Intellectual Property Rights Policy, Competition and Innovation

Intellectual Property Rights Policy, Competition and Innovation*

Daron Acemoglu
Massachusetts Institute of Technology

Ufuk Akcigit
University of Pennsylvania

May 8, 2011

Abstract

To what extent and in what form should the intellectual property rights (IPR) of innovators be protected? Should a company with a large technology lead over its rivals receive the same IPR protection as a company with a more limited advantage? The analysis of these questions necessitates a dynamic framework for the study of the interactions between IPR and competition, in particular to understand the impact of such policies on future incentives. In this paper, we develop such a framework. The economy consists of many industries and firms engaged in cumulative (step-by-step) innovation. IPR policy regulates whether followers in an industry can copy (or license or build upon) the technology of the leader. With full patent protection, followers can catch up to the leader in their industry only by making the same innovation(s) themselves (or by full licensing). We prove the existence of a steady-state equilibrium in a baseline environment and characterize some of its properties. We then quantitatively investigate the implications of different types of IPR policy on the equilibrium growth rate and welfare. The most important result from this exercise is that full patent protection is not optimal (welfare maximizing); instead, optimal policy involves *state-dependent* IPR protection, providing *greater* protection to technology leaders that are further ahead than those that are close to their followers. This form of the optimal policy results from the impact of policy on dynamic incentives, in particular from a form of "*trickle-down*" effect: providing greater protection to firms that are further ahead of their followers than a certain threshold increases the R&D incentives also for all technology leaders that are less advanced than this threshold.

Keywords: competition, economic growth, endogenous growth, industry structure, innovation, intellectual property rights, licensing, patents, research and development, trickle-down.
JEL classification: O31, O34, O41, L16.

*We thank conference and seminar participants at the FBBVA Lecture at 2011 ASSA conference in Denver, the Canadian Institute of Advanced Research, EPFL Technology Policy Conference, European Science Days, MIT, National Bureau of Economic Research Economic Growth and Productivity Groups, Toulouse Information Technology Network, University of Pennsylvania, University of Toronto and Philippe Aghion, Gino Gancia, Bronwyn Hall, Sam Kortum and Suzanne Scotchmer for useful comments. Financial support from the Toulouse Information Technology Network is gratefully acknowledged.

1 Introduction

What is the optimal extent and form of intellectual property rights (IPR) protection? Should a firm with a large technology lead receive the same IPR protection as a company with a more limited technological lead, or should IPR policy be coupled with antitrust and used to limit the monopoly power of technology leaders? Despite broad consensus that innovation is central to the long-run performance of an economy, there is no consensus on the answers to such questions. A large literature on IPR (discussed below) focuses on the *static* trade-offs between the positive incentive benefits of IPR protection and its costs in terms of reducing competition and increasing markups. In this paper, we argue that *dynamic* trade-offs between IPR protection and competition, which have so far been overlooked, may be equally or more important for developing answers to these questions.

These issues and the importance of these questions are highlighted by several recent high-profile cases.[1] For example, motivated by antitrust concerns, a recent ruling of the European Commission ordered Microsoft to share secret information about its operating system and products with other software companies (New York Times, December 22, 2004). Similar issues were also central to the US Department of Justice (DOJ) case against Microsoft, which started on May 18th, 1998 and ultimately resulted in a ruling against Microsoft. Figures 1 and 2 show the evolution of R&D by Microsoft and by other top 10 publicly traded R&D investors in the IT sector relative to the sector average before and after the start of the DOJ case.[2]

[Figure 1 & 2 here]

The relative R&D spending by Microsoft and other industry leaders, which had been steadily—perhaps even exponentially—increasing since the mid-80s, appear to decline after the DOJ action. While one might expect R&D by Microsoft to slow down for a variety of reasons, it is not obvious why there should be a relative decline in the R&D of other top companies, since they partly benefited from the weakening of, and the restrictions imposed on, Microsoft. This relative decline may have been caused by a slowdown in the R&D activities of these other companies or an increase in the R&D investments of smaller IT firms (or by

[1] In addition to the Microsoft case, the issue of technological lead has been central in the Department of Justice investigations of Intel (New York Times, May 11, 2009) and the debates about Google's market share (New York Times, February 21, 2009).

[2] All data are from COMPUSTAT. Top 10 firms is determined by the highest 10 R&D investors (except Microsoft) in 1995. The patterns shown in Figures 1 and 2 are very similar if we use in that top 10 investors in 2000 or 1990, or if we benchmark it to the median of the industry rather than the mean. Top 10 investors in Figure 2 are: CA Inc, Continuum Inc, Intergraph Corp, Sterling Software Inc, Oracle Inc, Adobe Inc, Symantec Corp, Electronic Arts Inc, Sybase Inc, Intuit Inc.

entirely different and unrelated factors). To investigate these issues more systematically, we need a dynamic equilibrium framework where R&D activities of different types of firms might be affected by a change in IPR and competition policy.

Our framework builds on and extends the step-by-step innovation models of Aghion, Harris and Vickers (1997) and Aghion, Harris, Howitt and Vickers (2001), where a number of (typically two) firms engage in price competition within an industry and undertake R&D in order to improve their production technology. The technology gap between the firms determines the extent of the monopoly power of the leader, and hence the price markups and profits. The purpose of R&D by the follower is to catch up and surpass the leader (as in standard Schumpeterian models of innovation, e.g., Reinganum, 1981, 1985, Aghion and Howitt, 1992, Grossman and Helpman, 1991), while the purpose of R&D by the leader is to escape the competition of the follower and increase its markup and profits. Despite the dynamic nature of these models, their policy implications are still mostly based on the same static trade-off mentioned above. For this reason, for example, Aghion, Harris, Howitt and Vickers (2001, p. 481) conjecture that IPR protection should be limited and particularly so for firms with larger technological leads over their rivals (which face less competition and thus have greater monopoly power).

We extend these existing models in several directions. Most importantly, we explicitly introduce *state-dependent* patent/IPR protection policy, meaning a policy that makes the extent of patent or intellectual property rights protection conditional on the technology gap between different firms in the industry. As in racing-type models in general (e.g., Harris and Vickers, 1985, 1987, Budd, Harris and Vickers, 1993), a large gap between the leader and the follower discourages R&D by both. Consequently, overall R&D and technological progress are greater when the technology gap between the leader and the follower is relatively small.[3] One may then expect that full patent protection may be suboptimal in a world of step-by-step competition and permitting followers to copy or use the leaders' technologies would be particularly beneficial in industries where there is a large technology gap between leaders and followers.[4] However, crucially, this reasoning ignores the dynamic incentive effects, which are our main focus in this paper and emerge more clearly when IPR policy is explicitly state-dependent.

Our analysis establishes that the *opposite* of the above conjecture is always true in such

[3] Aghion, Bloom, Blundell, Griffith and Howitt (2005) provide empirical evidence from British industries consistent with the view R&D increases when there is a smaller technological gap between firms. See also Aghion and Griffith (2007).

[4] This is indeed the basis of Aghion, Harris, Howitt and Vickers's conjecture.

a dynamic equilibrium framework: optimal IPR policy should provide *greater* protection to technologically more advanced leaders. Underlying this result is what we refer to as the *trickle-down of incentives*: providing relatively low protection to firms with limited leads and greater protection to those that have greater leads not only improves the incentives of firms that are technologically advanced, but also encourages R&D by those that have limited leads because of the prospect of reaching levels of technology gaps associated with greater protection. A corollary of this result is that full IPR protection is not optimal, and there should be limited (but state-dependent) IPR protection for firms with only limited technology leads over their rivals.

More specifically, we show that in contrast to the standard disincentive effects of uniform relaxation of IPR policy, state-dependent relaxation that provides *greater* protection to technologically more advanced firms creates *a positive incentive effect*. This is because when a particular state for the technology leader (say being n^* steps ahead of the follower) becomes more profitable, this increases the incentives to perform R&D not only for leaders that are $n^* - 1$ steps ahead, but for *all* leaders with a lead of size $n \leq n^* - 1$. It is this trickle-down effect that generates the positive incentive effect and makes state-dependent IPR, with greater protection for firms that are technologically more advanced than their rivals, preferable to uniform IPR.

We start with a partial equilibrium model, which under some simplifying assumptions allows an explicit characterization of the trickle-down effect. We then provide a richer dynamic general equilibrium framework which allows a variety of different assumptions on how innovation depends on R&D by technology leaders and followers. Our baseline model focuses on *quick catch-up*, meaning that a follower can catch up with the technology leader with a single innovation regardless of the size of the gap between them. For this environment, we establish the existence of a stationary equilibrium and characterize some of its properties. We then study the form of optimal (welfare maximizing) IPR and competition policy quantitatively. The same effects as in the partial equilibrium analysis make state-dependent relaxation of IPR optimal. Quantitatively, we find that optimal state-dependent IPR policy can increase the growth rate of the economy from 1.86% to 2.04%, and does so with *fewer* workers employed in the R&D sector (because R&D workers are reallocated towards firms where their efforts directly lead to productivity growth). In contrast, uniform relaxation of IPR policy reduces both welfare and growth. These patterns are quite robust to different parameter values.

We next show how the framework can be extended to study these issues under alternative assumptions, in particular, assuming *slow catch-up* so that followers close the gap between

themselves and technology leaders only gradually. The presence of slow catch-up also enables us to introduce different types of R&D efforts and different dimensions of IPR policy, in particular, licensing and patent infringement fees.[5] We show that the trickle-down effect and the result that optimal IPR policy should be state-dependent and provide *greater protection* to technologically more advanced firms are robust in these alternative environments. In most cases, optimal IPR policy also increases growth by a similar magnitude to our baseline model (though in some cases, it increases welfare but not necessarily growth).[6]

Our paper is a contribution both to the IPR protection and the endogenous growth literatures. Previous work has focused on the static trade-off between ex-post monopoly rents and ex-ante R&D incentives (e.g., Arrow, 1962, Reinganum, 1981, Tirole, 1988, Romer, 1990, Grossman and Helpman, 1991, Aghion and Howitt, 1992, Green and Scotchmer, 1995, Scotchmer, 1999, Gallini and Scotchmer, 2002, O'Donoghue and Zweimuller, 2004).[7] Much of the literature discusses the trade-off between these two forces to determine the optimal length and breadth of patents. For example, Klemperer (1990) and Gilbert and Shapiro (1990) show that optimal patents should have a long duration in order to provide inducement to R&D, but a narrow breadth so as to limit monopoly distortions. A number of other papers, for example, Gallini (1992) and Gallini and Scotchmer (2002), reach opposite conclusions.

Another branch of the literature, including the seminal paper by Scotchmer (1999) and the recent interesting papers by Llobet, Hopenhayn and Mitchell (2006) and Hopenhayn and Mitchell (2001, 2011), adopts a mechanism design approach to the determination of the optimal patent and intellectual property rights protection system. For example, Scotchmer (1999) derives the patent renewal system as an optimal mechanism in an environment where the cost and value of different projects are unobserved and the main problem is to decide which projects should go ahead. Llobet, Hopenhayn and Mitchell (2006) consider optimal patent policy in the context of a model of sequential innovation with heterogeneous quality and private information. They show that allowing for a choice from a menu of patents will be optimal in this context.

[5] In particular, in this regime, we allow firms to undertake frontier as well as catch-up R&D. With frontier R&D, they can build on the technology leader's knowledge base and, if successful, they immediately surpass the leader, but might be liable for a *patent infringement fee*.

We also allow followers to *license* the innovation of the technology leader by paying a prespecified license fee—i.e., a "compulsory licensing" where the license fee is determined by IPR policy. We also show that voluntary licensing agreements would not achieve the same results, so our analysis establishes a potential need for compulsory licensing policy. Previous work emphasizing importance of compulsory licensing includes Tandon (1982), Gilbert and Shapiro (1990), and Kremer (2002). See Moser and Voena (2011) for a recent empirical investigation.

[6] We also show that both licensing and the possibility of frontier R&D (subject to infringement fees) contributes to growth and welfare.

[7] Boldrin and Levine (2004, 2008) or Quah (2003) argue that patent systems are not necessary for innovation.

Hopenhayn and Mitchell (2011) build on an earlier version of our paper, Acemoglu and Akcigit (2006), and derive a form of trickle-down effect using a mechanism design approach in a model with recurring innovations.

Our paper also extends Aghion, Harris and Vickers (1997) and Aghion, Harris, Howitt and Vickers (2001).[8] Although our model builds on these papers, it also differs from them in a number of significant ways. First and most importantly, we introduce state-dependent IPR policy. Second, we also introduce and analyze the slow catch-up regime, and in this context, we allow for compulsory licensing and for leapfrogging, which makes the followers directly contribute to the economic growth. We provide a full quantitative analysis of state-dependent IPR policy under these different scenarios. Third, our economy is a full general equilibrium model with competition between production and R&D for scarce labor.[9] Finally, we provide a general existence result and a number of analytical results for the general model (with or without IPR policy), while previous literature has focused on the special cases where innovations are either "drastic" (so that the leader never undertakes R&D) or very small, and has not provided existence or general characterization results for steady-state equilibria.

Lastly, our results are also related to the literature on tournaments and races, for example, Fudenberg, Gilbert, Stiglitz and Tirole (1983), Harris and Vickers (1985, 1987), Choi (1991), Budd, Harris and Vickers (1993), Taylor (1995), Fullerton and McAfee (1999), Baye and Hoppe (2003), and Moscarini and Squintani (2004). This literature considers the impact of endogenous or exogenous prizes on effort in tournaments, races or R&D contests. In terms of this literature, state-dependent IPR policy can be thought of as "state-dependent handicapping" of different players (where the state variable is the gap between the two players in a dynamic tournament). To the best of our knowledge, these types of schemes have not been considered in this literature.

The rest of the paper is organized as follows. Section 2 introduces the partial equilibrium model and analytically demonstrates the trickle-down effect. Section 3 presents our baseline environment (where a successful innovation by followers closes the entire gap with technology leaders in one step, i.e., there is quick catch-up). Section 4 proves the existence of a steady-state equilibrium and characterizes some of its key properties under both uniform and state-dependent IPR policy. Section 5 defines the social welfare objective and outlines our

[8] Segal and Whinston (2007) analyze the impact of anti-trust policy on economic growth in a related model of step-by-step innovation.

[9] This general equilibrium aspect is introduced to be able to close the model economy without unrealistic assumptions and makes our economy more comparable to other growth models (Aghion, Harris, Howit and Vickers, 2001, assume a perfectly elastic supply of labor). We show that the presence of general equilibrium interactions does not significantly complicate the analysis and it is still possible to characterize the steady-state equilibrium.

quantitative methods. Section 6 characterizes the structure of optimal IPR policy quantitatively. Section 7 extends the model to allow for slow catch-up, compulsory license fees and leapfrogging, and quantitatively characterizes the structure of optimal IPR policy under different combinations of these policies. Section 8 concludes, while the Appendix contains the proofs of all the results stated in the text.

2 A Partial Equilibrium Illustration

We first illustrate the main economic force in this paper, the trickle-down effect, using a partial equilibrium model. Consider the following infinite horizon, step-by-step R&D race between two competing firms in continuous time. Each firm maximizes the expected net present discounted value of "net profits," defined as operating profit minus R&D cost,

$$\mathbb{E}_t \int_t^\infty \exp\left(-r\left(s-t\right)\right)\left[\pi_i\left(s\right)-\Phi_i\left(s\right)\right]ds,$$

where $\mathbb{E}_t$ denotes expectation at time t, $r>0$ is the interest rate, $\pi_i(t)$ is the instantaneous operating profit flow and $\Phi_i(t)$ represents the R&D cost of firm i at time t. In this game, firm $i \in \{1,2\}$ invests in R&D to advance its position relative to its rival $i' \neq i$. Suppose that the positions of both firms in this race can be characterized by integer values on the real line, and denote the distance of firm i from its rival at time t by $n_i(t)$. In the partial equilibrium model, we simplify the analysis by following Aghion, Harris, Howitt and Vickers (2001) and Aghion, Bloom, Blundell, Griffith and Howitt (2005) in assuming that the maximum technology gap between a leader and a follower is 2; this assumption is relaxed in the full general equilibrium model analyzed in the rest of the paper. For now it simplifies the analysis by ensuring that the relative position of firm i can take five possible values, $n_i(t) \in \mathbf{N_I} \equiv \{-2,-1,0,1,2\}$. Let us denote the absolute gap between the firms by $n(t) \equiv \max\{n_i(t), n_{-i}(t)\}$, and suppress the time subscripts to simplify notation.

The payoffs in this game are assumed to be stationary and only a function of the relative distance between the firms, thus represented by $\boldsymbol{\pi}: \mathbf{N_I} \to \mathbb{R}_+$ (see equation (20) in Section 3). Here $\pi_{n_i} \geq 0$ is simply the instantaneous payoff that firm i obtains when its distance from its competitor is n_i at time t and assumed to be a strictly increasing function of n_i. To advance its relative position, firm i invests in R&D, which determines the Poisson rate of arrival of innovation, $x_i \in \mathbb{R}_+$. Let us also assume that the cost of R&D is linear in the arrival rate of innovation, i.e., $\Phi(x_i) = \phi x_i$, with $\phi > 0$ (again see below for more general formulations). Each successful innovation is patented and advances firm i's state (relative position) by one

step, so that following a successful innovation by firm i at time t we have: $n_i(t+) = n_i(t) + 1$ (where $n_i(t+)$ stands for n_i immediately following time t).

IPR policy governs the expected length of a patent. For simplicity, we model patent length by assuming that it terminates at a Poisson rate. Crucially for our focus, IPR policy is state dependent, and we represent it by the function: $\boldsymbol{\eta} : \mathbf{N_I} \rightarrow \mathbb{R}_+$. Here $\eta(n) \equiv \eta_n < \infty$ is the flow rate at which the patent terminates (patent protection is removed) for a technology leader that is n steps ahead. When $\eta_n = 0$, this implies that there is full protection at technology gap n, in the sense that patent protection will never be removed. In contrast, $\eta_n \rightarrow \infty$ implies that patent protection is removed immediately once technology gap n is reached. When the patent protection is removed, the firm that is behind copies the technology of its competitor and both firms end up *neck-and-neck*, i.e., $n = 0$.

Finally, we take the interest rate r as exogenous and assume that it satisfies $r < (\pi_n - \pi_{n-1})/4\phi$ for each $n \in \mathbf{N_I}$. This assumption ensures positive R&D by each firm when $\eta_n = 0$. Throughout we will focus on (stationary) Markov Perfect Equilibria (MPE), where strategies (R&D decisions) are only functions of the payoff-relevant state, which is $n \in \mathbf{N_I}$. A more formal definition of the MPE in the general equilibrium environment is given below.

The MPE can be characterized by writing the value functions of each firm as a function of the state $n \in \mathbf{N_I}$. These value functions are given by the following recursions:

$$
\begin{aligned}
rv_2 &= \pi_2 + x_{-2}[v_1 - v_2] + \eta_2[v_0 - v_2], && (1)\\
rv_1 &= \max_{x_1 \geq 0}\{\pi_1 - \phi x_1 + x_1[v_2 - v_1] + x_{-1}[v_0 - v_1] + \eta_1[v_0 - v_1]\}, && (2)\\
rv_0 &= \max_{x_0 \geq 0}\{\pi_0 - \phi x_0 + x_0[v_1 - v_0] + \tilde{x}_0[v_{-1} - v_0]\}, && (3)\\
rv_{-1} &= \max_{x_{-1} \geq 0}\{\pi_{-1} - \phi x_{-1} + x_{-1}[v_0 - v_{-1}] + x_1[v_{-2} - v_{-1}] + \eta_1[v_0 - v_{-1}]\}, && (4)\\
rv_{-2} &= \max_{x_{-2} \geq 0}\{\pi_{-2} - \phi x_{-2} + x_{-2}[v_{-1} - v_{-2}] + \eta_2[v_0 - v_{-2}]\}. && (5)
\end{aligned}
$$

In all equations, the first term represents current profits. In equations (2)-(5), the second term substracts R&D costs from current profits, the third term represents the fact that the firm will successfully innovate at the flow rate x_n and increase its position by one step. The fourth term incorporates the change in value due to an innovation by the rival firm. In equations (1) and (2) the last term is the change in value for the leader due to patent expiration, which takes place at the rate η_n, while in (4) and (5) is the change in value for the follower. Finally, equation (3) has the same interpretation except that now $n = 0$ and the two firms are neck-and-neck and thus there is no IPR policy (and the flow rate of innovation of the other firm is denoted by $\tilde{x}_0$, and naturally, in a symmetric equilibrium, we will have $x_0 = \tilde{x}_0$). Note also that in equations

(1) and (5), we used the fact that a two-step ahead firm does not undertake any R&D since it has already achieved the maximum feasible lead.

We will now characterize the MPE under two different policy environments: uniform and state-dependent IPR policy.

Uniform IPR Policy. Uniform IPR policy corresponds to the case where $\eta_n = \eta < \infty$. Consequently, optimal R&D decisions in equations (2)-(5) can be solved out as (see the Appendix):

$$\begin{aligned} x^*_{-2} &= \max\left\{-4\eta + \frac{\pi_2 - \pi_{-2}}{\phi} - 4r, 0\right\},\ x^*_{-1} = \max\left\{-3\eta + \frac{\pi_1 - \pi_{-2}}{\phi} - 3r, 0\right\}, \\ x^*_0 &= \max\left\{-2\eta + \frac{\pi_0 - \pi_{-2}}{\phi} - 2r, 0\right\},\ \text{and } x^*_1 = \max\left\{-\eta + \frac{\pi_{-1} - \pi_{-2}}{\phi} - r, 0\right\}. \end{aligned}$$

Inspection of these expressions immediately establishes the following result:

Proposition 1 *Under uniform IPR policy regime, any relaxation of IPR policy (away from $\eta = 0$) creates a "disincentive effect" and reduces all R&D levels.*

State-dependent IPR Policy. We next consider state-dependent policy where the patent protection of a technology leader depends on the technology gap, n. Optimal R&D decisions can now be written out as (see the Appendix):

$$\begin{aligned} x^*_{-2} &= \max\left\{-4\eta_2 + \frac{\pi_2 - \pi_{-2}}{\phi} - 4r, 0\right\},\ x^*_{-1} = \max\left\{-\eta_1 - 2\eta_2 + \frac{\pi_1 - \pi_{-2}}{\phi} - 3r, 0\right\}, \\ x^*_0 &= \max\left\{-2\eta_2 + \frac{\pi_0 - \pi_{-2}}{\phi} - 2r, 0\right\},\ \text{and } x^*_1 = \max\left\{\eta_1 - 2\eta_2 + \frac{\pi_{-1} - \pi_{-2}}{\phi} - r, 0\right\}. \end{aligned}$$

Inspection of these expressions shows that, in contrast to the uniform IPR case, relaxing patent protection can increase the R&D effort of the one-step leader, x^*_1. In particular, this can be accomplished by providing a lower protection in the current state (higher η_1) and/or a higher protection upon a successful innovation (lower η_2).

Proposition 2 *Under state-dependent IPR policy regime, relaxing IPR policy (away from $\eta_n = 0$) by weakening current protection (i.e., increasing η_1) creates a "positive incentive effect" and increases x^*_1.*

Whether optimal IPR policy will involve $\eta_1 > 0$ and/or $\eta_2 > 0$ now depends on the social returns from x_n's. For example, if x_1 is socially more beneficial than x_{-1}, $\eta_1 > 0$ will always be preferred. In the context of our general equilibrium model, this will always be the case. Proposition 2 provides a preview of these results.

3 General Equilibrium Framework

We now describe our baseline dynamic general equilibrium model. To maximize continuity with the previous literature and to provide the sharpest theoretical characterization results, our baseline model assumes *quick catch-up*, meaning that one innovation by a follower is sufficient to close the gap with the technology leader in the industry. The characterization of the equilibrium in this environment under the different policy regimes is presented in the next section. Alternative assumptions on the form of catch-up are investigated in Section 7.

3.1 Preferences and Technology

Consider the following continuous time economy with a unique final good. The economy is populated by a continuum of 1 individuals, each with 1 unit of labor endowment, which they supply inelastically. Preferences at time t are given by

$$\mathbb{E}_t \int_t^\infty \exp\left(-\rho\left(s-t\right)\right) \log C\left(s\right) ds, \tag{6}$$

where $\mathbb{E}_t$ denotes expectations at time t, $\rho > 0$ is the discount rate and $C\left(t\right)$ is consumption at date t. The logarithmic preferences in (6) facilitate the analysis, since they imply a simple relationship between the interest rate, growth rate and the discount rate (see (7) below).

Let $Y\left(t\right)$ be the total production of the final good at time t. We assume that the economy is closed and the final good is used only for consumption (i.e., there is no investment), so that $C\left(t\right) = Y\left(t\right)$. The standard Euler equation from (6) then implies that

$$g\left(t\right) \equiv \frac{\dot{C}\left(t\right)}{C\left(t\right)} = \frac{\dot{Y}\left(t\right)}{Y\left(t\right)} = r\left(t\right) - \rho, \tag{7}$$

where this equation defines $g\left(t\right)$ as the growth rate of consumption and thus output, and $r\left(t\right)$ is the interest rate at date t.

The final good Y is produced using a continuum 1 of intermediate goods according to the Cobb-Douglas production function

$$\ln Y\left(t\right) = \int_0^1 \ln y\left(j,t\right) dj, \tag{8}$$

where $y\left(j,t\right)$ is the output of jth intermediate at time t. Throughout, we take the price of the final good as the numeraire and denote the price of intermediate j at time t by $p\left(j,t\right)$. We also assume that there is free entry into the final good production sector. These assumptions, together with the Cobb-Douglas production function (8), imply that the final good sector has

the following demand for intermediates

$$y(j,t) = \frac{Y(t)}{p(j,t)}, \quad \forall j \in [0,1]. \tag{9}$$

Intermediate $j \in [0,1]$ comes in two different *varieties*, each produced by one of two infinitely-lived firms. We assume that these two varieties are perfect substitutes and these firms compete a la Bertrand.[10] Firm $i = 1$ or 2 in industry j has the following technology

$$y(j,t) = q_i(j,t)\, l_i(j,t) \tag{10}$$

where $l_i(j,t)$ is the employment level of the firm and $q_i(j,t)$ is its level of technology at time t. Each consumer in the economy holds a balanced portfolio of the shares of all firms. Consequently, the objective function of each firm is to maximize expected profits.

The production function for intermediate goods, (10), implies that the marginal cost of producing intermediate j for firm i at time t is

$$MC_i(j,t) = \frac{w(t)}{q_i(j,t)} \tag{11}$$

where $w(t)$ is the wage rate in the economy at time t.

When this causes no confusion, we denote the *technology leader* in each industry by i and the follower by $-i$, so that we have:

$$q_i(j,t) \geq q_{-i}(j,t).$$

Bertrand competition between the two firms implies that all intermediates will be supplied by the leader at the "limit" price:[11]

$$p_i(j,t) = \frac{w(t)}{q_{-i}(j,t)}. \tag{12}$$

Equation (9) then implies the following demand for intermediates:

$$y(j,t) = \frac{q_{-i}(j,t)}{w(t)} Y(t). \tag{13}$$

[10] A more general case would involve these two varieties being imperfect substitutes, for example, with the output of intermediate j produced as

$$y(j,t) = \left[\varphi y_1(j,t)^{\frac{\sigma-1}{\sigma}} + (1-\varphi)\, y_2(j,t)^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}},$$

with $\sigma > 1$. The model analyzed in the text corresponds to the limiting case where $\sigma \to \infty$. Our results can be easily extended to this more general case with any $\sigma > 1$, but at the cost of additional notation. We therefore prefer to focus on the case where the two varieties are perfect substitutes. It is nonetheless useful to bear this formulation with imperfect substitutes in mind, since it facilitates the interpretation of "distinct" innovations by the two firms (when the follower engages in "catch-up" R&D).

[11] If the leader were to charge a higher price, then the market would be captured by the follower earning positive profits. A lower price can always be increased while making sure that all final good producers still prefer the intermediate supplied by the leader i rather than that by the follower $-i$, even if the latter were supplied at marginal cost. Since the monopoly price with the unit elastic demand curve is infinite, the leader always gains by increasing its price, making the price given in (12) the unique equilibrium price.

3.2 Technology, R&D and IPR Policy under Quick Catch-up

R&D by the leader or the follower stochastically leads to innovation. We assume that when the leader innovates, its technology improves by a factor $\lambda > 1$.

The follower, on the other hand, can undertake R&D to catch up with the frontier technology. We will call this type of R&D as *catch-up R&D*.[12] Catch-up R&D can be thought of R&D to discover an *alternative* way of performing the same task as the current leading-edge technology. Because this innovation applies to the follower's variant of the product (recall footnote 10) and results from its own R&D efforts, we assume in our baseline framework that it does not constitute infringement on the patent of the leader.[13]

R&D by the leader and follower may have different costs and success probabilities. We simplify the analysis by assuming that both types of R&D have the same costs and the same probability of success. In particular, in all cases, we assume that innovations follow a controlled Poisson process, with the arrival rate determined by R&D investments. Each firm (in every industry) has access to the following R&D technology:

$$x_i(j,t) = F(h_i(j,t)), \tag{14}$$

where $x_i(j,t)$ is the flow rate of innovation at time t and $h_i(j,t)$ is the number of workers hired by firm i in industry j to work in the R&D process at t. This specification implies that within a time interval of Δt, the probability of innovation for this firm is $x_i(j,t)\Delta t + o(\Delta t)$.

We assume that F is twice continuously differentiable and satisfies $F'(\cdot) > 0$, $F''(\cdot) < 0$, $F'(0) < \infty$ and that there exists $\bar{h} \in (0,\infty)$ such that $F'(h) = 0$ for all $h \geq \bar{h}$. The assumption that $F'(0) < \infty$ implies that there is no Inada condition when $h_i(j,t) = 0$. The last assumption, on the other hand, ensures that there is an upper bound on the flow rate of innovation (which is not essential but simplifies the proofs). Recalling that the wage rate for labor is $w(t)$, the cost for R&D is therefore $w(t)G(x_i(j,t))$ where

$$G(x_i(j,t)) \equiv F^{-1}(x_i(j,t)), \tag{15}$$

and the assumptions on F immediately imply that G is twice continuously differentiable and satisfies $G'(\cdot) > 0$, $G''(\cdot) > 0$, $G'(0) > 0$ and $\lim_{x\to\bar{x}} G'(x) = \infty$, where

$$\bar{x} \equiv F\left(\bar{h}\right) \tag{16}$$

[12] This contrasts with *frontier R&D* introduced in Section 7, which will allow the follower to leapfrog the leader.

[13] We allow for infringement in Section 7.

is the maximal flow rate of innovation (with $\bar{h}$ defined above).

We next describe the evolution of technologies within each industry. Suppose that leader i in industry j at time t has a technology level of

$$q_i(j,t) = \lambda^{n_{ij}(t)}, \tag{17}$$

and that the follower $-i$'s technology at time t is

$$q_{-i}(j,t) = \lambda^{n_{-ij}(t)}, \tag{18}$$

where $n_{ij}(t) \geq n_{-ij}(t)$ and $n_{ij}(t)$, $n_{-ij}(t) \in \mathbb{Z}_+$ denote the technology rungs of the leader and the follower in industry j. We refer to $n_j(t) \equiv n_{ij}(t) - n_{-ij}(t)$ as the *technology gap* in industry j. If the leader undertakes an innovation within a time interval of Δt, then its technology increases to $q_i(j, t+\Delta t) = \lambda^{n_{ijt}+1}$ and the technology gap rises to $n_j(t+\Delta t) = n_j(t) + 1$ (the probability of two or more innovations within the interval Δt will be $o(\Delta t)$, where $o(\Delta t)$ represents terms that satisfy $\lim_{\Delta t \to 0} o(\Delta t)/\Delta t$).

In our baseline model, we assume that there is *quick catch-up* between followers and leaders. Namely, when the follower is successful in catch-up R&D within the interval Δt, then its technology improves to

$$q_{-i}(j, t+\Delta t) = \lambda^{n_{ijt}},$$

and thus it catches up with the leader immediately (regardless of how large the technology gap was). In this case, the technology gap variable becomes $n_{jt+\Delta t} = 0$ upon a successful innovation by the follower.[14]

In addition to catching up with the technology frontier with their own R&D, followers can also copy the technology frontier because IPR policy is such that some patents expire. In particular, we assume that patents expire at some policy-determined Poisson rate η, and after expiration, followers can costlessly copy the frontier technology, jumping to $q_{-i}(j, t+\Delta t) = \lambda^{n_{ijt}}$.[15] As in the partial equilibrium model in Section 2, IPR policy governs the length of the patent and we allow it to be state dependent, so it is represented by the following function:

$$\boldsymbol{\eta} : \mathbb{N} \to \mathbb{R}_+$$

Here $\eta(n) \equiv \eta_n < \infty$ is the flow rate at which the patent protection is removed from a technology leader that is n steps ahead of the follower. When $\eta_n = 0$, this implies that there is

[14] In Section 7, we will replace this assumption with *slow catch-up* where one innovation enables the follower to proceed by one step.

[15] Alternative modeling assumptions on IPR policy, such as a fixed patent length of $T > 0$ from the time of innovation, are not tractable, since they lead to value functions that take the form of delayed differential equations.

full protection at technology gap n, in the sense that patent protection will never be removed. In contrast, $\eta_n \to \infty$ implies that patent protection is removed immediately once technology gap n is reached. Our formulation imposes that $\boldsymbol{\eta} \equiv \{\eta_1, \eta_2, ...\}$ is time-invariant. Given this specification, we can now write the law of motion of the technology gap in industry j as follows:

$$n_j(t+\Delta t) = \begin{cases} n_j(t)+1 & \text{with probability} \quad x_i(j,t)\Delta t + o(\Delta t) \\ 0 & \text{with probability} \quad \left(x_{-i}(j,t) + \eta_{n_j(t)}\right)\Delta t + o(\Delta t) \\ n_j(t) & \text{with probability} \quad 1 - \left(x_i(j,t) + x_{-i}(j,t) + \eta_{n_j(t)}\right)\Delta t - o(\Delta t)) \end{cases}. \tag{19}$$

Here $o(\Delta t)$ again represents second-order terms, in particular, the probabilities of more than one innovations within an interval of length Δt. The terms $x_i(j,t)$ and $x_{-i}(j,t)$ are the flow rates of innovation by the leader and the follower; and $\eta_{n_j(t)}$ is the flow rate at which the follower is allowed to copy the technology of a leader that is $n_j(t)$ steps ahead. Intuitively, the technology gap in industry j increases from $n_j(t)$ to $n_j(t)+1$ if the leader is successful. The firms become "neck-and-neck" when the follower comes up with an alternative technology to that of the leader (flow rate $x_{-i}(j,t)$) or the patent expires at the flow rate η_{n_j}.

3.3 Profits

We next write the instantaneous "operating" profits for the leader (i.e., the profits exclusive of R&D expenditures). Profits of leader i in industry j at time t are

$$\begin{aligned} \Pi_i(j,t) &= [p_i(j,t) - MC_i(j,t)]\, y_i(j,t) \\ &= \left(\frac{w(t)}{q_{-i}(j,t)} - \frac{w(t)}{q_i(j,t)}\right)\frac{Y(t)}{p_i(j,t)} \\ &= \left(1 - \lambda^{-n_j(t)}\right) Y(t) \end{aligned} \tag{20}$$

where $n_j(t) \equiv n_{ij}(t) - n_{-ij}(t)$ is the technology gap in industry j at time t. The first line simply uses the definition of operating profits as price minus marginal cost times quantity sold. The second line uses the fact that the equilibrium limit price of firm i is $p_i(j,t) = w(t)/q_{-i}(j,t)$ as given by (12), and the final equality uses the definitions of $q_i(j,t)$ and $q_{-i}(j,t)$ from (17) and (18). The expression in (20) also implies that there will be zero profits in neck-and-neck industries, i.e., in those with $n_j(t) = 0$. Also clearly, followers always make zero profits, since they have no sales.

The Cobb-Douglas aggregate production function in (8) is responsible for the form of the profits (20), since it implies that profits only depend on the technology gap of the industry

and aggregate output. This will simplify the analysis below by making the technology gap in each industry the only industry-specific payoff-relevant state variable.

The objective function of each firm is to maximize the net present discounted value of "net profits" (operating profits minus R&D expenditures). In doing this, each firm will take the sequence of interest rates, $[r(t)]_{t\geq 0}$, the sequence of aggregate output levels, $[Y(t)]_{t\geq 0}$, the sequence of wages, $[w(t)]_{t\geq 0}$, the R&D decisions of all other firms and policies as given.

3.4 Equilibrium

Let $\boldsymbol{\mu}(t) \equiv \{\mu_n(t)\}_{n=0}^{\infty}$ denote the distribution of industries over different technology gaps, with $\sum_{n=0}^{\infty} \mu_n(t) = 1$. For example, $\mu_0(t)$ denotes the fraction of industries in which the firms are neck-and-neck at time t. Throughout, we focus on Markov Perfect Equilibria (MPE), where strategies are only functions of the payoff-relevant state variables.[16] This allows us to drop the dependence on industry j, thus we refer to R&D decisions by x_n for the technology leader that is n steps ahead and by x_{-n} for a follower that is n steps behind. Let us denote the list of decisions by the leader and the follower with technology gap n at time t by $\xi_n(t) \equiv \langle x_n(t), p_i(j,t), y_i(j,t)\rangle$ and $\xi_{-n}(t) \equiv \langle x_{-n}(t)\rangle$.[17] Throughout, $\boldsymbol{\xi}$ will indicate the whole sequence of decisions at every state, so that $\boldsymbol{\xi}(t) \equiv \{\xi_n(t)\}_{n=-\infty}^{\infty}$. We define an allocation as follows:

Definition 1 ***(Allocation)*** *Let $\boldsymbol{\eta}$ be the IPR policy sequence. Then an allocation is a sequence of decisions for a leader that is $n = 0, 1, 2, ...$ step ahead, $[\xi_n(t)]_{t\geq 0}$, a sequence of R&D decisions for a follower that is $n = 1, 2, ...$ step behind, $\left[\xi_{-n}(t)\right]_{t\geq 0}$, a sequence of wage rates $[w(t)]_{t\geq 0}$, and a sequence of industry distributions over technology gaps $[\boldsymbol{\mu}(t)]_{t\geq 0}$.*

For given IPR sequence $\boldsymbol{\eta}$, MPE strategies, which are only functions of the payoff-relevant state variables, can be represented as follows

$$\mathbf{x} : \mathbb{Z} \times \mathbb{R}_+^2 \times [0,1]^{\infty} \to \mathbb{R}_+.$$

[16] MPE is a natural equilibrium concept in this context, since it does not allow for implicit collusive agreements between the follower and the leader. While such collusive agreements may be likely when there are only two firms in the industry, in most industries there are many more firms and also many potential entrants, making collusion more difficult. Throughout, we assume that there are only two firms to keep the model tractable.

[17] The price and output decisions, $p_i(j,t)$ and $y_i(j,t)$, depend not only on the technology gap, aggregate output and the wage rate, but also on the exact technology rung of the leader, $n_{ij}(t)$. With a slight abuse of notation, throughout we suppress this dependence, since their product $p_i(j,t)\, y_i(j,t)$ and the resulting profits for the firm, (20), are independent of $n_{ij}(t)$, and consequently, only the technology gap, $n_j(t)$, matters for profits, R&D, aggregate output and economic growth.

This mapping represents the R&D decision of a firm (both when it is the follower and when it is the leader in an industry) as a function of the technology gap, $n \in \mathbb{Z}$, the aggregate level of output and the wage, $(Y, w) \in \mathbb{R}^2_+$, and R&D decision of the other firm in the industry, $\tilde{\mathbf{x}} \in [0,1]^\infty$. Consequently, we have the following definition of equilibrium:

Definition 2 ***(Equilibrium)*** *Given an IPR policy sequence $\boldsymbol{\eta}$, a Markov Perfect Equilibrium is given by a sequence $[\boldsymbol{\xi}^*(t), w^*(t), Y^*(t)]_{t\geq 0}$ such that (i) $[p_i^*(j,t)]_{t\geq 0}$ and $[y_i^*(j,t)]_{t\geq 0}$ implied by $[\boldsymbol{\xi}^*(t)]_{t\geq 0}$ satisfy (12) and (13); (ii) R&D policy $[\mathbf{x}^*(t)]_{t\geq 0}$ is a best response to itself, i.e., $[\mathbf{x}^*(t)]_{t\geq 0}$ maximizes the expected profits of firms taking aggregate output $[Y^*(t)]_{t\geq 0}$, wages $[w^*(t)]_{t\geq 0}$, government policy $\boldsymbol{\eta}$ and the R&D policies of other firms $[\mathbf{x}^*(t)]_{t\geq 0}$ as given; (iii) aggregate output $[Y^*(t)]_{t\geq 0}$ is given by (8); and (iv) the labor market clears at all times given the wage sequence $[w^*(t)]_{t\geq 0}$.*

3.5 The Labor Market

Since only the technology leader produces, labor demand in industry j with technology gap $n_j(t) = n$ can be expressed as

$$l_n(t) = \frac{\lambda^{-n} Y(t)}{w(t)} \quad \text{for} \quad n \in \mathbb{Z}_+. \tag{21}$$

In addition, there is demand for labor coming for R&D from both followers and leaders in all industries. Using (14) and the definition of the G function, we can express industry demands for R&D labor as

$$h_n(t) = \quad G(x_n(t)) + G(x_{-n}(t)) \quad \text{for} \quad n \in \mathbb{Z}_+ \ , \tag{22}$$

where $G(x_n(t))$ and $G(x_{-n}(t))$ refer to the demand of the leader and the follower in an industry with a technology gap of n. Note that in this expression, $x_{-n}(t)$ refers to the R&D effort of a follower that is n steps behind.

The labor market clearing condition can then be expressed as:

$$1 \geq \sum_{n=0}^{\infty} \mu_n(t) \left[\frac{1}{\omega(t) \lambda^n} + G(x_n(t)) + G(x_{-n}(t)) \right], \tag{23}$$

and $\omega(t) \geq 0$, with complementary slackness, where

$$\omega(t) \equiv \frac{w(t)}{Y(t)} \tag{24}$$

is the labor share at time t. The labor market clearing condition, (23), uses the fact that total supply is equal to 1, and demand cannot exceed this amount. If demand falls short of 1, then

the wage rate, $w(t)$, and thus the labor share, $\omega(t)$, have to be equal to zero (though this will never be the case in equilibrium). The right-hand side of (23) consists of the demand for production (the terms with ω in the denominator), the demand for R&D workers from the neck-and-neck industries ($2G(x_0(t))$ when $n=0$) and the demand for R&D workers coming from leaders and followers in other industries ($G(x_n(t))+G(x_{-n}(t))$ when $n>0$).

Defining the index of aggregate quality in this economy by the aggregate of the qualities of the leaders in the different industries, i.e.,

$$\ln Q(t) \equiv \int_0^1 \ln q_i(j,t)\, dj, \tag{25}$$

the equilibrium wage can be written as:[18]

$$w(t) = Q(t)\, \lambda^{-\sum_{n=0}^{\infty} n\mu_n(t)}. \tag{26}$$

3.6 Steady State and the Value Functions under Quick Catch-up

Let us now focus on steady-state (Markov Perfect) equilibria, where the distribution of industries $\boldsymbol{\mu}(t) \equiv \{\mu_n(t)\}_{n=0}^{\infty}$ is stationary, $\omega(t)$ defined in (24) and g, the growth rate of the economy, are constant over time. We will establish the existence of such an equilibrium and characterize a number of its properties. If the economy is in steady state at time $t=0$, then by definition, we have $Y^*(t) = Y_0 e^{g^* t}$ and $w^*(t) = w_0 e^{g^* t}$, where g^* is the steady-state growth rate. These two equations also imply that $\omega(t) = \omega^*$ for all $t \geq 0$. Throughout, we assume that the parameters are such that the steady-state growth rate g^* is positive but not large enough to violate the transversality conditions. This implies that net present values of each firm at all points in time will be finite. This enables us to write the maximization problem of a leader that is $n>0$ steps ahead recursively.

First note that given an optimal policy $\hat{x}$ for a firm, the net present discounted value of a leader that is n steps ahead at time t can be written as:

$$V_n(t) = \mathbb{E}_t \int_t^{\infty} \exp(-r(s-t)) \left[\Pi(s) - w(s)\, G(\hat{x}(s))\right] ds$$

where $\Pi(s)$ is the operating profit at time $s \geq t$ and $w(s)\, G(\hat{x}(s))$ denotes the R&D expenditure at time $s \geq t$. All variables are stochastic and depend on the evolution of the technology gap within the industry.

[18] Note that $\ln Y(t) = \int_0^1 \ln q_i(j,t)\, l(j,t)\, dj = \int_0^1 \left[\ln q_i(j,t) + \ln \frac{Y(t)}{w(t)} \lambda^{-n_j}\right] dj$, where the second equality uses (21). Thus we have $\ln Y(t) = \int_0^1 \left[\ln q_i(j,t) + \ln Y(t) - \ln w(t) - n_j \ln \lambda\right] dj$. Rearranging and canceling terms, and writing $\exp \int n_j \ln \lambda dj = \lambda^{-\sum_{n=0}^{\infty} n\mu_n(t)}$, we obtain (26).

Next taking as given the equilibrium R&D policy of other firms, $x^*_{-n}(t)$, the equilibrium interest and wage rates, $r^*(t)$ and $w^*(t)$, and equilibrium profits $\{\Pi^*_n(t)\}_{n=1}^{\infty}$ (as a function of equilibrium aggregate output), this value can be written as (see the Appendix for the derivation of this equation):[19]

$$r^*(t)V_n(t) - \dot{V}_n(t) = \max_{x_n(t)\geq 0}\left\{\begin{array}{c}[\Pi^*_n(t) - w^*(t)G(x_n(t))] + x_n(t)[V_{n+1}(t) - V_n(t)] \\ +\left(x^*_{-n}(t) + \eta_n\right)[V_0(t) - V_n(t)]\end{array}\right\}, \quad (27)$$

where $\dot{V}_n(t)$ denotes the derivative of $V_n(t)$ with respect to time. The first term is current profits minus R&D costs, while the second term captures the fact that the firm will undertake an innovation at the flow rate $x_n(t)$ and increase its technology lead by one step. The remaining terms incorporate changes in value due to quick catch-up by the follower (flow rate $x^*_{-n}(t)+\eta_n$ in the second line).

In steady state, the net present value of a firm that is n steps ahead, $V_n(t)$, will also grow at a constant rate g^* for all $n \in \mathbb{Z}_+$. Let us then define the normalized values as

$$v_n(t) \equiv \frac{V_n(t)}{Y(t)} \quad (28)$$

for all $n \in \mathbb{Z}$, which will be independent of time in steady state, i.e., $v_n(t) = v_n$.

Using (28) and the fact that from (7), $r(t) = g(t)+\rho$, the recursive form of the steady-state value function (27) can be written as:

$$\rho v_n = \max_{x_n\geq 0}\left\{\left(1-\lambda^{-n}\right) - \omega^* G(x_n) + x_n[v_{n+1} - v_n] + \left[x^*_{-n} + \eta_n\right][v_0 - v_n]\right\} \text{ for } n \in \mathbb{N}, \quad (29)$$

where x^*_{-n} is the equilibrium value of R&D by a follower that is n steps behind, and ω^* is the steady-state labor share (while x_n is now explicitly chosen to maximize v_n).

Similarly the value for neck-and-neck firms is

$$\rho v_0 = \max_{x_0\geq 0}\left\{-\omega^* G(x_0) + x_0[v_1 - v_0] + x^*_0[v_{-1} - v_0]\right\}, \quad (30)$$

while the values for followers are given by

$$\rho v_{-n} = \max_{x_{-n}\geq 0}\left\{-\omega^* G(x_{-n}) + [x_{-n} + \eta_n][v_0 - v_{-n}] + x^*_n[v_{-n-1} - v_{-n}]\right\} \text{ for } n \in \mathbb{N}. \quad (31)$$

For neck-and-neck firms and followers, there are no instantaneous profits, which is reflected in (30) and (31). In the former case this is because neck-and-neck firms sell at marginal cost, and in the latter case, this is because followers have no sales. These normalized value functions

[19] Clearly, this value function could be written for any arbitrary sequence of R&D policies of other firms. We set the R&D policies of other firms to their equilibrium values, $x^*_{-n}(t)$, to reduce notation in the main body of the paper.

emphasize that, because of growth, the effective discount rate is $r(t) - g(t) = \rho$ rather than $r(t)$.

The maximization problems in (29)-(31) immediately imply that any steady-state equilibrium R&D policies, $\mathbf{x}^*$, must satisfy:

$$x_n^* = \max\left\{G'^{-1}\left(\frac{[v_{n+1} - v_n]}{\omega^*}\right), 0\right\} \tag{32}$$

$$x_{-n}^* = \max\left\{G'^{-1}\left(\frac{[v_0 - v_{-n}]}{\omega^*}\right), 0\right\} \tag{33}$$

$$x_0^* = \max\left\{G'^{-1}\left(\frac{[v_1 - v_0]}{\omega^*}\right), 0\right\}, \tag{34}$$

where the normalized value functions, the vs, are evaluated at the equilibrium, and $G'^{-1}(\cdot)$ is the inverse of the derivative of the G function. Since G is twice continuously differentiable and strictly concave, G'^{-1} is continuously differentiable and strictly increasing. These equations therefore imply that innovation rates, the x_n^*s, will increase whenever the incremental value of moving to the next step is greater and when the cost of R&D, as measured by the normalized wage rate, ω^*, is less. Note also that since $G'(0) > 0$, these R&D levels can be equal to zero, which is taken care of by the max operator.

The response of innovation rates, x_n^*, to the increments in values, $v_{n+1} - v_n$, is the key economic force in this model. A policy that reduces the patent protection of leaders that are $n+1$ steps ahead (by increasing η_{n+1}) will make being $n+1$ steps ahead less profitable, thus reduce $v_{n+1} - v_n$ and x_n^*. This corresponds to the standard *disincentive effect* of relaxing IPR policy. This result corresponds to fact (1) in the toy model. In contrast to existing models, however, here relaxing IPR policy can also create a *positive incentive effect*. Somewhat paradoxically, lower protection for technology leaders that are $n+1$ steps ahead will tend to reduce v_{n+1}, thus increasing $v_{n+2} - v_{n+1}$ and x_{n+1}^*. This result is very similar to fact (2) in the toy model. We will see this positive incentive effect plays an important role in the form of optimal state-dependent IPR policy. In addition to the incentive effects, relaxing IPR protection may also create a beneficial *composition effect*; this is because, typically, $\{v_{n+1} - v_n\}_{n=0}^{\infty}$ is a decreasing sequence, which implies that x_{n-1}^* is higher than x_n^* for $n \geq 1$ (see, e.g., Proposition 4). Weaker patent protection (in the form of shorter patent lengths) will shift more industries into the neck-and-neck state and potentially increase the equilibrium level of R&D in the economy. Finally, weaker patent protection also creates a beneficial "level effect" by influencing equilibrium markups and prices (as shown in equation (12) above) and by reallocating some of the workers engaged in "duplicative" R&D to production. This level effect will also feature

in our welfare computations. The optimal level and structure of IPR policy in this economy will be determined by the interplay of these various forces.

Given the equilibrium R&D decisions $\mathbf{x}^*$, the steady-state distribution of industries across states $\boldsymbol{\mu}^*$ has to satisfy the following accounting identities:

$$\left(x^*_{n+1} + x^*_{-n-1} + \eta_{n+1}\right)\mu^*_{n+1} = x^*_n \mu^*_n \text{ for } n \in \mathbb{N}, \tag{35}$$

$$\left(x^*_1 + x^*_{-1} + \eta_1\right)\mu^*_1 = 2x^*_0 \mu^*_0, \tag{36}$$

$$2x^*_0 \mu^*_0 = \sum_{n=1}^{\infty} \left(x^*_{-n} + \eta_n\right)\mu^*_n. \tag{37}$$

The first expression equates exit from state $n+1$ (which takes the form of the leader going one more step ahead or the follower catching up the leader) to entry into the state (which takes the form of a leader from state n making one more innovation). The second equation, (36), performs the same accounting for state 1, taking into account that entry into this state comes from innovation by either of the two firms that are competing neck-and-neck. Finally, equation (37) equates exit from state 0 with entry into this state, which comes from innovation by a follower in any industry with $n \geq 1$.

The labor market clearing condition in steady state can then be written as

$$1 \geq \sum_{n=0}^{\infty} \mu^*_n \left[\frac{1}{\omega^* \lambda^n} + G\left(x^*_n\right) + G\left(x^*_{-n}\right)\right] \text{ and } \omega^* \geq 0, \tag{38}$$

with complementary slackness.

The next proposition characterizes the steady-state growth rate. As with all the other results in the paper, the proof of this proposition is provided in the Appendix.

Proposition 3 *Let the steady-state distribution of industries and R&D decisions be given by* $< \boldsymbol{\mu}^*, \ \mathbf{x}^* >$*, then the steady-state growth rate is*

$$g^* = \ln \lambda \left[2\mu^*_0 x^*_0 + \sum_{n=1}^{\infty} \mu^*_n x^*_n\right]. \tag{39}$$

This proposition clarifies that the steady-state growth rate of the economy is determined by two factors: (1) R&D decisions of industries at different levels of technology gap, $\mathbf{x}^* \equiv \{x^*_n\}^{\infty}_{n=-\infty}$; (2) The distribution of industries across different technology gaps, $\boldsymbol{\mu}^* \equiv \{\mu^*_n\}^{\infty}_{n=0}$. IPR policy affects these two margins in different directions as illustrated by the discussion above.

4 Existence and Characterization of Steady-State Equilibria

We now define a steady-state equilibrium in a more convenient form, which will be used to establish existence and derive some of the properties of the equilibrium.

Definition 3 ***(Steady-State Equilibrium)*** *Given an IPR policy $\boldsymbol{\eta}$, a steady-state equilibrium is a tuple $< \boldsymbol{\mu}^*,\ \mathbf{v},\ \mathbf{x}^*,\ \omega^*,\ g^* >$ such that the distribution of industries $\boldsymbol{\mu}^*$ satisfy (35), (36) and (37), the values $\mathbf{v} \equiv \{v_n\}_{n=-\infty}^{\infty}$ satisfy (29), (30) and (31), the R&D decision $\mathbf{x}^*$ is given by (32), (33) and (34), the steady-state labor share ω^* satisfies (38) and the steady-state growth rate g^* is given by (39).*

We next provide a characterization of the steady-state equilibrium, starting first with the case in which there is uniform IPR policy.

4.1 Uniform IPR Policy

Let us first focus on the case where IPR policy is uniform, i.e. $\eta_n = \eta < \infty$ for all $n \in \mathbb{N}$ and we denote this by $\boldsymbol{\eta}^{uni}$. In this case, (31) implies that the problem is identical for all followers, so that $v_{-n} = v_{-1}$ for $n \in \mathbb{N}$. Consequently, (31) can be replaced with the following simpler equation:

$$\rho v_{-1} = \max_{x_{-1} \geq 0} \left\{ -\omega^* G\left(x_{-1}\right) + \left[x_{-1} + \eta\right]\left[v_0 - v_{-1}\right] \right\}, \tag{40}$$

implying optimal R&D decisions for all followers of the form

$$x_{-1}^* = \max\left\{ G'^{-1}\left(\frac{\left[v_0 - v_{-1}\right]}{\omega^*} \right), 0 \right\}. \tag{41}$$

Let us denote the sequence of value functions under uniform IPR as $\{v_n\}_{n=-1}^{\infty}$. We next establish the existence of a steady-state equilibrium under uniform IPR and characterize some of its most important properties. Establishing the existence of a steady-state equilibrium in this economy is made complicated by the fact that the equilibrium allocation cannot be represented as a solution to a maximization problem. Instead, as emphasized by Definition 3, each firm maximizes its value taking the R&D decisions of other firms as given; thus an equilibrium corresponds to a set of R&D decisions that are best responses to themselves and a labor share (wage rate) ω^* that clears the labor market. Nevertheless, there is sufficient structure in the model to guarantee the existence of a steady-state equilibrium and monotonic behavior of values and R&D decisions.

Proposition 4 *Consider a uniform IPR policy* $\boldsymbol{\eta}^{uni}$ *and suppose that* $G'^{-1}\left(\left(1-\lambda^{-1}\right)/\left(\rho+\eta\right)\right)>0$. *Then a steady-state equilibrium* $<\boldsymbol{\mu}^*,\ \mathbf{v},\ \mathbf{x}^*,\ \omega^*,\ g^*>$ *exists. Moreover, in any steady-state equilibrium* $\omega^*<1$. *In addition, if either* $\eta>0$ *or* $x^*_{-1}>0$, *then* $g^*>0$. *For any steady-state R&D decisions* $\mathbf{x}^*$, *the steady-state distribution of industries* $\boldsymbol{\mu}^*$ *is uniquely determined.*

In addition, we have the following results:

- $v_{-1}\leq v_0$ *and* $\{v_n\}_{n=0}^{\infty}$ *forms a bounded and strictly increasing sequence converging to some* $v_\infty\in(0,\infty)$.
- $x^*_0>x^*_1$, $x^*_0\geq x^*_{-1}$, *and* $x^*_{n+1}\leq x^*_n$ *for all* $n\in\mathbb{N}$ *with* $x^*_{n+1}<x^*_n$ *if* $x^*_n>0$. *Moreover, provided that* $G'^{-1}\left(\left(1-\lambda^{-1}\right)/\left(\rho+\eta\right)\right)>0$ *and* $x^*_0>x^*_{-1}$.

Proof. See the Appendix. ■

Remark 1 The condition that $G'^{-1}\left(\left(1-\lambda^{-1}\right)/\left(\rho+\eta\right)\right)>0$ ensures that there will be positive R&D in equilibrium. If this condition does not hold, then there exists a trivial steady-state equilibrium in which $x^*_n=0$ for all $n\in\mathbb{Z}_+$, i.e., an equilibrium in which there is no innovation and thus no growth (this follows from the fact that $x^*_0\geq x^*_n$ for all $n\neq 0$, see the Appendix for more details). Moreover, if $\eta>0$, then this equilibrium would also involve $\mu^*_0=1$, so that in every industry two firms with equal costs compete a la Bertrand and charge price equal to marginal cost, leading to zero aggregate profits and a labor share of output equal to 1. The assumption that $G'^{-1}\left(\left(1-\lambda^{-1}\right)/\left(\rho+\eta\right)\right)>0$, on the other hand, is sufficient to rule out $\mu^*_0=1$ and thus $\omega^*=1$. If, in addition, the steady-state equilibrium involves some probability of catch-up or innovation by the followers, i.e., either $\eta>0$ or $x^*_{-1}>0$, then the growth rate is also strictly positive.

In addition to the existence of a steady-state equilibrium with positive growth, Proposition 4 shows that the sequence of values $\{v_n\}_{n=0}^{\infty}$ is strictly increasing and converges to some v_∞, and more importantly that $\mathbf{x}^*\equiv\{x^*_n\}_{n=1}^{\infty}$ is a decreasing sequence, which implies that technology leaders that are further ahead undertake less R&D. Intuitively, the benefits of further R&D are decreasing in the technology gap, since greater values of the technology gap translate into smaller increases in the equilibrium markup (recall (20)). Moreover, the R&D level of neck and-and-neck firms, x^*_0, is greater than both the R&D level of technology leaders that are one step ahead and followers that are one step behind (i.e., $x^*_0>x^*_1$ and $x^*_0\geq x^*_{-1}$). This implies that with uniform policy neck-and-neck industries are "most R&D intensive," while

industries with the largest technology gaps are "least R&D intensive". This is the basis of the conjecture mentioned in the Introduction that reducing protection given to technologically advanced leaders might be useful for increasing R&D by bringing them into the neck-and-neck state.

4.2 State-Dependent IPR Policy

We now extend the results from the previous section to the environment with state-dependent IPR policy, though results on monotonicity of values and R&D efforts no longer hold.[20]

Proposition 5 *Consider the state-dependent IPR policy* $\boldsymbol{\eta}$ *and suppose that* $G'^{-1}\left(\left(1-\lambda^{-1}\right)/\left(\rho+\eta_1\right)\right) > 0$. *Then a steady-state equilibrium* $< \boldsymbol{\mu}^*,\ \mathbf{v},\ \mathbf{x}^*,\ \omega^*,\ g^* >$ *exists. Moreover, in any steady-state equilibrium* $\omega^* < 1$. *In addition, if either* $\eta_1 > 0$ *or* $x^*_{-1} > 0$, *then* $g^* > 0$.

Proof. See the Appendix. ■

Unfortunately, it is not possible to determine the optimal (welfare- or growth-maximizing) state-dependent IPR policy analytically. For this reason, in Section 5, we undertake a quantitative investigation of the form and structure of optimal state-dependent IPR policy using plausible parameter values.

5 Optimal IPR Policy: Towards A Quantitative Investigation

In the remainder of the paper, we investigate the implications of various different types of IPR policies on R&D, growth and welfare using numerical computations of the steady-state equilibrium. Our purpose is not to provide a detailed calibration of the model economy but to highlight its qualitative implications for optimal IPR policy under plausible parameter values. We focus on optimal policy, defined as steady-state welfare-maximizing choice of policy (growth-maximizing policies give very similar results and are omitted to save space). In this section, we introduce the measure of steady-state welfare and describe our quantitative methodology. Results are reported in the subsequent sections.

[20] This is because IPR policies could be very sharply increasing at some technology gap, making a particular state very unattractive for the leader. For example, we could have $\eta_n = 0$ and $\eta_{n+1} \to \infty$, which would imply that $v_{n+1} - v_n$ is negative.

5.1 Welfare

Our focus so far has been on steady-state equilibria (mainly because of the very challenging nature of transitional dynamics in this class of models). In our quantitative analysis, we continue to focus on steady states and thus look at steady-state welfare. In a steady-state equilibrium, welfare at time $t = 0$ can be written as

$$\begin{aligned} \text{Welfare}\,(0) &= \int_0^\infty e^{-\rho t} \ln\left(Y\,(0)\, e^{g^* t}\right) dt \\ &= \frac{\ln Y\,(0)}{\rho} + \frac{g^*}{\rho^2}, \end{aligned} \tag{42}$$

where the first-line uses the facts that all output is consumed, utility is logarithmic (recall (6)), output and consumption at date $t = 0$ are given by $Y\,(0)$, and in the steady-state equilibrium output grows at the rate g^*. The second line simply evaluates the integral. Next, note that

$$\begin{aligned} \ln Y\,(t) &= \int_0^1 \ln y\,(j,t)\, dj \\ &= \int_0^1 \ln\left(\frac{q_{-i}\,(j,t)\, Y\,(t)}{w\,(t)}\right) dj \\ &= \int_0^1 \ln q_{-i}\,(j,t)\, dj - \ln \omega\,(t) \\ &= \ln Q\,(t) - \ln \lambda \left(\sum_{n=0}^{\infty} n \mu_n\,(t)\right) - \ln \omega\,(t)\,, \end{aligned} \tag{43}$$

where the first line simply uses the definition in (8), the second line substitutes for $y\,(j,t)$ from (13), the third line uses the definition of the labor share $\omega\,(t)$, and the final line uses the definition of $Q\,(t)$ from (25) together with the fact that in the steady state $q_i\,(j,t) = \lambda^n q_{-i}\,(j,t)$ in a fraction $\mu_n\,(t)$ of industries. The expression in (43) implies that output simply depends on the quality index, $Q\,(t)$, the distribution of technology gaps, $\boldsymbol{\mu}\,(t)$ (because this determines markups), and also on the labor share, $\omega\,(t)$. In steady-state equilibrium, the distribution of technology gaps and labor share are constant, while output and the quality index grow at the steady-state rate g^*. Therefore, for steady-state comparisons of welfare across economies with different policies, it is sufficient to compare two economies with the same level of $Q\,(0)$, but with different policies. We can then evaluate steady-state welfare with the distribution of industries given by their steady-state values in the two economies, and output and the quality index growing at the corresponding steady-state growth rates. Expression (43) also makes it clear that only the aggregate quality index $Q\,(0)$ needs to be taken to be the same in the different economies. Given $Q\,(0)$, the dispersion of industries in terms of the quality levels

has no effect on output or welfare (though, clearly, the distribution of industries in terms of technology gaps between leaders and followers, $\boldsymbol{\mu}$, influences the level of markups and output, and thus welfare).

However, note one difficulty with welfare comparisons highlighted by equations (42) and (43); proportional changes in steady-state welfare due to policy changes will depend on the initial level of $Q(0)$, which is an arbitrary number. Therefore, proportional changes in welfare are not informative, though this has no effect on ordinal rankings and thus welfare-maximizing policy is well defined and independent of the level of $Q(0)$. Equations (42) and (43) also make it clear that changes in steady-state welfare will be the sum of two components: the first is the *growth effect,* given by g^*/ρ^2, whereas the second is due to changes in $\ln\lambda\left(\sum_{n=0}^{\infty} n\mu_n\right)/\rho - \ln\omega(0)$. Since changes in the labor share $\omega(0)$ are largely driven by the distribution of industries, we refer to this as the *distribution effect.* Policies will typically affect both of these quantities. In what follows, we give the welfare rankings of different policies and then report the relative magnitudes of the growth and the distribution effects. This will show that the growth effects will be one or two orders of magnitude greater than the distribution effects and dominate welfare comparisons. So if the reader wishes, he or she may think of the magnitudes of the changes in welfare as given by the proportional changes in growth rates.

5.2 Quantitative Methods and Parameter Choices

For our quantitative exercise, we take the annual discount rate as 5%, i.e., $\rho_{year} = 0.05$. In all our computations, we work with the monthly equivalent of this discount rate in order to increase precision, but throughout the tables, we convert all numbers to their annual counterparts to facilitate interpretation.

The theoretical analysis considered a general production function for R&D given by (14). The empirical literature typically assumes a Cobb-Douglas production function. For example, Kortum (1993) considers a function of the form

$$\text{Innovation}(t) = B_0 \exp(\kappa t)(\text{R\&D inputs})^{\gamma}, \tag{44}$$

where B_0 is a constant and $\exp(\kappa t)$ is a trend term, which may depend on general technological trends, a drift in technological opportunities, or changes in general equilibrium prices (such as wages of researchers etc.). The advantage of this form is not only its simplicity, but also the fact that most empirical work estimates a single elasticity for the response of innovation rates to R&D inputs. Consequently, they essentially only give information about the parameter γ in terms of equation (44). A low value of γ implies that the R&D production function is more

concave. For example, Kortum (1993) reports that estimates of γ vary between 0.1 and 0.6 (see also Pakes and Griliches, 1980, or Hall, Hausman and Griliches, 1988). For these reasons, throughout, we adopt a R&D production function similar to (44):

$$x = Bh^{\gamma} \tag{45}$$

where $B, \gamma > 0$. In terms of our previous notation, equation (45) implies that $G(x) = [x/B]^{\frac{1}{\gamma}} w$, where w is the wage rate in the economy (thus in terms of the above function, it is captured by the $\exp(\kappa t)$ term).[21] Equation (45) does not satisfy the boundary conditions we imposed so far and can be easily modified to do so without affecting any of the results, since in all numerical exercises only a finite number of states are reached.[22] Following the estimates reported in Kortum (1993), we start with a benchmark value of $\gamma = 0.35$, and then report sensitivity checks for $\gamma = 0.1$ and $\gamma = 0.6$. The other parameter in (45), B, is chosen so as to ensure an annual growth rate of approximately 1.9%, i.e., $g^* \simeq 0.019$, in the benchmark economy which features indefinitely-enforced patents. This growth rate together with $\rho_{year} = 0.05$ also pins down the annual interest rate as $r_{year} = 0.069$ from equation (7).

We choose the value of λ using a reasoning similar to Stokey (1995). Equation (39) implies that if the expected duration of time between any two consecutive innovations is about 3 years in an industry, then a growth rate of about 1.9% would require $\lambda = 1.05$.[23] This value is also consistent with the empirical findings of Bloom, Schankerman and Van Reenen (2005).[24] We take $\lambda = 1.05$ as the benchmark value. We then check the robustness of the results to $\lambda = 1.01$ and $\lambda = 1.2$ (expected duration of 8 months and 13 years, respectively). Finally, without loss

[21] More specifically, (45) can be alternatively written as

$$\text{Innovation}(t) = Bw(t)^{-\gamma} (\text{R\&D expenditure})^{\gamma},$$

thus would be equivalent to (44) as long as the growth of $w(t)$ can be approximated by constant rate.

[22] For example, we could add a small linear term to the production function for R&D, (45), and also make it flat after some level $\bar{h}$. For example, the following generalization of (45),

$$x = \min\left\{Bh^{\gamma} + \varepsilon h; B\bar{h}^{\gamma} + \varepsilon\bar{h}\right\}$$

for ε small and $\bar{h}$ large, makes no difference to our simulation results.

[23] In particular, in our benchmark parameterization with full protection without licensing, 24% of industries are in the neck-and-neck state. This implies that improvements in the technological capability of the economy is driven by the R&D efforts of the leaders in 76% of the industries and the R&D efforts of both the leaders and the followers in 24% of the industries. Therefore, the growth equation, (39), implies that $g \simeq \ln\lambda \times 1.24 \times x$, where x denotes the average frequency of innovation in a given industry. A major innovation on average every three years implies a value of $\lambda \simeq 1.05$.

[24] The production function for the intermediate good, (10), can be written as $\log(y(j,t)) = n(j,t)\log(\lambda) + \log(l(j,t))$, where $n(j,t)$ is the number of innovations to date in sector j and represents the "knowledge stock" of this industry. Bloom, Schankerman and Van Reenen (2005) proxy the knowledge stock in an industry by the stock of R&D in that industry and estimate the elasticity of sales with respect to the stock of R&D to be approximately 0.06. In terms of the exercise here, this implies that $\log(\lambda) = 0.06$, or that $\lambda \approx 1.06$.

of generality, we normalize labor supply to 1. This completes the determination of all the parameters in the model except the IPR policy.

As noted above, we begin with the full patent protection regime, i.e., $\boldsymbol{\eta} = \{0, 0, ...\}$. We then move to a comparison of the optimal (welfare-maximizing) uniform IPR policy $\boldsymbol{\eta}^{uni}$ to the optimal state-dependent IPR policy. Since it is computationally impossible to calculate the optimal value of each η_n, we limit our investigation to a particular form of state-dependent IPR policy, whereby the same η applies to all industries that have a technology gap of $n = 5$ or more. In other words, the IPR policy can be represented as:

IPR policy $\rightarrow$	*none*	η_1	η_2	η_3	η_4	η_5								
Technology gap: $n \rightarrow$	0	1	2	3	4	5	6	7	8	9	10	11	.	. ∞

We checked and verified that allowing for further flexibility (e.g., allowing η_5 and η_6 to differ) has little effect on our results.

The numerical methodology we pursue relies on *uniformization* and value function iteration. The details of the uniformization technique are described in the proof of Lemma 1 in the Appendix (for details of value function iteration, see Judd, 1999). In particular, we first take the IPR policy $\boldsymbol{\eta}$ as given and make an initial guess for the equilibrium labor share ω^*. Then for a given ω^*, we generate a sequence of values $\{v_n\}_{n=-\infty}^{\infty}$, and we derive the optimal R&D policies, $\{x_n^*\}_{n=-\infty}^{\infty}$ and the steady-state distribution of industries, $\{\mu_n^*\}_{n=0}^{\infty}$. After convergence, we compute the growth rate g^* and welfare, and then check for market clearing in the labor market from equation (23). Depending on whether there is excess demand for or supply of labor, ω^* is varied and the numerical procedure is repeated until the entire steady-state equilibrium for a given IPR policy is computed. The process is then repeated for different IPR policies.

In the state-dependent IPR case, the optimal (welfare-maximizing) IPR policy sequences, $\boldsymbol{\eta}$, are computed one element at a time, until we find the welfare-maximizing value for that component, for example, η_1. We then move the next component, for example, η_2. Once the welfare-maximizing value of η_2 is determined, we go back to optimize over η_1 again, and this procedure is repeated recursively until convergence.[25]

6 Optimal IPR Policy

In this section, we present a quantitative analysis of our baseline model.

[25] After we find a maximizer ($\boldsymbol{\eta}^*$), we also evaluate several random policy combinations around the maximizer to verify the solution.

6.1 Full IPR Protection

We start with the benchmark with full protection, which is the case that the existing literature has considered so far (e.g., Aghion, Harris, Howitt and Vickers, 2001). In terms of our model, this corresponds to $\eta_n = 0$ for all n. We choose the parameter B in terms of (45), so that the benchmark economy has an annual growth rate of 1.86%.

[Figure 3 & 4 & 5 here]

The value function for this benchmark case is shown in Figure 3 (solid line). The value function has decreasing differences for $n \geq 0$, which is consistent with the results in Proposition 4, and features a constant level for all followers (since there is no state dependence in the IPR policy). Figure 4 shows the level of R&D efforts for leaders and followers in this benchmark (again solid line). Again consistent with Proposition 4, this figure also shows that the R&D level of a leader declines as the technology gap increases and that the highest level of R&D is for firms that are neck-and-neck (i.e., at the technology gap of $n = 0$). Since there is no state-dependent IPR policy, all followers undertake the same level of R&D effort, which is also shown in the figure.

Figure 5 shows the distribution of industries according to technology gaps (again the solid line refers to the benchmark case). The mode of the distribution is at the technology gap of $n = 1$, but there is also a significant concentration of industries at technology gap $n = 0$, because innovations by the followers take them to the "neck-and-neck" state.

[Table 1 here]

The first column of Table 1 also reports the results for this benchmark simulation. As noted above, in each case B is chosen such that the annual growth rate is equal to 0.0186, which is recorded at the bottom of Table 1 together with the initial consumption and welfare levels according to (42) and (43). The table also shows the R&D levels x_0^*, x_{-1}^* and x_1^* (0.35, 0.22 and 0.29), the frequencies of industries with technology gaps of 0, 1 and 2. The steady-state value of ω is 0.95. Since labor is the only factor of production in the economy, ω^* should not be thought of as the labor share in GDP. Instead, $1 - \omega^*$ measures the share of pure monopoly profits in value added. In the benchmark parameterization, this corresponds to 5% of GDP, which is reasonable.[26] Finally, the table also shows that in this benchmark parameterization 3.2% of

[26] Bureau of Economic Analysis (2004) reports that the ratio of before-tax profits to GDP in the US economy in 2001 was 7% and the after-tax ratio was 5%.

the workforce is working as researchers, which is also consistent with US data.[27] These results are encouraging for our simple quantitative exercise, since with very few parameter choices, the model generates reasonable numbers, especially for the share of the workforce allocated to research.[28]

6.2 Optimal Uniform IPR Protection

For reference, we now characterize optimal uniform IPR policy, that is, we impose that $\eta_n = \eta$ for all n, and look for values of η that maximizes the welfare in the economy. Column 2 of Table 1 shows that the welfare-maximizing value of η is not different from zero at the three-digit level. Therefore the results of the full protection case carries over to uniform policy as well. The main reason for this result is the quick catch-up assumption. Recall that the uniform IPR policy discourages innovation, but generates a potential benefit because of the composition effect (bringing more firms into neck-and-neck position). In the quick catch-up regime, firms come into neck-and-neck position at a Poisson rate of 0.22, which results in 35% of sectors being in state 0 and 77% at two-step gap or below. This implies that there are only limited composition gains. In this light, it is not surprising that relaxing the IPR protection uniformly is not beneficial; it generates a significant disincentive effect and little benefit. Therefore, optimal IPR policy is to set full protection, $\boldsymbol{\eta}^* = 0$, and thus the value functions, innovation rates and industry distributions under optimal uniform IPR policy are given by the solid lines in Figures 3-5.

6.3 Optimal State-Dependent IPR

We next turn to our major question; whether state-dependent IPR makes a significant difference relative to the uniform IPR. In particular, we look for the combination of $\{\eta_1, ..., \eta_5\}$ that maximizes the welfare. The new value function, innovation rates and industry distribution are plotted in Figure 3-5 and the numerical results are shown in column 3 of Table 1.

Two features are worth noting. First, the growth rate increases noticeably relative to column 1; it is now 2.04% instead of 1.86%. Second and more important, we see the key pattern that will be present in all of our quantitative results: optimal state-dependent policy $\{\eta_1, ..., \eta_5\}$

[27] According to National Science Foundation (2006), the ratio of scientists and engineers in the US workforce in 2001 is about 4%.

[28] Most endogenous growth models imply that a significantly greater fraction of the labor force should be employed in the research sector and one needs to introduce various additional factors to reduce the profitability of research or to make entry into research more difficult. In the current model, the step-by-step nature of innovation and competition plays this role and generates a plausible allocation of workers between research and production.

provides *greater* protection to technology leaders that are further ahead. In particular, we find that the optimal policy involves $\eta_1 = 0.71$, $\eta_2 = 0.08$, and $\eta_3 = \eta_4 = \eta_5 = 0$. This corresponds to very little patent protection for firms that are one step ahead of the followers. In particular, since $\eta_1 = 0.71$ and $x^*_{-1} = 0.12$, in this equilibrium firms that are one step behind followers are more than *six times* as likely to catch up with the technology leader because of the expiration of the patent of the leader as they are likely to catch up because of their own successful R&D. Then, there is a steep increase in the protection provided to technology leaders that are two steps ahead, and η_2 is 1/12th of η_1. Perhaps even more remarkably, after a technology gap of three or more steps, optimal IPR involves full protection, and patents never expire.

This pattern of greater protection for technology leaders that are further ahead may go against a naïve intuition that state-dependent IPR policy should try to boost the growth rate of the economy by bringing the industries with largest technology gaps (where leaders engage in little R&D) into neck-and-neck competition. This composition effect is present, but dominated by another, more powerful force, the *trickle-down* effect. The intuition for the trickle-down effect is as follows: by providing secure patent protection to firms that are three or more steps ahead of their rivals, optimal state-dependent IPR increases the R&D effort of leaders that are one and two steps ahead as well. This is because technology leaders that are only one or two steps ahead now face greater returns to R&D, which will not only increase their profits but also the security of their intellectual property. Mechanically, high levels of η_1 and η_2 reduce v_1 and v_2, while high IPR protection for more advanced firms increases v_n for $n \geq 3$, and this increases the R&D incentives of leaders at $n = 1$ or at $n = 2$.[29]

Providing more secure patent protection through less frequent catch-up benefits an n-step leader more than $(n+1)$-step leader since the preserved profit is higher for a more advanced firm. This results in a steeper value function as illustrated in Figure 3. The slope of the value function is the key determining factor for R&D decisions and this increase in slope reflects itself in overall higher R&D effort by the leaders in Figure 4. It is also notable that state-dependent IPR introduces positive incentive effect while gaining also from the composition. Figure 5 shows that the mode of the new distribution is at $n = 0$. The average innovation rate is higher (as reflected on a higher growth rate, $g = 2.04\%$) and the average mark-up is lower ($C(0)$ increases by 52%). This pattern of greater R&D investments under state-dependent

[29] An alternative intuition, suggested by an anonymous referee, is that when the technology gap is greater, leaders will lose more from a relaxation of IPR. However, this intuition can only be partial, since, as shown in Section 2, state-dependent relaxation of IPR in this form creates a positive incentive effect, which is central to our results (and this is independent of how much technology leaders lose as a result of the relaxation of IPR). As a result, we believe that the trickle-down of incentives is the more correct intuition for our results.

IPR contrasts with uniform IPR, which always reduces R&D of all firms. The possibility that imperfect state-dependent IPR protection can *increase* (rather than reduce) R&D incentives is a novel feature of our approach and has also been shown explicitly in the partial equilibrium model of Section 2.

6.4 Robustness

The patterns shown in Figures 3-5 and Table 1 are highly robust. In the working paper version, we repeated this entire exercise for various combinations of values of γ and λ (in particular, varying γ to $\gamma = 0.1$ and $\gamma = 0.6$, and λ to $\lambda = 1.01$ and $\lambda = 1.2$). The overall pattern and in fact the quantitative magnitudes are remarkably similar to the baseline reported here. We do not report these robustness checks to save space (they are available upon request); instead, we focus on the results in the slow catch-up regime.

7 Optimal IPR Policy in the Slow Catch-up Regime

In this section, we extend our analysis to an environment where followers close the gap with technology leaders also step by step. This environment will further allow us to introduce different types of R&D efforts by followers and study several different dimensions of IPR policy.

7.1 Value Functions

The environment is the same as in Section 3, except that we now assume that successful R&D by followers close is the gap between themselves and the technology leader by one step. We will allow for different types of R&D below. The equivalent expressions for the value functions (29)-(31) in this case are

$$\rho v_n = \max_{x_n \geq 0} \left\{ \left(1 - \lambda^{-n}\right) - \omega^* G\left(x_n\right) + x_n \left[v_{n+1} - v_n\right] + x^*_{-n} \left[v_{n-1} - v_n\right] + \eta_n \left[v_0 - v_n\right] \right\} \text{ for } n \in \mathbb{N}, \tag{46}$$

$$\rho v_0 = \max_{x_0 \geq 0} \left\{ -\omega^* G\left(x_0\right) + x_0 \left[v_1 - v_0\right] + x^*_0 \left[v_{-1} - v_0\right] \right\}, \tag{47}$$

and

$$\rho v_{-n} = \max_{x_{-n} \geq 0} \left\{ -\omega^* G\left(x_{-n}\right) + x_{-n} \left[v_{-n+1} - v_{-n}\right] + x^*_n \left[v_{-n-1} - v_{-n}\right] + \eta_n \left[v_0 - v_{-n}\right] \right\} \text{ for } n \in \mathbb{N}. \tag{48}$$

These expressions are intuitive in light of those presented in Section 3, in particular, (29)-(31). The only difference from equations (29)-(31) is that, when a follower innovates, an

n-step leader's value changes from v_n to v_{n-1} instead of dropping all the way to v_0, since this innovation closes the technology gap only by one step. Similarly, in this event, the follower's value changes from v_{-n} to v_{-n+1} instead of increasing all the way to v_0. The rest of the analysis mirrors that in Section 4. In particular, existence of stationary equilibria can be proved using an analogous argument to that provided in the Appendix, but we are not able to prove the analogue of the second part of Proposition 4.

7.2 Quantitative Results

We next investigate the form of optimal IPR policy in the baseline slow catch-up regime.

7.2.1 Full IPR Protection

Under the slow catch-up regime, setting $\eta_n = 0$, that is, providing full protection via infinite patent length generates too little catch-up by the followers. Consequently, the steady-state distribution has little mass at or around the "neck-and-neck" state ($n = 0$). To generate a more plausible distribution with a non-zero share of industries in the neck-and-neck state, we instead impose $\eta_n = 0.02$, which implies an expected length of patent protection of 50 years (as the full protection benchmark) under slow catch-up regime.

[Table 2 here]

The first column of Table 2 reports the results under this scenario. Even with 50 years of protection, the share of industries that are neck-and-neck is only 2%, and the total share of industries that have a gap of less than two steps is only 8%. One implication of this pattern is that a relaxation of IPR policy may now be more powerful because it can affect the composition of industries, reduce the average mark-up in the economy, and perhaps have a large effect on average R&D. Therefore, this is a particularly relevant environment for investigating whether the trickle-down of incentives identified in the previous section is present and robust in different and perhaps more realistic environments.

7.2.2 Optimal Uniform IPR Protection

The second column of Table 2 shows optimal uniform IPR policy in this case. Consistent with Proposition 1 in Section 2, relaxing IPR protection creates a powerful disincentive effect. However, it also generates a beneficial composition effect by bringing more and more firms into neck-and-neck competition. For this reason, optimal uniform IPR policy is no longer full protection.

The results in the table show that the optimal policy reduces patent length from $\eta = 0.02$ (average protection of 50 years) to $\eta = 0.11$ (average protection of 9 years). This involves a lower innovation rate for technology leaders that are one-step ahead (from 1.1 to 0.15). Similarly average R&D is also reduced and the aggregate growth rate declines from 2.5% to 2.3%. However, because of the increase in the share of neck-and-neck industries (from 2% to 16%) and the increase in the total share of industries that are in the first 3 states (from 8% to 49%), the average mark-up in the economy decreases. This enables a large (19-fold) increase in initial consumption $C(0)$ (which is the reason why this policy is optimal even though it reduces growth).

7.2.3 Optimal State-Dependent IPR

Once again, the most interesting case is when IPR policy is state dependent. In this case, the optimal policy not only benefits from the composition effect, but can do so without sacrificing growth (by exploiting the positive incentive and the trickle-down effects highlighted in Proposition 2 in Section 2).

The optimal state-dependent policies shown in column 3 of Table 2. Under this optimal policy, the share of the first three states increases by an additional 6 percentage point (55%) and the initial consumption further increases relative to the uniform IPR policy by 30%. More interestingly, the innovation rate of a one-step leader now increases from 0.15 to 0.51 (relative to the uniform policy case) and the growth rate increases back to 2.5%. It is noteworthy that these gains are achieved by providing stronger protections to more advanced firms, and thus exploiting the trickle-down effect. For example, under the optimal policy one-step leader is caught up seven times more frequently than a five-step leader due to patent expiration.

7.3 Compulsory Licensing

In this subsection, we introduce (compulsory) licensing. Several recent empirical papers suggest that licensing has a significant positive impact on firm innovation (e.g., Moser and Voena, 2011, Almeida and Fernandes, 2008). Consistent with these findings, we model licensing as a way of generating knowledge spillovers to the licensee firm. In particular, in addition to independent R&D to proceed one step in the quality ladder, followers can also close all intervening steps by reverse-engineering the current leading-edge technology. But this is only possible by making use of the knowledge generated by the leading-edge technology, and the follower will have to pay a prespecified license fee $\hat{\zeta}_n(t) \geq 0$ to the leader. The licensing decision of the follower $-i$ is denoted by $a_{-i}(j,t) = 1$ ($a_{-i}(j,t) = 0$ corresponds to independent R&D). Throughout,

we allow $a_{-i}(j,t) \in [0,1]$ for mathematical convenience. The fees in question are compulsory license fees imposed by policy and are state-dependent, and thus we represent them as:

$$\hat{\boldsymbol{\zeta}}(t) : \mathbb{N} \to \mathbb{R}_+ \cup \{+\infty\}.$$

Note that $\hat{\boldsymbol{\zeta}}(t) \equiv \left\{\hat{\zeta}_1(t), \hat{\zeta}_2(t), ...\right\}$ is a function of time. This is natural, since in a growing economy, license fees should not remain constant. As in (28), in what follows we assume that license fees are also scaled up by GDP, so that $\zeta_n \equiv \hat{\zeta}_n(t)/Y(t)$, to keep the equilibrium stationary. We discuss voluntary licensing below.

7.3.1 Value Functions with Compulsory Licensing

With a similar reasoning to before, relevant value functions in this case can be written as

$$\rho v_n = \max_{x_n \geq 0} \left\{ \begin{array}{c} \left(1-\lambda^{-n}\right) - \omega^* G(x_n) + x_n [v_{n+1} - v_n] \\ +a^*_{-n} x^*_{-n} [v_0 - v_n + \zeta_n] + \left(1 - a^*_{-n}\right) x^*_{-n} [v_{n-1} - v_n] + \eta_n [v_0 - v_n] \end{array} \right\} \text{ for } n \in \mathbb{N},$$

where a^*_{-n} is the equilibrium value of licensing decision by a follower that is n steps behind, and ζ_n is the license fee that it has to pay. The value for neck-and-neck firms remain unchanged while the values for followers becomes

$$\rho v_{-n} = \max_{x_{-n} \geq 0, a_{-n} \in [0,1]} \left\{ \begin{array}{c} -\omega^* G(x_{-n}) + a_{-n} x_{-n} [v_0 - v_{-n} - \zeta_n] \\ +(1 - a_{-n}) x_{-n} [v_{-n+1} - v_{-n}] + x^*_n [v_{-n-1} - v_{-n}] + \eta_n [v_0 - v_{-n}] \end{array} \right\} \text{ for } n \in \mathbb{N},$$

Note that licensing $a_{-n} \in [0,1]$ is the new additional decision variable of the follower.

Full IPR protection in this case corresponds to prohibitively high licensing fees, i.e., $\zeta_n = \infty$ for all n, and as in the previous subsection, patent protection has expected duration of 50 years ($\eta = 0.02$). Therefore, the results in this case will be identical to those reported for full protection in the previous subsection (column 1 of Table 2). This is indeed the case; these results are repeated in column 1 of Table 3 for ease of comparison with the remaining results in this table.

[Table 3 here]

7.3.2 Optimal Uniform IPR Protection

Uniform compulsory licensing policy now corresponds to $\zeta_n = \zeta^* \geq 0$. The results under the optimal choice of such uniform compulsory licensing policy are reported in the second column of Table 3. This optimal policy involves $\zeta^* = 1.61$, which is more than half of the surplus that a three-step follower generates from licensing, $v_0 - v_{-3} = 2.9$.

Since this type of licensing allows for more frequent catch-up by followers, a greater share of industries are now in tight competition. In particular, the total share of industries with one or two step gaps goes up to 66% (this number was 8% under full protection). This again corresponds to a powerful composition effect and generates a significant reduction in the average mark-up and a corresponding increase in initial consumption. However, consistent with our previous results, this type of uniform licensing again generates a significant disincentive effect on technology leaders. In particular, more frequent catch-up implies a shorter durations of positive profits. As a result, innovation incentives are reduced; the innovation rate of a one-step leader is now 0.43 instead of 1.1 and the average growth rate declines from 2.5% to 2.1%.

7.3.3 Optimal State-Dependent IPR

As in our previous exercises, the negative incentive effects of uniform relaxations of IPR protection are rectified when policy is state dependent. Optimal state-dependent policy has in fact qualitatively very similar pattern to those reported above. Most importantly, column 3 of Table 3 shows that optimal state-dependent policy provides greater protection to technology leaders that are more advanced. For example, while a two-step leader receives a license fee of $\zeta_2^* = 1.5$, a five-step leader receives more than its double, $\zeta_5^* = 3.3$. Given this pattern, the *trickle-down effect* is again at work and generates positive innovation incentives: the innovation rate of a one-step leader increases to $x_1^* = 0.46$ and the aggregate growth rate goes back to 2.5% from 2.1%. This positive gain is generated without sacrificing the composition effect. Under this policy, 50% of total industries operate with a technology gap less than two and the initial consumption $C(0)$ is now even higher than under uniform policy (by 40%).

7.3.4 Compulsory Versus Bargained License Fees

The analysis so far has characterized the steady-state equilibrium for a given sequence of license fees $\boldsymbol{\zeta}$, implicitly assumed to be determined by IPR policy—i.e., these fees correspond to compulsory licensing fees for intellectual property that has been patented. This, therefore, corresponds to a world in which once a company patents an innovation, the knowledge embedded in this innovation can be used by its competitors as long as they pay a prespecified license fee.

One may also wish to consider an alternative world in which license fees are determined by bilateral bargaining. To characterize the equilibrium in such a world, one must first conduct exactly the same analysis as we have done in this subsection. In other words, one must

characterize the equilibrium for a given sequence of license fees, and then taking the license fees agreed by other firms as given, one can consider the bargaining problem between a leader and a follower. In general, there may or may not exist feasible voluntary license fees that the follower and the leader can bargain to (such voluntary agreements may be infeasible even if compulsory licensing is beneficial, since consumers also benefit from licensing).

[FIGURE 6 HERE]

Figure 6 plots the value of licensing to a follower in an industry with an n-step gap, $v_0 - v_{-n+1}$, and absolute value of the loss to the leader in the same industry, $|v_{n-1} - v_0|$ (with full protection as in column 1 of Tables 2-5).[30] The overall pattern is that the latter number is unambiguously greater than the former, which implies that voluntary licensing will not be beneficial in this environment. Therefore, compulsory licensing plays a useful role that bilateral licensing agreements between leaders and followers could not achieve, and is thus a useful policy tool. In addition, our analysis shows that compulsory licensing will be useful for welfare precisely when it is state dependent.

7.4 Leapfrogging and Infringement under Slow Catch-up

Finally, we allow the follower to engage in frontier R&D and "leapfrog" the technology leader. This exercise is useful for two reasons. First, the models analyzed so far do not allow R&D by followers to directly contribute to aggregate growth. One might conjecture that this feature strengthens the *trickle-down* effect. Second, frontier R&D and leapfrogging by followers will allow us to introduce another relevant and important dimension of IPR policy, patent infringement fees.

Suppose, now, that followers can undertake two types of R&D. The first, which is what we have focused on so far, is *catch-up R&D*, corresponding to R&D directed at discovering an *alternative* way of performing the same task as the current leading-edge technology. Catch-up R&D improves the technology of the follower by one step as before. The alternative, *frontier R&D*, involves followers improving the current leading-edge technology. If this type of R&D succeeds, the follower will have improved the leading-edge technology. However, following such an event, the follower will be judged (e.g., by courts) to have infringed the patent of technology

[30] Without licensing, the change in follower's value is $v_{-n+1} - v_{-n}$. Since licensing takes the follower to v_0, the change due to licensing is $v_0 - v_{-n+1}$. Similar reasoning applies to the leader's loss.

Note also that Figure 1 has no value for $n = 0, 1$ since neck-and-neck firms and one-step followers have no surplus to generate through licensing.

leader with probability $\tau \in (0,1)$ and will be required to pay a prespecified infringement penalty (fee) $\hat{\vartheta}_n \geq 0$ to the leader. The infringement fees are also state dependent and represented by:

$$\hat{\boldsymbol{\vartheta}}(t) : \mathbb{N} \to \mathbb{R}_+ \cup \{+\infty\},$$

and we again adopt the normalization $\vartheta_n \equiv \hat{\vartheta}_n(t)/Y(t)$, and denote the Poisson arrival rate of innovation by catch-up R&D and frontier R&D by x_n^c and x_n^f, respectively. Then the new value of an n-step leader takes the following form:

$$\rho v_n = \max_{x_n \geq 0} \left\{ \begin{array}{c} \left(1-\lambda^{-n}\right) - \omega^* G(x_n) + x_n [v_{n+1} - v_n] \\ +x_{-n}^{c*} [v_{n-1} - v_n] + x_{-n}^{f*} [v_{-1} - v_n + \tau\vartheta_n] + \eta_n [v_0 - v_n] \end{array} \right\} \text{ for } n \in \mathbb{N},$$

The main difference in this equation is that the follower has two different arrival rates of innovation. If the follower is successful with frontier R&D, the current leader falls one step behind the follower. However, in this event, with probability τ, it receives an infringement fee of ϑ_n. With a similar reasoning, the value of an n-step follower now becomes:

$$\rho v_{-n} = \max_{x_{-n}^c \geq 0, x_{-n}^f \geq 0} \left\{ \begin{array}{c} -\omega^* G(x_{-n}) + x_{-n}^c [v_{-n+1} - v_{-n}] + x_{-n}^f [v_1 - v_{-n} - \tau\vartheta_n] \\ +x_n^* [v_{-n-1} - v_{-n}] + \eta_n [v_0 - v_{-n}] \end{array} \right\} \text{ for } n \in \mathbb{N}.$$

The value of a neck-and-neck firm is unchanged.

The quantitative analysis requires an empirical estimate for τ. Lanjouw and Schankerman (2001) report that around 10% of the US utility patents are filed for infringement. We therefore set $\tau = 0.1$.

Note also that since the followers now improve the technology frontier through frontier R&D, the aggregate growth rate becomes

$$g^* = \ln \lambda \left[2\mu_0^* x_0^* + \sum_{n=1}^{\infty} \mu_n^* \left(x_n^* + x_{-n}^{f*} \right) \right]. \tag{49}$$

Full protection in this case corresponds to infinite patent infringement fees, i.e., $\vartheta_n = \infty$, and given the same parameter choices as before, will be identical to column 1 of Table 2. We repeat these results in column 1 of Table 4 for ease of comparison with the rest of the table.

[Table 4 here]

7.4.1 Optimal Uniform IPR Protection

In the uniform policy case, we set $\vartheta_n = \vartheta \geq 0$. Column 2 of Table 4 shows that the optimal uniform policy in this case is $\vartheta^* = 14$. Recall that when a follower undertakes frontier innovation, the probability the that it will have to make this payment is $\tau = 0.1$. Therefore the

expected infringement payment is $\tau \times \vartheta^* = 1.4$ which is more than half of the surplus that a three-step follower generates out of leapfrogging, $v_1 - v_{-3} = 2.7$.

Column 2 also shows that under this policy, followers undertake more frontier R&D ($x_{-1}^{*f} = 0.23$) than catch-up R&D $\left(x_{-1}^{*c} = 0.15\right)$. Parallel to the previous uniform policies, the shorter duration of monopoly position resulting from innovation reduces innovation incentives. For example, one-step leaders now innovate at the rate 0.3 instead of 1.1. However, despite this disincentive effect, the growth rate increases slightly because leapfrogging allows followers to directly contribute to aggregate growth, as shown by equation (49).

Column 2 also shows that the share of industries in one-step gap is now much larger, $\mu_1 = 0.42$. This is because leapfrogging puts the follower one-step ahead of the previous leader. Thanks to this effect, optimal uniform IPR protection achieves lower average mark-up and higher initial consumption as well as higher growth.

7.4.2 Optimal State-Dependent IPR

State-dependent IPR policy once again exploits the trickle-down effect and creates positive incentive effects on innovation. The form of state-dependent policy is the same as before: technologically more advanced leaders receive more protection in the form of higher fees when followers infringe their patents. While a two-step leader receives $\vartheta_1^* = 18.1$ in case of infringement, a five-step leader receives more than double of this fee, $\vartheta_5^* = 43.7$. In expectation, a three-step follower pays almost 3/4th of the surplus that it generates from leapfrogging ($\tau \times \vartheta_3^* = 3.1$ versus $v_1 - v_{-3} = 4.1$). As a result of this pattern, state-dependent policy not only generates a greater welfare gain in terms of the initial consumption ($C(0)$ is now approximately twice the level under the optimal uniform policy), but it also exploits the trickle-down effect and increases the equilibrium growth rate by an additional 0.5 percentage point relative to the uniform policy.

7.5 Patent Length, Compulsory Licensing and Infringements Fees under Slow Catch-up

In this subsection, we investigate the slow catch-up environment when all three IPR policies are simultaneously present. We do not repeat the value functions to save space.

[Table 5 here]

Table 5 first shows our benchmark full protection economy in the first column. The second and third columns report the optimal uniform and state-dependent policies with all three types

of policies present. The results are very similar to those reported in subsection 7.4 (with only leapfrogging), except that the patent lengths are now set to infinity ($\eta_n = 0$). The optimal IPR policy in this case involves infinitely long patents with prohibitively high compulsory license fees. The only dimension in which IPR protection is not full is because of moderate infringement fees, which permit followers to undertake frontier R&D and leapfrog technology leaders.

Most importantly for our focus, column 3 again shows the benefits of state-dependent IPR policy. This policy again provides greater protection for technology leaders and exploits the trickle-down effect. As a result, initial consumption is approximately twice the level under uniform IPR and innovation incentives are stronger, and the long-run growth rate increases from 2.5% to 3.3%.

7.6 Robustness Checks

Table 6 shows that the patterns documented in Table 5, particularly the gains from state-dependent policy and the major role played by the trickle-down effect, are robust for reasonable changes in parameter values.

[Table 6 here]

In this table, in each column we change one of the two parameters λ and γ (increasing or reducing λ to 1.2 or 1.01, and increasing or reducing γ to 0.6 or 0.1). In each case, we also change the parameter B in equation (45) to ensure the growth rate of the benchmark economy with full IPR protection is the same as in our initial baseline economy, $g^* = 1.86\%$.

To save space, we only show the results from the optimal state-dependent policies. Table 6 shows that the qualitative patterns in Table 5 are relatively robust. In all cases, optimal state-dependent IPR is shaped by the trickle-down effect. In all of the various parameterizations we have considered (and with different combinations of policies), there is little protection provided to technology leaders that are one-step ahead, but IPR protection grows as the technology gap increases. This is the typical pattern implied by the trickle-down effect. In addition, in all cases when all three forms of policy are incorporated, optimal IPR policy provides patents of infinite duration and prohibitively high compulsory licensing fees, but deviates from full IPR protection by imposing moderate levels of infringement fees. Most importantly for us, in all cases, these infringements fees are state dependent and provide greater protection to technologically more advanced leaders.

8 Conclusions

In this paper, we emphasized the importance of dynamic interactions between IPR protection and competition for understanding the structure of optimal IPR policy. Our main result highlights the importance of a new and powerful effect, *the trickle-down effect*, which implies that protection given to companies with significant technological leads over their rivals also dynamically incentivizes companies with more limited technological leads—as further innovation will not only increase their productivity but also grant them additional IPR protection. This new effect implies that optimal IPR policy should be *state-dependent* and provide greater protection to companies with significant technological leads and only limited IPR protection for those without.

To systematically investigate these issues, we developed a dynamic general equilibrium framework with cumulative (step-by-step) innovations. In each industry, technology leaders innovate in order to widen the gap between themselves and the followers, which enables them to charge higher markups. Followers innovate to catch up with or surpass the technology leaders in their industry (by undertaking "frontier R&D"), and can also license the technology of leaders. IPR policy regulates the length of patents, whether licensing is possible and the size of patent infringement fees.

We provided existence and characterization results, and a quantitative analysis of the form of "optimal" (welfare-maximizing) IPR policy. In several different environments and under different parameter values, we consistently found that the trickle-down effect is present and powerful. It implies that optimal IPR should be state-dependent and should provide greater protection to firms with greater technological lead over their rivals. In our benchmark parameterization, for example, optimal IPR policy increases the growth rate of the economy from 1.86% to 2.04%, and does so while also significantly increasing initial consumption (and in fact reducing the overall amount of resources allocated to the R&D sector). We also showed that similar qualitative and quantitative results are obtained when followers catch up with technology leaders only slowly. In this extended environment, we also investigated the form of optimal compulsory licensing fees and patent infringement fees, and found them to be similarly state-dependent (in a way that provides greater protection to firms that are technologically more advanced relative to their rivals). These extensions further showed that compulsory licensing, which allows followers to build on the leading-edge technology in return of a license fee, also has a major impact on the equilibrium growth rate.

Our main results go against a naïve intuition that providing less protection to techno-

logically more advanced firms is socially beneficial because it would exploit a composition effect (bringing firms that are furthest apart into "neck and neck" competition to both reduce markups and increase R&D which results from tight competition). This naïve intuition is not correct precisely because of the trickle-down effect we emphasized above. The trickle-down effect implies that providing greater protection to sufficiently advanced technology leaders not only increases their R&D efforts but also raises the R&D efforts of all technology leaders that are less advanced than this level. This is because the reward to innovation now includes the greater protection that they will receive once they reach this higher level of technology. Our analysis and results suggest that in addition to the reasoning based on the static trade-off between IPR protection and competition, the trickle-down effect should also be factored into policy analysis, and naturally calls for future empirical work to estimate its empirical magnitude.

In this context, it should be emphasized that our objective in this paper has not been to derive practical policy prescriptions. There is little doubt that our model is simplified, excludes a whole host of important factors, and ignores potential limitations on the form and complexity of IPR policies. Nevertheless, we believe that our results demonstrate a range of robust and new effects that should be further investigated in future work.

More generally, the analysis in this paper suggests that a move to a richer menu of IPR policies, in particular, a move towards optimal state-dependent policies, may significantly increase innovation, economic growth and welfare. The results also show that the form of optimal IPR policy may depend on the industry structure (and the technology of catch-up within the industry). The next step in this line of research should be to investigate the robustness of these effects in different models of industry dynamics. It would also be useful to study whether the relationship between the form of optimal IPR policy and industry structure suggested by our analysis also applies when variation in industry structure has other sources (for example, differences in the extent of fixed costs or demand structure causing differential gaps between technology leaders and followers across industries). The most important area for future work is a detailed empirical investigation of the form of optimal IPR policy, using both better estimates of the effects of IPR policy on innovation rates and also structural models that would enable the evaluation of the effects of different policies on equilibrium growth and welfare.

Appendix: Proofs

Derivation of Optimal R&D Decisions in the Partial Equilibrium Model

Since the costs are linear, optimal R&D decisions imposed that, in equilibrium,

$$v_{n+1} - v_n = \phi, \text{ for each } n \in \{-2, -1, 0, 1\}. \tag{50}$$

Combining this result with equation (1) gives the value of a two-step follower is

$$v_{-2} = \frac{\pi_{-2} + 2\phi\eta_2}{r}.$$

The previous equation, together with (50) implies

$$v_n = v_{-2} + \phi(n+2) = \frac{\pi_{-2} + 2\phi\eta_2}{r} + \phi(n+2), \text{ for each } n \in \{-1, 0, 1, 2\}. \tag{51}$$

Now we can use the value of v_2 to solve for x^*_{-2} from equation (1). Similarly, combining (51) with (2) gives the value of x^*_{-1}; (51) with (3) gives x^*_0. Finally, combining (51) with (4) gives the equilibrium value of x^*_1. ■

Derivation of Equation (27)

Fix the equilibrium R&D policies of other firms, $x^*_{-n}(t)$, the equilibrium interest and wage rates, $r^*(t)$ and $w^*(t)$, and equilibrium profits $\{\Pi^*_n(t)\}_{n=1}^{\infty}$. Then the value of the firm that is n steps ahead at time t can be written as:

$$\begin{aligned} V_n(t) \quad = \quad & \max_{x_n(t)} \{[\Pi^*_n(t) - w^*(t) G(x_n(t))] \Delta t + o(\Delta t) \\ & + \exp(-r^*(t+\Delta t)\Delta t) \left[\begin{array}{c} (x_n(t)\Delta t + o(\Delta t)) V_{n+1}(t+\Delta t) \\ + (\eta_n \Delta t + x^*_{-n}(t)\Delta t + o(\Delta t)) V_0(t+\Delta t) \\ + (1 - x_n(t)\Delta t - \eta_n \Delta t - x^*_{-n}(t)\Delta t - o(\Delta t)) V_n(t+\Delta t) \end{array} \right] \Bigg\}. \end{aligned} \tag{52}$$

The first part of this expression is the flow profits minus R&D expenditures during a time interval of length Δt. The second part is the continuation value after this interval has elapsed. $V_{n+1}(t)$ and $V_0(t)$ are defined as net present discounted values for a leader that is $n+1$ steps ahead and a firm in an industry that is neck-and-neck (i.e., $n = 0$). The second part of the expression uses the fact that in a short time interval Δt, the probability of innovation by the leader is $x_n(t)\Delta t + o(\Delta t)$, where $o(\Delta t)$ again denotes second-order terms. This explains the first line of the continuation value. For the remainder of the continuation value, note that the probability that the follower will catch up with the leader is $[\eta_n + x^*_{-n}(t)]\Delta t + o(\Delta t)$. Finally, the last line applies when no R&D effort is successful and patents continue to be enforced, so that the technology gap remains at n steps. Now, subtract $V_n(t)$ from both sides, divide everything by Δt, and take the limit as $\Delta t \to 0$ to obtain (27). ■

Proof of Proposition 3

Equations (24) and (26) imply

$$Y(t) = \frac{w(t)}{\omega(t)} = \frac{Q(t)\lambda^{-\sum_{n=0}^{\infty} n\mu^*_n(t)}}{\omega(t)}.$$

Since $\omega(t) = \omega^*$ and $\{\mu_n^*\}_{n=0}^{\infty}$ are constant in steady state, $Y(t)$ grows at the same rate as $Q(t)$. Therefore,

$$g^* = \lim_{\Delta t \to 0} \frac{\ln Q(t + \Delta t) - \ln Q(t)}{\Delta t}.$$

Now note the following: during an interval of length Δt (i) in the fraction μ_n^* of the industries with technology gap $n \geq 1$ the leaders innovate at a rate $x_n^* \Delta t + o(\Delta t)$; (ii) in the fraction μ_0^* of the industries with technology gap of $n = 0$, both firms innovate, so that the total innovation rate is $2x_0^* \Delta t + o(\Delta t)$); and (iii) each innovation increase productivity by a factor λ. Combining these observations, we have

$$\ln Q(t + \Delta t) = \ln Q(t) + \ln \lambda \left[2\mu_0^* x_0^* \Delta t + \sum_{n=1}^{\infty} \mu_n^* x_n^* \Delta t + o(\Delta t) \right].$$

Subtracting $\ln Q(t)$, dividing by Δt and taking the limit $\Delta t \to 0$ gives (39). ■

Proof of Proposition 4

We prove this proposition in four parts. (1) Existence of a steady-state equilibrium. (2) Properties of the sequence of value functions. (3) Properties of the sequence of R&D decisions. (4) Uniqueness of an invariant distribution given R&D policies.

Part 1: Existence of a Steady-State Equilibrium.

First, note that each x_n belongs to a compact interval $[0, \bar{x}]$, where $\bar{x}$ is the maximal flow rate of innovation defined in (16) above. Now fix a labor share $\tilde{\omega} \in [0, 1]$ and a sequence $\langle \tilde{\mathbf{x}} \rangle$ of (Markovian) steady-state strategies for all other firms in the economy, and consider the dynamic optimization problem of a single firm. Our first result characterizes this problem and shows that given some $\mathbf{z} \equiv \langle \tilde{\omega}, \tilde{\mathbf{x}} \rangle$, the value function of an individual firm is uniquely determined, while its optimal R&D choices are given by a convex-valued correspondence. In what follows, we denote sets and correspondences by uppercase letters and refer to their elements by lowercase letters, e.g., $x_n(\mathbf{z}) \in X_n[\mathbf{z}]$.

Lemma 1 *Consider a uniform IPR policy $\boldsymbol{\eta}^{uni}$, and suppose that the labor share and the R&D policies of all other firms are given by $\mathbf{z} = \langle \tilde{\omega}, \tilde{\mathbf{x}} \rangle$. Then the dynamic optimization problem of an individual firm leads to a unique value function $\mathbf{v}[\mathbf{z}] : \{-1\} \cup \mathbb{Z}_+ \to \mathbb{R}_+$ and optimal R&D policy $\hat{\mathbf{X}}[\mathbf{z}] : \{-1\} \cup \mathbb{Z}_+ \rightrightarrows [0, \bar{x}]$ is compact and convex-valued for each $\mathbf{z} \in \mathbf{Z}$ and upper hemi-continuous in $\mathbf{z}$ (where $\mathbf{v}[\mathbf{z}] \equiv \{v_n[\mathbf{z}]\}_{n=-1}^{\infty}$ and $\hat{\mathbf{X}}[\mathbf{z}] \equiv \left\{\hat{X}_n[\mathbf{z}]\right\}_{n=-1}^{\infty}$).*

Proof. Fix $\mathbf{z} = \left\langle \tilde{\omega}, \{\tilde{x}_n\}_{n=-1}^{\infty} \right\rangle$, and consider the optimization problem of a representative firm, written recursively as:

$$\begin{aligned}
\rho v_n &= \max_{x_n \in [0, \bar{x}]} \left\{ \left(1 - \lambda^{-n}\right) - \tilde{\omega} G(x_n) + x_n [v_{n+1} - v_n] + \tilde{x}_{-1} [v_0 - v_n] + \eta [v_0 - v_n] \right\} \text{ for } n \in \mathbb{N} \\
\rho v_0 &= \max_{x_0 \in [0, \bar{x}]} \left\{ -\tilde{\omega} G(x_0) + x_0 [v_1 - v_0] + \tilde{x}_0 [v_{-1} - v_0] \right\} \\
\rho v_{-1} &= \max_{x_{-1} \in [0, \bar{x}]} \left\{ -\tilde{\omega} G(x_0) + x_{-1} [v_0 - v_{-1}] + \eta [v_0 - v_{-1}] \right\}.
\end{aligned}$$

We now transform this dynamic optimization problem into a form that can be represented as a contraction mapping using the method of "uniformization" (see, for example, Ross, 1996, Chapter 5). Let $\tilde{\boldsymbol{\xi}} = \{\tilde{x}_n\}_{n=-1}^{\infty}$ and $p_{n,n'}\left(\boldsymbol{\xi} \mid \tilde{\boldsymbol{\xi}}\right)$ be the probability that the next state will be n' starting with state n when the firm in question chooses policies $\boldsymbol{\xi} \equiv \{x_n\}_{n=-1}^{\infty}$ and the R&D policy of other firms is given by $\tilde{\boldsymbol{\xi}}$. Using the fact that, because of uniform IPR policy, $x_{-n} = x_{-1}$ for all $n \in \mathbb{N}$, these transition

probabilities can be written as:

$p_{-1,0}\left(\boldsymbol{\xi} \mid \tilde{\boldsymbol{\xi}}\right) = \frac{x_{-1}+\eta}{x_n+x_{-1}+\eta}$	$p_{n,0}\left(\boldsymbol{\xi} \mid \tilde{\boldsymbol{\xi}}\right) = \frac{\tilde{x}_{-1}+\eta}{x_n+\tilde{x}_{-1}+\eta}$
$p_{0,-1}\left(\boldsymbol{\xi} \mid \tilde{\boldsymbol{\xi}}\right) = \frac{\tilde{x}_0}{x_0+\tilde{x}_0}$	$p_{n,n+1}\left(\boldsymbol{\xi} \mid \tilde{\boldsymbol{\xi}}\right) = \frac{x_n}{x_n+\tilde{x}_{-1}+\eta}$
$p_{0,1}\left(\boldsymbol{\xi} \mid \tilde{\boldsymbol{\xi}}\right) = \frac{x_0}{x_0+\tilde{x}_0}$	

.

Uniformization involves adding fictitious transitions from a state into itself, which do not change the value of the program, but allow us to represent the optimization problem as a contraction. For this purpose, define the transition rates ψ_n as

$$\psi_n\left(\boldsymbol{\xi} \mid \tilde{\boldsymbol{\xi}}\right) = \begin{cases} x_n + x_{-1} + \eta \text{ for } n \in \{1, 2, ...\} \\ x_{-1} + \eta \text{ for } n = -1 \\ 2x_n \text{ for } n = 0 \end{cases}.$$

These transition rates are finite since $\psi_n\left(\boldsymbol{\xi} \mid \tilde{\boldsymbol{\xi}}\right) \leq \psi \equiv 2\bar{x} + \eta < \infty$ for all n, where $\bar{x}$ is the maximal flow rate of innovation defined in (16) in the text (both $\bar{x}$ and η are finite by assumption).

Now following equation (5.8.3) in Ross (1996), we can use these transition rates and define the new transition probabilities (including the fictitious transitions from a state to itself) as:

$$\tilde{p}_{n,n'}\left(\boldsymbol{\xi} \mid \tilde{\boldsymbol{\xi}}\right) = \begin{cases} \frac{\psi_n(\boldsymbol{\xi}|\tilde{\boldsymbol{\xi}})}{\psi} p_{n,n'}\left(\boldsymbol{\xi} \mid \tilde{\boldsymbol{\xi}}\right) & \text{if } n \neq n' \\ 1 - \frac{\psi_n(\boldsymbol{\xi}|\tilde{\boldsymbol{\xi}})}{\psi} & \text{if } n = n' \end{cases}.$$

This yields equivalent transition probabilities

$\tilde{p}_{-1,-1}\left(\boldsymbol{\xi} \mid \tilde{\boldsymbol{\xi}}\right) = 1 - \frac{x_{-1}+\eta}{2\bar{x}+\eta}$	$\tilde{p}_{-1,0}\left(\boldsymbol{\xi} \mid \tilde{\boldsymbol{\xi}}\right) = \frac{x_{-1}+\eta}{2\bar{x}+\eta}$	$\tilde{p}_{0,1}\left(\boldsymbol{\xi} \mid \tilde{\boldsymbol{\xi}}\right) = \frac{x_0}{2\bar{x}+\eta}$
$\tilde{p}_{0,-1}\left(\boldsymbol{\xi} \mid \tilde{\boldsymbol{\xi}}\right) = \frac{\tilde{x}_0}{2\bar{x}+\eta}$	$\tilde{p}_{0,0}\left(\boldsymbol{\xi} \mid \tilde{\boldsymbol{\xi}}\right) = 1 - \frac{x_0+\tilde{x}_0}{2\bar{x}+\eta}$	$\tilde{p}_{n,n+1}\left(\boldsymbol{\xi} \mid \tilde{\boldsymbol{\xi}}\right) = \frac{x_n}{2\bar{x}+\eta}$
$\tilde{p}_{n,0}\left(\boldsymbol{\xi} \mid \tilde{\boldsymbol{\xi}}\right) = \frac{\tilde{x}_{-1}+\eta}{2\bar{x}+\eta}$	$\tilde{p}_{n,n}\left(\boldsymbol{\xi} \mid \tilde{\boldsymbol{\xi}}\right) = 1 - \frac{x_n+\tilde{x}_{-1}+\eta}{2\bar{x}+\eta}$	

,

and also defines an effective discount factor β given by

$$\beta \equiv \frac{\psi}{\rho + \psi} = \frac{2\bar{x} + \eta}{\rho + 2\bar{x} + \eta}.$$

Also let the per period return function (profit net of R&D expenditures) be

$$\hat{\Pi}_n(x_n) = \begin{cases} \frac{1-\lambda^{-n}-\tilde{\omega}G(x_n)}{\rho+2\bar{x}+\eta} \text{ if } n \geq 1 \\ \frac{-\tilde{\omega}G(x_n)}{\rho+2\bar{x}+\eta} \text{ otherwise} \end{cases}. \tag{53}$$

Using these transformations, the dynamic optimization problem can be written as:

$$\begin{aligned} v_n &= \max_{x_n}\left\{\hat{\Pi}_n(x_n) + \beta\sum_{n'}\tilde{p}_{n,n'}\left(\xi_n|\,\tilde{\boldsymbol{\xi}}\right)\tilde{v}_{n'}\right\}, \text{ for all } n \in \mathbb{Z}, \\ &\equiv T\tilde{v}_n, \text{ for all } n \in \mathbb{Z}. \end{aligned} \tag{54}$$

where $\mathbf{v} \equiv \{v_n\}_{n=-1}^{\infty}$ and the second line defines the operator T, mapping from the space of functions $\mathbf{V} \equiv \{\mathbf{v} : \{-1\} \cup \mathbb{Z}_+ \to \mathbb{R}_+\}$ into itself. T is clearly a contraction mapping. The innovation rates $\{\tilde{x}_n\}_{n=-1}^{\infty}$ are upper hemi-continuous therefore $\tilde{\mathbf{p}} : \{-1\} \cup \mathbb{Z}_+ \times \{-1\} \cup \mathbb{Z}_+ \rightrightarrows [0,1]$ is upper-hemicontinuous and forms a multivalued stochastic kernel. Then Proposition 2.2 in Blume (1982)

implies that $\tilde{\mathbf{p}}$ has the Feller property. Thus, for given $\mathbf{z} = \langle \tilde{\omega}, \{\tilde{x}_n\}_{n=-1}^{\infty} \rangle$, T possesses a unique fixed point $\mathbf{v}^* \equiv \{v_n^*\}_{n=-1}^{\infty}$ (e.g., Stokey, Lucas and Prescott, 1989).

Moreover, $x_n \in [0, \bar{x}]$, and v_n for each $n = -1, 0, 1, ...$ given by the right-hand side of (54) is continuous in x_n, so Berge's Maximum Theorem (Aliprantis and Border, 1999, Theorem 16.31, p. 539) implies that the set of maximizers $\left\{\hat{X}_n\right\}_{n=-1}^{\infty}$ exists, is nonempty and compact-valued for each $\mathbf{z}$ and is upper hemi-continuous in $\mathbf{z} = \langle \tilde{\omega}, \{\tilde{x}_n\}_{n=-1}^{\infty} \rangle$. Moreover, concavity of v_n in x_n for each $n = -1, 0, 1, ...$ implies that $\left\{\hat{X}_n\right\}_{n=-1}^{\infty}$ is also convex-valued for each $\mathbf{z}$, completing the proof. ■

Now let us start with an arbitrary $\mathbf{z} \equiv \langle \tilde{\omega}, \tilde{\mathbf{x}} \rangle \in \mathbf{Z} \equiv [0,1] \times [0, \bar{x}]^{\infty}$. From Lemma 1, this $\mathbf{z}$ is mapped into optimal R&D decision sets $\hat{\mathbf{X}}[\mathbf{z}]$, where $\hat{x}_n[\mathbf{z}] \in \hat{X}_n[\mathbf{z}]$. From R&D policies $\tilde{\mathbf{x}}$, we calculate $\boldsymbol{\mu}[\tilde{\mathbf{x}}] \equiv \{\mu_n[\tilde{\mathbf{x}}]\}_{n=0}^{\infty}$ using equations (35), (36) and (37). Then we can rewrite the labor market clearing condition (38) as

$$\begin{aligned} \omega &= \min\left\{\sum_{n=0}^{\infty} \mu_n \left[\frac{1}{\lambda^n} + G(\tilde{x}_n)\,\tilde{\omega} + G(\tilde{x}_{-n})\right]\tilde{\omega}; 1\right\}, \\ &\equiv \varphi(\tilde{\omega}, \tilde{\mathbf{x}}) \end{aligned} \tag{55}$$

where due to uniform IPR, $\hat{x}_{-n} = \hat{x}_{-1}$ for all $n > 0$. Next, define the mapping (correspondence)

$$\Phi[\mathbf{z}] \equiv \left(\varphi(\mathbf{z}),\ \hat{\mathbf{X}}[\mathbf{z}]\right)$$

, which maps $\mathbf{Z}$ into itself, that is,

$$\Phi\colon \mathbf{Z} \rightrightarrows \mathbf{Z}. \tag{56}$$

That Φ maps $\mathbf{Z}$ into itself follows since $\mathbf{z} \in \mathbf{Z}$ consists of $\tilde{\mathbf{x}} \in [0, \bar{x}]^{\infty}$ and $\tilde{\omega} \in [0,1]$, and the image of $\mathbf{z}$ under Φ consists of $\hat{\mathbf{x}} \in [0, \bar{x}]^{\infty}$, and moreover, (55) is clearly in $[0,1]$ (since the right-hand side is nonnegative and bounded above by 1). Finally, from Lemma 1, $\hat{X}_n[\mathbf{z}]$ is compact and convex-valued for each $\mathbf{z} \in \mathbf{Z}$, and also upper hemi-continuous in $\mathbf{z}$, and φ is continuous. Using this construction, we can establish the existence of a steady-state equilibrium as follows.

We first show that the mapping $\Phi\colon \mathbf{Z} \rightrightarrows \mathbf{Z}$ constructed in (56) has a fixed point, and then establish that when $G'^{-1}\left(\left(1-\lambda^{-1}\right)/(\rho+\eta)\right) > 0$ this fixed point corresponds to a steady state with $\omega^* < 1$. First, it has already been established that Φ maps $\mathbf{Z}$ into itself. We next show that $\mathbf{Z}$ is compact in the product topology and is a subset of a locally convex Hausdorff space. The first part follows from the fact that $\mathbf{Z}$ can be written as the Cartesian product of compact subsets, $\mathbf{Z} = [0,1] \times \prod_{n=-1}^{\infty} [0, \bar{x}]$. Then by Tychonoff's Theorem (e.g., Aliprantis and Border, 1999, Theorem 2.57, p. 52; Kelley, 1955, p. 143), $\mathbf{Z}$ is compact in the product topology. Moreover, $\mathbf{Z}$ is clearly nonempty and also convex, since for any $\mathbf{z}, \mathbf{z}' \in \mathbf{Z}$ and $\lambda \in [0,1]$, we have $\lambda\mathbf{z} + (1-\lambda)\mathbf{z}' \in \mathbf{Z}$. Finally, since $\mathbf{Z}$ is a product of intervals on the real line, it is a subset of a locally convex Hausdorff space (see Aliprantis and Border, 1999, Lemma 5.54, p. 192).

Next, φ is a continuous function from $\mathbf{Z}$ into $[0,1]$ and from Lemma 1, $\hat{X}_n(\mathbf{z})$ for $n \in \{-1\} \cup \mathbb{Z}_+$ is upper hemi-continuous in $\mathbf{z}$. Consequently, $\Phi \equiv \left\langle \varphi[\mathbf{z}],\ \hat{\mathbf{X}}[\mathbf{z}] \right\rangle$ has closed graph in $\mathbf{z}$ in the product topology. Moreover, each one of $\varphi(\mathbf{z})$ and $\hat{X}_n(\mathbf{z})$ for $n = -1, 0, ...$ is nonempty, compact and convex-valued. Therefore, the image of the mapping Φ is nonempty, compact and convex-valued for each $\mathbf{z} \in \mathbf{Z}$. The Kakutani-Fan-Glicksberg Fixed Point Theorem implies that if the function Φ maps a convex, compact and nonempty subset of a locally convex Hausdorff space into itself and has closed graph and is nonempty, compact and convex-valued $\mathbf{z}$, then it possesses a fixed point $\mathbf{z}^* \in \Phi(\mathbf{z}^*)$ (see Aliprantis and Border, 1999, Theorem 16.50 and Corollary 16.51, p. 549-550). This establishes the existence of a fixed point $\mathbf{z}^*$ of Φ.

To complete the proof, we need to show that the fixed point, $\mathbf{z}^*$, corresponds to a steady state equilibrium. First, since $\hat{x}_n\left(\omega^*, \{x_n^*\}_{n=-1}^{\infty}\right) = x_n^*$ for $n \in \{-1\} \cup \mathbb{Z}_+$, we have that given a labor share of ω^*, $\{x_n^*\}_{n=-1}^{\infty}$ constitutes an R&D policy vector that is best response to itself, as required

by steady-state equilibrium (Definition 3). Next, we need to prove that the implied labor share ω^* leads to labor market clearing. This follows from the fact that the fixed point involves $\omega^* < 1$, since in this case (55) will have an interior solution, ensuring labor market clearing. Suppose, to obtain a contradiction, that $\omega^* = 1$. Then, as noted in the text, we must have $\mu_0^* = 1$. From (35), (36) and (37), this implies $x_n^* = 0$ for $n \in \{-1\} \cup \mathbb{Z}_+$. However, we have shown above that this is not possible when $G'^{-1}\left(\left(1-\lambda^{-1}\right)/\left(\rho+\eta\right)\right) > 0$. Consequently, (55) cannot be satisfied at $\omega^* = 1$, implying that $\omega^* < 1$. When $\omega^* < 1$, the labor market clearing condition (38) is satisfied at ω^* as an equality, so ω^* is an equilibrium given $\{x_n^*\}_{n=-1}^{\infty}$, and thus $\mathbf{z}^* = \left(\omega^*, \{x_n^*\}_{n=-1}^{\infty}\right)$ is a steady-state equilibrium as desired.

Finally, if $\eta > 0$, then (37) implies that $\mu_0^* > 0$. Since $x_0^* > 0$, equation (39) implies $g^* > 0$. Alternatively, if $x_{-1}^* > 0$, then $g^* > 0$ follows from (39). This completes the proof of the existence of a steady-state equilibrium with positive growth.

Part 2: Properties of the Sequence of Value Functions.

Let $\{x_n\}_{n=-1}^{\infty}$ be the R&D decisions of the firm and $\{v_n\}_{n=-1}^{\infty}$ be the sequence of values, taking the decisions of other firms and the industry distributions, $\{x_n^*\}_{n=-1}^{\infty}$, $\{\mu_n^*\}_{n=-1}^{\infty}$, ω^* and g, as given. By choosing $x_n = 0$ for all $n \geq -1$, the firm guarantees $v_n \geq 0$ for all $n \geq -1$. Moreover, since flow profit satisfy $\pi_n \leq 1$ for all $n \geq -1$, $v_n \leq 1/\rho$ for all $n \geq -1$, establishing that $\{v_n\}_{n=-1}^{\infty}$ is a bounded sequence, with $v_n \in [0, 1/\rho]$ for all $n \geq -1$.

Proof of $v_1 > v_0$: Suppose, first, $v_1 \leq v_0$, then (34) implies $x_0^* = 0$, and by the symmetry of the problem in equilibrium (30) implies $v_0 = v_1 = 0$. As a result, from (33) we obtain $x_{-1}^* = 0$. Equation (29) implies that when $x_{-1}^* = 0$, $v_1 \geq \left(1-\lambda^{-1}\right)/\left(\rho+\eta\right) > 0$, yielding a contradiction and proving that $v_1 > v_0$. □

Proof of $v_{-1} \leq v_0$: Suppose, to obtain a contradiction, that $v_{-1} > v_0$.

If $v_1 \leq v_0$, (33) yields $x_{-1}^* = 0$. This implies $v_{-1} = \eta v_0/\left(\rho+\eta\right)$, which contradicts $v_{-1} > v_0$ since $\eta/\left(\rho+\eta\right) < 1$. Thus we must have $v_1 > v_0$. The value function of a neck-and-neck firm can be written as:

$$\begin{aligned}
\rho v_0 &= \max_{x_0}\left\{-\omega^* G\left(x_0\right) + x_0\left[v_1 - v_0\right] + x_0^*\left[v_{-1} - v_0\right]\right\}, \qquad (57)\\
&\geq \max_{x_0}\left\{-\omega^* G\left(x_0\right) + x_0\left[v_1 - v_0\right]\right\},\\
&\geq -\omega^* G\left(x_{-1}^*\right) + x_{-1}^*\left[v_1 - v_0\right],\\
&\geq -\omega^* G\left(x_{-1}^*\right) + x_{-1}^*\left[v_0 - v_{-1}\right] + \eta\left[v_0 - v_{-1}\right],\\
&= \rho v_{-1},
\end{aligned}$$

which contradicts the hypothesis that $v_{-1} > v_0$ and establishes the claim. □

Proof of $v_n < v_{n+1}$: Suppose, to obtain a contradiction, that $v_n \geq v_{n+1}$. Now (32) implies $x_n^* = 0$, and (29) becomes

$$\rho v_n = \left(1-\lambda^{-n}\right) + x_{-1}^*\left[v_0 - v_n\right] + \eta\left[v_0 - v_n\right] \qquad (58)$$

Also from (29), the value for state $n+1$ satisfies

$$\rho v_{n+1} \geq \left(1-\lambda^{-n-1}\right) + x_{-1}^*\left[v_0 - v_{n+1}\right] + \eta\left[v_0 - v_{n+1}\right]. \qquad (59)$$

Combining the two previous expressions, we obtain

$$\begin{aligned}
&\left(1-\lambda^{-n}\right) + x_{-1}^*\left[v_0 - v_n\right] + \eta\left[v_0 - v_n\right]\\
\geq\ & 1 - \lambda^{-n-1} + x_{-1}^*\left[v_0 - v_{n+1}\right] + \eta\left[v_0 - v_{n+1}\right].
\end{aligned}$$

Since $\lambda^{-n-1} < \lambda^{-n}$, this implies $v_n < v_{n+1}$, contradicting the hypothesis that $v_n \geq v_{n+1}$, and establishing the desired result, $v_n < v_{n+1}$. Consequently, $\{v_n\}_{n=-1}^{\infty}$ is nondecreasing and $\{v_n\}_{n=0}^{\infty}$ is (strictly) increasing. Since a nondecreasing sequence in a compact set must converge, $\{v_n\}_{n=-1}^{\infty}$ converges to its limit point, v_∞, which must be strictly positive, since $\{v_n\}_{n=0}^{\infty}$ is strictly increasing and has a nonnegative initial value. □

The above results combined complete the proof that values form an increasing sequence. ■

Part 3: **Properties of the Sequence of R&D Decisions.**

Proof of $x^*_{n+1} < x^*_n$: From equation (32),

$$\delta_{n+1} \equiv v_{n+1} - v_n < v_n - v_{n-1} \equiv \delta_n \tag{60}$$

would be sufficient to establish that $x^*_{n+1} < x^*_n$ whenever $x^*_n > 0$. We next show that this is the case.

Let us write:

$$\bar{\rho} v_n = \max_{x_n} \left\{ \left(1 - \lambda^{-n}\right) - \omega^* G\left(x_n\right) + x^*_n \left[v_{n+1} - v_n\right] + x^*_{-1} v_0 + \eta v_0 \right\}, \tag{61}$$

where $\bar{\rho} \equiv \rho + x^*_{-1} + \eta$. Since x^*_{n+1}, x^*_n and x^*_{n-1} are maximizers of the value functions v_{n+1}, v_n and v_{n-1}, (61) implies:

$$\begin{aligned}
\bar{\rho} v_{n+1} &= 1 - \lambda^{-n-1} - \omega^* G\left(x^*_{n+1}\right) + x^*_{n+1}\left[v_{n+2} - v_{n+1}\right] + x^*_{-1} v_0 + \eta v_0, \\
\bar{\rho} v_n &\geq 1 - \lambda^{-n} - \omega^* G\left(x^*_{n+1}\right) + x^*_{n+1}\left[v_{n+1} - v_n\right] + x^*_{-1} v_0 + \eta v_0, \\
\bar{\rho} v_n &\geq 1 - \lambda^{-n} - \omega^* G\left(x^*_{n-1}\right) + x^*_{n-1}\left[v_{n+1} - v_n\right] + x^*_{-1} v_0 + \eta v_0, \\
\bar{\rho} v_{n-1} &= 1 - \lambda^{-n+1} - \omega^* G\left(x^*_{n-1}\right) + x^*_{n-1}\left[v_n - v_{n-1}\right] + x^*_{-1} v_0 + \eta v_0.
\end{aligned} \tag{62}$$

Now taking differences with $\bar{\rho} v_n$ and using the definitions of δ_ns, we obtain

$$\begin{aligned}
\bar{\rho}\delta_{n+1} &\leq \lambda^{-n}\left(1 - \lambda^{-1}\right) + x^*_{n+1}\left(\delta_{n+2} - \delta_{n+1}\right) \\
\bar{\rho}\delta_n &\geq \lambda^{-n+1}\left(1 - \lambda^{-1}\right) + x^*_{n-1}\left(\delta_{n+1} - \delta_n\right).
\end{aligned}$$

Therefore,

$$\left(\bar{\rho} + x^*_{n-1}\right)\left(\delta_{n+1} - \delta_n\right) \leq -k_n + x^*_{n+1}\left(\delta_{n+2} - \delta_{n+1}\right), \tag{63}$$

where

$$k_n \equiv (\lambda - 1)^2 \lambda^{-n-1} > 0.$$

Now to obtain a contradiction, suppose that $\delta_{n+1} - \delta_n \geq 0$. From (63), this implies $\delta_{n+2} - \delta_{n+1} > 0$ since k_n is strictly positive. Repeating this argument successively, we have that if $\delta_{n'+1} - \delta_{n'} \geq 0$, then $\delta_{n+1} - \delta_n > 0$ for all $n \geq n'$. However, we know from Part 2 of the proposition that $\{v_n\}_{n=0}^{\infty}$ is strictly increasing and converges to a constant v_∞. This implies that $\delta_n \downarrow 0$, which contradicts the hypothesis that $\delta_{n+1} - \delta_n \geq 0$ for all $n \geq n' \geq 0$, and establishes that $x^*_{n+1} \leq x^*_n$. To see that the inequality is strict when $x^*_n > 0$, it suffices to note that we have already established (60), i.e., $\delta_{n+1} - \delta_n < 0$, thus if equation (32) has a positive solution, then we necessarily have $x^*_{n+1} < x^*_n$.

We next prove that $x^*_0 \geq x^*_{-1}$ and then show that under the additional condition $G'^{-1}\left(\left(1 - \lambda^{-1}\right)/\left(\rho + \eta\right)\right) > 0$, this inequality is strict.

Proof of $x^*_0 \geq x^*_{-1}$: Equation (30) can be written as

$$\rho v_0 = -\omega^* G\left(x^*_0\right) + x^*_0\left[v_{-1} + v_1 - 2v_0\right]. \tag{64}$$

We have $v_0 \geq 0$ from Part 2 of the proposition. Suppose $v_0 > 0$. Then (64) implies $x^*_0 > 0$ and

$$\begin{aligned}
v_{-1} + v_1 - 2v_0 &> 0 \\
v_1 - v_0 &> v_0 - v_{-1}.
\end{aligned} \tag{65}$$

This inequality combined with (34) and (41) yields $x^*_0 > x^*_{-1}$. Suppose next that $v_0 = 0$. Inequality (65) now holds as a weak inequality and implies that $x^*_0 \geq x^*_{-1}$. Moreover, since $G(\cdot)$ is strictly convex and x^*_0 is given by (34), (64) then implies $x^*_0 = 0$ and thus $x^*_{-1} = 0$.□

We now have the following intermediate lemma.

Lemma 2 *Suppose that* $G'^{-1}\left(\left(1-\lambda^{-1}\right)/\left(\rho+\eta\right)\right)>0$*, then* $x_0^*>0$ *and* $v_0>0$*.*

Proof. Suppose, to obtain a contradiction, that $x_0^*=0$. The first part of the proof then implies that $x_{-1}^*=0$. Then (29) implies

$$\rho v_1 \geq 1-\lambda+\eta\left[v_0-v_1\right].$$

Equation (30) together with $x_0^*=0$ gives $v_0=0$, and hence

$$v_1-v_0\geq\frac{1-\lambda^{-1}}{\rho+\eta}.$$

Combined with this inequality, (34) implies

$$\begin{aligned} x_0^* &\geq \max\left\{G'^{-1}\left(\frac{1-\lambda^{-1}}{\omega^*\left(\rho+\eta\right)}\right),0\right\}, \\ &\geq \max\left\{G'^{-1}\left(\frac{1-\lambda^{-1}}{\rho+\eta}\right),0\right\}, \end{aligned}$$

where the second inequality follows from the fact that $\omega^*\leq 1$. The assumption that $G'^{-1}\left(\left(1-\lambda^{-1}\right)/\left(\rho+\eta\right)\right)>0$ then implies $x_0^*>0$, thus leading to a contradiction and establishing that $x_0^*>0$. Strict convexity of $G\left(\cdot\right)$ together with $x_0^*>0$ then implies $v_0>0$. ■

Proof of $x_0^*>x_{-1}^*$ *when* $G'^{-1}\left(\left(1-\lambda^{-1}\right)/\left(\rho+\eta\right)\right)>0$ *:* Given Lemma 2, $G'^{-1}\left(\left(1-\lambda^{-1}\right)/\left(\rho+\eta\right)\right)>0$ implies that $x_0^*>0$. Then (64) implies

$$v_1-v_0>v_0-v_{-1}$$

and as a result $x_0^*>x_{-1}^*$. □

Proof of $x_0^*>x_1^*$: To prove that $x_0^*>x_1^*$, let us write the value functions v_2, v_1 and v_0 as in (62):

$$\begin{aligned} \bar{\rho}v_2 &= 1-\lambda^{-2}-\omega^*G\left(x_2^*\right)+x_2^*\left[v_3-v_2\right]+x_{-1}^*v_0+\eta v_0, \\ \bar{\rho}v_1 &\geq 1-\lambda^{-1}-\omega^*G\left(x_2^*\right)+x_2^*\left[v_2-v_1\right]+x_{-1}^*v_0+\eta v_0, \\ \bar{\rho}v_1 &\geq 1-\lambda^{-1}-\omega^*G\left(x_0^*\right)+x_0^*\left[v_2-v_1\right]+x_{-1}^*v_0+\eta v_0, \\ \bar{\rho}v_0 &= -\omega^*G\left(x_0\right)+x_0^*\left[v_1-v_0\right]+\eta v_0+x_{-1}^*v_0+x_0^*\left[v_{-1}-v_0\right]. \end{aligned}$$

Now taking differences with $\bar{\rho}v_n$ and using the definitions of δ_ns as in (60), we obtain

$$\begin{aligned} \bar{\rho}\delta_2 &\leq \lambda^{-1}\left(1-\lambda^{-1}\right)+x_2^*\left(\delta_3-\delta_2\right), \\ \bar{\rho}\delta_1 &\geq \left(1-\lambda^{-1}\right)+x_0^*\left(\delta_2-\delta_1\right)+x_{-1}^*\left[v_0-v_0\right]-x_0^*\left[v_{-1}-v_0\right], \\ \bar{\rho}\delta_1 &\geq \left(1-\lambda^{-1}\right)+x_0^*\left(\delta_2-\delta_1\right)-x_0^*\left[v_{-1}-v_0\right], \\ \bar{\rho}\delta_1 &\geq \left(1-\lambda^{-1}\right)+x_0^*\left(\delta_2-\delta_1\right)-x_0^*\left[v_{-1}-v_0\right]. \end{aligned} \tag{66}$$

Next recall from Part 2 that $v_{-1}-v_0\leq 0$. Moreover, the first part of the first part of the proof has established that $x_{-1}^*-x_0^*\leq 0$. Therefore $\left[x_{-1}^*-x_0\right]\left[v_{-1}-v_0\right]\geq 0$, and the last inequality then implies

$$\bar{\rho}\delta_1\geq\left(1-\lambda^{-1}\right)+x_0^*\left(\delta_2-\delta_1\right).$$

Now combining this inequality with the first inequality of (66), we obtain

$$\left(\bar{\rho}+x_0^*\right)\left(\delta_2-\delta_1\right)\leq-\left(1-\lambda^{-1}\right)^2+x_2^*\left(\delta_3-\delta_2\right). \tag{67}$$

Part 2 has already established $\delta_2>\delta_3$, so that the right-hand side is strictly negative, therefore, we must have $\delta_2-\delta_1<0$, which implies that $x_0^*>x_1^*$ and completes the proof. □

The above results together complete the proof of Part 3. ■

Part 4: Uniqueness of the Invariant Distribution.

Lemma 3 *Consider a uniform IPR policy* $\boldsymbol{\eta}^{uni}$ *and a corresponding steady-state equilibrium* $\langle \boldsymbol{\mu}^*,\ \mathbf{v},\ \ \mathbf{x}^*,\ \omega^*,\ g^* \rangle$. *Then, there exists* $n^* \in \mathbb{N}$ *such that* $x_n^* = 0$ *for all* $n \geq n^*$.

Proof. The first-order condition of the maximization of the value function (29) implies:

$$G'(x_n) \geq \frac{v_{n+1} - v_n}{\omega^*} \text{ and } x_n \geq 0,$$

with complementary slackness. $G'(0)$ is strictly positive by assumption. If $(v_{n+1} - v_n)/\omega^* < G'(0)$, then $x_n = 0$. The second part of the proposition implies that $\{v_n\}_{n=-1}^{\infty}$ is a convergent and thus a Cauchy sequence, which implies that there exists $\exists n^* \in \mathbb{N}$ such that $v_{n+1} - v_n < \omega^* G'(0)$ for all $n \geq n^*$. ■

An immediate consequence of Lemma 3, combined with (35) is that $\mu_n = 0$ for all $n \geq n^*$ (since there is no innovation in industries with technology gap greater than n^*). Thus the law of motion of an industry can be represented by a finite Markov chain. Moreover, because after an innovation by a follower, all industries jump to the neck-and-neck state, this Markov chain is irreducible (and aperiodic), thus converges to a unique steady-state distribution of industries. More formally, there exists n^* such that $x_{n^*}^* = 0$ and $x_n^* = 0$ for all $n > n^*$. Combined with the fact $G'^{-1}\left(\left(1 - \lambda^{-1}\right)/(\rho + \eta)\right) > 0$ and that either $\eta > 0$ or $x_{-1}^* > 0$, this implies that the states $n > n^*$ are transient and can be ignored. Consequently, $\{\mu_n^*\}_{n=0}^{\infty}$ forms a finite and irreducible Markov chain over the states $n = 0, 1, ..., n^*$. To see this, let $n^* = \min_{n \in \{0,...,n^{**}\}} \{n \in \mathbb{N} : v_{n+1} - v_n \leq \omega^* G'(0)\}$. Such an n^* exists, since the set $\{0, ..., n^{**}\}$ is finite and nonempty because of the assumption that $G'^{-1}\left(\left(1 - \lambda^{-1}\right)/(\rho + \eta)\right) > 0$. Then by construction $x_n^* > 0$ for all $n < n^*$ and $x_{n^*}^* = 0$ as desired. Now denoting the probability of being in state $\tilde{n}$ starting in state n after τ periods by $P^\tau(n, \tilde{n})$, we have that $\lim_{\tau \to \infty} P^\tau(n, \tilde{n}) = 0$ for all $\tilde{n} > n^*$ and for all n. Thus we can focus on the finite Markov chain over the states $n = 0, 1, ..., n^*$, and $\{\mu_n^*\}_{n=0}^{n^*}$ is the limiting (invariant) distribution of this Markov chain. Given $\{x_n^*\}_{n=-1}^{n^*}$, $\{\mu_n^*\}_{n=0}^{n^*}$ is uniquely defined. Moreover, the underlying Markov chain is irreducible (since $x_n^* > 0$ for $n = 0, 1, ..., n^* - 1$, so that all states communicate with $n = 0$ or $n = 1$). Therefore, by Theorem 11.2 in Stokey, Lucas and Prescott (1989, p. 62) there exists a unique stationary distribution $\{\mu_n^*\}_{n=0}^{\infty}$. ■

Proof of Proposition 5

We prove this proposition using two crucial lemmas.

Lemma 4 *Consider the state-dependent IPR policy* $\boldsymbol{\eta}$, *and suppose that* $\langle \boldsymbol{\mu}^*,\ \mathbf{v},\ \ \mathbf{x}^*,\ \omega^*,\ g^* \rangle$ *is a steady-state equilibrium. Then there exists a state* $n^* \in \mathbb{N}$ *such that* $\mu_n^* = 0$ *for all* $n \geq n^*$.

Proof. There are two cases to consider. First, suppose that $\{v_n\}_{n \in \mathbb{Z}_+}$ is strictly increasing. Then it follows from the proof of Lemma 3 that there exists a state $n^* \in \mathbb{N}$ such that $x_n^* = 0$ for all $n \geq n^*$, and as in the proof of Part 4 of Proposition 4, states $n \geq n^*$ are transient (i.e., $\lim_{\tau \to \infty} P^\tau(n, \tilde{n}) = 0$ for all $\tilde{n} > n^*$ and for all n), so $\mu_n^* = 0$ for all $n \geq n^*$.

Second, in contrast to the first case, suppose that there exists some $n^{**} \in \mathbb{Z}_+$ such that $v_{n^{**}} \geq v_{n^{**}+1}$. Then, let $n^* = \min_{n \in \{0,...,n^{**}\}} \{n \in \mathbb{N} : v_{n+1} - v_n \leq \omega^* G'(0)\}$, which is again well defined. Then, optimal R&D decision (32) immediately implies that $x_n^* > 0$ for all states with $n < n^*$, and since $x_{n^*}^* = 0$, all states $n > n^*$ are transient and $\lim_{\tau \to \infty} P^\tau(n, \tilde{n}) = 0$ for all $\tilde{n} > n^*$ and for all n, completing the proof. ■

Lemma 5 *Consider the state-dependent IPR policy* $\boldsymbol{\eta}$ *and suppose that the labor share and the R&D policies of all other firms are given by* $\mathbf{z} = \langle \tilde{\omega}, \tilde{\mathbf{x}} \rangle$. *Then the dynamic optimization problem of an individual firm leads to a unique value function* $\mathbf{v}[\mathbf{z}] : \mathbb{Z} \to \mathbb{R}_+$ *and optimal R&D policy* $\hat{\mathbf{X}}[\mathbf{z}] : \mathbb{Z} \rightrightarrows [0, \bar{x}]$ *are compact and convex-valued for each* $\mathbf{z} \in \mathbf{Z}$ *and upper hemi-continuous in* $\mathbf{z}$ *(where* $\mathbf{v}[\mathbf{z}] \equiv \{v_n[\mathbf{z}]\}_{n=-1}^{\infty}$, $\hat{\mathbf{X}}[\mathbf{z}] \equiv \left\{\hat{X}_n[\mathbf{z}]\right\}_{n=-1}^{\infty}$ *).*

Proof. The proof follows closely that of Lemma 1. In particular, again using uniformization, the maximization problem of an individual firm can be written as a contraction mapping similar to (54) there. The finiteness of the transition probabilities follows, since $\psi_n\left(\boldsymbol{\xi} \mid \tilde{\boldsymbol{\xi}}\right) \leq \psi \equiv 2\bar{x} + \max_n \{\eta_n\} < \infty$ (this is a consequence of the fact that $\bar{x}$ defined in (16) is finite and $\max_n \{\eta_n\}$ is finite, since each $\eta_n \in \mathbb{R}_+$ and by assumption, there exists $\bar{n} < \infty$ such that $\eta_n = \eta_{\bar{n}}$). This contraction mapping uniquely determines the value function $\mathbf{v}[\mathbf{z}] : \mathbb{Z} \to \mathbb{R}_+$.

Berge's Maximum Theorem (Aliprantis and Border, 1999, Theorem 16.31, p. 539) again implies that each of $\hat{X}_n(\mathbf{z})$ for $n \in \mathbb{Z}$ is upper hemi-continuous in $\mathbf{z} = \langle \tilde{\omega}, \tilde{\mathbf{x}} \rangle$, and moreover, since v_n for $n \in \mathbb{Z}$ is concave in x_n, the maximizer of $\mathbf{v}[\mathbf{z}]$, $\hat{\mathbf{X}} \equiv \left\{\hat{X}_n\right\}_{n=-\infty}^{\infty}$, are nonempty, compact and convex-valued. ■

Now using the previous two lemmas, we can establish the existence of a steady-state equilibrium. This part of the proof follows that of Proposition 4 closely. Fix $\mathbf{z} = \left\langle \tilde{\omega}, \{\tilde{x}_n\}_{n=-\infty}^{\infty} \right\rangle$, and define $\mathbf{Z} \equiv [0,1] \times \prod_{n=-\infty}^{\infty} [0, \bar{x}]$. Again by Tychonoff's Theorem, $\mathbf{Z}$ is compact in the product topology. Then consider the mapping Φ: $\mathbf{Z} \rightrightarrows \mathbf{Z}$ constructed as $\Phi \equiv \left(\varphi, \hat{\mathbf{X}}\right)$, where φ is given by (55) and $\hat{\mathbf{X}}$ is defined in Lemma 5. Clearly Φ maps $\mathbf{Z}$ into itself. Moreover, as in the proof of Proposition 4, $\mathbf{Z}$ is nonempty, convex, and a subset of a locally convex Hausdorff space. The proof of Lemma 5 then implies that Φ has closed graph in the product topology and is nonempty, compact and convex-valued in $\mathbf{z}$. Consequently, the Kakutani-Fan-Glicksberg Fixed Point Theorem again applies and implies that Φ has a fixed point $\mathbf{z}^* \in \Phi(\mathbf{z}^*)$. The argument that the fixed point $\mathbf{z}^*$ corresponds to a steady-state equilibrium is identical to that in Proposition 4, and follows from the fact that within argument identical to that of Lemma 2, $G'^{-1}\left(\left(1 - \lambda^{-1}\right) / \left(\rho + \eta_1\right)\right) > 0$ implies $x_0^* > 0$. The result that $\omega^* < 1$ then follows immediately. Finally, as in the proof of Proposition 4, either $\eta_1 > 0$ or $x_{-1}^* > 0$ is sufficient for $g^* > 0$. ■

References

Acemoglu, Daron and Ufuk Akcigit (2006) "State-Dependent Intellectual Property Right Policy" NBER working paper, No. 12775.

Aghion, Philippe, Nick Bloom, Richard Blundell, Rachel Griffith and Peter Howitt (2005) "Competition and Innovation: an Inverted-U Relationship," *Quarterly Journal of Economics*, 120, pp. 701-728.

Aghion, Philippe and Rachel Griffith (2007) *Competition and Growth: Reconciling Theory and Evidence*, MIT Press, Cambridge.

Aghion, Philippe, Christopher Harris, Peter Howitt and John Vickers (2001) "Competition, Imitation and Growth with Step-by-Step Innovation," *Review of Economic Studies,* 68, pp. 467-492 .

Aghion, Philippe, Christopher Harris and John Vickers (1997) "Step-by-step Competition, Innovation and Growth: An Example," *European Economic Review,* 41, pp. 771-782.

Aghion, Philippe and Peter Howitt (1992) "A Model of Growth Through Creative Destruction," *Econometrica*, 110, 323-351.

Aliprantis, Charalambos D. and Kim C. Border (1999) *Infinite Dimensional Analysis: Hitchhiker's Guide*, Springer-Verlag, New York.

Almeida, Rita and Ana Margarida Fernandes (2008). "Openness and Technological Innovations in Developing Countries: Evidence from Firm-Level Surveys," *The Journal of Development Studies*, 44(5), 701-727.

Arrow, Kenneth (1962) "Economic Welfare and the Allocation of Resources for Invention," in R.R. Nelson, ed. *Universities-National Bureau of Economic Research Conference Series; The Rate and Direction of Economic Activities: Economic and Social Factors,* Princeton University Press, New York.

Baye, Michael R., and Heidrun C. Hoppe (2003) "The Strategic Equivalence of Rent-seeking, Innovation, and Patent-race Games," *Games and Economic Behavior*, 44, 217-226.

Bloom, Nick, Mark Schankerman and John Van Reenen (2005) "Identifying Technology Spillovers and Product Rivalry," CEP Discussion Paper, No. 675.

Blume, Lawrence E. (1982) "New Techniques for the Study of Stochastic Equilibrium Processes," *Journal of Mathematical Economics*, 9, 61-70.

Boldrin, Michele and David K. Levine (2008) "Perfectly Competitive Innovation,"

Journal of Monetary Economics, 55, 435-453.

Boldrin, Michele and David K. Levine (2004) "IP and Market Size," Univesity of Minnesota and UCLA mimeo.

Budd, Christopher, Christopher Harris and John Vickers (1993) "A Model of the Evolution of Duopoly: Does the Asymmetry Between Firms Tend to Increase or Decrease?" *Review of Economic Studies*, 60, 540-573.

Bureau of Economic Analysis (2004) "Annual Revision of the National Income and Product Accounts" *Bureau of Economic Analysis News.* BEA 04-31. www.bea.gov/bea/rels.htm

Choi, Jay P. (1991) "Dynamic R&D Competition under Hazard Rate Uncertainty," *Rand Journal of Economics*, 22, 596-610.

Fudenberg, Drew, Richard Gilbert, Joseph Stiglitz and Jean Tirole (1983) "Preemption, leapfrogging and competition in patent races," *European Economic Review* 22, 3-31.

Fullerton, Robert and Preston McAfee (1999) "Auctioning Entry into Tournaments," *Journal of Political Economy*, 107, 573-605.

Gallini, Nancy T. (1992) "Patent Policy and Costly Imitation," *RAND Journal of Economics*, 23, 52-63.

Gallini, Nancy T. and Suzanne Scotchmer (2002) "Intellectual Property; When is it the Best Incentive Mechanism?" *Innovation Policy and the Economy,* 2, 51-78.

Gilbert, Richard. and Carl Shapiro (1990) "Optimal Patent Length and Breadth," *Journal of Industrial Economics*, 21(1), 106-112.

Green, Jerry and Suzanne Scotchmer (1995) "On the Division of Profit in Sequential Innovation," *Journal of Industrial Economics*, 26(1), 20-33.

Grossman, Gene and Elhanan Helpman (1990) *Innovation and Growth in the Global Economy,* Cambridge, MA, MIT Press.

Hall, Bronwyn H., Jerry Hausman and Zvi Griliches (1988) "Patents and R&D: Is There a Lag?" *International Economic Review*, 27, pp. 265-83.

Harris, Christopher and John Vickers (1985) "Perfect Equilibrium in a Model of Race," *Review of Economic Studies,* 52, 193-209.

Harris, Christopher and John Vickers (1987) "Racing with Uncertainty," *Review of Economic Studies,* 54, 1-21.

Hopenhayn, Hugo and Matthew Mitchell (2001) "Innovation, Variety and Patent Breadth," *RAND Journal of Economics*, 32, 152-166.

Hopenhayn, Hugo and Matthew Mitchell (2011) "Optimal Patent Policy with Recurrent Innovators," mimeo.

Hugo Hopenhayn, Llobet, Gerard and Matthew Mitchell (2006) "Rewarding Sequential Innovators: Prizes, Patents and Buyouts," *Journal of Political Economy*, 114, 1041-1068.

Judd, Kenneth (1999) *Numerical Methods in Economics*, MIT Press, Cambridge.

Kelley, John L. (1955) *General Topology*, Springer-Verlag, New York.

Klemperer, Paul (1990) "How Broad Should the Scope of Patent Protection Be?" *Journal of Industrial Economics,* 21, 113-130.

Kremer, Michael (2002). "Pharmaceuticals and the Developing World," *Journal of Economic Perspectives,* 16(4), 67-90.

Kortum, Samuel (1993) "Equilibrium R&D and the Patent-R&D Ratio: US Evidence," *American Economic Review,* 83, pp. 450-457.

Lanjouw, Jean O and Mark Schankerman (2001) "Enforcing Intellectual Property Rights," NBER Working Paper, No. 8656.

Moscarini, Giuseppe and Francesco Squintani (2004) "Competitive Experimentation with Private Information," Cowles Foundation Discussion Paper, No. 1489.

Moser, Petra and Alessandra Voena (2011) "Evidence from the Trading with the Enemy Act," *American Economic Review,* forthcoming.

National Science Foundation (2006) "Industrial R&D Employment in the United States and in U.S. Multinational Corporations" *Division of Science Resource Statistics.* NSF 05-302. Arlington, VA: National Science Foundation. www.bsf.gov/sbe/srs/stats.htm

O'Donoghue, Ted, Suzanne Scotchmer and Jacques François Thisse (1998) "Patent Breadth, Patent Life and the Pace of Technological Progress," *Journal of Economics and Management Strategy*, 7, 1-32.

O'Donoghue, Ted and Joseph Zweimuller (2004) "Patents in a Model of Endogenous Growth," *Journal of Economic Growth*, 9, 81-123.

Pakes, Ariel and Zvi Griliches (1980) "Patents and R&D at the Firm Level: A First Look," *Economic Letters*, 5, pp. 377-381.

Quah, Danny (2004) "24/7 Competitive Innovation," LSE Mimeo.

Reinganum, Jennifer (1981) "Dynamic Games of Innovation," *Journal of Economic Theory,* 25, 21-24.

Reinganum, Jennifer (1985) "Innovation and Industry Evolution," *Quarterly Journal of Economics,* 100, 81-100.

Romer, Paul M. (1990) "Endogenous Technological Change," *Journal of Political Economy*, 98, S71-S102.

Ross, Sheldon (1996) *Stochastic Processes,* John Wiley & Sons, Inc.

Scotchmer, Suzanne (1999) "On the Optimality of the Patent System," *The Rand Journal of Economics,* 30, 181-196.

Scotchmer, Suzanne (2005) *Innovations and Incentives,* MIT Press, Cambridge, Massachusetts.

Segal, Ilya and Michael Whinston (2007) "Anti-Trust in Innovative Industries," *American Economic Review*, 97, 1703-1730.

Tandon, Pankaj (1982) "Optimal Patents with Compulsory Licensing," *Journal of Political Economy* 90(3), 470-486.

Taylor, Curtis (1995) "Digging for Golden Carrots: An Analysis of Research Tournaments," *American Economic Review*, 85, 872-890.

Tirole, Jean (1988) *The Theory of Industrial Organization*, MIT Press, Cambridge, Massachusetts.

Stokey, Nancy (1995) "R&D and Economic Growth" *Review of Economic Studies,* 62, pp. 469-489.

Stokey, Nancy and Robert Lucas with E. Prescott (1989) *Recursive Methods in Economic Dynamics*, Harvard University Press.

Tables

Note: Tables 1-5 give the results of the numerical computations with $\rho = 0.05$ under three different IPR policy regimes. Tables 1-4 consider a different environment (quick catch-up, slow catch-up, licensing and leapfrogging) at a time, whereas Table 5 combines the latter three environments (slow catch-up, licensing and leapfrogging). Table 6 reports the robustness checks of the state-dependent results of Table 5 with alternative step sizes and R&D elasticity parameters. Depending on the applicability and necessity, Tables 1-6 report the steady-state equilibrium values of the difference in the values $v_1 - v_{-3}$ and $v_0 - v_{-3}$; the (annual) catch-up and frontier R&D rates of a follower that is one step behind, $(x_{-1}^{c*}, x_{-1}^{f*})$; the (annual) R&D rate of neck-and-neck competitors, x_0^*; the (annual) R&D rate of one-step leader, x_1^*; fraction of industries in neck-and-neck competition, μ_0^*; fraction of industries at a technology gap of $n = 1, 2$; the value of "labor share," ω^*; the ratio of the labor force working in research; log of initial (annual) consumption, $\ln C(0)$; the annual growth rate, g^*; and the welfare level according to equation (42). It also reports the welfare-maximizing uniform and state-dependent IPR policies. See text for details.

$\lambda = 1.05,\ \gamma = 0.35,$ $B = 0.1$	Full IPR	Optimal Uniform IPR	Optimal State-dependent IPR
η_1	0	0	0.71
η_2	0	0	0.08
η_3	0	0	0
η_4	0	0	0
η_5	0	0	0
x_{-1}^*	0.22	0.22	0.12
x_0^*	0.35	0.35	0.25
x_1^*	0.29	0.29	0.41
μ_0^*	0.24	0.24	0.46
μ_1^*	0.33	0.33	0.19
μ_2^*	0.20	0.20	0.13
ω^*	0.95	0.95	0.96
Researcher ratio	0.032	0.032	0.028
$\ln C\,(0)$	33.78	33.78	34.20
g^*	0.0186	0.0186	0.0204
Welfare	683.0	683.0	692.1

Table 1. Optimal Patent Length in Quick Catch-up Regime

$\lambda = 1.05,\ \gamma = 0.35$ $B = 0.1,\ \zeta_n = \infty,\ \vartheta_n = \infty$	Full IPR	Optimal Uniform IPR	Optimal State-dependent IPR
η_1	0.02	0.11	0.69
η_2	0.02	0.11	0.20
η_3	0.02	0.11	0.14
η_4	0.02	0.11	0.12
η_5	0.02	0.11	0.08
x^*_{-1}	0.75	0.27	0.17
x^*_0	0.99	0.14	0.32
x^*_1	1.10	0.15	0.51
μ^*_0	0.02	0.16	0.30
μ^*_1	0.03	0.19	0.15
μ^*_2	0.03	0.14	0.10
ω^*	0.56	0.90	0.90
Researcher ratio	0.150	0.055	0.059
$\ln C(0)$	31.31	34.31	34.57
g^*	0.025	0.023	0.025
Welfare	636.3	695.3	701.2

TABLE 2. OPTIMAL PATENT LENGTH IN SLOW CATCH-UP REGIME

$\lambda = 1.05,\ \gamma = 0.35$ $B = 0.1,\ \eta_n = 0.02,\ \vartheta_n = \infty$	Full IPR	Optimal Uniform IPR	Optimal State-dependent IPR
ζ_1	∞	1.61	0
ζ_2	∞	1.61	1.54
ζ_3	∞	1.61	2.45
ζ_4	∞	1.61	2.92
ζ_5	∞	1.61	3.32
$v_0 - v_{-3}$	10.1	2.9	3.2
x^*_{-1}	0.75	0.27	0.31
x^*_0	0.99	0.39	0.45
x^*_1	1.10	0.43	0.46
μ^*_0	0.02	0.21	0.18
μ^*_1	0.03	0.25	0.20
μ^*_2	0.03	0.20	0.12
ω^*	0.56	0.94	0.91
Researcher ratio	0.150	0.043	0.071
$\ln C(0)$	31.31	34.13	34.47
g^*	0.025	0.021	0.025
Welfare	636.3	690.9	699.3

TABLE 3. LICENSING IN SLOW CATCH-UP REGIME

$\lambda = 1.05,\ \gamma = 0.35$ $B = 0.1,\ \eta_n = 0.02,\ \zeta_n = \infty$	Full IPR	Optimal Uniform IPR	Optimal State-dependent IPR
ϑ_1	∞	14	0
ϑ_2	∞	14	18.1
ϑ_3	∞	14	31.3
ϑ_4	∞	14	36.6
ϑ_5	∞	14	43.7
$v_1 - v_{-3}$	21.4	2.7	4.1
x_{-1}^{c*}	0.75	0.15	0.14
x_{-1}^{f*}	0	0.23	0.33
x_0^*	0.99	0.30	0.29
x_1^*	1.10	0.30	0.39
μ_0^*	0.02	0.14	0.12
μ_1^*	0.03	0.42	0.35
μ_2^*	0.03	0.22	0.17
ω^*	0.56	0.95	0.94
Researcher ratio	0.150	0.028	0.058
$\ln C(0)$	31.31	35.48	36.17
g^*	0.025	0.026	0.031
Welfare	636.3	720.0	735.9

TABLE 4. LEAPFROGGING IN SLOW CATCH-UP REGIME

$\lambda = 1.05,\ \gamma = 0.35$ $B = 0.1$	Full IPR	Optimal Uniform IPR	Optimal State-dependent IPR
η_1	0.02	0	0
η_2	0.02	0	0
η_3	0.02	0	0
η_4	0.02	0	0
η_5	0.02	0	0
ζ_1	∞	∞	∞
ζ_2	∞	∞	∞
ζ_3	∞	∞	∞
ζ_4	∞	∞	∞
ζ_5	∞	∞	∞
ϑ_1	∞	16.6	0
ϑ_2	∞	16.6	21.7
ϑ_3	∞	16.6	34.6
ϑ_4	∞	16.6	39.4
ϑ_5	∞	16.6	51.3
$v_1 - v_{-3}$	21.4	3.1	4.6
$v_0 - v_{-3}$	10.1	0.8	2.7
x_{-1}^{c*}	0.75	0.16	0.16
x_{-1}^{f*}	0	0.23	0.34
x_0^*	0.99	0.32	0.29
x_1^*	1.10	0.32	0.41
μ_0^*	0.02	0.11	0.36
μ_1^*	0.03	0.41	0.17
μ_2^*	0.03	0.22	0.10
ω^*	0.56	0.94	0.94
Researcher ratio	0.150	0.031	0.065
$\ln C(0)$	31.31	35.62	36.38
g^*	0.025	0.027	0.033
Welfare	636.3	723.0	740.5

Table 5. All Three Policies in Slow Catch-up Regime

$\kappa = 0.1$	Optimal State-dependent IPR $\gamma = 0.1$ $B = 0.04$	Optimal State-dependent IPR $\gamma = 0.6$ $B = 0.2$	Optimal State-dependent IPR $\lambda = 1.01$ $B = 0.35$	Optimal State-dependent IPR $\lambda = 1.20$ $B = 0.024$
η_1	0	0	0	0
η_2	0	0	0	0
η_3	0	0	0	0
η_4	0	0	0	0
η_5	0	0	0	0
ζ_1	∞	∞	∞	∞
ζ_2	∞	∞	∞	∞
ζ_3	∞	∞	∞	∞
ζ_4	∞	∞	∞	∞
ζ_5	∞	∞	∞	∞
ϑ_1	0	0	3.6	0
ϑ_2	16.7	12.6	5.5	33.4
ϑ_3	35.6	20.0	8.9	82.3
ϑ_4	44.8	23.0	12.3	100.7
ϑ_5	54.0	32.9	15.8	128.7
$v_1 - v_{-3}$	4.1	4.0	1.3	17.5
$v_0 - v_{-3}$	2.6	1.6	0.8	5.2
x_{-1}^{c*}	0.23	0.06	0.69	0.02
x_{-1}^{f*}	0.28	0.61	0.88	0.10
x_0^*	0.27	0.42	0.97	0.10
x_1^*	0.29	0.98	1.24	0.10
μ_0^*	0.20	0.01	0.10	0.06
μ_1^*	0.47	0.15	0.26	0.48
μ_2^*	0.19	0.09	0.11	0.22
ω^*	0.95	0.90	0.97	0.11
Researcher ratio	0.011	0.161	0.048	0.087
$\ln C(0)$	36.08	37.54	13.80	118.35
g^*	0.027	0.044	0.026	0.035
Welfare	732.4	768.3	286.3	2381.2

TABLE 6. ALL THREE POLICIES IN SLOW CATCH-UP REGIME - ROBUSTNESS CHECKS

Figures

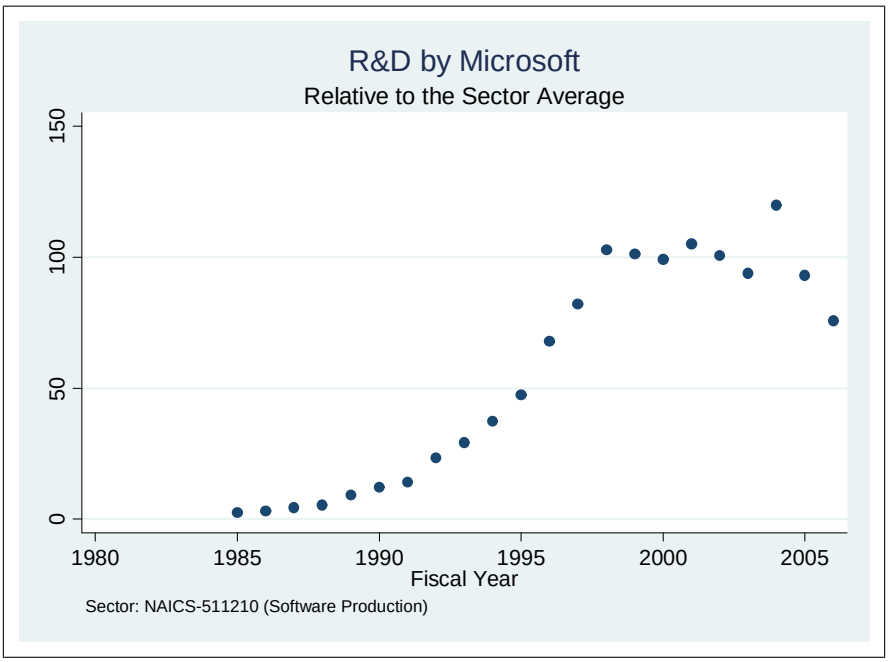

FIGURE 1

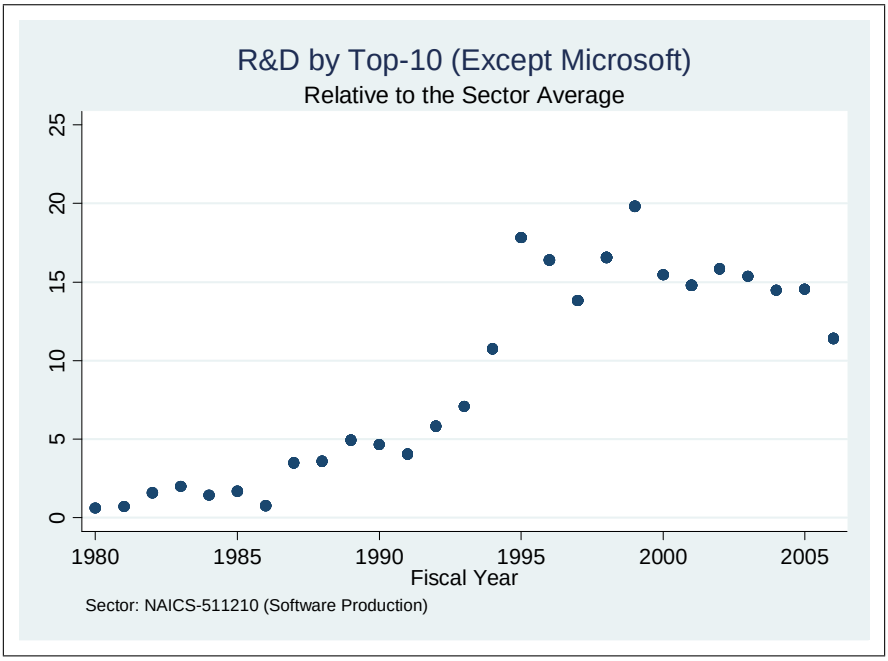

FIGURE 2

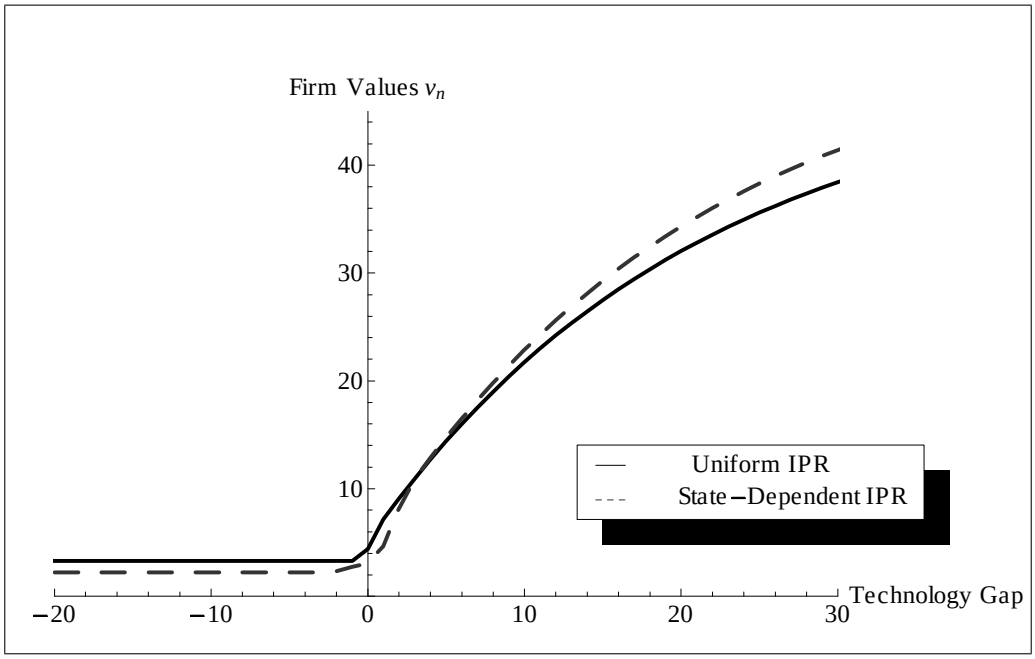

FIGURE 3. VALUE FUNCTIONS

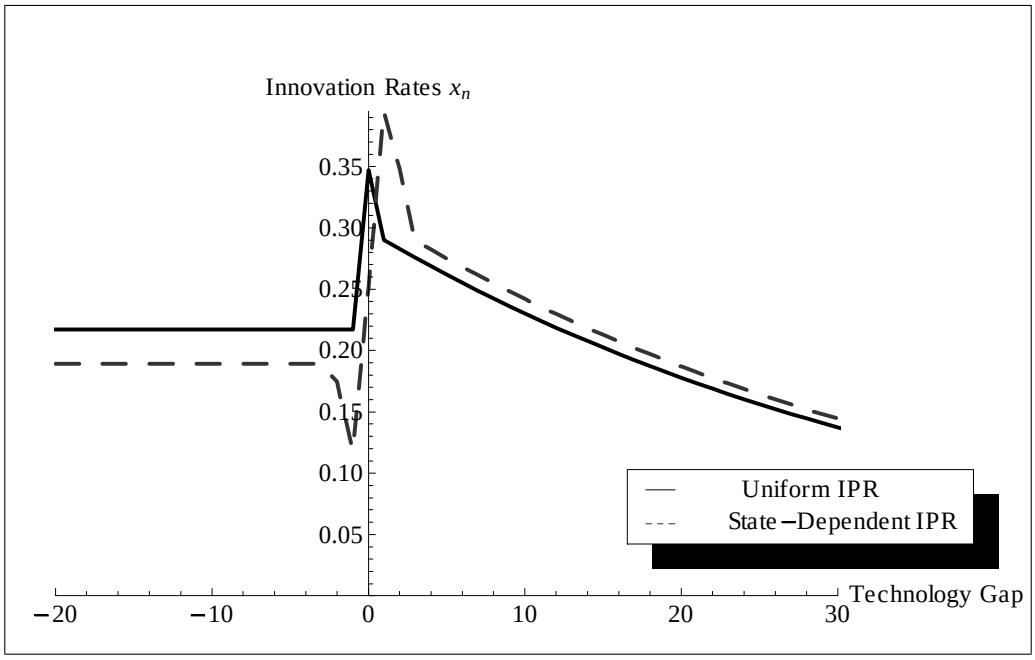

FIGURE 4. R&D EFFORTS

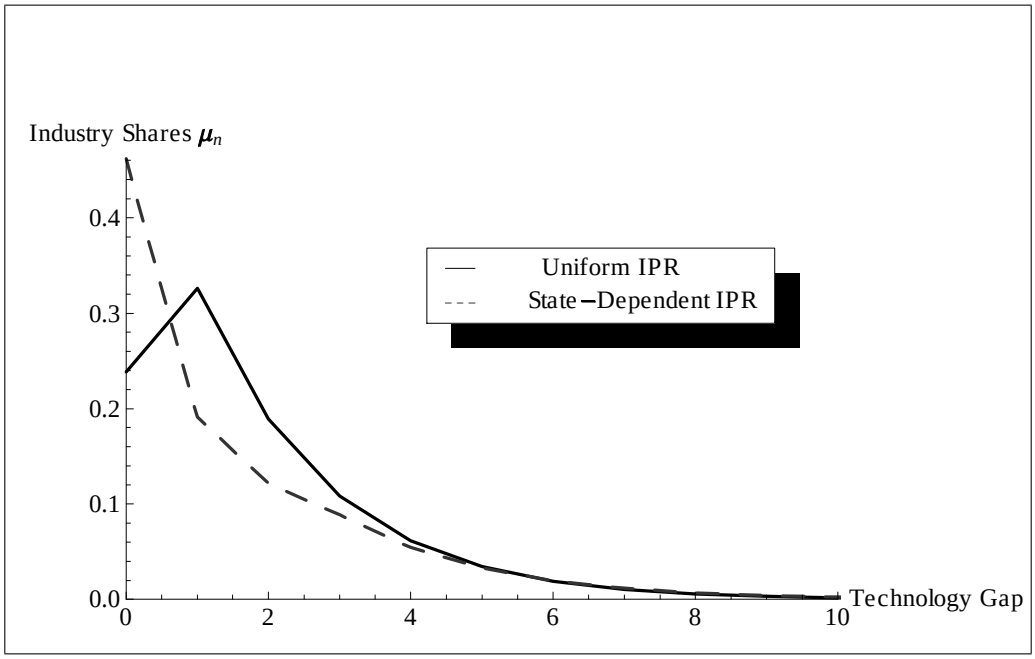

FIGURE 5. INDUSTRY SHARES

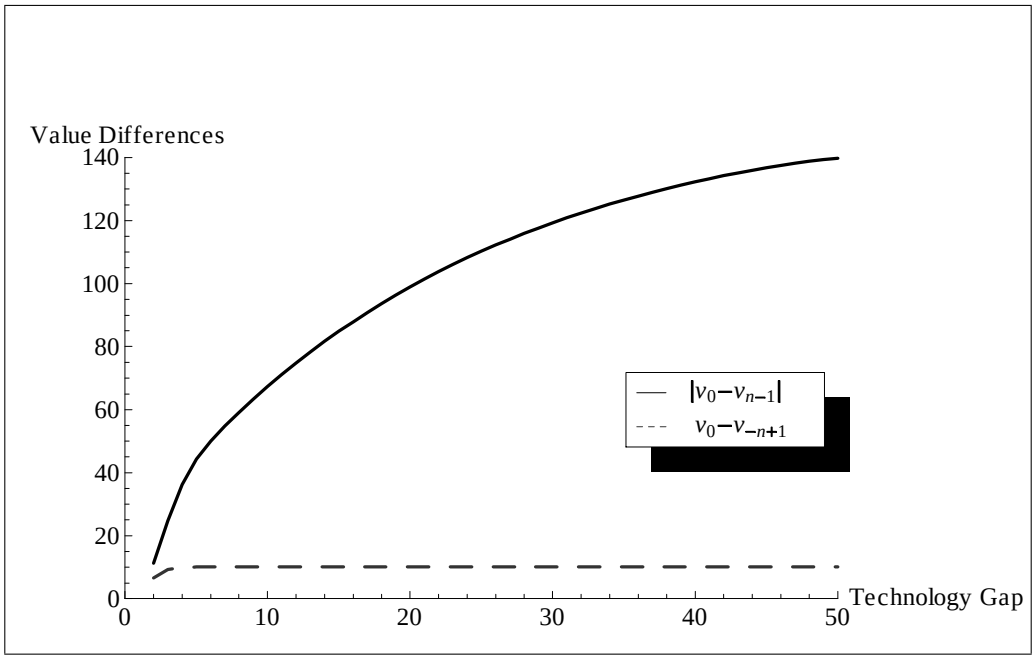

FIGURE 6. VALUE DIFFERENCES UNDER FULL PROTECTION & SLOW CATCH-UP