



University of Pennsylvania
ScholarlyCommons

Business Economics and Public Policy Papers

Wharton Faculty Research

1-2003

An R&D Race with Knowledge Accumulation

Ulrich Doraszelski
University of Pennsylvania

Follow this and additional works at: https://repository.upenn.edu/bepp_papers



Part of the [Economics Commons](#)

Recommended Citation

Doraszelski, U. (2003). An R&D Race with Knowledge Accumulation. *The RAND Journal of Economics*, 34 (1), 20-42. <http://dx.doi.org/10.2307/3087441>

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/bepp_papers/94
For more information, please contact repository@pobox.upenn.edu.

An R&D Race with Knowledge Accumulation

Abstract

I develop a model of an R&D race with knowledge accumulation. My model does not inherit the memorylessness property of the exponential distribution that troubles existing models of R&D races. Hence, firms' knowledge stocks are no longer irrelevant to their behavior during the R&D race, and knowledge accumulation has strategic implications. In this more general setting, I obtain results that stand in marked contrast to the previous literature. In particular, under some conditions, the firm that is behind in the race engages in catch-up behavior. This pattern of strategic interactions (action-reaction) is consistent with empirical research.

Disciplines

Economics

An R&D Race with Knowledge Accumulation*

Ulrich Doraszelski[†]
Hoover Institution

March 2002

Abstract

I develop a model of an R&D race with knowledge accumulation. The knowledge that a firm has acquired as a by-product of its investment in R&D is valuable even if success is not immediate because it helps the firm to make the discovery later on. My model does not inherit the memorylessness property of the exponential distribution that troubles existing models of R&D races. Hence, firms' knowledge stocks are no longer irrelevant to their behavior during the R&D race, and knowledge accumulation has strategic implications. Unlike multi-stage race models, knowledge is productive in my model. In this more general setting, I obtain results that stand in marked contrast to the previous literature. In particular, under some conditions, the firm who is behind in the race engages in catch-up behavior. The resulting pattern of strategic interactions (action-reaction) is consistent with empirical research on R&D races.

The model does not in general allow for an analytical solution, and I show how numerical techniques (projection methods) can be used instead. An accuracy check indicates that the approximations yield a good description of the equilibrium payoffs and strategies.

*An earlier version of this paper was circulated under the title "An R&D Race with Learning and Forgetting". I would like to thank David Besanko, Michaela Draganska, Dino Gerardi, Heidrun Hoppe, Johannes Hörner, Ken Judd, Rosa Matzkin, Rob Porter, Mark Satterthwaite, Katja Seim, Mihkel Tombak, and Sevin Yeltekin for helpful discussions. Furthermore, I am extremely grateful to the editor Kyle Bagwell and two anonymous reviewers for their constructive suggestions. This paper has benefited from meetings of the Computational Group at Northwestern University.

[†]Hoover Institution, Stanford University, Stanford, CA 94305-6010, U.S.A., doraszelski@hoover.stanford.edu

1 Introduction

In an R&D race, firms compete to be the first to make a discovery by investing in R&D. The firm which makes the discovery first is awarded a prize, often in the form of a patent, whereas the remaining firms either receive nothing at all or a smaller prize. As a by-product of its investment in R&D, a firm accumulates knowledge. This raises the question whether and how this knowledge accumulation affects firms' behavior. In particular, what are the strategic implications of knowledge accumulation in an R&D race?

Suppose, for example, that the competing firms have unequal knowledge stocks. Does the lagging firm succumb to the leading firm or does it step up its investment in R&D in an attempt to make up for the leader's advantage? Casual observation suggests that the laggard strives to catch up with the leader. When Transmeta unveiled its power-stingy Intel-compatible Crusoe chip in 2000, Intel pledged to introduce a version of its Pentium III processor that matches Crusoe's power consumption in the first half of 2001 and announced a new set of technologies for 2002 or 2003 that would give it the lead over Transmeta.¹ Similarly, after Celera Genomics in 1998 challenged the Human Genome Project to be the first to sequence the human genome, the Human Genome Project announced that it will move up its target date from 2005 to 2003 and indeed dramatically stepped up its own pace during 1999. And yet, although Celera Genomics started the race as the underdog, it completed a draft of the human genome in 2000 and beat the Human Genome Project.² The existing models of R&D races cannot explain this pattern of strategic interactions. In fact, in so-called memoryless races, knowledge accumulation has no strategic implications, whereas the laggard gives way to the leader in multi-stage races.³

In memoryless R&D race models, the knowledge that firms have acquired as a result of their past R&D efforts is irrelevant to their current R&D efforts (Loury 1979, Lee & Wilde 1980, Reinganum 1981, Reinganum 1982). This is driven by the assumption that the time of a successful innovation is exponentially distributed. Because of the *memorylessness property* of the exponential distribution, memoryless R&D race models by design cannot capture history dependence. In other words, knowledge accumulation has no strategic implications

because there is no value to knowledge.

Multi-stage race models attempt to circumvent the memorylessness property in order to allow for history dependence. These models account for the possibility that one firm may be ahead of another by introducing intermediate steps into the research process. Thus, to win the race, a firm must be the first to complete all stages of the R&D project. Deterministic multi-stage race models assume that firms transit from one stage to the next in a deterministic fashion (Fudenberg, Gilbert, Stiglitz & Tirole 1983, Harris & Vickers 1985a, Harris & Vickers 1985b, Lippman & McCardle 1988). The outcome is ϵ -preemption: The slightest advantage of one firm causes the other to immediately drop out of the race. This strong result is moderated somewhat when the stage-to-stage transitions are assumed to be probabilistic. Grossman & Shapiro (1987), Harris & Vickers (1987), and Lippman & McCardle (1987) adopt the stochastic structure of the memoryless R&D race models and assume that the time to completion of each stage is exponentially distributed. Consequently, while a firm's investment in R&D depends on the number of stages it and its rival have left to complete, within each stage the memorylessness property renders firms' current R&D efforts independent of their past R&D efforts. In these models the follower devotes *less* resources to R&D than the leader. The follower therefore tends to fall further behind as the race progresses whereas the leader tends to build up its advantage. This leads to *increasing dominance*.⁴

ϵ -preemption and increasing dominance imply that the R&D race is effectively decided once a firm falls behind. Thus, the R&D race consists of a short but intense battle at the outset of game, followed by a phase during which the winner of the battle completes the R&D project essentially unimpaired by competitive pressures. In contrast, casual observation as well as empirical evidence (Grabowski & Baxter 1973, Scherer 1992, Khanna 1995, Lerner 1997) indicates that, at least in some cases, the firm that is behind engages in catch-up behavior, thereby leading to drawn-out battles between firms. Hence, the evidence suggests a pattern of strategic interaction that is more like *action-reaction* than increasing dominance.

I develop a general model of an R&D race with knowledge accumulation that is consistent with the evidence. In particular, I provide conditions under which the follower invests *more*

in R&D than the leader, so that a pattern of action-reaction results. As the game can no longer be solved analytically, I apply projection methods to solve the partial differential equation that characterizes a firm's value function. Special considerations arise since I need not only a good approximation of the value function but also good approximations of its partial derivatives to compute the Nash equilibrium in feedback strategies.

The memoryless R&D race models are a special case of my model. I show that, apart from this special case, knowledge accumulation shapes firms' equilibrium payoffs and strategies, and therefore has strategic implications. This demonstrates how restrictive the memoryless R&D race models are. The main findings are as follows. First, I show that a firm is inclined to scale back its investment in R&D as its knowledge stock increases. Second, as the race progresses, the follower eventually works harder than the leader. Underlying these results is what I call the *pure knowledge effect*. The source of this effect is that a firm's past R&D efforts contribute to its chances of winning the R&D race. In other words, *knowledge is productive* in my model and therefore valuable to firms. On the other hand, it is also possible that the follower works less than the leader in the early stages of the race provided that there are increasing returns to knowledge accumulation. Third, a firm can either respond aggressively or submissively to an increase in its rival's knowledge stock. I show that a firm responds aggressively if it has a sufficiently large knowledge stock and submissively otherwise. Simulations of the evolution of the R&D race suggest that these strategic considerations are dominated by the pure knowledge effect. Fourth, competition (as measured by the sum of firms' R&D expenditures) is not necessarily fiercest when firms are neck-and-neck. This again contrasts with multi-stage race models. In multi-stage race models, competition is most intense in these situations because the race is effectively decided once a firm pulls ahead of its rival.

The remainder of this paper is organized as follows. Section 2 develops the model. Section 3 outlines the computational strategy. Section 4.1 analyzes the impact of knowledge accumulation on the equilibrium payoffs and strategies. I also describe how the equilibrium changes with the value of the patent and the degree of patent protection. Section 4.2

discusses how the race unfolds over time. While I study a small number of examples in Sections 4.1 and 4.2, I show in Section 4.3 that the economic intuition underlying these examples generalizes. Section 5 concludes.

2 Model

Consider an R&D race in which two firms compete to be the first to make a discovery. As a firm invests in R&D, its chances to immediately make the discovery increase and, in addition, the firm adds to its knowledge stock. On the other hand, the firm's knowledge stock may depreciate over time. Its knowledge stock is a measure of the firm's past R&D efforts, and is valuable to the extent that, even if success is not immediate, it helps the firm to make the discovery later on. Knowledge accumulation thus gives rise to history dependence. Incorporating knowledge accumulation into the model is important because it allows me to capture phenomena like learning and organizational forgetting. As I explain below in greater detail, history dependence also arises if R&D is done through experimentation.

For simplicity, I assume that firms may differ in the knowledge stocks they possess at the outset of the R&D race, but are identical in every other respect. Time is continuous and the horizon is infinite.

Knowledge Accumulation. Let $z_i(t)$ denote firm i 's accumulated knowledge and $u_i(t)$ its rate of knowledge acquisition at time t . $z_i(t)$ is a measure of the firm's past R&D efforts and $u_i(t)$ represents its current R&D effort. To simplify notation, I write z_i and u_i instead of $z_i(t)$ and $u_i(t)$, respectively. Firm 1's accumulated knowledge evolves according to

$$\dot{z}_1 = u_1 - \delta z_1, \quad z_1(0) = z_1^0 \geq 0, \quad \delta \geq 0.$$

If $\delta > 0$, the firm's knowledge stock depreciates over time.

Distribution of Success Times. Firm 1’s hazard rate of successful innovation is given by

$$h_1 = \lambda u_1 + \gamma z_1^\psi, \quad \lambda \geq 0, \quad \gamma \geq 0, \quad \psi > 0.$$

The hazard rate represents the rate at which the discovery is made at a certain point in time given that it has not been made before. λ measures the effectiveness of current R&D effort in making the discovery and γ the effectiveness of past R&D efforts. The parameter ψ determines the marginal impact of past R&D efforts. Depending on the value of ψ , h_1 is concave ($\psi < 1$), linear ($\psi = 1$), or convex ($\psi > 1$) in z_1 .⁵

The special case of $\gamma = 0$ corresponds to the memoryless R&D race models analyzed by Reinganum (1981, 1982). These models assume an exponential distribution of success times, which implies that the hazard rate is independent of past R&D efforts. If $\gamma > 0$, the model allows for history dependence. Hence, the knowledge stocks that firms have acquired as a result of their past R&D efforts are no longer irrelevant to firms’ current R&D efforts and thus to the outcome of the race. This allows me to model learning and organizational forgetting and R&D through experimentation.

Learning and Organizational Forgetting. Many empirical studies have documented learning in a production context (“learning-by-doing”). In general, learning means that a firm’s past experiences add to its current capabilities. Learning may occur when the practices of the organization as a whole are altered in light of past experiences. It may also occur when an R&D project is cumulative in the sense that researchers need to draw on intermediate results in order to make a discovery. I capture this in my model by setting $\gamma > 0$.

More recently, it has been shown that firms not only learn but also forget (Argote, Beckman & Epple 1990, Benkard 2000). To the extent that a firm’s experience is embodied in its workers, organizational forgetting happens because of turnover and layoffs. If there is organizational forgetting, a firm’s stock of experience depreciates over time. This implies that a firm’s recent experiences are more relevant for making the discovery than its distant experiences. I allow for organizational forgetting by setting $\delta > 0$.

R&D through Experimentation. History dependence also arises if R&D is done through experimentation.⁶ To see this, suppose that R&D is conducted by growing cultures of bacteria. Because it takes time for the culture to grow, the results from a culture that is started in the present will not be available for some time. To capture this, let u_1 be the number of experiments that firm 1 starts in the present (i.e., the firm's current R&D effort) and z_1 the stock of experiments that are still in progress (i.e., the firm's past R&D efforts). If results become available at a hazard rate of δ , then the law of motion for the stock of experiments is $\dot{z}_1 = u_1 - \delta z_1$. The firm makes the discovery and thus wins the race once an experiment ends in a success. Hence, if success occurs among the results that become available at a hazard rate of $\frac{\gamma}{\delta}$, then the hazard rate of successful innovation is $h_1 = \delta \left(\frac{\gamma}{\delta}\right) z_1 = \gamma z_1$.

Value of Innovation and Imitation. The firm which makes the innovation first is awarded a patent of positive value $\bar{P} > 0$, whereas its rival receives nothing if patent protection is perfect. If patent protection is imperfect, the loser receives a positive payoff \underline{P} , where $\bar{P} > \underline{P} > 0$. \bar{P} is understood to be the expected net present value of all future revenues from marketing the innovation net of any costs the firm incurs in doing so. Similarly \underline{P} is the expected net present value of all future cash flows from imitating the discovery.⁷

Bellman Equation. Let $V^1(z_1, z_2)$ denote the value of the race to firm 1 when firm 1 has accumulated $z_1 \geq 0$ units of knowledge and firm 2 has accumulated $z_2 \geq 0$ units of knowledge. The Bellman equation which characterizes the value function under the presumption that firms behave optimally is given by

$$\begin{aligned} rV^1(z_1, z_2) = & \max_{u_1 \geq 0} h_1 (\bar{P} - V^1(z_1, z_2)) + h_2 (\underline{P} - V^1(z_1, z_2)) - c(u_1) \\ & + V_1^1(z_1, z_2)\dot{z}_1 + V_2^1(z_1, z_2)\dot{z}_2, \end{aligned} \quad (1)$$

where V_i^1 denotes the partial derivative of V^1 with respect to z_i . $r > 0$ is the interest rate and the cost incurred to acquire knowledge at rate u_1 is $c(u_1) = \frac{1}{\eta} u_1^\eta$. $\eta > 1$ measures the elasticity of the cost function. The value function $V^1(z_1, z_2)$ can be interpreted as the asset

or option value to firm 1 of participating in the race. This option is priced by requiring that the opportunity cost of holding it, $rV^1(z_1, z_2)$, equals the current cash flow, $-c(u_1)$, plus the expected capital gain or loss flow. The latter is composed of three parts, namely the capital gain from winning the race, $\bar{P} - V^1(z_1, z_2)$, times the likelihood of doing so, h_1 , the capital loss from losing the race, $\underline{P} - V^1(z_1, z_2)$, times the likelihood of doing so, h_2 , and the capital gain or loss flow attributable to changes in the knowledge stocks $V_1^1(z_1, z_2)\dot{z}_1 + V_2^1(z_1, z_2)\dot{z}_2$.

Current R&D Effort. Differentiating the right-hand side of the Bellman equation (1) yields firm 1's FOC for an interior solution. Since $\eta > 1$, the FOC is also sufficient. Hence,

$$u^{1*}(z_1, z_2) = (\lambda(\bar{P} - V^1(z_1, z_2)) + V_1^1(z_1, z_2))^{\frac{1}{\eta-1}}. \quad (2)$$

The firm has two incentives to engage in R&D. First, as the firm invests an additional dollar in R&D, its chances of making the discovery at this point in time increase by λ . Since the capital gain from winning the race is $\bar{P} - V^1(z_1, z_2)$, the marginal benefit accruing to the firm is $\lambda(\bar{P} - V^1(z_1, z_2))$. Second, the firm adds to its knowledge stock which carries a marginal benefit of $V_1^1(z_1, z_2)$.

Equilibrium. Tractability requires a restriction of the notion of a strategy when analyzing differential games. I focus on symmetric stationary Nash equilibria in feedback strategies as given by equation (2) (e.g., Basar & Olsder (1999)). A firm's strategy thus maps the accumulated knowledge of firms 1 and 2 into a rate of knowledge acquisition. The feedback strategy in equation (2) is constructed using the Bellman equation (1), and thus ensures that the firm behaves optimally at all points (z_1, z_2) of the state space, irrespective of whether these knowledge stocks are on or off the equilibrium path.

Define $V^1(z_1, z_2) = V(z_1, z_2)$ and $u^{1*}(z_1, z_2) = u^*(z_1, z_2)$. Then, using symmetry, the value of the race to firm 2 when firm 1 has accumulated z_1 units of knowledge and firm 2 has accumulated z_2 units of knowledge is given by $V^2(z_1, z_2) = V(z_2, z_1)$ and firm 2's equilibrium strategy is $u^{2*}(z_1, z_2) = u^*(z_2, z_1)$. At a symmetric Nash equilibrium in feedback strategies,

the Bellman equation (1) at $u_1 = u^*(z_1, z_2)$ and $u_2 = u^*(z_2, z_1)$ can be rewritten as the operator equation

$$\mathcal{N}(V) = 0, \quad (3)$$

where

$$\begin{aligned} \mathcal{N}(V)(z_1, z_2) = & \left(\lambda u^*(z_1, z_2) + \gamma z_1^\psi \right) \bar{P} + \left(\lambda u^*(z_2, z_1) + \gamma z_2^\psi \right) \underline{P} - \frac{1}{\eta} u^*(z_1, z_2)^\eta \\ & - \left(r + \lambda u^*(z_1, z_2) + \gamma z_1^\psi + \lambda u^*(z_2, z_1) + \gamma z_2^\psi \right) V(z_1, z_2) \\ & + V_1(z_1, z_2)(u^*(z_1, z_2) - \delta z_1) + V_2(z_1, z_2)(u^*(z_2, z_1) - \delta z_2). \end{aligned} \quad (4)$$

The operator equation (3) defines a nonlinear first-order partial differential equation (PDE). This PDE does not in general allow for an analytic solution, and in the next section, I present a numerical method for solving it.

3 Computation

I employ projection methods or, more precisely, collocation techniques (Judd 1992, Judd 1998, chap. 11) to solve the PDE defined by the operator equation (3). The idea underlying projection methods is to convert the infinite-dimensional problem of solving the PDE for the unknown value function into a finite-dimensional problem of finding a zero of a system of equations. I accomplish this by approximating the value function by a high-order polynomial. Hence, instead of having to solve for the unknown value function, I only need to solve for the unknown coefficients of the polynomial approximation. The unknown coefficients in turn are chosen such that the polynomial approximation satisfies the PDE at some appropriately chosen points in the state space.

In a previous attempt to model an R&D race without the restrictive assumption of an exponential distribution of success times, Judd (1985) employs perturbation methods. However, the validity of Judd's (1985) approximations hinges on the value of the innovation being sufficiently small. The reason is that he uses the degenerate case of $\bar{P} = 0$ as a starting

point to approximate the solution to “nearby” cases in which the value of the patent is small.⁸ In contrast, the projection methods I use do not rely on $\bar{P} \approx 0$, and allow me to solve the PDE for an arbitrary value of the patent.

Approximating the Value Function. I approximate $V(z_1, z_2)$ using a tensor product basis⁹ of univariate Chebyshev polynomials,

$$\hat{V}(z_1, z_2) = \sum_{k_1=0}^K \sum_{k_2=0}^K \theta_{k_1, k_2} T_{k_1}(z_1) T_{k_2}(z_2),$$

where $T_{k_i}(z_i)$ is a k_i th order Chebyshev polynomial in z_i (Judd 1998, pp. 204) and $\theta = (\theta_{k_1, k_2})$ is a vector of $(K + 1)^2$ unknown coefficients. K is the order of the approximation.

Determining the Unknown Coefficients. Substituting $\hat{V}(z_1, z_2)$ for $V(z_1, z_2)$ in the operator equation (3), I define the residual function

$$\Delta^V(z_1, z_2; \theta) = \frac{1}{r\bar{P}} \mathcal{N}(\hat{V})(z_1, z_2), \quad (5)$$

where I divide by $r\bar{P}$ to make the residual function unit-free. While the unknown true value function by construction satisfies $\mathcal{N}(V)(z_1, z_2) = 0$ at all points (z_1, z_2) of the state space, the polynomial approximation will generally not. Hence, there will generally be a nonzero residual. By choosing the unknown coefficients θ such that the residual is “small”, I ensure that the polynomial approximation is “close to” the value function.

Collocation methods choose θ to ensure that the residual function is zero at $(K + 1)^2$ so-called collocation points $(z_1^{k_1}, z_2^{k_2})$.¹⁰ Hence, θ is the solution to the $(K + 1)^2$ nonlinear equations $\Delta^V(z_1^{k_1}, z_2^{k_2}; \theta) = 0$. While it is possible to use any collection of $(K + 1)^2$ points, the Chebyshev interpolation theorem suggests using the Cartesian product of the zeros of a $(K + 1)$ th order univariate Chebyshev polynomial. It turns out that the “fit” is improved by employing the expanded Chebyshev array (as defined by Judd (1998), p. 222) instead.

Thus I use

$$z_i^{k_i} = \sec\left(\frac{\pi}{2(K+1)}\right) \left(1 - \cos\left(\frac{2(k_i+1)-1}{2(K+1)}\pi\right)\right) \left(\frac{\bar{z}-z}{2}\right) + z, \quad k_i = 0, \dots, K,$$

which implies that $z_i^{k_i} \in [z, \bar{z}]$. Consequently, I am able to approximate $V(z_1, z_2)$ in the region $[z, \bar{z}]^2$. Since the value of the race depends on the knowledge stocks of both firms, $V(z_1, z_2)$ has domain $[0, \infty)^2$, and I set $z = 0$. Clearly, \bar{z} has to be large enough so that the R&D race will have ended with probability close to one by the time that either $z_1(t)$ or $z_2(t)$ reaches \bar{z} given its initial value of z_1^0 and z_2^0 , respectively. Of course, the appropriate choice of \bar{z} is only known ex-post and thus must be determined by experimentation. I use $\bar{z} = 1$ in what follows.¹¹

Accuracy Check. The quality of the polynomial approximation of the value function depends on the order of the approximation K . As K increases, the polynomial approximation becomes more flexible and “fits” the value function better. Indeed, I can make the residuals arbitrarily small (i.e., zero up to machine precision) by choosing K large enough. On the other hand, as K increases, so does the computational burden because a larger number of unknown coefficients has to be determined. I choose an intermediate value of K that ensures that the residuals $\Delta^V(z_1, z_2; \theta)$ are on the order of 10^{-7} to 10^{-10} (depending on the parameter values). That is, the error a firm makes in computing the value of the race is 10^7 to 10^{10} times smaller than the value of the patent at all points of the state space. Put differently, if the patent is worth a million dollars, then the firm errs by less than a cent. In addition to being small, the residuals come close to exhibiting the equioscillation property necessary for a best polynomial approximation. This indicates that there are no systematic errors in the polynomial approximation. Details are presented in Appendix A.

A firm’s behavior during an R&D race is described by its policy function $u^*(z_1, z_2)$, which in turn depends on the value function $V(z_1, z_2)$ and on the partial derivatives of the value function. Hence, it is important to obtain good approximations of $V_1(z_1, z_2)$ and $V_2(z_1, z_2)$ in addition to a good approximation of $V(z_1, z_2)$. I take the partial derivatives of the

approximation $\hat{V}_1(z_1, z_2)$ and $\hat{V}_2(z_1, z_2)$ as approximations of the partial derivatives $V_1(z_1, z_2)$ and $V_2(z_1, z_2)$. In general, even if $\hat{V}(z_1, z_2)$ converges to $V(z_1, z_2)$, this does not imply that $\hat{V}_1(z_1, z_2)$ and $\hat{V}_2(z_1, z_2)$ converge to $V_1(z_1, z_2)$ and $V_2(z_1, z_2)$, respectively. This raises the question of how good the partial derivatives of the approximation are as approximations of the partial derivatives, a question that has received scant attention in the literature to date.

I develop a general method that allows me to answer this question. The idea is that although the partial derivatives of the true value function are unknown, expressions for them can be derived.¹² Analogously to the residual function $\Delta^V(z_1, z_2; \theta)$ in equation (5), I define two additional residual functions

$$\begin{aligned}\Delta^{V_1}(z_1, z_2; \theta) &= \frac{\bar{z}}{r\bar{P}} \frac{\partial}{\partial z_1} \mathcal{N}(\hat{V})(z_1, z_2), \\ \Delta^{V_2}(z_1, z_2; \theta) &= \frac{\bar{z}}{r\bar{P}} \frac{\partial}{\partial z_2} \mathcal{N}(\hat{V})(z_1, z_2)\end{aligned}$$

corresponding to the two partial derivatives of the value function.¹³ If these residuals are small, then the partial derivatives of the polynomial approximation should be close to the unknown true partial derivatives. Given the intermediate value of K that I have chosen, it turns out that the residuals $\Delta^{V_1}(z_1, z_2; \theta)$ and $\Delta^{V_2}(z_1, z_2; \theta)$ are on the order of 10^{-4} to 10^{-8} (again depending on the parameter values, see Appendix A). That is, the error a firm makes in computing the value of additional knowledge is 10^4 to 10^8 times smaller than the value of the patent at all points of the state space. In short, the collocation techniques I employ deliver an extremely accurate approximation not only of the value function but also of the policy function.

Simulating the Time Paths. Once approximations of the value function and its partial derivatives have been obtained, they can be used to simulate the evolution of the R&D race over time. I employ Euler’s method (Judd 1998, chap. 10) to solve the system of ordinary differential equations (ODEs)

$$\dot{z}_i(t) = \hat{u}^{i*}(z_1(t), z_2(t)) - \delta z_i(t), \quad z_i(0) = z_i^0, \quad (6)$$

over the interval $[0, T]$, where T is constrained by the requirement that $\underline{z} \leq z_i \leq \bar{z}$.¹⁴ z_i^0 is firm i 's initial knowledge stock and $\hat{u}^{i*}(z_1, z_2)$ is given by equation (2) (with $V(z_1, z_2)$ replaced by $\hat{V}(z_1, z_2)$).

Solving the system of ODEs given by equation (6) yields $h_i(t)$, the hazard rate of a successful innovation by firm i , as a by-product. While the hazard rates are not of immediate interest themselves, they allow me to compute the probability that the R&D race has ended at or before time t , which is given by

$$p(t) = \text{Prob}(\tau \leq t) = 1 - \exp\left(-\int_0^t h_1(s) + h_2(s) ds\right), \quad (7)$$

where $\tau = \min\{\tau_1, \tau_2\}$ and τ_i is the random date of a successful innovation by firm i . In addition, I compute the probability that firm i has won the race given that the race has ended at or before time t ,

$$p_i(t) = \text{Prob}(\tau_i = \tau | \tau \leq t) = \frac{\int_0^t h_i(s)(1 - p(s)) ds}{p(t)}. \quad (8)$$

I numerically integrate using the trapezoid rule (Judd 1998, Chap. 7).¹⁵

Parameterization. Since the emphasis of this paper is on the role of knowledge accumulation in an R&D race, the functional form of the hazard rate is of primary importance. In Sections 4.1 and 4.2, I focus on five scenarios that capture a wide range of functional forms (see Table 1). First, I look at the special case in which the hazard rate depends on current R&D effort alone ($\gamma = 0$). Next, I analyze the polar case in which the hazard rate depends on past R&D effort alone ($\lambda = 0$). In this case, a firm's current R&D effort does not directly help it to win the race, but indirectly helps it to win the race by adding to its knowledge stock. Between these extreme cases are parameterizations in which both current and past R&D efforts enter the hazard rate. My starting point is $\lambda = \gamma = \psi = 1$, leading to a model in which the hazard rate is linear in the firm's knowledge stock. Then I allow for a nonlinear influence of past R&D efforts, and set $\psi = \frac{1}{2}$ and $\psi = 2$ to obtain a concave or a convex

hazard rate, respectively. Moreover, I set $\delta = 0.2$, $r = 0.105$, $\eta = 2$, and $\underline{P}/\bar{P} = 0.2$ in all five scenarios. Hence, the degree of patent protection is intermediate (I also examine the extreme cases of $\underline{P} = 0$ and $\underline{P} = \bar{P}$ in Section 4.1).

By changing the functional form of the hazard rate, I change the distribution of success times and thereby the duration of the race. It is to be expected that knowledge accumulation is *per se* of greater importance in a longer race than in a shorter race. Thus, in order to isolate the role of the functional form of hazard rate, I hold the expected duration of the race constant. To this end, I choose the remaining parameter \bar{P} such that the expected duration of the race is 3 years in all five scenarios. Clearly, the expected duration also depends on the initial knowledge stocks of the competing firms, and I set $z_1^0 = z_2^0 = 0$ for now (I allow for an initial asymmetry in firms' knowledge stocks in Section 4.2). The expected duration of the race along with the values for δ , r , η , and \underline{P}/\bar{P} are chosen to be somewhat representative of the R&D process in a wide range of industries.¹⁶

4 Results

In this section, I discuss the strategic implications of knowledge accumulation in an R&D race. The results of the numerical analysis are presented as follows: I first look at firms' behavior given their knowledge stocks as implied by the value and policy functions (Section 4.1). To track the evolution of the race, I then analyze the time paths of some variables of interest (Section 4.2). While Sections 4.1 and 4.2 focus on the five scenarios listed in Table 1, I solve the model for a wide range of parameter values in Section 4.3 to arrive at general results.

4.1 Equilibrium Payoffs and Strategies

In what follows, I discuss the properties of the value and policy functions. I focus on R&D expenditures $c(u^*(z_1, z_2))$, an increasing and convex transformation of current R&D effort $u^*(z_1, z_2)$, because R&D expenditures are measured in terms of monetary units divided by time units rather than “knowledge units” per unit of time. Figure 1 depicts $c(u^*(z_1, z_2))$ for

the scenarios listed in Table 1.

I start with the case of $\gamma = 0$, where the hazard rate is a function of current R&D effort alone. This case corresponds to the memoryless R&D race models analyzed by Reinganum (1981, 1982). The equilibrium payoffs and strategies are constant and thus independent of firms' knowledge stocks.

The Pure Knowledge Effect. Once I relax the restrictive assumption of an exponential distribution of success times, the equilibrium payoffs and strategies are no longer constant. The reason is that in this more general setting *knowledge is productive*. Indeed, it can be shown that $\lim_{z_1 \rightarrow \infty} V(z_1, z_2) = \bar{P}$ and $\lim_{z_2 \rightarrow \infty} V(z_1, z_2) = \underline{P}$. That is, a firm benefits in the limit as the size of its own knowledge stock approaches infinity because it wins the race for sure whereas the firm loses the race for sure as the size of its rival's knowledge stock approaches infinity.

The fact that knowledge is productive influences firms' behavior. Indeed, as Figure 1 shows, a firm's R&D expenditures $c(u^*(z_1, z_2))$ are decreasing in its own knowledge stock z_1 in the case of past R&D effort alone (top left panel), the case of a linear hazard rate (top right panel), and the case of a concave hazard rate (bottom left panel). Underlying this result is what could be called the *pure knowledge effect*. The pure knowledge effect is independent of strategic considerations. It comes about because a firm's past R&D efforts contribute to its chances of winning the R&D race. Due to the pure knowledge effect, the firm can afford to scale back its investment in R&D as its knowledge stock increases.

In the case of a convex hazard rate (bottom right panel), the increasing returns nature of the hazard rate gives a firm a strong and growing incentive to invest in R&D that initially offsets the pure knowledge effect. However, the pure knowledge effect eventually dominates, i.e., $c(u^*(z_1, z_2))$ first increases in z_1 , then decreases.

Leader versus Follower. One of the most important aspects of a race is how the strategic interactions between the racing firms depend on their relative positions. In memoryless R&D races, the leading firm invests the same in R&D as the lagging firm because, as noted

previously, the equilibrium strategies are independent of firms' knowledge stocks. Hence, the distinction between leader and follower is irrelevant in these models, and there is no sense in which one can properly speak of one competitor being ahead of another, or of the two competitors being neck-and-neck.

In multi-stage models, the follower devotes *less* resources to R&D than the leader. That is, the follower tends to give way to the leader, and a pattern of increasing dominance (or its stronger form – ϵ -preemption) arises. The reason lies in the *pure progress effect* which, independently of strategic considerations, causes a firm to increase its investment in R&D as it gets closer to the finishing line (Grossman & Shapiro 1987). The intuition is simply that the gain from winning the race from the intermediate stage of a two-stage race is larger than the gain from making the intermediate discovery from the initial stage.

In my model, a firm is able to conserve on its R&D investment as its knowledge stock increases due to the pure knowledge effect. Moreover, the firm with the larger knowledge stock is able to conserve more on its investment in R&D than the firm with the smaller knowledge stock. The follower thus devotes *more* resources to R&D than the leader, i.e., $c(u^*(z_1, z_2)) > c(u^*(z_2, z_1))$ if and only if $z_1 < z_2$. The pattern of strategic interactions among the racing firms is thus more like action-reaction than increasing dominance.

As can be seen from the levels of R&D expenditures in Figure 1, the follower tries to catch up with the leader in the case of past R&D effort alone, the case of a linear hazard rate, and the case of a concave hazard rate. The case of a convex hazard rate is again somewhat different because the increasing returns nature of the hazard rate initially offsets the pure knowledge effect. Consequently, the follower at first devotes less resources to R&D than the leader, i.e., $c(u^*(z_1, z_2)) < c(u^*(z_2, z_1))$ if z_1 is small and $z_1 < z_2$, whereas the follower later on devotes more resources to R&D than the leader, i.e., $c(u^*(z_1, z_2)) > c(u^*(z_2, z_1))$ if z_1 is large and $z_1 < z_2$. A closer inspection shows that a sufficient condition for $c(u^*(z_1, z_2)) > c(u^*(z_2, z_1))$ is $z_1 \geq 0.6446 - z_2$ and $z_1 < z_2$. In sum, in the case of a convex hazard rate, the laggard strives to catch up with the leader once he has a sufficiently large knowledge stock himself.

The extent to which the follower's current R&D effort exceeds the leader's depends on the functional form of the hazard rate. In the case of a linear hazard rate, for example, the difference between the follower's and the leader's current R&D effort is larger than in the case of past R&D effort alone. To see the reason for this, note that the follower works harder than the leader if and only if $c(u^*(z_1, z_2)) > c(u^*(z_2, z_1))$ for all $z_1 < z_2$ which is equivalent to

$$-\lambda V(z_1, z_2) + V_1(z_1, z_2) > -\lambda V(z_2, z_1) + V_1(z_2, z_1)$$

for all $z_1 < z_2$. In the case of $\lambda = 0$, a necessary and sufficient condition is thus $V_1(z_1, z_2) > V_1(z_2, z_1)$. In the case of $\lambda > 0$, $V_1(z_1, z_2) > V_1(z_2, z_1)$ and $V(z_1, z_2) < V(z_2, z_1)$ are jointly sufficient. But since the leader's chances of winning the race *ceteris paribus* exceed that of the follower, the value of the race to the follower is less than the value to the leader, i.e., $V(z_1, z_2) < V(z_2, z_1)$ if and only if $z_1 < z_2$. Hence, the effect of the slope of the value function is reinforced by the effect of its level.¹⁷

Aggressive versus Submissive Response: Size of Knowledge Stock. A further question of interest is how a firm's R&D expenditures depends on its rival's knowledge stock. There are two possibilities. First, the firm may decide to invest more as its rival's knowledge stock increases, i.e., the firm may respond *aggressively*. Alternatively, it may reduce its investment in R&D in response to an increase in its rival's knowledge stock, i.e., it may respond *submissively*.

Multi-stage models suggest a submissive response on the part of the follower, thus reinforcing the pattern of increasing dominance. Specifically, the follower slows down as he falls further behind whereas the leader may or may not speed up as he gets further ahead. In contrast, I find that a firm can either respond aggressively or submissively to an increase in its rival's knowledge stock, i.e., $c(u^*(z_1, z_2))$ can be either increasing or decreasing in z_2 . In my model, an aggressive or submissive response is not tied to a firm's relative position.

To see what determines whether a firm responds aggressively or submissively, consider the case of past R&D effort alone. Inspection of the partial derivatives of $c(u^*(z_1, z_2))$ reveals

that $\frac{\partial}{\partial z_2} c(u^*(z_1, z_2)) > 0$ whenever $z_1 \geq 0.6886 + 0.4856z_2$. Hence, a firm with a sufficiently large knowledge stock increases its own R&D effort as its rival accumulates knowledge. This aggressive response is confined to the leading firm (i.e., to points below the diagonal of the state space) whereas the lagging firm always responds submissively.

For other functional forms of the hazard rate, also the lagging firm may respond aggressively. Compared to the case of past R&D effort alone, the region of the state space where $c(u^*(z_1, z_2))$ increases in z_2 expands to include points above the diagonal in the case of a linear hazard rate. Indeed, a closer inspection shows that a sufficient condition for $\frac{\partial}{\partial z_2} c(u^*(z_1, z_2))$ to be positive is $z_1 \geq 0.2157 + 0.1894z_2$. To see the reason for this, note that the firm responds aggressively if and only if

$$-\lambda V_2(z_1, z_2) + V_{12}(z_1, z_2) > 0.$$

From the above expression it is clear that in the case of $\lambda = 0$, a necessary and sufficient condition for an aggressive response is $V_{12}(z_1, z_2) > 0$. In the case of $\lambda > 0$, $V_{12}(z_1, z_2) > 0$ and $V_2(z_1, z_2) < 0$ are jointly sufficient. But since the value of the race to the firm is decreasing in its rival's knowledge stock, the first effect is reinforced by the second. Hence, unlike the case of past R&D effort alone, the laggard may respond aggressively.

To summarize, in my model a firm responds aggressively if it has a sufficiently large knowledge stock and submissively otherwise. The finding that the lagging firm may respond aggressively seems counterintuitive at first glance because one would expect the laggard to realize that there is no point in trying even harder as its rival advances. On the other hand, each firm invests in R&D up to a point where the expected marginal benefit equals the marginal cost. Since the probability that firm 2 wins the R&D race in the next short interval of time dt is increasing in z_2 , the probability that firm 1 has to sustain whatever current R&D effort it chooses beyond time $t + dt$ is decreasing in z_2 . Hence, the point where the expected marginal benefit of investment in R&D equals its marginal cost is reached for an increasing level of R&D investment. Phrased differently, as firm 2 advances, firm 1 takes

its chances and invests heavily but briefly in R&D. Moreover, since the expected duration of the remainder of the race is also decreasing in z_1 , an aggressive response from firm 1 becomes more likely as it accumulates knowledge.

Aggressive versus Submissive Response: Value of Patent. Whether a firm responds aggressively or submissively to an increase in its rival's knowledge stock may also depend on the value of the patent. Figure 2 illustrates this for the case of a linear hazard rate. It presents the policy functions for $\bar{P} = 0.00435$ (left panel) and $\bar{P} = 0.435$ (right panel). Current R&D effort and the value of the race are increasing in \bar{P} . Moreover, as the value of the patent increases and competition becomes fiercer, a firm tends to respond aggressively rather than submissively as its rival accumulates knowledge. In fact, if $\bar{P} = 0.00435$, then $\frac{\partial}{\partial z_2} c(u^*(z_1, z_2)) < 0$ (submissive response), whereas, if $\bar{P} = 0.435$, then $\frac{\partial}{\partial z_2} c(u^*(z_1, z_2)) > 0$ (aggressive response). Put differently, in this example, $\frac{\partial}{\partial z_2} c(u^*(z_1, z_2))$ switches sign as \bar{P} increases. This demonstrates the limitations of Judd's (1985) analysis, who assumes that the value of the patent is close to zero in order to obtain $\frac{\partial}{\partial z_2} c(u^*(z_1, z_2)) < 0$.

Intensity of Competition. I measure the intensity of competition as the sum of firms' R&D expenditures $c(u^*(z_1, z_2)) + c(u^*(z_2, z_1))$. As Figure 3 shows for the scenarios listed in Table 1, $c(u^*(z_1, z_2)) + c(u^*(z_2, z_1))$ is quasi-concave in the case of a convex hazard rate and convex in the remaining cases involving $\gamma > 0$.

In multi-stage races, once a firm has gained an advantage over its rival, it tends to win the race. Consequently, the outcome of the race is "decided" while firms are neck-and-neck, and competition is fiercest in these situations. I obtain a similar result in the case of a convex hazard rate: Since $c(u^*(z_1, z_2)) + c(u^*(z_2, z_1))$ is quasi-concave, competition is equally intense on an ellipsis with center on the diagonal. This implies that, holding the combined amount of knowledge constant, competition is more intense among firms with equal knowledge stocks than among firms with unequal knowledge stocks. In contrast, $c(u^*(z_1, z_2)) + c(u^*(z_2, z_1))$ is convex for the other functional forms of the hazard rate. Hence, holding their combined knowledge constant, competition is most intense when firms have unequal knowledge stocks.

Degree of Patent Protection. Patent protection is perfect if $\underline{P} = 0$ and completely ineffective if $\underline{P} = \bar{P}$. Figure 4 illustrates the impact of the degree of patent protection for the case of a linear hazard rate. It describes the equilibrium strategies for $\underline{P} = 0$ (left panel) and $\underline{P} = \bar{P}$ (right panel). Current R&D effort is decreasing and the value of the race to a firm is increasing in \underline{P}/\bar{P} . As the value of imitation increases from 0 to \bar{P} , competition softens and, as Figure 4 shows, the policy function change its shape. Moreover, a firm's value function is no longer decreasing in its rival's knowledge stock. The reason is that the character of the R&D race changes from a preemption game into a waiting game as the degree of patent protection decreases.

To see this, I compare the planner's solution (collusive solution) to the noncooperative outcome. In a winner-take-all situation in which the winning firm is awarded a patent of positive value whereas the losing firm receives nothing ($\underline{P} = 0$), there is rent dissipation because each firm invests excessively in R&D. Since the two firms compete for the same discovery, each additional dollar invested in R&D brings a firm closer to winning the race and, at the same time, brings its rival closer to losing the race. Hence, its R&D efforts impose a *negative* externality on its rival, and the firm consequently invests excessively in R&D. In the polar case in which the loser can costlessly and immediately imitate the winner and thus both firms receive the same payoff ($\underline{P} = \bar{P}$), there is again rent dissipation, but the reason is now that firms invest too little in R&D. In contrast to a winner-take-all situation, each additional dollar invested in R&D brings both firms closer to the finish line. Hence, a firm's R&D efforts impose a *positive* externality on its rivals, which causes the firm to underinvest in R&D. In other words, depending on whether patent protection is perfect or completely ineffective, a firm's R&D efforts impose a negative or a positive externality on its rival, and the character of the R&D race changes from a preemption game into a waiting game.¹⁸

4.2 Time Paths

Given firms' knowledge stocks, the value and policy functions provide a “snapshot” of the R&D race. But since firms' knowledge stocks are changing over time, it is crucial to know how these “snapshots” fit together. Thus, in this section, I look at the time paths of some variables of interest and show how knowledge accumulation affects the evolution of the race.

Knowledge Stocks. Figure 5 shows the vector field for the case of a convex hazard rate (see Table 1). The vector field indicates the direction and speed of movement of the state variables as given by (\dot{z}_1, \dot{z}_2) for each point (z_1, z_2) in the state space. It thus summarizes the time paths of firms' knowledge stocks for all possible initial positions. Note that since the equilibrium is symmetric and accumulated knowledge evolves deterministically, firms are neck-and-neck at all times during the R&D race if they start from the same position. Moreover, starting from different positions, firms' knowledge stocks eventually reach a symmetric steady state.¹⁹ Hence, an initial asymmetry in firms' knowledge stocks vanishes over time. The fact that the gap between firms' knowledge stocks closes over time contrasts with the pattern of increasing dominance that emerges in multi-stage races.

Figure 6 illustrates the evolution of the race when firm 1 starts with 50% of the steady-state knowledge stock and firm 2 starts from scratch. I use the system of ODEs given by equation (6) to compute the time path of firm i 's accumulated knowledge, $z_i(t)$, and then obtain the time path of its R&D expenditures, $c(u_i^*(t))$.²⁰ A dash-dotted (dashed) line designates firm 1 (firm 2). As can be seen in the top panel of Figure 6, the follower catches up with the leader as the knowledge stocks monotonically approach the steady state over time, i.e., $z_i(t) \rightarrow 0.5435$ as $t \rightarrow \infty$. There are two reasons for this: First, at all times the impact of depreciation is worse for the leader than it is for the follower. Second, as the middle panel of Figure 6 shows, the leader decreases its R&D expenditures over time whereas the follower initially increases its R&D expenditures markedly and eventually decreases them slightly. Overall, while the follower at first invests less in R&D than the leader, he eventually invests more, which causes the gap to shrink.

The Pure Knowledge Effect versus Strategic Considerations. The pure knowledge effect provides a firm with an incentive to decrease its R&D expenditures as its own knowledge stock becomes larger. A firm may, however, also decide to be aggressive and increase its R&D expenditures as its rival's knowledge stock becomes larger. The question then is which incentive dominates.

Consider again the case of a convex hazard rate and recall that firm 1's R&D expenditures decrease over time (middle panel of Figure 6). A closer inspection shows that the time paths of $z_1(t)$ and $z_2(t)$ traverse a region of the state space where $\frac{\partial}{\partial z_1}c(u^{1*}(z_1, z_2)) < 0$ and $\frac{\partial}{\partial z_2}c(u^{1*}(z_1, z_2)) > 0$. This gives rise to two contradicting incentives: First, as firm 1 increases its knowledge stock, the pure knowledge effect provides it with an incentive to decrease its R&D expenditures. Second, as firm 2 increases its knowledge stock, firm 1 responds aggressively and thus has an incentive to increase its R&D expenditures. The pure knowledge effect prevails, and firm 1 decreases its investment in R&D as the race unfolds.

The pure knowledge effect also prevails in the case of a linear hazard rate (see again Table 1). For example, if firm 1 begins the race with 150% of the steady-state knowledge stock and firm 2 starts from scratch, then $z_1(t)$ is monotonically decreasing and $z_2(t)$ is monotonically increasing (top panel of Figure 7). Since $\frac{\partial}{\partial z_2}c(u^{2*}(z_1, z_2)) < 0$ and $\frac{\partial}{\partial z_1}c(u^{2*}(z_1, z_2)) < 0$ along these time paths, as firm 2 extends its knowledge stock, it has an incentive to decrease its R&D expenditures whereas, it has an incentive to increase its R&D expenditures as firm 1 extends its knowledge stock. The middle panel of Figure 7 shows that firm 2's R&D expenditures decrease over time. Hence, the pure knowledge effect dominates over strategic considerations

Conditional Probability of Winning the Race. In the case of $\gamma = 0$, where the hazard rate is a function of current R&D effort alone, both firms have equal chances of winning the race. In particular, a firm's chances of winning are independent of the knowledge stocks at the outset of the race. In contrast, once the restrictive assumption of an exponential distribution of success times is relaxed, an initial asymmetry in the knowledge stocks affects

a firm's chances of winning the race.

To see this, consider again the case of a linear hazard rate. I use equation (7) to compute $p(t)$, the probability that one of the firms wins the race at or before time t , and equation (8) to compute $p_i(t)$, the probability that firm i has won the race given that the race has ended at or before time t . A solid line refers to $p(t)$ and a dash-dotted (dashed) line designates firm 1 (firm 2). Firm 2 works harder than firm 1 at all times, thereby narrowing the gap between the firms' knowledge stocks over time. As a result firm 2's conditional probability of winning increases from 0.1437 to 0.2817 whereas firm 1's conditional probability decreases from 0.8563 to 0.7183 (bottom panel of Figure 7).

In sum, although the gap between the firms' knowledge stocks closes over time, this does not quite suffice to make up for an initial asymmetry. The laggard's conditional probability of winning remains less than the leader's. This also clarifies why the leader is willing to cut back on its R&D investment in the first place: Doing so simply does not severely reduce its chances of winning the race.

4.3 General Results

In this section, I solve the model for a wide range of parameter values in order to “establish” general results about the equilibrium payoffs and strategies and the induced time paths. I consider all parameterizations such that $\bar{P} \in \{0.00435, 0.0435, 0.435\}$, $\underline{P}/\bar{P} \in \{0, 0.2, 0.8\}$, $\lambda \in \{0.1, 1, 10\}$, $\gamma \in \{0.1, 1, 10\}$, $\psi \in \{1, 2\}$, $\delta \in \{0, 0.2, 2\}$, and $\eta \in \{2, 3, 4\}$. Note that the interest rate r merely determines the time scale, and is therefore not of interest by itself.²¹ This set of 1458 parameterizations contains the case of a linear hazard rate (see Table 1) that I have studied in Sections 4.1 and 4.2. There I have also highlighted the differences between a linear and a convex hazard rate. Consequently, I specify a linear hazard rate ($\psi = 1$) in 729 parameterizations and a convex hazard rate ($\psi = 2$) in the others.²²

To compute the value and policy functions, I proceed as follows. Given a parameterization, I first solve the model for different orders of approximation. Specifically, I use $K = 3, 5, 7, 9, 11$ for all parameterizations and, in addition, $K = 13, 15, 17, 19$ for some pa-

parameterizations. Then I pick the order of approximation that yields the smallest residuals of the value function and its partial derivatives. In 3% of the parameterizations, the residuals are unacceptably large²³, and I ignore these parameterizations. There is a small number (less than 1%) of parameterizations that contradict the results stated below. A closer inspection shows that these contradictions arise because of numerical problems (e.g., round-off error or numerical instability of the polynomial approximation). I therefore disregard these parameterizations in what follows.

Median Duration of the Race. I use equation (7) to compute $p(t)$, the probability that the race ends at or before time t . Judging from the median durations (defined as the solution to $p(t) = 0.5$), the remaining parameterizations capture a wide range of R&D processes: In 10% of parameterizations the median duration of the race is below 0.1 years, in 50% it is below one year, and in 90% it is below 30.2 years. The following result summarizes the comparative statics properties of the median duration.

Result 1 *The median duration is decreasing in \bar{P} , λ , γ , and η .*

The effect of a change in \underline{P}/\bar{P} , ψ , and δ is ambiguous.

Knowledge Stocks. I argued in Section 4.2 that the knowledge stocks approach a steady state provided that the rate of depreciation is positive. I compute the steady-state knowledge stock $z^\infty = \lim_{t \rightarrow \infty} z_i(t)$. The following result summarizes its comparative statics properties.

Result 2 *The steady-state knowledge stock z^∞ is increasing in \bar{P} and λ and decreasing in \underline{P}/\bar{P} and δ .*

To get a sense for how fast the steady-state knowledge stock is approached, I linearize the system of ODEs given by equation (6) around z^∞ and compute its eigenvalues.

Result 3 *The eigenvalues of the linearized system are real and negative. The eigenvalues are decreasing in δ .*

Since the eigenvalues are real and negative, the system of ODEs given by equation (6) is locally stable. The largest (in absolute value) eigenvalue determines the rate of convergence in the vicinity of the steady state. The above result indicates that the steady-state knowledge stock is approached faster as the rate of depreciation increases.

Recall that I approximate the value and policy functions in the region $[0, 1]^2$. For some parameterizations the steady-state knowledge stock lies outside this region, so that I am unable to exactly determine its value. In what follows, I restrict attention to the region $[0, \hat{z}]^2$, where $\hat{z} = \min\{z^\infty, 1\}$. The reason is that, if $z^\infty \leq 1$, then firm i 's knowledge stock $z_i(t)$ remains below \hat{z} provided that its initial knowledge stock z_i^0 is sufficiently small. Otherwise, the race has ended with high probability before $z_i(t)$ reaches \hat{z} .

My next result is that the laggard adds to its knowledge stock as long as its knowledge stock is below the steady state.

Result 4 *If $z_1 \leq z_2 < \hat{z}$, then $\dot{z}_1 > 0$.*²⁴

In contrast, $z_1 \leq z_2 < \hat{z}$ does not imply $\dot{z}_2 > 0$. That is, the leader may not add to its knowledge stock. In fact, if the firm is far ahead of its rival, then the leader may relax and allow himself to fall back ($\dot{z}_2 < 0$) while the laggard builds up its knowledge stock. Then, once the gap between the firms has closed somewhat, the leader resumes ($\dot{z}_2 > 0$). It follows that the convergence to the steady state need not be monotonic.

The Pure Knowledge Effect. Recall that the pure knowledge effect arises because a firm's past R&D efforts contribute to its chances of winning the race. Hence, as I argued in Section 4.1, a firm is inclined to decrease its R&D expenditures as its knowledge stock increases. In general, the pure knowledge effect determines the shape of the policy function if the hazard rate is linear.

Result 5 *If $\psi = 1$, then $\frac{\partial}{\partial z_1} c(u^*(z_1, z_2)) < 0$.*

On the other hand, there is a counteracting force if the hazard rate is convex and there are increasing returns to knowledge accumulation. However, as the following result establishes, once the pure knowledge effect starts to dominate, it continues to do so.

Result 6 *If $z'_1 > z_1$ and $\frac{\partial}{\partial z_1}c(u^*(z_1, z_2)) < 0$, then $\frac{\partial}{\partial z_1}c(u^*(z'_1, z_2)) < 0$.*

To see that the above result has content, note that $\frac{\partial}{\partial z_1}c(u^*(0, 0)) < 0$ in 51% of parameterizations, $\frac{\partial}{\partial z_1}c(u^*(\frac{\hat{z}}{2}, \frac{\hat{z}}{2})) < 0$ in 79% of parameterizations, and $\frac{\partial}{\partial z_1}c(u^*(\hat{z}, \hat{z})) < 0$ in 83% of parameterizations. Hence, it frequently happens that a firm decreases its investment in R&D in response to an increase in its knowledge stock, so that the above result applies.

Leader versus Follower. Perhaps the most important implication of the pure knowledge effect is that the follower devotes more resources to R&D than the leader, thus giving rise to a pattern of action-reaction. The economic intuition developed for the case of a linear hazard rate in Section 4.1 again generalizes.

Result 7 *If $\psi = 1$ and $z_1 < z_2$, then $c(u^*(z_1, z_2)) > c(u^*(z_2, z_1))$.*

A corollary to the above result is that $\psi = 1$ and $z_1 < z_2$ imply $\dot{z}_1 > \dot{z}_2$. Consequently, the gap between the leader's and the follower's knowledge stocks closes over time.

In contrast, $z_1 < z_2$ implies $c(u^*(z_1, z_2)) > c(u^*(z_2, z_1))$ in 13% of parameterizations with $\psi = 2$ and $z_1 < z_2$ implies $\dot{z}_1 > \dot{z}_2$ in 68% of parameterizations with $\psi = 2$. As I pointed out in Section 4.1, the reason is that the laggard may not strive to catch up with the leader until he has a sufficiently large knowledge stock himself.

Aggressive versus Submissive Response. A firm can either respond aggressively or submissively to an increase in its rival's knowledge stock. As I claimed in Section 4.1, a firm responds aggressively if it has a sufficiently large knowledge stock and submissively otherwise. The following result pertains to all parameterizations with $\psi = 1$ and establishes that, once a firm starts to respond aggressively, it continues to respond aggressively.

Result 8 *If $\psi = 1$, $z'_1 > z_1$, and $\frac{\partial}{\partial z_2}c(u^*(z_1, z_2)) > 0$ then $\frac{\partial}{\partial z_2}c(u^*(z'_1, z_2)) > 0$.*

Moreover, the above result extends to 98% of parameterization with $\psi = 2$. Again this has content because $\frac{\partial}{\partial z_2}c(u^*(0, 0)) > 0$ in 49% of parameterizations, $\frac{\partial}{\partial z_2}c(u^*(\frac{\hat{z}}{2}, \frac{\hat{z}}{2})) > 0$ in 57% of parameterizations, and $\frac{\partial}{\partial z_2}c(u^*(\hat{z}, \hat{z})) > 0$ in 65% of parameterizations.

Intensity of Competition. My final result shows that, holding the combined amount of knowledge constant, competition is fiercest when firms have unequal knowledge stocks; it pertains to all parameterizations with $\psi = 1$.

Result 9 *If $\psi = 1$ and $0 < \lambda < 1$, then*

$$\begin{cases} c(u^*(z, 0)) + c(u^*(0, z)) > c(u^*(\lambda z, (1 - \lambda)z)) + c(u^*((1 - \lambda)z, \lambda z)), & z \leq \hat{z}, \\ c(u^*(\hat{z}, z - \hat{z})) + c(u^*(z - \hat{z}, \hat{z})) > c(u^*(\hat{z}_1, \hat{z}_2)) + c(u^*(\hat{z}_2, \hat{z}_1)), & z > \hat{z}, \end{cases}$$

where $\hat{z}_1 = (2\lambda - 1)\hat{z} + (1 - \lambda)z$ and $\hat{z}_2 = (1 - 2\lambda)\hat{z} + \lambda z$.

In contrast, in 74% of parameterizations with $\psi = 2$, competition is fiercest when firms have equal knowledge stocks.

5 Conclusions

I develop a general model of an R&D race that incorporates knowledge accumulation. The model does not allow for an analytic solution, and I use projection methods to obtain the equilibrium payoffs and strategies. I propose a method to check the accuracy of the approximations. The approximations seem to be close to both the true value function and its partial derivatives, and thus give an accurate description not only of the equilibrium payoffs but also of the equilibrium strategies. This suggests that projection techniques are a promising tool for the analysis of differential games.

My model does not inherit the memorylessness property of the exponential distribution that troubles existing models of R&D races. Indeed, the memoryless R&D race models analyzed by Reinganum (1981, 1982) are a special case of my model, and I show how relaxing the restrictive assumption of an exponential distribution of success times affects the equilibrium payoffs and strategies. In the more general setting of my model, knowledge accumulation has strategic implications.

In particular, I show that a firm has an incentive to reduce its R&D expenditures as its knowledge stock increases. Underlying this result is the pure knowledge effect. The source

of this effect is that a firm's past R&D efforts contribute to its chances of winning the R&D race because the firm's knowledge stock enters its hazard rate. If the hazard rate is concave or linear, the pure knowledge effect determines the shape of the policy function and implies that the follower works harder than the leader. If the hazard rate is convex and there are increasing returns to knowledge accumulation, then a firm has a strong and growing incentive to invest in R&D. While this gives rise to a counteracting force, the pure knowledge effect gathers force as the race unfolds. Hence, once he has a sufficiently large knowledge stock himself, the laggard strives to catch up with the leader. The pattern of strategic interactions among the racing firms is thus more like action-reaction than increasing dominance, the pattern that emerges in multi-stage race models.

In multi-stage race models, the follower slows down as he falls further behind whereas the leader may or may not speed up as he gets further ahead. In contrast, I find that a firm can either respond aggressively or submissively to an increase in its rival's knowledge stock. In my model, an aggressive or submissive response is not tied to a firm's relative position. Rather, a firm responds aggressively if it has a sufficiently large knowledge stock and submissively otherwise. These strategic considerations appear to be dominated by the pure knowledge effect in the sense that the response in the firm's investment in R&D to a change in its own knowledge stock swamps the response to a change in its rival's knowledge stock.

Also in contrast to multi-stage race models, I show that competition is not necessarily fiercest when firms are neck-and-neck. If the hazard rate is linear, competition among firms is most intense when their knowledge stocks are of unequal size and least intense when they are of equal size, whereas this need not be the case if the hazard rate is convex.

Despite the abundance of anecdotal evidence, empirical research on R&D races is sparse. In line with my model, the available studies observe patterns of strategic interactions that are more akin to action-reaction than to increasing dominance. Grabowski & Baxter (1973) find that in the chemical industry firms increase R&D expenditures in response to rivals' outlays. Based on a sample of 28 U.S. manufacturing industries, exploratory research by

Richard Caves²⁵ indicates that leaders and followers react positively to each other's increases in R&D expenditures, and Scherer (1992) finds that firms with greater domestic sales in more concentrated U.S. markets were likely to react much more aggressively to increasing import competition than smaller firms or firms in less concentrated markets. More recent studies attempt to operationalize the notion of a knowledge stock. Khanna (1995) proposes to measure a firm's technological position relative to its rivals by constructing a technological frontier for the high-end computer industry. He shows that firms that fall behind the technological frontier engage in catch-up behavior. In his study of the disk drive industry, Lerner (1997) attempts to directly measure a firm's technological position relative to its rivals and demonstrates that the firms that trail the industry leader display a greater propensity to innovate. This again lends more support to action-reaction than to increasing dominance.

At the same time, empirical research points to a number of extensions of my model. Studies of the pharmaceutical industry in particular highlight the role of spillovers in the R&D process (Cockburn & Henderson 1994, Henderson & Cockburn 1994, Henderson & Cockburn 1996). Moreover, the studies by Cohen & Levinthal (1989) and Adams & Jaffe (1996) indicate that a firm has to itself engage in R&D in order to absorb spillovers. Extending the notion of absorptive capacity to "spillovers" from a firm's own knowledge stock requires specifying a nonadditive hazard rate. A multiplicative hazard rate, for example, implies that the firm's past R&D efforts are not productive unless combined with current R&D effort. Spillovers and a nonadditive hazard rate are readily incorporated into the model. The studies by Khanna & Iansiti (1997) and Lerner (1997) finally testify to the important role product market competition and repeated interactions between competing firms play in determining the outcome of an R&D race. This suggests modelling the two-way relationship between innovative activity and market structure, a task that is left to future research.

Appendix A Computational Details

All programs are written in Matlab 5.3. I use the equation solver c05nbf from the NAG toolbox, a Newton method, to solve the $(K+1)^2$ nonlinear equations $\Delta^V(z_1^{k_1}, z_2^{k_2}; \theta) = 0$ for θ . My starting values for the elements of θ are zeros except for $\theta_{0,0}$ which is chosen to solve (with $\theta_{0,0}$ replacing V)

$$0 = \lambda u^*(\bar{P} + \underline{P}) - \frac{1}{\eta} (u^*)^\eta - (r + 2\lambda u^*)V,$$

where

$$u^* = (\lambda(\bar{P} - V))^{\frac{1}{\eta-1}}.$$

The above equations characterize the unique constant solution to the PDE defined by the operator equation (3) in the special case of $\gamma = 0$. Good starting values are essential to ensure convergence of the equation solver.

To validate the approximations of the value function and its partial derivatives, I compute $\Delta^V(z_1, z_2; \theta)$, $\Delta^{V_1}(z_1, z_2; \theta)$, and $\Delta^{V_2}(z_1, z_2; \theta)$ at 30^2 equidistant grid points in $[\underline{z}, \bar{z}]^2$. Since the grid points differ from the collocation points, $\Delta^V(z_1, z_2; \theta)$ is not necessarily zero at these grid points. Moreover, there is no reason for $\Delta^{V_1}(z_1, z_2; \theta)$ and $\Delta^{V_2}(z_1, z_2; \theta)$ to be zero anywhere in $[\underline{z}, \bar{z}]^2$. Nevertheless, as I argued before, if all three sets of residuals are small, then the polynomial approximation and its partial derivatives should be close to the value function and its partial derivatives. I summarize each residual function by picking its largest (absolute) value and denote their sup-norms as L_∞^V , $L_\infty^{V_1}$, and $L_\infty^{V_2}$, respectively.

The system of equations generally admits multiple zeros. It turns out, however, that exactly one of these zeros gives rise to small residuals; the residuals that correspond to the other zeros are larger by several orders of magnitude. Moreover, the equation solver converges to this zero whenever I use the starting values described above. This suggests that there exists a unique solution to the PDE.

Among the five scenarios listed in Table 1 on page 44, the ones involving a nonlinear hazard rate are the most troublesome, and I devote particular attention to them in what follows. In general, a nonlinear hazard rate adds curvature to the model, thereby making it harder to approximate the value function. Moreover, a hazard rate of the form $h_i = \lambda u_i + \gamma z_i^\psi$ does not have a derivative at $z_i = 0$ if $\psi < 1$. This gives rise to a singularity in the case of a concave hazard rate. I move this singularity from $z_i = 0$ to $z_i = -\tilde{z}$ by respecifying firm i 's hazard rate as

$$h_i = \lambda u_i + \gamma \left((z_i + \tilde{z})^\psi - \tilde{z}^\psi \right)$$

if $\psi < 1$. Since the curvature of the hazard rate decreases as \tilde{z} increases, a small value for \tilde{z} is appropriate, I and set $\tilde{z} = 0.1$ in what follows.

Table 2 lists the sup-norms L_∞^V , $L_\infty^{V_1}$, and $L_\infty^{V_2}$ for various orders of approximation K for the case of a convex hazard rate and Table 3 does so for the case of a concave hazard rate. The repercussions of the singularity in the case of a concave hazard rate are still fairly pronounced. In particular, for fixed K , the residuals for the concave hazard rate are by several orders of magnitude larger than the residuals for the convex hazard rate and, as K increases, they approach zero at a much slower rate. Nevertheless, the approximation improves steadily for both parameterizations and is already quite good for intermediate

values of K . Moreover, the elements of θ that correspond to higher order terms are extremely small (in absolute value) for both parameterizations.

Figure 8 displays the residual functions $\Delta^V(z_1, z_2; \theta)$, $\Delta^{V_1}(z_1, z_2; \theta)$, and $\Delta^{V_2}(z_1, z_2; \theta)$ for the case of a convex hazard rate with $K = 11$. The residuals $\Delta^V(z_1, z_2; \theta)$ oscillate around zero, and are close to exhibiting the equioscillation property necessary for a best (with respect to the sup-norm) polynomial approximation. The residuals $\Delta^{V_1}(z_1, z_2; \theta)$ and $\Delta^{V_2}(z_1, z_2; \theta)$ are large at the “upper” and “lower” boundaries of $[\underline{z}, \bar{z}]^2$. However, the residuals are much smaller in the interior of $[\underline{z}, \bar{z}]^2$ and, as a closer inspection reveals, fluctuate around zero. Despite the fact that the partial derivatives of the value function are approximated somewhat worse than the value function itself, overall the choice of $K = 11$ leads to extremely accurate approximations.

It remains to choose the order of approximation for the remaining parameterizations. I set $K = 19$ in the case of a concave hazard rate and $K = 11$ in the case of past R&D effort alone and in the case of a linear hazard rate. For the case of past R&D effort alone, I obtain $L_\infty^V = 2.01 \times 10^{-10}$, $L_\infty^{V_1} = 2.52 \times 10^{-8}$, and $L_\infty^{V_2} = 1.75 \times 10^{-8}$ and for the case of a linear hazard rate I obtain $L_\infty^V = 5.81 \times 10^{-10}$, $L_\infty^{V_1} = 5.32 \times 10^{-8}$, and $L_\infty^{V_2} = 6.35 \times 10^{-8}$.

References

- Adams, J. & Jaffe, A. (1996). Bounding the effects of R&D: An investigation using matched establishment-firm data, *Rand Journal of Economics* **24**(7): 700–721.
- Argote, L., Beckman, S. & Epple, D. (1990). The persistence and transfer of learning in an industrial setting, *Management Science* **36**(2): 140–154.
- Athey, S. & Schmutzler, A. (2001). Investment and market dominance, *Rand Journal of Economics* **32**(1): 1–26.
- Basar, T. & Olsder, J. (1999). *Dynamic noncooperative game theory*, 2nd edn, Society for Industrial and Applied Mathematics, Philadelphia.
- Benkard, L. (2000). Learning and forgetting: The dynamics of aircraft production, *American Economic Review* **90**(4): 1034–1054.
- Cockburn, I. & Henderson, R. (1994). Racing to invest? The dynamics of competition in ethical drug discovery, *Journal of Economics and Management Strategy* **3**(3): 481–519.
- Cohen, W. & Levinthal, D. (1989). Innovation and learning: The two faces of R&D, *Economic Journal* **99**: 569–596.
- Doraszelski, U. & Draganska, M. (2002). Rent dissipation in R&D races, *Working paper*, Hoover Institution, Stanford.
- Fudenberg, D., Gilbert, R., Stiglitz, J. & Tirole, J. (1983). Preemption, leapfrogging and competition in patent races, *European Economic Review* **22**: 3–31.
- Grabowski, H. & Baxter, N. (1973). Rivalry in industrial research and development, *Journal of Industrial Economics* **21**: 209–235.
- Grabowski, H. & Vernon, J. (1994). Returns to R&D on new drug introductions in the 1980s, *Journal of Health Economics* **13**: 383–406.
- Grossman, G. & Shapiro, C. (1987). Dynamic R&D competition, *Economic Journal* **97**(386): 372–387.
- Harris, C. & Vickers, J. (1985a). Patent races and the persistence of monopoly, *Journal of Industrial Economics* **33**(4): 461–481.
- Harris, C. & Vickers, J. (1985b). Perfect equilibrium in a model of a race, *Review of Economic Studies* **52**(2): 193–209.
- Harris, C. & Vickers, J. (1987). Racing with uncertainty, *Review of Economic Studies* **54**(1): 1–21.
- Henderson, R. & Cockburn, I. (1994). Measuring competence? Exploring firm effects in pharmaceutical research, *Strategic Management Journal* **15**: 63–74.
- Henderson, R. & Cockburn, I. (1996). Scale, scope, and spillovers: The determinants of research productivity in drug discovery, *Rand Journal of Economics* **7**(1): 32–59.

- Judd, K. (1985). Closed-loop equilibrium in an innovation race with experience, *Working paper*, Northwestern University, Evanston.
- Judd, K. (1992). Projection methods for solving aggregate growth models, *Journal of Economic Theory* **58**: 410–452.
- Judd, K. (1998). *Numerical methods in economics*, MIT Press, Cambridge.
- Judd, K., Schmedders, K. & Yeltekin, S. (2002). Optimal policies for patent races, *Working paper*, Hoover Institution, Stanford.
- Khanna, T. (1995). Racing behavior: Technological evolution in the high-end computer industry, *Research Policy* **24**: 933–958.
- Khanna, T. & Iansiti, M. (1997). Firm asymmetries and sequential R&D: Theory and evidence from the mainframe computer industry, *Management Science* **43**(4): 405–421.
- Lee, T. & Wilde, L. (1980). Market structure and innovation: A reformulation, *Quarterly Journal of Economics* **94**(2): 429–436.
- Lerner, J. (1997). An empirical exploration of a technology race, *Rand Journal of Economics* **28**: 228–247.
- Lippman, S. & McCardle, K. (1987). Dropout behavior in R&D races with learning, *Rand Journal of Economics* **18**(2): 287–295.
- Lippman, S. & McCardle, K. (1988). Preemption in R&D races, *European Economic Review* **32**: 1661–1669.
- Loury, G. (1979). Market structure and innovation, *Quarterly Journal of Economics* **93**: 395–410.
- Meron, A. & Caves, R. (1991). Rivalry among firms in research and development outlays, *unpublished manuscript*, Harvard University, Cambridge.
- Reinganum, J. (1981). Dynamic games of innovation, *Journal of Economic Theory* **25**: 21–41.
- Reinganum, J. (1982). A dynamic game of R and D: Patent protection and competitive behavior, *Econometrica* **50**: 671–688.
- Reinganum, J. (1989). The timing of innovation: Research, development, and diffusion, in R. Schmalensee & R. Willig (eds), *Handbook of industrial organization*, North-Holland, Amsterdam.
- Scherer, F. (1992). *International high-tech competition*, Harvard University Press, Cambridge.
- Vedenov, D. & Miranda, M. (2001). Numerical solution of dynamic oligopoly games with capital investment, *Economic Theory* **18**: 237–261.

Notes

¹*New York Times*, January 20, 2000 and October 11, 2000.

²*New York Times*, September 15, 1998 and *Business Week*, June 12, 2000.

³Reinganum (1989) refers to memoryless R&D race models as symmetric R&D race models. There is also a third class, models of repeated races. Since these models address quite different questions, I ignore them here. See Reinganum (1989) for a survey.

⁴Unfortunately, the terms increasing dominance and action-reaction have been used differently by different authors. I take them to describe the relationship between state variables and firms' actions. That is, increasing dominance (action-reaction) means that the firm that is ahead in the race invests more (less) in R&D than its rival. This has also been called weak increasing dominance by Athey & Schmutzler (2001).

⁵If $\psi < 1$, a computational difficulty requires me to respecify firm 1's hazard rate as $h_1 = \lambda u_1 + \gamma((z_1 + \tilde{z})^\psi - \tilde{z}^\psi)$, where $\tilde{z} > 0$. Details are given in Appendix A.

⁶I would like to thank an anonymous referee for suggesting this interpretation.

⁷In general, further investments in R&D may be needed to arrive at a marketable product after the discovery has been made. See Judd, Schmedders & Yeltekin (2002) for a model that endogenizes the value of innovation along these lines.

⁸Judd (1985) assumes that patent protection is perfect and thus sets $\underline{P} = 0$. Hence, if the benefits to innovation are zero, then the equilibrium strategy is not to invest in R&D, $u^*(z_1, z_2) = 0$, which yields a payoff of $V(z_1, z_2) = 0$.

⁹A comparison between tensor products and complete polynomials indicated little difference.

¹⁰I explored two alternatives to Chebyshev collocation. First, I chose the unknown coefficients to minimize the residual sum of squares at $M > (K + 1)^2$ collocation points. Second, I turned to Galerkin methods. Galerkin methods integrate the residual against the basis functions and solve the resulting system of $(K + 1)^2$ nonlinear equations for the unknown coefficients. The differences were negligible.

¹¹The exact value of \bar{z} has little or no effect on the value and policy functions once I compensate for increasing \bar{z} by increasing K .

¹²This idea has also been exploited by Vedenov & Miranda (2001) in the context of a discrete-time stochastic game.

¹³The resulting expressions can be simplified using the envelope theorem. The application of the envelope theorem is justified since $\hat{u}^*(z_1, z_2)$ by construction satisfies the FOC exactly.

¹⁴To apply Euler's method, I specify a grid of points on the time axis, $t^l = ldt$, $l =$

$0, 1, \dots$, where the step size dt is small but positive, and approximate the system of differential equations with a system of difference equations, $z^{l+1} = z^l + \dot{z}^l dt$, where $z^l = (z_1(t^l), z_2(t^l))$ and $\dot{z}^l = (\dot{z}_1(t^l), \dot{z}_2(t^l))$. Starting from $z^0 = (z_1^0, z_2^0)$ this allows me to compute $z^1 = (z_1^1, z_2^1)$, which in turn enables me to compute $z^2 = (z_1^2, z_2^2)$, and so on. The error of Euler's method is proportional to dt , thus displaying linear convergence.

¹⁵The accuracy of the numerical integration procedure is readily checked using the fact that $p_1(t) + p_2(t) = 1$. For $dt = 10^{-2}$, the difference between $p_1(t) + p_2(t)$ and unity is on the order of 10^{-6} (or better depending on the exact parameterization), for $dt = 10^{-1}$ ($dt = 10^{-3}$) the error is on the order of 10^{-4} (10^{-8}). I therefore choose $dt = 10^{-2}$ in what follows.

¹⁶The empirical literature has studied the pharmaceutical industry in great detail (e.g., Cockburn & Henderson (1994), Grabowski & Vernon (1994), and Henderson & Cockburn (1996)), and I use these studies to pick parameter values.

¹⁷While $V(z_1, z_2)$ determines $u^*(z_1, z_2)$, the value function also takes the equilibrium strategy into account. Hence, the nature of my argument is more suggestive than formal.

¹⁸See Doraszelski & Draganska (2002) for details.

¹⁹Of course, the existence of a steady state presupposes nonzero depreciation.

²⁰ $c(u_i^*(t))$ is shorthand for $c(u^{i*}(z_1(t), z_2(t)))$.

²¹I continue to set $r = 0.105$.

²²To double check my conclusions, I analyze an additional set of 1458 parameterizations that entails $\bar{P} \in \{0.01035, 0.1035, 1.035\}$ instead of $\bar{P} \in \{0.00435, 0.0435, 0.435\}$. This set contains the case of a convex hazard rate.

²³In the notation of Appendix A, a parameterization is deemed acceptable if $L_\infty^V < 10^{-4}$, $L_\infty^{V_1} < 10^{-2}$, and $L_\infty^{V_2} < 10^{-2}$.

²⁴I verified this result at 30^2 equidistant grid points in $[0, \hat{z}]^2$ for each parameterization. A similar remark pertains to the subsequent results.

²⁵Personal communication. See also Meron & Caves (1991).




Figure 1: R&D expenditures $c(u^*(z_1, z_2))$ in the case of past R&D effort alone (top left panel), linear hazard rate (top right panel), concave hazard rate (bottom left panel), and convex hazard rate (bottom right panel). In the case of current R&D effort alone, $c(u^*) = 0.0139$ (omitted).




Figure 2: Value of patent. R&D expenditures $c(u^*(z_1, z_2))$ for $\bar{P} = 0.00435$ (left panel) and $\bar{P} = 0.435$ (right panel). Linear hazard rate.

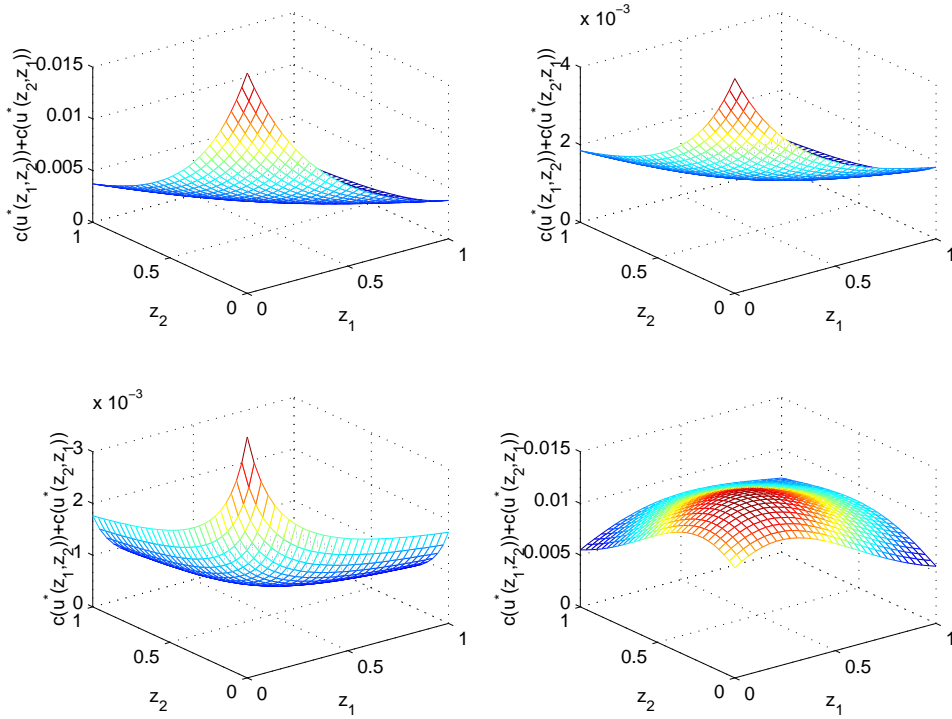


Figure 3: Intensity of competition $c(u^*(z_1, z_2)) + c(u^*(z_2, z_1))$ in the case of past R&D effort alone (top left panel), linear hazard rate (top right panel), concave hazard rate (bottom left panel), and convex hazard rate (bottom right panel). In the case of current R&D effort alone, $2c(u^*) = 0.0227$ (omitted).




Figure 4: Degree of patent protection. R&D expenditures $c(u^*(z_1, z_2))$ for $\underline{P} = 0$ (left panel) and $\underline{P} = \bar{P}$ (right panel). Linear hazard rate.




Figure 5: Vector field. Convex hazard rate.




Figure 6: Time paths of $z_1(t)$ and $z_2(t)$ (top panel), $c(u_1^*(t))$, $c(u_2^*(t))$, and $c(u_1^*(t)) + c(u_2^*(t))$ (middle panel), as well as $p_1(t)$, $p_2(t)$, and $p(t)$ (bottom panel). A dash-dotted (dashed) line designates firm 1 (firm 2). Vertical lines indicate the time by which the race has ended with a probability of 0.5, 0.9, and 0.99, respectively. Initial knowledge stocks are $z_1^0 = 0.2717$ and $z_2^0 = 0$. Convex hazard rate.




Figure 7: Time paths of $z_1(t)$ and $z_2(t)$ (top panel), $c(u_1^*(t))$, $c(u_2^*(t))$, and $c(u_1^*(t)) + c(u_2^*(t))$ (middle panel), as well as $p_1(t)$, $p_2(t)$, and $p(t)$ (bottom panel). A dash-dotted (dashed) line designates firm 1 (firm 2). Vertical lines indicate the time by which the race has ended with a probability of 0.5, 0.9, and 0.99, respectively. Initial knowledge stocks are $z_1^0 = 0.3673$ and $z_2^0 = 0$. Linear hazard rate.

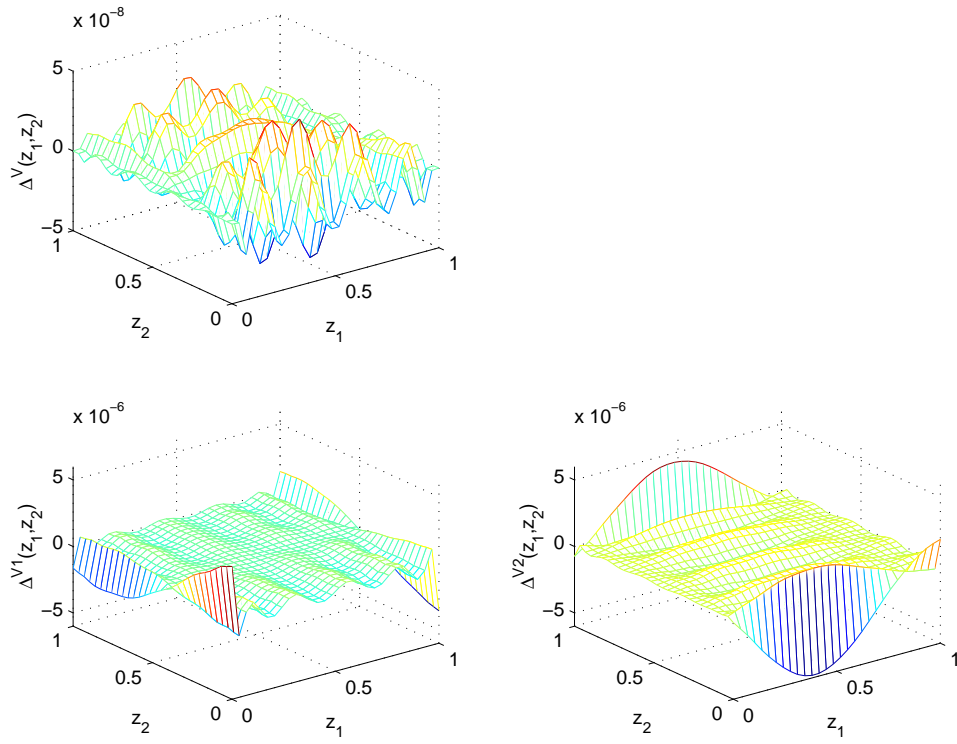


Figure 8: Residuals $\Delta^V(z_1, z_2; \theta)$ (top left panel), $\Delta^{V_1}(z_1, z_2; \theta)$ (bottom left panel), and $\Delta^{V_2}(z_1, z_2; \theta)$ (bottom right panel). Convex hazard rate and $K = 11$.

	λ	γ	ψ	\bar{P}
<i>polar cases:</i>				
current R&D effort alone	1	0	–	0.248
past R&D effort alone	0	1	1	0.137
<i>intermediate cases:</i>				
linear hazard rate	1	1	1	0.0435
concave hazard rate	1	1	$\frac{1}{2}$	0.0350
convex hazard rate	1	1	2	0.1035

Table 1: Parameter values.

K	L_∞^V	$L_\infty^{V_1}$	$L_\infty^{V_2}$
3	5.54×10^{-2}	1.15	3.81×10^{-1}
7	1.92×10^{-5}	1.09×10^{-3}	8.43×10^{-4}
11	4.92×10^{-8}	3.99×10^{-6}	5.81×10^{-6}
15	1.37×10^{-10}	1.48×10^{-8}	3.58×10^{-8}
19	7.51×10^{-13}	8.88×10^{-11}	1.55×10^{-10}

Table 2: Sup-norms L_∞^V , $L_\infty^{V_1}$, and $L_\infty^{V_2}$ for different orders of approximation K . Convex hazard rate.

K	L_∞^V	$L_\infty^{V_1}$	$L_\infty^{V_2}$
3	1.40×10^{-1}	3.27	2.12
7	2.87×10^{-3}	2.70×10^{-1}	1.74×10^{-1}
11	8.86×10^{-5}	2.31×10^{-2}	1.51×10^{-2}
15	4.98×10^{-6}	2.01×10^{-3}	1.35×10^{-3}
19	1.78×10^{-7}	1.76×10^{-4}	1.20×10^{-4}
23	1.66×10^{-8}	1.54×10^{-5}	1.06×10^{-5}
27	7.40×10^{-10}	1.34×10^{-6}	9.31×10^{-7}

Table 3: Sup-norms L_∞^V , $L_\infty^{V_1}$, and $L_\infty^{V_2}$ for different orders of approximation K . Concave hazard rate.