

THE TIMING OF INNOVATION: RESEARCH, DEVELOPMENT, AND DIFFUSION

JENNIFER F. REINGANUM*

The University of Iowa

Contents

1. Introduction	850
2. Symmetric models	853
3. Asymmetric models	868
4. Licensing	884
5. Adoption and diffusion of innovations	893
6. Conclusions	904
References	905

*The financial support of the National Science Foundation (Grant No. SES-8216407), the Alfred P. Sloan Foundation and the Center for the Study of the Economy and the State is gratefully acknowledged.

1. Introduction

The analysis of the timing of innovation posits a particular innovation (or sequence of innovations) and examines how the expected benefits, the cost of R&D and interactions among competing firms combine to determine the pattern of expenditure across firms and over time, the date of introduction, and the identity of the innovating firm. In the case of a sequence of innovations, the expected lifetime of a given innovation and the pattern of technological leadership are also determined endogenously. Given that an innovation has been perfected, the extent and timing of its dissemination into use may be examined. Again this may depend upon a number of factors, including the existence of rival firms and institutions which may facilitate or retard the dissemination of innovations.

Section 2 addresses issues of innovation production in the context of symmetric noncooperative models. The important questions which have been examined in this context include: What is the aggregate noncooperative investment in research and development and how is it distributed across firms and across time? How many firms enter the race, and what is the resulting equilibrium date of innovation? The answers to these questions can then be compared with various benchmarks, such as their cooperative or surplus-maximizing counterparts. The typical outcome of these comparisons is that aggregate expenditure on R&D is too high relative to the cooperative optimum; there are too many firms and each invests too much. These problems can be attributed to two types of market failure. At each date, each firm considers only its own marginal benefit from investment and does not take into account the reduction it imposes on the expected value of the other firms' investments; consequently, each firm invests too much. Moreover, since entry into the race is unrestricted, it will continue until all expected profits are dissipated; entering firms do not take into account the loss of intertemporal efficiencies by rival firms when they decide to enter. Thus, the firms (collectively) forego intertemporal efficiencies which could be realized by investing at a lower rate over a longer planning horizon. By analogy to the problem of the commons, there is "over-grazing" in the industry. Comparison with the surplus-maximizing investment is more difficult, since the innovator is typically unable to appropriate the full surplus. This will tend to depress investment in research and development, making the comparison ambiguous. See Hirshleifer and Riley (1979) for a more complete discussion of this issue.

Section 3 considers asymmetric models, in which the issues of primary interest have been the effects on investment incentives provided by current market power, anticipated future innovation, and the possession of a technological advantage

(e.g. being closer to completion). Results in this area seem particularly sensitive to the presence or absence of technological uncertainty in the production of the innovation. When innovation is uncertain, a firm which currently enjoys a large market share will invest at a lower rate than a potential entrant, for an innovation which promises the winner a large share of the market. When innovation is deterministic, the opposite is true. Moreover, this same dichotomy extends to the case of a sequence of drastic innovations. When innovation is stochastic, the role of technological leader tends to circulate around the industry, while deterministic innovation results in a single, persistent technological leader. The effect of anticipated future innovation will also differ in these two cases; in the former, it reduces the value of winning the current race, since today's winner is likely to lose the next race; in the latter case, winning today is all-important, since today's winner also wins all future races. In a multi-stage game, the impact of having a technological lead is, as one might expect, to increase that firm's likelihood of winning the overall race, all else equal. Indeed, with all else equal and deterministic innovation, a very small head start is sufficient to ensure that the leading firm will win. However, if all else is not equal – that is, one firm anticipates greater benefits or faces lower costs – then an absolute disadvantage in terms of distance from completion will be overcome by the increased investment occasioned by the lagging firm's greater desire or effectiveness. When innovation is stochastic, although firms with a technological lead invest at a higher rate than their lagging rivals, a lucky laggard may still win the race.

Section 4 investigates the extent of dissemination of the innovation, where this dissemination is achieved through licensing. Several recent papers have examined optimal fixed-fee licensing for a patentholder selling to an oligopolistic industry. This patentholder may be an independent researcher or a joint venture of a subset of the industry's firms. One robust finding is that, in the absence of involuntary spillovers, firms who are not members of the research joint venture are left worse off as a result of innovation. In the case of an independent research lab, all members of the downstream industry are worse off; the patentholder reaps more than the total cost savings attributable to the innovation. Research joint ventures tend to restrict the dissemination of an innovation relative to an independent researcher; moreover, incentives to develop an innovation are weaker the larger is the joint venture. Thus, joint ventures would tend to restrict both the development and dissemination of an innovation. Of course, in evaluating the desirability of a joint venture, one would have to weigh against these restrictive tendencies any benefits (such as a reduction in the duplication of effort, or scale efficiencies) which might be generated. The motivation for licensing in these models is the cost savings which are generated by the use of the innovation, which can be at least partially appropriated by the patentholder, even under a less than optimal licensing agreement. When future technical advance is possible, another incentive for licensing arises; a firm with a superior technology can

license it to a rival firm in order to make further investment in R&D an unattractive strategy for the rival. Thus, the dissemination of the current technology (or the output of the current product) is enhanced at the expense of slower development of the next innovation.

Section 5 examines the timing of the adoption of an innovation, and summarizes recent work which provides alternative theoretical explanations for the observed diffusion of innovations into use. One explanation involves firms with differing initial priors about the profitability of the innovation. Information which accumulates over time is used to update these priors, and firms with more optimistic prior beliefs become “convinced” of an innovation’s profitability sooner than those with less optimistic beliefs. The combination of adjustment costs which decline with the period of adjustment and benefits of adoption which decline with the number of other adopters (such as arise when firms enjoy some degree of market power) results in firms planning to adopt an innovation in sequence. Since each firm ignores the fact that its adoption decreases the value of adoption for all subsequent adopters, each firm adopts too early from the perspective of the industry as a whole; that is, industry profits would be higher if firms coordinated their adoption plans, resulting in a slower diffusion of the innovation. Similar results obtain for standard specifications of the social good. Finally, if a firm perceives network externalities (its own benefits from adoption increase with the number of other adopters), it may find it optimal to wait until its more eager colleagues have adopted it. In this case, it may be that “excess inertia” exists, so that despite the fact that all firms prefer life with the innovation to life without it, no one is sufficiently eager to initiate the adoption process. If the process is begun, however, the existence of network externalities generates a “bandwagon effect”, since the value of adoption increases with the number of previous adopters.

Throughout this chapter I will focus on recent theoretical work, which is primarily game theoretic in nature. Recent empirical work in this area is surveyed in Chapter 19 of this Handbook. For more comprehensive surveys, see the monographs by Kamien and Schwartz (1980) and Stoneman (1983). In brackets, I note the source of each proposition. However, since I have made some modifications and re-interpretations, do not take these references too literally. To find out what the authors *really* claimed (and how they proved it) see the original papers.

An equally important – though less thoroughly investigated – aspect of technological change is the extent of innovation. This is typically examined in the context of cost reduction [e.g. Dasgupta and Stiglitz (1980a), Flaherty (1980), Telser (1982), Reinganum (1982b, 1983a), Spence (1984), Tandon (1984) and Katz (1986)], although an alternative measure for product innovation is the extent of product diversity [e.g. Spence (1976), Dixit and Stiglitz (1977) and Judd (1985b)]. The issue of the extent of innovation will not be dealt with here,

primarily due to space constraints. Another relevant strand of the literature which will not be discussed here is the work on learning by doing [e.g. Arrow (1962b), Spence (1981) and Fudenberg and Tirole (1983)], in which cost reduction is achieved as a result of production experience.

The timing of innovation has been examined in two basic paradigms: (1) a deterministic “auction” model, which can be traced to Barzel (1968) and Scherer (1967) and appears subsequently in Dasgupta and Stiglitz (1980b), Gilbert and Newbery (1982) and Katz and Shapiro (1985b); and (2) a stochastic “racing” model, which was analyzed for the single-firm case by Lucas (1971), and Kamien and Schwartz (1971), and subsequently was generalized by Fethke and Birch (1982) and Grossman and Shapiro (1986). Kamien and Schwartz (1972) generalized the single-firm model to include a partial account of the effects of rivalry, and the stochastic racing model appears as a full equilibrium model in Loury (1979), Dasgupta and Stiglitz (1980b), Lee and Wilde (1980) and Reinganum (1981a, 1982a). Since both these paradigms have been used repeatedly to address problems of innovation timing and related issues, both will be described in some detail for the case of symmetric games of research and development. This will make it easier to address related issues such as pre-emptive innovation and the persistence of monopoly, and licensing and the diffusion of innovations, which rely to a considerable extent upon variations of these two basic paradigms. We will begin with the most restrictive environment, and relax assumptions as we go along.

A third paradigm for examining investment in R&D is that used by Futia (1980), Hartwick (1982) and Rogerson (1982). This model assumes that there is one innovation per period and the innovator is determined as a random function of the firms’ investments. For example, if firm i invests x_i , it wins with probability $p_i = x_i / \sum x_j$. Thus, the game is not one of timing; it is rather a “contest” model. However, it does predict that firms overinvest in R&D and that each invests more the better is the patent protection. This sort of model has been examined in a laboratory experiment by Isaac and Reynolds (1985) who find that the experimental data are consistent with predicted Nash equilibrium play, and that the aforementioned predictions are indeed borne out in an experimental market.

2. Symmetric models

The first environment we wish to consider is one in which a particular invention is sought simultaneously by a number of identical potential inventors, which we will refer to as firms. The firm which produces the invention first is awarded a patent, which completely protects it from imitation or duplication. Invention is a costly activity, with the cost of invention by any given date being a decreasing

convex function of the time prior to invention. One justification for this cost function would be that the firm is optimally allocating its efforts on the invention throughout this period but suffers from classical diseconomies of scale; hence postponing the date of invention allows the firm to reduce its invention costs, but at a decreasing rate. For a computational example, see Kamien and Schwartz (1974); for additional arguments in favor of costs having this form, see Scherer (1967).

An alternative but essentially equivalent formulation assumes that a commitment of funds today determines the eventual date of invention according to a decreasing and convex function. Since invention is completely deterministic, the firm which commits the greatest expenditure (today) will obtain the invention first. A simple formal model which summarizes this situation is the following "auction" model. Let P denote the value of winning to the inventor; by assumption, all others receive 0. Assume that P is constant, and that all firms use the same discount rate r . If firm i were to spend x_i on research and development, it would complete the invention at date $T_i = T(x_i)$, where $T(\cdot)$ is decreasing and convex. A *strategy* for firm i is a bid x_i , and a *Nash equilibrium* is a vector of bids with the property that no firm wishes to unilaterally change its bid. The firm with the largest bid wins; if more than one firm wins, then the patent is awarded randomly among the winning firms. The key assumption about this model is that *no real resources are expended until the winner is determined*. That is, the firms "bid" what they would spend, a winner is determined, and only the winning firm actually develops the invention (spends the amount of its bid). If real resources were committed, no Nash equilibrium exists in pure strategies; however, the result of Proposition 1 below can be sustained as a Stackelberg equilibrium in which the firm with the highest bid moves first.

Under these assumptions, firm i 's payoff when the strategy (bid) vector is x is

$$V^i(x) = \begin{cases} P e^{-rT(x_i)} - x_i, & \text{if } x_i > x_j, \text{ for all } j \neq i, \\ \frac{1}{n(x)} [P e^{-rT(x_i)} - x_i], & \text{if } x_i = x_j > x_k, \\ & \text{for } n(x) - 1 \text{ } j \text{'s and remaining } k \text{'s}, \\ 0, & \text{if } x_i < x_j, \text{ for some } j \neq i. \end{cases}$$

Proposition 1 [Dasgupta and Stiglitz (1980b)]

A Nash equilibrium for this game involves two or more firms bidding x^* , where x^* is the largest value of x such that $P e^{-rT(x)} - x = 0$. Only one firm will actually invest x^* , and all firms make zero profits in equilibrium.

The intuition behind this result is clear; if any losing firm could pre-empt the winner and make positive profits, it would do so by bidding more. Hence, the winning firm must make zero profits given that it must spend its winning bid. The losing firms invest nothing and receive nothing.

A joint venture involving all firms in the industry would involve a single firm investing an amount x^{**} which maximizes $P e^{-rT(x)} - x$. To see that this involves less investment than under noncooperative play, it suffices to note that for $x > x^*$, $P e^{-rT(x)} - x < 0$. Thus, $x^{**} \leq x^*$, with a strict inequality so long as there is any investment level which yields strictly positive profits.

It was remarked earlier that a similar equilibrium configuration results if we specify one firm as a dominant firm, or first mover (Stackelberg leader). In this case, the dominant firm commits its expenditure at x^* , knowing that only this will keep rival firms from out-investing it and succeeding earlier. The lesson of this model is that one need not actually observe a "race" in progress; *potential* competition may make one run equally hard.

Alongside the auction model of invention is the second paradigm, in which the stochastic nature of invention is incorporated. In this framework, an investment of x_i still buys firm i a date of success, say $\tau_i = \tau(x_i)$, but now this date is regarded as random, indicating that success by any given date is only stochastically related to expenditure. A *strategy* in this framework is an investment level x_i . In earlier work on the management of research and development, Lucas (1971) and Kamien and Schwartz (1971) employed relatively general distribution functions. However, the addition of strategic rivals required some simplifications of other aspects of the problem, leading authors of recent game theoretic models [e.g. Loury (1979), Dasgupta and Stiglitz (1980b) and Lee and Wilde (1980)] to adopt the exponential distribution. That is, the probability that firm i is successful by date t is $\Pr\{\tau(x_i) \leq t\} = 1 - \exp\{-h(x_i)t\}$, where $h(x_i)$ is twice differentiable, strictly increasing, and satisfies:

- (i) $h(0) = 0 = \lim_{x \rightarrow \infty} h'(x)$,
- (ii) $h''(x) \geq (\leq) 0$ as $x \leq (\geq) \bar{x}$,
- (iii) $h(x)/x \geq (\leq) h'(x)$ as $x \geq (\leq) \tilde{x}$,

for some (\bar{x}, \tilde{x}) such that $0 \leq \bar{x} \leq \tilde{x} < \infty$.

Again let P denote the value of winning (assumed stationary) and r the common discount rate. Since success is not equivalent to winning, we need to compute the probability that firm i wins at any date t . Here the virtue of the exponential assumption becomes evident. Assuming that there are n firms whose processes are stochastically independent, the probability density that firm i wins

at t is

$$\begin{aligned} & \Pr\{\tau(x_i) \in [t, t + dt] \text{ and } \tau(x_k) > t \text{ for all } k \neq i\} \\ &= h(x_i) \exp\left\{-\sum h(x_j)t\right\} dt, \end{aligned}$$

where the summation is taken over $j = 1, 2, \dots, n$. For the exponential distribution, the expected date of invention is simply $1/\sum h(x_j)$. Let $a_i = \sum_{j \neq i} h(x_j)$ denote the aggregate rival hazard rate. Since the game is completely symmetric, we can write the payoff to any one firm as a function of its own investment x and the aggregate rival hazard rate a :

$$\begin{aligned} V(x, a) &= \int_0^\infty P e^{-rt} h(x) \exp\{-(a + h(x))t\} dt - x \\ &= \frac{Ph(x)}{a + h(x) + r} - x. \end{aligned} \quad (1)$$

A *best response function* for firm i to the aggregate rival hazard rate a is a function $\hat{x}(\cdot)$ such that for all a , $V(\hat{x}(a), a) \geq V(x, a)$ for all x . A *symmetric Nash equilibrium* for a given number of firms n will be denoted $x^*(n)$ and satisfies the relation $x^* = \hat{x}(a^*)$, where $a^* = (n - 1)h(x^*)$.

Proposition 2 [Loury (1979)]

As the number of firms in the industry increases, the equilibrium level of firm investment declines: $dx^*/dn < 0$.

This also implies that the expected success time for each firm rises, since $E\tau_i = 1/h(x^*(n))$. However, the time of invention, denoted $\tau(n)$, is the time of the *first* success: $\tau(n) = \min_i \{\tau_i\}$. With more firms there are more chances for early success, even though each one is less likely to yield early success. Under the following stability condition, Loury shows that the expected time to invention, $E\tau(n) = 1/nh(x^*(n))$, falls with an increase in n .

Assumption 1

Assume that $-h'(x^*)\hat{x}'(a^*) < 1$.

That is, in equilibrium a marginal increase in investment by any single rival firm causes the investment of a given firm to fall by a smaller amount. To

understand this interpretation, suppose that one rival firm increases its investment by dx ; then $da = h'(x^*)dx$, and the condition above says that $-h'(x^*)d\hat{x} < h'(x^*)dx$, or $-d\hat{x} < dx$.

Proposition 3 [Loury (1979)]

Under Assumption 1, increasing the number of firms reduces the expected date of invention.

Once equilibrium is characterized for an arbitrary number of firms, one can permit this number to be determined endogenously through entry (i.e. via a zero profit condition). It is easy to show that equilibrium expected profits decrease with increasing n ; thus entry continues unless $V(x^*, a^*) = 0$.

Consider the alternative problem of firms investing cooperatively and sharing equally in the reward. There are several reasons why the solution to this problem does not coincide with the noncooperative equilibrium. First, the joint venture may not value the innovation at P ; for example, antitrust regulation may prevent the venture from extracting the innovation's full value; on the other hand, perhaps the members of the joint venture can do better than P by a creative use of market segmentation. Second, noncooperative firms may not operate at the jointly optimal scale, and finally, the noncooperative equilibrium with entry may not result in the jointly optimal number of firms. If we assume that each firm values the innovation at P (whether investment is done noncooperatively or cooperatively), and that the research joint venture can also restrict entry, then it can be shown that for a given number of noncooperative firms, each invests too much. Moreover, in a noncooperative equilibrium with unrestricted entry, there will be too many firms relative to the joint optimum. Combining these two results with the fact that $dx^*/dn < 0$, while $d(nx^*)/dn > 0$ implies that in a free-entry noncooperative equilibrium there is too much investment in aggregate terms (relative to the joint optimum). This aggregate investment is produced by too many firms operating at less than efficient scale. Cooperation would involve fewer firms operating at efficient scale, but investing less in aggregate terms.

Proposition 4 [Loury (1979)]

Suppose that the function $V(x, (n-1)h(x))$ is a single-peaked (that is, increasing, then decreasing) function of x . Given a fixed number of firms, in industry equilibrium each firm invests more in R&D than is jointly optimal.

Each firm ignores its impact on its rivals' payoffs; consequently there is an excessive duplication of effort in the noncooperative equilibrium.

Proposition 5 [Loury (1979)]

If $\bar{x} > 0$, then unrestricted entry results in too many firms, each of which invests too little (relative to the cooperative optimum).

Lee and Wilde (1980) argued that the formalism which assumed that research and development expenditures were committed up front was inappropriate; since expenditure actually occurs over time, firms can stop investing once someone has succeeded. Instead, they permitted firms to choose a research intensity; once this intensity is fixed, the firm must either sustain this level of investment or cease investment altogether. Thus, the probability of success by any date t is still exponential, but the parameter now depends upon the *intensity* of research, rather than the *scale* of the lab. The new formula for firm i 's payoff is:

$$\begin{aligned} V^i(x) &= \int_0^\infty \exp\left\{-\sum h(x_j)\right\} [P e^{-rh(x_i)} - x_i] dt - F \\ &= \frac{Ph(x) - x}{a + h(x) + r} - F, \end{aligned} \quad (2)$$

where F denotes some possible fixed cost which is involved in entering the industry.

Lee and Wilde showed that this modification can substantially alter some of the model's implications. In particular, in equilibrium an increase in the number of rivals is associated with an increase in the intensity of research and development investment.

The fact that in Nash equilibrium $x^*(n) = \hat{x}(a^*(n))$, where $a^*(n) = (n-1)h(x^*(n))$, can be used to determine dx^*/dn . Under the following stability condition, Lee and Wilde show that a firm's noncooperative equilibrium research intensity is an increasing function of the number of firms.

Assumption 2

Assume that $1 - (n-1)h'(x^*)\hat{x}'(a^*) > 0$.

This requires that if a firm's competitors all increase their investments just enough to generate a unit increase in rivalry, then the remaining firm must respond with less than a full unit increase in investment. If all other firms increase their investment rates by the amount dx , the remaining firm must increase its investment rate by less than dx . To see this, note that an increase of dx by all others implies $da = (n-1)h'(x)dx$. The requirement that $\hat{x}'(a) < 1/(n-1)h'(x)$ is then equivalent to $d\hat{x} < da/(n-1)h'(x)$ or $d\hat{x} < dx$.

Proposition 6 [Lee and Wilde (1980)]

Under Assumption 2, as the number of firms in the industry increases, the equilibrium rate of investment per firm increases; a fortiori the aggregate rate of investment increases.

Since the expected invention date is $1/nh(x^*(n))$, an increase in the number of competing firms is associated with an earlier invention date on average; there are more firms, and each firm invests at a higher rate. Again it is easy to show that $V(x^*, a^*)$ decreases with increasing n . Free entry then occurs until equilibrium expected profits are zero.

Proposition 7 [Lee and Wilde (1980)]

Suppose that the function $V(x, (n-1)h(x))$ is a single-peaked (that is, increasing, then decreasing) function of x . Given a fixed number of firms, in industry equilibrium each firm invests at a higher rate than is jointly optimal.

Proposition 8 [Lee and Wilde (1980)]

If $\bar{x} > 0$, then unrestricted entry results in too many firms, each of which invests at too high a rate (relative to the cooperative solution).

Although Lee and Wilde have shown that the noncooperative firm's research intensity is an increasing function of n , this does not directly contradict the Loury result. Instead, one would like to compare total expected investment in the Lee and Wilde model with total lump-sum investment in the Loury model. It turns out that no general ranking emerges, but plausible examples can be devised in which expected total investment declines with an increase in the number of competing firms. Thus, it is quite possible that these results are consistent. The remaining inconsistency is the effect of n on each noncooperative firm's expected success date; in Loury (1979), success by any one firm is delayed (on average) by an increase in n , while in Lee and Wilde (1980) it is hastened. However, the *first* success date is hastened by an increase in n in both models. Since this is the only *observable* indicator of success, from a positive perspective it is immaterial which of these two models is the more empirically representative. Similarly, if one were to interpret P as both the private and the social value of the innovation, then normative prescriptions are essentially the same for both models; restrict entry to the optimal number of firms, then adjust the patent value to eliminate over-investment.

Mortensen (1982) has shown that one aspect of the externality which competing firms impose on one another can be internalized by the following institution:

the winning firm receives the value P less a compensation paid to each losing firm which is equal to the foregone value of continued play. This institution induces noncooperative firms to select the optimal cooperative investment level *given the number of firms* n . Since this institution raises individual firm profits relative to the noncooperative equilibrium without this institution, more firms would prefer to enter the industry when this institution is in place. Thus, the cooperating firms must also be able to limit entry in order to fully internalize the externality. Since the foregone value of continued play is stationary in this model, equilibrium under this institution is relatively easy to characterize. Although this argument works equally well for asymmetric firms, the case of symmetric firms is simpler to describe and has been independently and more extensively analyzed by Stewart (1983).

Suppose that the winning firm receives the value P but must compensate the remaining $n - 1$ firms in the amount $P(1 - \sigma)/(n - 1)$ each; thus the winning firm retains the amount σP . In this case, one can write the expected profit to firm i if it invests at rate x while the aggregate rival hazard rate is a , as

$$V(x, a) = \frac{P[\sigma h(x) + a(1 - \sigma)/(n - 1)] - x}{a + h(x) + r} - F. \quad (3)$$

At a symmetric Nash equilibrium for this game, which we will denote by $x^*(n, \sigma)$, it must be that

$$\begin{aligned} \frac{\partial V(x^*, a^*)}{\partial x} &= \frac{(nh(x^*) + r)(P\sigma h'(x^*) - 1) - (Ph(x^*) - x^*)h'(x^*)}{(nh(x^*) + r)^2} \\ &= 0. \end{aligned} \quad (4)$$

For comparison, joint profits $nV(x, (n - 1)h(x))$ are maximized for $x^{**}(n)$ such that

$$\frac{(r + nh(x^{**}))(Ph'(x^{**}) - 1) - (Ph(x^{**}) - x^{**})nh'(x^{**})}{(nh(x^{**}) + r)^2} = 0. \quad (5)$$

Thus, $x^*(n, \sigma) = x^{**}(n)$ if

$$\sigma = \sigma^*(n) \equiv \frac{Ph'(x^{**}(n)) + n - 1}{nPh'(x^{**}(n))}.$$

That is, there is a winner's share $\sigma^*(n)$ which induces noncooperative firms to invest at the cooperative level.

If members of an industry can credibly set up such an institution for sharing the reward for innovation, then each member has incentives for (noncooperatively) choosing the cooperative investment level. Stewart (1983) interprets this as a model of imperfect patent protection (as opposed to cooperative innovation), but in this case the reward is unlikely to be P ; it is likely that because of nonappropriability, some of the value of the innovation will be dissipated.

If one is willing to assume that the social value of the innovation is also P , then all of the above comparisons of noncooperative equilibrium and joint optimality are also applicable to the comparison between noncooperative equilibrium and social optimality.

The excessive investment in research which is implied by these models arises out of two sources. First, each firm wants to win the race, while society typically has no particular preferences regarding the identity of the winner, so long as there is one; this results in too much investment for a given number of firms. Second, because there is unrestricted access to the common pool of undiscovered innovations, too many firms will compete.

Reinganum (1981a, 1982a) undertook to generalize the earlier decision-theoretic work of Kamien and Schwartz (1972) to include explicit game-theoretic interactions among rivals. To this end, she posited that firms are free to react instantaneously to a number of features of the economic environment, including time, whether or not rival firms have already succeeded, their own and rival firms' accumulated investment. As in Lee and Wilde (1980), firms are assumed to be able to respond to a rival's success by ceasing investment; that is, we focus on the effects of flow expenditures upon the likelihood of success. However, unlike the Lee and Wilde formulation, the rate of expenditure is not restricted to be constant through time. Instead, firms may adjust the rate of expenditure in response to elapsed time and state variables which summarize rival progress. Reinganum (1982a) also considers the case of imperfect patent protection, which previous work eschewed.

Assume that a given number of firms n are competing to perfect a particular invention. Firm i succeeds if it perfects the invention, but firm i wins only if it succeeds before any other firm. If firm i wins, then firm i is designated the *innovator* and receives the "leader's" payoff P_L . However, assume that immediate reverse engineering may be possible and thus rival firms may also receive some benefit when firm i wins; the rival firms receive the *imitator's* or "follower's" payoff $P_F \leq P_L$. Use of the terms "leader" and "follower" does not connote any behavioral differences.

Each firm accumulates knowledge relevant to the innovation by investing resources on knowledge acquisition. This knowledge accumulates according to the differential equation $\dot{z}_i(t) = u_i(t, z(t))$, where $z_i(t)$ denotes firm i 's knowledge stock at t , $z(t) = (z_1(t), \dots, z_n(t))$ and $u_i(t, z)$ denotes firm i 's rate of knowledge acquisition at (t, z) . We assume that $z_i(0) = 0$. The date of successful

innovation is a random function of the amount of accumulated knowledge. Specifically, we assume that the probability of success given a knowledge stock of z or less is $F(z) = 1 - \exp\{-\lambda z\}$.

Thus, the amount of knowledge needed to succeed is exponentially distributed with mean $1/\lambda$. Since knowledge is accumulated over time, the distribution of firm i 's random success time t_i is

$$\Pr\{t_i \leq t\} = F(z_i(t)) = 1 - \exp\{-\lambda z_i(t)\},$$

and the conditional probability density of success, given no success to date is

$$\Pr\{t_i \in (t, t + dt] | t_i > t\} = \lambda u_i(t, z(t)) dt.$$

Thus, the conditional density of success depends only upon current investment. In the stochastic racing papers discussed above the date of invention was assumed to be exponentially distributed. In the current model, the amount of time needed for success need not be exponentially distributed. It will be so, of course, in the case of a constant rate of investment.

To complete the model, let $c(u) = (1/2)u^2$ represent the cost of acquiring new knowledge at the rate u . Let $[0, T]$ represent the planning horizon, and let r denote the discount rate. The assumption of a finite planning horizon is used primarily in order to allow us to use dynamic programming "backward". It may also be representative of many research situations in which funding will be terminated if concrete results are not forthcoming by a given date.

Previous studies have assumed that the patent value was constant independent of the date of success; more generally, one could argue that this patent value might grow over time as additional uses are discovered. Alternatively, some sort of exogenous obsolescence may be applicable, so that the patent value might decline over time. Thus, let g denote the rate of growth or decline of the patent value so that the values of winning and losing, respectively, at date t are (in present value terms) $P_L e^{gt}$ and $P_F e^{gt}$.

There are several possible formulations of strategies for differential games. The most commonly used are open-loop and feedback strategies. Open-loop strategies depend only upon the current date and the initial conditions of the problem. They may be obtained by applying standard optimal control arguments to the problem for each firm, and then solving the resulting linked systems of ordinary differential equations which characterize the equilibrium. Logically, open-loop strategies have the characteristic of precommitment; that is, one solves the problem from the initial conditions taking the time path of others' strategies as given. Feedback strategies are decision rules which are permitted to depend upon the current date and state variables, but not on the initial conditions. They are obtained by solving the problem from arbitrary (date, state) pairs; that is, by

dynamic programming. Thus, feedback equilibria (Nash equilibria in feedback strategies) will embody the no-commitment assumption associated with subgame perfection [Selten (1975)]; noncredible threats about what a firm will do off the equilibrium path are ruled out. In situations in which firms have the information and the flexibility envisioned here, feedback equilibrium is the preferred solution concept. Although we will show below that our assumption that the conditional density of success depends only on current investment renders this distinction moot in this particular case, the method described below applies more generally.

A *strategy* for firm i will be a function $u_i(t, z)$, where z denotes the vector of state variables $z = (z_1, \dots, z_n)$. Given the strategy vector $u(t, z) = (u_1(t, z), \dots, u_n(t, z))$, one can solve for the trajectories of the state variables by solving the system of ordinary differential equations:

$$\dot{z}_i = u_i(t, z), \quad z_i(0) = 0, \quad i = 1, 2, \dots, n.$$

Expected profits for firm i consist of three terms: costs are paid so long as no firm has yet succeeded. The probability that no firm has succeeded by t is

$$\Pr\{t_i > t \text{ for all } i\} = \exp\{-\lambda \sum z_i(t)\}.$$

If firm i wins at t , then it receives $P_L e^{gt}$, while if firm i loses at t , it receives $P_F e^{gt}$. The probability that firm i wins at t is

$$\Pr\{t_i \in (t, t + dt], t_j > t \text{ for all } j\} = \lambda u_i(t, z) \exp\{-\lambda \sum z_j(t)\}.$$

If $a_i(t, z) \equiv \sum_{j \neq i} \lambda u_j(t, z)$ denotes the aggregate rival hazard rate at t , then the probability that firm i loses at t (i.e. the probability that any rival firm wins at t) is

$$\begin{aligned} \Pr\{t_k \in (t, t + dt] \text{ for any } k \neq i, t_j > t \text{ for all } j\} \\ = a_i(t, z) \exp\{-\lambda \sum z_j(t)\}. \end{aligned}$$

Combining these terms and discounting to the present implies that firm i 's payoff for any given strategy u_i and aggregate rival hazard rate a_i is

$$J^i(u_i, a_i) = \int_0^T e^{-rt} 3 \exp\{-\lambda \sum z_j\} [P_L e^{gt} \lambda u_i + P_F e^{gt} a_i - (1/2)(u_i)^2] dt.$$

A *Nash equilibrium* is a vector of strategies $u^*(\cdot, \cdot)$ such that, for $i = 1, 2, \dots, n$, $J^i(u_i^*, a_i^*) \geq J^i(u_i, a_i^*)$ for all u_i .

Ideally we would like to solve this problem for arbitrary rates g of growth or decline in the patent value; failing this, we should at least select an interesting value for g . One interesting and computationally convenient value for g is $g = r$;

that is, the patent value grows at the rate of discount. The assumption of growth in the patent value seems plausible and, although this is an extremely high growth rate, it gives us an idea of the qualitative impact of patent value growth upon equilibrium investment and does so in a tractable manner. Qualitative features of the equilibrium should be the same for lower rates of growth in the patent value. One can conjecture on the basis of these results about the impact of declining patent values as well. Thus, from here on we assume $g = r$.

To characterize the Nash equilibrium for this game by using standard dynamic programming techniques, we must be concerned about two things: first, there should be a unique solution to the system $\dot{z} = u(t, z)$ through the boundary condition $z(t) = z$ for each $(t, z) \in [0, T] \times [0, \infty)^n$. Second, in order to assert the necessity of the partial differential equation of dynamic programming, we need the value functions to be continuously differentiable in (t, z) . Both these requirements place restrictions upon the admissible set of strategies. It is relatively straightforward to find sufficient restrictions upon the strategies to guarantee that a unique solution exists for the system $\dot{z} = u(t, z)$. However, weak sufficient conditions for the continuous differentiability of the equilibrium value functions are unknown at this time. Our method of dealing with this problem is to forge ahead assuming sufficient smoothness and to argue subsequently that the solution we obtain is in fact a Nash equilibrium for the specified game.

Proposition 9 [Reinganum (1982a)]

A feedback Nash equilibrium strategy for each firm is

$$u_i^*(t, z; P_L, P_F) = \frac{2\lambda P_L(P_L - P_F)(n-1)e^{rt}}{(2n-1)P_L - [P_L + 2(n-1)P_F]\exp\{m(t)\}}.$$

The feedback Nash equilibrium payoff to each firm is

$$J^i(u_i^*, a_i^*) = P_L - \frac{2(P_L - P_F)P_L(n-1)}{(2n-1)P_L - [P_L + 2(n-1)P_F]\exp\{m(0)\}},$$

where $m(t) = (P_L - P_F)(n-1)\lambda^2(e^{rt} - e^{rT})/r$.

Notice that due to the "memorylessness" property of the exponential distribution of required knowledge—that is, the expected amount of additional knowledge given no success is independent of accumulated knowledge—the equilibrium strategies are independent of the state variables z . Since this is due to a special feature of the exponential distribution, it is not likely to carry over to games which involve alternative distribution functions. It is in some ways undesirable,

since no firm can gain a convincing lead on its rivals, but it also likely accounts for the relative tractability of this model.

Proposition 10 [Reinganum (1982a)]

For $P_L > P_F$, an increase in P_L stimulates each firm to acquire knowledge at a higher rate, while an increase in P_F causes each firm to reduce its equilibrium rate of knowledge acquisition.

For the case of perfect patent protection, the equilibrium strategies are

$$u_i^*(t, z; P, 0) = \frac{2\lambda(n-1)Pe^{rt}}{2n-1 - \exp\{P\lambda^2(n-1)(e^{rt} - e^{rT})/r\}}.$$

Proposition 11 [Reinganum (1982a)]

For the case of perfect patent protection, an increase in n increases the equilibrium rate of investment for each firm. Therefore an increase in n unambiguously decreases the expected time till innovation.

Next consider the limiting case $P_L = P_F$. In this case,

$$u_i^*(t, z; P_F, P_F) = \frac{2P_F\lambda e^{rt}}{2 - (2n-1)P_F\lambda^2(e^{rt} - e^{rT})/r}.$$

When patent protection is imperfect, the impact of an increase in the number of rival firms is much more complicated. In this case, it seems plausible that the payoffs P_L and P_F should both depend upon n . The determination of $P_L(n)$ and $P_F(n)$ might be regarded as the outcome of a subsequent licensing or oligopoly game. In order to determine the effect of increasing rivalry upon the equilibrium rate of investment, we now need to know the sign of

$$\frac{du_i^*}{dn} = \frac{\partial u_i^*}{\partial n} + \frac{\partial u_i^*}{\partial P_L} P'_L(n) + \frac{\partial u_i^*}{\partial P_F} P'_F(n).$$

Under the most plausible circumstances, the sign of this expression is ambiguous, and will depend upon the specific nature of the institution or process which determines the relative payoffs to innovator and imitators (or licensor and licensees). For illustrative purposes, we could consider the case in which rewards are completely nonappropriable and are generated by a symmetric Cournot oligopoly. If P is the value of a monopoly on the innovation, then $P_L(n) =$

$P_F(n) = 4P/(n+1)^2$. In this case, it is straightforward to show that increasing n results in a uniform decrease in the rate of investment; that is, $du_i^*/dn < 0$ for all t .

A final comparison highlights the importance of the extent of appropriability in the determination of equilibrium investment. In the model developed above, there are no fixed costs associated with entering the industry. Thus, if entry is unrestricted, the equilibrium number of firms will be infinite. When patent protection is perfect:

$$\lim_{n \rightarrow \infty} u_i^*(t, z; P, 0) = P\lambda e^{rt}.$$

On the other hand, when patent protection is completely ineffective, then both $P_L(n)$ and $P_F(n)$ approach 0 as n gets large. In this case,

$$\lim_{n \rightarrow \infty} u_i^*(t, z; P_L(n), P_F(n)) = 0.$$

When patent protection is ineffective, no firm finds research and development a worthwhile undertaking.

In a companion piece, Reinganum (1981a) compares noncooperative and cooperative investment, and considers another form of nonappropriability which is characterized by spillovers in knowledge. To focus on these issues, assume that $n = 2$ and that patent protection is perfect.

Cooperation among firms involves coordinating research strategies, but it may also involve the cooperative exchange of knowledge. Thus, firms are able to operate on the lower portions of their cost curves while still generating the same aggregate amount of new knowledge. For cooperative firms, $\dot{z}_i = u_i(t, z) + \gamma u_j(t, z)$, where γ represents the fraction of new knowledge which can be shared with rival firms; γ may be less than unity because some knowledge may not be transferable or there may be some duplication in the knowledge acquired. Since the problem is symmetric, we can ignore subscripts: $z_i = z$, $u_i(t, z) = u(t, z)$, and $\dot{z}_i = \dot{z} = (1 + \gamma)u(t, z)$. The payoff to the joint venture is the sum of the individual firms' payoffs (with the understanding that $\dot{z} = (1 + \gamma)u(t, z)$; knowledge as well as profits are shared). Thus, joint profits are

$$J(u) = \int_0^T e^{-2\lambda z} [2P\lambda u(t, z) - e^{-rt}(u(t, z))^2] dt.$$

Proposition 12 [Reinganum (1981a)]

The cooperative rate of knowledge acquisition is

$$u^{**}(t, z) = \frac{(1 + \gamma)P\lambda e^{rt}}{1 - P\lambda^2(1 + \gamma)^2(e^{rt} - e^{rT})/r}.$$

The joint payoff is

$$J(u^{**}) = P - \frac{P}{1 - P\lambda^2(1 + \gamma)^2(1 - e^{rT})/r}.$$

In order to compare the timing of innovation under cooperative and noncooperative behavior, we need to compare the individual rates of knowledge acquisition. For noncooperative and cooperative rivals these are, respectively,

$$u^*(t, z) = \frac{2P\lambda e^{rt}}{3 - \exp\{m(t)\}}$$

and

$$\dot{z}^{**}(t; \gamma) = (1 + \gamma)u^{**}(t, z) = \frac{(1 + \gamma)^2 P\lambda e^{rt}}{1 - (1 + \gamma)^2 m(t)},$$

where $m(t) \equiv P\lambda^2(e^{rt} - e^{rT})/r$.

Proposition 13 [Reinganum (1981a)]

For $\gamma = 0$, $u^*(t, z) \geq \dot{z}(t; 0) = u^{**}(t, z)$ with equality only at $t = T$. That is, noncooperative rivals will (on average) succeed sooner than cooperative firms who are unable to share knowledge.

The cooperative rate of knowledge acquisition $\dot{z}^{**}(t; \gamma)$ is an increasing function of γ , so for $\gamma > 0$, it is typically the case that the rate of knowledge acquisition is higher for noncooperative rivals over the first portion, but higher for cooperative rivals over the latter portion of the planning horizon.

There are at least two ways in which a rival firm can benefit from a particular firm's investment in research and development. By imitating (or licensing) the innovation, the rival may be able to capture some of the benefits. However, even with perfect patent protection, the rival may benefit if some of the knowledge which is generated spills over to the rival firm. In this case, $\dot{z}_i = u_i(t, z) + \rho u_j(t, z)$, where $\rho \in [0, 1]$ denotes the extent of knowledge spillovers. For simplicity, we will suppose that spillovers in knowledge are complete; anything learned by firm i is also learned by firm j . Moreover, suppose that there is effectively no duplication. Then $\rho = 1$ and firm i 's payoff is

$$J^i(u) = \int_0^T e^{-\lambda(z_1 + z_2)} [P\lambda(u_1 + u_2) - e^{-rt}(1/2)(u_1)^2] dt.$$

Analysis of this case proceeds as before, yielding the following results.

Proposition 14 [Reinganum (1981a)]

Let $u_i^*(t, z; \rho)$ denote the feedback Nash equilibrium strategies. For $\rho = 1$,

$$u_i^*(t, z; 1) = \frac{P\lambda e^{rt}}{1 - 3P\lambda^2(e^{rt} - e^{rT})/r}.$$

The feedback Nash equilibrium payoffs are

$$J^i(u^*(t, z; 1)) = P/2 - \frac{P/2}{1 - 3P\lambda^2(1 - e^{rT})/r}.$$

Several comparisons are possible. The problem for a central planner is equivalent to that of a joint venture with the substitution of the social value Q for the private value P . Thus, the socially optimal path is as in Proposition 12 with Q substituted for P .

Proposition 15 [Reinganum (1981a)]

Suppose that the social value of the innovation Q exceeds $P/2$, all knowledge is transferable ($\gamma = 1$) and spillovers are complete ($\rho = 1$). Then the noncooperative rate of knowledge acquisition is less than is socially optimal. Consequently, innovation will be delayed on average relative to the socially optimal date. Innovation by Nash rivals will occur later on average than the cooperative date.

Recall that for the opposite extreme case ($\gamma = 0, \rho = 0$), Nash rivals could be expected to innovate at an *earlier* date than cooperative firms. Thus, again we see the crucial effect of appropriability; that market structure which most promotes innovation depends critically upon the extent of spillovers.

It is also interesting to note that the existence of spillovers need not adversely affect the timing of innovation under noncooperative play. While it is true that each firm invests at a lower rate in the presence of spillovers, each also benefits from the investment of the other. For some parameter values, the existence of spillovers results in stochastically earlier innovation.

3. Asymmetric models

A topic of long-standing interest in industrial organization is the effect of current monopoly power upon a firm's incentives to engage in innovative activity [e.g. Schumpeter (1942)]. Arrow (1962a) argued that for a drastic innovation (one which leaves the inventor a monopolist), an incumbent monopolist would have less incentive to invent than would an inventor who currently has no share in the

market. Gilbert and Newbery (1982) use the auction model to examine this question when an incumbent firm is faced with potential entrants who also compete for the innovation. They assume that each competing firm enters a bid which represents the maximum amount that firm will spend on research and development. The firm which bids the most is conceded to be the winner, and is required to invest the amount of its bid. For simplicity, assume that if the current monopolist, or incumbent, ties with one or more potential entrants, then the patent is awarded to the incumbent. If the incumbent wins with a bid of x_i , it receives $P^m e^{-rT(x_i)} - x_i$ in present value terms, where $T(x_i)$ is the date of completion and P^m is the capitalized value of the innovation if the relevant product market is monopolized. If a potential entrant wins with a bid of x_e , then the firms must share the market somehow. It is assumed that the entrant receives $P^e e^{-rT(x_e)} - x_e$, while the incumbent receives $P^i e^{-rT(x_e)}$. In this case P^e and P^i represent the entrant's and incumbent's portions of the value of the innovation. Under most plausible specifications, $P^m \geq P^e + P^i$; that is, there is some dissipation of rents when the market is noncooperatively shared rather than monopolized.

Proposition 16 [Gilbert and Newbery (1982)]

If $P^m \geq P^e + P^i$, then the current incumbent will win the bidding game with a bid of x^* , where x^* is the largest solution of

$$P^e e^{-rT(x)} - x = 0. \quad (6)$$

Thus, a firm which currently enjoys monopoly power will pre-emptively patent the innovation and persist as a monopolist.

Following the same line of argument as in Proposition 1, competing potential entrants will bid up to x^* as described above. If the incumbent is willing to bid at least x^* , then it will win. By bidding less, the incumbent would receive $P^i e^{-rT(x^*)}$; by matching the potential entrants' bid, the incumbent would receive $P^m e^{-rT(x^*)} - x^*$. The latter option is preferred to the former if and only if $P^m e^{-rT(x^*)} - x^* \geq P^i e^{-rT(x^*)}$. Substituting from equation (6) for x^* , this reduces to $[P^m - P^i - P^e] e^{-rT(x^*)} \geq 0$.

Thus, the incentive for pre-emptive patenting and persistent monopoly arises from the dissipation of industry profits which one anticipates will accompany a less concentrated market structure. Notice that the incumbent firm need not use the innovation (e.g. implement the new technology or produce the new product); even a product or technology which is inferior to the incumbent's current one will elicit pre-emptive investment from the incumbent. In this case, one may find "sleeping patents", which are used solely to preserve the incumbent's monopoly position. While the use of inferior technologies is inefficient, so is monopoly; the

industry composed of two competing firms, one of which employs an inefficient technology, may be welfare-preferred to the more concentrated but cost-efficient industry.

Gilbert and Newbery remarked that in the event that the above inequality is strict, an incumbent with a relative cost disadvantage in innovation would still pre-empt the potential entrant. Salant (1984) argued that this result is based on the assumption that there is no possibility of ex post licensing. If licensing is permitted, Salant shows that the firm which is most efficient at innovation will always win the patent, but may sell it to the other firm if the other firm is a more efficient producer. In any event, optimal licensing will still result in a monopolized market.

Katz and Shapiro (1987) examine pre-emption in a somewhat more general version of the auction model. They permit two active firms from the outset, and allow for the possibility of licensing or imitation following innovation. In addition, they envision an exogenous decline in the costs of developing an innovation due to ongoing and freely available basic research. Their analysis involves lengthy arguments, and we briefly summarize their model and results here.

Let π_0^i denote firm i 's flow profits prior to development of the innovation. If firm i develops the innovation, its profits become π_i^i and its rival's profits become π_i^j . Let $\pi_i = \pi_i^i + \pi_i^j$ denote industry profits when firm i wins, and suppose without loss of generality that $\pi_1 \geq \pi_2$.

Each firm has two incentives to innovate. First, firm i has an incentive to win because its profits (are assumed to) rise if it develops the innovation; that is, $\pi_i^i \geq \pi_0^i$. Second, firm i has an incentive to win (to avoid losing) because its profits are (assumed to be) higher when it wins than when the rival wins; that is, $\pi_i^i \geq \pi_j^i$. The former incentive, $\pi_i^i - \pi_0^i$, is the "stand-alone" incentive, while the latter incentive, $\pi_i^i - \pi_j^i$, is a measure of the "incentive to pre-empt". Katz and Shapiro show that if firm 1 has both a larger stand-alone incentive and a larger incentive to pre-empt, then firm 1 will win the race. If firm 2 has the larger stand-alone incentive, but firm 1 has the larger incentive to pre-empt, then either firm may win.

Reinganum (1983b) addresses the question of the effect of current monopoly profit upon an incumbent firm's incentives to invest in research and development in the context of the stochastic racing model. This is done to compare the results with those of Gilbert and Newbery for the auction model and to provide a theoretical explanation of some stylized facts about the sources of innovation.

According to Scherer (1980, pp. 437–438):

There is abundant evidence from case studies to support the view that actual and potential new entrants play a crucial role in stimulating technical progress, both as direct sources of innovation and as spurs to existing industry

members...new entrants contribute a disproportionately high share of all really revolutionary new industrial products and processes.

Gilbert and Newbery's analysis captures some of this in the sense that potential entrants do act as a spur to the current incumbent; on the other hand, potential entrants do not contribute directly. In the context of the Lee and Wilde (1980) stochastic racing model, Reinganum (1983b) shows that when the first successful innovator captures a sufficiently high share of the post-innovation market, then in a Nash equilibrium the incumbent firm invests less on a given project than does the potential entrant, or challenger. Thus, the incumbent is less likely to be the innovator than is the challenger.

The intuition for this result is straightforward, at least for the case in which the innovation is drastic; that is, when the innovator captures the entire post-innovation market. When innovation is uncertain, the incumbent firm receives flow profits before successful innovation. This period is of random length, but is stochastically shorter the more the incumbent (or the challenger) invests. The incumbent has relatively less incentive than the challenger to shorten the period of its incumbency.

This model provides a framework in which equilibrium play generates the stylized facts mentioned above: potential entrants stimulate progress both through their own investment and by provoking incumbents to invest more. In equilibrium, potential entrants contribute a disproportionate share of large innovations.

To illustrate this model and its results, consider a cost-reducing innovation in an industry with constant returns to scale. Let \bar{c} denote the incumbent firm's current unit costs, and let c be the unit cost associated with the new technology. Let R be the current flow rate of profit; let $\Pi(c)$ denote the present value of monopoly profits under the new technology, which is also the value of the reward to the incumbent if it invents the new technology; finally, let $\pi_I(c)$ and $\pi_C(c)$ denote the present value of Cournot-Nash profits to the incumbent and challenger, respectively, if the challenger invents the new technology and the incumbent retains use of the current technology.

We will assume that the functions $\Pi(\cdot)$, $\pi_I(\cdot)$ and $\pi_C(\cdot)$ are continuous and piecewise continuously differentiable; $\Pi(\cdot)$ and $\pi_C(\cdot)$ are nonincreasing, while $\pi_I(\cdot)$ is nondecreasing. The innovation will be termed *drastic* if $c \leq c^0$, where c^0 is the largest value of c such that $\pi_I(c) = 0$. That is, c^0 is the largest unit cost for the challenger which induces the incumbent to leave the post-innovation product market.

The assumption of constant returns to scale is important because it ensures that output is zero when profits are zero. Thus, for drastic innovations, the challenger becomes a monopolist and $\Pi(c) = \pi_C(c)$. Note that $\Pi(c) \geq \pi_I(c) + \pi_C(c)$ with a strict inequality whenever the innovation is not drastic. The

assumption that $c < \bar{c}$ ensures that $\Pi(c) > R/r$; that is, the present value of post-innovation monopoly profits exceeds the present value of pre-innovation monopoly profits. Moreover, $R/r > \pi_I(c)$ for all $c < \bar{c}$; this follows from the fact that $R/r = \Pi(\bar{c}) > \pi_I(\bar{c}) \geq \pi_I(c)$ for $c < \bar{c}$.

Let x_i , $i = I, C$ denote the rate of investment for the incumbent and the challenger, respectively. This generates the hazard rate $h(x_i)$ for firm i . Let a_i , $i = I, C$ denote the rival hazard rate for firm i ; for instance, $a_I = h(x_C)$. For simplicity we will assume that the hazard function $h(\cdot)$ is twice continuously differentiable with $h'(\cdot) > 0$, $h''(\cdot) < 0$, $h(0) = 0$, and satisfies the conditions $\lim_{x \rightarrow 0} h'(x) = \infty$ and $\lim_{x \rightarrow \infty} h'(x) = 0$.

Assuming that patent protection is perfect, the race terminates with the first success. The expected profit to the incumbent as a function of its own investment rate and its rival's hazard rate is

$$\begin{aligned} V^I(x_I, a_I) &= \int_0^\infty e^{-rt} e^{-(h(x_I) + a_I)t} [h(x_I)\Pi(c) + a_I\pi_I(c) + R - x_I] dt \\ &= \frac{h(x_I)\Pi(c) + a_I\pi_I(c) + R - x_I}{r + h(x_I) + a_I}. \end{aligned}$$

The challenger's payoff is analogous:

$$\begin{aligned} V^C(x_C, a_C) &= \int_0^\infty e^{-rt} e^{-(h(x_C) + a_C)t} [h(x_C)\pi_C(c) - x_C] dt \\ &= \frac{h(x_C)\pi_C(c) - x_C}{r + h(x_C) + a_C}. \end{aligned}$$

The differences between these payoffs are due to the fact that the incumbent receives a flow payoff of R so long as no one has succeeded, and the incumbent receives a (possibly positive) payoff $\pi_I(c)$ if the challenger wins.

As usual, a *strategy* for firm i is an investment rate x_i ; a *best response function* for firm i is a function $\hat{x}_i(a)$ such that for all a , $V^i(\hat{x}_i(a), a) \geq V^i(x_i, a)$ for all x_i . A *Nash equilibrium* is a pair (x_I^*, x_C^*) such that $x_I^* = \hat{x}_I(a_I^*)$, where $a_I^* = h(x_C^*)$, and $x_C^* = \hat{x}_C(a_C^*)$, where $a_C^* = h(x_I^*)$. That is, each firm plays a best response against the other's strategy.

Proposition 17 [Reinganum (1983b)]

There exists a Nash equilibrium pair $(x_I^*(c, R), x_C^*(c, R))$; $x_i^*(c, R)$ is continuous in (c, R) for $i = C, I$.

Proposition 18 [Reinganum (1983b)]

The incumbent's best response function is upward-sloping; thus the existence of the challenger provokes the incumbent to invest more than it otherwise would. If the innovation is drastic and $R > 0$, then in a Nash equilibrium the incumbent invests less than the challenger. That is, $x_I^*(c, R) < x_C^*(c, R)$.

An immediate corollary of Proposition 18 and the continuity of the Nash equilibrium strategies in the parameters (c, R) is that if $R > 0$, then there exists an open neighborhood of c^0 (which may depend on R), denoted $N(c^0; R)$, such that if the technology is not drastic, but $c \in N(c^0; R)$, then $x_I^*(c, R) < x_C^*(c, R)$. That is, there is a set of nondrastic innovations for which the incumbent firm will still invest less than the challenger. Since the incumbent invests less than the challenger, the challenger is more likely to win the asymmetric patent race. Thus, one would empirically observe that challengers contribute disproportionately more large innovations.

This model has been extended to an arbitrary number of firms and a sequence of innovations in Reinganum (1985), in order to generate a model of the Schumpeterian "process of creative destruction". For simplicity, we assume that each innovation is drastic. Then the model remains symmetric among all challengers.

Consider a market in which an incumbent monopolist competes with $n - 1$ identical challengers for a new innovation. The firms are assumed to be symmetric in all other respects; that is, they face the same innovation production possibilities. Each innovative success initiates a new stage; within each stage firms compete for the next generation. The game with t stages to go is constructed recursively from shorter horizon games under the assumption of subgame perfect Nash equilibrium play.

Nash equilibria are found to be symmetric among the challengers, with each challenger investing more than the incumbent. Thus, the incumbent firm enjoys temporary monopoly power, but is soon overthrown by a more inventive challenger.

The basic model is now familiar. It is essentially that of Lee and Wilde, except that we now specify the values of winning and losing the current race as v^W and v^L , respectively. These represent the values of continuing on in a Nash equilibrium fashion when one fewer innovations remain. The values are ultimately endogenous to the model, but at each stage they may be treated parametrically because they are independent of actions taken in the current stage.

Thus, for any given stage, the game is summarized by n , the number of competing firms; x_i , the investment rate of firm i ; F , a fixed cost of entry; $h(\cdot)$, the hazard function; R , the current profit flow to the incumbent; r , the common discount rate; and v^W and v^L , the continuation values.

Let $a_i = \sum_{j \neq i} h(x_j)$ be the aggregate rival hazard rate. Again the payoff to firm i can be written as a function of its own research intensity x_i and its aggregate rival hazard rate a_i . Suppose, without loss of generality, that firm 1 is the incumbent. Then

$$\begin{aligned} V^1(x_1, a_1) &= \int_0^\infty e^{-rt} e^{-(h(x_1)+a_1)t} [h(x_1)v^W + a_1v^L + R - x_1] dt - F \\ &= \frac{h(x_1)v^W + a_1v^L + R - x_1}{r + h(x_1) + a_1} - F. \end{aligned}$$

The payoffs to the challengers are analogous except that they accrue no flow profits. For $i = 2, 3, \dots, n$,

$$\begin{aligned} V^i(x_i, a_i) &= \int_0^\infty e^{-rt} e^{-(h(x_i)+a_i)t} [h(x_i)v^W + a_iv^L - x_i] dt - F \\ &= \frac{h(x_i)v^W + a_iv^L - x_i}{r + h(x_i) + a_i} - F. \end{aligned}$$

Our induction hypothesis is that $v^W > v^L$. We shall show that if this hypothesis is true for some stage, then it is also true for the previous stage. In the last stage, only one innovation remains, so $v^W = R_0/r$, where R_0 is the flow rate of profit on the last innovation, and $v^L = 0$. Thus, the hypothesis is true for the last stage.

Assumption 3

There exists x^0 such that for each challenger i , $V^i(x^0, a) + F \geq v^L$ for all a . This reduces to: there exists x^0 such that $h(x^0)(v^W - v^L) - x^0 \geq rv^L$.

Assumption 3 says that there always exists an investment level for a challenger for which gross profits exceed the value of losing immediately. Note that Assumption 3 holds trivially at the last stage, in which $v^L = 0$. Since v^W and v^L are parameters for the current stage, the assumed existence of such an x^0 is a restriction on the function $h(\cdot)$.

Proposition 19 [Reinganum (1985)]

The Nash equilibrium in the current stage is symmetric among the challengers; that is, $x_i^* \equiv x_C$ for all $i \neq 1$. The incumbent invests less than each challenger in the current stage. That is, $x_1^* \equiv x_1 < x_C$.

This proposition highlights the dynamic evolution of the market. The current incumbent, since it invests at a lower rate, is least likely to win the current race.

Thus, the industry is characterized by a turnover of the technological leadership rather than a single continuing leader. It is in this sense that the equilibrium process resembles Schumpeter's "process of creative destruction".

Proposition 20 [Reinganum (1985)]

Let V^I and V^C denote the equilibrium expected profit for the incumbent and each challenger, respectively. Then each firm would prefer to be the incumbent in the current stage than a challenger. That is $V^I > V^C$.

Under the hypothesis that $v^W > v^L$, we have deduced that $V^I > V^C$. But these are simply the continuation values for the previous stage. This completes the induction argument.

In this model the length of the current stage – and hence the reward to the incumbent over the current stage – is affected by each firm's investment. Since the challenger firms do not forfeit any current stream of profit by inventing, they have a greater marginal incentive to invest in research and development. We accorded the incumbent no advantage which was due to incumbency per se. If the incumbent were to enjoy (for example) a marginal cost advantage in the conduct of research, the conjecture is that the incumbent might then invest more. Thus, a sufficiently large incumbent advantage may reverse the main result of this model. However, by focusing on the no-advantage case, we are able to isolate this inertial tendency of the incumbent to invest less than each challenger.

Vickers (1984) has addressed similar questions with a sequence of process innovations in the context of the auction model. In particular, he wants to discover how the product market structure evolves over time; does one firm become increasingly dominant by winning most or all of the races, or is there a process of "action–reaction", in which market leadership is constantly changing hands? Using a two-firm model, he finds that when the product market is very competitive (e.g. Bertrand) then there is increasing dominance; but when it is not very competitive (e.g. Cournot) then there is action–reaction.

Vickers assumes a sequence of not-so-drastic innovations, so that the profit flows of the two firms typically depend upon the levels of technology represented by each firm's most recent patent. There are T periods in the game and we label them backwards; at t , there are t periods (and hence t innovations) to go. Each innovation is associated with a cost level c_t , with $c_1 < c_2 < \dots < c_T$. At the beginning of period t there is a race for the innovation with cost level c_t , which takes the form of a simple auction in which the winner pays its bid (or alternatively, the maximum bid the loser would have been willing to make) and the loser does not forego its bid. Let $\pi(s, t)$ denote the flow profit (gross of research and development expenses) of a firm with cost level c_s when its rival has cost level c_t . This function is assumed to be non-negative for s and t , decreasing

in s and increasing in t . Let $\Pi(s, t) = \pi(s, t) + \pi(t, s)$ be joint profits. For simplicity, firms are assumed not to discount the future.

Proposition 21 [Vickers (1984)]

If $\Pi(t, t + 1) > \Pi(t, t + k)$ for all t, k , then the evolution of the market has an “action–reaction” character; that is, firms alternate being the technological leader.

The reverse of the hypothesis of Proposition 21 is not sufficient to cause increasing dominance (that is, for the same firm to win all races). However, the following proposition gives a sufficient condition for increasing dominance.

Proposition 22 [Vickers (1984)]

If $\pi(s + k, s) = 0$ for all s and $k \geq 1$, then the evolution has an increasing dominance character; that is, the same firm wins every race.

Note that this result is for drastic innovations; thus this result and that of Reinganum (1985) parallel the results of Gilbert and Newbery (1982) and Reinganum (1983b) for a sequence of innovations. The use of the auction model again gives opposite results from the stochastic racing model. To understand why we obtain these disparate results, it is useful to recall the incentives for investment described by Katz and Shapiro (1987); the stand-alone incentive represents the difference between the firm’s profits after versus before it innovates, while the incentive to pre-empt represents the difference between the firm’s profits if it innovates instead of its rival. In the deterministic model, so long as the stand-alone incentive is non-negative, the incentive to pre-empt dominates the firm’s decision (and an incumbent monopolist has a greater incentive to pre-empt than does a challenger). But when the date of rival success is drawn from a continuous distribution as in the stochastic racing model, concern about pre-emption is much less acute. Moreover, for drastic innovations, the pre-emption incentive is the same for both firms (both get monopoly profits if successful and nothing if unsuccessful), while the stand-alone incentive is greater for the challenger. Even for less drastic innovations, in which the pre-emption incentive is greater for the incumbent, the fact that pre-emption is only probabilistic means that both incentives come into play, with the result that for some less than drastic innovations, it is the greater stand-alone incentive for the challenger which carries the day.

The papers discussed so far in this section involved asymmetrically placed firms. However, the differences among firms did not confer an ex ante advantage

upon any particular firm. That is, for a given level of investment in research and development, all firms were equally likely to become the winner. Differences in incentives generated ex post advantages, since in equilibrium firms chose to invest different amounts on innovative activity. The papers to be discussed in the remainder of this section describe models in which the nature of the asymmetry confers a *strategic* advantage upon one firm. These are essentially multi-stage models which culminate in a single innovation; however, a firm's position at an intermediate stage affects the effectiveness of its investment in research and development. Thus, at any intermediate stage (in which firms' positions differ) firms are not equally likely to become the winner even if they (from now on) invest the same amount.

Fudenberg et al. (1983) and Harris and Vickers (1985) have devised very similar models of such a multi-stage race. In Fudenberg et al. (1983), firms are envisioned as suffering from information and/or response lags regarding the research activities of their rivals. Lack of information or the inability to respond quickly allows firms that are only slightly behind to catch up before the leading firm can act to prevent it. The existence of these lags effectively makes time discrete for this model. In period t , firms are informed about their rivals' research activities up through period $t - 1$. Invention is assumed to occur as soon as one firm has accumulated enough knowledge, as measured by total research and development spending. Firms may elect to learn at a high or a low rate in each period, and the costs of learning are strictly convex. Within the current period, each firm must choose its rate of knowledge acquisition without knowing its rival's choice.

They find that firms will choose the high rate only if they are sufficiently close together in terms of accumulated experience. If a firm lags by a sufficiently large amount, then it drops out of the race, allowing the remaining firm to proceed at the low rate. As the information lag becomes arbitrarily short, the lagging firm drops out immediately; only if the firms remain tied is there any competition. Thus, if firms begin with equal experience there is a short intense battle followed by the emergence of a single firm. If firms begin with unequal experience, the firm which is at an initial disadvantage simply never enters the race. They go on to show that as the length of the period of information lag decreases, the lag in experience for which the follower still competes also decreases. In the limit as the period length approaches zero, an arbitrarily small headstart in terms of knowledge is sufficient to cause the lagging firm to drop out immediately.

A somewhat more general version of this model appears in Harris and Vickers (1985). In Fudenberg et al. (1983), both players valued the patent equally, and both faced the same cost conditions. Thus, distance from completion could be measured as the difference between accumulated knowledge and required knowledge. Harris and Vickers allow firms to place different values on the reward and to face different cost functions. They too find that if one player is far enough

ahead, then the other gives up. However, being "far enough ahead" in this case is not measured in terms of literal distance; it depends upon the value placed on winning and the costs of achieving a win. The Harris and Vickers (1985) model is also cast in discrete time, but players are assumed to move in alternate periods. In the limit as the length of the period approaches zero, the firm which has the opportunity to move first pre-empts the other completely. The equilibrium is somewhat easier to characterize due to the alternating moves assumption (there will be no mixed strategies).

Four significant factors combine to determine which player has a strategic advantage. First, firms may differ in their valuations of the patent. Second, they may discount the future to different degrees. Third, firms may differ in the efficiency (i.e. the cost) of performing research and development. Finally, firms may differ in the amount of knowledge and experience they have already acquired.

Formally, two players, denoted A and B, are competing for a single prize. They value the prize at P_A and P_B , respectively, with $P_i > 0$ for $i = A, B$. At the beginning of the game, A and B are distances x_0 and y_0 from the finish line (i.e. A requires x_0 more units of knowledge, B requires y_0 more units of knowledge). Firm A is assumed to move first, then firm B, and so on. The first firm to reach 0 wins the prize. Progress toward the goal depends upon the amount a firm invests in each period. In particular, if firm i invests z he moves a distance of $w_i(z)$ toward the goal, where $w_i(0) = 0$ and $w_i(\cdot)$ is continuous and strictly increasing. Thus, after firm A has made his k th investment a_k , the positions are $x_{2k-1} = x_{2k-2} - w_A(a_k)$ and $y_{2k-1} = y_{2k-2}$. After firm B has made his k th investment b_k , the positions are $x_{2k} = x_{2k-1}$ and $y_{2k} = y_{2k-1} - w_B(b_k)$. Let N denote the smallest integer such that either $x_N \leq 0$ or $y_N \leq 0$; thus N is endogenously determined. The prize is awarded to the firm which first reaches 0; since firms move alternately, they will not reach 0 simultaneously. If no firm ever reaches 0, then no firm wins the prize. Let ρ_B and ρ_A denote the firms' discount rates. If A wins the prize with his k th investment, he receives $\rho_A^{k-1}P_A - \sum_{i=1}^{\infty} \rho_A^{i-1}a_i$, where a_i is understood to drop to zero after one firm wins. If A does not win the prize, its payoff is $-\sum_{i=1}^{\infty} \rho_A^{i-1}a_i$. Firm B's payoff is analogously defined.

A *strategy* for firm i is an infinite sequence of investment levels which may be chosen contingent upon the sequence of previous bids. The notion of equilibrium to be employed is subgame perfect Nash equilibrium [Selten (1975)]. A strategy pair is a *subgame perfect equilibrium* if its restriction to any subgame is a Nash equilibrium. The following convention will be maintained: if a player is indifferent between winning the prize with an overall payoff of zero and not winning the prize, then he will choose to win the prize.

Harris and Vickers define a sequence of critical distances from the finish line for A and B, denoted by $\{C_n\}_{n=0}^{\infty}$ and $\{D_n\}_{n=0}^{\infty}$. Heuristically, C_1 is the maximum distance that A can cover with one bid and obtain a non-negative

payoff overall. C_2 is the maximum distance that A can cover subject to covering at least $C_2 - C_1$ with his first bid, and obtain a non-negative payoff overall. C_n is the maximum distance that A can cover with a sequence of non-negative investments, subject to moving within C_{n-1} with the first investment, and without spending more than the present value of the prize.

The sequence $\{C_n\}$ has the following properties: (1) the sequence $\{C_n\}$ is nondecreasing; (2) if $\{C_n\}$ ever fails to be strictly increasing it remains constant thereafter; and (3) it is possible for A to cover distance h and obtain a non-negative payoff overall if and only if $h \leq C_n$ for some n . For the formal definition of these sequences and the proof that they have these properties, the reader is referred to Harris and Vickers.

Proposition 23 [Harris and Vickers (1985)]

Suppose that A and B are respectively at distances x and y from the finish line. Then in perfect equilibrium, there are four mutually exclusive and exhaustive possibilities.

(i) For some $n \geq 1$, $x \leq C_n$ and $y > D_n$. Then firm A wins; his investments are those he would make in the absence of rivalry from firm B; firm B always invests 0. The point (x, y) belongs to A's "safety zone".

(ii) For some $n \geq 1$, $x > C_n$ and $y \leq D_n$. Then firm B wins; his investments are those he would make absent any rivalry from firm A; firm A always invests 0. The point (x, y) belongs to B's "safety zone".

(iii) For some $n \geq 0$, $C_n < x \leq C_{n+1}$ and $D_n < y \leq D_{n+1}$. Then if it is firm A's turn to move, firm A wins; his investments are those he would make if (absent rivalry) he were required to move to within C_n of the finish line with his first investment; B always invests 0. Conversely, if it is firm B's turn to move, then firm B wins; his investments are those which he would make if (absent rivalry) he were required to move to within D_n of the finish line with his first investment; A always invests 0. The point (x, y) belongs to a "trigger zone".

(iv) For all $n \geq 0$, $x > C_n$ and $y > D_n$; then neither firm wins and both always invest 0.

It is apparent that the equilibrium outcome depends upon the initial point (x_0, y_0) . To show that it also depends upon the other parameters of the model, Harris and Vickers show that C_n is strictly increasing in P_A for $n \geq 1$ and that C_n is increasing in ρ_A for $n \geq 2$.

Again it is possible to determine what happens to this equilibrium as the reaction times shrink. In the limit, the trigger zone collapses to a curve; the safety zones for A and B lie on opposite sides of this curve. The fact that the curve depends upon more than just the distance to the finish line (e.g. the valuations,

discount factors and cost functions) implies that this curve need not be the 45° line. Harris and Vickers give a specific example in which it is linear, but does not have unitary slope.

In either the discrete game or the limiting case, the equilibrium has similar features. If the game begins in one firm's safety zone, then the winner is already determined and that firm proceeds as though no rival existed; the rival invests nothing. If the game begins in a trigger zone, then the firm which is accorded the first move jumps immediately to its safety zone, after which it proceeds as though no rival existed and again the rival invests nothing.

Park (1984) and Grossman and Shapiro (1987) have analyzed a two-stage version of Lee and Wilde (1980) in order to investigate the impact of position (leading or lagging) upon equilibrium investment. They assume two stages with identical (stochastic) technologies for producing success. Completion of the intermediate stage does not result in a prize, but brings one closer to it; the first firm to complete both stages wins a prize worth P . The stationarity of the problem implies that one need only characterize four investment levels: the symmetric equilibrium investment level when both firms have completed 0 stages, denoted x_{00} ; the investment levels for the case where (say) firm 1 has completed the first stage and firm 2 has not, denoted x_{10} and x_{01} , respectively; and the symmetric equilibrium investment level when both firms have completed the first stage, denoted x_{11} . Let V_{00} , V_{10} , V_{01} and V_{11} denote the corresponding Nash equilibrium profits.

By dynamic programming backward, one can characterize the subgame perfect Nash equilibrium rates of investment. Consider first the case where both firms have completed the first stage, but neither has completed the second stage. This is identical to the original Lee and Wilde (1980) case; the payoff to each firm can be written:

$$V^{11}(x, a) = \frac{h(x)P - x}{r + h(x) + a},$$

where a represents the rival firm's hazard rate.

Consider next the case in which firm 1 has succeeded with the first stage, but firm 2 has not. Since the same hazard function $h(\cdot)$ applies, we can write profits to firm 1 and 2, respectively, as

$$V^{10}(x, a) = \frac{h(x)P + aV_{11} - x}{r + h(x) + a}$$

and

$$V^{01}(x, a) = \frac{h(x)V_{11} - x}{r + h(x) + a}.$$

Finally, consider the case in which no firm has yet succeeded with the first stage. Each firm's expected payoff can be written:

$$V^{00}(x, a) = \frac{h(x)V_{10} + aV_{01} - x}{r + h(x) + a}$$

Proposition 24 [Grossman and Shapiro (1987)]

The rate of investment when both have succeeded in the first stage exceeds that of the leading firm which exceeds that of the lagging firm when only one firm has succeeded in the first stage; that is, $x_{11} > x_{10} > x_{01}$. Moreover, the rate of investment when both firms have succeeded with the first stage exceeds that when neither has succeeded; that is, $x_{11} > x_{00}$. The relationships between x_{10} and x_{00} and that between x_{01} and x_{00} are ambiguous.

The reason for the residual ambiguity is that success by one firm in the first stage has two effects. The lagging firm may reduce its rate of expenditure; this diminished rivalry induces the leading firm to reduce its expenditures as well. On the other hand, the fact that it is now closer to the prize causes the leading firm to increase its rate of expenditure. Success by its rival in the first stage would tend to cause the (now) lagging firm to reduce its investment rate, but an increase in the (now) leading firm's expenditure (due to its being closer to the prize) would tend to spur investment by the lagging firm. Simulations reported by Grossman and Shapiro indicated that the likely response to success by one firm in the first stage is for the leading firm to increase, and the lagging firm to decrease, its rate of expenditure.

Judd (1984) has formulated a more general version of the stochastic racing game of Lee and Wilde (1980) and Reinganum (1981a, 1982a) to include elements of feedback (recall that while Reinganum's method permitted feedback, the particular specification of the research and development process rendered the value of the state variable unimportant). The basic framework is that of Reinganum (1981a) with the exception that knowledge accumulates and depreciates according to the equation $\dot{z}_i = \gamma u_i - \delta z_i$ and the hazard rate is $\alpha u_i + \beta z_i$. If $\beta > 0$, firms' investment in R&D today increases their own current and future probabilities of success, and builds a stock of experience which may cause the rival firm to decrease its rate of investment tomorrow. Judd solves this problem for small values of the prize P by using perturbation methods and finds that indeed each firm's rate of investment does depend negatively upon its rival's accumulated experience.

The model in Judd (1985a) has the characteristic that firms have no exogenous strategic advantages (such as first moves or initial experience), but strategic

advantages are acquired endogenously over time through acquired knowledge and intermediate successes. This model incorporates uncertainty of two types, one of which can be characterized as "more risky" than the other. Thus, it allows one to examine whether rivalry in research and development causes firms to invest in projects which are insufficiently or excessively risky.

Assuming that the prize and social benefits are small or that the rate of time preference is large (enough to make approximations valid), Judd finds that if the prize equals the social benefit, then firms invest relatively too much in the riskier discovery process. Despite this, it is optimal to allow competition to proceed until one firm has completely finished rather than to award the prize to the leader at some earlier juncture; moreover, the prize ought to be nearly equal to the social benefit. He also characterizes the dependence of investment on the current positions. It turns out that if one player advances, the other reduces its effort on the riskier project, but may increase its effort on the less risky project. The description and manipulation of the formal model is somewhat tedious, and the reader is referred to Judd (1985a) for proofs.

Suppose two firms compete for a particular innovation. The position of firm 1 is denoted by a nonpositive scalar x (firm 2's position is denoted y), the absolute value of which could be regarded as the extent of additional knowledge required for success. There are two parallel projects in which the firm can invest in attempting to complete the innovation. The first is characterized by *gradual jumps* which have a probability of $F(a)$ of hitting zero (if a is the firm's current position) and otherwise have a probability $f(s, a)ds$ of landing in the interval $(s, s + ds)$. There is also a more risky process which never lands at an intermediate value, but hits 0 with probability $G(a)$ if a is the current position. This process is characterized by *leaps*, and is a more risky process than the one that involves gradual jumps. The firms choose intensities at which to operate these processes; these intensities affect the likelihood, but not the magnitude of the resulting jumps. The symbols x and y denote the state variable for firms 1 and 2, respectively; $u dt$ denotes the probability that the gradual jump process results in a jump of x if firm 1 chooses u ; $v dt$ is the probability of a jump of y if firm 2 chooses v . Let $f(s, a)ds$ denote the probability of a jump from a to $(s, s + ds)$ if a gradual jump occurs; $f(s, a) = 0$ if $s < a$ (firms only improve their positions). If $a' > a$, then $f(s, a')$ first-order stochastically dominates $f(s, a)$. $F(a)$ denotes the probability that the gradual jump process hits 0 from a given that a gradual jump occurs; $F(a)$ is increasing in a and is positive everywhere:

$$F(a) \equiv 1 - \lim_{\xi \rightarrow 0} \int_a^\xi f(s, a) ds.$$

Let $wG(x)dt$ symbolize the probability that firm 1 leaps to 0, where firm 1 chooses w ; $zG(y)dt$ is the probability that firm 2 leaps to 0, where firm 2

chooses z . $G(\cdot)$ is positive everywhere. Firm 1's costs are $\alpha u^2/2 + \beta w^2/2$, where α and β are positive scalars. Similarly, firm 2's costs are $\alpha v^2/2 + \beta z^2/2$. $P \geq 0$ represents the prize to the winner, and $\rho > 0$ is the common discount rate. Throughout Judd (1985a) uses infinitesimal notation; we will follow his convention here.

Consider first the research intensities of a joint venture between the two firms. In this case, the joint value function $W(x, y)$ satisfies the following dynamic programming equation:

$$\begin{aligned} W(x, y) = \max_{u, v, w, z} & \left\{ -(\alpha u^2/2 + \alpha v^2/2 + \beta w^2/2 + \beta z^2/2) dt \right. \\ & + u dt(1 - \rho dt) \left[\int_x^0 W(s, y) f(s, x) ds + PF(x) \right] \\ & + v dt(1 - \rho dt) \left[\int_y^0 W(x, s) f(s, y) ds + PF(y) \right] \\ & + P(1 - \rho dt)[wG(x) + zG(y)]dt \\ & \left. + (1 - \rho dt)[1 - (u + v + wG(x) + zG(y))dt] W(x, y) \right\}. \end{aligned} \quad (7)$$

To interpret this, the value of being at state (x, y) is the value of choosing (u, v, w, z) optimally for the next dt , and then continuing optimally. The choice of (u, v, w, z) incurs the costs on the first line; with probability $u dt$ firm 1 experiences a gradual jump, which has an associated expected present value (this term appears on the second line above); similarly, with probability $v dt$, firm 2 experiences a gradual jump, which has an associated expected present value (third line); there is also a probability $wG(x) + zG(y)$ that one of the firms will leap to success and an associated present value (fourth line); finally, there is a probability that neither firm experiences any advance at all, which event has expected present value $(1 - \rho dt)W(x, y)$. The probability that both firms experience gradual jumps and/or leaps is of order $(dt)^2$ and can safely be ignored.

Proposition 25 [Judd (1985a)]

There exists a unique solution $W(x, y)$ to the joint research problem, and $W(x, y)$ is C^∞ in P and ρ^{-1} .

Consider now a noncooperative version of this game. Assume that the position vector (x, y) is common knowledge. Then firms choose their research intensities contingent upon their current positions. Thus, the equilibrium concept used here is that of feedback equilibrium. Only symmetric equilibrium is considered. Let $V(x, y)$ represent firm 1's value function and $V(y, x)$ represent firm 2's value function. The equation of dynamic programming for firm 1 is

$$\begin{aligned}
 V(x, y) = \max_{u, w} \bigg\{ & - [\alpha u^2/2 + \beta w^2/2] dt \\
 & + u dt(1 - \rho dt) \left[\int_x^\infty V(s, y) f(s, x) ds + PF(x) \right] \\
 & + v dt(1 - \rho dt) \int_y^0 V(x, s) f(s, y) ds \\
 & + P(1 - \rho dt) wG(x) dt \\
 & + (1 - \rho dt) [1 - (u + v + wG(x) + zG(y)) dt] V(x, y) \bigg\}.
 \end{aligned} \tag{8}$$

Proposition 26 [Judd (1985a)]

There exists a $\bar{P} > 0$ such that for $P \in [0, \bar{P}]$, there is a symmetric feedback equilibrium value function $V(x, y)$ which is C^∞ in P and ρ^{-1} .

Proposition 27 [Judd (1985a)]

Noncooperative equilibrium play results in overinvestment relative to the joint optimum. Moreover, this excess is greater the closer is either firm to success. If P is small, joint profits would be increased if resources were shifted from the risky "leap" process to the less risky "jump" process. Thus, noncooperative firms undertake more risk than is jointly optimal.

If one can legitimately interpret P as the social value of the innovation, then the same proposition describes the relationship between the noncooperative equilibrium and the social optimum.

4. Licensing

In the work discussed above, the value of a patent was taken as given. But how is this value determined? Arrow (1962a) described the value of a patent on a

cost-reducing innovation as the revenue which an innovator could acquire by licensing the innovation to producing firms. He compared the value of licensing the innovation to a single producer versus members of a competitive industry, and found that the competitive environment yielded more revenue (even absent problems of bilateral monopoly, which might be expected to further reduce the value of licensing to a single producer). Thus, a competitive product market offered greater incentives for suppliers of innovations.

Kamien and Tauman (1984, 1986) performed a similar analysis when the downstream product market is oligopolistic, and members make their decisions to license in a strategic manner. By the term “firm” we refer only to producing firms; the patent holder is understood to be an independent researcher, not a current member of the industry. Assuming that firms are initially identical, with constant marginal costs c and a linear industry demand curve $p = 1 - bq$, Kamien and Tauman determined the maximum value to the patent holder from licensing a cost-reducing innovation to the industry. Given a license contract, which consists of a fixed fee and a linear royalty rate, the firms play a simultaneous-move game, where their strategies are either to license the innovation, or to forego licensing. The patent holder offers a licensing contract to maximize his profits, taking into account how the contract affects the subsequent Nash equilibrium among the firms.

The general game, which involves both a fixed fee and a linear royalty rate, is denoted game G . They examine two restrictions of the general game: G_1 , in which the royalty rate is constrained to be zero, while the fixed fee is subject to choice; and G_2 , in which the fixed fee is constrained to be zero, and the royalty rate is subject to choice. Their results are summarized in the three propositions below [Kamien, Tauman and Zang (1985) extend this analysis of licensing to the context of product innovation].

Proposition 28 [Kamien and Tauman (1984)]

(a) For any finite number n of firms, G_1 yields a higher payoff to the patent holder than G_2 , and consumers benefit more under G_1 than under G_2 . Firms make no more profit under G_1 , and no less profit under G_2 , than they did prior to the innovation. (b) The equilibrium of G_1 results in a monopoly if and only if the innovation is drastic. (A drastic innovation is here defined as one in which the monopoly price with the new technology does not exceed the competitive price under the old technology.) (c) If the innovation is not drastic, then in the limit (as the number of firms increases without bound), the patent holder makes the same profit in both G_1 and G_2 ; this profit is equal to the magnitude of the cost reduction times the original competitive output.

Since the game G permits the use of both a fixed fee and a royalty rate, the innovator must do at least weakly better under G than under either G_1 or G_2 .

Proposition 29 [Kamien and Tauman (1984)]

(a) The equilibrium of G results in a monopoly if and only if the innovation is drastic. In this case, the profit is the difference between monopoly profit under the new technology and the licensee's oligopoly profit under the old technology. (b) If the innovation is not drastic, then the number of licensees is never below $(n + 2)/2$. (c) In the limit (as the number of firms increases without bound), the profit of the patent holder in G coincides with his profits in G_1 and G_2 .

Finally, it is possible to compare the output levels, market prices and firms' profits before and after the innovation.

Proposition 30 [Kamien and Tauman (1984)]

In the (subgame) perfect Nash equilibrium of G : (a) total output increases and the market price falls as a result of the innovation; and (b) each firm is worse off relative to its profit prior to the innovation unless the patent is drastic and then only the monopoly breaks even.

Two key features of the Kamien and Tauman analysis are modified in Katz and Shapiro (1986). First, in Kamien and Tauman's model, the patent holder effectively *posts a contract*, which firms can either accept or reject. That is, it offers a pair consisting of a fixed fee and a royalty rate, and any firm which is willing to accept those terms may acquire a license. Of course, the optimal contract takes into account the subsequent equilibrium behavior of the potential buyers; that is, the patent holder computes its (equilibrium) demand function for licenses, and chooses its preferred point on that schedule. Second, in Kamien and Tauman (1984, 1986), the patent holder is understood to be an independent researcher.

Restricting attention to fixed fee contracts, Katz and Shapiro (1986) argue that when firms' demands for licenses are interdependent a superior selling strategy for an independent researcher involves offering a restricted number of licenses for auction with a minimum required bid. They also consider the optimal distribution strategy (within this class of auction-type strategies) for research joint ventures of arbitrary size. They find that dissemination of the technology is greater the smaller is the joint venture. Subsequently they examine the seller's incentives to develop the innovation given the feasibility of licensing. Again all downstream firms who are not members of the joint venture are worse off as a result of innovation.

Consider the case of a research lab which has developed an innovation which is potentially useful to the n member firms of a particular industry. Assume that each firm has need for a single license, and all n firms are identical.

Given this symmetry, firms' identities are irrelevant; the information which is relevant to the payoffs of the patent holder and the firms is the number of firms which will obtain a license. Let k denote the number of firms which obtain a license. Let $W(k)$ represent the profits of a firm that obtains a license when a total of k firms have done so, and let $L(k)$ denote the profits of a firm that does not obtain a license when a total of k firms have obtained licenses. These profits are gross of any licensing fees, which are assumed to be lump sums independent of subsequent output levels.

Assumption 4

(a) $L(k) \leq L(k-1)$ and (b) $L(k) < W(k)$, for $k = 1, 2, \dots, n-1$.

That is, a firm that has not obtained a license is worse off the greater is the number of firms which have obtained licenses and, given that k firms obtain licenses, profits are greater for those who have than for those who have not obtained a license.

The set of selling strategies open to the patent holder is the set of multiple-object sealed-bid first-price auctions with a minimum bid. That is, the patent holder makes available k licenses, but requires a minimum bid of \underline{b} . Each firm may submit a single bid b_i (to prevent anti-competitive hoarding, which might be individually profitable); the licenses go to the firms with the k highest bids (provided the bid is at least \underline{b}) at the bid values, and any ties are broken at random. Thus, a sales policy can be summarized by a pair (k, \underline{b}) . Katz and Shapiro refer to a policy of the form $(k, 0)$ as a *quantity* strategy and one of the form (n, \underline{b}) as a *price* strategy.

For a given policy (k, \underline{b}) , we need to characterize the Nash equilibrium of the bidding game for the n firms. Firm i 's willingness to pay clearly depends upon what it expects other firms to do. However, since there is complete and perfect information in this game, it is clear that in any bidding equilibrium all licensees pay the same price; if two licensees paid different prices, the one paying more could have lowered its bid and still received a license.

Consider first the case in which the patent holder is an independent research lab. If $k < n$ licenses are sold under the quantity strategy $(k, 0)$, then each firm knows that k licenses will be distributed, independent of his own actions. Then bidding for the licenses will drive the winning bid to $W(k) - L(k)$. The use of a price mechanism (n, \underline{b}) implies that each firm knows that one fewer licenses will be distributed if that firm refrains from buying one. In this case, the highest price obtainable for k licenses is $W(k) - L(k-1) \leq W(k) - L(k)$. Thus, a pure price strategy is strictly inferior to a pure quantity strategy when fewer than n licenses are sold and $L(\cdot)$ is strictly decreasing.

When n licenses are offered, a positive minimum bid is necessary since each firm will bid at most \underline{b} . To further characterize the outcome in this case, define the value of obtaining a license, given that $k - 1$ other firms also obtain licenses, to be $V(k) = W(k) - L(k - 1)$. Suppose that $V(k)$ decreases with k ; in this case, each firm finds a license less valuable the greater the number of other firms which are licensed. This would be typical of a cost-reducing innovation in a simple Cournot model with linear demand and constant marginal costs. If the patent holder licenses all firms, each firm compares $W(n)$ to $L(n - 1)$. The highest minimum bid which will still sell n licenses is $\underline{b} = V(n)$.

Proposition 31 [Katz and Shapiro (1986)]

If $V(\cdot)$ is strictly decreasing, then the optimal selling strategy (within the specified class) has one of two forms: (a) $(k, 0)$, where $k < n$ and the winning bid is $W(k) - L(k)$; or (b) (n, \underline{b}) , where $\underline{b} = V(n)$.

An alternative form of market organization would involve a number of firms maintaining a research lab as a joint venture. Call members of the joint venture *insiders* and nonmembers *outsiders*. Insiders now face a tradeoff between profits they receive from licensing the innovation to competitors and the profits they receive from production. Thus, a research joint venture is likely to have reduced incentives to license the innovation. Assume that the research joint venture has available the same class of licensing policies (k, \underline{b}) but is also free to distribute licenses at no cost to some or all of its members. Suppose that m firms participate and share equally in the profits of the research joint venture. In this case, all members share the same objective function and are thus unanimous regarding the preferred licensing policy.

If k licenses are issued, and \tilde{k} go to insiders, then profits to the insiders are $\tilde{k}W(k) + (m - \tilde{k})L(k) + R$, where R is revenue raised by licensing to outsiders. For given values of k and \tilde{k} , the joint venture will try to maximize R . From Proposition 31 we know that if $k < n$ licenses are distributed, the venture can extract a maximum of $W(k) - L(k)$ per license sold to outsiders, so $R = (k - \tilde{k})[W(k) - L(k)]$. Thus, insider profit is $k[W(k) - L(k)] + mL(k)$. This is independent of \tilde{k} , since (in equilibrium) the marginal revenue from a license equals the venture's opportunity cost.

Let $R^0(k)$ denote the licensing revenues that an independent researcher earns when it sells k licenses. When $k < n$ licenses are issued, $R^0(k) = k[W(k) - L(k)]$. For the joint venture of size m , $R^m(k) = k[W(k) - L(k)] + mL(k)$, or $R^m(k) = R^0(k) + mL(k)$. If the joint venture were to issue licenses to all firms, and $V(\cdot)$ is strictly decreasing, then total insider profit is $R^m(n) = mW(n) + (n - m)[W(n) - L(n - 1)]$, since the minimum bid which induces all outsiders to buy is $[W(n) - L(n - 1)]$.

Proposition 32 [Katz and Shapiro (1986)]

Suppose that $V(\cdot)$ is strictly decreasing. The m -firm joint venture's optimal selling strategy (within the specified class) has one of two forms: (a) $(k, 0)$, where $k < n$ and the winning bid is $W(k) - L(k)$; or (b) (n, \underline{b}) , where the winning bid is $\underline{b} = W(n) - L(n - 1)$.

Let k^m denote the number of licenses issued by the m -firm venture, where $m = 0, 1, 2, \dots, n$; that is, k^m maximizes $R^m(k)$. To determine whether or not to issue an additional license, the m -firm venture examines:

$$\begin{aligned}\Delta R^m(k) &\equiv R^m(k) - R^m(k - 1) \\ &= [R^0(k) - R^0(k - 1)] + m[L(k) - L(k - 1)].\end{aligned}$$

For two ventures of sizes m and $m - 1$, respectively, $\Delta R^m(k) - \Delta R^{m-1}(k) = L(k) - L(k - 1) \leq 0$, so the m -firm venture has less incentive to sell the k th license than does the $(m - 1)$ -firm venture. If the m -firm venture issues n licenses, the comparison is between $\Delta R^m(n)$ and $\Delta R^{m-1}(n)$. If $V(\cdot)$ is strictly decreasing, then Proposition 32 may be applied to obtain $R^m(n) = R^0(n) + mL(n - 1)$, and $\Delta R^m(n) = \Delta R^0(n)$, which is independent of m . Thus, the incentives of the independent researcher and the m -firm joint venture coincide for the n th license.

Proposition 33 [Katz and Shapiro (1986)]

Suppose that $V(\cdot)$ is strictly decreasing. Then an $(m - 1)$ -firm venture issues at least as many licenses as does an m -firm venture, for $m = 1, 2, \dots, n$.

Thus, research joint ventures tend to restrict the distribution of licenses relative to an independent researcher, and the extent of the restriction increases with the size of the venture. Moreover, an outsider cannot be better off (and is strictly worse off whenever $m > 0$ and $L(\cdot)$ is strictly decreasing) as a result of the innovation. To see why, suppose $k^m < n$. In equilibrium, an outsider is indifferent about buying a license, and thus has profits of $L(k^m) \leq L(0)$. If $k^m = n$, then $\underline{b} \geq W(n) - L(m)$ since this would induce all outsiders to obtain a license. Thus, an outsider has net profits that are no greater than $W(n) - [W(n) - L(m)] \leq L(0)$.

Now consider a research lab (either independent or a joint venture) deciding whether or not to develop the innovation. If we identify a researcher's incentive to develop the innovation with its profits from licensing and production net of its previous profits from production, then for an m -firm venture this incentive is

$R^m(k^m) - R^m(0) = R^m(k^m) - mL(0)$ (note that this formula is equally valid for $m = 0$).

Proposition 34 [Katz and Shapiro (1986)]

Suppose that $V(\cdot)$ is strictly decreasing. Then an m -firm venture has greater incentives to develop the innovation than has an $(m - 1)$ -firm venture.

In another paper, Katz and Shapiro (1985b) examine Nash equilibrium licensing and development behavior in a two-firm industry. One goal of this paper is to determine how the pattern of licensing depends upon the magnitude of the innovation. Assuming fixed fee licensing, they find that major innovations will not be licensed, but minor innovations will be licensed if firms are approximately equally efficient prior to innovation. If at least one firm would exclude the other (by refusing to license the innovation), then licensing will not occur because, in equilibrium, an excluding firm will be the innovator.

The aforementioned papers deal with a case in which no further innovation is anticipated; if another technology with equal or lower costs is possible, then an additional incentive to license the current innovation arises. Gallini (1984) has shown that an incumbent firm may choose to license a potential entrant to use its technology in order to forestall innovation by the entrant. That is, it will offer to share its market in order to make further innovation less attractive to the potential entrant.

Consider a homogeneous good market which consists of a single incumbent firm and a single potential entrant. These firms (and only these firms) may compete in the research and development of new production technologies. For simplicity, suppose that the incumbent currently has constant unit cost of c_3 and that two other cost levels exist: c_2 and c_1 with $c_1 < c_2 < c_3$. However, there may be a large number of technologies associated with each of these cost levels; thus discovery and patenting of a c_2 technology does not preclude the rival from discovering another route to the same cost. Research and development is therefore represented as sampling with replacement from a known discrete distribution over the cost levels $\{c_1, c_2, c_3\}$, with p_i denoting the probability that a technology with unit cost c_i is observed on any one draw. The results of each draw are revealed to both firms, but a patent prevents firms from immediately imitating the rival's technology. Each technology is assumed to be drastic in relation to the one with next highest cost; thus if one firm has c_3 and the other c_2 , the low-cost firm is the current incumbent and is free to price its output at the monopoly price. Finally, production takes place once research and development has ceased.

The analysis begins with one incumbent firm, which possesses a technology with unit costs of c_2 and one potential entrant, which currently has unit costs of

c_3 ; thus the entrant must discover a c_2 or c_1 technology to enter. The incumbent must decide whether to license its c_2 technology to the potential entrant, thus granting it a permanent share of the market. It is assumed that when firms have the same costs, they share cooperative profits equally (an equivalent analysis applies if they are noncooperative in the product market). Let Π_i denote industry profits if both firms have cost c_i , $i = 1, 2$. If a license is agreed upon, research is terminated; if not, the firms decide noncooperatively and simultaneously whether or not to engage in further research. Their actions are to continue (C) or to terminate (T) research.

Under the following assumption, the discovery of a c_1 technology makes further research unprofitable for the rival firm. Thus, one only needs to determine when the *current* (i.e. c_2) technology will be licensed to induce a rival to terminate research. The c_1 technology would never be licensed; because the innovation is drastic and firms have constant unit costs, there are no efficiency gains to having more than one firm producing at the same time with the same technology. Thus, all research terminates with the discovery of a c_1 technology.

Assumption 5

Assume that $p_1\Pi_1/2 - D < 0$, where D represents the cost per observation.

Suppose that if the entrant discovers a c_2 technology, then it would prefer to produce in this market and share the cooperative profits rather than to continue researching alone until a c_1 technology is discovered.

Assumption 6

Assume that $\Pi_2/2 > \Pi_1 - D/p_1$.

In this case, the entrant expects more from continuing research alone than does the incumbent (since the incumbent only benefits from discovering a c_1 technology, while the entrant benefits from discovering either a c_1 or a c_2 technology). Moreover, the entrant has less incentive to stop research when its rival continues than does the incumbent.

For licensing to be an equilibrium, each firm's profit must be as great as it could achieve in a game without licensing. This is because by not offering or by rejecting a license, either firm can bring about this outcome. Refer to the equilibrium in the game with no licensing as the *alternative equilibrium*. Let R_{ij}^I and R_{ij}^E denote the payoffs to the incumbent and entrant, respectively, when the pair of actions ij with $i, j \in \{C, T\}$ are taken, assuming that the firms continue on in a Nash equilibrium fashion. If a c_1 technology has been discovered, both firms terminate research by Assumption 5. Thus, for cost pairs of the form

(c_1, c_j) , $j = 1, 2, 3$ or (c_2, c_1) , the equilibrium action pair is TT . At (c_2, c_2) , TT is always a Nash equilibrium because Assumption 6 implies that if the incumbent terminates, the entrant would prefer to terminate rather than to continue researching alone; because the incumbent's incentives to continue researching alone are always weaker than the entrant's, if the entrant terminates at (c_2, c_2) , then the incumbent will do so as well. However, if a firm continues to search at this cost pair, then its rival can stop and receive 0 or compete for a c_1 technology and receive (in expected value) $R' = \Pi_1/2 - D/p_1(2 - p_1)$. If $R' > 0$, then CC will also be a Nash equilibrium at (c_2, c_2) . TT is Pareto superior to CC and is selected as the relevant equilibrium at this point, but the alternative selection would yield the same results.

Now it is possible to describe the payoffs from various strategy pairs at (c_2, c_3) . If both firms terminate, $R_{TT}^I = \Pi_2$ and $R_{TT}^E = 0$. If both continue,

$$R_{CC}^I = R_{CC}^E = \frac{p_1(2 - p_1)\Pi_1/2 + p_2(1 - p_1)\Pi_2/2 - D}{1 - (1 - p_1)p_3}.$$

If the incumbent continues but the entrant terminates,

$$R_{TC}^I = \frac{p_2\Pi_2}{2(1 - p_3)} \quad \text{and} \quad R_{TC}^E = \frac{p_1\Pi_1 + p_2\Pi_2/2 - D}{1 - p_3}.$$

Finally, if the incumbent terminates and the entrant continues, $R_{CT}^I = \Pi_1 - D/p_1$ and $R_{CT}^E = 0$. Note that CT cannot be an equilibrium because this requires $R_{CT}^I \geq \Pi_2$, which contradicts Assumption 6. The remaining three pairs of actions can be alternative equilibria for some parameter values.

A licensing equilibrium requires that there must exist a share of profits using the current technology such that both firms earn at least as much as they would in the alternative equilibrium. Moreover, each firm must receive as much under the licensing agreement as it would receive from continuing research alone (otherwise it will subsequently deviate from the agreement not to continue research). These two conditions will be met if the cooperative profits from the c_2 technology are at least as large as the sum of the firms' maximum profits from the alternative equilibrium or from continuing research alone. Let R_a^i denote the alternative equilibrium payoff to agent i , $i = I, E$. Then a licensing equilibrium requires that

$$\Pi_2 \geq \max\{R_a^I, R_{CT}^I\} + \max\{R_a^E, R_{TC}^E\}. \quad (9)$$

When the alternative equilibrium is TT , then no license will be offered. When the alternative equilibrium is CC , equation (9) reduces to $\Pi_2 \geq R_{CT}^I + R_{TC}^E$ or

$\Pi_2 \geq 2[\Pi_1 - D/p_1]$, which is always true under Assumption 6. When the alternative equilibrium is TC , there are two possibilities. If $R_{TC}^I \geq R_{CT}^I$, then equation (9) becomes $\Pi_2 \geq R_{TC}^I + R_{TC}^E = \Pi_1 - D/p_1$, which always holds. If $R_{TC}^I < R_{CT}^I$, then licensing requires $\Pi_2 > 2[\Pi_1 - D/p_1]$ as above. Assume that licensing always occurs when firms are indifferent.

Thus under Assumption 6, a licensing contract will always be struck to terminate research that would take place absent licensing. From the incumbent's perspective, licensing protects against the risk of discovery of a lower cost technology by the entrant. Moreover, resources which would have been devoted to research (by the entrant and possibly also by the incumbent) are saved.

One can relax Assumption 6 so that both firms have an incentive to continue research until a c_1 technology is obtained.

Assumption 7

Assume that $\Pi_2/2 \leq \Pi_1 - D/p_1$.

Under this complementary assumption, both firms face the same incentives to continue and terminate research. In this case, in order for a licensing equilibrium to exist, it must be that $\Pi_2 \geq 2[\Pi_1 - D/p_1]$; but this contradicts Assumption 7. Thus, in this case there will be no equilibrium with licensing.

Gallini and Winter (1985) extend the analysis of this strategic incentive for licensing to more general environments including nondrastic innovations. They find that licensing encourages additional research when the firms' current production costs are close and discourages further research when current production costs are relatively far apart. This is because there are two effects of licensing. First, having developed a superior technology, a firm can license it to its rival; this is the incentive which was pointed out by Salant (1984) in the context of Gilbert and Newbery's preemption model, and it is greatest when current costs are close together. Second, when costs are far apart the low cost firm has an incentive to offer a license to the high cost firm in order to make further research by the high-cost firm unattractive; this minimizes the erosion of the low-cost firm's market share while economizing on development expenditures.

5. Adoption and diffusion of innovations

In the previous section the extent of licensing was examined, but in a timeless framework; all licensing was assumed to be completed at once. However, an important empirical observation regarding the adoption of innovations is that adoption is typically delayed and that firms do not adopt an innovation simultaneously. Instead, innovations "diffuse" into use over time.

The general pattern for economic models of diffusion is concisely described by David (1969, ch. II, p. 10):

whenever or wherever some stimulus variate takes on a value exceeding a critical level, the subject of the stimulation responds by instantly determining to adopt the innovation in question. The reasons such decisions are not arrived at simultaneously by the entire population of potential adopters lies in the fact that at any given point of time either the "stimulus variate" or the "critical level" required to elicit an adoption is described by a distribution of values, and not a unique value appropriate to all members of the population. Hence, at any point in time following the advent of an innovation, the critical response level has been surpassed only in the cases of some among the whole population of potential adopters. Through some exogenous or endogenous process, however, the relative positions of stimulus variate and critical response level are altered as time passes, bringing a growing proportion of the population across the "threshold" into the group of actual users of the innovation.

The heterogeneity posited here may involve any firm characteristic which is relevant to the adoption decision. For instance, David (1969) offers both theoretical and empirical arguments in favor of the use of firm size. Other explanations, such as differential access to information regarding the innovation's profitability and/or managerial willingness to take risk, are also common. A combination of these two latter features generates the diffusion of innovation described in Jensen (1982), which provides a formal model of the type described by David.

When an innovation is first announced, a firm may be uncertain regarding its profitability should it adopt the innovation. However, this uncertainty may be reduced over time as information regarding the innovation accumulates. Jensen's formal model assumes that at any decision point in time, the firm has two options: it can adopt the innovation, which involves a fixed cost and is irreversible; or it can wait. If the firm waits, then it receives additional information regarding the innovation's profitability, but of course it foregoes for one period any profit it might have made by adopting the innovation. The firm begins with a prior estimate of the likelihood that the innovation will be profitable, and "learns" over time, updating its estimate in a Bayesian fashion. Thus, the firm's decision problem can be modelled as an optimal stopping problem. An optimal adoption rule has the following form: if the posterior estimate of the likelihood that the innovation would be profitable is sufficiently high, adopt; otherwise, wait. If an industry consists of firms who differ in their initial assessments of the innovation, then they will typically reach this critical level of estimated profitability at different times. Thus, the innovation will be observed to diffuse into use.

Suppose that a firm is currently at equilibrium in its industry; normalize its current profits to zero for simplicity. Suppose that an alternative production process (an innovation) is exogenously developed and may be acquired at a cost

of C . This process has some stochastic features in the sense that with probability θ the firm earns (in present value terms) $R_1 = r_1/(1 - \beta)$, where r_1 is the rate of flow profit and β is the discount factor; with probability $1 - \theta$ the firm earns $R_0 = r_0/(1 - \beta)$. It is assumed that $R_0 < R_1$. Thus, $1 - \theta$ might be interpreted as the fraction of “down time” associated with the process. The parameter θ is unknown to the firm, but it is known to be one of two possible values, θ_1 or θ_2 , with $1 > \theta_1 > \theta_2 > 0$. Assume that

$$\theta_1 R_1 + (1 - \theta_1) R_0 - C > 0 > \theta_2 R_1 + (1 - \theta_2) R_0 - C, \quad (10)$$

so that if $\theta = \theta_1$, the innovation can be classified as “profitable”, while if $\theta = \theta_2$, the innovation can be termed “unprofitable”. If the firm does not adopt the innovation in period i , it is assumed to receive a costless signal, representable as a Bernoulli random variable Z_i , which takes on the value 1 if the information is favorable and 0 if it is unfavorable. The probability that $Z_i = 1$ is the unknown parameter θ .

Given a sequence Z_1, \dots, Z_n , the firm can construct an estimate of the parameter θ as follows. If p is the firm’s prior probability that $\theta = \theta_1$, then its estimate of θ is

$$q(p) \equiv p\theta_1 + (1 - p)\theta_2. \quad (11)$$

Its posterior probability that $\theta = \theta_1$ is $h_1(p) \equiv p\theta_1/q(p)$ if the observation is favorable and $h_0(p) \equiv p(1 - \theta_1)/(1 - q(p))$ if the observation is unfavorable. Assuming that the firm’s initial prior probability that $\theta = \theta_1$ is g , then after n observations, k of which were favorable, the firm’s posterior probability that $\theta = \theta_1$ is

$$p(n, k, g) \equiv \left[1 + (\theta_2/\theta_1)^k ((1 - \theta_2)/(1 - \theta_1))^{n-k} (1 - g)/g \right]^{-1}. \quad (12)$$

Beginning from the initial prior g , the state variable for the decision process is $p(n, k, g)$, the firm’s current probabilistic belief that the innovation is profitable. Assuming that an infinite number of decision periods exists, then $V(p)$, the maximum expected return when the current state is p , is defined as the solution to the following functional equation of dynamic programming:

$$V(p) = \max\{V^a(p), V^w(p)\}, \quad (13)$$

where the expected value of adoption is

$$V^a(p) \equiv q(p)R_1 + (1 - q(p))R_0 - C, \quad (14)$$

and the expected value of waiting one period and continuing optimally is

$$V^w(p) \equiv \beta [q(p)V(h_1(p)) + (1 - q(p))V(h_0(p))]. \quad (15)$$

Proposition 35 [Jensen (1982)]

There exists a unique $p^* \in (0, 1)$ such that $V^a(p) \geq V^w(p)$ if and only if $p \geq p^*$.

Thus, the optimal adoption rule is to adopt the innovation at the first date n for which $p(n, k, g) \geq p^*$. Moreover, the probability of adoption at or before a given stage N is an increasing function of g , k , r_1 , r_0 and β and is a decreasing function of C .

It is also easy to see that immediate adoption may not be optimal, but a profitable innovation will eventually be adopted with probability 1 if $g \neq 0$. If $g < p^*$, then the firm will wait at least one period to gather additional information about the innovation's profitability. However, by the law of large numbers, the Bayesian estimate of θ will eventually converge to its true value. If this is θ_1 , then the firm will eventually adopt the innovation.

A firm is more likely to adopt by a given date the more favorable is its initial assessment of the innovation; thus a firm which begins by being sufficiently skeptical will delay adoption; if it is willing to learn, however, it will not forego a profitable innovation indefinitely. Clearly, the analogous result for unprofitable innovations is not true; some unprofitable innovations will be adopted due to optimistic initial beliefs or the receipt of favorable information. The length of the delay prior to adoption will be shorter (on average) the more optimistic the initial belief, the more favorable the information received, the higher the discount factor, and the higher the rate of flow profits; the length of delay will be greater (on average) the higher are the adoption costs.

Suppose now that there is an industry composed of a continuum of these firms; each receives the same information about the innovation, but they may begin with different prior beliefs about it. In this case, firms with different prior beliefs will adopt the innovation at different times. The traditional S-shaped diffusion curve can be obtained by means of appropriate assumptions regarding the distribution of prior beliefs within the industry.

McCardle (1985) has generalized this model to include explicit costs of information gathering. When information is costly, a firm may elect to reject the innovation (i.e. terminate sampling without adopting). In this case, the optimal decision function is "cone-shaped". That is, for sufficiently high posterior beliefs, the firm stops sampling and adopts the innovation; for sufficiently low posterior beliefs, it stops sampling and rejects the innovation; finally, for intermediate beliefs, it continues sampling. Jensen (1984a, 1984b) has also considered alterna-

tive specifications of sampling costs, as well as asymmetry in information processing capacity [Jensen (1984c)]. Mamer and McCardle (1985) extend the results of McCardle (1985) to a two-firm game in which the firms receive private signals at a fixed cost per signal. Roberts and Weitzman (1981) present a single-agent sequential decision model which is applicable to the innovation adoption problem and which uses a more general specification of uncertainty.

One difficulty with this formulation is that no firm anticipates any future technological improvements. Balcer and Lippman (1984) develop a one-firm model of the timing of adoption assuming that the firm has perfect information regarding the current best available technology, but is uncertain about the rate and magnitude of future improvements. They find that there is a critical technological lag beyond which the firm immediately adopts the best available technology; otherwise, it postpones adoption. The critical lag length increases with the anticipated rate of future innovation.

Another difficulty is that no firm perceives the impact of other firms' adoption decisions upon its own profits. Reinganum (1981b, 1981c) has argued that no *ex ante* heterogeneity among firms, nor any imperfect information regarding the innovation's profitability, is necessary to obtain diffusion of an innovation; under some circumstances, a certain amount of rationality and foresight on the part of firms is sufficient. In particular, suppose that an industry of n identical firms produces and markets a homogeneous good in a Cournot–Nash manner. When a cost-reducing, capital-embodied process innovation is announced, each firm must decide when to adopt it, accounting for the costs and benefits of the innovation itself, and for the effects of rival firms' adoption decisions. It is assumed that each firm must commit itself to an adoption date at once and without knowledge of its rival's decisions. The justification for this assumption is that adoption of a process innovation is a time-consuming activity, with installation and adjustment costs a function of the planned adjustment path. Thus, the choice of an "adoption date" really represents a time at which adoption will be completed (assuming it begins immediately); it may be very costly to alter the planned path of adjustment once it has been selected. That is, the whole *path* of adjustment, not just a delivery date, would have to be changed. We assume that such alterations of plans are prohibitively costly.

Let $\pi_0(m)$ be the rate of profit flow to firm i when m firms have adopted the innovation, but firm i has not. Next let $\pi_1(m)$ be the flow of profit to firm i when m firms have adopted and i is among them. We assume that $\pi_0(m)$ and $\pi_1(m)$ are known with certainty.

Assumption 8

Profit rates are non-negative, and the increase in profit rates due to adopting $(m - 1)$ th is greater than due to adopting m th. That is, $\pi_0(m - 1) \geq 0$

and $\pi_1(m) \geq 0$ with $\pi_1(m-1) - \pi_0(m-2) > \pi_1(m) - \pi_0(m-1) > 0$ for all $m \leq n$.

Let τ_i denote firm i 's adoption date and let $p(\tau_i)$ represent the combined purchase price plus adjustment costs (in present value terms) required to bring the new technology on line by date τ_i . The function $p(\cdot)$ is assumed to be twice differentiable and convex.

Assumption 9

(a) $\lim_{t \rightarrow 0} p(t) = -\lim_{t \rightarrow 0} p'(t) = \infty$; (b) $\lim_{t \rightarrow \infty} p'(t) > 0$; (c) $p''(t) > re^{-rt}[\pi_1(1) - \pi_0(0)]$ for all t .

In keeping with the adjustment costs story, Assumption 9(a) implies that instantaneous adjustment is prohibitively costly, but costs drop off sharply as the adjustment period is lengthened. Assumption 9(b) states that there is an "efficient scale" or cost-minimizing period of adjustment; finally, Assumption 9(c) states that adjustment costs increase at a sufficiently fast rate as the adjustment period is compressed. This assumption ensures that firm i 's objective function will be (locally) strictly concave in its choice variable.

Let $\tau = (\tau_1, \dots, \tau_n)$ denote the vector of adoption times in increasing order of adoption and let τ_{-i} denote this vector without the i th element τ_i . Thus, $\tau = (\tau_i, \tau_{-i})$. Let $V^i(\tau)$ be the i th adopter's profit (in present value terms) when the vector of adoption dates is τ . Then

$$V^i(\tau) = \sum_{m=0}^{i-1} \int_{\tau_m}^{\tau_{m+1}} \pi_0(m) e^{-rt} dt + \sum_{m=i}^n \int_{\tau_m}^{\tau_{m+1}} \pi_1(m) e^{-rt} dt - p(\tau_i),$$

where $\tau_0 \equiv 0$ and $\tau_{n+1} \equiv \infty$.

Proposition 36 [Reinganum (1981c)]

The n -tuple of adoption dates τ^* defined by system (16) is a Nash equilibrium and $\tau_{i-1}^* < \tau_i^* < \tau_{i+1}^*$, $i = 1, 2, \dots, n$:

$$\partial V^i / \partial \tau_i = [\pi_0(i-1) - \pi_1(i)] e^{-r\tau_i^*} - p'(\tau_i^*) = 0. \quad (16)$$

Thus, an equilibrium for this game is asymmetric, implying a "diffusion" of innovation over time, despite the facts that information is perfect and the firms are identical.

Reinganum's work focused on situations in which a firm was committed to its adoption date, regardless of any subsequent information it might receive regarding the adoption decisions of rival firms. This seems plausible under the adjustment costs interpretation given above. However, in many instances firms will be able to respond to the actions of rival firms without significant lags or associated costs of changing plans. Judd (1983) and Fudenberg and Tirole (1985) have examined this situation, and find that the pattern of adoption will still be characterized by diffusion, but that firms will be forced to adopt the innovation faster due to the threat of pre-emption by rival firms. In some cases, there may also exist a continuum of simultaneous adoption equilibria.

Fudenberg and Tirole (1985) show that, for the model described above, a firm's equilibrium payoff declines monotonically with its rank in the order of adoption. Judd (1983) argues that if firms are able to respond to the choices and/or actions of rival firms, such a situation cannot occur in equilibrium. Instead, firms would compete to be the first firm, knowing that rival firms would adjust their adoption plans in response. In order to characterize equilibrium when firms are able to respond quickly to rivals' behavior it is necessary to give an explicitly dynamic description of strategies. Following Judd (1983), let $d_i(t, k)$ be the decision rule for firm i ; it specifies whether or not firm i adopts the innovation at t if k other firms have already done so; thus $d_i(t, k) = 1$ if firm i decides to adopt the innovation, and $d_i(t, k) = 0$ if firm i decides not to adopt the innovation at t . A decision rule $\{d_i\}_{i=1}^n$ is a (subgame perfect) equilibrium if and only if at each t , d_i maximizes the profits of firm i given the decision rules of the other firms and the value of k at t .

For simplicity, Judd makes the following assumptions [these turn out to involve some loss of generality, as shown by Fudenberg and Tirole (1985)]. He assumes that time is discrete, and that firms move in alternate periods (since periods are of very short duration, approximately simultaneous adoption is possible). Let us denote the subgame perfect equilibrium by $\{T_i^*\}_{i=1}^n$, and let $V^i(T_i^*, T_{-i}^*)$ denote firm i 's payoff, where T_i^* denotes the equilibrium adoption time for firm i assuming that equilibrium adoption decision rules are used. Suppose without loss of generality that the i th firm is also the i th adopter. Assuming that optimal continuation play is independent of the identities of those firms who have already adopted, it can be shown that all firms must make the same profit in equilibrium; that is, $V^i(T^*) = V^j(T^*)$ for all i and j . To see this, suppose otherwise; suppose that $V^i(T^*) > V^j(T^*)$ for some $i \neq j$. If $i < j$ (that is, i adopts before j), then j cannot be using an optimal decision rule, because j could essentially "become" i by adopting slightly before i , causing firms $i, i+1, \dots, n$ to respond optimally, and leaving j with profits arbitrarily close to $V^i(T^*)$. If $i > j$ (that is, j adopts before i), firm j could postpone adoption and again receive almost $V^i(T^*)$ by just pre-empting firm i . The adoption dates can be computed by backwards induction.

Proposition 37 [Judd (1983)]

There is a unique set of (subgame perfect) equilibrium innovation times $\{T_i^*\}_{i=1}^n$, and $T_i^* \leq \tau_i^*$ for all i .

Thus, the innovation still diffuses into use, but its diffusion is more rapid when firms are able to respond quickly to rivals' actions than when they are unable to respond.

Fudenberg and Tirole (1985) show that the assumption of discrete time is not without loss of generality. In their continuous time model (the analysis of which necessitates the use of some quite technical arguments), there may also exist a continuum of equilibria involving simultaneous adoption. These equilibria can be Pareto ranked, with later adoption being preferred. Finally, if one is free to assign different continuation values (that is, different subgame equilibria) to different adoption histories, then profits need not be equalized across all firms when $n > 2$.

Quirmbach (1986) has analyzed the case of coordinated adoption behavior in the framework of Reinganum (1981b, 1981c), and finds that diffusion is characteristic of optimal adoption under a variety of alternative objective functions. The key elements which combine to generate diffusion are (1) declining incremental benefits for later adopters, and (2) declining adoption costs. Thus, although there is no ex ante heterogeneity among firms, these characteristics of the market (along with firms' abilities to perceive them) generate ex post heterogeneity in the form of diffusion. For example, consider the case of coordinated adoption by all n members of the industry. Assuming noncooperative production, joint industry profits can be written:

$$W(\tau) = \sum_{m=0}^n \int_{\tau_m}^{\tau_{m+1}} [m\pi_1(m) + (n-m)\pi_0(m)] e^{-rt} dt - \sum_{m=1}^n p(\tau_m),$$

where again $\tau_0 \equiv 0$ and $\tau_{n+1} \equiv \infty$. Recall Assumption 8; defining $\Delta\pi(m) \equiv \pi_1(m) - \pi_0(m-1)$, Assumption 8 says that $\Delta\pi(m-1) > \Delta\pi(m) > 0$; that is, the incremental benefit of adoption declines with the number of previous adopters. For the cooperative firms, the analogs to Assumptions 8 and 9 are as follows.

Assumption 10

$\Delta B(m-1) > \Delta B(m) > 0$, where

$$\begin{aligned} \Delta B(m) \equiv & [m\pi_1(m) + (n-m)\pi_0(m)] \\ & - [(m-1)\pi_1(m-1) + (n-m+1)\pi_0(m-1)]. \end{aligned}$$

Assumption 11

(a) $\lim_{t \rightarrow 0} p(t) = -\lim_{t \rightarrow 0} p'(t) = \infty$; (b) $\lim_{t \rightarrow \infty} p'(t) > 0$; (c) $p''(t) > re^{-rt}[\Delta B(1)]$ for all t .

Given these assumptions, we can characterize the jointly optimal adoption dates $\{\tau_i^0\}_{i=1}^n$ by differentiating $W(\tau)$ to obtain, for $i = 1, 2, \dots, n$:

$$\partial W(\tau^0)/\partial \tau_i = -\Delta B(i)e^{-r\tau_i^0} - p'(\tau_i^0) = 0. \quad (17)$$

Proposition 38 [Quirmbach (1986)]

The cooperative optimum is characterized by a diffusion of innovation which is uniformly slower than the noncooperative diffusion; that is, $\tau_i^0 > \tau_{i-1}^0$ and $\tau_i^0 > \tau_i^*$, for $i = 1, 2, \dots, n$.

The innovations envisioned above were those which exerted negative externalities upon the remaining members of the industry. However, in some cases the adoption of an innovation confers positive externalities upon all users. Examples include communications systems which allow agents to converse with one another, provided both employ compatible systems, and products such as video cassette players (or personal computers) where more movies (or software) will be available if more units are in operation. Unlike innovations which exert negative externalities on rival firms, these innovations suffer from an individual's unwillingness to adopt unilaterally; expectations about whether others will follow are crucial to the behavior of initial adopters. An interesting recent contribution to the literature on this subject is Farrell and Saloner (1985); for related work, see Dybvig and Spatt (1983) and Katz and Shapiro (1985a).

Consider an industry composed of two firms. When an innovation, consisting of a new standard, is announced, firms noncooperatively decide whether to adopt it. Adoption is considered to be an irreversible decision. For firm j , define $B_j(1, Y)$ to be the net benefit to firm j of unilaterally switching from the old standard X to the new standard Y . $B_j(2, Y)$ is the net benefit to j if both firms switch to Y . Status quo profits are normalized to zero so that firm j will be in favor of a change by the entire industry if and only if $B_j(2, Y) > 0$. Let $B_j(1, X)$ denote j 's payoff if j unilaterally remains with the old standard X , while the other firm switches to the new standard Y . By the normalization assumption $B_j(2, X) = 0$. The assumption that adoption of the new standard confers positive externalities upon other adopters is formalized below.

Assumption 12

For $k = X$ or Y , $B_j(1, k) \leq B_j(2, k)$. That is, whatever choice j makes, he prefers to have the other firm make the same choice.

Proposition 39 [Farrell and Saloner (1985)]

Suppose that $B_i(2, Y) > B_i(2, X)$ and $B_j(2, Y) > B_j(1, X)$ for some i and $j \neq i$. Then the unique perfect equilibrium involves all firms switching.

To see why, assume that there are two decision periods. Since the firms have complete information, each firm can foresee whether or not the other firm will follow; thus if $B_i(2, Y) > B_i(2, X)$ and $B_j(2, Y) > B_j(1, X)$ for $j \neq i$ and $i = 1$ or 2 , then the unique (subgame perfect) equilibrium involves all firms switching. This is because the firm which satisfies the hypotheses can ensure adoption by adopting unilaterally in the first period, knowing that the other firm will follow.

Since the assumption that $B_j(2, Y) > B_j(1, X)$ is weaker than the assumption that $B_j(2, Y) > B_j(2, X)$, an immediate corollary of Proposition 39 is that if $B_j(2, Y) > B_j(2, X)$ for $j = 1, 2$, then the unique perfect equilibrium involves all firms switching. Thus, firms need not be unanimous in their desire for the entire industry to adopt the new standard; it suffices for firm j to prefer company at the new standard than to maintain the old one alone.

Suppose now that each firm is uncertain about the other firm's evaluation of the new standard. This evaluation is simply indexed by a superscript i denoting the other firm's "type", with higher values of i indicating stronger preferences for the new standard Y . Let $B^i(1, k)$ denote the net benefits to the firm of type i from maintaining the standard k alone, and $B^i(2, k)$ the net benefits to the firm of type i from having the industry standardized at k , for $k = X$ or Y . The distribution of types is assumed to be uniform on the interval $[0, 1]$.

There are again two decision periods, and each firm can switch (irreversibly) either in period 1 or in period 2. In the first period, each firm must decide whether or not to switch based on its own type; in the second period, each firm must decide whether or not to switch based on its own type and the actions taken in period 1.

Assumption 13

$B^i(2, k) > B^i(1, k)$, $k = X, Y$. That is, networks are beneficial.

Assumption 14

$B^i(2, Y)$ and $B^i(1, Y)$ are continuous and strictly increasing in i ; that is, higher types are uniformly more eager to switch to Y .

Assumption 15

$B^1(1, Y) > 0$ and $B^0(2, Y) < B^0(1, Y)$. At least one type is willing to switch unilaterally, and at least one type is willing to remain alone at the old standard.

Given Assumption 15 and incomplete information regarding the other firm's type, a firm which switches early cannot be assured that the other firm will follow. For intermediate values of i , the firm's decision will depend nontrivially upon the decision of its predecessor.

Assumption 16

$B^i(2, Y) - B^i(1, X)$ is monotone increasing in i . Thus, if a type i firm would prefer an industry switch to Y to remaining alone at X , so would any type $i' > i$.

Definition

A *bandwagon strategy* for a firm is defined by a pair (i^*, \bar{i}) with $i^* > \bar{i}$ such that

- (a) if $i \geq i^*$, the firm switches in period 1;
- (b) if $i^* > i \geq \bar{i}$, the firm does not switch in period 1, and switches in period 2 if and only if the other firm switched in period 1;
- (c) if $i < \bar{i}$, the firm never switches.

A *bandwagon equilibrium* is defined to be a subgame perfect Bayesian Nash equilibrium in which each firm plays a bandwagon strategy. Farrell and Saloner characterize symmetric bandwagon equilibria; that is, those for which the pair (i^*, \bar{i}) is the same for both firms. They show that there is a unique such equilibrium, and that there are no equilibria which are not bandwagon equilibria.

Proposition 40 [Farrell and Saloner (1985)]

A unique symmetric bandwagon equilibrium exists.

The equilibrium in this game has two essential features: one is that it exhibits "bandwagon effects"; that is, some firm types would move early in hopes of inducing the other firm to follow (even though they would not prefer to switch unilaterally; that is, $B^i(1, Y) < 0$). Conversely, some firms who prefer the combined switch ($B^i(2, Y) > 0$) but are of types $i < i^*$ will wait until the other firm has switched before switching themselves. If both firms fall into this set, then the industry remains at the old standard, even though both firms prefer the combined switch to the new one, and even though such a switch would occur if information were complete. As Farrell and Saloner (p. 16) picturesquely put it, "both firms are fence-sitters, happy to jump on the bandwagon if it gets rolling but insufficiently keen to set it rolling themselves". There will be some combinations of firm types such that the sum of benefits is positive, yet the switch will not be made, and other combinations for which the sum of benefits is negative and the switch *will* be made (because one firm favors the switch enough to adopt unilaterally,

and the other prefers company at the new standard to maintaining the old one alone).

6. Conclusions

From the collection of symmetric models discussed in Section 2, we come away with an appreciation of the extent to which rivalry and appropriability interact to determine the incentives for individual firm investment in research and development. For instance, we have seen that whether or not entry results in increased or decreased investment by a given firm can depend critically upon the extent to which the rewards to innovation are appropriable. Similarly, when rewards are sufficiently appropriable, firms will overinvest relative to the cooperative optimum; on the other hand, when rewards are sufficiently inappropriable, firms will underinvest relative to that benchmark.

The models in Section 3 focused upon situations in which firms are asymmetrically placed. This asymmetry might be inherited, as in the incumbent/challenger models, or it might have developed over time as a consequence of intermediate successes. In comparing the auction and stochastic racing paradigms, we found that the associated equilibria were sometimes qualitatively different. In view of the possible differences in results, it seems important to choose the appropriate paradigm. The stochastic racing model seems to more accurately capture what we think of as research or "invention"; an activity that might or might not yield a worthwhile end-product, and one that may take more or less time and money than expected. The auction model may well be preferred for the case of development or new product introduction, in which any substantial technological uncertainties have already been resolved. Both research and development are significant aspects of innovative activity, and although the dividing line between them is by no means clear, some attempt should be made to match the appropriate paradigm to the specific application.

The models described in Section 4 examined restricted forms of licensing (primarily fixed-fee contracts) in the context of oligopolistic production. Under this assumption, it is generally concluded that firms who are not members of the venture which holds the patent are worse off as a consequence of the innovation, whether or not they obtain a license to use it. Thus, one incentive for licensing is this redistribution of wealth away from nonmembers to members. When additional research may yield an equivalent or better innovation, another incentive to license arises. In this situation, it may be optimal to license a drastic innovation to a potential entrant so as to decrease incentives for future research.

In Section 5 we described models of the diffusion of innovation over time. Although an information-based model of diffusion has definite appeal in terms of

realism, imperfect information is by no means necessary to explain why an innovation might diffuse relatively slowly into use. In the case of an innovation which is known to be profitable, the trick is to discover why firms might delay adoption. The key determinants prove to be declining adoption costs and the perception that the benefits of adoption decline with the number of previous adopters. This latter perception stems from strategic interactions in the product market; for example, oligopoly or multi-plant monopoly. It is interesting that a similar diffusion curve can be derived for the case of an innovation with positive external effects. When the value of adoption increases with the number of previous adopters, early adopters stimulate the subsequent diffusion of the innovation.

Most of this work has treated the process of research and development and the dissemination of its outcome as two separate issues. Invoking subgame perfection and dynamic programming suggests that this separation is legitimate and that one may fruitfully combine the results of these separate analyses. If various sorts of long-term commitments are possible, however, this approach will rule out some potentially interesting strategic features of the problem. For instance, a firm which could make credible a policy of never licensing to R&D rivals might be able to restrict the competition it faces in R&D. In this case, a model which simultaneously addresses both aspects of innovative activity would be required.

One important goal of future research should be to develop testable models of industry equilibrium behavior. The papers summarized here have used stark models in order to identify the significant characteristics of firms, markets and innovations which are likely to affect incentives to invest and/or adopt. But since it is largely restricted to these special cases (e.g. deterministic innovations, drastic innovations, two firms, symmetric firms), this work has not yet had a significant impact on the applied literature in industrial organization; its usefulness for policy purposes should also be considered limited. For these purposes, one needs a predictive model which encompasses the full range of firm, industry and innovation characteristics.

References

- Arrow, K.J. (1962a) 'Economic welfare and the allocation of resources for invention', in: NBER conference no. 13, *The rate and direction of inventive activity: Economic and social factors*. Princeton: Princeton University Press.
- Arrow, K.J. (1962b) 'The economic implications of learning by doing', *Review of Economic Studies*, 29:155–173.
- Balcer, Y. and Lippman, S.A. (1984) 'Technological expectations and adoption of improved technology', *Journal of Economic Theory*, 34:292–318.
- Barzel, Y. (1968) 'Optimal timing of innovations', *Review of Economics and Statistics*, 50:348–355.

- Dasgupta, P. and Stiglitz, J. (1980a) 'Industrial structure and the nature of innovative activity', *Economic Journal*, 90:266-293.
- Dasgupta, P. and Stiglitz, J. (1980b) 'Uncertainty, industrial structure and the speed of R&D', *Bell Journal*, 11:1-28.
- David, P.A. (1969) 'A contribution to the theory of distribution', Research Memorandum no. 71, Research Center in Economic Growth, Stanford University.
- Dixit, A.K. and Stiglitz, J.E. (1977) 'Monopolistic competition and optimum product diversity', *American Economic Review*, 67:297-308.
- Dybvig, P.H. and Spatt, C.S. (1983) 'Adoption externalities as public goods', *Journal of Public Economics*, 20:231-247.
- Farrell, J. and Saloner, G. (1985) 'Standardization, compatibility and innovation', *Rand Journal of Economics*, 16:70-83.
- Fethke, G.C. and Birch, J.J. (1982) 'Rivalry and the timing of innovation', *Bell Journal*, 13:272-279.
- Flaherty, M.T. (1980) 'Industry structure and cost-reducing investment', *Econometrica*, 48:1187-1209.
- Fudenberg, D. and Tirole, J. (1983) 'Learning-by-doing and market performance', *Bell Journal*, 14:522-530.
- Fudenberg, D. and Tirole, J. (1985) 'Preemption and rent equalization in the adoption of new technology', *Review of Economic Studies*, 52:383-401.
- Fudenberg, D., Gilbert, R.J., Stiglitz, J. and Tirole, J. (1983) 'Preemption, leapfrogging and competition in patent races', *European Economic Review*, 22:3-31.
- Futia, C.A. (1980) 'Schumpeterian competition', *Quarterly Journal of Economics*, 93:675-695.
- Gallini, N.T. (1984) 'Deterrence by market sharing: A strategic incentive for licensing', *American Economic Review*, 74:931-941.
- Gallini, N.T. and Winter, R.A. (1985) 'Licensing in the theory of innovation', *Bell Journal*, 16:237-252.
- Gilbert, R.J. and Newbery, D.M.G. (1982) 'Preemptive patenting and the persistence of monopoly', *American Economic Review*, 72:514-526.
- Grossman, G.M. and Shapiro, C. (1986) 'Optimal dynamic R&D programs', *Rand Journal of Economics*, 17:581-593.
- Grossman, G.M. and Shapiro, C. (1987) 'Dynamic R&D competition', *Economic Journal*, 97:372-387.
- Harris, C. and Vickers, J. (1985) 'Perfect equilibrium in a model of a race', *Review of Economic Studies*, 52:193-209.
- Hartwick, J.M. (1982) 'Efficient prizes in prototype development contests', *Economics Letters*, 10:375-379.
- Hirshleifer, J. and Riley, J.G. (1979) 'The analytics of uncertainty and information - An expository survey', *Journal of Economic Literature*, 17:1375-1421.
- Isaac, R.M. and Reynolds, S.S. (1985) 'Innovation and property rights in information: An experimental approach to testing hypotheses about private R&D behavior', discussion paper no. 85-30, University of Arizona.
- Jensen, R. (1982) 'Adoption and diffusion of an innovation of uncertain profitability', *Journal of Economic Theory*, 27:182-193.
- Jensen, R. (1984a) 'Adoption of an innovation of uncertain profitability with costly information', working paper no. 84-8, Ohio State University.
- Jensen, R. (1984b) 'Innovation adoption with both costless and costly information', working paper no. 84-22, Ohio State University.
- Jensen, R. (1984c) 'Information capacity and innovation adoption', working paper no. 84-33, Ohio State University.
- Judd, K.L. (1983) 'Equilibrium diffusion of innovation with no pre-commitment', mimeo.
- Judd, K.L. (1984) 'Closed-loop equilibrium in an innovation race with experience', mimeo.
- Judd, K.L. (1985a) 'Closed-loop equilibrium in a multi-stage innovation race', discussion paper no. 647, Center for Mathematical Studies in Economics and Management Science, Northwestern University.
- Judd, K.L. (1985b) 'On the performance of patents', *Econometrica*, 53:567-585.
- Kamien, M.I. and Schwartz, N.L. (1971) 'Expenditure patterns for risky R&D projects', *Journal of Applied Probability*, 8:60-73.
- Kamien, M.I. and Schwartz, N.L. (1972) 'Timing of innovations under rivalry', *Econometrica*, 40:43-60.

- Kamien, M.I. and Schwartz, N.L. (1974) 'Patent life and R&D rivalry', *American Economic Review*, 64:183–187.
- Kamien, M.I. and Schwartz, N.L. (1980) *Market structure and innovation*. Cambridge: Cambridge University Press.
- Kamien, M.I. and Tauman, Y. (1984) 'The private value of a patent: A game theoretic analysis', *Zeitschrift für Nationalökonomie/Journal of Economics* (Supplement), 4:93–118.
- Kamien, M.I. and Tauman, Y. (1986) 'Fees versus royalties and the private value of a patent', *Quarterly Journal of Economics*, 101:471–492.
- Kamien, M.I., Tauman, Y. and Zang, I. (1985) 'Optimal license fees for a new product', mimeo.
- Katz, M.L. (1986) 'An analysis of cooperative research and development', *Rand Journal of Economics*, 17:527–543.
- Katz, M.L. and Shapiro, C. (1985a) 'Network externalities, competition and compatibility', *American Economic Review*, 75:424–440.
- Katz, M.L. and Shapiro, C. (1985b) 'On the licensing of innovations', *Rand Journal of Economics*, 16:504–520.
- Katz, M.L. and Shapiro, C. (1986) 'How to license intangible property', *Quarterly Journal of Economics*, 101:567–590.
- Katz, M.L. and Shapiro, C. (1987) 'R&D rivalry with licensing or imitation', *American Economic Review*, 77:402–420.
- Lee, T. and Wilde, L.L. (1980) 'Market structure and innovation: A reformulation', *Quarterly Journal of Economics*, 94:429–436.
- Loury, G.C. (1979) 'Market structure and innovation', *Quarterly Journal of Economics*, 93:395–410.
- Lucas, Jr., R.E. (1971) 'Optimal management of a research and development project', *Management Science*, 17:679–697.
- Mamer, J.W. and McCardle, K.F. (1985) 'Uncertainty, competition and the adoption of new technology', mimeo.
- McCardle, K.F. (1985) 'Information acquisition and the adoption of new technology', *Management Science*, 31:1372–1389.
- Mortensen, D.T. (1982) 'Property rights and efficiency in mating, racing and related games', *American Economic Review*, 72:968–979.
- Park, J. (1984) 'Innovation competition under certainty', Ph.D. dissertation, Northwestern University.
- Quirmbach, H.C. (1986) 'The diffusion of new technology and the market for an innovation', *Rand Journal of Economics*, 17:33–47.
- Reinganum, J.F. (1981a) 'Dynamic games of innovation', *Journal of Economic Theory*, 25:21–41; (1985) 'Corrigendum', 35:196–197.
- Reinganum, J.F. (1981b) 'On the diffusion of new technology: A game theoretic approach', *Review of Economic Studies*, 48:395–405.
- Reinganum, J.F. (1981c) 'Market structure and the diffusion of new technology', *Bell Journal*, 12:618–624.
- Reinganum, J.F. (1982a) 'A dynamic game of R and D: Patent protection and competitive behavior', *Econometrica*, 50:671–688.
- Reinganum, J.F. (1982b) 'Strategic search theory', *International Economic Review*, 23:1–15.
- Reinganum, J.F. (1983a) 'Nash equilibrium search for the best alternative', *Journal of Economic Theory*, 30:139–152.
- Reinganum, J.F. (1983b) 'Uncertain innovation and the persistence of monopoly', *American Economic Review*, 73:741–748.
- Reinganum, J.F. (1984) 'Practical implications of game theoretic models of R & D', *American Economic Review*, 74:61–66.
- Reinganum, J.F. (1985) 'Innovation and industry evolution', *Quarterly Journal of Economics*, 81–99.
- Roberts, K. and Weitzman, M.L. (1981) 'Funding criteria for research, development, and exploration projects', *Econometrica*, 49:1261–1288.
- Rogerson, W.P. (1982) 'The social costs of regulation and monopoly: A game-theoretic analysis', *Bell Journal*, 13:391–401.
- Salant, S.W. (1984) 'Preemptive patenting and the persistence of monopoly: Comment', *American Economic Review*, 74:247–250.
- Scherer, F.M. (1967) 'Research and development resource allocation under rivalry', *Quarterly Journal of Economics*, 81:359–394.

- Scherer, F.M. (1980) *Industrial market structure and economic performance*. Chicago: Rand McNally.
- Schumpeter, J.A. (1942) *Capitalism, socialism and democracy*. New York: Harper & Row.
- Selten, R. (1975) 'Reexamination of the perfectness concept for equilibrium points in extensive games', *International Journal of Game Theory*, 4:25–55.
- Spence, A.M. (1976) 'Product selection, fixed costs and monopolistic competition', *Review of Economic Studies*, 43:217–235.
- Spence, A.M. (1981) 'The learning curve and competition', *Bell Journal*, 12:49–70.
- Spence, A.M. (1984) 'Cost reduction, competition and industrial performance', *Econometrica*, 52:101–121.
- Stewart, M.B. (1983) 'Noncooperative oligopoly and preemptive innovation without winner-take-all', *Quarterly Journal of Economics*, 98:681–694.
- Stoneman, P. (1983) *The economic analysis of technological change*. Oxford: Oxford University Press.
- Tandon, P. (1984) 'Innovation, market structure, and welfare', *American Economic Review*, 74:394–403.
- Telser, L.G. (1982) 'A theory of innovation and its effects', *Bell Journal*, 13:69–92.
- Vickers, J. (1984) 'Notes on the evolution of market structure when there is a sequence of innovations', mimeo.