

Optimal Rules for Patent Races*

Kenneth L. Judd	Karl Schmedders
Hoover Institution	Kellogg School of Management
Stanford University	Northwestern University
judd@hoover.stanford.edu	k-schmedders@kellogg.nwu.edu

Şevin Yeltekin
Kellogg School of Management
Northwestern University
s-yeltekin@kellogg.nwu.edu

April 18, 2002

Abstract

There are two important rules in a patent race: what an innovator must accomplish to receive the patent and the allocation of the benefits that flow from the innovation. Most patent races end before R&D is completed and the prize to the innovator is often less than the social benefit of the innovation. We derive the optimal combination of prize and minimal accomplishment necessary to obtain a patent for a dynamic multi-stage innovation race. Competing firms are assumed to possess perfect information about each others' innovation state and cost structures. A planner, who cannot distinguish between the firms, chooses the innovation stage at which the patent is awarded and the magnitude of the prize to the winner. We examine both social surplus and consumer surplus maximizing patent race rules. A key consideration is the efficiency costs of transfers and of monopoly power to the patentholder. If efficiency costs are low and the planner maximizes social surplus, then races are undesirable. However, as efficiency costs of transfers associated with the patent rise, the optimal prize is reduced and the optimal policy spurs innovative effort through a race of nontrivial duration. Races are also used to filter out inferior innovators since a long race (i.e. a high minimal accomplishment requirement) in innovation makes it less likely for an inefficient firm to win through random luck. In general, races do serve to spur innovation.

*The authors are grateful to Mort Kamien, Ariel Pakes, Mark Satterthwaite, Christopher Sleet, and seminar audiences at Stanford University, MIT, Harvard University, and the University of Texas at Austin for helpful discussions. The authors also thank Ulrich Doraszelski for detailed comments on an earlier draft and Stuart Halliday for excellent editorial assistance.

1 Introduction

Races are designed to motivate agents. Firms involved in innovation race to develop a new product first and obtain exclusive rights to sell that product as in a patent, or race to develop a product that conforms to a buyer's specifications. For example, airplane producers compete to build new military planes which meet Department of Defense specifications. In both cases, there is a principal that designs a race in which innovators compete for a prize. There are many dimensions to the design of a patent race. Conventional discussions of patent policy focus on the optimal duration and breadth of patent protection. These discussions typically assume that a firm does not receive the patent until the R&D process is complete. This is not true of actual innovation processes where a firm often bears significant expenditures after it receives a patent. For example, drug firms can patent a drug before they have proven its efficacy and safety. The time at which a patent or exclusive contract is awarded to a firm is an important element of patent policy design. This paper focuses on optimal design of a patent race in a multistage model of innovation.

There are important trade-offs in constructing the rules of a patent race. A long race with a large prize stimulates innovators to work hard. However, much of the effort is duplicative and wasteful. A short race may reduce overall duplication of effort, but may lead to poor intertemporal resource allocation since firms would work very hard to win the patent and but then proceed much more slowly to finish R&D. Reducing the prize reduces all investment effort and delays the arrival of a socially valuable product. It is not obvious which effect dominates in choosing the optimal rules.

Consideration of these issues also highlights the importance of being explicit about the preferences of the designer of the race and the constraints he faces. We explicitly consider the efficiency costs which naturally arise in efforts to compensate innovators, such as the inefficiencies of monopoly pricing and the deadweight burden of cash prizes financed by distortionary taxes. We also consider the limits a designer faces; for example, there may be many positive externalities which make the profits from a patent small relative to the total social benefits. We examine two different objectives for the designer; social surplus (roughly defined as the social benefit of the innovation less the total cost of innovation and the distortionary costs of transfers to the patentholder) and consumer surplus (roughly defined as social benefit of the innovation less the prize to the winning firm and distortionary costs).

We use a simple multistage model of a race. In each stage of the game, a firm's position in the patent race represents the current state of its knowledge. Each firm's R&D investment determines the stochastic rate at which it advances in the race. The race is a game of perfect information where each firm knows its opponent's cost function and current state. Firms compete to reach the stage at which a patent or similar monopoly is awarded. At this stage, the laggard firms are forced to leave the race and the winner continues to invest in R&D until the innovation process is complete and a socially valuable product is produced. A planner chooses when a patent is awarded and the winner's

prize. We study how the primitives of our model, i.e. the technology of innovation, the preferences of the planner, and the heterogeneity of the firms, along with the inefficiencies associated with prizes affect the optimal patent rules.

R&D competition and optimal patent policies have been studied widely in the industrial organization literature. Our paper bridges the gap between two lines of research on R&D competition and optimal patent policies. The first line of research concerns the study of competition and investment in patent races. Some earlier examples in this literature are Kamien and Schwartz (1982), Loury (1979), Lee and Wilde (1980), Reinganum (1982a,b) and Dasgupta and Stiglitz (1980a,b). In these models, a firm's probability of success of obtaining a patent at a point in time depends only on that firm's current R&D expenditure and not on its past R&D experience. Competition takes place in "memoryless" or "Poisson" environments (see also the survey article by Reinganum (1989)). Subsequent work on races incorporates learning and experience. Fudenberg et al. (1983) and Harris and Vickers (1985a,b, 1987) formalize learning or experience effects in patent races by assuming that a firm's probability of discovery per unit of time depends not just on current R&D expenses, but on experience accumulated to date. The work of Harris and Vickers and Fudenberg et al. shows that competition in R&D may be strongly restricted by first-mover advantages and experience effects. These models display ϵ -preemption: once a firm attains a small leadership position, the laggard dramatically reduces its investment level and the leader wins with high probability. Doraszelski (2000) and Judd (1985) introduce experience effects in an extension of Reinganum (1982a,b). Doraszelski's model has no ϵ -preemption, showing that the specific modelling of dynamics dramatically affects the nature of the race. These papers also take patent policy as fixed.

The second line of research our work is related to focuses on issues regarding optimal patent policy, in particular optimal patent length and breadth; see Nordhaus (1969), Klemperer (1990), Gilbert and Shapiro (1990), Denicolo (1999, 2000), and Hopenhayn and Mitchell (2000). The question of when to issue a patent is generally ignored in the R&D literature.

Our dynamic game closely resembles games studied by Fudenberg et al. (1983), Harris and Vickers (1985a,b, 1987). However, we do not take the rules that define the patent policy as given. We consider the problem of a social planner who chooses the stage at which a patent is rewarded and the winning firm's prize once it has completed R&D. We present a simple model that shows which factors are important in designing rules for patent races. Our results indicate that optimal patent policy may involve both patents granted in early stages and in later stages of development, depending on the efficiency costs of transfers to patentholders and the degree of heterogeneity between innovating firms. For example, if the innovators differ in their cost of innovation but the planner cannot distinguish between them, then a race is a device for filtering out the inferior competitors. On the other hand, if all innovators have similar costs then the optimal policy is to grant a patent early, to avoid excessive rent dissipation. Another critical factor is the efficiency cost of transfers to the winner.

If the efficiency cost of monopoly is high, then the planner reduces the prize, which subsequently reduces innovative effort. To stimulate investment effort in this case, the planner varies the stage of development at which the patent is granted. The specification of the planner's preferences also affects the optimal rules. If the social planner maximizes social surplus, then short races with large prizes are commonly optimal whereas if he maximizes consumer surplus, long races with smaller prizes are chosen.

This paper basically asks what purposes do races serve. We find that the patent race serves two purposes in our model. First, it motivates the firms to invest heavily and complete the innovation process quickly when the prize alone cannot, due to inefficiencies or limitations, adequately serve to motivate innovation. Second, it serves as a filtering device for the planner who can verify when a firm has achieved the requirement for patenting, but cannot observe an individual firm's efficiency as an innovator.

This paper also makes a contribution to the literature on numerical solutions of dynamic games. The standard algorithms, such as that in Pakes and McGuire (1994), are too slow for our problem since we need to solve thousands of games to find optimal races. We develop an algorithm for solving multistage races that exploits their natural structure. Since this structure is common in dynamic games, the algorithm is applicable in a variety of strategic dynamic environments.

The remainder of the paper is organized as follows. Section 2 presents our model of a dynamic game for a patent race. In Section 3 we discuss in detail our computational method for the computation of equilibria. We present results from many computations in Section 4. Section 5 concludes the paper with a discussion of possible extensions of our research.

2 The Model

We assume two kinds of infinitely-lived agents: innovator firms and a social planner, which we call "the patent granting authority" (PGA). Innovation requires the completion of N stages of development. We assume that each firm controls a separate innovation process. Each firm begins at stage 0 and the firm that first reaches the stage $D \leq N$ obtains exclusive rights to continue. We call that exclusive right a "patent" even though we mean to model any institutional arrangement where a buyer constructs a race among potential seller-innovators. The choice of D corresponds to filing requirements for a patent. After winning the race, the patentholder completes the final $N - D$ stages without competition. When the patentholder reaches stage N the innovation process ends, the social benefits of the innovation become available and these benefits are allocated between the patentholder and the rest of society.

We let B denote the present value of the innovation's potential social benefits. This includes the potential social surplus of a new good as well as any technological or knowledge spillovers into other markets. We assume that the patentholder receives a fraction, γ , of these benefits as a prize

$\Omega = \gamma B$. The prize may be literally a cash prize or, like a patent, it may be a grant of a monopoly which produces a profit flow with present value Ω .

The PGA maximizes its objective by choosing γ , the fraction of potential social benefits that goes to the patentholder, and $D \in \{0, \dots, N\}$, the patent-granting stage. The case $D = 0$ represents the case in which there is no race. In the game, it formally corresponds to the PGA giving the patent at random to one of the firms. It also represents the case where the patent requirements are so minor that the patent goes to whomever, with trivial effort, first comes up with the barest notion of the innovation. The key assumption is that $D = 0$ corresponds to the case where each firm has equal chance of winning without having made any investment.

Firms compete in a multistage innovation game in discrete time. In each period, firms have perfect information about each other's cost structure and position and choose their investment levels simultaneously. The PGA, on the other hand, chooses the rules of the race before it begins. The simplicity of the PGA's actions corresponds to actual patent law and to the situations where it is impractical to continuously monitor the race. We do not intend to present a full optimal mechanism design analysis of the R&D problem. Instead, we analyze policy choices faced by actual policymakers such as patent law officials.

We use the framework of previous papers such as Fudenberg et al. (1983), and Harris and Vickers (1985a,b, 1987). We follow the standard race framework and ignore all possible transactions and cooperation between the two firms for several reasons. First, these transactions may be unprofitable for the firms for reasons not explicitly modelled here. Most research is done by multi-product firms. Such firms have substantial amounts of private information and intellectual property which they may not want to share with other firms. If a firm wins the patent, it may avoid any cooperation with another firm because such technical cooperation or transactions may lead to the leakage of other valuable private information. We do not model these considerations, but note that they clearly reduce the likelihood that firms would engage in transactions or cooperation.

Second, cooperation may hurt overall social welfare. For example, if Ford, GM, Honda, and Toyota all worked together on engine research, they may slow down the pace of technological progress to reduce obsolescence of their current technologies. Furthermore, cooperation in technology may lead to collusion in other matters such as pricing. Again, we do not explicitly model these issues, but note that they make it less likely that any social planner would allow cooperation. Third, we show that under certain circumstances, a social planner would shut down trades and use the race to stimulate innovation and to filter out inefficient firms.

For the remainder of this paper, we concentrate on the two-firm case for ease of exposition and reasons of tractability. We first present the details of the equilibrium behavior of the firms given D and Ω . We then more precisely describe the PGA's preferences.

2.1 The Firms: A multistage Model of Racing

The patent race with a specific Ω and D creates a dynamic game between the two firms. Let $x_{i,t}$ denote firm i 's stage at time t . We assume that each firm starts at stage 0; therefore, $x_{1,0} = x_{2,0} = 0$. If firm i is at stage n then it can either stay at n or advance to $n + 1$ ¹, where the probability of jumping to $n + 1$ depends on firm i 's investment, denoted $a_i \in A = [0, \bar{A}] \subset \mathbb{R}_+$. The upper bound \bar{A} on investment is chosen sufficiently large so that it never binds in equilibrium. Firm i 's state evolves according to

$$x_{i,t+1} = \begin{cases} x_{i,t}, & \text{with probability } p(x_{i,t}|a_{i,t}, x_{i,t}) \\ x_{i,t} + 1, & \text{with probability } p(x_{i,t} + 1|a_{i,t}, x_{i,t}). \end{cases}$$

There are many functional forms we could use for $p(x|a, x)$. We choose a probability structure so that innovation resembles search and sampling. Let $F(x|x) = p(x|1, x)$, that is, $F(x|x)$ is the probability that there is no change in the state if $a = 1$. For general values of a we assume.

$$\begin{aligned} p(x|a, x) &= F(x|x)^a \\ p(x + 1|a, x) &= 1 - F(x|x)^a. \end{aligned} \tag{1}$$

This structure can be motivated by a coin tossing analogy. For $a = 1$, equation (1) says we toss a coin and move ahead if heads comes up, a probability $F(x|x)$ event, but otherwise stay put. For integer values of a , equation (1) says that we move ahead one stage if and only if we flip a coins and at least one comes up heads. This specification is like hiring a people to work for one period and having them work independently on the problem of moving ahead one stage. While this specification is a special one, its simple statistical foundation helps us interpret our results.²

During R&D, firm i 's cost function is $C_i(a)$, $i = 1, 2$, assumed to be strictly increasing and weakly convex in a . For the remainder of the paper, we assume the cost function for firm i is

$$C_i(a) = c_i a^\eta, \quad \eta \geq 1, \quad c_i > 0, \quad i = 1, 2.$$

Firms discount future costs and revenues at the common rate of $\beta < 1$ and maximize their expected discounted payoffs.

2.2 Equilibrium

The patent race involves two phases. When one of the firms reaches stage D , it becomes the only innovator. The monopoly phase contains the set of states after the patent is awarded in stage

¹We have computed solutions to our model with firms being able to advance more than one stage in each period. These changes do not lead to any results that contradict the basic insights of this paper. Computational results with larger jumps can be obtained from the authors upon request.

²This specification allows only forward movement. While this is typical of most of the patent race literature, recent work by Doraszelski (2001) examines a model with “forgetting”, that is, x_{t+1} may be less than x_t .

$D, X^M = \{D, D + 1, \dots, N\}$. The duopoly phase is the set of states before the patent has been granted. During the duopoly phase the positions of the two firms are denoted by $x = (x_1, x_2)$. The set of states in the duopoly phase is

$$X^D = \{(x_1, x_2) | x_i \in \{0, \dots, D\}, i = 1, 2\}.$$

We first solve for the monopoly phase and then for the duopoly phase, since the monopoly phase can be solved independently of the duopoly phase, but not vice versa.

2.2.1 Monopoly Phase

We formulate the monopolist's problem recursively. At the terminal stage N , the innovation process is over and the monopolist receives a prize of Ω . In stages D through $N - 1$, the monopolist spends resources on investment. Let $V_i^M(x_i)$ denote the value function of firm i if it is a monopoly in state x_i . V_i^M solves the Bellman equation

$$\begin{aligned} V_i^M(x_i) &= \max_{a_i \in A} \left\{ -C_i(a_i) + \beta \sum_{x'_i \geq x_i} p(x'_i | a_i, x_i) V_i^M(x'_i) \right\}, \quad D \leq x_i < N \\ V_i^M(N) &= \Omega. \end{aligned} \quad (2)$$

The policy function of the monopolist is defined by

$$a_i^M(x_i) = \arg \max_{a_i \in A} \left\{ -C_i(a_i) + \beta \sum_{x'_i \geq x_i} p(x'_i | a_i, x_i) V_i^M(x'_i) \right\}, \quad D \leq x_i < N. \quad (3)$$

Proposition 1 *Firm i 's monopoly problem at state $x_i \in \{0, 1, \dots, N\}$ has a unique optimal solution $a_i^M(x_i)$. The value function V_i^M and the policy function a_i^M are nondecreasing in the state x_i .*

Proof. See Appendix. ■

2.2.2 Duopoly Phase

During the duopoly phase, firms compete to reach D first. We restrict attention to Markov strategies. A pure Markov strategy $\sigma_i : X^D \rightarrow A$ for firm i is a mapping from the state space X to its investment set A . We define the firms' value functions recursively. Let $\mathbb{V}_i(x)$ represent the value of firm i 's value function if the two firms are in state $x = (x_1, x_2) \in X^D$. If at least one of the firms has reached the patent stage D , firm i 's value function is defined as follows:

$$\mathbb{V}_i(x_i, x_{-i}) = \begin{cases} V_i^M(x_i), & \text{for } x_{-i} < x_i = D \\ V_i^M(x_i)/2, & \text{for } x_i = x_{-i} = D \\ 0, & \text{for } x_i < x_{-i} = D. \end{cases} \quad (4)$$

If neither firm has received the patent, the Bellman equations for the two firms are defined by

$$\mathbb{V}_i(x_i, x_{-i}) = \max_{a_i \in A} \left\{ -C_i(a_i) + \beta \sum_{x'_i, x'_{-i}} p(x'_i | a_i, x_i) p(x'_{-i} | a_{-i}, x_{-i}) \mathbb{V}_i(x'_i, x'_{-i}) \right\}, \quad (5)$$

for $x_i, x_{-i} < D$, for $i = 1, 2$.

The optimal strategy functions of the firms must satisfy

$$\sigma_i(x_i, x_{-i}) = \arg \max_{a_i \in A} \left\{ -C_i(a_i) + \beta \sum_{x'_i, x'_{-i}} p(x'_i | a_i, x_i) p(x'_{-i} | a_{-i}, x_{-i}) \mathbb{V}_i(x'_i, x'_{-i}) \right\}, \quad (6)$$

for $x_i, x_{-i} < D$, for $i = 1, 2$.

We now define the Markov perfect equilibrium of the race.

Definition 1 *A Markov perfect equilibrium (MPE) is a pair of value functions \mathbb{V}_i , $i = 1, 2$, and a pair of strategy functions σ_i^* , $i = 1, 2$, such that*

1. *Given σ_{-i}^* , the value function \mathbb{V}_i solves the Bellman equation (5), $i = 1, 2$.*
2. *For $a_{-i} = \sigma_{-i}^*$ the strategy function σ_i^* solves equation (6), $i = 1, 2$.*

A Markov perfect equilibrium always exists.

Theorem 1 *There exists a Markov perfect equilibrium.*

Proof. See the Appendix. ■

2.3 The Objective and Constraints of the Patent Granting Authority

The multistage race between the firms implicitly makes assumptions about what the PGA can observe. We assume that the PGA does not offer different prizes to different firms, corresponding to actual policy. However, we do assume that the PGA is aware of the technology of innovation. Specifically, we assume that the PGA knows the parameters of the two cost functions, but does not know any particular firm's costs; it must therefore offer the same incentives to all firms.³

³It may be possible to elicit information about a firm's costs. It may also be possible to hire firms to conduct R&D under the guidance of some central planner. However, that is not what a patent system does. Our analysis is a long way from being a fully specified mechanism design analysis; it represents instead the nature of feasible alternatives within a patent system. Our focus in this paper is on patent races, therefore we abstract from policies that would allow the PGA to conduct its own research and development by employing the firms in question.

Once a firm has invented a marketable good, the allocation of social benefits is governed by the firm's marketing policies and the terms of the patent. Figure 1 displays how the potential social benefit of a new good is allocated per period. Suppose that demand is given by DD and that there is a constant marginal cost of production. Figure 1 assumes that the patentholder can sell the new good at the monopoly price, but not engage in price discrimination, creating a profit Pf for the firm and leaving consumers with a surplus of CS . The area H represents the deadweight loss from monopoly pricing.

Once the patent has expired, the good is assumed to sell competitively at marginal cost, implying that consumers will receive all the social benefits, which equal $CS + Pf + H$. We assume that the PGA chooses the prize, denoted by Ω , received by the innovator once he has completed the R&D project. The PGA may have various tools at hand, such as direct payments and patent length and duration, but these decisions essentially fix Ω . We assume that the prize equals a proportion γ of the present value of potential social benefit; hence, $\Omega = \gamma B$. We focus on the fraction γ since profits from patents are proportional to demand, and, therefore, roughly proportional to social benefits B .

The PGA may face constraints on its choice of γ . For example, if the PGA faced the situation in Figure 1, then B equals the present value of $CS + Pf + H$, and even if the patent had infinite life, the present value of profits is at most equal to one-half of B . Furthermore, it may be difficult to protect a patent forever, reducing the practical size of Ω . More generally, γ may be reduced if firms are not able to charge the full monopoly price; for example, moral considerations (and fear of regulation) may lead drug manufacturers to restrain their prices. Therefore, we also specify an upper limit on the PGA's choice of γ , $\bar{\gamma} \leq 1$, which represents various constraints on what proportion of B can be transferred to the patentholder.

In Figure 1, the deadweight loss H represents the social cost of monopoly profits in patent system.⁴ More generally, we assume that the deadweight loss is proportional to the profits received by the innovator, and equals $\theta\Omega = \theta\gamma B$ for some $\theta \geq 0$. For example, $\theta = 0.5$ in Figure 1. This linear specification for deadweight loss captures the basic point that $\gamma > 0$ causes inefficiencies, and is an exact description of this loss when demand is linear and marginal costs are constant, and when demand has constant elasticity and marginal cost is zero. There are similar inefficiencies when Ω is a cash prize financed by distortionary taxes. In that case, θ represents the marginal efficiency cost of funds, a number which can plausibly be as low as .1 or as high as 1, depending on estimates of various elasticities, tax policy parameters, and the source of marginal funds; see Judd (1987) for a discussion

⁴Price controls may be used to reduce the deadweight loss, but they would also reduce monopoly profits and the prize. Long-lived patents will increase γ but at the expense of increasing the total deadweight losses of monopoly. Cash prizes may be granted by the PGA along with shorter duration patents. This will reduce the time during which the market experiences the deadweight loss H , but it only creates other inefficiencies since society bears the distortionary cost of the taxes used to finance the prize. Therefore, there will be inefficiencies no matter what financing scheme is used.

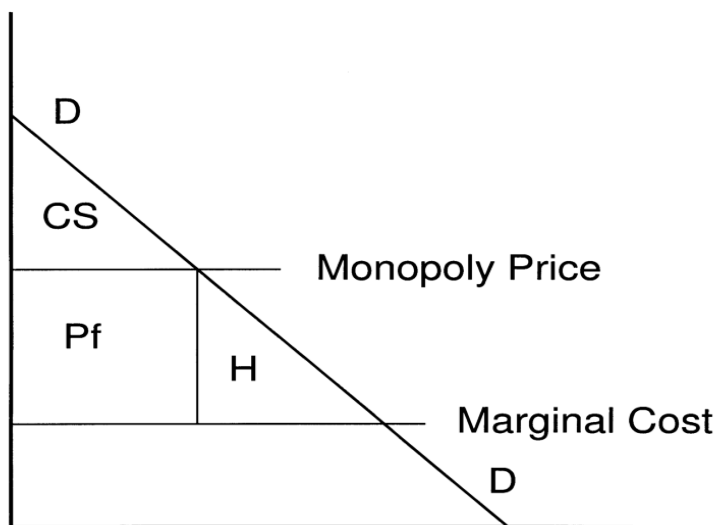


Figure 1: Allocation of Potential Social Benefits

of these factors. Therefore, the θ parameter represents either the relation between deadweight loss and profits for monopoly or the marginal efficiency cost of tax revenue.

In addition to γ , the patent authority also chooses D , the stage at which the race is ended to maximize some objective. We consider two different specifications of the PGA's preferences. In our first specification, the PGA maximizes total social surplus, which equals the present discounted value of the social benefit B minus $\theta\Omega$ minus total investment cost from the patent race. In our second specification, the PGA maximizes the present discounted value of consumer surplus, $(1 - \gamma)B - \theta\Omega$. We examine optimal patent policy when the PGA's preferences are of the latter type because it may represent the preferences of the median voter who is likely to be a consumer waiting for new goods. It may also represent the preferences of a buyer who is providing incentives to two suppliers that must engage in innovation to produce the desired product.

Given the equilibrium strategies $\sigma_i(x)$ of the race and optimal policy function $a_i^M(x)$ during the

monopoly phase, we can define the social surplus function W^S recursively as follows:

$$\begin{aligned}
W^{S,D}(x_1, x_2) &= -\sum_{i=1}^2 C_i(\sigma_i(x)) + \beta \sum_{x'_i, x'_{-i}} p(x'_1|\sigma_1(x), x_1)p(x'_2|\sigma_2(x), x_2)W(x'_1, x'_2), \quad x_1, x_2 < D \\
W(x_1, x_2) &= \begin{cases} W^{S,D}(x_1, x_2), & x_1, x_2 < D \\ \frac{1}{2}(W^{S,M}(1, D) + W^{S,M}(2, D)), & x_1 = x_2 = D \\ W^{S,M}(i, x_i), & x_i = D \text{ and } x_{-i} < D, \quad i = 1, 2 \end{cases} \\
W^{S,M}(i, x_i) &= -C_i(a_i^M(x)) + \beta \sum_{x'_i \geq x_i} p(x'_i|a_i^M(x_i), x_i)W^{S,M}(i, x'_i), \quad x_i < N, \quad i = 1, 2 \\
W^{S,M}(N) &= B - \theta\gamma B.
\end{aligned}$$

The initial social surplus at $t = 0$ equals

$$W^S(D, \gamma; \theta, B) = W^{S,D}(0, 0).$$

The consumer surplus function W^C is similarly defined as

$$\begin{aligned}
W^{C,D}(x_1, x_2) &= \beta \sum_{x'_i, x'_{-i}} p(x'_1|\sigma_1(x), x_1)p(x'_2|\sigma_2(x), x_2)W(x'_1, x'_2), \quad x_1, x_2 < D \\
W(x_1, x_2) &= \begin{cases} W^{C,D}(x_1, x_2), & x_1, x_2 < D \\ \frac{1}{2}(W^{C,M}(1, D) + W^{C,M}(2, D)), & x_1 = x_2 = D \\ W^{C,M}(i, x_i), & x_i = D \text{ and } x_{-i} < D, \quad i = 1, 2 \end{cases} \\
W^{C,M}(i, x_i) &= \beta \sum_{x'_i \geq x_i} p(x'_i|a_i^M(x_i), x_i)W^{C,M}(i, x'_i), \quad x_i < N, \quad i = 1, 2 \\
W^{C,M}(N) &= (1 - \gamma)B - \theta\gamma B.
\end{aligned}$$

Initial consumer surplus at $t = 0$ equals

$$W^C(D, \gamma; \theta, B) = W^{C,D}(0, 0).$$

Definition 2 *The social surplus maximizing patent policy is a pair (D^*, γ^*) that maximizes $W^S(D, \gamma; \theta, B)$ given (θ, B) . The consumer surplus maximizing patent policy is a pair (D^*, γ^*) that maximizes $W^C(D, \gamma; \theta, B)$ given (θ, B) .*

3 Computing Optimal Patent Policies

For any specific patent policy, (D, γ) , we need to compute the equilibrium of the race which involves solving two dynamic problems. First, we solve the dynamic optimization problem for each firm after it wins the patent. Second, we solve the patent race in the duopoly phase. We discuss the solution procedures for these two problems in detail.

3.1 Computing the Monopoly Phase

The monopoly phase begins after one of the firms reaches stage D , which can take any value between 0 and N . Therefore, we solve the monopoly problem for all $x_i \in [0, N]$, $i = 1, 2$. The successful firm's value function during the monopoly phase, V_i^M , solves the Bellman equation 2. We compute it by backward induction on states beginning at stage N and proceeding to the lower stages. At stage N , $V_i^M(N) = \Omega$ and $a_i^M(N) = 0$. Once we have computed $a_i^M(x')$ and $V_i^M(x')$ for $x' > x_i$, we can then compute the value functions $V_i^M(x_i)$ and policy functions $a_i^M(x_i)$ by using equations (2) and (3).

In addition to employing a standard value function iteration and implementing the Gauss-Seidel method for dynamic programming, see p. 418 in Judd (1998), we also occasionally use a second approach when the convergence criterion is very tight. This second approach solves a nonlinear system of first-order necessary and sufficient conditions. These conditions are necessary and sufficient given our assumption on the cost and Markov transition functions. The conditions are as follows:

$$V_i^M(x_i) = -C_i(a_i) + \beta \sum_{x'_i \geq x_i} p(x'_i | a_i, x_i) V_i^M(x'_i) \quad (7)$$

$$0 = -C'_i(a_i) + \beta \sum_{x'_i \geq x_i} \frac{\partial}{\partial a_i} p(x'_i | a_i, x_i) V_i^M(x'_i) + \lambda_i \quad (8)$$

$$0 = \lambda_i a_i \quad (9)$$

$$0 \leq \lambda_i, a_i. \quad (10)$$

To find the solution to (7)-(10), we convert it into a nonlinear system of equations that guarantees a_i to be nonnegative. For this purpose we define

$$a_i = \max\{0, \alpha_i\}^\kappa \quad \text{and} \quad \lambda_i = \max\{0, -\alpha_i\}^\kappa$$

where $\kappa \geq 3$ is an integer and $\alpha_i \in \Re$. Note that, by definition, equation (9) and inequalities (10) are immediately satisfied. Thus, the unique solution to the nonlinear system of the two equations (7) and (8) with $a_i = \max\{\alpha_i, 0\}^\kappa$ in the two unknowns $V_i^M(x_i)$ and α_i yields the optimal policy and the corresponding value function of the monopolist.⁵

3.2 Solving the Duopoly Phase by an Upwind Procedure

The duopoly game lives on a finite set of states and could be solved using the techniques of Pakes and McGuire (1994). However, we have a special structure which allows for much faster computation. Since the game is over when one firm reaches D , the monopoly phase solution provides the value for

⁵The constraint on the effort level a can only be binding when the cost function C is linear. Nevertheless we use the constrained-optimization approach involving a Lagrange multiplier even when we use strictly convex cost functions. This approach is numerically much more stable than solving the first-order conditions of the unconstrained problem.

each firm at all states (x_1, x_2) with $\max\{x_1, x_2\} = D$. The solution process for the remaining stages of the duopoly game utilizes a backward induction technique.

Figure 2 displays the critical features of equilibrium dynamics and computation. There are four possible outcomes in each period. The state could remain unchanged because no firm advances; the state could jump to the right indicating progress only by firm 1; similarly, the state could jump up indicating progress only by firm 2; and finally, the state could jump up and to the right indicating progress by both firms. The solid lines in Figure 2 display a possible sample path starting from state $(D - 3, D - 3)$. After staying in state $(D - 3, D - 3)$ for one period, the firms move through states $(D - 2, D - 3)$, $(D - 2, D - 2)$, and $(D - 1, D - 1)$ until the game ends in state $(D, D - 1)$. The information that is needed to compute equilibrium values flows in the opposite direction, indicated by the broken lines. For example, if we know the value at (D, D) , $(D - 1, D)$, and $(D, D - 1)$, then the game at $(D - 1, D - 1)$ reduces to a simple game where the only unknowns are the values and actions of each firm at $(D - 1, D - 1)$. The upwind method⁶ computes equilibrium values by traversing the nodes in a manner consistent with the direction of the broken lines until it has reached and solved the game at $(0, 0)$.

At each state (x_1, x_2) , we compute an equilibrium action pair $(\sigma_1(x_1, x_2), \sigma_2(x_1, x_2))$ and the corresponding values $(\mathbb{V}_1(x_1, x_2), \mathbb{V}_2(x_1, x_2))$ that satisfy equations (5, 6). This computational task is surprisingly difficult; we employ two different algorithms.

The first algorithm is a Gauss-Seidel iterated best reply approach. We choose a starting point of actions and values. Next, we solve the first firm's dynamic programming problem using value function iteration just as in the monopoly problem. We update the first firm's policy function and solve the second firm's dynamic programming problem. Then we update the second firm's policy function and solve the first firm's problem and continue executing these steps until convergence.

Although this Gauss-Seidel iterated best reply algorithm appears to be the natural approach for solving the stage game, it often does not converge. In particular, when the race is close (i.e., x_1 is very close to x_2) and both firms continue to invest, the algorithm typically cycles. We use a second algorithm when the two firms are close to each other. We formulate the equilibrium problem in state (x_1, x_2) as a nonlinear system of equations. The following conditions are necessary and sufficient for optimality. For $i = 1, 2$,

⁶Our solution approach for the duopoly race phase is an example of the general idea behind upwind procedures for dynamic problems. Suppose that there is a partial order \prec on the state space X with the following properties: if the state can change from x to x' , $x \neq x'$, then $x \prec x'$ and if $x \prec x'$ then there is no feasible sequence of states by which the game moves from x' to x . There is such a partial order for our game since the states x_1 and x_2 can never decline. In general, if we know that the game can only go to states in $Y \subset X$ from state $x \notin Y$ and we already know each player's value function in all states in Y , then we can directly compute the players' value functions for state x by solving a small set of equations for the equilibrium strategies in state x . We can thus compute an equilibrium for all states in $X \setminus Y$ from which the game can only move directly to a state in Y . We sweep through X in some convenient order and compute the equilibrium values for all states x .




Figure 2: Movement and information flow in the state space.

$$0 = -\nabla_i(x_i, x_{-i}) - C_i(a_i) + \beta \sum_{x'_i, x'_{-i}} p(x'_i|a_i, x_i)p(x'_{-i}|a_{-i}, x_{-i})\nabla_i(x'_i, x'_{-i}) \quad (11)$$

$$0 = -\frac{\partial}{\partial a_i}C_i(a_i) + \beta \sum_{x'_i, x'_{-i}} \frac{\partial}{\partial a_i}p(x'_i|a_i, x_i)p(x'_{-i}|a_{-i}, x_{-i})\nabla_i(x'_i, x'_{-i}) + \lambda_i \quad (12)$$

$$0 = \lambda_i a_i \quad (13)$$

$$0 \leq \lambda_i, a_i. \quad (14)$$

We transform this system of equations and inequalities into a nonlinear system of equations characterizing a Nash equilibrium at a state (x_1, x_2) with $x_i, x_{-i} < D$. We set $a_i = \max\{0, \alpha_i\}^\kappa$ and $\lambda_i = \max\{0, -\alpha_i\}^\kappa$ in equations (11) and (12) and omit the complementary slackness conditions (13) and the inequalities (14). The solutions to the resulting four nonlinear equations in the four unknowns $\nabla_i(x_i, x_{-i})$ and α_i for $i = 1, 2$, correspond to the Nash equilibrium of the stage game. Again we solve a constrained problem instead of an unconstrained problem since this choice results in a numerically much more stable procedure.

3.3 Optimal Patent Policy

The PGA maximizes its objective function W^S or W^C taking into consideration the effect of its policy (D, γ) on firms' investment. We parameterize the PGA's objective function in θ and B . We solve the dynamic equilibrium of the patent race for a large discrete set of (D, γ) pairs to find the optimal PGA policy (D^*, γ^*) . The ratio γ takes values from a discrete set $\Gamma \subset [0, \bar{\gamma}]$. We summarize all computational steps in the following algorithm.

Algorithm 1 (Computation of welfare-maximizing policy)

1. Select an objective function $W \in \{W^S, W^C\}$. Fix the parameters θ and B . Choose $\bar{\gamma}$ and a grid $\Gamma \subset [0, \bar{\gamma}]$.
2. For each $\gamma \in \Gamma$
 - (a) Set $\Omega = \gamma B$.
 - (b) Solve the monopoly problem given Ω .
 - (c) For $D = 0$, compute the expected planner surplus, $W(0, \gamma; \theta, B)$, of giving the patent monopoly to a firm chosen randomly with equal probabilities.
 - (d) For each $D \in \{1, 2, \dots, N\}$
 - i. Solve the duopoly game for $x_1, x_2 < D$.
 - ii. Compute the expected planner surplus, $W(D, \gamma; \theta, B)$
3. Find the optimal (D^*, γ^*) which maximizes $W(D, \gamma; \theta, B)$.

3.4 Optimal Policy for a Simple Case

There is one case where we can immediately derive the optimal policy. We mention this special case since it serves as a nice benchmark even though it is not robust. Suppose $\eta = 1$, $c = 1$, $\bar{\gamma} = 1.0$, and $\theta = 0$; that is, costs are linear and equal, there is no limit on the portion of B that can be transferred to the winning innovator, and the transfers cause no inefficiencies. Since $\eta = 1$, there is no advantage in having two firms working on innovation. If the PGA just gives the project to one firm, $D = 0$, and sets $\gamma = 1$, which is feasible, then the firm's profits equal social surplus and the firm chooses the social surplus maximizing innovation effort policy.

In this special case, there is no value to a race since the PGA's problem can be perfectly internalized in a firm's profit maximizing strategy. A race would speed up innovation but only through inefficiently excessive investment. Our numerical examples show that races are desirable when we make more reasonable choices for c , θ , and $\bar{\gamma}$, and when we examine consumer surplus maximizing policies.

4 Computations and Results

Except for the special parametric case discussed in Section 3.4, it is difficult to state analytical results. Thus we use the numerical procedures outlined in Section 3 to analyze optimal patent rules for a variety of parameterizations. Table 1 displays the set of parameters we use in our computations.

TABLE 1: Parameter Values	
$N \in \{5, 10\}$	number of stages of innovation
$D \in \{0, \dots, N\}$	winning stage
$B \in \{100, 1000\}$	total social benefit
$\beta \in \{0.96, 0.996\}$	discount factor
$\eta \in \{1, 1.5, 2\}$	elasticity of cost
$\gamma \in \{.02, .04, .06, \dots, 1.00\}$	possible γ choices
$\bar{\gamma} \in \{0.1, 0.3, 0.5, 1.0\}$	upper bound on the prize to total benefit ratio, Ω/B
$\theta \in \{0, 0.1, 0.25, 0.4, 1.0\}$	deadweight loss parameter
$c_1 = 1$	normalization on firm 1's cost parameter
$c \in \{1, 2, 3, \dots, 20\}$	ratio of firms' costs coefficients, c_2/c_1 .
$F(x x) = 0.5$	transition probability for unit investment

These parameter values represent a wide range of cases. We make two normalizations: $c_1 = 1$ and $F(x|x) = .5$. We examine 5- and 10-stage races because longer ones do not provide any additional insights and a shorter number of stages is not enough to display the trade-offs between alternative patent rules. The θ values are motivated by inefficiency costs of monopoly for standard demand curves and by the excess burden results in Judd (1987). We examine two values for β to model the unit of time. When $\beta = .996$, the unit of time is about a month, whereas $\beta = .96$ implies that the unit of time is about a year. The two values of B were chosen so that races are neither too short nor too long. In general, the parameter values in Table 1 are chosen to represent innovation processes lasting from several months to a few years.

4.1 A Sample Race

Figures 3 and 4 display a sample equilibrium path of motion and investment for $B = 100$, $\eta = 1.5$, $\theta = 0.4$, $\beta = 0.996$, $c = 2$, $D = 3$, and $\gamma = 0.12$; the policy parameters maximize social surplus. This particular game lasts 10 periods. Although Firm 2 has a larger average cost of investment and invests less at $t = 1$, it is lucky and gets ahead of Firm 1 in period 1. It then increases its investment level considerably from $a_2 = 0.02$ in state $x = (0, 0)$ to $a_2 = 0.37$ in state $x = (0, 1)$. Firm 1 also increases its investment effort in order to stay in the race. The firms stay in $x = (0, 1)$ for five periods until Firm 1 catches up with Firm 2 in period 7. At this stage, both firms decrease their level of

investment for one period. However, Firm 2 moves ahead again and both firms increase their effort for one period after which Firm 1 catches up. This puts the two firms one stage away from winning the race. Firm 1 increases its effort and Firm 2 slightly reduces its effort, but they continue to make large investments. Both firms reach stage 3 simultaneously at $t = 10$. Firm 2 wins the patent by a coin toss and substantially reduces its effort since it is now protected from competition.

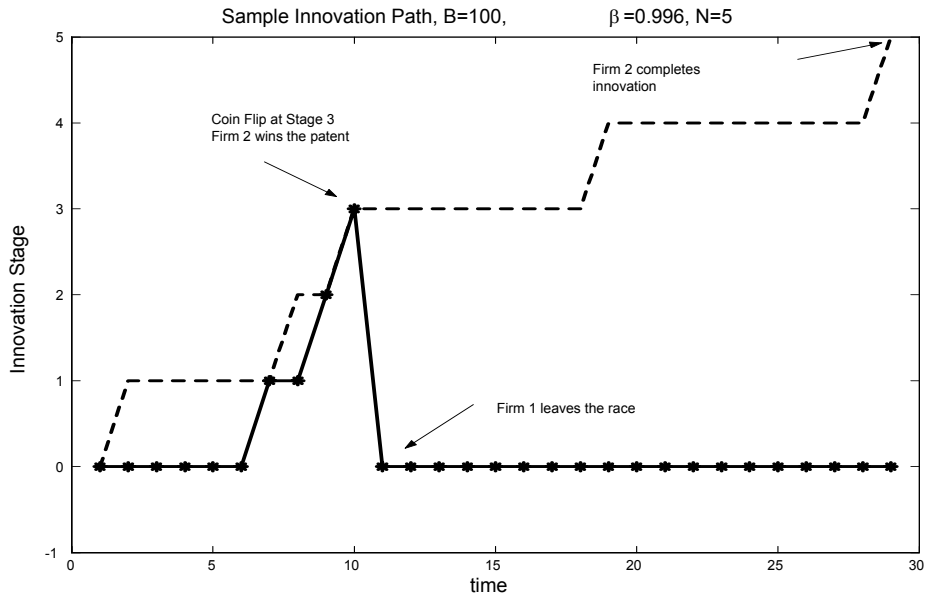


Figure 3: Sample innovation path: $c = 2$.




Figure 4: Sample innovation path investment levels: $c = 2$.

This example illustrates several points. First, the firms react strongly to each other’s movements. Second, once a firm wins the patent, innovation effort falls substantially. This tells us that effort is substantially affected by the competitive environment and concerns about duplicative investment and rent dissipation are important. Third, Firm 2 is an active competitor even though it is much less efficient. Fourth, competition is not always most fierce when the firms are tied. In general, the equilibrium behavior of this model is more complex than that in the models of Harris and Vickers (1985a, 1987), and Fudenberg et al. (1983).

4.2 Social Surplus Maximization

We now examine the case of social surplus maximization and the impact of cost heterogeneity and deadweight losses in optimal patent rules. These initial examinations assume strictly convex costs, $\eta > 1$. We then consider the case of linear costs since new equilibrium properties arise.

4.2.1 No Deadweight Loss

Figure 5 shows the optimal prize to benefit ratio γ^* , and expected discounted social surplus W^S as a function of the cost ratio of the two firms, c for a specific case with $\theta = 0$. Each line in Figure 5 corresponds to a different patent granting stage D . The maximized social surplus is the upper envelope of the three lines in Figure 5. Therefore the optimal patent granting stage is the D that corresponds to the highest line for a given cost ratio.

If the PGA maximizes social surplus and there is no deadweight loss (i.e. $\theta = 0$), the basic trade-off is between the total cost of innovation and its duration. The PGA would like firms to innovate quickly but with minimal investment. To motivate firms the PGA could set a high prize. However, a large prize also leads to fierce competition, wasteful duplication of investment and inefficient rent dissipation in a race. In Figure 5 we see that if $c = 1$, the choice of $D = 0$ and $\gamma = 1.0$ maximizes social surplus. Figure 5 shows that this result is robust to a nontrivial set⁷ of values for c . Although a coin toss may grant the patent to the less efficient firm, the resulting loss in social surplus is less than the inefficient rent dissipation during a race.

If $D = 0$, social surplus falls as c rises. At small cost ratios (less than 2.5 in Figure 5), the social surplus from races with $D > 0$ has a different pattern. For small c , social surplus decreases. The rising costs for Firm 2 result in even more rent dissipation during a race. But once the cost ratio is sufficiently large (above 1.4 in Figure 5), social surplus begins to rise. The race now serves as a mechanism to filter out the less efficient firm. In the lower stages of development, the presence of

⁷The range of c values for which the optimal policy is $(D^*, \gamma^*) = (0, 1.0)$ varies with some of the key parameters such as B , β , N and η . In Figure 3 this range is about $[1, 2.5]$.

Firm 2 motivates the more efficient Firm 1 to innovate quickly. Once Firm 1 has a sufficiently large lead⁸, Firm 2 reduces its investment level which lowers cost of duplication, and raises social surplus.




Figure 5: Social Surplus for $\bar{\gamma} = 1.0$ and $\theta = 0$

The role of the race as a mechanism to filter out the less efficient firm becomes more pronounced as the cost ratio increases. A coin toss, $D = 0$, delivers a lower surplus than a race with $D > 0$, because a patent granted to the less efficient firm leads to slow and costly innovation. A race, on the other hand, may lead to costly duplication of investment during the early stages. These costs are offset, however, when the more efficient firm takes a lead, and the laggard firm effectively drops out of the race. If the PGA chooses a small D , then competition is intense at the early stages of the race and the less efficient firm may win through luck. The PGA can discourage this by reducing the prize $\Omega = \gamma B$, but a reduction in the prize also reduces efficient firm's incentive to invest and innovates quickly after stage D^* . Thus the PGA must strike the right balance between the patent stage D and the prize level. When deadweight loss, θ , is 0, as in Figure 5, the PGA prefers to set $D^* = 5$, which enables it to filter out the less efficient firm. The PGA also chooses a large prize to

⁸A laggard firm reduces its investment considerably (effectively drops out of the race), when the probability of catching up to the leader is small and the investment cost of catching up is large. The sufficient gap between the two firms that induces such a behavior depends on the Markov process for transition from one stage to the next and the cost of investment.

provide the efficient firm with the proper incentives.

When the cost ratio increases further, the optimal γ^* increases for fixed $D > 0$ because the presence of the inefficient firm poses less of a threat to the efficient firm and so the PGA must motivate the efficient firm by giving it a larger prize. When the cost ratio becomes very large (above 15) the optimal social surplus for all values $D > 0$ converges to the same level as that for $c = 1$, $D = 0$ and $\gamma = 1.0$. Firm 2's cost of investment is very high, therefore it invests very little during the course of the race and Firm 1 can effectively act as a monopolist. In this case the PGA is indifferent among all positive values for D .

The patterns displayed in Figure 5 are robust, as stated in the following summary.

Summary 1 *The following results hold for all values of B, N, β listed in Table 1 with $\eta > 1$, $\bar{\gamma} = 1.0$ and $\theta = 0$.*

1. *When firms have similar costs the social surplus maximizing patent policy is $(D^*, \gamma^*) = (0, 1)$, that is, there is no race and the prize equals the full social benefit.*
2. *For a nontrivial race, $D \geq 1$, the optimal prize ratio γ^* is nondecreasing in the patent stage D .*
3. *As the cost ratio c rises to infinity,*
 - (a) *The investment level of the less efficient firm goes to zero and the more efficient firm proceeds as a monopolist.*
 - (b) *The PGA becomes indifferent between all positive D .*
 - (c) *The optimal prize to benefit ratio γ^* , goes to 1 for all positive D .*
 - (d) *The social surplus converges to the surplus in the case where firms have identical cost functions.*

4.2.2 Deadweight Loss

With positive θ , there is a deadweight loss, $\theta\Omega = \theta\gamma B$, associated with the patentholder's prize winnings. In this case, the choice of γ affects social surplus directly through the deadweight loss term in the objective function of the PGA. Consequently, a social surplus maximizing PGA now prefers a smaller prize. However, the trade-offs are now more complex. A small γ reduces the incentives for investment and slows down innovation. The PGA can influence the competition by varying D . The choice of D heavily depends on the degree of heterogeneity of the two firms.

Figure 6 shows the optimal prize γ^* and the social surplus W^S for a specific case with $\theta = 0.25$. When neither firm has a substantial cost advantage ($c < 1.5$ in Figure 6), competition in a race is fierce and it allows the PGA to set a small γ ($\gamma^* \in [0.1, 0.14]$) and a large D ($D^* = 5$). As the cost ratio c increases, the competition between firms is reduced because the efficient firm has a greater

cost advantage. The less efficient firm effectively drops out of the race⁹ and the more efficient firm's incentives for large investment and quick innovation are reduced. To remedy that, the PGA responds by increasing the prize at first. However, increasing the prize raises the deadweight loss and reduces the social surplus. When the cost ratio increases further, the PGA tries to induce Firm 2 to stay in the race and provide competition for Firm 1 by reducing D . For example, at cost ratios of 1.4 and 1.5, $D^* = 5$ and γ^* is 0.12 and 0.14 respectively. At cost ratios of 1.75, 2.0 and 2.5, $D^* = 4$ and γ^* is 0.14, 0.16, 0.18 respectively.




Figure 6: Social Surplus for $\bar{\gamma} = 1.0$ and $\theta = 0.25$.

Tables 2 and 3 display the optimal patent policy and the associated social surplus (as a percentage of benefit B) for a variety of θ and $\bar{\gamma}$ values. Table 2 reports the solutions for $\beta = 0.996$ and $N = 5$. The top half examines the symmetric cost case, $c = 1$, and the bottom half examines the asymmetric cost case of $c = 2$. Table 3 includes solutions from a smaller discount factor, $\beta = 0.96$, and symmetric costs.

⁹Note that with the given parameterization $C'(0) = 0$ and so the laggard prefers to invest a very small but positive amount.

		$\bar{\gamma} = 1.0$			$\bar{\gamma} = 0.5$			$\bar{\gamma} = 0.3$			$\bar{\gamma} = 0.1$		
c	θ	D^*	γ^*	W^S/B	D^*	γ^*	W^S/B	D^*	γ^*	W^S/B	D^*	γ^*	W^S/B
1	0	0	1.00	87.7	0	0.50	87.1	0	0.30	85.6	5	0.10	83.5
	0.1	0	0.34	83.0	0	0.34	83.0	0	0.30	82.9	5	0.10	82.6
	0.25	5	0.10	81.3	5	0.10	81.3	5	0.10	81.3	5	0.10	81.3
	0.4	5	0.10	79.9	5	0.10	79.9	5	0.10	79.9	5	0.10	79.9
	1.0	5	0.08	75.5	5	0.08	75.5	5	0.08	75.5	5	0.08	75.5
2	0	0	1.00	84.9	5	0.22	84.3	5	0.22	84.3	3	0.1	80.7
	0.1	5	0.20	82.4	5	0.20	82.4	5	0.20	82.4	3	0.1	79.9
	0.25	4	0.16	79.9	4	0.16	79.9	4	0.16	79.9	3	0.1	78.6
	0.4	3	0.12	78.0	3	0.12	78.0	3	0.12	78.0	3	0.1	77.3
	1.0	3	0.10	72.3	3	0.10	72.3	3	0.10	72.3	3	0.1	72.3

In all of the cases reported, social surplus and the optimal prize/benefit ratio, γ^* are decreasing in θ . When $c = 1$, firms are identical, a race is not used to filter out one of the firms; it is used to spur competition. Although a low D leads to intense competition between equal competitors during the duopoly phase, it creates much rent dissipation, followed by a slow innovation process during the monopoly phase unless the prize is large. With a higher D , competition is still fierce, but innovation is completed quickly, without the need for a high prize.

When firms have asymmetric costs, the race fills both an incentive and filtering role. For example, if $\theta = 0$ and firms are symmetric, the optimal D^* is 0 (unless γ^* is constrained by a very low $\bar{\gamma}$); there is no race. However, when firms are not symmetric, i.e. when $c = 2$, the optimal D^* can be 5 even in the case of no deadweight loss. As θ rises, D^* decreases because the PGA chooses to give innovation incentives to the more efficient firm by reducing the stages of the duopoly phase and allowing the less efficient firm to be a strong competitor for the efficient firm in the early stages of the race. Also, when a race is desired (which is when $\theta \in [0.25, 1.0]$ in Table 2), γ^* is higher in the asymmetric cost case than in the case with $c = 1$. Again, the presence of cost heterogeneity makes it less likely that a larger prize will lead to excessive duplication of effort. Although a combination of high θ and γ reduces social welfare substantially, in the case with asymmetric costs, a low γ reduces the less efficient firms' incentives to compete and the more efficient firm's incentives to innovate quickly.

Table 2 also shows what happens as the $\bar{\gamma}$ limit becomes binding on the choice of γ . The impact of $\bar{\gamma}$ on the optimal D^* is different in the symmetric and asymmetric cost cases. In the symmetric cost case, when $\bar{\gamma}$ is 0.1, D^* jumps to 5. Since a high prize is not available to provide incentives for quick innovation, investment effort is increased by forcing firms to compete for the full length of the innovation. In the asymmetric cost case, even if D is chosen to be 5, competition may not be very fierce if the efficient firm takes a strong lead and the prize is low. Therefore, with the exception in

the case of $\theta = 0$, the optimal D^* is lower (compared to the cases with higher $\bar{\gamma}$) to spur competition at the early stages of innovation, even though the duration of the monopoly phase is extended.

As the discount factor decreases, the present value of the prize decreases and thus dampens the incentive for a high investment level. Conditional on having a race of nontrivial duration the discount factor is inversely related to the optimal γ^* .

	$\bar{\gamma} = 1.0$			$\bar{\gamma} = 0.5$			$\bar{\gamma} = 0.3$			$\bar{\gamma} = 0.1$		
θ	D^*	γ^*	W^S/B	D^*	γ^*	W^S/B	D^*	γ^*	W^S/B	D^*	γ^*	W^S/B
0	0	1.00	60.9	0	0.50	59.4	5	0.26	58.6	5	0.1	46.2
0.1	5	0.24	56.9	5	0.24	56.9	5	0.24	56.9	5	0.1	45.7
0.25	5	0.22	54.6	5	0.22	54.6	5	0.22	54.6	5	0.1	45.0
0.4	5	0.20	52.5	5	0.20	52.5	5	0.20	52.5	5	0.1	44.2
1.0	5	0.16	45.8	5	0.16	45.8	5	0.16	45.8	5	0.1	41.2

Summary 2 *The following results hold for all values of B, N, β listed in Table 1 with $\eta > 1$.*

1. *The optimal prize ratio γ^* is non-increasing in the deadweight loss coefficient θ .*
2. *Social surplus is decreasing in θ .*
3. *Conditional on the presence of a race,*
 - (a) *the optimal patent stage D^* is nonincreasing in θ ,*
 - (b) *the optimal prize ratio γ^* is non-increasing in the discount factor.*

4.2.3 Linear Cost

In the case of linear cost $C'(0)$ is not zero. This implies that a firm may quit, setting $a = 0$, if it is sufficiently far behind. Table 4 reports the optimal patent policy and the resulting social surplus for linear cost functions. As in Table 2, the optimal patent policy for $\theta = 0$, $\bar{\gamma} = 1.0$, $\beta = 0.996$ and $c = 1$ is $(D^*, \gamma^*) = (0, 1)$. When θ increases, the PGA would like to reduce the prize ratio γ in order to avoid a large deadweight loss. A small γ , however, decreases the firms' incentive to innovate quickly. The PGA tries to motivate to firms to increase investment by lengthening the race and prolonging competition. With strictly convex costs the optimal patent stage is $D^* = N = 5$ because the cost from duplication of investment is less than the deadweight loss from a sufficient γ to motivate. In the linear cost case, the cost of duplication may outweigh the deadweight loss. Consequently the optimal D^* in this case is lower.¹⁰

¹⁰The optimal D^* is lower because optimal effort is less than 1 in most stages along the race in both the convex and the linear cost cases. Under alternative parameterizations the optimal effort may be greater than 1 and consequently lead to higher optimal D^* in the linear cost case.

TABLE 4: Optimal Patent Policy for $\eta = 1.0$.													
c	θ	$\bar{\gamma} = 1.0$			$\bar{\gamma} = 0.5$			$\bar{\gamma} = 0.3$			$\bar{\gamma} = 0.1$		
		D^*	γ^*	W^S/B	D^*	γ^*	W^S/B	D^*	γ^*	W^S/B	D^*	γ^*	W^S/B
1	0.0	0	1.00	86.3	0	0.50	85.9	0	0.30	85.1	4	0.1	81.8
	0.1	0	0.26	82.4	0	0.26	82.4	0	0.26	82.4	4	0.1	80.9
	0.25	3	0.12	79.7	3	0.12	79.7	3	0.12	79.7	4	0.1	79.6
	0.4	4	0.10	78.2	4	0.10	78.2	4	0.10	78.2	4	0.1	78.2
	1.0	4	0.10	72.8	4	0.10	72.8	4	0.10	72.8	4	0.1	72.8
2	0	5	0.28	85.0	5	0.28	85.0	5	0.28	85.0	1	0.1	79.0
	0.1	5	0.26	82.5	5	0.26	82.5	5	0.26	82.5	1	0.1	78.2
	0.25	2	0.20	79.6	2	0.2	79.6	2	0.20	79.6	1	0.1	76.9
	0.4	1	0.14	77.3	1	0.14	77.3	1	0.14	77.3	1	0.1	75.6
	1.0	1	0.12	70.5	1	0.12	70.5	1	0.12	70.5	1	0.1	70.4

Table 4 also displays results for $c = 2$. In this case, the PGA is concerned about the less efficient firm reaching the patent stage D first. When θ is small, the PGA can offer a high prize and set D high to motivate the firms to invest and to ensure that the more efficient firm wins. When θ is large, the deadweight loss, $\theta\Omega = \theta\gamma B$ becomes important, and the PGA becomes more concerned about the adverse effect of giving a large prize rather than the adverse effects of selecting the less inefficient firm. Thus the optimal patent stage and the optimal prize are reduced as θ increases. These results are consistent with the results from the strictly convex cost case.

4.3 Consumer Surplus Maximization

We next examine the case where the planner maximizes consumer surplus. In this case, the cost of innovation does not enter the PGA's objective function, so the PGA is only concerned about the duration of the race and the fraction of the benefit that consumers can retain. A reduction of the prize to the innovator increases consumer benefits, but slows the arrival of the innovation. One way to relieve this tension is to use races to stimulate investment.

We first examine a simple case. Figure 7 displays the optimal prize parameter γ^* and consumer surplus $W^C(\cdot)$ as a function of the cost ratio c for $\theta = 0$. Each line corresponds to a different D . The maximized consumer surplus is the upper envelope of the four lines in the figures.

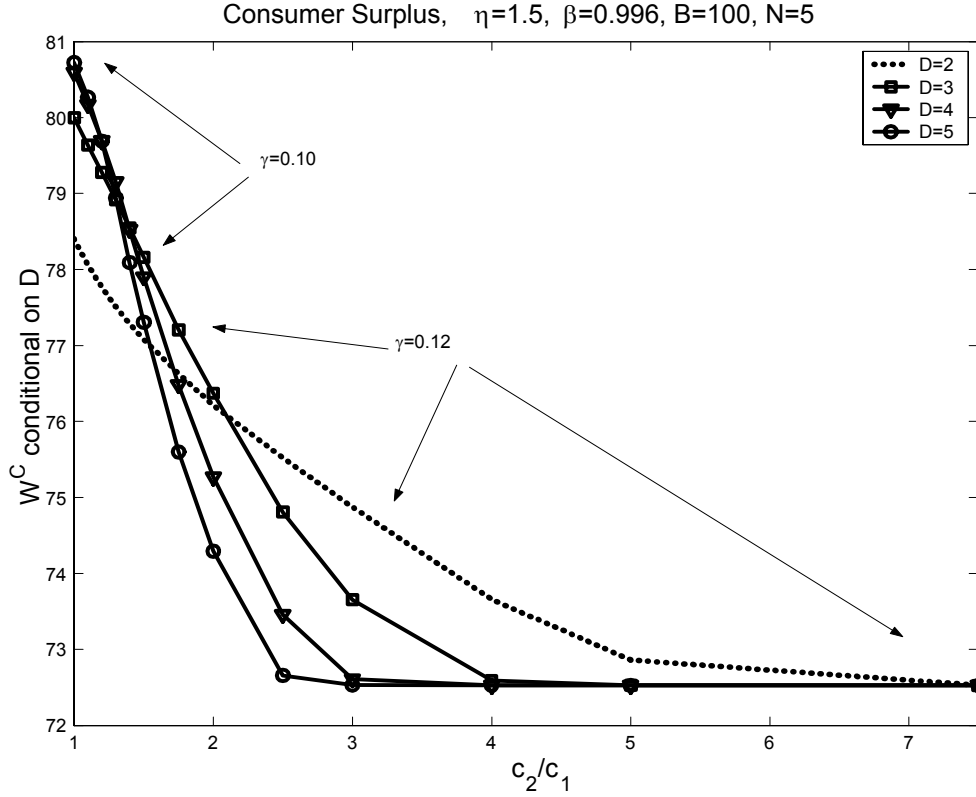


Figure 7: Consumer Surplus for $\bar{\gamma} = 1.0$ and $\theta = 0$.

Several patterns are apparent in Figure 7. Consumer surplus decreases as cost asymmetry rises. At small cost ratios the PGA can rely on the intense competition among the firms to ensure that the firms innovate quickly. Since the competition provides ample motivation for high investment levels, the PGA can set the prize-to-benefit ratio γ to be very low and the patent stage to $D = N = 5$. As c rises, the intensity of competition decreases since the inefficient firm reduces investment. The PGA remedies this by increasing γ and by choosing a lower D . These changes spur both firms to work harder in the duopoly phase without creating too much risk that the inferior firm wins. In Figure 7, γ^* increases from 0.10 to 0.12 and D^* decreases from 5 to 2. As c increases further, even a short duopoly phase is not enough to motivate the firms. Since the PGA is reluctant to increase γ , the race becomes, for all practical purposes, just a monopoly innovation process by the more efficient firm. Thus the PGA is indifferent between setting D to any value between 1 to N .

Tables 5 and 6 display results for sensitivity analysis with respect to the parameters η , N , B , and θ and confirm that the results displayed in Figure 7 are robust to changes in these parameters. The optimal γ^* is always much smaller than under the objective of social surplus maximization, and changes only slightly as deadweight loss, θ , and the cost ratio c change. Consumer surplus is decreasing in both of these parameters.

The pattern of the D^* values provides insights into the structure of our model and, in particular, highlights the difference between strictly convex costs and linear costs. As the cost of investment for Firm 2 increases, its investment level declines; Firm 2 poses less of a competitive threat to Firm 1. In order to motivate both firms, the PGA lowers the optimal patent stage D^* , but this policy only partially motivates the firms to choose higher investment levels. For the linear cost case reported in Table 6, $C'(0) > 0$ and Firm 2 reduces its investment level to zero when the cost ratio is sufficiently large. Consequently, the probability of this firm advancing is zero, and the optimal patent stage D^* can be equal to 1 as in the case of $\theta = 1$, $c = 3$. When the cost function is strictly convex, Firm 2 never chooses a zero investment level since $C'(0) = 0$, and always has a chance of reaching stage 1 before Firm 1. As a result, the optimal D^* is always greater than 1. In some cases, for example, at a cost ratio of $c = 3$ and $\theta = 0$, D^* may become as low as 2. As in the case of social surplus maximization, a further increase in the cost ratio transforms the race effectively into a monopoly and the PGA eventually becomes indifferent among all $D > 1$.

TABLE 5: Optimal Patent Policy for $\eta = 1.5, \beta = 0.996$.

		$N = 5$						$N = 10$					
		$B = 100$			$B = 1000$			$B = 1000$			$B = 100$		
θ	c	D^*	γ^*	W^C/B	D^*	γ^*	W^C/B	D^*	γ^*	W^C/B	D^*	γ^*	W^C/B
0	1	5	0.10	80.7	5	0.04	92.6	10	0.06	85.1	6	0.18	64.6
	1.5	3	0.10	78.2	5	0.04	91.7	6	0.06	83.3	3	0.18	61.8
	2	3	0.12	76.4	4	0.04	90.7	5	0.06	82.3	3	0.20	60.2
	3	2	0.12	74.9	3	0.04	89.8	4	0.08	81.4	2	0.20	59.0
1	1	5	0.06	73.7	5	0.02	90.5	10	0.04	80.8	10	0.12	54.0
	1.5	3	0.08	70.6	4	0.02	88.7	5	0.04	78.7	3	0.14	50.5
	2	2	0.08	68.7	3	0.04	87.5	4	0.04	77.6	2	0.14	48.9
	3	2	0.08	66.3	3	0.04	86.0	3	0.04	76.6	2	0.14	47.6

		$N = 5, \eta = 1$			$N = 10, \eta = 1$			$N = 5, \eta = 2$			$N = 10, \eta = 2$		
θ	c	D^*	γ^*	W^C/B	D^*	γ^*	W^C/B	D^*	γ^*	W^C/B	D^*	γ^*	W^C/B
0	1.0	5	0.02	93.6	10	0.04	87.7	5	0.02	93.6	10	0.04	85.5
	1.5	4	0.04	92.5	4	0.06	86.0	5	0.02	92.7	7	0.06	83.9
	2.0	3	0.04	91.7	3	0.06	85.2	5	0.04	91.8	6	0.06	82.9
	3.0	2	0.04	91.0	2	0.06	84.6	4	0.04	90.7	5	0.06	81.9
1	1.0	5	0.02	91.7	10	0.04	84.0	5	0.02	91.7	10	0.04	82.0
	1.5	3	0.02	89.8	3	0.04	81.9	5	0.02	90.8	7	0.04	80.1
	2.0	2	0.02	88.9	2	0.04	81.2	5	0.02	89.8	5	0.04	79.0
	3.0	1	0.02	88.1	2	0.04	80.8	4	0.02	88.4	4	0.04	77.9

The results from consumer surplus maximization can be summarized as follows.

Summary 3 *When the PGA maximizes consumer surplus the optimal patent policy exhibits the following properties for parameters listed in Table 1.*

1. *The optimal patent policy has a nontrivial race, $D^* > 0$.*
2. *The optimal prize to benefit ratio, γ^* , is smaller than when the PGA maximizes social surplus.*
3. *The expected duration of innovation process is longer due to lower investment level compared to the social surplus maximization case.*
4. *Consumer surplus is nonincreasing in the cost ratio.*
5. *For sufficiently small cost ratios, the optimal patent granting stage, D^* , is nonincreasing in the cost ratio.*
6. *As the cost ratio c rises to infinity,*
 - (a) *The less efficient firm essentially exits the race and the more efficient firm proceeds as if a monopolist.*
 - (b) *The PGA sets $D^* > 0$ but becomes indifferent between all positive D .*

5 Conclusions and Extensions

Patent races are an integral part of the R&D process, but they do not represent the complete innovation process. A firm that has been granted a patent typically needs to incur additional costs and develop the product further before it can be produced and sold. The parameters of the race –

the stage at which the patent or exclusive contract is awarded and the winning prize – are chosen by a social policymaker or a private organization to maximize its objective. Previous patent policy analyses have focussed on the nature of the prize – the length and breadth of the patent, and previous multistage race analyses have taken patent policy as given. We present an analysis of how both parameters, the innovation stage at which a patent or an exclusive contract is granted and the size of the prize, should be chosen in a simple multistage race. Thus we bridge some of the gap between the literature on patent races and the literature on optimal patent policies.

We find that there is no one dominant form. The choice between granting patents at an early stage of development versus a later stage depends on the social returns to innovation, the planner’s objective (social vs. consumer surplus), and the inefficiency costs of compensating the patent winner. The basic trade-off for a patent policy is between the speed of innovation and costly duplication of effort. In our setting, the patent race serves two purposes. First, it motivates the firms to invest and complete the innovation process quickly. When the prize causes inefficiencies, such as the monopoly grant implicit in a patent, using a race allows the planner to reduce the size of the prize and still give firms incentives to invest in innovation. Second, a race filters out inferior innovators since they cannot keep up with more efficient ones. This is important for the planner since he cannot observe firms’ costs. When the planner wants to maximize consumer surplus, the important trade-off in this case is the speed of innovation versus the prize needed to compensate the firms. If the planner maximizes consumer surplus, prizes are lower and patent stages higher compared to the social surplus maximization case.

Our model is simple but allows us to understand the fundamental issues of developing a patent policy. Also, it is straightforward to relax some of the assumptions that we made for our computations. We have already computed many examples of races where firms can advance more than one stage at a time. We did not report results from these examples in the present paper, since they give no substantial additional insights into the workings of the model. It’s equally easy to allow the probability distribution F to depend on the stage of the innovation process. Doing so would allow us to incorporate different degrees of difficulty for the various steps in the innovation process.

Our results indicate that once a firm receives protection from competition, it reduces its investment level and slows the innovation process. The PGA varies the patent granting stage and the prize to induce firms to innovate quickly. In actual patent policy, there is a time limit on how long a product is protected under a patent. If firms develop the product too late, then they may not receive any (substantial) prize. This time limit could also serve both as a filtering device and an incentive for quick innovation, and therefore the planner may not rely on a race to differentiate between firms and spur investment. In our current formulation of the problem, however, the time it takes for the firms to move from the patent-granting stage to the terminal innovation stage is short, thus the limit on innovation would not change our current results. However, it is possible to think of environments

or parameterizations where such an additional policy tool would become an important component of patent rules. Additionally, the planner may be able to stop the race once one firm is sufficiently far ahead. So far the planner ignores the distance between firms.

In addition to varying the technology of innovation and the policy tools available to the planner, the specification of informational asymmetry can also be altered. Currently there is an informational asymmetry between the planner and the firms. The firms have perfect information about each other's cost structures and innovation stages. A more realistic model may allow firms to observe each other's development stage only ever so often.

Several additional extensions are also worth mentioning. In the development process of a new product firms often face two uncertain issues. Typically it is unclear a priori how many research steps are necessary for the development of a new product. Also, a new technology may quickly become obsolete if a better, different technology is soon to be developed. Both issues could be modelled in our environment. The uncertainty about the final product stage could be incorporated with a probability distribution for the stage N , possibly one that is updated as the R&D process progresses. A possible extension of our model may also be more explicit about the market structure after the patent is granted or after the final product has been developed. For example, we may allow firms to merge or buy each other's services once the race has been terminated.

6 Appendix

Proof of Proposition 1. We present the proof of this proposition for the case of strictly convex costs. The proof easily extends to the linear cost case, but it gets messy due to the possibility of corner solutions. In the trivial case $\Omega = 0$ we have $V_i^M(x_i) = 0$ and $a^*(x_i) = 0$ for all $x_i \in \{0, 1, \dots, N\}$. Thus, we assume throughout the proof that $\Omega > 0$. The proof proceeds in four steps. First, we prove that there exists a solution to the Bellman equation. Second, we show that the value function is nondecreasing in the state. Third, we prove that there exists a unique optimal policy function. Finally, we show that the policy function is nondecreasing in the state.

Firm i 's monopoly problem is a dynamic programming problem with discounting that satisfies the standard assumptions for the existence of a solution, see Puterman (1994, Chapter 6) or Judd (1998, Chapter 12). The state space is finite. The discount factor satisfies $\beta < 1$. The cost function $C_i(\cdot)$ is continuous and thus bounded on the compact effort set A . The transition probability function $p(x_i'|\cdot, x_i)$ is also continuous on A for all $x_i \in \{0, 1, \dots, N\}$. Therefore, there exists a unique solution V_i^M to the Bellman equation and some optimal effort level $a^*(x_i)$ for each stage $x_i \in \{0, 1, \dots, N\}$.

Fix a state $x_i < N$ and an optimal effort level $a^*(x_i)$. The value $V_i^M(x_i)$ satisfies the equation

$$V_i^M(x_i) = \frac{-C_i(a^*(x_i)) + \beta p(x_i + 1 | a^*(x_i), x_i) V_i^M(x_i + 1)}{1 - \beta p(x_i | a^*(x_i), x_i)}.$$

Since $C_i(\cdot)$ is nonnegative, $\beta < 1$, and $V_i^M(x_i + 1) \geq 0$ it follows that $V_i^M(x_i) \leq V_i^M(x_i + 1)$.

For the remainder of the proof we make use of the special form of the transition probability function p . Without loss of generality we assume that F is independent of the state x_i and write $F(x_i | x_i) = F < 1$. Under all our assumptions ($\Omega > 0$, $C(0) = 0$, $C'(0) = 0$, and $p(x_i | x_i, a_i) = F^{a_i}$) it holds that $V_i^M(x_i) > 0$ and $a^*(x_i) > 0$ for all $x_i \in \{0, 1, \dots, N\}$. Note that the optimal effort level is always in the interior of the set A . Given the value function V_i^M , a necessary (and sufficient) first-order condition for the optimal effort level is

$$F^a \beta \ln F (V_i^M(x_i) - V_i^M(x_i + 1)) - C_i'(a) = 0.$$

This equation must have a least one solution according to the first step of this proof. The second derivative of the function on the left-hand side equals $F^a \beta (\ln F)^2 (V_i^M(x_i) - V_i^M(x_i + 1)) - C_i''(a) < 0$. Hence, there is a unique optimal effort $a^*(x_i)$.

Given the value $V_i^M(x_i + 1)$, the optimal effort $a^*(x_i)$ and value $V_i^M(x_i)$ must be the (unique) solution of the following system of two equations in the two variables a and V , respectively,

$$\begin{aligned} V(1 - \beta F^a) - \beta(1 - F^a)V_i^M(x_i + 1) + C(a) &= 0 \\ F^a \beta \ln F (V - V_i^M(x_i + 1)) - C_i'(a) &= 0 \end{aligned}$$

An application of the Implicit Function Theorem reveals that both variables in the solution are nondecreasing functions of the value $V_i^M(x_i + 1)$. The Jacobian of the function on the left-hand side at the solution equals

$$J = \begin{bmatrix} F^a \beta (\ln F)^2 (V - V_i^M(x_i + 1)) - C''(a) & 0 \\ -F^a (\beta \ln F) & 1 - \beta F^a \end{bmatrix}.$$

The gradient of the function on the left-hand side with respect to the parameter $V_i^M(x_i + 1)$ equals

$$\begin{pmatrix} -\beta(1 - F^a) \\ -F^a \beta \ln F \end{pmatrix}.$$

The Implicit Function Theorem yields

$$\begin{pmatrix} \frac{\partial V}{\partial V_i^M(x_i + 1)} \\ \frac{\partial a}{\partial V_i^M(x_i + 1)} \end{pmatrix} = -\frac{J}{D} \begin{pmatrix} -\beta(1 - F^a) \\ -F^a \beta \ln F \end{pmatrix} \geq 0,$$

where $D = (1 - \beta F^a)(F^a \beta (\ln F)^2 (V - V_i^M(x_i + 1)) - C''(a)) < 0$ is the determinant of the Jacobian. The value function V_i^M is nondecreasing in the state x_i and $a^*(x_i)$ is nondecreasing in the value $V_i^M(x_i + 1)$. Thus, the function a^* is nondecreasing in the state. ■

Proof of Theorem 1. For a given patent policy (D, γ) the strategy functions σ_i^* , $i = 1, 2$, constitute a Markov perfect equilibrium if they simultaneously solve equations (6). The proof is by backward induction. If $x_i = D$ for some i , then an optimal strategy pair $\sigma_i^*(x_i, x_{-i})$, $i = 1, 2$, and a pair of value functions \mathbb{V}_i , $i = 1, 2$, trivially exist. It is now sufficient to prove that for any state $(x_1, x_2) \in X$ with $x_i < D$, $i = 1, 2$, there exists a pure strategy Nash equilibrium (a_1^*, a_2^*) . To prove the existence of such an equilibrium we define a continuous function f on a convex and compact set such that any fixed point of this function is a pure strategy Nash equilibrium.

Given a state $(x_1, x_2) \in X$ with $x_i < D$, $i = 1, 2$, and values $\mathbb{V}_i(x_i + 1, x_{-i})$, $\mathbb{V}_i(x_i, x_{-i} + 1)$, $\mathbb{V}_i(x_i + 1, x_{-i} + 1)$ from the states that can be reached from (x_1, x_2) in one period. As in the proof of Proposition 1 we assume without loss of generality that the transition probability distribution is independent of the state and we write $F(x_i | x_i) = F$, $i = 1, 2$. We define a function f on a domain $S \equiv A \times [0, \gamma B] \times A \times [0, \gamma B]$. Choose an arbitrary element $(\hat{a}_i, V_i, \hat{a}_{-i}, V_{-i}) \in S$. Consider the equation

$$0 = -C'_i(a_i) \left(\frac{1}{F} \right)^{a_i} + \beta \ln F \cdot \left(F^{\hat{a}_{-i}} (V_i - \mathbb{V}_i(x_i + 1, x_{-i})) + (1 - F^{\hat{a}_{-i}}) (\mathbb{V}_i(x_i, x_{-i} + 1) - \mathbb{V}_i(x_i + 1, x_{-i} + 1)) \right)$$

with the one unknown a_i . If $\delta \equiv F^{\hat{a}_{-i}} (V_i - \mathbb{V}_i(x_i + 1, x_{-i})) + (1 - F^{\hat{a}_{-i}}) (\mathbb{V}_i(x_i, x_{-i} + 1) - \mathbb{V}_i(x_i + 1, x_{-i} + 1))$ is positive, then this equation has no solution. In this case we define $\hat{a}_i = 0$. If $\delta \leq 0$ then this equation has a unique solution $\hat{a}_i \geq 0$ (since $-C''_i(a_i) \left(\frac{1}{F} \right)^{a_i} + C'_i(a_i) \ln F \left(\frac{1}{F} \right)^{a_i} < 0$ for all $a_i \in A$). Note that $\hat{a}_i \in A$. We define $f_{i,1}(\hat{a}_i, V_i, \hat{a}_{-i}, V_{-i}) = \hat{a}_i$. Note that δ is continuous in V_i . An application of the Implicit Function Theorem shows that $f_{i,1}$ is continuous in V_i .

Next define \hat{V}_i by

$$\hat{V}_i = \frac{1}{1 - \beta F^{\hat{a}_i} F^{\hat{a}_{-i}}} \left(-C(\hat{a}_i) + \beta \left(F^{\hat{a}_i} (1 - F^{\hat{a}_{-i}}) \mathbb{V}_i(x_i, x_{-i} + 1) + (1 - F^{\hat{a}_i}) F^{\hat{a}_{-i}} \mathbb{V}_i(x_i + 1, x_{-i}) + (1 - F^{\hat{a}_i}) (1 - F^{\hat{a}_{-i}}) \mathbb{V}_i(x_i + 1, x_{-i} + 1) \right) \right).$$

Note that $\hat{V}_i \in [0, \gamma B]$ and define $f_{i,2}(\hat{a}_i, V_i, \hat{a}_{-i}, V_{-i}) = \hat{V}_i$. Clearly, the function $f_{i,2}$ is continuous.

In summary, we have defined a continuous function $f = (f_{1,1}, f_{1,2}, f_{2,1}, f_{2,2}) : S \rightarrow S$ mapping the convex and compact domain S into itself. Brouwer's fixed-point theorem implies that f has a fixed point $(a_1^*, V_1^*, a_2^*, V_2^*) \in S$. By construction of the function f this fixed point satisfies the equations (5) and (6). This completes the proof of the existence of a pure strategy Nash equilibrium in the state (x_1, x_2) . ■

References

- [1] Dasgupta, P. “Patents, Priority and Imitation or, the Economics of Races and Waiting Games.” *The Economic Journal*, Vol. 98 (1988), pp. 66–80.
- [2] Dasgupta, P. and Stiglitz, J. “Industrial Structure and the Nature of Innovative Activity.” *Economic Journal*, Vol. 90 (1980a), pp. 266–293.
- [3] Dasgupta, P. and Stiglitz, J. “Uncertainty, Market Structure and the Speed of Research.” *Bell Journal of Economics*, (1980b), pp. 1–28.
- [4] Denicolo, V. ”Two-stage patent races and patent policy.” *RAND Journal of Economics*, Vol. 31 (2000), pp. 488–501.
- [5] Denicolo, V. ”The Optimal Life of a Patent when Timing of Innovation is Stochastic.” *International Journal of Industrial Organization* Vol. 17 (1999), pp. 827–846.
- [6] Doraszelski, U. “An R&D Race with Learning and Forgetting,” Northwestern University, mimeo, 2001.
- [7] Fudenberg, D., Gilbert, R., Stiglitz, J. and Tirole, J. “Preemption, Leapfrogging, and Competition in Patent Races.” *European Economic Review*, Vol. 22 (1983), pp. 3–31.
- [8] Gilbert, R. and Shapiro, C. “Optimal Patent Length and Breadth.” *RAND Journal of Economics*, Vol. 21 (1990), pp. 106–112.
- [9] Grossman, G.M. and Shapiro, C. “Optimal Dynamic R&D Programs.” *RAND Journal of Economics*, Vol. 17 (1986), pp. 581–593.
- [10] Grossman, G.M. and Shapiro, C. “Dynamic R&D Programs.” *Economic Journal*, Vol. 97 (1987), pp. 372–387.
- [11] Harris, C. and Vickers, J. “Perfect Equilibrium in a Model of a Race.” *Review of Economic Studies*, Vol. 52 (1985a), pp. 193–209.
- [12] Harris, C. and Vickers, J. “Patent Races and the Persistence of Monopoly.” *Journal of Industrial Economics*, Vol. 33 (1985b), pp. 461–481.
- [13] Harris, C. and Vickers, J. “Racing with Uncertainty.” *Review of Economic Studies*, Vol. 54 (1987), pp. 1–21.

- [14] Hörner, J, “A Perpetual Race to Stay Ahead.” Kellogg School of Management, Northwestern University, mimeo, 2001.
- [15] Hopenhayn, H. and Mitchell, M. “Innovation Variety and Patent Breadth.” *The RAND Journal of Economics*, Vol. 32 (2001), no.1.
- [16] Judd, K. L. “Closed-loop equilibrium in an innovation race with experience,” Northwestern University, 1985a.
- [17] Judd, K. L. “Closed-Loop Equilibrium in a multistage Innovation Race.” The Center for Mathematical Studies in Economics and Management Science, Kellogg School of Management, Northwestern University, 1985b.
- [18] Judd, K. L. *Numerical Methods in Economics*. Cambridge, MA: MIT Press, 1998.
- [19] Kamien, M. and Schwartz, N. *Market Structure and Innovation*. New York, NY: Cambridge University Press, 1982.
- [20] Klemperer, P. ”How Broad should the Scope of a Patent Protection Be?” *RAND Journal of Economics* Vol. 21 (1990), pp. 113–130.
- [21] Lee, T. and Wilde, L. “Market Structure and Innovation: A Reformulation.” *Quarterly Journal of Economics*, Vol. 94 (1980), pp. 429–436.
- [22] Nordhaus, W. D., *Invention, Growth and Welfare*, Cambridge: M.I.T. Press, 1969.
- [23] Pakes, A. and McGuire, P. “Computing Markov-perfect Nash Equilibria: Numerical Implications of a Dynamic Differentiated Product Model.” *RAND Journal of Economics*, Vol. 25 (1994), pp. 555–589.
- [24] Puterman, M. L. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. New York, NY: John Wiley & Sons, Inc., 1994.
- [25] Reinganum, J. “Dynamic Games of Innovation.” *Journal of Economic Theory*, Vol. 25 (1981), pp. 21–41.
- [26] Reinganum, J. “A Dynamic Game of R&D: Patent Protection and Competitive Behaviour.” *Econometrica*, Vol. 50 (1982), pp. 671–688.
- [27] Reinganum, J. “The Timing of Innovation: Research, Development, and Diffusion.” in *Handbook of Industrial Organization*, Vol. 1 eds. R. Schmalensee and R. D. Willig. Amsterdam: The Netherlands: Elsevier, 1989.