Optimal Rules for Patent Races*

Kenneth L. Judd
Hoover Institution
Stanford University
judd@hoover.stanford.edu

Karl Schmedders
Kellogg School of Management
Northwestern University
k-schmedders@northwestern.edu

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

October 8, 2003

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 study the optimal combination of prize and minimal accomplishment necessary to obtain a patent in a dynamic multistage 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. We show that races are undesirable only when efficiency costs are low, firms have similar technologies, and the planner maximizes social surplus. However, in all other circumstances, 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.

*The authors are grateful to Ulrich Doraszelski, Mort Kamien, Ariel Pakes, Mark Satterthwaite, Christopher Sleet, and seminar audiences at Stanford University, MIT, Harvard University, the University of Texas at Austin, Northwestern University, the Clarence W. Tow Conference on Computational Economics at the University of Iowa, the 8th CEF conference in Aix-en-Provence, the First Brazilian Workshop of the Game Theory Society at the University of Sao Paulo, Washington University in St. Louis, the University of Bielefeld, and at IESE/Barcelona for helpful discussions.

1 Introduction

Patent systems (i.e. patent rules and rewards) and procurement contracts use races to spur innovation. Firms involved in innovation race to be the first to develop a new product and obtain a patent and monopoly rights to sell that product, or to be the first to develop a product satisfying a buyer's specifications. It's been long argued (see Wright (1983) for an earlier criticism) that patent systems may be suboptimal mechanisms because they spark races and generate wasteful duplication of effort. As Loury (1979) noted, races also have offsetting benefits: increased investment leads to quicker innovation. The use of races as a part of an optimal patent system relies crucially on their costs and benefits. This paper provides a model of optimal patent rules that endogenizes the choice of races in designing incentives for innovation.

We model races as multistage stochastic games between heterogenous firms that possess perfect information about each others' innovation state and cost structures. Firms proceed through several stages of progress (e.g. rough idea, blue print, prototype), with the final stage culminating in successful innovation and a marketable product. For a given patent policy, our dynamic, stochastic innovation race resembles those studied by Fudenberg et al. (1983) and Harris and Vickers (1985a,b, 1987). However, we endogenize patent policy by embedding this game into the problem of a central patent authority who selects the rules that govern the races.

We consider a patent system parametrized by a finite set of instruments. The patent authority, who has imperfect information about the competing firms' cost structures, chooses their values. First, it chooses when to award the patent, namely, it chooses the innovation stage at which a patent or an exclusive contract is awarded to a firm and the race is terminated. This represents the minimal accomplishment necessary to obtain a patent and sets the length[1] of a patent race. Second, the authority chooses rules such as patent length and breadth that determine how the benefits from the completed innovation is allocated between the winner and the rest of society. Firms race as long as no firm has achieved the minimal accomplishment level, but as soon as one has met the requirements of the patent rules, it proceeds with exclusive rights to develop the product.

Our policy instruments differ from those employed in conventional discussions of patent policy which focus on the optimal duration and breadth of patent protection. They assume that a firm does not receive the patent for a particular innovation until its R&D process is complete. Thus the potential length of a patent race is assumed to be equal to the length of the innovation process.[2] In our setting, we distinguish between the stage a patent is granted and the end of the innovation

[1] Throughout the paper we use the term short (long) race to describe a low (high) minimal accomplishment necessary for a patent. It is possible for a "long" race to lead to "slow" innovation and for a "short" race to lead to "quick" innovation. The terms slow and quick are used to describe the time it takes for firms to complete the innovation.

[2] This is true for papers that look at R&D processes involving a single invention, which is the case we examine in this paper.

process. This distinction allows us to evaluate the desirability of races and to analyze the effects of the race length on the pace and cost of innovation preceding and following a patent. If races are undesirable because they generate wasteful duplication of effort under certain circumstances, then the central authority awards the patent at the initial stage; there is no race and the innovation process is completed by a single firm without competition.

Our assumption concerning the separation between the time a patent is awarded and the time that particular innovation is completed is consistent with actual innovation processes. Firms often receive patents for a product in its early development stages, bear significant expenditures afterwards and reap financial benefits only when R&D is finished and the product is marketed. The history of innovation in the U.S. clearly shows that patents are often granted long before the patentholder has a marketable product. This feature is seen in many of the case histories of important inventions examined in Jewkes et al. (1969). For example, the first patents for xerography were granted many years before the first copy machine, and far more money was spent on development of the transistor after the patent was granted than before. When a patent should be granted represents a fundamental question in the design of patent policy, and has important implications for the costs and benefits of innovation. Thus it is very important to understand its role in optimal patent systems and how it can be used to balance the trade-offs a patent authority faces.

In our environment, these trade-offs are the following. The patent authority must weigh the benefits from early completion of the innovation process and early introduction of the product against the total discounted cost of innovation. A long race requiring a high level of accomplishment combined with a large prize stimulates firms to work hard, to innovate quickly, but leads to wasteful duplication and rentseeking. A shorter race reduces the aggregate investment level, but leads to poor intertemporal resource allocation; each firm works very hard to win the patent and but then the winner proceeds more slowly to finish the process and introduce the good. Reducing the prize level reduces all investment effort and delays the arrival of a socially valuable product. A priori, it is not obvious which effect dominates in choosing the optimal rules and to what degree each instrument should be used.

We show that both policy instruments, minimal accomplishment and rewards to the winner, will generally be used to spur competition and innovation, and the optimal mix depends on the primitives of the model and on the costs, benefits and limitations associated with using one or both of these instruments. In most circumstances, it is optimal for the patent authority to grant a patent after some progress has been made by the firms. In other words, races are desirable. They serve two important purposes in our model. First, the patent authority uses races to motivate innovators when the prize alone cannot, due to inefficiencies or limitations, provide adequate incentives. Second, a patent race serves as a filtering device. Since the patent authority can verify partial success but cannot observe an individual firm's efficiency, a race is used to increase the chance that the patent is rewarded to the most efficient innovator.

An important factor that influences the optimal mix of the two policy instruments concerns the preferences of the patent authority. Our model is explicit about these preferences. We consider two different specifications: social and consumer surplus. We examine optimal patent policy when the patent authority'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 multiple suppliers that must engage in innovation to produce the desired product.

Our results indicate that in an environment with no inefficiency costs of transfers to the winner[3], if the patent authority is maximizing social surplus, the race length is typically increasing in the degree of cost heterogeneity amongst firms. A patent is awarded after substantial progress has been made in innovation. If firms are homogenous, on the other hand, a shorter race is preferred, to avoid excessive investment; in fact, in the limit, the patent authority would just grant an immediate patent monopoly to one of the firms at random at the initial stage. In an environment with inefficient transfer mechanisms, the trade-offs faced by the patent authority are more complex. These inefficiencies constrain the patent authority's ability to use prizes (either cash or monopoly rights) to spur innovation. Consequently, races are the instruments used more heavily as incentives for innovation, but unlike the case with no efficiency costs, longer races are preferred when firms are homogenous and shorter races are chosen otherwise. This result overturns the conventional wisdom that when firms are likely to compete fiercely, i.e., when they possess identical technologies, races are undesirable because they avoid excessive investment. Our analysis shows that this is true only when there is very little constraint on the prizes a patent authority can give. When there are limitations present, races are preferred because they fulfill the role of prizes in spurring innovation. However, when giving incentives to innovate requires a short race, the effectiveness of the filtering role of races is compromised. In a short race, the probability of a firm with high R&D costs winning the patent is higher than in a long race. In this case, the adverse effect in question is not excessive investment; it is the tighter constraints on the PGA's ability to identify the right firm.

The results from consumer surplus maximization with and without inefficiency costs are similar to the results from social surplus maximization constrained by such costs. In all of these cases, races of nontrivial duration are part of the optimal patent rules. A consumer surplus maximizing patent authority chooses long races when firms are homogenous and short races when they differ in their cost structures. Since consumer surplus is defined by the residual social benefit after the prize to the winner has been allocated, the optimal patent rule is to avoid using the prize as incentive for investment and to rely on races and the competition they create to facilitate innovation.

In general, in environments with higher inefficiency costs and externalities that restrict the re-

[3]Such as monopoly distortions or distortionary cost of taxes used to finance these transfers.

wards to the patentholder, it is optimal for patents to be awarded earlier in the innovation process. This is more pronounced if firms are heterogenous; they face smaller minimal accomplishment levels, which lead to competitive, but short races of nontrivial duration.

The paper is organized as follows. Section 2 presents our model of a patent race and describes the patent authority's problem. Section 3 discusses the relevant literature and elaborates on the differences between our approach and those contained in related papers. Section 4 presents the results. Section 5 discusses optimal patent policies when the patent authority has little information about the costs and benefits of innovation. Section 6 concludes the paper with a discussion of possible extensions of the present analysis. All proofs and a detailed description of our numerical method are included in the Appendices.

2 A Model of Multistage R&D Processes

We assume two kinds of infinitely-lived agents: innovator firms and a social planner, which we call "the patent granting authority" (hereafter, 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 planner organizes a race among potential seller-innovators and grants special rights to the winner. 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.

Firms compete in a multistage stochastic innovation race in discrete time. In each period, they have perfect information about each other's cost structure and position and choose their investment levels simultaneously. We let B denote the potential value to society from the invention. 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. We expand on our interpretation of γ in Section 2.3 below.

The PGA chooses the rules that govern races prior to any investment by the firms. This is similar to Congress passing patent legislation. We consider two different specifications of the PGA's preferences: social and consumer surplus. The premise is that the PGA cannot continuously monitor all races and has imperfect information about the cost structures of firms involved. Thus it is assumed to have a limited set of policy instruments. The PGA maximizes its objective (social surplus or consumer surplus) by choosing γ, the fraction of potential social benefits that goes to the patentholder, and $D \in \{0, \ldots, N\}$, the patent-granting stage. If $D = 0$, then there is no race. 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. In the game, it formally corresponds to the PGA giving the patent at random to one of the firms. The key assumption is that in this case each firm has equal chance of winning without having made any investment.

For the remainder of the 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 and then 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 $n + 1$[4], 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 $p\left(x|a,x\right)$. We choose a probability structure so that innovation resembles search and sampling. Let $F(x|x) = p\left(x|1,x\right)$, that is, $F\left(x|x\right)$ is the probability that there is no change in the state if $a = 1$. For general values of a we assume

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

This specification is analogous to 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.[5]

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

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

[4] We 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.

[5] This 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.

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 is awarded the patent and becomes the only innovator. We refer to the subsequent innovation stages as the monopoly phase and denote it by $X^M = \{D, D+1, \ldots, N\}$. Prior to the monopoly phase the position of the two firms is described by $x = (x_1, x_2)$. We refer to the set of states before the patent is granted as the duopoly phase and denote it by $X^D = \{(x_1, x_2) | x_i \in \{0, ..., D\}, i = 1, 2\}$. Since we employ backward induction to solve for the equilibrium of the game, we first solve for the monopoly phase and then for the duopoly phase.

2.2.1 Monopoly Phase

Firm i precedes as a monopoly after it receives the patent. We formulate firm i's monopoly problem recursively. At the terminal stage N, the innovation process is over and firm i receives a prize of Ω. In stages D through $N-1$, it 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

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

The policy function a_i^M of a firm i monopolist is defined by

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

Proposition 1 *Firm i's monopoly problem at state $x_i \in \{0, 1, \ldots, N\}$ 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 A. ■

2.2.2 Duopoly Phase

We formulate the competition between the firms before stage D as a duopoly game. In the analysis of this game we restrict attention to Markov strategies. A pure Markov strategy $\sigma_i : X^D \to A$ 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 $x = (x_1, x_2) \in X^D$. We use the conventional notation that x_{-i} (a_{-i}) denote the state (action) of firm i's opponent. If at least one of the firms has reached the patent stage D, firm i's value function is defined as follows:

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

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

$$\mathbb{V}_i(x) = \max_{a_i \in A} \left\{ -C_i(a_i) + \beta \sum_{x_i', x_{-i}'} p(x_i'|a_i, x_i) p(x_{-i}'|a_{-i}, x_{-i}) \mathbb{V}_i(x_i', x_{-i}') \right\}, \; x_1, x_2 < D. \tag{5}$$

The optimal strategy functions of the firms must satisfy

$$\sigma_i(x) = \arg\max_{a_i \in A} \left\{ -C_i(a_i) + \beta \sum_{x_i', x_{-i}'} p(x_i'|a_i, x_i) p(x_{-i}'|a_{-i}, x_{-i}) \mathbb{V}_i(x_i', x_{-i}') \right\}, \; x_1, x_2 < D. \tag{6}$$

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. *Given the value functions $\mathbb{V}_i$, and the strategy function of his opponent, the strategy function σ_i^* for player i solves equation (6), $i = 1, 2$.*

A Markov perfect equilibrium always exists.

Theorem 1 *There exists a Markov perfect equilibrium.*

Proof. See Appendix A. ■

2.3 The Constraints and Goals of the Patent Granting Authority

The PGA faces many constraints in its choices and has several possible objectives. We now discuss our treatment of the informational constraints of the PGA, the efficiency costs of its policy instruments, and its possible objectives.

2.3.1 Informational Constraints of the PGA

In essence, the patent authority must choose race rules that apply across a broad range of industries and products, rather than rules tailored to specific firms or innovations. Our view is that the PGA has some beliefs about the possible social benefits and costs of innovation, but that these beliefs are rather diffuse. Specifically, we assume that the PGA faces some distribution of benefits B and the cost ratio c_2/c_1 of the two innovators (with $c_1 = 1$ as a normalization). However, there is very limited empirical documentation of these two distributions, so we have very little precise information about them. Consequently, we proceed in two stages. First, we assume that the PGA knows the exact, ex-ante social value of the innovation, B, and the innovation technologies of the two firms but not their identity. 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.[6] We derive the optimal patent race rules conditional on knowing the cost ratio c_2/c_1 and social benefits B. Then we analyze several parameterizations of this conditional optimal patent policy problem and precisely identify the trade-offs the PGA faces as the cost structures of the firms and the social value of the product change. In particular, we identify conditions under which races are desirable. In Section 5, we turn to the analysis of the PGA's problem when it only knows the distribution of ex-ante social values and innovation costs. In this analysis, we consider a variety of distributions for cost ratios and social benefits, and examine optimal policy rules that maximize ex-ante expected consumer/social surplus. These policies represent the basic patent rules that apply across industries and products. However, we find that they are essentially averages of the policy rules derived from the analysis that conditions on specific values for the cost ratio and the social benefit. Thus, we focus on the basic trade-offs that determine optimal rules conditional on specific cost ratios and social benefits from innovation.

2.3.2 Efficiency Costs of the PGA's Policy Instruments

Any PGA has tools for encouraging innovation, but those tools are often imperfect and cause inefficiencies. For example, a patent transfers wealth to the patentholder only at the cost of inefficient monopoly pricing. Also, prizes issued by a government are typically financed by distortionary taxation. Our analysis of the PGA's problem takes into account the constraints these inefficiencies impart on the PGA's choices.

We assume that the PGA has two policy tools: γ, the share of benefits going to the patentholder,

[6] It 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.

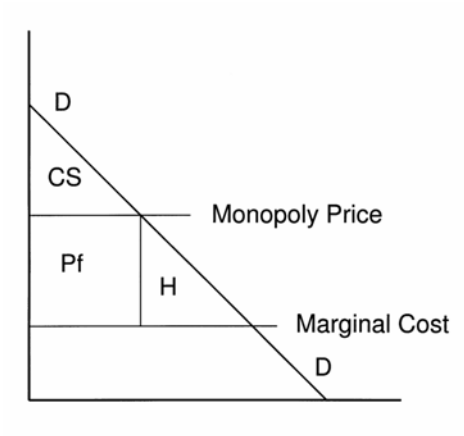

Figure 1: Division of Social Value

and D, the stage at which a patent is awarded. 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 represented by γ. The parameter γ is meant to represent many features of a patent. For example, γ is small if the patent life is short or if it has small breadth. We do not make any distinction among these features of a patent in this paper since the pace of R&D is determined by the prize γB and not by its decomposition.

To motivate the PGA's loss function, we examine how the allocation of social benefits B is affected by patent policies. Figure 1 displays the per-period allocation before the patent has expired. 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 it 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 γ. 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, $\overline{\gamma} \leq 1$, on the PGA's choice of γ.

In Figure 1, the deadweight loss H represents the social cost of monopoly profits in the patent system.[7] More generally, we assume that the deadweight loss is proportional to the profits received by the innovator, and is equal to $\theta\Omega = \theta\gamma B$ for some $\theta \geq 0$. For example, $\theta = 0.5$ in Figure 1. This linear specification for the 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 0.1 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 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.

While any patent is in effect, the firm receives profits, the consumers receive some benefit, but some of B is wasted in the transfer process. We assume that the patentholder's profits are γB but that the deadweight loss due to inefficiencies is $\theta\gamma B$, leaving consumers with $B - \gamma B - \theta\gamma B$.

2.3.3 Possible Objectives of the PGA

The PGA chooses both γ and D to maximize some objective. Most analyses assume that social surplus is the appropriate objective, but it is not clear that this should always hold. For example, in a procurement context one suspects that the PGA does not care about the profits of the participants in the race. Also, some argue that Congress should choose patent policy to maximize social surplus, but it is also imaginable that Congress will consider maximizing consumers' welfare if the median voter is more of a consumer than a producer. Therefore, we consider two different possible objectives for the PGA. The first objective we examine is total social surplus, $W^S(D, \gamma; \theta, B)$, which equals the present discounted value of the social benefit, B, minus the deadweight losses of transfers to the patentholder, $\theta\Omega$, and minus the total investment expenditures of all innovators in the patent race. Our second specification assumes the PGA maximizes the present discounted value of consumer surplus, $(1 - \gamma)B - \theta\Omega$, denoted by $W^C(D, \gamma; \theta, B)$. The precise statements of the PGA's objectives are stated in Appendix B.3, along with our numerical method that solves for the optimal policy.

[7] Price 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.

3 Related Literature

This paper contributes to a growing literature that studies the implications of R&D competition for the design of optimal patent policy and bridges a gap between two distinct lines of research. The first of these lines focuses on R&D competition, taking patent policy as given. The second endogenizes patent policy, but largely abstracts from the R&D competition that precedes the award of the patent.

Early contributions to the first line of research include Kamien and Schwartz (1982), Loury (1979), Lee and Wilde (1980), Reinganum (1981, 1982) and Dasgupta and Stiglitz (1980a,b). In these models, the probability that a firm successfully obtains a patent at each date depends only on the 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)). This first generation of models was subsequently extended by Fudenberg et al. (1983) and Harris and Vickers (1985a,b, 1987) to incorporate learning or experience effects. Throughout, patent policy was taken as given. In contrast, contributions to the second line of research, including Nordhaus (1969), Klemperer (1990), Gilbert and Shapiro (1990), O'Donaghue, Scotchmer, Thisse (1998), Denicolo (1999, 2000), and Hopenhayn and Mitchell (2001), Llobet, Hopenhayn and Mitchell (2002), endogenize the policymaker's choice of patent length and/or breadth. However, these papers largely abstract from the R&D races that precede the award of a patent. Rather, they focus upon the tradeoff between providing current incentives for innovations and minimizing post-innovation distortions.

For a given patent policy, our dynamic, stochastic innovation race resembles those studied by Fudenberg et al., Harris and Vickers. However, we endogenize patent policy by embedding this race into the problem of a social planner who chooses when a patent is rewarded and the winning firm's prize once it has completed R&D and is producing a marketable product. We analyze how changes in these features of the patent policy affect the equilibrium behavior of firms in the dynamic race.

Our paper is complementary to a growing body of work that studies the effects of patent policy on the strategic interaction between competing firms in alternative environments. Amongst these are Scotchmer and Green (1990, 1995). The former paper focusses upon the novelty requirement and disclosure aspects of patent law and their implications for the pace of innovation. The focus of the latter paper is on the role of patent licensing in environments with successive innovations. Both papers consider a truncated R&D phase, with two stages of innovation only. We consider a multistage innovation race which enables us to make a distinction between the stage at which a patent can be awarded and the stage at which an innovation is successfully completed. This distinction, which is absent from previous work on patent races and optimal patent policy, allows us to characterize the pace of innovation before and after one of the firms obtains the patent.

Our main focus is on the management of innovation races by two instruments: timing of patent awards and prizes. Thus, in contrast to Scotchmer and Green (1990) which assume successful firms receive the full social benefit of their product, prizes are endogenously determined in our setting. We also assume that prizes are constrained by externalities from monopoly distortions or deadweight loss of taxes used to finance them. Explicit considerations of these externalities allow us to study the effectiveness of races in spurring innovation when the ability of the patent authority to use prizes is limited. We do, however, abstract from possible transactions and cooperation between innovators modeled in Scotchmer and Green (1995), even where such cooperation may be socially beneficial. In an extension of this paper, we have examined equilibrium with technology trades and found conditions under which they will occur. We have found that the results for optimal patent rules considered in this paper are not significantly affected by these trades.

Our patent instruments are different from those utilized in multiple innovation settings considered in O'Donaghue, Scotchmer and Thisse and Llobet, Hopenhayn and Mitchell. In O'Donaghue et.al., the focus is on increasing profits for successive improvements in innovation while minimizing monopoly distortions. The patent authority is assumed to have several instruments that control the length and the breadth of the patent. The optimal mixture of these instruments is not identified, instead two types of policies, with different lengths and breadths are compared in terms of the effects they have on the diffusion and costs of innovation. In Llobet et.al., the focus is on optimal rewards for successive innovation and a system of buyouts and licensing that can implement these rewards. Again, the policy instruments are the length and the scope of the patent. Our model is about a single innovation, not a sequence of innovations. Our main emphasis is on the role of races: During the innovation phase of a single good, at what stage should a patent-type protection be given and races be terminated? We are not concerned with the decomposition of the prize to the winner into patent breadth or depth. Patent breadth issues, protection from imitation, length of patents are all represented by the prize in our setting. We choose our parameters of the model so that the imposition of a reasonable time limit on the duration of innovation from the time a firm gets the patent to the time it reaches the final stage is not binding.

Patents races are not the only mechanism for spurring innovation. Research tournaments, where contestants compete to find the innovation with the highest value to the sponsor and receive a prespecified prize if successful, can be and are used to achieve a similar goal. Research tournaments are particularly useful when research inputs are unobservable and research outcomes cannot be verified by courts. Taylor (1995), and Fullerton and McAfee (1999) are amongst the papers that study such tournaments, the former in an environment with identical firms, the latter with firms of heterogenous ability. Innovation races and research tournaments differ both in institutional and model details. In a single, well-defined innovation race, the quality requirement is fixed, the time of innovation is variable. The focus in on the pace of innovation and the competition between the firms. In a research tournament, quality is variable, the terminal date, on the other hand, is fixed. In research tournaments, the emphasis is on the quality of the product, not on the pace of innovation. In the McAfee-Fullerton model, an entry auction filters out less efficient firms, in our model, filtering is achieved by varying the quality requirement. Our goal is to endogenize the choice of races and to study the changes in the innovation pace and intensity when quality requirements and prizes are chosen optimally. Thus we find patent races to be the most natural environments[8] in which to achieve this goal. Although we concentrate on two policy instruments, which are central to any patent system, our framework and computational method is suited to analyze other instruments such as, prizes that depend on the relative performance of the firms, intermediate subsidies contingent on innovation stages, and technology trades between firms. Our main objective is to identify the basic trade-offs a patent authority faces, so that the choice of additional instruments that may remedy some of the inefficiencies that arise in the innovation race is not made arbitrarily. One of the larger goals of our paper is to make progress in identifying the right set of instruments for the patent authority. We elaborate more on some of these instruments in our conclusion.

4 Results and Computations

We begin our analysis with a simple, canonical example with linear technology, identical firms, and no deadweight loss associated with the price. In this case, we also set $\bar{\gamma} = 1.0$; the prize to the winner can equal the full social value of the product. This is the case that's been extensively studied in the innovation race literature. In this special case, there is obviously no value to a race since the PGA's problem can be perfectly internalized in a firm's profit maximizing strategy. 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 investment level. A race would speed up innovation but only through inefficiently excessive investment. In the following sections, we report the changes in optimal patent rules as we move away from this special case and look at more realistic and interesting examples with heterogenous firms, non-linear technology and inefficiencies associated with rewards and show that the policy implications of this simple canonical example are typically not robust to reasonable perturbations in the environment. In particular, races are part of optimal patent rules and serve multiple purposes under different parameterizations of our model. Table 1 displays the set of parameters we use in our numerical computations.

[8] See Scotchmer (1999) for an environment in which the patent system is optimal. She shows that if a direct mechanism cannot use ex-post information on value or costs, the only feasible incentive mechanisms are patent renewals systems with fees.

Table 1: Parameter Values	
$N \in \{5, 10\}$	number of stages of innovation
$B \in \{100, 250, 1000\}$	total social benefit
$\beta \in \{0.96, 0.996\}$	discount factor
$\eta \in \{1, 1.5, 2\}$	elasticity of cost
$\overline{\gamma} \in \{0.1, 0.3, 0.5, 1.0\}$	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
$c \in \{1, 1.25, ..., 4.75, 5, 6, ..., 20\}$	ratio of firms' costs coefficients, c_2/c_1. ($c_1 = 1$)
$F(x\|x) = 0.5$	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 report detailed results from a 5-stage race only. Races with more stages do not provide any additional insights or change our qualitative results. The θ values are motivated by inefficiency costs of monopoly for standard demand curves and by the excess burden results in Judd (1987). 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. Optimal choices for γ are restricted to lie in a discrete set $\{0.02, 0.04, 0.06, ..., 1.0\}$.

4.1 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.

4.1.1 No Deadweight Loss

Figure 2 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, when $\eta = 1.5$ and $\theta = 0$. Each line in Figure 2 corresponds to a different patent granting stage D. The maximized social surplus is the upper envelope of the three lines in Figure 2. 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 the firms, the PGA could set a high prize or a long race. However, a large prize or a long race could lead to fierce competition, wasteful duplication of investment and inefficient rent dissipation. As the degree of heterogeneity between the firms changes, so does the balance in this trade-off and the optimal mix of the two instruments of the patent policy: race and prize.

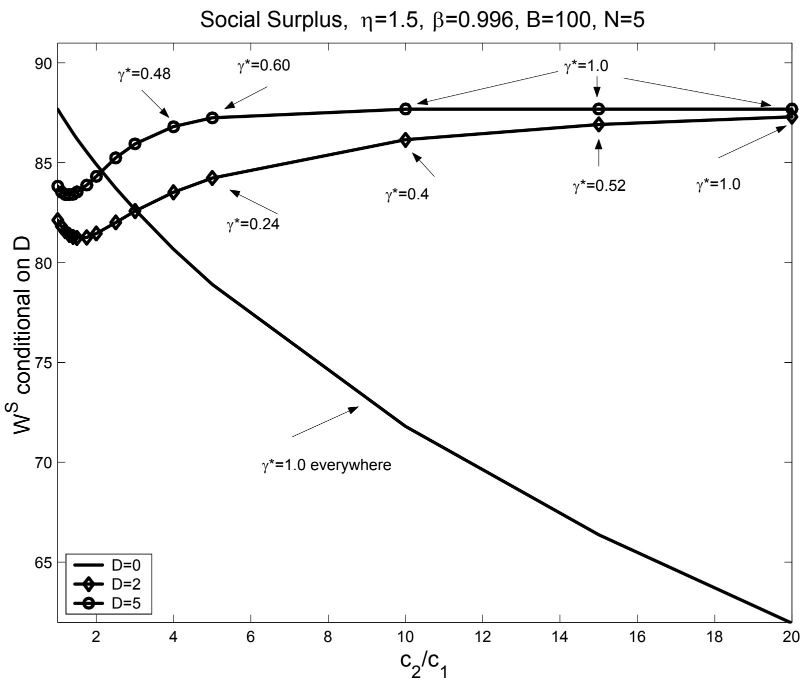

Figure 2: Social surplus for $\bar{\gamma} = 1$ and $\theta = 0.0$.

When firms are identical in their technology, our computations confirm the results from our canonical example: the optimal patent rules comprise of no race and full social surplus as a prize in this case. The PGA awards the patent by a coin toss at the very beginning of the innovation phase, i.e., $D^* = 0$. Figure 2 shows that this result is robust to a set[9] of small values of 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. This result is overturned as c continues to rise: now the cost of assigning the less inefficient firm the patent by a coin toss exceeds the cost of overinvestment in a race. Thus the optimal patent rule changes to a full length race ($D^* = 5$) and a lower prize ($\gamma^* = (0.48, 0.60)$ for $c = (4, 5)$ respectively) and social surplus begins to rise. At these intermediate cost ratios, the purpose of the race is twofold: to spur competition in the initial stages of innovation and to filter out the less efficient firm (Firm 2). At the beginning of the race, the presence of Firm 2 motivates Firm 1 to innovate quickly. Once Firm 1 has a sufficiently large lead[10], Firm 2 reduces its investment level which lowers cost of duplication, and raises social surplus. If the PGA chooses

[9] The 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 $[1, 2.5]$.

[10] A 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. This result is analogous to the well known $\varepsilon - preemption$ result in patent races.

a small D, then competition is intense at the early stages of the race and Firm 2 may win through luck. The PGA can discourage this by reducing the prize $\Omega = \gamma B$, but a large reduction in the prize also lowers Firm 1's incentive to invest and innovate quickly after stage D^*. Thus the optimal mix of races and prizes reflect the careful balance between the benefits of quick innovation, information extraction and inefficiency costs of competition and prizes. When the deadweight loss of a prize, θ, is 0, as in Figure 2, the PGA prefers sets a long race ($D^* = 5$), which enables it to filter out the less efficient firm.

When the cost ratio increases further, the optimal γ^* increases for fixed $D > 0$ because the presence of Firm 2 poses less of a threat to Firm 1, so the PGA must motivate Firm 1 by giving it a larger prize. When the cost ratio becomes very large, the optimal social surplus for all values $D > 0$ converges to the same level as that for $c = 1$, $D = 0$ and $\gamma = 1.0$. The reason is straightforward. 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 2 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 $\eta > 1$, $\overline{\gamma} = 1.0$ 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.*

4.1.2 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, all else equal, a social surplus maximizing PGA now prefers a smaller prize.

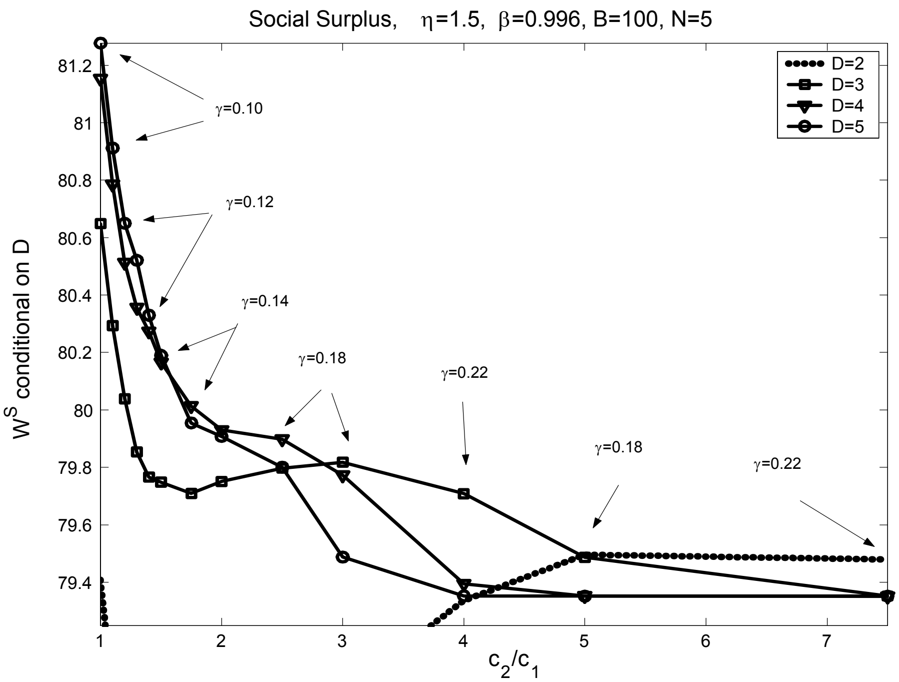

Figure 3: Social Surplus for $\bar{\gamma} = 1.0$ and $\theta = 0.25$.

Figure 3 shows the optimal patent rules (D^*, γ^*) and the social surplus W^S for $\theta = 0.25$. When neither firm has a substantial cost advantage ($c < 1.5$ in Figure 3), competition in a race is fierce and therefore a large prize is not necessary to provide incentives for quick innovation. Thus the PGA set a small prize ($\gamma^* \in [0.1, 0.14]$) and a long race ($D^* = 5$). As the cost ratio c increases, the competition between the two firms is reduced because the ability of the less efficient firm (Firm 2) to compete with Firm 1 is reduced. Firm 2 effectively drops out of the race[11] and Firm 1's incentives for large investment and quick innovation are reduced. To remedy that, the PGA responds by increasing the prize at first, which 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^*. Once the cost ratio reaches 5, even a short race can filter out the less efficient firm, thus social surplus does not decline further, despite a rise in the prize.

As it is clear from the policies displayed in Figure 3, the trade-offs the PGA faces are more complex in the presence of deadweight loss and firm heterogeneity. With $\theta = 0$, the PGA could use prizes to balance the benefits of quick innovation and total cost of investment. Its ability to vary the prize to spur innovation and competition is more constrained, the higher the deadweight loss. Thus the PGA is compelled to use races for quick innovation. A long race (high D) can stimulate

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

investment and competition, when firms are alike. If firms are sufficiently heterogeneous, long races must be coupled with a large prize to motivate the inefficient firm to compete. Large prizes are very costly in this environment, unlike the previous case with $\theta = 0$, consequently a shorter race is chosen to motivate Firm 2 by increasing its chances of winning the patent. However, this has an adverse effect; a short race is a less effective mechanism for filtering out Firm 2, thus social surplus declines as the cost ratio increases and races become shorter. This adverse effect persists until Firm 2's R&D costs are so large that it does not participate even in a short race. At that point, a short race can filter Firm 2 out, but works less effectively as an incentive device: a decrease in competition allows Firm 1 to reduce its investment level and leads to slower innovation. The comparative statics results reported in Table 2 confirm the intuition from the single example depicted in Figure 3.

Table 2: Optimal Patent Policy for $\beta = 0.996$

		$\bar{\gamma} = 1.0$			$\bar{\gamma} = 0.5$			$\bar{\gamma} = 0.1$		
c	θ	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	5	0.10	83.5
	0.25	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
	1.0	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	3	0.10	80.7
	0.25	4	0.16	79.9	4	0.16	79.9	3	0.10	78.6
	0.4	3	0.12	78.0	3	0.12	78.0	3	0.10	77.3
	1.0	3	0.10	72.3	3	0.10	72.3	3	0.10	72.3

In all of the cases reported, social surplus, the optimal prize/benefit ratio, γ^* and race length, D^* are nonincreasing in θ, for $\theta > 0$. When $c = 1$, firms are identical, a race is not used to filter out Firm 2; 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 $c = 2$, races are shorter, and social surplus is smaller.

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.

Another interesting comparative static is with respect to the discount factor. Our numerical exercises[12] have shown that as the discount factor decreases, the present value of the prize decreases and that 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^*: higher prize levels are needed to motivate the firms.

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.*

We solve our benchmark model also with linear cost of investment. There are two reasons for considering linear costs. First, in the case with strictly convex costs, since the marginal cost of effort is zero, firms invest very small, but positive amounts even when they are behind. Although they effectively drop out of the race, this behavior does not correspond to a true exit form the race. With a linear cost function, $C'(0)$ is not zero, thus a firm that is sufficiently behind would set its investment level to 0 and exit out of the race. Second, in the linear cost case, merging the R&D technologies of the firms would not decrease the overall cost of duplication of effort. Despite these important differences, the optimal patent rules with linear costs are very similar to the optimal rules with strictly convex costs when the PGA maximizes social surplus. Consequently, we do not provide detailed results of this case. We revisit the linear cost case when we discuss results from consumer surplus maximization.

4.2 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 down the arrival of the innovation. One way to relieve this tension is to use races to stimulate investment.

[12] For brevity, we do not include these results in the paper. Detailed numerical results can be obtained from the authors on request.

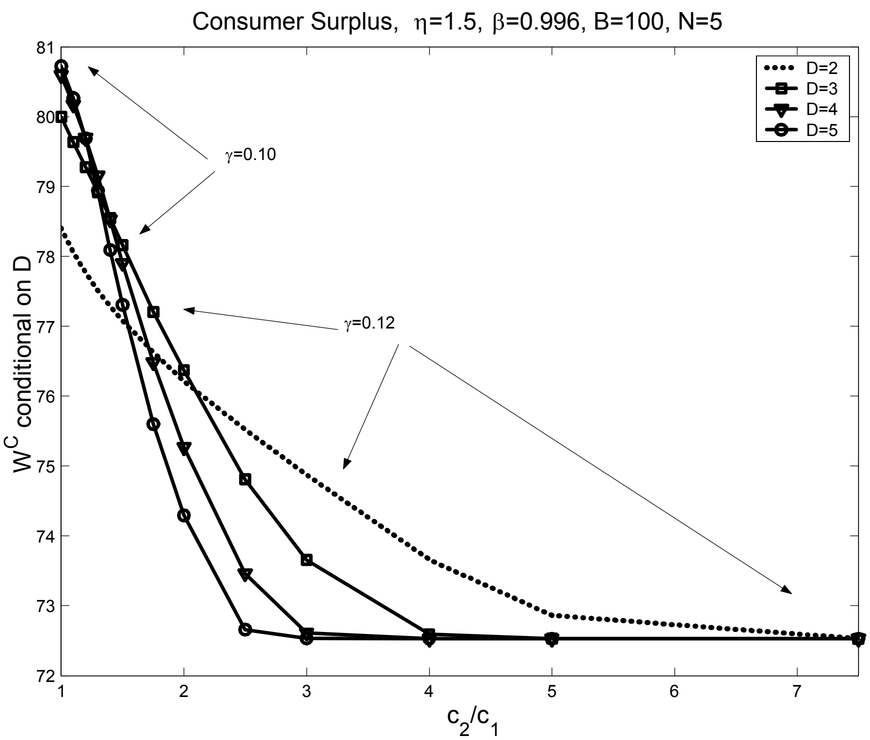

Figure 4: Consumer Surplus for $\bar{\gamma} = 1.0$ and $\theta = 0.0$.

Figure 4 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.

Several patterns are apparent in Figure 4. Consumer surplus decreases as the 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 4, 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.

Table 3 displays results for sensitivity analysis with respect to the parameters η, N, B, and θ and confirm that the results displayed in Figure 4 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. When the cost of innovation is linear in investment effort and the cost ratio is sufficiently large, Firm 2 reduces its investment level to zero and exits the race. Consequently, the probability of this firm advancing is zero, and the optimal patent stage D^* is set to 1. 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 $D \geq 1$.

TABLE 3: Optimal Patent Policy for $B = 100, \beta = 0.996$.

		$\eta = 1.5$			$\eta = 1.0$		
θ	c	D^*	γ^*	W^C/B	D^*	γ^*	W^C/B
0	1	5	0.10	80.7	4	0.10	81.0
	1.5	3	0.10	78.2	1	0.14	78.3
	2	3	0.12	76.4	1	0.12	77.9
	3	2	0.12	74.9	1	0.12	77.9
1	1	5	0.06	73.7	4	0.10	72.0
	1.5	3	0.08	70.6	1	0.10	68.9
	2	2	0.08	68.7	1	0.10	68.9
	3	2	0.08	66.3	1	0.10	68.9

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 Average Patent Rules

In Section 4, we assumed that the PGA knows the ex-ante social value of the innovation, B, as well as the cost ratio between the firms. We have identified the trade-offs the PGA and the firms face as the patent policy, innovation costs and benefits change. Motivated by our discussion in Section 2.3, we now consider the case where the PGA's information about the social value of the invention and the firms' technologies is poor; it only knows the distributions of social values and cost ratios and their support, but not their exact values. Thus, it must set rules which apply to a large set of races. Specifically, we assume that the PGA's beliefs about B are given by the probability density function $g(B)$, that its beliefs about $c = c_1/c_2$ are represented by the density $f(c)$, and that its beliefs across B and c are independent. Given these beliefs, a social surplus-maximizing PGA maximizes the expected discounted social surplus, $\sum_{c,B} W^S(D, \gamma; \theta, B) f(c) g(B)$, and a consumer surplus-maximizing PGA maximizes the expected discounted consumer surplus, $\sum_{c,B} W^C(D, \gamma; \theta, B) f(c) g(B)$.

In order to study this problem, we need to specify $f(c)$ and $g(B)$. We do not aim to execute a carefully calibrated exercise since the necessary data is not available. However, we do want to compute some "average" patent rules, illustrate the ease with which we can incorporate this into our analysis, and demonstrate the robustness of our previous results to this more general case. Therefore, we use the little data available to construct interesting examples. Pakes (1986) provides some documentation on the benefits of innovation for some European countries, and shows that their distribution is highly skewed: most innovations have very little or no social value and a few innovations have very large values. However, these are ex-post realized values on innovations. Since the patent rules we analyze are chosen before any social value is realized, the relevant data for our model is the ex-ante distribution of values held by firms when they enter a patent race. Nevertheless, we assume that Pakes's empirical evidence represents an approximation for the distribution of ex-ante social values, so we use highly skewed distributions for B in our numerical results.

The following table displays the two distributions, B_1 and B_2, for social values we use in our computations. B_1 and B_2 differ in their supports. The first row in Table 4 displays the support of B_1 and the second row the support for B_2. The third row presents the probabilities for the possible values of B_i. For example, the probability that $B_1 = 10$ is 0.5 and the probability that $B_2 = 100$ is 0.5.

Table 4: Distribution of B

Support for B_1	10	32.5	57.5	85	120	160	210	277.5	380	600
Support for B_2	100	325	575	850	1200	1600	2100	2775	3800	6000
$\Pr(B)$	0.5	0.25	0.125	0.0625	0.0315	0.016	0.008	0.004	0.002	0.001

We use four possible distributions for the cost ratio c. We have little data on this, so we examine four cases. We assume that $c_1 = 1$; this is a normalization. First, we examine two possible uniform distributions over a finite set of possible values denoted $Z = \{z_1, .., z_9\}$. The first, denoted U_1 assumes $z_i = 1 + (i - 1).25$. The second, denoted U_2 assumes $z_i = i$. We look at both cases since they represent different degrees of heterogeneity in costs and different lack of information for the PGA. We do not want the results to strongly depend on the uniform specification. Therefore, we also consider two triangular distributions for c. More precisely, these distributions assume that the probability that $c = z_i$ is $(10 - i)/45$. We look at two possibilities for $Z = \{z_1, .., z_9\}$. The first, denoted T_1, is $z_i = 1 + (i - 1).25$, and the second, denoted by T_2, is $z_i = i$.

With these probability densities, we now compute the PGA's optimal policies. Table 5 reports our numerical results for the social surplus maximizing policy when $\theta = 0$, $\bar{\gamma} = 1$. The two rightmost columns represent the two possible beliefs about B, and the four bottom rows represent the four possible beliefs about c. As we move down the table, the mean and variance of the belief about c increases and as we move right the mean and variance of the belief about B increases. These results confirm the generality of our previous insights. When the PGA maximizes expected social surplus, full length races (i.e. $D^* = N$) are part of the optimal policy if there is sufficient variability in firms' costs. As we saw before, in these cases races are useful to filter out the inefficient firms and motivate the efficient firm. However, the optimal γ is far less than one when there are races, reflecting our earlier finding that small γ's are desirable to avoid excessive duplication of effort. As we move right in Table 5, the range of possible values for B increases and we see a tendency towards no race and towards setting $\gamma = 1$. This reflects that fact that when B is large enough, it is preferable for some firm to proceed with the innovation quickly and efficiently, so there is no race and the firm that is awarded the patent receives all social benefits. We find that no race is desirable under B_2, unless the cost heterogeneity is so large that a race is needed to filter out inefficient firms. We have also computed the optimal policies for $\theta > 0$; as θ increases, races become more desirable, for the same reasons cited in Section 4.

Table 5: Social Surplus: Average Policies (D^*, γ^*)			
	B :	B_1	B_2
c			
T_1		(0,1)	(0,1)
U_1		(0,1)	(0,1)
T_2		(5,0.38)	(0,1)
U_2		(5,0.40)	(5,0.16)

We next compute the policies that maximize expected discounted consumer surplus. Table 6 reports the results when $\theta = 0$, $\bar{\gamma} = 1$. As before, full length races (i.e. $D^* = N$) with small γ are part of the optimal policy if firms' costs are not too variable and the benefits from innovation are large. In these cases, the competitive efforts of the firms are sufficient to provide consumers with benefits even when the prize γ is small. In particular, as the mean and variance of firms' costs increases, races are used to filter out the inefficient firms but only after their presence has motivated the efficient firms to work hard. The resulting duplication of effort is of no concern to consumers.

Table 6: Consumer Surplus: Average Policies (D^*, γ^*)			
	B	B_1	B_2
c			
T_1		(3,0.18)	(5,0.02)
U_1		(3,0.20)	(5,0.04)
T_2		(3,0.20)	(3,0.06)
U_2		(2,0.20)	(3,0.06)

The results in Tables 5 and 6 present a few examples, but show that the results from the conditional analyses in Section 4 are robust to the more general case where the PGA must choose rules that apply over a wide variety of races.

6 Extensions and Conclusions

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. We present an analysis of how the two parameters of the race – when the patent or exclusive contract is awarded and the winning prize – should be chosen in a simple multistage race.

We find that races of nontrivial duration are part of an optimal policy under most circumstances. The choice between short and long races depends on the social returns to innovation, the planner's objective (social vs. consumer surplus), and the inefficiency costs of compensating the patent winner.

In general, in environments with high inefficiency costs and externalities that restrict the rewards to the patentholder and firm heterogeneity, it is optimal for patents to be awarded early in the innovation process, but not at the very initial stage. Thus races are short, but not of trivial duration. 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 it cannot observe firms' costs. When the planner maximizes consumer surplus, the important trade-off is the speed of innovation versus the prize needed to compensate the firms. In this case, prizes are lower and patent stages higher compared to the social surplus maximization case.

We show that in an environment with inefficient transfer mechanisms, longer races are preferred when firms are homogenous and shorter races are chosen otherwise. This result overturns the conventional wisdom that when firms are likely to compete fiercely, i.e., when they possess identical technologies in a simultaneous-move race, short races are preferable because they avoid excessive investment. Our analysis shows that this is true only when there is very little constraint on the prizes a patent authority can give. When there are limitations present, longer races are preferred because they fulfill the role of prizes in providing incentives for innovation.

Our model allows us to understand the fundamental issues of developing a patent policy and identifying the complex trade-offs a patent authority faces. The environment we consider is a simple one, but our subsequent work indicates that the results are robust to many possible extensions. For example, one immediate extension is to consider races where firms can advance more than one stage at a time. We computed many such examples; they do not provide any substantial additional insights into the workings of the model. We have also studied cases where the technology of investment, i.e. the distribution F, depends on the stage of the innovation process. Again, no additional insights in terms of the trade-offs a patent authority faces were delivered by the modified technologies.

Another interesting extension is to allow firms to trade their technologies. It is straightforward to allow firms in our model to negotiate technology trades at each stage, similar to the trades examined in Green and Scotchmer (1995). In the context of our model, the technology leader may want to sell its technology to the laggard. We have studied this extension and found it to have no significant impact on the results for optimal patent policies.

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. However, in all of the examples we computed, we chose parameters so that the time it takes for the firms to move from the patent-granting stage to the terminal innovation stage is short. Thus, the time limit of a patent would not significantly change any result.

It may be possible to devise other additional policy instruments that may remedy some of the inefficiencies that arise in the innovation race. One of the contributions of this paper is to identify the trade-offs the patent authority and firms face as the two fundamental features of patent policy-when a patent is granted and its associated prize- change, so that the choice of additional instruments is not made arbitrarily.

A Proofs

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 $x_i \in \{0, 1, \ldots, N\}$. 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 $x_i \in \{0, 1, \ldots, N\}$. Therefore, there exists a unique solution V_i^M to the Bellman equation and some optimal effort level $a^*(x_i)$ for each stage $x_i \in \{0, 1, \ldots, N\}$.

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 $x_i \in \{0, 1, \ldots, N\}$. 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

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

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

$$\begin{pmatrix} \frac{\partial V}{\partial V_i^M(x_i+1)} \\ \frac{\partial a}{\partial V_i^M(x_i+1)} \end{pmatrix} = -\frac{1}{D} \begin{bmatrix} F^a\beta(\ln F)^2(V - V_i^M(x_i+1)) - C''(a) & 0 \\ -F^a(\beta \ln F) & 1-\beta F^a \end{bmatrix} \begin{pmatrix} -\beta(1-F^a) \\ -F^a\beta \ln F \end{pmatrix} \geq 0,$$

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 $a^*(x_i)$ in nondecreasing in the value $V_i^M(x_i+1)$. Thus, the function a^* in nondecreasing in the state. ■

Proof of Theorem 1. we present again the proof for the case of strictly convex costs. For a given patent policy (D, γ) the strategy functions $\sigma_i^*, i = 1, 2$, constitute a Markov perfect equilibrium if they simultaneously solve equations (**??**). 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 are 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

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

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 $\acute{a}_i = 0$. If $\delta \leq 0$ then this equation has a unique solution $\acute{a}_i \geq 0$ (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 $\acute{a}_i \in A$. We define $f_{i,1}(\hat{a}_i, V_i, \hat{a}_{-i}, V_{-i}) = \acute{a}_i$. 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 $\acute{V}_i$ by

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

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

In summary, we have defined a continuous function $f = (f_{1,1}, f_{1,2}, f_{2,1}, f_{2,2}) : S \to S$ 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 (**??**) and (**??**). This completes the proof of the existence of a pure strategy Nash equilibrium in the state (x_1, x_2). ■

B 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.

B.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.[13]

B.2 Solving the Duopoly Phase by an Upwind Procedure

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

[13] The 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.

of the duopoly game utilizes a backward induction technique. 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 $(D-1, D-1)$.

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 conditions (5, 6). This computational task is surprisingly difficult; a Gauss-Seidel iterated best reply approach, a natural choice in such dynamic games that solves each firm's problem sequentially and updates their best responses to each other's actions, typically does not converge in our setting. Consequently we employ an alternative algorithm. 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 $i = 1, 2$,

$$0 = -\mathbb{V}_i(x_i, x_{-i}) - C_i(a_i) + \beta \sum_{x_i', x_{-i}'} p(x_i'|a_i, x_i) p(x_{-i}'|a_{-i}, x_{-i}) \mathbb{V}_i(x_i', x_{-i}') \tag{11}$$

$$0 = -\frac{\partial}{\partial a_i} C_i(a_i) + \beta \sum_{x_i', x_{-i}'} \frac{\partial}{\partial a_i} p(x_i'|a_i, x_i) p(x_{-i}'|a_{-i}, x_{-i}) \mathbb{V}_i(x_i', x_{-i}') + \lambda_i \tag{12}$$

$$0 = \lambda_i a_i \tag{13}$$

$$0 \leq \lambda_i, a_i. \tag{14}$$

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 $\mathbb{V}_i(x_i, x_{-i})$ 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.

B.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. 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:

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

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

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

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 $(D^*,\ \gamma^*)$ that maximizes $W^S(D,\gamma;\theta,B)$ given (θ,B). The consumer surplus maximizing patent policy is a pair $(D^*,\ \gamma^*)$ that maximizes $W^C(D,\gamma;\theta,B)$ given (θ,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* $D \in \{1, 2, \ldots, N\}$

      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* $W(D, \gamma; \theta, B)$.

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] Fullerton,R. and McAfee, P. "Auctioning Entry into Tournaments." *Journal of Political Economy*, Vol. 107, (1999), pp. 573-605.

[9] Gilbert, R. and Shapiro, C. “Optimal Patent Length and Breadth.” *RAND Journal of Economics,* Vol. 21 (1990), pp. 106–112.

[10] Green, J. and Scotchmer, S. " Novelty and Disclosure in Patent Law." *RAND Journal of Economics,* Vol. 21 (1990), pp. 131–146.

[11] Green, J. and Scotchmer, S. " On the Division of Profit in Sequential Innovation" *RAND Journal of Economics,* Vol. 26 (1995), pp. 20–33.

[12] Grossman, G.M. and Shapiro, C. “Optimal Dynamic R&D Programs.” *RAND Journal of Economics,* Vol. 17 (1986), pp. 581–593.

[13] Grossman, G.M. and Shapiro, C. “Dynamic R&D Programs.” *Economic Journal,* Vol. 97 (1987), pp. 372–387.

[14] Harris, C. and Vickers, J. “Perfect Equilibrium in a Model of a Race.” *Review of Economic Studies,* Vol. 52 (1985a), pp. 193–209.

[15] Harris, C. and Vickers, J. “Patent Races and the Persistence of Monopoly.” *Journal of Industrial Economics,* Vol. 33 (1985b), pp. 461–481.

[16] Harris, C. and Vickers, J. “Racing with Uncertainty.” *Review of Economic Studies,* Vol. 54 (1987), pp. 1–21.

[17] Hopenhayn, H. and Mitchell, M. “Innovation Variety and Patent Breadth.” *The RAND Journal of Economics*, Vol. 32 (2001), pp. 152-166.

[18] Jewkes, John, David Sawers, and Richard Stillerman. *The Sources of Invention*, New York, N.Y: Norton, 1969.

[19] Judd, K. L. “Closed-loop equilibrium in an innovation race with experience,” Northwestern University, 1985a.

[20] 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.

[21] Judd, K. L. *Numerical Methods in Economics.* Cambridge, MA: MIT Press, 1998.

[22] Kamien, M. and Schwartz, N. *Market Structure and Innovation.* New York, NY: Cambridge University Press, 1982.

[23] Klemperer, P. ”How Broad should the Scope of a Patent Protection Be?” *RAND Journal of Economics* Vol. 21 (1990), pp. 113–130.

[24] Lee, T. and Wilde, L. “Market Structure and Innovation: A Reformulation.” *Quarterly Journal of Economics,* Vol. 94 (1980), pp. 429–436.

[25] Llobet, G., Hopenhayn, H. and Mitchell, M. " Rewarding Sequential Innovators: Prizes, Patents and Buyouts." Federal Reserve Bank of Minneapolis, Staff Report 273.

[26] Loury, G. " Market Structure and Innovation. " *The Quarterly Journal of Economics*, Vol. 93, (1979), pp. 395-410.

[27] Nordhaus, W. D., *Invention, Growth and Welfare*, Cambridge: M.I.T. Press, 1969.

[28] O’Donoghue, T., Scotchmer, S. and Thisse, J-F. "Patent Breadth, Patent Life, and the Pace of Technological Progress." *Journal of Economics, Management and Strategy*, Vol 7 (1998), pp. 1-32.

[29] Pakes, A. "Patents as Options: Some Estimates of the Value of Holding European Patent Stocks." Econometrica, Vol 54 (1986), pp. 755-784.

[30] 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.

[31] Puterman, M. L. *Markov Decision Processes: Discrete Stochastic Dynamic Programmming.* New York, NY: John Wiley & Sons, Inc., 1994.

[32] Reinganum, J. “Dynamic Games of Innovation.” *Journal of Economic Theory,* Vol. 25 (1981), pp. 21–41.

[33] Reinganum, J. “A Dynamic Game of R&D: Patent Protection and Competitive Behaviour.” *Econometrica,* Vol. 50 (1982), pp. 671–688.

[34] 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.

[35] Scotchmer, S. "Protecting Early Innovators: Should Second Generation Porducts be Patentable? " *RAND Journal of Economics,* Vol. 27 (1996), pp. 117–126.

[36] Scotchmer, S. " On the Optimality of the Patent Renewal System." *RAND Journal of Economics,* Vol. 30 (1999), pp. 181–196.

[37] Taylor, C. "Digging for Golden Carrots: An Analysis of Research Tournaments." *American Economic Review*, Vol 85 (1995), pp. 872-889.

[38] Wright, B. D. " The Economics of Invention Incentives: Patents, Prizes and Research Contracts." *American Economic Review*, Vol 73 (1983), pp. 691-707.