s > 0, \quad t = 0.$$

From this we can see that while each party faces the same competitive threat, the incumbent has the smaller profit incentive because it is only for him that the innovation will replace a positive current flow of profits. Thus we get $y^* > x^*$ and so Reinganum's conclusion that the outcome will be one of creative destruction or action/reaction.[3]

It is important to notice, however that drasticness is crucial to this argument. For suppose instead that we had process innovation, Bertrand competition, but non-drastic innovation. We would still satisfy the Reinganum conditions that $B = E = t = 0$, and that $s > 0$, so the incumbent has greater current profits to protect than the challenger. However, as we noticed in the last section, $A > D$, so the incumbent has the greater competitive threat. Moreover, it is perfectly possible that $(Ar - s) > Dr$, so that the incumbent also has the greater profit incentive. In this case the introduction of uncertainty, far from overturning the presumption of dominance that emerges from the model with certainty, would reinforce it.

There have been a series of papers by Harris and Vickers (1985, 1986, 1987) in which they look at the race as a multi-stage one. We will discuss their model in more detail in the next section, but for the moment the features we wish to note are that they assume there are no profits earned until the innovation takes place, and that, while the payoffs to innovation are finite, there is no discounting. Thus $r = s = t = 0$, and so there are no profit incentives: $x_0 = y_0 = 0$. The winner will be the firm which faces the greater competitive threat.

Hence, while Reinganum's result is driven by the asymmetry of profit incentives, Harris and Vickers' results are entirely a matter of competitive threats. Other features of their model guarantee that these always favour the incumbent.

[3]This has, of course, to be understood in the probabilistic sense that the challenger is more likely to win, rather than that they actually will win.

3.4 Summary

The introduction of uncertainty into the analysis is important, not because, when incorporated in the natural way of uncertainty over the timing of innovation, it is patently 'more realistic'—we obtained many useful insights from the models with no uncertainty—but because it introduces an important extra incentive into firms' R&D decisions. In this section we have examined a relatively general model incorporating both the competitive threat and the profit incentive, and have shown that many models in the literature are special cases of it. While it is certainly true that in the case of drastic innovations, where the certainty model was indeterminate, this extra consideration tells in favour of action/reaction, this is by no means a general result, and in some circumstances the introduction of uncertainty can reinforce the persistence of dominance.

We saw in the last section that it is essential to study the issue of dominance vs. action/reaction in sequential models, but that the certainty models we examined were somewhat inadequate since the loser commits no resources to R&D. We now have a model in which, in general, both firms undertake R&D, so we are now in a position to combine these two major strands in the modelling of strategic innovation.

4. Sequential innovation under uncertainty

4.1 Introduction

Integrating the models of §§2 and 3 is, in principle, fairly straightforward. Thus suppose, once again, that the two firms start off with unit costs c_1 and c_2, and that there is a sequence of T cost-reducing innovations producing the sequence of unit costs, $c_3, \ldots, c_{T+2}$, where $c_{\tau+1} > c_{\tau+2}$, $\tau = 1, \ldots, T$. The first race the firms enter is that for the patent on c_3. As soon as one of them succeeds in discovering the technology for c_3, they get an infinitely-lived patent on it, and they both start to undertake R&D to discover the technology for c_4, and so on.

The way in which we could determine the outcome of the race is then as follows. First, recalling the definitions of the functions $V(c_{\tau+2}, c_k)$ etc. in §2, we clearly have $\forall k = 1, \ldots, T+1$,

$$V(c_{T+2}, c_k) = \pi(c_{T+2}, c_k)/r; \quad V(c_k, c_{T+2}) = \pi(c_k, c_{T+2})/r. \qquad (4.1)$$

Now consider the final race. The two firms enter this with one holding the patent on c_{T+1}, the other with the least-cost technology on which it holds a patent being c_k, $k = 1, \ldots, T$. Let

$$A = V(c_{T+2}, c_k); \qquad B = V(c_{T+1}, c_{T+2})$$
$$D = V(c_{T+2}, c_{T+1}); \quad E = V(c_k, c_{T+2})$$
$$s = \pi(c_{T+1}, c_k); \qquad t = \pi(c_k, c_{T+1}).$$

Using the model employed in §3, solve for the equilibrium values x^* and y^*, and substitute the values of these and of the above parameters into the formulae for V and W given in (3.1) and (3.2). Then

$$V(c_{T+1}, c_k) = V; \quad V(c_k, c_{T+1}) = W.$$

By repeating the process of backward recursion we can, in principle, solve the value functions and associated levels of hazard rates (R&D expenditures), in every state, and in every race. We can then start in the first race, work out the likely winner, and, assuming this is the winner, move on to the second race, work out the likely winner, and in this way, determine the entire outcome of the sequence of races.

There are, however, a number of issues to confront before reporting the results of this analysis.

The first is that some of the assumptions made in the previous section, which seemed perfectly natural in the context of a single innovation, need not hold in a sequential framework. For example, there is no guarantee in general that $Ar > s$. This is simply because while winning an innovation may bring greater current profits π, the firm may have to incur such large R&D expenditures in the attempt to maintain their leadership in future races that the present value of their net future income is less than their current gross profits. In itself this causes no great difficulties in characterizing the incumbent's reaction function. For example, if $Ar < s$ but $A > B$, so the incumbent has a zero profit incentive, but positive competitive threat, then all that happens is that there exists a $\widehat{y} > 0$ such that the incumbent's optimal choice of x is zero if $y \leq \widehat{y}$, and positive and increasing in y for $y > \widehat{y}$. This is illustrated in Fig. 4.1.

The real difficulty is that the equilibrium now need not be unique. This is illustrated in Fig. 4.2, where the challenger also faces a competitive threat that is greater than its profit incentive. Recalling that when the competitive threat exceeds the profit incentive then a firm's profits are strictly decreasing in its rival's hazard rate, it is clear that in this particular case it would be natural to choose as equilibrium the point $(0, y_0)$, since this Pareto-dominates the other equilibria. However, in other cases there may be no obvious equilibrium to select.

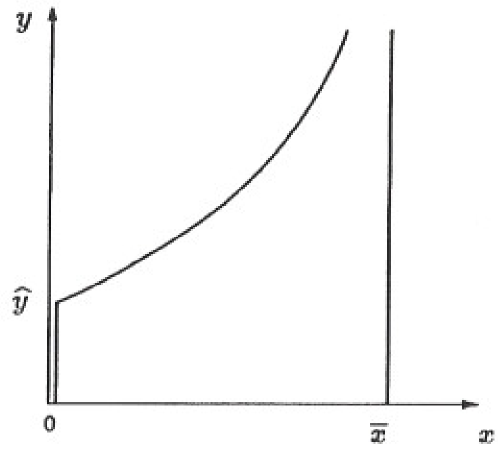

Figure 4.1

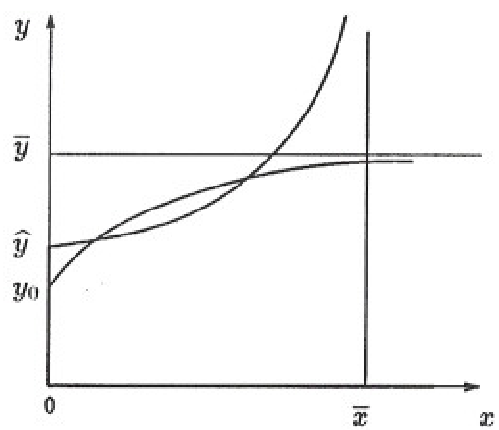

Figure 4.2

Notice also that this non-uniqueness does not arise because of the nature of the hazard cost function $\gamma(x)$, and cannot be removed by strengthening the assumptions on it.

Even were there no problem of non-uniqueness, analytical results are extremely hard to obtain with these models.[4] Because the equilibrium values of R&D expenditures in one race enter the present values A, B, D, and E that appear in the previous R&D race, we need to have a fairly explicit link between these equilibrium levels of R&D

[4] Analytical results have been obtained by Harris and Vickers (1987) for the case of a one-dimensional tug-of-war, where the firm that commits most resources to R&D will be more likely to move one stage nearer their goal, simultaneously taking their rival one step futher away from theirs. Moreover, in their model there is a sequence of stages in a single innovation, and firms earn no profits until the innovation is discovered. In addition, although the prizes are finite, there is no discounting.

expenditure and the parameters of the associated R&D race in order to be able to link together what is happening in the various races. In the certainty model this link was extremely simple, but nevertheless we found that analytical results were fewer than in the single-period model. In the model with uncertainty the link between equilibrium levels of R&D expenditure and the parameters of the model cannot be expressed in a simple closed form, even for the quadratic form of the cost function $\gamma(\cdot)$, and this is why analytical results are so hard to obtain.

For both these reasons the results we report below come from computer simulation models. For the cases we report, we have been able to check that all equilibria are unique, so nothing depends on an arbitrary choice of a particular equilibrium. The loss of generality from resorting to explicit calculation is compensated for by a fairly rich set of results which give us a good feel for how some of the essential features of the models affect their outcomes.

Another issue we must confront concerns the nature of the innovative process. In the description given above, and implicit in all the models discussed so far, is the assumption that at all times firms are competing for the same innovation. This means that although one firm has a patent on a technology that prevents everyone from using it, nevertheless all the relevant technological information has become common knowledge and, moreover, those who were unsuccessful in the previous race are just as able to exploit this common knowledge as the firm which discovered the new technology. This latter assumption rules out learning-by-doing *between* races—we have already had to rule it out *within* races. These assumptions imply that however large a gap a firm opens up with its rival it always faces the danger of being *leapfrogged* by it.

An alternative, equally strong, assumption that we could make is that either because technological information does not become common knowledge, or else because there are extremely strong learning-by-doing effects, each firm essentially has to discover each technology by themselves before they can move on to discover the next. In this case, at any one time, the incumbent and the challenger will be competing for *different* innovations. If the incumbent firm succeeds before the follower it will pull ahead of its rival, while if the challenger succeeds first it will simply move one stage closer. In this case all it can do is *catch up*—it cannot *leapfrog*. It is precisely this *catch up* framework that has been employed in the recent models of Harris and Vickers (1985, 1987), albeit in a somewhat different setting where firms are engaged in a single multi-stage race rather than a sequence of different races. In what follows we will examine the difference between models in which leapfrogging is a possibility, and those in which the structure of moves is that of catching up, but, to maintain comparability, will do so in the context of a sequence of innovations.

Consider then our sequence of T cost-reducing innovations. To make sense of the move structure we are going to employ, it will help to think of c_{T+2} as some absolute minimum cost level beyond which costs cannot be further reduced by any subsequent innovation. This means that if one firm has succeeded in lowering its costs to c_{T+2} it undertakes no further R&D. The other firm may continue to do R&D until its costs too have fallen to c_{T+2}, but it is possible that this firm will have given up at some stage before this. In this way the actual number of races run is endogenous to the model. Formally the structure of the *catch up* model is specified as follows:

(i) $V(c_{T+2}, c_{T+2}) = \pi(c_{T+2}, c_{T+2})/r$

(ii) If we are now in the position where the incumbent has c_{T+2} and the challenger has c_{T+1}, then the only issue is whether the challenger does any R&D. So consider a race of the kind studied in the previous section in which $x^\star = 0$; $B = D = V(c_{T+2}, c_{T+2})$; $s = \pi(c_{T+2}, c_{T+1})$; $t = \pi(c_{T+1}, c_{T+2})$. Then, if $Dr \leq t$, $y^\star = 0$, while if $Dr > t$ then $y_0 > 0$, and $y^\star = y_0$. Insert these values into (3.1) and (3.2) to obtain

$$V(c_{T+2}, c_{T+1}) = V; \quad V(c_{T+1}, c_{T+2}) = W. \tag{4.2}$$

(iii) Proceeding by backward recursion we can in a similar fashion define $V(c_{T+2}, c_k)$and $V(c_k, c_{T+2})$ for all $k = 1, \ldots, T+1$.

(iv) Suppose now the incumbent has costs c_j, $j = 2, \ldots, T+1$, and the challenger has costs c_k, $1, \ldots, j$ and that the value functions have been defined for all greater cost levels $c_{j'}$, $j' > j$ for the incumbent, and $c_{k'}$, $k' > k$, for the challenger. Let $A = V(c_{j+1}, c_k)$; $B = V(c_j, c_{k+1})$; $D = V(c_{k+1}, c_j)$; $E = V(c_k . c_{j+1})$; $s = \pi(c_j, c_{k+1})$; $t = \pi(c_k, c_j)$.

Using the ideas of the previous section, determine the equilibrium values $x^\star$, $y^\star$, and insert all these values into (3.1) and (3.2) and so determine

$$V(c_j, c_k) = V : V(c_k, c_j) = W. \tag{4.3}$$

(v) Proceeding by backward recursion, we can determine in this way all the value functions and associated hazard rates for all possible cost combinations of the incumbent and rival.

(vi) We can now work forwards to determine the outcome of the sequence of races.

Intuitively we would expect that models with a *leapfrog* structure will produce very different outcomes from those with a *catch up* structure. In a *catch up* model, if the incumbent can open up a sufficiently large gap then even if the rival firm puts in a lot of R&D and innovates ahead of its rival, it will still be a long way behind, and will face the prospect of having to continue to put in a lot of R&D to try to close the gap further. Its profits if it wins may not be much greater, if at all, than if it loses, so it faces little competitive threat. Similarly the incumbent will recognize that its profits if it loses will not be much lower than those if it wins, and so will also have a small competitive threat. Thus R&D effort on everyone's part should fall as the gap widens. However, anticipating this at an earlier race when the gap may be smaller, the incumbent faces lower future R&D costs if it wins the race than does the challenger, and while this by no means guarantees persistent dominance, it certainly makes it more likely. What we would expect then, is that if the incumbent and challenger start off fairly close together, and the rate of technical change is not very great so the incumbent cannot rapidly get too far ahead, then a far-sighted challenger may realize that if it can *catch up*, while R&D levels will be high, they will be more or less the same for both incumbent and challenger, and in this case the outcome is more likely to be action/reaction.

By contrast, with *leapfrogging* the leader can always be overtaken however large the gap, and may face the prospect of greater competitive threats and hence more R&D expenditure, the greater the gap gets. In this case we expect action/reaction to be the more likely outcome.[5]

To test out these intuitions we turn to the results of our computer simulations.

4.2 The results

The results we report were based on the following model. The two firms produce an identical product. The demand for this is given by the constant elasticity demand curve

$$p = S^{\varepsilon} q^{-\varepsilon}, \tag{4.3}$$

where p is the price, Q total output, ε, $0 < \varepsilon < 1$, is the inverse elasticity of demand, and S is a parameter reflecting the size of the market.

[5] Fudenberg *et al.* (1983) have stressed the important role that allowing the *possibility* of leapfrogging plays in determining whether the outcome of the innovation process is one of persistent dominance.

There are T cost reducing innovations. Technical progress occurs at the constant rate g, where

$$g = (c_k/c_{k+1}) - 1 \quad k = 1, \ldots, T+1. \tag{4.4}$$

The R&D cost function is quadratic, i.e. $\gamma(x) = x^2$, $r > 0$ is the rate of interest.

In addition to varying these parameters we also varied two main structural features of the model: whether the nature of market competition was Cournot or Bertrand; and whether the structure of moves in the sequence was *leapfrog* or *catch up*.

It turns out that the latter tend to be more significant, so we will report results for the cases where the parameters S, T and ε are fixed at 10000, 15 and 0.5 respectively, and where the pair of parameters r and g take on the two sets of values $(0.05, 0.05)$ and $(0.01, 0.01)$. The results we obtained from the simulations were as follows:

If we start with the case where the market equilibrium is that of Bertrand competition, and where the move structure is *catch up*, then the outcome is one of persistent dominance. Moreover, while the hazard rates chosen by the two firms in the initial race were very high ($x^\star = 7554$, $y^\star = 3780$ for the case where $r = g = 0.05$), they rapidly fell until, after six races, the challenger had given up completely and the hazard rate chosen by the incumbent was 33. This is precisely what we would expect from the work of Harris and Vickers.

In contrast, if the move sequence is *leapfrog*, and the market equilibrium is again Bertrand, then while the last races in the sequence are again characterized by persistent dominance, in the early stages there is action/reaction as the firms jockey for the persistent leadership of the later races. The higher growth and interest rate are associated with dominance over a smaller number of races. However, the striking feature is that as the gap widens between the two firms throughout the phase of persistent dominance, the levels of R&D spending rise sharply as well, reflecting the greater threat the leader faces of being leapfrogged by its rival.

If we retain the assumption that the move structure is *leapfrog* but now introduce Cournot competition, then the outcome is one of action/reaction in every race, and this is true for both sets of values for r and g.

However, probably the most interesting set of outcomes arose when the move structure was *catch up* and the market equilibrium Cournot. Here, in the case where the rate of technical progress (and the rate of interest) was 5 per cent, the incumbent wins the first 9 races, opening up a considerable gap between the firms. The challenger then completely closes the gap, and thereafter whenever one firm pulls ahead the other immediately puts in more R&D effort and closes the gap again. However, when the rate of technical change and rate of interest are at 1 per cent, the challenger always closes any gap that opens up, and in this sense we have action/reaction all the time.

Thus it still remains true that Bertrand competition produces 'more' dominance than does Cournot. However, unlike the case of certainty, where, with a *leapfrog* move structure Bertrand competition always produced dominance, with the introduction of uncertainty this result is no longer true, and it can produce action/reaction, at least over some of the races in a sequence.

Similarly, while in the work of Harris and Vickers where the sequence of races were stages in a single race, it was shown that the *catch up* move structure always produced dominance, in the case of a sequence of races this is no longer the case, and we can get action/reaction when there is Cournot competition and low rates of technical change.

However, drawing on a range of simulations not reported here, we can say that the combination of *leapfrog* moves and Cournot competition almost always produces action/reaction, while the combination of *catch up* moves and Bertrand competition almost always produces persistent dominance.

4.3 Summary

We have seen that in modelling a sequence of races under uncertainty it is important to distinguish cases where there is rapid dissemination of information and no strong learning-by-doing effects, from those where there is poor dissemination or strong learning-by-doing. While the latter by no means guarantees that there will be persistent dominance, it certainly makes it more likely.

However, there are very few analytical results available for models of this kind, and a great deal still remains to be learned.

5. Conclusions

The literature on strategic innovation is somewhat bewildering with a variety of results being produced by a host of special models each built on rather different assumptions. Moreover, as is clear from the interchange between Reinganum (1983) and Gilbert and Newbery (1984), it is frequently difficult to decide just which features of these individual models are responsible for producing the results. In this chapter we have tried to provide a coherent framework within which these various models can be set, thus making it considerably easier to understand what is really driving the various results.

In undertaking this exercise, however, we have had to narrow the focus of much of the work covered. For example, we have conducted the entire discussion in the context of having just two firms involved in technological competition. Though in many cases this is not essential, and results have been obtained in a multi-firm context, there are some areas, such as sequential product innovation under uncertainty, where we have as yet no formal models, let alone results. In addition, we have focused exclusively on the positive predictions of the theories and have neglected some of the normative features which have often been the motivation for the work.

It is also important to be aware of some of the limitations of the literature we have surveyed. The competition we have been examining has taken the form of a *tournament* where there is a single prize to be won by whoever comes first. While this is no doubt a useful description of some kinds of technological competition, there are important lessons to be learned from non-tournament models of the kind surveyed in Dasgupta (1986), where there are still strategic issues of using innovation as an entry barrier. Even within a tournament framework, much of the literature, by implicitly using a *leapfrog* move structure, has focused exclusively on the *incentive* to innovate, rather than the *opportunity* to do so. Although the *catch up* move structure introduced by Harris and Vickers goes some way to redressing this, it is equally special, and more attention needs to be paid to information dissemination and imitation. In particular, it would be important to allow firms to choose between a strategy of making a number of small innovations where information may be rapidly disseminated and *leapfrogging* might take place, and that of going for a larger, riskier innovation which could take longer to develop, but where the *catch up* move structure may better describe the rivals' opportunities.

Finally, the literature has concentrated almost exclusively on conditions in the product market, and much remains to be done to understand how innovatory success is linked to factors such as the internal organization of the firm, access to financial markets, and the nature of the labour markets firms face.[6]

[6] The paper by Bhattacharya and Ritter (1983), discusses some of the issues that arise in the financing of innovation, while the paper by Ulph and Ulph (1988) explores the way in which labour market conditions can affect a firm's success in innovation.

References

Arrow, K. J. (1962). Economic welfare and the allocation of resources for inventions. In *The rate and direction of inventive activity*, (ed. R. R. Nelson). Princeton University Press.

Beath, J., Katsoulacos, Y., and Ulph, D. (1987). Sequential product innovation and industry evolution. *Economic Journal*, **97**, 32–43.

Beath, J., Katsoulacos, Y., and Ulph, D. (1989). Sequential product innovation with endogenous market structure. In *Market Structure and Market Strategy*, (ed. A. Gee and G. Norman). Wheatsheaf Books, London. Forthcoming.

Bhattacharya, S., and Ritter, J. (1983). and communication: Signalling with partial disclosure. *Review of Economic Studies*, **50**, 331–46.

Dasgupta, P. (1986). The theory of technological competition. In *Economic Organizations as Games*, (ed. K. Binmore and P. Dasgupta). Blackwell, Oxford.

Dasgupta, P., and Stiglitz, J. (1980a). Industrial structure and the nature of innovative activity. *Economic Journal*, **90**, 266–93.

Dasgupta, P., and Stiglitz, J. (1980b). Uncertainty, industrial structure and the speed of R&D. *Bell Journal of Economics*, **11**, 1–28.

Dasgupta, P., and Stiglitz, J. (1981). Entry, innovation, exit: Towards a dynamic theory of oligopolistic industrial structure. *European Economic Review*, **18**, 137–58.

Fudenberg, D., Gilbert, R., Stiglitz, J., and Tirole, J. (1983). Preemption, leapfrogging and competition in patent races. *European Economic Review*, **22**, 3–31.

Gilbert, R., and Newbery, D. (1982). Preemptive patenting and the persistence of monopoly. *American Economic Review*, **72**, 514–26.

Gilbert, R., and Newbery, D. (1984). Uncertain innovation and the persistence of monopoly: Comment. *American Economic Review*, **74**.

Harris, C., and Vickers, J. (1985). Perfect equilibrium in a model of a race. *Review of Economic Studies*, **52**, 193–209.

Harris, C., and Vickers, J. (1986). Racing with uncertainty. Mimeograph. Nuffield College, Oxford.

Harris, C., and Vickers, J. (1987). Racing with uncertainty. *Review of Economic Studies*, **54**, 1–21.

Katz, M. L., and Shapiro, C. (1987). R&D rivalry with licensing or imitation. *American Economic Review*, **77**, 402–20.

Lee, T., and Wilde, L. (1980). Market structure and innovation: A reformulation. *Quarterly Journal of Economics*, **43**, 429–36.

Loury, G. (1979). Market structure and innovation. *Quarterly Journal of Economics*, **93**, 395–410.

Phillips, A. (1971). *Technology and market structure: A study of the aircraft industry*. Lexington, Lexington, Mass.

Reinganum, J. (1983). Uncertain innovation and the persistence of monopoly. *American Economic Review*, **73**, 741–48.

Reinganum, J. (1985). Innovation and industry evolution. *Quarterly Journal of Economics*, **100**, 81–99.

Rosenberg, N. (1974). Science, innovation and economic growth. *Economic Journal*, 84, 90–108.

Shaked, A., and Sutton, J. (1986). Natural oligopolies. In *Economic Organizations as Games*, (ed. K. Binmore and P. Dasgupta). Blackwell, Oxford.

Ulph, A., and Ulph, D. (1988). Bargaining and the delay in innovation. *Scandinavian Journal of Economics*, **90**, 475–491.

Vickers, J. (1986). The evolution of industry structure when there is a sequence of innovations. *Journal of Industrial Economics*, **35**, 1–12.