Economic Theory **21**, 673–695 (2003)

Economic Theory

© Springer-Verlag 2003

Closed-loop equilibrium in a multi-stage innovation race*

Kenneth L. Judd

Hoover Institution, Stanford, CA 94305, USA (e-mail: judd@hoover.stanford.edu)

Received: January 3, 2002; revised version: June 14, 2002

Summary. We examine a multistage model of an R&D race where players have multiple projects. We also develop perturbation methods for general dynamic games that can be expressed as analytic operators in a Banach space. We apply these perturbation methods to solve races with a small prize. We compute second-order asymptotically valid solutions for equilibrium and socially optimal decisions to determine qualitative properties of equilibrium. We find that innovators invest relatively too much on risky projects. Strategic reactions are ambiguous in general; in particular, a player may increase expenditures as his opponent moves ahead of him.

Keywords and Phrases: Multistage races, Perturbation methods, Dynamic games.

JEL Classification Numbers: C630, 0310.

1 Introduction

Innovation processes are important for economic growth and has been the subject of much study. The early work of Kamien and Schwartz, summarized in Kamien and Schwartz (1982), concentrated on the decision-theoretic problems associated with innovation and lead to the study of equilibrium of competition in innovation contained in Loury (1979) Lee and Wilde (1980), Reinganum (1981–1982), and Dasgupta and Stiglitz (1980a,b). These analyses examined one-shot innovation processes – as long as no competitor won, all competitors were equal. Also, they

* This is the final version of Judd (1985). The author gratefully acknowledges the comments of anonymous referees, Paul Milgrom, seminar participants at Northwestern University, the University of Chicago, the 1984 Summer Meetings of the Econometric Society, University of California at Berkeley, Stanford University, and Yale University, and the financial support of the National Science Foundation (SES-8409786, SES-8606581)

This content downloaded from 128.135.100.108 on Fri, 11 Oct 2013 15:51:58 PM
All use subject to JSTOR Terms and Conditions

assumed that there was just one available innovation technology. More recently, Fudenberg et al. (1983), Lee (1982), Telser (1982), and Harris and Vickers (1985) examined multi-stage games. Bhattacharya and Mookherjee (1986) examine allocation of innovative effort across alternative projects in a one-shot, simultaneous move game. However, the specifications of R&D processes used in those models were limited for reasons of tractability.

This paper has two purposes. First, we examine the equilibrium of a race for a prize where each of two agents controls independent R&D projects. At each moment, each agent works to advance his own state of knowledge while knowing that of his opponent. The race ends when one of the firms has achieved a critical state of knowledge, here called "success." There is a social gain realized at that time and some of that gain is paid to the winner as a prize. This model is intended to be a stylized representation of a multi-stage R&D race, and we use it to address questions concerning firms' strategies and the allocation of resources across alternative investments.

Second, we use approximation techniques to precisely examine the nature of the subgame perfect equilibrium in some cases of our game. Global closed-form solutions to our general model are not known. The approximation techniques used below provide precise answers to interesting questions for an open set of games. While such an approach does not yield a global resolution of the issues, it does provide guidance as to what is true in some cases and points out the critical factors. The presentation of this analysis is itself a second independent purpose of this paper since it represents a general way to analyze subgame perfect equilibria of dynamic games without imposing economically unmotivated restrictions on critical model elements.

Even though this paper uses perturbation methods to analyze only multistage R&D races, the methods are quite general since they are based on a version of the implicit function theorem for analytic functions in Banach spaces. The perturbation methods described below can be used to solve many other dynamic games with state variables. We introduce the technique using R&D games since the application is clear and not excessively encumbered with complex notation. Budd et al. (1993), in independent work, also presented a perturbative analysis (using different asymptotic theorems) of a different dynamic game. These papers are just a couple of examples of the potential of perturbation and asymptotic methods for economic analysis.

Other numerical methods, such as those presented in Judd (1992), could be used to solve our R&D game. For example, Doraszelski (2002) uses projection methods from Judd (1992) to solve a generalization of Reinganum (1981). Perturbation methods and projection methods are complementary and both have a role. Perturbation methods need not make functional form assumptions and the results are theorems about an open set of cases, but that set may be small and not include some empirically interesting cases. Numerical procedures, such as projection methods, can examine a much more varied range of cases of a model but must make functional form assumptions, can examine only a finite number of instances, and suffers from numerical error. In general one would like to use both methods when studying a general model. See Judd (1997) for a more extensive discussion of the trade-offs between perturbation and alternative numerical methods. We focus on perturba-

This content downloaded from 128.135.100.108 on Fri, 11 Oct 2013 15:51:58 PM
All use subject to JSTOR Terms and Conditions

tion methods in this paper and the kind of qualitative results such an approach can produce.

Specifically, we find that if the prize to the innovator and the net social benefits are "small" (in a sense specified below) the model yields several results. First, if the prize equals the benefits, there is excessive innovation effort, a result common to innovation models of this nature. Second, since agents can be at differing levels of knowledge in our model, we would like to compare the relative efficiency of resource allocation across firms. We find that lagging firms are less efficient in that if there is to be a momentary subsidy of innovation effort, the first dollar of such a subsidy should go to the leading firm.

Since agents choose how to allocate resources across projects of varying riskiness, we examine the allocative efficiency of investment within firms. We find that there is relatively excessive investment in the riskier projects. We also want to know how each innovator reacts to his rival's advances. We find that if one firm advances, the other will surely increase its effort in risky projects, a movement contrary to the socially optimal reaction. However, we find that he may increase or decrease effort in less risky projects. This contrasts with the arguments in Fudenberg et al. (1983) and Harris and Vickers (1985) which find that each firm's effort is a decreasing function of his opponent's position. This finding shows that the special assumptions used in previous multistage models are critical for their results, and that some of their results do not hold up in general. This is a good example of how the perturbation method can find results missed by methods relying on closed-form solutions.

We also use this model to examine the nature of optimal R&D policy. First, we find that, in spite of the relative inefficiency of the lagging firm, it is optimal to let competition continue until some firm enjoys complete success. Second, we also find that the optimal prize asymptotically equals the social benefit when the social benefit is small.

Some of our results hold because the multi-stage nature of the game disappears if the prize is small. However, other features, particularly the nature of firms' reactions and the risk allocation decisions, are related critically to the multi-stage subgame perfect nature of our analysis. This indicates that we have successfully peeked into the nature of subgame perfect equilibrium in innovation races. Furthermore, we indicate how other approximations could be carried out, showing that the perturbation approach does not rely on the small prize assumption. The only requirement for the application of the approximation techniques used below is some example with a known closed-form solution. Since the key theorems and techniques are general, it is clear that our perturbation approach to closed-loop subgame perfect equilibrium analysis is of general applicability for game-theoretic analysis of dynamic strategic interaction.

Some of the features of our analysis will initially appear odd. In particular, we assume that the prize and social benefits are small, and we examine second-order terms of a Taylor series expansion instead of the more common linear terms. However, neither feature negates the usefulness of the perturbation approach. Perturbation analysis is often based on a degenerate case, particularly in the physical sciences where perturbation methods are commonly used. For example, the Ein-

This content downloaded from 128.135.100.108 on Fri, 11 Oct 2013 15:51:58 PM
All use subject to JSTOR Terms and Conditions

stein equations of general relativity theory are generally intractable. However, many of that theory's powerful implications, such as gravitational radiation, have come from computing the high-order terms of power series solutions to the general relativity field equations around the case of a universe with no matter and no gravity. If physicists find nearly empty universes to be informative then the case of zero prize in R&D races should be informative for economists. This kind of "nearly degenerate" approach combined with first-order approximations has often been useful in economics. In macroeconomics, we often implicitly assume that shocks are nearly zero and use linearizations of dynamic systems around their steady states to examine dynamic stochastic economies. In public finance, we often implicitly assume that taxes are nearly zero and say that the excess burden of a tax is proportional to the square of the tax rate. The usefulness of high-order approximations around the case of a zero prize will be apparent below.

Section 2 describes the general model and compares it with the multiperiod models of Fudenberg et al. (1983), Lee (1982), Telser (1982), and Harris and Vickers (1985). Section 3 gives an overview of the approximation technique which we utilize below and Section 4 demonstrates it in detail for a useful special case. We then examine the nature of our problem for the case of a small net social value, discussing in Section 5 the social optimum and in Section 6 the competitive outcome. Section 7 compares the optimum and equilibrium outcomes and Section 8 examines some implications for optimal social policy given rivalrous innovation. Section 9 discusses the relation of our analysis with other approaches, arguing that our approach gives a method to generalize solutions to problems which generate closed-form solutions. Section 10 concludes.

2 A multistage model of a race

We investigate a simple model of multi-state innovation with two firms. Competition takes the form of a race. The position of each firm is denoted by a scalar with firm one at $x < 0$ and two at $y < 0$. Success is defined by one firm crossing 0; therefore we assume x and y are initially both negative and that the current state of the race, (x, y), is represented by a point in the third quadrant of the plane. A firm can attempt to improve its position by investments which determine the probability of jumping to a better state of knowledge.

Jumps occur in two ways. There are *partial jumps* which will move a firm closer to the goal. If firm one (two) is at point $x < 0$ ($y < 0$) and invests at rate u (v) on partial jump investment, then there is a partial jump with probability udt (vdt) during a dt time interval. If a partial jump occurs, there is a probability of $F(x)dt$ ($G(y)\,dt$) of hitting 0 and otherwise there is a probability of $f(s, x)ds$ ($f(s, y)ds$) of landing in the interval $(s, s + ds)$, $s < 0$. We assume that the distributions of the partial jumps are ordered by first-order stochastic dominance, that is, if $x' > x$, then $f(s, x')$ first-order stochastically dominates $f(s, x)$. We assume that f is bounded above and there are no moves backwards; hence, if $s < x$, then $f(s, x) = 0$. Note that $F(x) = 1 - \int_{[x,0)} f(x, a)ds$. $F(x)$ is increasing in x, by the stochastic ordering of f in x. We also assume that F is positive everywhere;

This content downloaded from 128.135.100.108 on Fri, 11 Oct 2013 15:51:58 PM
All use subject to JSTOR Terms and Conditions

this says that there is always some chance of jumping to the finish from any state x. Our procedures could handle the more general case, but at substantial notational cost and little substantive gain.

There are also *leaps* from x (y) to 0. If firm one (two) is at point $x < 0$ ($y < 0$) and invests at rate w (z) on leap investment, then firm one (two) leaps to $x = 0$ ($y = 0$) with probability $wG(x)dt$, $(zG(y)dt)$ during a time interval dt. The leaps will be called more risky since whenever investment is such that leaps and partial jumps have the same expected jump, the expected gain in the value of any convex function of position is greater for leaps. For the sake of simplicity, we assume square cost functions. That is, firm one's costs and the social costs associated with its choice of u and w are $\alpha u^2/2 + \beta w^2/2$, where α, $\beta > 0$. The costs associated with firm two's choices equal $\alpha v^2/2 + \beta z^2/2$. The first firm to succeed receives a prize of P, with no prize for the loser. We assume that the social benefit from any success is $B > 0$ and that $\rho > 0$ is the social and private discount rate.

This model differs from earlier multi-stage models in substantial ways. In Lee (1982) and in Telser (1982), a firm may pull away in the sense that it may achieve an increasingly superior cost structure, but the leading firm has no advantage in achieving any other low level of costs. In this model, a firm may pull away from its competition and final success is easier to achieve the more advanced it is. The ability to pull away and attain some dynamic advantage is present in models analyzed in Fudenberg et al. (1983) and in Harris and Vickers (1985) but they both assume very special structures for innovation costs and limit the investment choices of innovators. In particular, innovation is a natural monopoly in Harris and Vickers' model in that society would only want one innovation project commanding resources, a feature which limits the ability to address issues in patent policy and the structuring of incentives for innovation. Under our assumptions, however, there is a social value to having resources allocated to each innovation project since the marginal cost of effort is zero when the effort level is zero for each project.

Both Fudenberg et al. (1983) and Harris and Vickers (1985) focus on conditions under which a firm will surely win the patent race once it has any small advantage over its competitor. The information lag model studied in Fudenberg et al. paper is closely related to our model. In both models no firm knows what the other firm is currently doing, but both know the position of its opponent at the beginning of each period. The models differ in that the state of each firm responds stochastically to his efforts whereas Fudenberg et al. assume a deterministic response. They also make an increasing cost assumption concerning the relationship between effort and progress, but must make restrictive assumptions to render the analysis tractable.

All previous dynamic models have assumed only one kind of research investment. By permitting alternatives of varying riskiness, we can compare the relative allocation of resources among projects of varying riskiness. Finally, we also determine how relative efficiency of the two firms is related to their relative position, finding that the lagging firm is less efficient. We address the issue of when a competition should be ended and a winner granted the monopoly right to the innovation, a question previously ignored.

This general model can be used to address several issues in the economics of races. Before analyzing our model we will first present our approximation approach.

This content downloaded from 128.135.100.108 on Fri, 11 Oct 2013 15:51:58 PM
All use subject to JSTOR Terms and Conditions

3 Banach space approximations

The model described above is far too general to hope for a closed-form solution. Nor will the structure be sufficiently tractable so as to allow for comparative static analysis as in previous work. We will instead use basic approximation techniques to study our general model for cases near some tractable case. This section reviews the basic mathematics underlying our approach and discusses its usefulness.

The primary tool used below is the generalization of Taylor's theorem and the Implicit Function theorem in R to Banach spaces. Taylor's theorem for a real-valued function over R, says that if $f(x;z)$ is C^{n+1} in x on $[0,\ b]$ for all z, where we think of x as the variable in R and z as a parameter, then for any z and any $x \in (0,\ b)$ there is a $\xi \in (0,\ x)$ such that

$$f(x;z) = \sum_{k=0}^{n} f^{(k)}(0;z)\frac{x^k}{k!} + f^{(n+1)}(\xi;z)\frac{x^{n+1}}{(n+1)!}.$$

This states that the n-th degree polynomial in Taylor's Theorem is an $o(x^n)$ approximation of $f(x;z)$ for x near zero. In particular, properties such as positivity and convexity which hold for this approximating polynomial near zero also hold for $f(x;z)$ when x is near zero.

Since equilibria in our games will be expressed as a collection of functional equations of the equilibrium strategies, we will use the Implicit Function Theorem to compute equilibria for games close to games for which solutions are known. Generally, the Implicit Function Theorem states that f can be uniquely defined for x near zero by the equation $H(x, f(x); z) = 0$ if $H_1(0, f(0); z)$ exists and $H_2(0, f(0); z)$ is invertible. This allows us to implicitly compute the derivatives of f with respect to x as a functions of x and z, leading to a polynomial approximation for f.

However, our strategies are not going to be vectors of real numbers, but rather functions of the state variable, objects from infinite-dimensional spaces. It is necessary, therefore, to first introduce some terminology from nonlinear functional analysis and state the generalized Implicit Function Theorem for functions and power series over Banach spaces. Suppose that X and Y are Banach spaces, i.e., normed complete vector spaces. A map $M : X^k \to Y$ is *k-linear* if it is linear in each of its k arguments. It is a *power* map if it is symmetric and k-linear, in which case it is denoted by $Mx^k \equiv M(x, x, \ldots, x)$. The norm of M is constructed from the norms on X and Y, and is defined by

$$||M|| = \sup_{||x_i||=1,\ i=1,2,\ldots,k} ||M(x_1, x_2, \ldots, x_k)||$$

For any fixed x_0 in X, consider the infinite sum in Y:

$$Tx = \sum_{k=1}^{\infty} M_k(x - x_0)^k$$

where each of the M_k is a k-linear power map from X to Y. When the infinite series converges, T is a map from X to Y. It will be convenient to associate a real

This content downloaded from 128.135.100.108 on Fri, 11 Oct 2013 15:51:58 PM
All use subject to JSTOR Terms and Conditions

valued series, called its *majorant series*, with T

$$\sum_{k=0}^{\infty} ||M_k||\,||x - x_0||^k$$

The important connection between the power series for T and its majorant series is that T will converge whenever its majorant series does.

Definition 1. *T is analytic at x_0 if and only if it is defined for some neighborhood of x_0 and its majorant series converges for some neighborhood of x_0.*

With these definitions, we can now state the analytic operator version of the Implicit Function Theorem.

Theorem 2. *Implicit Function Theorem for Analytic Operators: Suppose that*

$$F(\varepsilon, x) = \sum_{n,k=0}^{\infty} \varepsilon^n M_{nk}\, x^k \tag{1}$$

defines an analytic operator, $F : U(0,0) \subset R \times X \to Y$, where $U(0,0)$ is a neighborhood of $(0,0)$ in $R \times X$. Furthermore, assume that $F(0,0) = 0$ and that the operator $M_{01} : X \to Y$, representing the Frechet cross-partial with respect to x at $(0,0)$, is invertible. Consider the equation

$$F\left(\varepsilon, x(\varepsilon)\right) = 0 \tag{2}$$

implicitly defining a function $x(\varepsilon) : R \to X$. The following are true:

1. *There is a neighborhood, $\mathcal{V}$, of $0 \in R$, and a number, $r > 0$, such that (2) has a unique solution of $||x|| < r$ for each $\varepsilon \in \mathcal{V}$.*
2. *The solution, $x(\varepsilon)$, of (2) is analytic at $\varepsilon = 0$, and, for some sequence of x_n in X, can be expressed as*

$$x(\varepsilon) = \sum_{n=1}^{\infty} x_n\, \varepsilon^n \tag{3}$$

where the coefficients x_n can be determined by substituting (3) into (1) and equating coefficients of like powers of ε.

3. *The radius of convergence of the power series representation in (3) is no less than that of the analytic map, $z(\varepsilon) : R \to R$, defined implicitly for some neighborhood of 0 by*

$$0 = \sum_{n,k=0}^{\infty} \varepsilon^n\, ||M_{nk}||\, z(\varepsilon)^k \tag{4}$$

Furthermore, for some sequence z_n of real numbers,

$$z(\varepsilon) = \sum_{n=0}^{\infty} \varepsilon^n\, z_n$$

represents the solution to (4) and $|z_n| > ||x_n||$.

This content downloaded from 128.135.100.108 on Fri, 11 Oct 2013 15:51:58 PM
All use subject to JSTOR Terms and Conditions

Proof. See Zeidler (1986). □

The actual execution of the mathematics in Theorem 2 turns out to be elementary since our task is reduced to recursive computation of x_n terms. The term-by-term approach alluded to in item 2 in Theorem 2 will be illustrated in the next section.

However, we should first discuss the value of such an approximation approach. Our objective below is to apply it to examine subgame-perfect equilibria in our model. In most of the analysis below, we will express equilibrium strategies and values as functions of the prize, P, social benefit, B, and the position, (x, y) and examine approximations for them around the case of a zero prize and no social value. At first blush, approximations based on such cases may appear useless since the case of a zero prize is degenerate. A number of considerations justify the effort and indicate the general value of this approach.

First, the approximations can provide counterexamples to conjectures. Suppose $g_1(P)$ and $g_2(P)$ are functions of interest, and it is initially conjectured that $g_1(P) > g_2(P)$. If we can show that $g_1(0) = g_2(0)$ and $g_1'(0) < g_2'(0)$, then there must be an interval of $P > 0$ where $g_1 < g_2$, contradicting the conjecture. This in fact will occur below when we discuss equilibrium reaction functions.

Furthermore, suppose g_1 depends on some function F, i.e., $g_1(P; F)$. More generally, one could identify conditions on F which lead to the "$g_1 > g_2$" conjecture failing. In models of dynamic competition, we often make special assumptions about the functional form of such F's. After deriving our results, however, we usually don't know exactlywhat general feature of the functional form was crucial. Our approach below will find exactly what features of all structural elements are critical for any results for the case of a small prize. Whenever the intuition gathered from such an analysis does not depend on P being nearly 0, then we have perhaps discovered a robust feature of the model. Generally, we study such approximations not because they are valid for nearly degenerate cases, but rather that they likely indicate patterns which continue to hold more generally.

Second, any analytical investigation of this model must focus on cases which are degenerate in some ways. Note that the models of Lee and Wilde, Reinganum, and Fudenberg et al. are all special cases of this general model (or some slightly different general model) which are degenerate in some dimension. For example, Lee and Wilde, and Reinganum implicitly assume that the success probability function $G(x)$ is independent of the position x, making position irrelevant. Also, $F(x)$ is essentially absent in their models, as if α were infinite. Each of these special cases are of interest despite their degeneracies. However, if we are interested, for example, in a precise look at how innovators react to each other's successes, it is valuable to look at cases in which there are as few unmotivated restrictions on the underlying stochastic structure as possible. It is unfortunate that we may have to assume a small prize, but that is the price we pay here to attain this particular goal. Finally, the technique that is exploited below can be used generally to develop a robustness analysis for all the special cases studied previously.

This content downloaded from 128.135.100.108 on Fri, 11 Oct 2013 15:51:58 PM
All use subject to JSTOR Terms and Conditions

4 An example: the case of a single firm

In this section we analyze the case of a single firm. This will illustrate the analysis used below and will also be used later when we examine the optimal stage at which to end the race. Also, to cut down on inessential clutter, we will examine here only the simple case when β is infinite. The general solution will be displayed at a later point.

The case of a single innovator is a dynamic programming problem. If $M(x)$ is the value of position x to the firm, then the Bellman equation for M is

$$M(x) = \max_u \left\{ -\frac{\alpha u^2}{2} dt + M(x)(1-\rho dt)\,(1-udt) + (1-\rho dt)udt \left(\int_x^0 M(s) f(s,x) ds \right) + udtPF(x) \right\} \tag{5}$$

where dt is the infinitesimal unit of time.[1] The individual terms of the maximand represent the expected value of innovative effort. If the rate of effort is u, R&D expenditures during dt equal $-(1/2)\alpha u^2 dt$. With probability $1 - udt$ there will be no success, implying that the state of knowledge dt units of time in the future will remain x and the value will remain $M(x)$. The current unconditional expected value of that event is $(1-\rho dt)(1-udt)M(x)$. With probability udt there will be a jump to some $s \in (x, 0]$. If x jumps to 0, an event with probability $F(x)$ conditional on a jump occurring, the immediate reward is P. Since the reward is immediate, no discounting occurs. If x jumps to a point $x' \in (s, s+ds)$, an event with a conditional probability of $f(s,x)ds$, the value becomes $M(s)$ in the next period. In the foregoing, $\int_x^0 \dots ds$ will represent $\int_{[x,0)} \cdots ds$, thereby ignoring the atom at $x = 0$. We use this notation to distinguish reaching an intermediate stage from that of winning. Equation (5) therefore states that the value of a position equals the maximum expected current value of future positions net of current costs.

Solving the maximization problem in (5) shows that

$$\alpha u = \int_x^0 M(s) f(s,x) ds + PF(x) - M(x) \tag{6}$$

Substituting this first-order condition into the control equation yields

$$0 = \frac{1}{2\alpha} \left(\int_x^0 M(s) f(s,x) ds + PF(x) - M(x) \right)^2 - \rho M(x) \tag{7}$$

By standard dynamic optimization methods, there exists a unique such M.

We cannot generally find a closed-form solution for M in (7). We will instead use Theorem 2 to give us precise information about M for any F and an open set of values for P. Note that this fits our discussion above. If we assume that the value function M is in the Banach space of real-valued functions on the negative reals with the supremum norm, then the RHS of (7) is the sum of a linear and a bilinear

[1] We will employ the intuitive infinitesimal notation of equation (5). All the dynamic programming equations can be derived formally, as in Bryson and Ho.

This content downloaded from 128.135.100.108 on Fri, 11 Oct 2013 15:51:58 PM
All use subject to JSTOR Terms and Conditions

operator acting on M and the real parameter P. To proceed in this fashion one should examine dimensionless versions of a problem since the concept of "small" should not depend on the choice of units. Define $m \equiv M/P$ to be the value of problem (5) relative to the prize. m is a dimensionless quantity representing the value of the game which will yield a substantive concept of small.

Rewritten in terms of m, (7) becomes the equation

$$m(x;p) = p\left(\int_x^0 m(s)f(s,x)ds + F(x) - m(x)\right)^2 \tag{8}$$

where $p \equiv P/2\alpha\rho$ is the size of the prize relative to the marginal cost of innovation and the cost of capital. Since the dimension of ρ is (time)$^{-1}$ and that of α is (dollars)$\times$(time), p is dimensionless and will be our measure of the prize. Since m, p, f, and F are all dimensionless, (8) is a dimensionless representation of (7). When p is zero, (8) yields the obvious solution, $m(x) = 0$. p may be zero either because P is zero or because $\alpha\rho$, the costs, are infinite. Focussing on p makes clear that we are not assuming that the prize itself is small but rather it is small compared to the rate of increase in marginal cost. This will imply that the prize is to the first order equal to the costs and that the net profits of an innovator are small relative to the prize. The interpretation that the prize just covers the opportunity costs of innovative activity makes our focus on small p more plausible.

Once we transform (7) into a dimensionless equation, we also must transform other variables of interest; in particular, the control variable, u. However, u is not dimensionless since it measures effort per unit of time and depends on the time unit. We can rewrite (6) into the dimensionless form:

$$\tilde{u} \equiv \frac{u}{\rho} = 2p\left(\int_x^0 m(s)f(s,x)ds + F(x) - m(x)\right) \tag{9}$$

where $\tilde{u}$ is the dimensionless rate of effort per normalized unit of time.

We now illustrate computing a local solution to (8). If $p = 0$, then $m = 0$. Applying the Implicit Function Theorem tells us that $m(x;p)$ is smooth in p for p near zero, and that we can approximate $m(x;p)$ for such p up to $O(p^n)$

$$m(x;p) \approx m(x;0) + pk^1(x) + p^2k^2(x) + \ldots + p^nk^n(x) \tag{10}$$

where we define $k^n(x) \equiv \frac{1}{n!}\frac{\partial^n m}{\partial p^n}(x,0)$. First note that $m(x;0) = 0$ since a zero prize makes the optimal value of the problem zero.

Differentiating (8) with respect to p and evaluating at $p = 0$ shows that

$$k^1(x) = F(x)^2 \tag{11}$$

Taking a second derivative of (8) with respect to p, evaluating it at $p = 0$, and using the fact that $\partial m/\partial p(x;0) = k^1(x) = F(x)^2$, we find that[2]

$$k^2(x) = 2F(x)\left(\int_x^0 F(s)^2 f(s,x)ds - F(x)^2\right) \tag{12}$$

[2] Our notation will be burdened with many superscripts. Superscripts to functional names, as in $k^2(x)$, will represent distinct functions, and will *never* represent iteration as in $k(k(x))$. Superscripts to functional evaluations represent powers. Hence, $k^2(x)^3$ is the cube of the value of the function k^2 evaluated at x.

This content downloaded from 128.135.100.108 on Fri, 11 Oct 2013 15:51:58 PM
All use subject to JSTOR Terms and Conditions

Continuing in this fashion, one can recursively compute $k^n(x)$ for any n justified by the known smoothness of m in terms of p. Note that *no* smoothness of m in x need be assumed.

It is usually quite tedious to do all the differentiation explicitly. A standard trick in perturbation analysis is to take the polynomial approximation for m in terms of p in (10), insert it into (8), and conduct the algebraic operations indicated in (8) to get an approximate polynomial representation of (8). Equation (8) then becomes

$$\begin{aligned} &pk^1(x) + p^2k^2(x) + \ldots = \\ &pF(x)^2 + 2p^2\left(\int_x^0 k^1(s)f(s,x)ds) - k^1(x)) + \ldots\right) \end{aligned} \tag{13}$$

If we equate terms linear in p in(13), we find that $k^1(x) = F(x)^2$. Combining p^2 terms and using the computed solution for k^1 demonstrates (12). Continuing in this fashion will yield all k^n functions. Since this approach yields the terms of the Taylor series more efficiently, we will use it below.

From these expressions we may infer several obvious properties of the optimal control for small p. For example, if p is small, effort increases as one is closer to the finish. This follows from the observation that the $pF(x)$ term dominates in (9) since m is $O(p)$ implying that u rises as $F(x)$, and hence x, rises. Also, u falls and as α and ρ rise, an intuitive result since both represent costs. Using this approach, we next examine the total social optimum when we have two separate projects and two firms.

5 The social optimum

Let $W(x,y)$ be the social value function when current states are x and y. Then the Bellman equation becomes

$$\begin{aligned} W(x,y) = \max_{u,v,w,z}\ &\frac{1}{2}(-\alpha u^2 - \alpha v^2 - \beta w^2 - \beta z^2)dt \\ &+udt\left(\int_x^0 W(s,y)f(s,x)ds + BF(x)\right)(1-\rho dt) \\ &+vdt\left(\int_y^0 W(x,s)f(s,y)ds + BF(y)\right)(1-\rho dt) \\ &+(wG(x)+zG(y))(1-\rho dt)Bdt \\ &+(1-\rho dt)(1-(u+v+wG(x)+zG(y))dt)W(x,y) \end{aligned} \tag{14}$$

The first-order conditions of (14) imply

$$\begin{aligned} \alpha u &= \textstyle\int_x^0 W(s,y)f(s,x)ds + BF - W(x,y) \\ \beta w &= G(x)(B - W(x,y)) \end{aligned} \tag{15}$$

αv and βz may be expressed similarly. Using the first-order conditions, (15), for u and w, and the corresponding conditions for v and z, (14) becomes

$$\begin{aligned} 0 = &\left(E_x\{W(s,y)\} - W(x,y)\right)^2/2\alpha + \left(E_y\{W(x,s)\} - W(x,y)\right)^2/2\alpha \\ &+ \left(G(x)(B - W(x,y))\right)^2/2\beta + \left(G(y)(B-W(x,y))\right)^2/2\beta - \rho W(x,y) \end{aligned} \tag{16}$$

This content downloaded from 128.135.100.108 on Fri, 11 Oct 2013 15:51:58 PM
All use subject to JSTOR Terms and Conditions

where

$$E_x\{W(s,y)\} \equiv \int_x^0 W(s,y)f(s,x)ds + BF(x)$$
$$E_y\{W(x,s)\} \equiv \int_y^0 W(x,s)f(s,y)ds + BF(y)$$

Theorem 3. *There exists a unique solution, $W(x, y)$, to the social optimum problem, and $W(x, y)$ is analytic in B, α, β, and ρ.*

Proof. The RHS of (16) is an analytic operator on bounded functions over the nonpositive reals. When $P = 0$, the unique solution is $W = 0$. Furthermore, the cross Frechet derivative, first with respect to P then with respect to W, is $-\rho$, which is an invertible operator on bounded functions. Therefore, the conclusions follow from Theorem 2. □

We next compute an approximation for W. Suppose $W(x, y) = B(bh^1(x, y) + b^2h^2(x, y) + \ldots)$ is the approximating series for W around $B = 0$, which exists by the Implicit Function Theorem. We let $b = B/2\alpha\rho$ be a dimensionless measure of the social value and $\gamma = \alpha/\beta$ be the dimensionless ratio of costs across projects, and use them to create a dimensionless representation of W/B. The linear term, h^1, is computed to be

$$h^1(x,y) = F(x)^2 + F(y)^2 + \gamma\left(G(x)^2 + G(y)^2\right) \tag{17}$$

and the investment rules are approximated to $O(b^2)$ by

$$\frac{u}{\rho} \approx 2bF(x) + 2b^2\left(\int_x^0 h^1(s,y)f(s,x)ds - h^1(x,y)\right) \tag{18}$$
$$\frac{w}{\rho} \approx 2\left(b - b^2h^1(x,y)\right)\gamma G(x)$$

and similarly for v and z. The first-order approximations for u and w are as if the current hazard rate of immediate success was common to all stages since $\alpha u \approx BF(x)$ and $\beta w \approx BG(x)$ to $O(B)$. This indicates that the first-order behavior of this multi-stage game at any stage reduces to the behavior of a single-stage game. In particular, to a first order, the presence of other projects has no impact on investment rules. Intuitively, this is because for small B, effort levels are "small," the probability of success for any one project is "small," and by independence the probability of success by two projects is "small squared," hence negligible. Therefore, most of the interesting multi-stage questions we ask below will require examination of the h^2 function that appears in the $O(b^2)$ term.

Straightforward combinations of (17) and (18) prove Theorem 4.

Theorem 4. *For small B, the following hold for the optimal innovation policy:*

1. *as $x(y)$ increase, $u(v)$ and $w(z)$ increase and $v(u)$ and $z(w)$ fall;*
2. *$w(z)$ is increasing and concave in B;*
3. *$u(v)$ is increasing in B but may be convex or concave in B;*

This content downloaded from 128.135.100.108 on Fri, 11 Oct 2013 15:51:58 PM
All use subject to JSTOR Terms and Conditions

4. *W is increasing and convex in (x, y) if $F(x)$ and $G(y)$ are convex;*
5. *u and v (w and z) are decreasing in ρ and $\alpha(\beta)$;*
6. *and w and z are decreasing in α.*

Particularly note that, if the two firms were managed in a socially optimal fashion, each firm would increase its efforts on both projects as it advances, and the other would decrease its effort. Also, the magnitude of these reactions are on the order of B^2. These features will be substantially different in the equilibrium of the R&D race.

6 Equilibrium of the innovation game

We next solve for the symmetric subgame-perfect equilibrium of the corresponding game. We are implicitly assuming that the current states of both firms are common knowledge since if we had assumed that no firm could observe the position of his competitor then the open-loop solution would be the correct equilibrium concept. While this common knowledge aspect is certainly valid in sports races, it may appear awkward here. It asserts that a firm may know how much its opponent knows without knowing exactly what its opponent knows. This is not an unrealistic description of matters in knowledge-intensive activities. Academics, for example, should not be uncomfortable with this assumption since they often judge colleagues' relative levels of knowledge about a subject without having an equivalent level of expertise in the area. In sum, we are assuming that firms may determine their relative positions without actually having access to each other's knowledge.

Let $V(x, y)$ represent the value to firm one of state (x, y). We will examine symmetric equilibria, implying that $V(y, x)$ will represent the value to firm two of state (x, y). We also limit our examination to equilibria which depend only on the current state of the game. The Bellman equation for firm one is

$$\begin{aligned} V(x,y) = \max_{u,w} \{ & - \left(\alpha u^2/2 + \beta w^2/2\right)\, dt + wG(x)P(1-\rho\, dt)\, dt \\ & + u\, dt \left(\int_x^0 V(s,y) f(s,x) ds + PF(x) \right) (1-\rho\, dt) \\ & + v\, dt \left(\int_y^0 V(x,s) f(s,y) ds \right) (1-\rho\, dt) \\ & + (1-\rho\, dt)\left(1-(u+v+wG(x)+zG(y))\, dt\right) V(x,y) \} \end{aligned} \tag{19}$$

The first-order conditions from (19) allow us to express firm one's strategy in terms of the value function:

$$\begin{aligned} \alpha u(x,y) &= \int_x^0 V(s,y) f(s,x) ds + PF(x) - V(x,y) \\ \beta w(x,y) &= (P - V(x,y))\, G(x) \end{aligned} \tag{20}$$

This content downloaded from 128.135.100.108 on Fri, 11 Oct 2013 15:51:58 PM
All use subject to JSTOR Terms and Conditions

By symmetry, the strategies of firm two are

$$\alpha v(x,y) = \int_y^0 V(s,x)f(s,y)ds + PF(y) - V(y,x) \quad (21)$$
$$\beta z(x,y) = (P - V(y,x))\, G(y)$$

Equilibrium is characterized by substituting the equations for strategies into the Bellman equation, which then reduces to

$$\begin{aligned} 0 = {} & \frac{1}{2\alpha}\left(\int_x^0 V(s,y)f(s,x)ds + PF(x) - V(x,y)\right)^2 \\ & + \frac{1}{2\beta}\left(P - V(x,y)\right)^2 G(x)^2 \\ & + \frac{1}{\alpha}\left(\int_y^0 V(s,x)f(s,y)ds + PF(y) - V(y,x)\right) \\ & \times \left(\int_y^0 V(x,s)f(s,y)ds - V(x,y)\right) \\ & - \left(\rho + \frac{(P - V(y,x))G(y)^2}{\beta}\right) V(x,y) \end{aligned} \quad (22)$$

Theorem 5. *There exists a $\bar{P} > 0$ such that for $P \in [0, \bar{P}]$, there is a unique symmetric subgame perfect equilibrium $V(x,y)$, which is analytic in P, α, β, and ρ, and represented as a solution to (22).*

Proof. Same as Theorem 3. □

Theorem 5 is a strong result, but one which fits the focus on equilibria which depend on only the current state. Implicitly, we are ruling out equilibria where current actions depend on past history. This eliminates some reputation effects, trigger strategies, and other phenomena which can support implicit collusion in such infinite-horizon dynamic games. This is reasonable in the case of leap investment since such investments are unobserved and any cheating could be inferred only when a leap occurred, which would be too late. Some implicit collusion in partial jump investment may be possible since, as long as neither had won, each could infer cheating if the other seemed to be moving too quickly. The uniqueness result in Theorem 5 does not rule out the existence of reputational equilibria since it just says that there is a unique function $V(x,y)$ that is analytic in (x,y) and solves the equilibrium equations for small P.

Suppose $V(x,y) = P(pg^1(x,y) + p^2g^2(x,y) + \ldots)$ is a Taylor series approximation of $V(x,y)$ for small p. By Theorem 5, such a representation exists and is unique for small p. Even though (22) is not expressed in p, it can be straightforwardly rewritten so that V/P, the dimensionless value of the game, depends on P, α, β, and ρ only through p and the dimensionless ratio $\gamma = \alpha/\beta$. By substituting

This content downloaded from 128.135.100.108 on Fri, 11 Oct 2013 15:51:58 PM
All use subject to JSTOR Terms and Conditions

this representation for V in (22) and equating coefficients of like powers, we find

$$\begin{aligned} g^1(x,y) &= F(x)^2 + \gamma G(x)^2 \\ g^2(x,y) &= 2F(x)\left(\int_x^0 (F(s)^2+\gamma G(s)^2) f(s,x)ds - F(x)^2 - \gamma G(x)^2\right) \\ &\quad -2\left(\gamma G(x)^2 + \gamma G(y)^2 + F(y)^2\right)\left(F(x)^2+\gamma G(x)^2\right) \end{aligned} \tag{23}$$

The equilibrium strategies are therefore approximated to $O(p^3)$ by

$$\begin{aligned} \frac{u(x,y)}{\rho} &\approx 2pF(x) + 2p^2\left(\int_x^0 g^1(s,y)f(s,x)ds - g^1(x,y)\right) \\ &\quad +2p^3\left(\int_x^0 g^2(s,y)f(s,x)ds - g^2(x,y))\right) \\ \frac{w(x,y)}{\rho} &\approx 2\gamma p\left(1 - pg^1(x,y) - p^2 g^2(x,y)\right)G(x) \end{aligned} \tag{24}$$

and similarly for $v(x,y)$ and $z(x,y)$. This solution and its approximation now allow us to compare equilibrium with the social optimum and evaluate the competitive equilibrium allocation of resources.

7 Comparisons of the optimal and equilibrium outcomes

We next will compare the levels of innovative activity under social control with those levels in the game equilibrium. If $P = B$, the difference between innovative effort under competition, u^c, w^c, and the socially optimal levels, u^s, w^s, is expressed, up to $O(p^2)$, by

$$\rho^{-1}(u^s - u^c) \approx -2p^2\left(F(y)^2 + \gamma G(y)^2\right)F(x) \tag{25}$$

$$\rho^{-1}(w^s - w^c) \approx -2\gamma p^2\left(F(y)^2 + \gamma G(y)^2 + \gamma G(y)^2\right)G(x) \tag{26}$$

The difference between firm two's choices, v^c and z^c, and the optimal controls v^s and z^s, are similarly expressed. First note that there is excessive investment in all projects under competition, a conclusion common in these models. The excess is greater as either firm is closer to success. Also the excess investment relative to the socially optimal investment increases for each firm as the other firm is closer to success. These results are expected since each firm ignores the social value of the other's presence in the innovation process.

We also note that it is not clear which firm is more excessive in R&D investment. If E_{uv} is the difference, $(u^c - u^s) - (v^c - v^s)$, between the two competitor's excessive investment in their partial jump processes, then

$$E_{uv}/\left(2\rho p^2\right) \approx F(y)F(x)\left(F(y) - F(x)\right) + \gamma\left(G(y)^2 F(x) - G(x)^2 F(y)\right)$$

to $O(p^2)$. If there are no leaps, then $G \equiv 0$ and $E_{uw} < 0$ if $x > y$, that is, the laggard's investment is more excessive than the leader's. This holds also if the leap and partial jump processes are sufficiently similar, in particular if $G = \lambda F$ for some

This content downloaded from 128.135.100.108 on Fri, 11 Oct 2013 15:51:58 PM
All use subject to JSTOR Terms and Conditions

$\lambda > 0$. However, if $F(y)$ is small but $G(y)$ is not, then $E_{uw} > 0$, and the leader invests more excessively in partial jumps.

In relative terms, however, we can be more precise since

$$\frac{u^c - u^s}{u^s} \approx p\left(F(y)^2 + \gamma G(y)^2\right) \tag{27}$$

is increasing in y. $(w^c - w^s)/w^s$ is similarly found to be increasing in y. The dependence of $v^c - v^x$ and $z^c - z^s$ on x are symmetrically expressed. Therefore, the laggard's excess investment in both partial jumps and leaps expressed as a fraction of the socially optimal investment is greater. Theorem 6 summarizes these comparisons.

Theorem 6. *If B is small and $P = B$ then*

$$\frac{u^c - u^s}{u^s} > \frac{v^c - v^s}{v^s} \text{ and } \frac{w^c - w^s}{w^s} > \frac{s^c - z^s}{z^s}$$

if and only if $x < y$.

These comparisons do not necessarily say anything about the efficiency of resource allocation given that there is competition. For example, in deciding whether to subsidize the current leader a social planner should consider its impact on the future nature of the distorted allocation of resources due to the competition. We next address this issue for the case $P = B$.

If $P = B$, the social value of the game is $V(y, x) + V(x, y)$ since all benefits of innovation are appropriated by the firms. At any position, the net social marginal values, $NSMV_u$ and $NSMV_v$, of u and w per dollar of expenditure equal the ratio of the net contribution to the social value and the marginal cost:

$$\begin{aligned} NMSV_u &= \frac{\int_x^0 V(y,s)f(s,x)ds - V(y,x)}{\int_x^0 V(s,y)f(s,x)ds + PF(x) - V(x,y)} \\ NSMV_w &= -\frac{V(y,x)}{P - V(x,y)} \end{aligned} \tag{28}$$

where we use (20) to simplify expressions. Using our expansion for $V(x, y)$, (28) implies that, as p converges to 0,

$$\begin{aligned} p^{-1}NMSV_u &\approx -g^1(y,x) = -F(y)^2 - \gamma G(y)^2 \\ p^{-1}NMSV_w &\approx -F(y)^2 \end{aligned} \tag{29}$$

Symmetric expressions for $NMSV_v$ and $NMSV_z$ hold. If $x > y$ then $F(x) > F(y)$ and $G(x) > G(y)$, implying that $NMSV_z < NMSV_w$, and $NMSV_v < NMSV_u$. Therefore, the social value of more investment in either project is greater at the leading firm, even when we consider the distortions implicit in the competition.

Theorem 7. *If $P = B$ and P is small, social welfare at any stage would be increased by shifting innovation effort from the laggard to the leader. That is, if (x, y) is the current state and $x > y$, $V(x, y) + V(y, x)$ is increased if $u(x, y)$ is increased and $w(x, y)$ is decreased, and similarly for $z(x, y)$ and $w(x, y)$.*

This content downloaded from 128.135.100.108 on Fri, 11 Oct 2013 15:51:58 PM
All use subject to JSTOR Terms and Conditions

Theorem 7 shows that any small temporary subsidy/tax scheme which reallocates effort towards the leader is socially desirable since combinations of subsidies and taxes can induce such a switch and the objective of $V(x,y) + V(y,x)$ ignore any redistributive component of such a policy. Therefore, in this limited sense, policy should favor the current leader over the laggard.

Another interesting issue which we can address in this model is that of the efficiency of the allocation of resources between the risky leaps and the less risky partial jumps. The social efficiency of the portfolio choice by firm one is determined by comparing the net social marginal values of u and v. $NSMV_u > NSMV_w$ iff $g^1(x,y) - \int_x^0 g^1(s,y)f(s,x)ds < F(x)g^1(x,y)$ which is true since $g^1(x,y)$ is increasing in x. Hence, there is an excessive share of resources allocated to the "risky" project. To get an intuitive grasp on this result, we should compare the social valuation of the intermediate stages with the equilibrium valuation by firm one. Since the difference between g^1 and h^1 is independent of x, we need to compare g^2 with h^2 to study differences relevant for one's portfolio choice between u and w. Straightforward manipulation of the expansions for V and W shows that, ignoring terms which are of $o(P^3)$,

$$V(x,y)-W(x,y) \approx 2p^2 \left(F(x)^2 + \gamma G(x)^2\right) \left(F(y)^2 + \gamma G(y)^2\right) P+Z(y) \quad (30)$$

where $Z(y)$ depends only on y. Therefore, $V - W$ is increasing in x for small p. First, this implies that investment is even more excessive than indicated by p^2 terms since the gap between social and private values of R&D is increasing at $O(P^3)$. Second, it indicates a bias towards risky R&D projects. Since this excess increases in x, those projects which are more likely to yield big jumps, holding the expected jump constant, will find their private value to be more excessive relative to their social value.

Theorem 8. *If $P = B$ and P is small, social welfare would be increased if resources were shifted from the risky R&D projects to the less risky projects.*

The last comparison we will make is between the optimal and equilibrium reactions of firms to each other's partial successes. Before using our approximations, note that our expression for firm one's equilibrium choice of w in equation (20), differs substantially from the expression for the social choice in equation (15), despite their formal similarity. In (15), it is clear that the optimal choice of w falls if the social value of the social position (x,y) increases but x, the position of firm one, remains unchanged. In particular, an advance in firm two's position will increase the social value, and hence lead to a reduction of expenditure at firm one on the leap investments. In (20), we find that expenditure on w will rise as the value of the game to firm one falls, which is the expected response to an advance by firm two. Hence, if the social and private value functions vary with position in the intuitive fashion, firm one will increase leap investments in response to an advance by firm two, even though the socially optimal response would be a reduction in effort.

Proving these conjectures globally would be quite difficult given the nonlinear nature of the expression for the equilibrium value functions. However, our approximations will immediately confirm them. Since $g^2(x,y)$ is independent of y, the

This content downloaded from 128.135.100.108 on Fri, 11 Oct 2013 15:51:58 PM
All use subject to JSTOR Terms and Conditions

dependence of u and w on y for small P, is determined by the dependence of g^3 on y, and is summarized in

$$\begin{aligned}\rho^{-1}u^c &= \ldots + 2p^3\left(F(y)^2+\gamma G(y)^2\right) \\ &\qquad \times\left(\int_x^0 \left(F(s)^2+\gamma G(s)^2\right)F(s,x)ds - F(x)^2-\gamma G(x)^2\right) \\ \rho^{-1}w^c &= \ldots + 2\gamma p^3\left(F(y)^2+\gamma G(y)^2\right)\left(F(x)^2+\gamma G(x)^2\right)G(x)\end{aligned} \tag{31}$$

where we have displayed all terms of $O(p^3)$ that depend only on y.

Theorem 9. *If $P = B$ and P is small,*

$$0 < \left|\frac{\partial u^c}{\partial y}\right| < -\frac{\partial u^s}{\partial y}$$

$$\frac{\partial w^c}{\partial y} > 0 > \frac{\partial w^s}{\partial y},$$

that is, firm one's equilibrium reactions are less than the optimal reactions in magnitude. Furthermore, $\partial u^c/\partial y$ is always positive and $\partial w^c/\partial y$ is of ambiguous sign. Symmetric results for firm two hold.

Proof. The comparisons of magnitude follow from the fact that $\partial u^c/\partial y$ is $O(p^3)$ by (31) but $\partial u^s/\partial y$ and $\partial w^s/\partial y$ are $O(p^2)$ by (18). The sign conditions for w^c and z^c follow from (31). If $F(s)$ and $G(s)$ are large relative to $F(x)$ and $G(x)$ for $s > x$, then the integral in (31) dominates and $\partial u^c/\partial y > 0$. However, if $F(s) \approx F(x)$ and $G(x) \approx G(x)$ for $s > x$, then $\int_x^0 (F(s)^2+\gamma G(x)^2) f(s,x)ds \approx (F(x)^2+\gamma G(x)^2)\,(1-F(x))$ and $\partial u^c/\partial y < 0$ in (31). □

In comparing the dependence of strategies on the positions of the firms, first note that there is no reaction of one firm to another's position to $O(p^2)$. Hence, the equilibrium reactions of the firms to each are smaller than the optimal reactions. Furthermore the direction may be wrong. In the case of leap investment, the reaction will always be in the wrong direction. This is intuitively seen from (20): we expect that as firm two advances, the value of the game to firm one, $V(x, y)$ decreases, thereby raising firm one's choice of w. In the social control case, the value increases as firm two advances, reducing the social choice for w.

However, the reaction of u is ambiguous. The reaction of a partial jump's control to the other firm's movement depends on just how different the stages are. If the stages are similar in that the probability of winning immediately per unit of effort with a leap, $G(x)$, or partial jump, $F(x)$, is nearly as large at x as at any later stage, then u will fall. On the other hand, if later stages have substantially greater likelihoods of getting one to success, then a firm's effort in partial jumps may increase as its opponent moves ahead. In the latter case, the improvement in the opponent's prospects prompts one to work harder, as if one must either work hard or concede the race. Also note that if a mean preserving spread in the probability weights $f(s, x)$ will increase the likelihood of a perverse reaction for u since the integral in (31) has a convex integrand.

This content downloaded from 128.135.100.108 on Fri, 11 Oct 2013 15:51:58 PM
All use subject to JSTOR Terms and Conditions

This result about the reaction of u differs substantially from previous papers. Fudenberg et al. (1983) and Harris and Vickers (1985) make very specific assumptions about the R&D process and arrive at more definitive results. They made special assumptions since they were necessary to arrive at conclusions given their approach. The perturbation approach used here can handle models that are much more general in many dimensions. Of course, we assume that the prize is small, a loss of generality. The perturbation method does not uniformly dominate alternative analytical approaches, but the results here show that it can investigate new territory.

At this point we should expand on the appropriate interpretation of our juggling of these various orders of magnitude. For example, the fact that the reaction of u^c to y is zero at $O(p^2)$ and possibly nonzero only at $O(p^3)$ does not imply that reactions are generally unimportant and uninteresting when compared to the effects which show up at $O(p^2)$. In fact, in many games where reactions are generally important we would find that, as the payoffs go to zero, the reactions go to zero faster than other elements of equilibrium strategies. Only for nearly degenerate games does the order reflect the relative importance of various effects. Since our objective is to gain more general insight, we make no comparisons. On the other hand, one cannot infer that an $O(p^3)$ effect will eventually dominate any $O(p^2)$ effect since other, even higher, orders also contribute. Our goal in these calculations is to sign various effects and determine the critical structural elements for an open set of games, hoping to elicit general qualitative insights about the nature of the subgame equilibria. Arguments which mix various orders of magnitudes are either illegitimate or focus too tightly on the small p nature of the analysis.

8 Implications for social innovation policy

We next examine the optimal values of two parameters of social innovation policy, the portion of social benefits to be awarded to the winner and the stage at which a patent is to be granted. We will find that when B is small, the difference between the optimal P and B is negligible relative to B. This result validates our focus on the case $P = B$ in the previous section since it implies that all those results continue to hold for an optimally chosen P. In particular, this shows that the misallocation of resources between projects of varying riskiness will not change with an optimally chosen P. While these results are not surprising, it is instructive to show how to rigorously demonstrate them within our approach.

Let $P = \theta B$, i.e., θ is the portion of social benefits of innovation which the innovator is allowed to appropriate. We are making the simplifying assumption that this allocation of social benefits to the innovator can be made in a nondistortionary fashion. In the case of patents this is only valid if demand is inelastic. If a prize is awarded, this assumes that it is financed by nondistortionary revenue sources.

Presumably, θ is a parameter at least partially chosen by policy markers. Given that we found that there was excessive allocation of resources for innovation in the equilibrium of the innovation game, the optimal θ is never unity. Let W again represent the social value function except now we make explicit the dependence on

This content downloaded from 128.135.100.108 on Fri, 11 Oct 2013 15:51:58 PM
All use subject to JSTOR Terms and Conditions

θ and B. Then

$$W(x,y,\theta,B) = -\left[\alpha(u^2+v^2)+\beta(w^2+z^2)\right]\frac{1}{2}dt \qquad (32)$$
$$+(1-\rho dt)\ (uF(x)+vF(y)+wG(x)+zG(y))\,Bdt$$
$$+(1-\rho dt)\ (1-(u+v+wG(x)+zG(y))\,dt)\ W(x,y)$$
$$+(1-\rho dt)\left(u\int_x^0 W(z,y)f(z,x)+v\int_y^0 W(x,s)f(s,y)ds\right)dt$$

where u, v, and z are the equilibrium policy functions if the prize is θB.

We can use the characterization in (32) to generate some information about the optimal θ, $\theta^*(B)$, when B is small. $\theta^*(B)$ is defined by the equation $W_\theta(x,y,\theta^*(B),B) = 0$ This is not a completely trivial calculation since any θ is optimal when $B=0$. Therefore we compute $\theta^*(0^+)$, the limit of θ^* as B falls to zero.

First, $\theta^*(B)$ is implicitly defined by $W_\theta(x,y,\theta^*(B),B)=0$. For sufficiently small B, $\theta^*(B)$ is continuously differentiable by the Implicit Function Theorem applied to the equation $W_\theta(x,y,\theta^*(B),B) = 0$, since direct calculation shows that $W_{\theta\theta}$ is not zero and $W_{\theta B}$ exists for B close to zero. Since $\theta^*(B)$ is optimal for the initial position (x,y),

$$\lim_{B\to 0^+}\frac{W(x,y,\theta^*(0^+),B)-W(x,y,\theta,B)}{B^2}\geq 0 \qquad (33)$$

for all θ. Since $W(x,y,\theta,B)$ and $W_B(x,y,\theta,B)$ both converge to 0 as B converges to 0, by l'Hospital's rule (33) equals

$$\lim_{B\to 0^+}\frac{W_{BB}(x,y,\theta^*(0^+),B)-W_{BB}(x,y,\theta,B)}{2}$$

for all θ. Therefore, $W_{BB}(x,y,\theta^*(0^+),0)-W_{BB}(x,y,\theta,0)\geq 0$ for all θ, implying that $\theta^*(0^+)\in \arg\max_\theta W_{BB}(x,y,\theta,0)$ and

$$W_{BB\theta}(x,y,\theta^*(0^+),0)=0 \qquad (34)$$

Direct substitution of the asymptotic equilibrium strategies into (34) shows that

$$\alpha\rho W_{BB}(x,y) = 4\left(\theta-\theta^2/2\right)\left(F(x)^2+F(y)^2+\gamma\left(G(x)^2+G(y)^2\right)\right) \qquad (35)$$

which is maximized at $\theta^*(0^+)=1$. Therefore, when the prize is small, the social surplus maximizing policy gives nearly all of the social benefits to the innovator.

Note that this does not contradict our earlier result that innovation is excessive whenever the prize equals the benefit, just that the difference between the optimal prize and the social benefit goes to zero faster than the social benefit. This is not surprising since it just says that the externalities due to the competition over the rents fall more rapidly than B as B goes to zero. The primary point of this exercise is to illustrate how to determine the limit.

Second, further expansion of the social value function and application of l'Hospital's rule shows that the optimal θ falls more rapidly as B increases when

This content downloaded from 128.135.100.108 on Fri, 11 Oct 2013 15:51:58 PM
All use subject to JSTOR Terms and Conditions

$F(x)$ and $G(x)$, the probability of an immediate success from the current position (assuming the firms begin at the same position), rises. This implies that the shorter the race, the smaller should be the winner's share under competition. Since the details entail only repeated applications of the foregoing calculations, they are omitted here.

Another crucial aspect of patent policy is the stage at which a patent is granted. A patent may be granted before final and complete success is achieved. In fact, in the existing patent system, a patent is granted when a description of an invention has been completed, before the development stages leading to a workable and commercial prototype have been achieved. This may be socially optimal if the effort of followers is so excessive and wasteful that it is better to force them out of the race, bearing the possible inefficiencies that may result when an innovator is given the monopoly early. In our model, this can be modeled by assuming that a patent is granted to the first firm which crosses $c \leq 0$. If $c = 0$, the firm must complete the project before acquiring a patent worth P. If $c < 0$, then a firm receives a patent at c and may finish development without any competition.

Proceeding as in the $c = 0$ case, we find that the equilibrium value function for the firms satisfies

$$\begin{aligned} 0 = \frac{1}{2\alpha}\left(\int_x^c V(s,y)f(s,x)ds + \int_c^0 M(s)f(s,x)ds + PF(x) - V(x,y)\right)^2 \quad (36) \\ + \left(\int_y^c V(s,x)f(s,y)ds + \int_c^0 M(s)f(s,y)ds + PF(y) - V(y,x)\right) \\ \times \left(\int_y^c V(x,s)f(s,y)ds - V(x,y)\right)\frac{1}{\alpha} - \rho V(x,y) \end{aligned}$$

where $M(\cdot)$ is the monopoly value function computed in Section 4 with the extension to two instruments, u and v or w and z. We expand (36) as before for the case of a small social benefit and prize. We find that when B is small and $P = B$, the value of $V(x,y) + V(y,x)$, the social surplus value function, for $c < 0$ minus the value when $c = 0$ is approximated by

$$-F(y)\int_c^0 g^1(x,s)f(s,y)ds - F(x)\int_c^0 g^1(y,s)f(s,x)ds < 0 \quad (37)$$

Hence, the major factor is that if $c < 0$, the contest is ended early and the resulting loss in total effort is excessive relative to the cost savings. Note that this strict inequality depends critically on our standing assumption that $F(x) > 0$ for all x.

Theorem 10 summarizes our findings concerning optimal policy.

Theorem 10. *When B is small, the optimal policy is to award a prize only when the race is completely won and the prize should be nearly the entire social value of the innovation. Furthermore, the closer the innovators are to final success when competition begins, the less should be their share in the social benefit.*

While these conclusions are surely not globally true, we have shown their validity for an open set of problems and that contrary conjectures cannot be globally

This content downloaded from 128.135.100.108 on Fri, 11 Oct 2013 15:51:58 PM
All use subject to JSTOR Terms and Conditions

true. More important, we have shown how to address these questions for that set of problems. Other exercises, such as the impact of suboptimal innovation resource allocation on the optimal prize, can be conducted by straightforward examination of the higher-order terms of our expansion for W, the social planner's objective. In the interest of space, we leave such extensions to the reader.

9 Generalizations

There are many other exercises which could be used to demonstrate the applicability of perturbation methods. We examined one that most clearly illustrates the general approach advanced here. To indicate that the perturbation approach is not too specialized, we will now discuss some other possible applications.

All models with closed-form solutions are degenerate in some sense. When we use them we hope that the features that these tractable models ignore are not important. Perturbation methods can be used to test this presumption. Take, for example, the model used by Loury (1979) and Reinganum (1981). While it yields closed-form solutions for the quadratic cost specification, it abstracts from the possibility of intermediate stages, our focus here. Recall that our model with F and G equal to constant functions is exactly that model. To examine the importance of intermediate stages on the nature of equilibrium, we could have assumed that F and G deviated slightly from constant functions. This alternative would have allowed us to determine the nature of equilibrium for arbitrary prize but with only a small deviation from the implicit stage-independence assumption sometimes used.

Another possible generalization is allowing intermediate payments. The perturbation analysis conducted above could also allow intermediate payoffs since nothing we did used the absence of intermediate payoffs in an essential fashion; we focussed on the more simple payoff structure since our purpose was to present a robust analysis of the positional dynamics among competitors for one kind of race. A more general analysis with intermediate payoffs could generate insights, for example, into strategic implications of the learning curve; one approach would be to approximate the slow-learning case by knowing the solution to the no-learning case. However, we leave such an analysis to another study.

While this is certainly not an exhaustive list, it shows that perturbation methods are useful in examining the robustness of simple models generally, allowing us to add otherwise intractable elements to the analysis of a problem. While our analysis got started by examining the trivial case of no payoff, generally one can begin with any tractable case, making perturbation analysis a generally valuable tool for dynamic strategic analysis.

10 Conclusion

We have analyzed a simple closed-loop subgame perfect model of multi-stage innovation. We found the usual result of excessive innovative effort when the prize equals the social value. Under the assumption that the net social value of innovation is small, we have also found that there will be excessive risk-taking, that at any

This content downloaded from 128.135.100.108 on Fri, 11 Oct 2013 15:51:58 PM
All use subject to JSTOR Terms and Conditions

moment the following firm is a less efficient innovator relative to the leader, that the prize to the innovator should nearly equal social benefits, and that the competition should not be ended before one of the competitors has succeeded completely. While these results have obvious limitations on their generality, they do tell us that the contrary propositions cannot be generally true. While many of the results, e.g., the excessive investment when the prize equals the social benefit, follow naturally from the fact that these subgame perfect equilibria are close to some open-loop equilibria others, in particular the computation of the equilibrium reactions, are specific to the subgame-perfect solution. They have therefore given us a peek into the nature of subgame perfect equilibrium in such innovation models.

References

Bhattacharya, S., Mookherjee, D.: Portfolio choice in research and development. Rand Journal of Economics **17**, 594–605 (1986)

Budd, C., Harris, C., Vickers, J.: A model of the evolution of duopoly: Does the asymmetry between firms tend to increase or decrease? Review of Economic Studies **60**, 543–573 (1993)

Bryson, A. E., Ho, Y.: Applied optimal control. New York: Hemisphere Publishing 1975

Dasgupta, P., Stiglitz, J.: Uncertainty, market structure and the speed of research. Bell Journal of Economics **11**, 1–28 (1980a)

Dasgupta, P., Stiglitz, J.: Industrial structure and the nature of innovative activity. Economics Journal **90**, 266–293 (1980b)

Doraszelski, U.: An R&D race with knowledge accumulation. Hoover Institution, mimeo (2002)

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

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

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

Judd, K. L.: Closed-loop equilibrium in a multi-stage innovation race. Discussion Paper No. 647, Kellogg Graduate School of Management, Northwestern University (1985)

Judd, K. L.: Projection methods for solving aggregate growth models. Journal of Economic Theory **58**, 410–452 (1992)

Judd, K. L.: Computational economics and economic theory: substitutes or complements? Journal of Economic Dynamics and Control **21**, 907–942 (1997)

Kamien, M., Schwartz, N.: Market structure and innovation. Cambridge: Cambridge University Press 1982

Lee, T.: On a fundamental property of research and development rivalry. University of California, San Diego, mimeo (1982)

Lee, T.,Wilde, L.L.: Market structure and innovation: a reformulation. Quarterly Journal of Economics **94**, 429–436 (1980)

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

Reinganum, J.: Dynamic games of innovation. Journal of Economic Theory **25**, 21–41 (1981)

Reinganum, J. F. A dynamic game of R&D: patent protection and competitive behavior. Econometrica **50**, 671–688 (1982)

Selten, R.: Reexamination of the perfectness concept for equilibrium points in extensive games. International Journal of Game Theory **4**, 25–55 (1974)

Telser, L.: A theory of innovation and its effects. Bell Journal of Economics **13**(1), 69–92 (1982)

Zeidler, E.: Nonlinear functional analysis and its applications. Part I: Fixed point theorems. New York: Springer 1986

This content downloaded from 128.135.100.108 on Fri, 11 Oct 2013 15:51:58 PM
All use subject to JSTOR Terms and Conditions