Symbiotic Competition and Intellectual Property

Rafael R. Guthmann and David Rahman

July 2022

Abstract

What is the optimal policy regarding patents? To answer this question, we develop a model of dynamic competition and incorporate this model into a tractable endogenous growth model. Within this framework, patent laws function as an institutional means to manipulate the degree of competitiveness in industries. Patents reward the creation of novel technologies, which might increase the equilibrium stock of technologies. Still, they slow down growth in process productivity in each technology, which is driven by competition. Depending on the parametrization, the optimal policy might or might not feature patent protection at all. However, when we calibrate the model with parameters consistent with the empirical evidence, the model suggests that the optimal policy is that patents should last approximately fifteen to twenty years.

"We believe in help for the underdog, but we want him to stay under."

Cry, The Beloved Country, Alan Paton, 1948.

1 Introduction

Economists have yet to reach a consensus regarding the costs and benefits of the institution of patents. Since Nordhaus [1969], economists have studied the economic consequences of patents as a trade-off between the increase in innovation promoted by longer or more strictly enforced patents versus the costs in terms of loss of economic efficiency due to the increase in market power that is generated by patents. Some economists have argued that the current state of the evidence shows that there is no reason for governments to grant patents for new inventions (Boldrin and Levine [2013]), while others argue that this puzzle has not been solved yet (Williams [2017] and Bryan and Williams [2021]). This paper represents a contribution to this debate by studying patents in a general equilibrium framework that incorporates endogenous growth with dynamic competition. Our conclusion is that patents provide an important social role, and our calibration suggests that an optimal patent policy that equates the marginal costs of less competition with the marginal benefits of increased incentives to innovate might involve a slightly shorter patent duration than the current centuries-long standard of 20 years.

We find that competition contributes to economic development through the channel of technological spillovers: if we model competition dynamically, firms learn from each other, which increases the growth rate in total factor productivity of the industry when the government grants a monopoly right to certain inventions (a patent) this provides an incentive for the development of new varieties but reduces the growth rate of total factor productivity in the monopolized industry.

We consider an economy where accumulated knowledge takes two forms: there is an endogenously determined set of varieties available for production (it is an expanding variety growth model), and the productivity of the firms in an industry for any given variety also grows over time. We consider the simplest cases: when there is a monopolist and when there are two firms in each sector.

Our contribution is twofold: first, we provide an endogenous growth model where we can track the costs and benefits of patent policy. Second, our theoretical framework has shed light on another source of total factor productivity growth, which we call symbiotic competition between firms.

Symbiotic competition occurs when there are knowledge spillovers. That is, the productivity laggard can learn faster by imitating the productivity leader rather than learning by themselves, and when there are stochastic shocks to the firm's productivity. These two factors combined mean that shocks induce productivity laggards to often leapfrog the productivity leader, and imitation then increases the growth rate of the technology of the former productivity leader. Therefore, the expected technology of the firms in the industry grows faster than under a monopoly. Symbiotic competition formalizes the concept presented by Hayek [1968] who argued that a major component in the social usefulness of competition is it enables the discovery of novel knowledge that would not be possible to discover without it: "is useful to recall that wherever we make use of competition, this can only be justified by our not knowing the essential circumstances that determine the behavior of the competitors. In sporting events, examinations, the awarding of government contracts, or the bestowal of prizes for poems, not to mention science, it would be patently absurd to sponsor a contest if we knew in advance who the winner would be. Therefore, as the title of this lecture suggests, I wish now to consider competition systematically as a procedure for discovering facts which, if the procedure did not exist, would remain unknown or at least would not be used."

In addition, we argue that monopolies might not be able to replicate symbiotic competition by simulating the effects of competition internally. The basic reason for this result is that the growth in total factor productivity increases returns in the future, so if the monopolist has a high enough discount rate, he will not choose to implement symbiotic competition. In comparison, even perfectly myopic firms in competition generate the gains of symbiotic competition. Therefore, competition induces impatient agents to generate productivity growth investments that would only occur in a monopoly if the monopolist were patient.

We study an economy in the balanced growth path where the optimal patent policy maximizes the representative household's utility. When implementing a patent policy, there are significant trade-offs to be considered: increasing patent length increases the degree of market power from a novel technology which provides incentives for its development, but an industry under monopoly suffers from two inefficiencies compared to a duopoly: the loss of the beneficial effects of symbiotic competition and the increased deadweight loss from increased market power. Our calibration suggests that the optimal policy features the enforcement of patents lasting between 15 to 20 years, and this result is robust to substantial changes in the parameters.

1.1 Related literature

Given the substantial literature in the area, it should be helpful to give a brief overview of the works closely related to the present paper. Schumpeter [1980] was among the first economists to focus on innovation in his theory of economic development. Schumpeter also recognized that innovation does not mean only the invention of new technologies but the market introduction of a technical or organizational novelty, not just its invention (Schumpeter [1980]). Arrow [1962] made the systematic argument regarding how the accumulation of knowledge can be incorporated into the theory of economic development in explicit models of learning-by-doing. The model presented here is part of the literature on endogenous models of growth, heavily influenced by models presented in Romer [1986], Lucas [1988], and Stokey [1988]. It departs from these models by the incorporation of an explicit model of competition with technological spillovers.

Maskin and Tirole [1988a,b] provided a foundational contribution to the formal analysis of dynamic competition by introducing the solution concept of Markov Perfect Equilibrium. In this paper, we study dynamic competition with learning-by-doing and extend the concept present in papers such as Besanko et al. [2010, 2014, 2019] to environments with technological spillovers. In our model, technological spillovers play a critical role in allowing markets to approximate perfect competition. Other firms in a sector can learn the technology of the technologically leading firm, allowing the sector to converge to perfect competition over time.

Aghion et al. [2005] study the relationship between competition and innovation. The difference between the model they study and our model is that we incorporate two distinct notions of innovation. We consider innovation in two dimensions: First, creating new products through the application of conscious effort for research and development. Second, the improvements of technology through unexpected learning by operating in the market. The presence of learning spillovers means that multiple firms can learn from each other, and technological improvement inside an industry increases as the intensity of competition increases.

To incorporate the concept that previous experience allows a firm to improve its technology, as in Besanko et al. [2010, 2014, 2019], the technology available to the firms depends on their own experience in the industry. However, in the present paper, the firms can also learn from their competitor's productive activity (for example, reverse engineering). Therefore, learning-by-doing in this model includes a technological spillover effect. A technological spillover implies that equilibrium prices become more competitive the higher the degree of learning-externalizes in the market.

In papers such as Boldrin and Levine [2013], it is argued that patent law, which is a form of intellectual property law that is designed to grant a monopoly in the use of certain technologies for a period of time, is ineffective in promoting innovation and growth. In this paper, we aim to study the optimal design of intellectual property rights under a general theoretical framework. According to our framework, the optimal patent policy agrees with Boldrin and Levine [2013] if the learning spillovers are large enough so that expected productivity growth under competition is very high enough relative to monopoly. In that case, the loss in total factor productivity from grating monopoly rights to some firms becomes large enough so that even at the corner solution of no patents, the costs of the monopoly's market power exceed the benefits of increased innovation in novel technologies produced by patents.

2 The Model

The model below is a standard Dixit-Stiglitz economy in continuous time with endogenous growth following Romer's expanding variety model, that is, driven by a growing population that invents new output varieties. However, we dispense with the typical assumption of monopolistic competition, and instead assume that production in each variety takes place subject to duopolistic competition with learning-by-doing and knowledge spillovers among firms producing each variety. Thus, our description contains a detailed explanation of the operation of competition for the intermediate input varieties.

2.1 Physical Environment

We consider an economy populated by a continuum of identical households with mass $N_t > 0$ at each moment in (continuous) time $t \in [0, \infty)$ that grows at an exogenously given rate $g > 0$. Households consume output and supply labor. They discount future consumption according to the common discount rate $r > 0$.

There is an homogeneous consumption good that is produced by final good producers from different varieties of intermediate goods. There is a continuous set of varieties of intermediate goods which is indexed by the interval $[0, B_t]$ where B_t is the stock of blueprints available at time t.

A household's utility over consumption (also indexed by i) and labor is given by $U(c, \ell) = c - v(\ell)$, where v is a continuously differentiable, increasing strictly convex function (so v' has an inverse). Each household is endowed with one divisible unit of labor that can be allocated between (i) research and development of blueprints for new intermediate product varieties, and (ii) the production of output of existing varieties. Let $e_t \in [0, 1]$ denote the fraction of labor allocated to research and development, and $1 - e_t$ the fraction dedicated to production.

New blueprints are produced by workers employed in research and development according to the linear technology ce_tN_t, for some $c > 0$. Blueprints become obsolete at rate $\delta > 0$. Thus, the stock of blueprints evolves according to

$$\frac{dB_t}{dt} = ce_t N_t - \delta B_t. \tag{1}$$

Blueprints are patented. Patent law attributes the property of the blueprint to the household that developed it. Patents last for a length of time $T \in \mathbb{R}_+ \cup \{+\infty\}$ so an intermediate good producer has to purchase the patent to produce the intermediate good before the expiration date of the patent.

2.2 Intermediate good sector

Intermediate good varieties are produced one input, labor, according to the linear technology $y_i = a_i \ell_i$, where y_i denotes output of the ith variety, ℓ_i its labor input, and a_i its marginal product of labor.

At time t, for each variety $i \in [0, B_t]$, there are potentially two firms which can produce the variety. For the duration of patent policy, only one firm call it $i1$, the firm that purchased the patent to use the blueprint operates as a monopolist, after the patent expires the competing firm $i2$ enters the industry and competes in prices with the former monopolist. Let $a_i 1, a_i 2$ denote the technologies of these firms.

With Bertrand competition, the firm with lowest marginal cost prevails in each sector; the market price equals the minimum of the opponent's marginal cost and own monopoly price. Let the marginal cost of a firm $ij, j \in \{1, 2\}$ be γ_{ij} . As two firms populate each sector so market prices in the Markov perfect equilibrium (See the Appendix Section A for a detailed discussion of why this is the equilibrium price.) are

$$p_i = \begin{cases} \gamma_{i1}/\alpha & \text{if } \alpha a_{i2} < a_{i2} \leq \alpha a_{i1} < a_{i1} \\ \gamma_{i2} & \text{if } \alpha a_{i2} \leq \alpha a_{i1} \leq a_{i2} \leq a_{i1} \\ \gamma_{i1} & \text{if } \alpha a_{i1} \leq \alpha a_{i2} \leq a_{i1} \leq a_{i2} \\ \gamma_{i2}/\alpha & \text{if } \alpha a_{i1} < a_{i1} \leq \alpha a_{i2} < a_{i2}. \end{cases}$$

Suppose that $a_{i2} = \varphi a_{i1}$ for every variety i, with $\varphi \in [\alpha, 1]$. In this case, firm $i1$ is more productive, yet prices are set to equal $i2$'s higher marginal cost. But since these higher costs are a constant fraction of the lower costs across sectors, allocative efficiency prevails.

We assume that each firm's productivity is lognormal as follows:

$$z_{ij} = \ln a_{ij} \sim N(\mu_j, \sigma_j^2).$$

We make this assumption because, later, we will distinguish type $i1$ firms from type $i2$ firms.

2.2.1 Symbiotic Competition in the industry for an intermediate good variety

Consider an industry with two firms, 1 and 2, competing in continuous time. Firm $i \in \{1, 2\}$ has linear production technology that converts one unit of labor into A_{it} units of output at time $t \in [0, \infty)$; let $Z_{it} = \ln(A_{it})$ denote the firm's log productivity, with initial condition $Z_{10} \geq Z_{20}$. Each firm i's productivity is a stochastic process that satisfies

$$dZ_{it} = \begin{cases} (\mu + \theta)dt + \sigma dW_{it} & \text{if } Z_{it} < Z_{jt} \text{ and} \\ \mu dt + \sigma dW_{it} & \text{if } Z_{it} \geq Z_{jt}, \end{cases} \tag{2}$$

where j is the other firm in this industry and W_1, W_2 are independent Wiener processes.[1]

This model of productivity offers a simple way to understand the effect of spillovers on growth. We think of a monopolist's productivity growth following the law of motion

$$dZ_t = \mu dt + \sigma dW_t.$$

Crucially, the duopoly model of (1) above has the drift in a firm's productivity depend on whether it leads or lags the other firm. A leading firm's productivity grows exactly like a monopolist's would, with productivity growth due to learning by doing or research in process innovations captured by the assumption that $\mu > 0$. At the same time, the productivity of a lagging firm grows at an even faster rate, due to technological or other knowledge spillovers

[1] The two-dimensional stochastic process of productivity above is characterized by having rank dependent parameters, in this case each dimension's drift. For references on rank dependent stochastic processes, see bibliography.

from the leading firm. This is captured by assuming that $\theta > 0$. Eventually, the laggard catches up and becomes the new leader, and any innovation by this leader is in turn absorbed by the new laggard. As we will see, this yields overall growth beyond the leader's rate.

The parameter θ captures how quickly firms catch up to their leaders when they lag. The limit of $\theta \to 0$ can be interpreted as there being no spillovers, where firms grow through innovations just as they would in case of monopoly. On the other hand, as $\theta \to \infty$, the laggard firm's productivity converges towards the leader's arbitrarily quickly, until, in the limit, each infinitesimal improvement is immediately absorbed by both firms.

This generalizes the learning models of Boldrin and Levine [2013] and Lucas and Moll [2014] by acknowledging that technology absorption takes time. In their models, when two individuals meet, whoever has lower productivity immediately learns, or absorbs, the technology of the more productive individual and is henceforth able to produce at the higher rate. In our model, however, an individual's productivity can improve by at most θ per unit time.

In addition, unlike Lucas and Moll [2014], our model allows for the productivity of both leader and laggard to be subject to variation during their interaction. This simple extension generates rich economic dynamics. We interpret $\mu > 0$ as a deterministic rate of learning by doing and $\sigma > 0$ as a parameter describing how a firm's productivity fluctuates randomly via process innovation. However, $\sigma > 0$ together with $\theta > 0$ opens the way for firms to outpace—and thus learn from—each other. We refer to this virtuous cycle as *symbiotic competition*.

To begin to illustrate the effects of this virtuous cycle, Figure 1 below depicts sample paths of (log) productivity for each firm in the dynamic duopoly above together with the path of a monopolist. The blue path shows productivity growing at rate $\mu = 0.1$ with $\sigma = 0.2$. The red and yellow paths reflect the system in (1) with $\theta = 0.25$. To facilitate comparison, the innovations in Z_1 are made identical to those in Z.

Three immediate observations may be drawn from Figure 1. First, there is significant catch-up by the laggard firm to the leading firm. Secondly, productivity grows faster in duopoly than in monopoly. Thirdly, duopolists' productivity rates are usually close together, although they do sometimes veer away. Thus, even if firms usually compete neck-and-neck, they enjoy

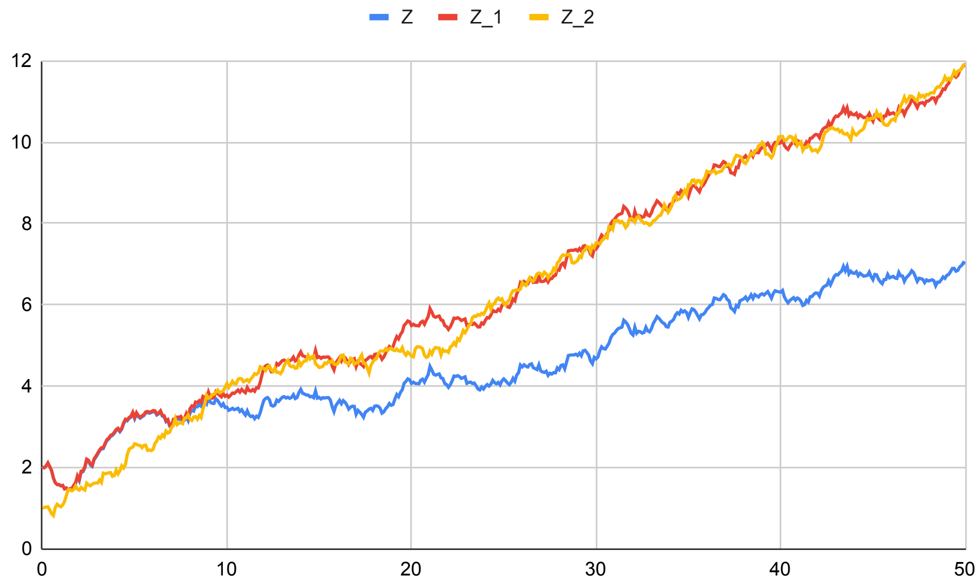

Figure 1: Sample path of productivity—monopoly versus duopoly.

some periods of relative advantage on their competitors.

These observations are formalized below by studying the transition density of each firm's productivity. The productivity process (Z_1, Z_2) above is naturally decomposed into sum and difference processes that submit to analysis more easily. Let us study them in turn.

Proposition 1. Average productivity $X = \frac{1}{2}(Z_1 + Z_2)$ *obeys the following law of motion:*

$$dX_t = (\mu + \tfrac{1}{2}\theta)dt + \sigma dW_{xt}$$

and $W_{xt} = \frac{1}{2}(W_{1t} + W_{2t})$ *is a Wiener processes. Thus, average productivity in a duopoly grows at rate* $\mu + \frac{1}{2}\theta$*, in contrast to the growth rate of monopoly,* μ*.*

Proposition 1 states that average productivity follows a relatively simple arithmetic Brownian motion. Therefore, its transition density into x at time t is given by

$$f_t(x) = \varphi_t\left(x - x_0 - (\mu + \tfrac{1}{2}\theta)t\right), \qquad \text{where} \qquad \varphi_t(z) = \tfrac{1}{\sqrt{2\pi\sigma^2 t}} \exp\left\{-\tfrac{1}{2\sigma^2 t} z^2\right\}$$

is the PDF of a normal random variable with mean 0 and variance $\sigma^2 t$ and $x_0 = X_0$ is the initial value of X. Therefore, average productivity grows faster with knowledge spillovers than without them by $\frac{1}{2}\theta$. In this sense, knowledge spillovers increase productivity growth.

Proposition 2. *Firms'* productivity gap $Y = \frac{1}{2}(Z_1 - Z_2)$ *obeys the law of motion*

$$dY_t = -\tfrac{1}{2}\theta sgn(Y_t)dt + \sigma dW_{yt}$$

and $W_{yt} = \frac{1}{2}(W_{1t} - W_{2t})$ *is a Wiener process independent of* W_{xt}. *Thus, the productivity gap between firms in a duopoly with knowledge spillovers drifts towards zero at rate* $\frac{1}{2}\theta$.

The productivity gap between firms describes the extent of competition in an industry. By Proposition 6, this gap is a diffusion whose drift depends on its sign. If the difference in productivity is positive, it tends to diminish, whereas if it's negative it will tend to increase. This is the "catch-up" effect due to spillovers, which leads firms to engage in "neck-and-neck" competition more frequently than they would without spillovers (i.e., if θ were equal to 0). It turns out that the productivity gap Y has a transition density that can also be written in closed form, as the next result shows.

Lemma 1. *The productivity gap* Y *has transition density*

$$g_t(y) = \hat{g}_t(y) + \Psi_t(|y|), \qquad \text{where} \qquad \hat{g}_t(y) = \begin{cases} \varphi_t(y - y_0 + \frac{1}{2}\theta t) & \text{if } y > 0, \\ e^{-\frac{\theta}{\sigma^2}|y_0|}\varphi_t(y - y_0 - \frac{1}{2}\theta t) & \text{if } y \leq 0, \end{cases}$$

φ_t *is the normal PDF above,* $y_0 = Y_0 \geq 0$ *is the initial value of* Y,

$$\Psi_t(|y|) = \frac{\theta}{2\sigma^2}e^{-\frac{\theta}{\sigma^2}|y|}\Phi_t(\tfrac{1}{2}\theta t - |y| - |y_0|), \qquad \text{and} \qquad \Phi_t(z) = \int_{-\infty}^{z} \varphi_t(x)dx$$

is the normal CDF of φ_t. *For* $y_0 < 0$, *the density* $g_t(y)$ *is anti-symmetric about the origin.*

Lemma 1 characterizes transition dynamics of the productivity gap. The gap's probability density has two components: a term corresponding to the direct tendency from y_0 to y of an arithmetic Brownian motion with drift $-\frac{1}{2}\theta$ sgn(Y), and a correction term for time spent hovering around zero due to catch-up (also known as Y's *local time* around 0) before escaping to y. (See the Appendix for a detailed discussion of these technicalities.)

To gain intuition for the dynamics of Y, it is useful to consider certain limits. As $y_0 \to 0$, the density g_t becomes symmetric, so it's just as likely that either firm is the leader or the laggard at every point in time. As $t \to \infty$, the industry's productivity gap has a stochastic steady state with a "double exponential" probability density g_∞, as the next result shows. This density describes the relative frequency of neck-and-neck competition in the long run, as well as the time firms spend ahead of their competitors, hence earning significant profit.

Proposition 3. *In the long run, firms' productivity gap has probability density*

$$g_\infty(y) = \lim_{t\to\infty} g_t(y) = \tfrac{\theta}{2\sigma^2} e^{-\frac{\theta}{\sigma^2}|y|}.$$

Thus, the relative time during which firm 1's productivity lead exceeds $y \geq 0$ equals $\frac{1}{2}e^{-\frac{\theta}{\sigma^2}y}$. This duration converges to 0 at rate $\frac{\theta}{\sigma^2}$ as $y \to \infty$.

Together, average productivity X and the gap Y completely characterize firms' productivity, since $Z_1 = X + Y$ and $Z_2 = X - Y$. Our next result shows that X and Y are independent, therefore the transition density of (Z_1, Z_2) equals the product of those of X and Y above.

Proposition 4. *The probability density function of (Z_1, Z_2) at time t is given by*

$$h_t(z_1, z_2) = f_t(x) g_t(y)$$

where $x = \frac{1}{2}(z_1 + z_2)$ and $y = \frac{1}{2}(z_1 - z_2)$.

Proposition 4 above delivers a closed form for the probability density function of firm productivity when laggards tend to catch up to leaders. Figure 2 below illustrates what this looks like for fixed $y_0 > 0$ and $t > 0$. Firms' average productivity traverses the fold in the figure, whereas firms' productivity gap proceeds perpendicularly. Interestingly, even though the productivity gap is stationary, average productivity is not, so in the long run there can be big differences in productivity across—but not within— industries.

2.3 Final goods sector

There is a continuum of measure one of producers of the final good which is consumed by the population. The representative producer of the final good purchases the intermediate inputs and combines them to produce the final good according to the following production function

$$F(\boldsymbol{x}_t) = \left[\int_0^{B_t} (x_{it})^\alpha di\right]^{1/\alpha}.$$

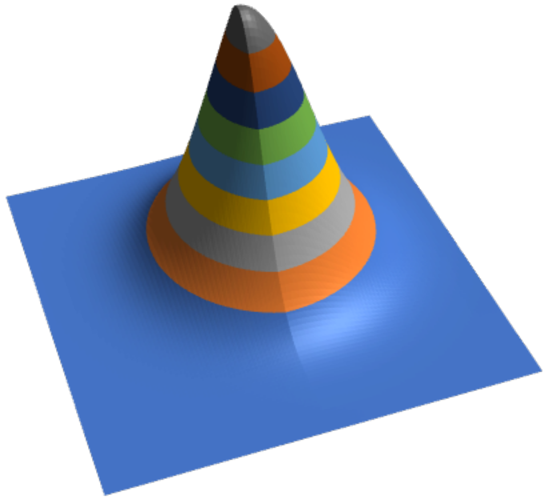

Figure 2: Probability density of (Z_j, Z_k)

The producers maximize profits by choosing inputs. Thus, a final good producer seeks to minimize production cost given a desired output level y_t, therefore

$$\min_{\boldsymbol{x}_t} \int_0^{B_t} p_{it} x_{it} di$$
$$\text{s.t. } F(\boldsymbol{x}_t) \geq y_t.$$

where p_{it} is the minimum price posted by the producers of intermediate input i (which is the same for both producers in the Markov perfect equilibrium) and $B_t > 0$ is the measure of intermediate input varieties in production. The price for the final consumption good is normalized to 1, as it functions as the numeraire. Cost minimization implies that the demand for each intermediate input x_{it} as a function of its price p_{it} and total quantity produced y_t, it satisfies

$$x_{it}(p_{it}) = y_t / p_{it}^{1/(1-\alpha)}. \tag{3}$$

Let L_t be labor demand, labor demand is equal to

$$L_t = \int_0^{B_t} [x_{it}(p_{it})/A_{1t}^i] di, \tag{4}$$

in equilibrium labor demand is equal to labor supply, for every period, thus, in equilibrium $L_t = M_t(1 - e_t)$ and, therefore,

$$M_t(1 - e_t) = \int_0^{B_t} [x_{it}(p_{it})/A^i_{1t}]di. \tag{5}$$

Given wages, the relative level of output of the intermediate good are determined by their relative prices, which are given by

$$x_{it}(p_{it}) = x_{0t}(p_{0t})(p_{0t}/p_{it})^{1/(1-\alpha)}, \tag{6}$$

substituting the right hand side of 6 in 5 and solving for x_{0t} we have

$$x_{0t}(p_{0t}) = M_t(1 - e_t) \int_0^{B_t} A^i_{1t}(p_{it}/p_{0t})^{1/(1-\alpha)}di \tag{7}$$

In the Markov refined equilibrium the quantity purchased by the final good producer satisfies 3, substituting we have that $x_{0t}(p_{0t})$ can be written as

$$x_{0t}(p_{0t}) = M_t(1 - e_t) \int_0^{B_t} A^i_{1t}(\min\{1/A^i_{2t}, 1/\alpha A^i_{1t}\}/\min\{1/A^0_{2t}, 1/\alpha A^0_{1t}\})^{1/(1-\alpha)}di, \tag{8}$$

which determines the output and therefore the prices and wages, allowing us to solve the model. With 6 we have that output is determined by:

$$y_t = \left[\int_0^{B_t} (x_{0t}(p_{0t})(p_{0t}/p_{it})^{1/(1-\alpha)})^{\alpha}di\right]^{1/\alpha},$$

and wages are equal to the marginal product of labor divided by the markup level in each sector.

2.4 Stationary Equilibrium

Our definition of equilibrium for this economy incorporates the notion that the distribution of industry ages is stationary. To incorporate that notion the distribution of blueprint ages is stationary if, at each point in time t, the distribution of blueprints by age $F(a, t)$ is stationary so $F(a, t) = F(a, t')$ for any two dates of blueprint ages $t, t' \geq 0$.

Definition 1. Definition 6. An equilibrium is an allocation and prices such that for each point in time t: (1) Each industry for variety $j \in [0, B_t]$ is in Markov Refined Equilibrium. (2) The choice of effort in development of new blueprints compared to supplying labor to the market is optimal for the workers, and wages w_t equal labor demand with labor supply. (3) The distribution of blueprint ages is stationary.

One important property of this model is that all investment consists of blueprint production effort. The capacity of the economy to produce investments is determined by two factors: demand for innovation, which depends on the degree of market power in the industries that produce intermediate goods and the size of the population.

As patents last for a period of length T, there is only one firm active in the industry for good j from the time t when it has been patented to time $t + T$. After time $t + T$, the patent is expired and after the patent expires we assume there are two firms supplying the market with the same variety and competing in prices.

2.5 Balanced Growth Path

Definition 2. A balanced growth path equilibrium (BGPE) is an economy in stationary equilibrium where output grows at a constant rate.

As labor supply is constrained in reality (each individual is constrained to supplying less than 168 hours in a week), in the balanced growth path labor supply should not explode to infinity as productivity increases, so we consider a class of utility functions that yields a constant labor supply over time, this class of utility functions is given by

$$U_t(c, \ell) = c - K * B_t^{\frac{1-\alpha}{\alpha}} * v(\ell)$$

where $K > 0$ is a constant and v is an increasing continuously differentiable convex function. In this case, the disutility of work depends on the stock of blueprints in the same way as the output of the final good (as a consumer might have more options for entertainment as technology develops, which means technology increases the opportunity cost of work in proportion to the total factor productivity) and so technological progress does not change the supply of labor in equilibrium.

Therefore, in the balanced growth path, as the population grows at a constant rate, the labor supply and the stock of blueprints also grow at a constant rate, the distribution of firm ages is constant and the growth rate of output per capita is constant as well. As blueprint depreciate at the rate δ, the balanced growth path requires a constant growth rate g_b in the output of blueprints, so the stock of blueprints also grows at the rate g_b. Population grows at a constant rate g_p, and production of blueprints is linear in the population, as feasible effort levels are in $[0, 1]$, to maintain a balanced growth path the growth rate in the effort level must be zero, therefore the balanced growth path is characterized by a constant effort level e^*, which implies that blueprint stock grows at the same rate as the population. Therefore, the growth rate in balanced growth path is determined by the growth rate of population, g, and the distribution of growth rates in productivity of the industries for each specific variety. Therefore the stock of blueprints is of size $B(t) = b(t)/(\delta + g)$ and the distribution of blueprints by age at time t, is given by $F(a, t') = 1 - \exp(-(g + \delta)(t - t'))$.

From the previous section we know the expected growth rate of total factor productivity of an industry under neck-and-neck competition is $\mu + (1/2)\theta$ while the growth rate of an industry where the leader is at a significant distance from the other firm (due to patents, for example) is μ. The aggregate growth rate in the balanced growth path is therefore an increasing function of r_p, and it is also an increasing function of the proportion of industries under neck-and-neck competition relative to the industries where the productivity leader is at considerable distance of the other firm.

Proposition 5. *There exists a balanced growth path equilibrium.*

Proof. In a balanced growth path stationary equilibrium there is a constant effort level e^*, with a stationary distribution of ages for product varieties, as the stock of blueprints grows at a constant rate, as well as the population.

If the effort level e^* is optimal in some period t, then consider a period $t' \neq t$, the distribution of varieties by age is the same, the population, and the stock of blueprints are in fixed proportion. Thu, the labor share is the same and therefore e^* is optimal choice of effort level for period t'.

Thus, to prove that our balanced growth path equilibrium exists it suffices to show that there exists an effort level e^* that is optimal. To show that, suppose that the effort converges goes to one, then the stock of blueprints increases while the labor supply converges to zero, as the distribution of varieties by age is constant, which implies that labor share of output is constant, therefore wages diverge to infinity. If the effort level converges to zero, the stock of blueprints converges to zero and therefore wages fall to zero. Thus, the BGPE effort level is lower than 1 and higher than 0. In an interior equilibrium, the effort level $e^* \in (0, 1)$ satisfies the condition where the marginal cost of effort, given by the wage rate, is equal to the present value of profits accrued by a new blueprint times the productivity of blueprint production (given by the constant c). As wages and present value of profits vary continuously on the stock of blueprints which is determined by the effort level, by the intermediate value theorem there exists an effort level e^* such that wages are equal to the present value of profits accrued by a new blueprint times the productivity of blueprint production. Thus an BGPE exists. □

2.6 Intellectual property rights policy

Intellectual property policy in our framework implies that in the balanced growth path the optimal policy is a patent length $T \geq 0$ that maximizes the level of consumption: in the balanced growth path the growth rate of the economy does not vary with the level of effort employed in the production of varieties or with the level of competition across sectors. As in the balanced growth path the distribution of industry ages is constant and therefore varying the degree of competition across industries by changing T will vary the distribution of the productivity parameters Z across industries, but as the age distribution is constant, there is not growth in the average productivity parameter Z in the balanced growth path. Therefore, the optimal patent policy only has to maximize the level of consumption in the balanced growth path at a given date t.

Consider an economy in balanced growth path at some date t, is it reasonable to think that the optimal patent policy is to set $T = 0$ to maximize the benefits of competition? That is the optimal policy if the degree of profits that can be obtained under competition are high enough to yield an equilibrium level of effort for blueprint production under $T = 0$ that is not too low as to negate the benefits of increased competition. Simulations of the model, however, show that to obtain these high enough levels of profit under competition imply that firm's productivities should vary substantially to allow firms to obtain high profits under Bertrand competition and the simulations shows that the shocks to the firm's productivities that yield levels of variation in productivity across firms are much higher than suggested by empirical evidence. But our model suggests the possibility that optimal patent policy is for patents to not exist at all under certain parameter values, thus formalizes the argument in Boldrin and Levine [2013].

3 Simulation

The impact of a more strict patent law on innovation and its effects on higher deadweight losses due to market power and slower total factor productivity growth can be quantified in this simulation. We consider a functional form for the representative household's utility function $U(c, l)$ as follows

$$U(c, \ell) = c - K * B_t^{\frac{1-\alpha}{\alpha}} * \ell^2$$

where $K > 0$ is a constant. This functional form, which implicitly assigns a value for leisure that is an increasing function of the stock of blueprints, implies that the labor supply is a linear function of the labor share, which is constant in the balanced growth rate. Thus, labor supply is constant in the balanced growth path and changes only when the parameters of the model or the patent policy change the labor share of the economy.

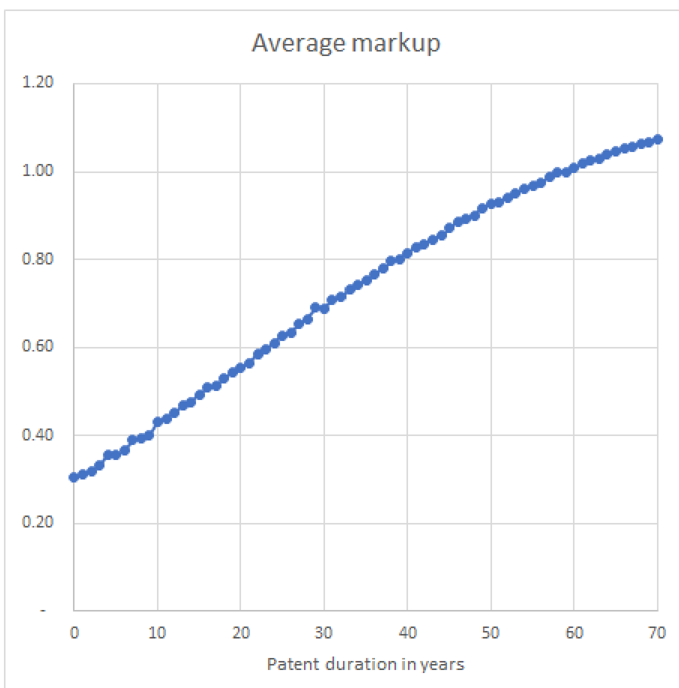

Figure 3: Mean economy wide markups, weighted by sales.

We calibrate the parameters so that under status quo policy (patent law grants 20 years of monopoly to inventors), the average economy-wide markup is between 1.5 to 1.6, the standard deviation of labor productivity across firms is 1/3, the discount rate is 4% per year, the firm's exit rate is 8% per year, and the expected growth rate of productivity of laggard firms is 2.5 times the rate of leaders.

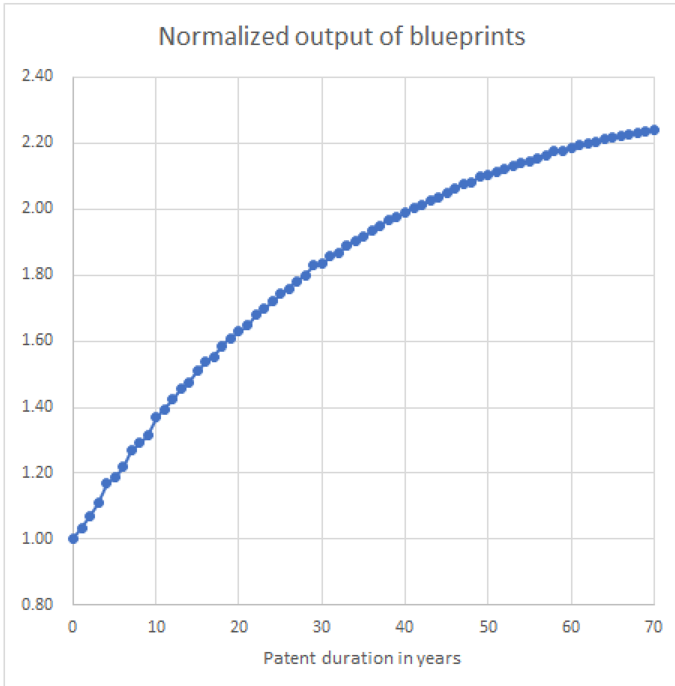

Figure 4: The output of blueprints is strictly increasing on strictness of patent law.

The results of our calibration are that substantial production of novel technologies ("blueprints") exists in equilibrium even when patents do not exist (that is, the policy that sets the patent duration to zero), and shifting from a policy of no patents to a policy where patents last 20 years increases the equilibrium production of blueprints by approximately two-thirds. Thus, our model implies that patents promote technological innovation but are not strictly required for a substantial innovative activity to exist (as argued by various authors such as Moser [2013]). Thus, if monopolies had no social costs, the optimal patent law would be for patents to last indefinitely.

While patents increase the incentives for innovation in novel blueprints, their effects are substantial both in terms of increasing the average markups as well as decreasing the average productivity of industries for existing blueprints.

In terms of output and utility, both the output and the representative household's utility approximately double from a regime of no patents to a regime of patents lasting 20 years.

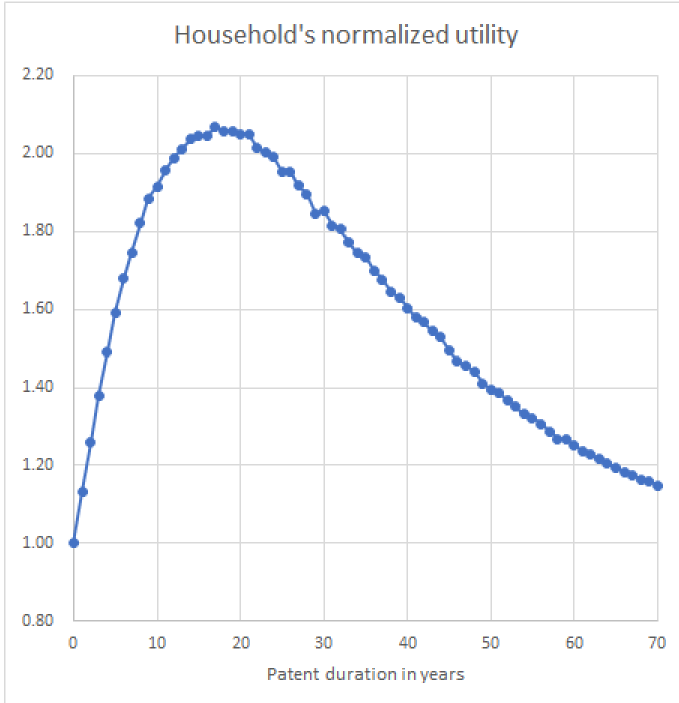

Figure 5: Utility of representative household varying the patent duration policy from 1 year to 70 years, utility is maximized at the patent duration policy of 17 years.

Therefore, our model refutes the claim in Boldrin and Levine [2013] that patents do not play an important social role. However, it does not conclude that our current policy of patents lasting 20 years is optimal. Instead, the optimal policy appears to consist of a slightly shorter patent duration, varying around 15 to 18 years.

The results of this calibration are robust to different specifications of the degree of knowledge spillovers from the productivity leaders. The reason why the degree of spillovers does not substantially change the optimal policy is that increasing θ decreases the durability of an advantageous position which reduces the present value of profits and, therefore, the incentives to produce new blueprints, decreasing welfare, but also it increases the intensity of the productivity gains from symbiotic competition which increases total factor productivity and therefore increases welfare. These two factors mostly cancel each other in our calibration.

It is possible for the model to endorse the idea from Boldrin and Levine [2013] that abolishing patents is a welfare-improving policy over our current patent policy of 20 years. To achieve this result, we have to increase the knowledge spillovers to much higher levels than which are suggested by the empirical evidence: the productivity growth rate of laggards needs to be

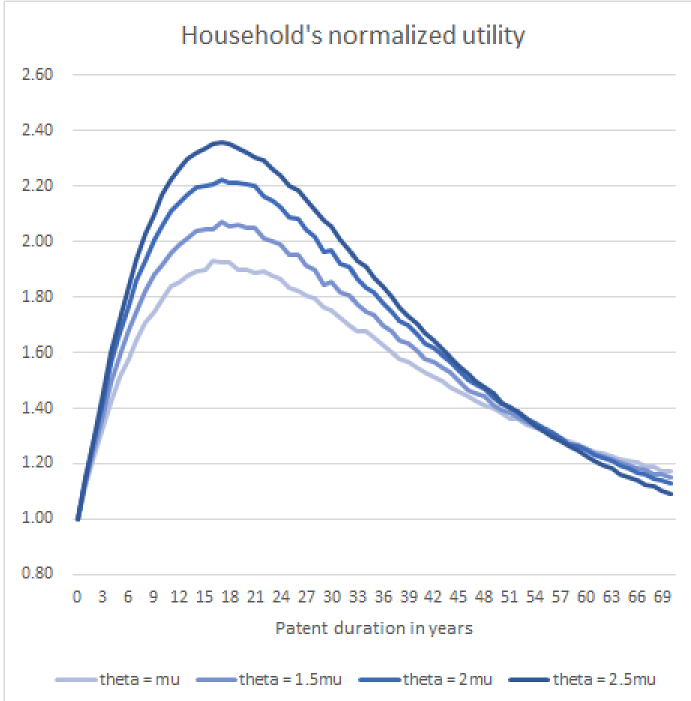

Figure 6: Robustness of the results to changes in the patent policy, the optimal patent policy is in the interval from 15 to 17 years varying θ from 50% lower than the calibration and two-thirds higher than calibration.

approximately 8 to 10 times larger than the productivity leaders. These very high knowledge spillovers produce increased TFP growth under competition such that the costs of a patent policy lasting 20 years exceed the benefits.

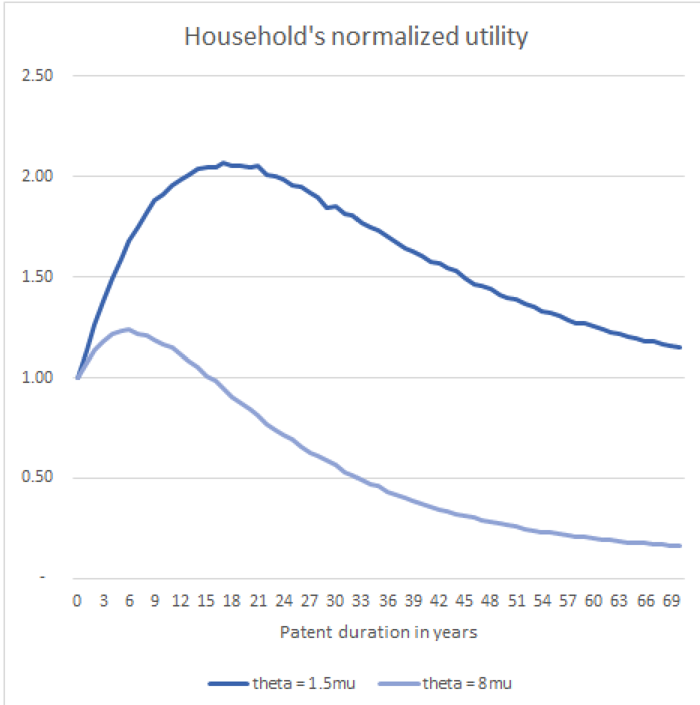

Figure 7: Technology spillovers have to be very high for a no-patent policy to be close to optimal as in this case the productivity of laggard firms grows at 9 times the rate of leaders.

4 Endogeneizing the technology

We have assumed that when we have two firms in an industry, each firm's productivity is subject to independent stochastic shocks and that if one firm is lagging behind the other firm, that firm can learn faster than the productivity leader. These two assumptions imply that the expected average productivity of the firms in the duopolized industry grows at the rate $\mu+\theta/2$, but if the industry is monopolized, the expected productivity of the monopolist grows only at the rate μ, which is strictly lower.

However, profits would increase if production costs decreased due to improved productivity. Thus, it would be rational for the monopolist to emulate the evolution of technology under duopoly that allows a faster rate of productivity growth. Therefore, consider the case where the monopolist can do so by hiring an additional manager to run a part of his company as a separate division, which means his firm runs as two division in "pseudo-competition" where one unit learns from the other as if they were a duopoly. The monopolist uses the maximum the the technology of the two divisions to produce the output and, to replicate the technology of the symbiotic competition in our environment, he commits to using the additional division forever after starting it.

In this situation, the expected growth of productivity of the monopolist is the same as the expected growth in productivity of the industry under duopoly: expected productivity will grow at the rate μ. The gain in expected profits from having two divisions instead of one is then given by the gain in productivity over time: Let $\Pi_{jt}^{m,t'}$ be the profits at time t if the monopolist has two divisions instead of one, which the second division being created in time $t' \leq t$. Clearly, if $t = t'$, then $\Pi_{jt}^{m,t'} = \Pi_{jt}^{m}$, as the gain in productivity from symbiotic competition between the two division only accrues over time.

Let time t be the time when the second division is created, then expected profits at time Δ after t are $E_t[\Pi_{jt+\Delta}^{m,2}]$ which satisfy

$$E_t[\Pi_{jt+\Delta}^{m,2}] = y_t(\frac{w_t}{\alpha})^{-\lambda}(1-\alpha)\exp(\lambda z_{jt} + \Delta(\mu + \theta/2)).$$

The expected profits of the monopolist without the additional division are

$$E_t[\Pi_{jt+\Delta}^{m}] = y_t(\frac{w_t}{\alpha})^{-\lambda}(1-\alpha)\exp(\lambda z_{jt} + \Delta\mu).$$

Let $G_t(\Delta)$ be the gain in expected profits of our monopolist after a length of time Δ from starting the additional division at time t, $G_t(\Delta)$ is given by

$$\begin{aligned} G_t(\Delta) &= E_t[\Pi_{jt+\Delta}^{m,2}] - E_t[\Pi_{jt+\Delta}^{m}] & (9)\\ &= y_t(\frac{w_t}{\alpha})^{-\lambda}(1-\alpha)(\exp(\lambda z_{jt} + \Delta(\mu+\theta/2)) - \exp(\lambda z_{jt} + \Delta\mu)) & (10)\\ &= y_t(\frac{w_t}{\alpha})^{-\lambda}(1-\alpha)\exp(\lambda z_{jt} + \Delta\mu)[\exp(\Delta\theta/2) - 1]. & (11) \end{aligned}$$

Let us assume, however, that the monopolist, however, cannot set up this additional division for free: he has to hire a manager to run this division and let C_{jt} be the flow cost the monopolist has to pay to the manager to manage his additional division in industry for variety j at period t. This cost represents the opportunity cost to keep the manager from exiting the firm and setting up his own firm in the industry, competing with the former monopolist. This opportunity cost means that we are assuming that the monopolist has all the bargaining power and makes a take-it-or leave it offer to the manager: the manager will accept managing the division for a value C_{jt} if and only if C_{jt} is equal or higher than the expected flow profits of one firm under duopoly conditional on the assumption that the average technology of the two firms in the industry under duopoly is the same as the technology the monopolist has at the time t, which is the maximum of the productivity parameters of the two divisions.

Thus, the monopolist's decision problem to set up the additional division, given our assumption the monopolist is committed after he decides to do so, is solved if the expected present value of the costs of setting the additional division are lower than the present value of the increase in expected profit flow.

Proposition 6. *There exists a discount rate* $\underline{r} > 0$ *such that for all discount rates higher than* $\underline{r}$ *the monopolist does not replicate symbiotic competition in its present profit maximizing equilibrium.*

Proof. The monopolist will choose to internally replicate symbiotic competition if and only if

$$\int_0^\infty \exp(-r\Delta)\left[G_t(\Delta) - E_t[C_{j,t+\Delta}]\right] d\Delta > 0.$$

Note that as $C_{j,t+\Delta}$ is the expected profit flow the monopolist expects that the manager obtains if he runs his own firm as one of the firms under a duopoly in time $t + \Delta$, this expected profit is always strictly positive and it grows according to the growth in productivity of the sector in duopoly (which grows at the expected rate $\mu + \theta/2$, the same as the expected growth rate of productivity under monopoly with two divisions) and according to the growth of the economy (as described by the parameter y_t).

Therefore, $C_{j,t+\Delta} > 0$ for every $\Delta \geq 0$ while $G_t(\Delta)$ converges to zero as Δ converges to zero. Thus, there exists a time interval $[0, \overline{\Delta}]$, with $\overline{\Delta} \in (0, \infty)$, such that the monopolist expects negative profits by hiring the manager. Note that $\overline{\Delta} > 0$ cannot be infinity: as Δ diverges to infinity, the faster rate of productivity increase implies that the increase in monopoly profits from shifting to the two division-system a relative increase in profits that is infinitely many times higher than under single division-system. Since profits under monopoly are higher than profits under duopoly given the same industry-average technology, therefore the increase in expected monopoly profits of a shift from one division to two divisions becomes higher than profits under competition for $\Delta > 0$ large enough. Therefore, profits are lower than the costs of manager from $[0, \overline{\Delta}]$ and higher than the costs of the manager during the period $[\overline{\Delta}, \infty)$. Thus, if the discount rate $r > 0$ is high enough the monopolist never finds it profit maximizing to replicate the effects of symbiotic competition under duopoly. □

This result implies that competition allows the benefits of patience to be replicated by myopic agents: if firms are perfectly patient they have natural incentives to value the future benefits of faster technological progress, when they are not patient the value of these benefits is heavily discounted, and so firms do not take them into consideration. Competition, however, induces the faster technological growth that occurs under a patient monopolist even if the firms are perfectly myopic.

5 Concluding Remarks

In this paper, we have developed a theoretical framework to study the interaction of competition with economic development to analyze the macroeconomic consequences of patent policy. As it is broadly understood, patents provide an incentive for innovation in providing monopoly rights for novel technologies, but our analysis also provides a framework to estimate the costs of these monopolies along two dimensions: first, by suppressing competition, monopolies reduce the speed of improvement in productivity for existing technologies and second, the increased market power of monopolists creates deadweight losses by lowering the quantity supplied in the market relative to the optimal level.

Our framework shows that learning spillovers across firms are a major factor in driving economic development: when firms compete neck-and-neck, learning spillovers drive the productivity frontier and enable faster growth than in industries that are monopolized by a single firm. This analysis shows that the benefits of competitive markets are not only allocative efficiency in a static sense but also dynamic: competition produces a sustained increase in total factor productivity growth.

While under certain assumptions our model can endorse the claim of some economists that shifting from our current patent system to a policy of no-patents would be welfare improving, calibration of the model according to the evidence provided by empirical studies suggests that the optimal patent policy is between 15 and 20 years and this result is robust to substantial changes in the parameters. Therefore we conclude that patent policy as it is currently practiced is likely superior to the abolition of patents.

Appendix

A Industry equilibrium

We study price competition between two firms. As we will see, it is without loss of generality to assume that only two firms operate in the industry as the goods are perfect substitutes. When costs are asymmetric, in standard Nash Equilibrium of a static competition model there is a continuum of equilibrium in the model, as stated in Lemma below:

Lemma 2. *Consider the case where $a_{1i} > a_{1i}$, then the equilibrium in a static price competition game (Bertrand competition) is given by the set of price pairs $\{(p,p) : p \in [\gamma_{i_1}, \gamma_{i_2}]\}$.*

Proof. Without loss of generality normalize wage in period t to 1. To see that both firms posting $p^* \in [1/A_{1t}, 1/A_{2t}]$ is an equilibrium note that: if both firms post p^* the firm 1, with better technology, thus lower cost, has profits $p - 1/A_{1t} \geq 0$, it does not find advantageous to increase prices and earn zero profits or to lower prices and earn lower profits. The firm with higher costs makes zero profits, raising prices means it will earn zero profits, lowering prices means it will earn negative profits.

To see that set exhausts all possible equilibrium, note that any other pair of prices besides the set (p^*, p^*) where $p^* \in [1/A_{1t}, 1/A_{2t}]$, is not consistent with equilibrium: the firm with better technology never finds it optimal to post a price strictly lower than the price posted by the firm with less advanced technology. Thus, all equilibrium involve both firms posting the same price, unless one firms posts a price lower than the marginal cost of both firms.

Then the best response by the other firm is to post any higher price. The firm making a loss will not find it advantageous to post that price. Thus, equilibrium involves both firms posting a price equal or higher than the marginal cost of the lowest cost firm. In addition, it is not an equilibrium for both firms to post a price higher than the marginal cost of the higher cost firm: in that situation the higher cost firm has an incentive to deviate. Thus, the set of all equilibria is for both firms to post $p^* \in [1/A_{1t}, 1/A_{2t}]$. □

The existence of a continuum of different profit levels that can be supported as a static Nash Equilibrium is an issue for the dynamic analysis of competition. It means that continuation values are undefined even in Markov Perfect Equilibrium, which implies that the equilibrium of the model is not uniquely determined. To restrict the set of equilibrium prices a reasonable strategy for refinement, we consider the assumption that firms have captive customers, that is, that there is a probability that some final goods producers do not have access to the competing firm and therefore are willing to purchase from that specific supplier at any posted price. This assumption is common in models of endogenous equilibrium price dispersion such as Butters Butters [1977], Varian Varian [1980], and Burdett and Judd Burdett and Judd [1983]. Several recent papers use the concept of captive buyers (in the sense of buyers who only have access to a single seller), for example Burdett and Menzio Burdett and Menzio [2018] (their calibration of the model according to the empirical evidence also implied that the fraction of captive customers should be circa 8%). Suppose that a fraction of final goods producer $\epsilon \in (0, 1/2)$ are captive of each firm which means that this fraction of final goods producers only have access to that firm when purchasing the intermediate input. Then, the Nash Equilibrium for the static competition game is unique and given by a randomized strategy described by a cumulative distribution function F_e, as e approximates 0 then the distribution of prices F_e converges in probability to $\min\{1/A_{1t}, 1/A_{2t}\}$. Thus, in the dynamic environment we study we use the notion of Markov Perfect Equilibrium augmented with captive customer refinement, in what we call Markov Refined Equilibrium, which yields an unique dynamic equilibrium that allows us to have a unique solution when we study the effects of technological spillovers.

Our solution concept is a Markov perfect equilibrium that is the limit of a sequence of e-equilibrium as $e \downarrow 0$, where firm's price posting strategies take into consideration a probability $\epsilon > 0$ that customers are captive (that is, they do not have the competing firm in their consideration set). The proposition below shows that with Markov perfect equilibrium we reduce the number of equilibria to one:

Proposition 7. *The Markov perfect equilibrium is unique, with equilibrium price for variety* *j given by*

$$p_j = \begin{cases} \gamma_{i_1}/\alpha & \textit{if } \alpha a_{i_2} < a_{i_2} \leq \alpha a_{i_1} < a_{i_1} \\ \gamma_{i_2} & \textit{if } \alpha a_{i_2} \leq \alpha a_{i_1} \leq a_{i_2} \leq a_{i_1} \\ \gamma_{i_1} & \textit{if } \alpha a_{i_1} \leq \alpha a_{i_2} \leq a_{i_1} \leq a_{i_2} \\ \gamma_{i_2}/\alpha & \textit{if } \alpha a_{i_1} < a_{i_1} \leq \alpha a_{i_2} < a_{i_2}. \end{cases}$$

Proof. Consider a firm $i \in \{i_1, i_2\}$. With the probability

$$\epsilon > 0$$

of the customer being captive a firm's profits take the form

$$\Pi_i(\boldsymbol{p}, \boldsymbol{X}_t) = \begin{cases} p_j - \gamma_i & \text{if } p_i < p_{-i} \text{or } p_j = p_{-j} \text{ and } X_{jt} > X_{-j,t} \\ (1/2 + \epsilon/2)(p_j - 1/X_{jt}) & \text{if } p_j = p_{-j} \text{ and } X_{jt} = X_{-j,t} \\ \epsilon(p_j - 1/X_{jt}) & \text{if } p_j > p_{-j} \text{or } p_j = p_{-j} \text{ and } X_{jt} < X_{-j,t} \end{cases}.$$

Equilibrium in mixed strategies in period T has the following structure: Without loss of generality assume $X_{1T} > X_{2T}$. The unique Nash equilibrium in period T consists of a pair of cumulative distribution functions (F_1, F_2) both with support $[\underline{p}, 1]$, where 1 is the monopoly price and $\underline{p}$ is the lower bound of the support.

The lower bound satisfies the following equal profit condition

$$\underline{p} - 1/X_{2T} = \epsilon(1 - 1/X_{2T}).$$

The mixed strategy F_{-j} satisfies the following equal profit condition

$$p[\epsilon + (1-\epsilon)F_{-j}(p)] = \underline{p} - 1/X_{jT},$$

which implies that to maintain the equal profit condition on the support of F_1, the cumulative distribution function F_2 has an atom at 1, as expected profits from posting $\underline{p}$ are $\underline{p} - 1/X_{jT}$ which are higher than $\epsilon(1 - 1/X_{2T})$. As $\epsilon \downarrow 0$, the lower bound of the support $\underline{p}$ converges to $1/X_{2T}$ and for $j \in \{1, 2\}$, $F_j(p)$ converges to 1 for any $p > 1/X_{2T}$. Thus for both firms the equilibrium distribution of prices converges in probability to $1/X_{2T}$. □

B Computation of Profits

Assuming that firms engage in price competition within an industry, prices will equal the minimum of the more productive firm's monopoly price and the less productive firm's marginal cost. Together with Proposition 7, the Dixit-Stiglitz framework gives tractable expressions for expected future prices and profits. In this section of the Appendix we derive these expressions.

B.1 Flow Profits

A monopolist's profit is given by

$$\begin{aligned}\Pi_{it}^m = y_t p_t^{-1/(1-\alpha)}(p_t - w_t/a_{it}) &= y_t \left(\frac{w_t}{\alpha a_{it}}\right)^{-1/(1-\alpha)} \left[\left(\frac{w_t}{\alpha a_{it}}\right) - \frac{w_t}{a_{it}}\right] \\ &= y_t \left(\frac{w_t}{\alpha a_{it}}\right)^{-\alpha/(1-\alpha)} (1-\alpha) \\ &= y_t \left(\frac{w_t}{\alpha}\right)^{-\lambda} (1-\alpha) e^{\lambda z_{it}},\end{aligned}$$

where $\lambda = \alpha/(1-\alpha)$ and $a_{it} = e^{z_{it}}$. Here, y_t and w_t correspond to aggregate variables, namely aggregated output and wages.

Proposition 8. *Consider a monopolist whose log productivity Z obeys $dZ_t = \mu dt + \sigma dW_t$, where W is a Wiener process. The monopolist's profit at time t is given by*

$$\begin{aligned}\mathbb{E}[\Pi_{it}^m] &= y_t \left(\frac{w_t}{\alpha}\right)^{-\lambda} (1-\alpha)\mathbb{E}[e^{\lambda z_t}] \\ &= y_t \left(\frac{w_t}{\alpha}\right)^{-\lambda} (1-\alpha) e^{\lambda[x_0 + \frac{1}{2}(\lambda+\mu)t]}.\end{aligned}$$

If $y_t = y$ and $w_t = w$ are constant over time, the normalized present value of profit equals

$$\begin{aligned}\int_0^\infty re^{-rt}\mathbb{E}[\Pi_{it}^m]dt &= y\left(\frac{w}{\alpha}\right)^{-\lambda}(1-\alpha)\int_0^\infty re^{-rt}e^{\lambda[x_0+\frac{1}{2}(\lambda\sigma^2+2\mu)t]}dt \\ &= y\left(\frac{w}{\alpha}\right)^{-\lambda}(1-\alpha)\frac{re^{\lambda x_0}}{r-\mu-\frac{1}{2}\lambda\sigma^2}\end{aligned}$$

assuming, of course, that the denominator is positive: $r > \mu + \frac{1}{2}\lambda\sigma^2$.

This simple result, which follows from Lemma 6 below, establishes the present value of profit for a monopolist whose productivity grows at rate μ. This serves as a benchmark for duopoly profit, which we study next.

A duopolist's profit depends on both its own and its competitors' productivity, since duopoly prices satisfy $p_t = \min\{w_t/(\alpha a_{it}), w_t/a_{jt}\}$ when $a_{it} \geq a_{jt}$. If $p_t = w_t/(\alpha a_{it})$, the duopololist's productivity advantage is big enough to warrant monopoly price and profit. If not, $p_t = w_t/a_{jt}$, and firm i earns

$$
\begin{aligned}
\Pi_{it}^d = y_t p_t^{-1/(1-\alpha)}(p_t - w_t/a_{it}) &= y_t \left(\frac{w_t}{a_{jt}}\right)^{-1/(1-\alpha)} \left[\frac{w_t}{a_{jt}} - \frac{w_t}{a_{it}}\right] \\
&= y_t \left(\frac{w_t}{a_{jt}}\right)^{-1/(1-\alpha)} \left[\frac{w_t}{a_{jt}} - \frac{w_t}{a_{it}}\right] \\
&= y_t \left(\frac{w_t}{a_{jt}}\right)^{-\alpha/(1-\alpha)} \left[1 - \frac{a_{jt}}{a_{it}}\right] \\
&= y_t w_t^{-\lambda} \left[e^{\lambda z_{jt}} - e^{(\lambda+1)z_{jt} - z_{it}}\right].
\end{aligned}
$$

A firm's flow profit is either Π_{it}^d or Π_{it}^m depending on whether $p_t = w_t/(\alpha a_{it})$ or w_t/a_{jt}, that is, the monopoly price or i's competitor's marginal cost.

From the point of view of time 0, a firm's expected profit will vary with the productivity of both firms. We calculate this expected profit next. Without loss, consider firm 1's profit. Of course, for firm 1 to produce, it must be the case that $z_{1t} \geq z_{2t}$. Moreover, if $z_{1t} \leq z_{2t} - \ln(\alpha)$, the prevailing market price will be firm 2's marginal cost, whereas if $z_{1t} > z_{2t} - \ln(\alpha)$ then the price will equal firm 1's monopoly level.

Therefore, firm 1's expected profit flow, $E[\Pi_{1t}]$, satisfies

$$\mathbb{E}[\Pi_{1t}] = \mathbb{E}[\Pi_{it}^d, z_1 + \ln\alpha \leq z_2 \leq z_1] + \mathbb{E}[\Pi_{it}^m, z_1 + \ln\alpha \geq z_2], \tag{*}$$

where $\mathbb{E}[R, A] = \mathbb{E}[R\mathbf{1}_A]$ and each of the terms on the right hand side have a closed form.

To find this closed form, we rely on two preliminary results that derive closed-form expressions for the Laplace transforms of X and Y, respectively. Let $\gamma = \mu + \frac{1}{2}\theta$ and

$$F_t(\lambda) = \int_{-\infty}^{\infty} e^{\lambda x} f_t(x) dx.$$

Lemma 3. *The Laplace transform of X, denoted by $F_t(\lambda)$, satisfies*

$$F_t(\lambda) = e^{\lambda[x_0 + \frac{1}{2}(\lambda\sigma^2 + 2\gamma)t]}.$$

Proof. Completing the square yields

$$\begin{aligned}
&\lambda x - \frac{1}{2\sigma^2 t}[(x-x_0)-\gamma t]^2 = -\frac{1}{2\sigma^2 t}\left[(x-x_0)^2 - 2(x-x_0)\gamma t + (\gamma t)^2 - 2\lambda x \sigma^2 t\right] \\
&= -\frac{1}{2\sigma^2 t}\left[(x-x_0)^2 - 2(x-x_0)(\gamma+\lambda\sigma^2)t + (\gamma t)^2 - 2\lambda x_0 \sigma^2 t\right] \\
&= -\frac{1}{2\sigma^2 t}\left[(x-x_0)^2 - 2(x-x_0)(\gamma+\lambda\sigma^2)t + ((\gamma+\lambda\sigma^2)t)^2 - 2\lambda x_0\sigma^2 t - \lambda^2\sigma^4 t^2 - 2\lambda\sigma^2\gamma t^2\right] \\
&= -\frac{1}{2\sigma^2 t}\left[(x-x_0) - (\gamma+\lambda\sigma^2)t\right]^2 + \lambda x_0 + \tfrac{1}{2}\lambda^2\sigma^2 t + \lambda\gamma t \\
&= -\frac{1}{2\sigma^2 t}\left[(x-x_0) - (\gamma+\lambda\sigma^2)t\right]^2 + \lambda[x_0 + \tfrac{1}{2}(\lambda\sigma^2 + 2\gamma)t]
\end{aligned}$$

Therefore, plugging this identity into the Laplace transform above gives

$$\begin{aligned}
\int_{-\infty}^{\infty} e^{\lambda x} f_t(x)dx &= \frac{1}{\sqrt{2\pi\sigma^2 t}}\int_{-\infty}^{\infty} e^{\lambda x} e^{-\frac{1}{2\sigma^2 t}[(x-x_0)-\gamma t]^2}dx \\
&= \frac{1}{\sqrt{2\pi\sigma^2 t}}\int_{-\infty}^{\infty} e^{\lambda[x_0+\frac{1}{2}(\lambda\sigma^2+2\gamma)t]} e^{-\frac{1}{2\sigma^2 t}[(x-x_0)-(\gamma+\lambda\sigma^2)t]^2}dx \\
&= e^{\lambda[x_0+\frac{1}{2}(\lambda\sigma^2+2\gamma)t]},
\end{aligned}$$

since $\frac{1}{\sqrt{2\pi\sigma^2 t}}\int_{-\infty}^{\infty} e^{-\frac{1}{2\sigma^2 t}[(x-x_0)-(\gamma+\lambda\sigma^2)t]^2}dx = 1.$ □

B.2 Present Value Profit

Let $V(X,Y)$ be firm 1's present value profit when the average productivity equals X and the productivity gap equals Y. Bertrand competition implies three regimes for firm 1: when it sets a monopoly price, when it charges firm 2's marginal cost, and when it is priced out by firm 2. Formally, flow profit is given by

$$\Pi_{1t} = \begin{cases} \Pi_{1t}^m & \text{if } Y > \bar{y}, \\ \Pi_{1t}^d & \text{if } 0 < Y < \bar{y}, \\ \Pi_{1t}^0 & \text{if } Y < 0, \end{cases}$$

where $\bar{y} = -\frac{1}{2}\ln\alpha$ and $\Pi_{1t}^0 = 0$.

Our goal is to compute V, which, formally, is given by

$$V_t = V(X_t, Y_t) = \mathbb{E}_t\left[\int_t^{\infty} e^{-\rho(\tau-t)}\Pi_{1\tau}d\tau\right].$$

Lemma 4. *$V(X_t, Y_t) = F(X_t)G(Y_t)$ is multiplicatively separable, with*

$$F(X_t) = yw^{-\lambda}e^{\lambda X_t} \quad \text{and} \quad G(Y_t) = \mathbb{E}_t\left[\int_t^\infty e^{-\beta(\tau-t)}g(Y_\tau)d\tau\right]$$

for some parameters y and w, where $\beta = \rho - \lambda\gamma - \frac{1}{2}\lambda^2\sigma^2$ and g is defined as

$$g(Y) = \begin{cases} g_m(Y) = (1-\alpha)\alpha^\lambda e^{\lambda Y} & \text{if } Y > \bar{y}, \\ g_d(Y) = e^{-\lambda Y} - e^{-(\lambda+2)Y} & \text{if } 0 < Y < \bar{y}, \\ g_0(Y) = 0 & \text{if } Y < 0. \end{cases}$$

Proof. First, flow profit Π_{1t} is clearly separable. To see this, let $f(X) = yw^{-\lambda}e^{\lambda X}$, where y and w are general equilibrium objects to be found later. By definition, $\Pi_{1t} = f(X_t)g(Y_t)$, and, since X_τ and Y_τ are independent,

$$\begin{aligned} V(X_t, Y_t) &= \mathbb{E}_t\left[\int_t^\infty e^{-\rho(\tau-t)}\Pi_{1\tau}d\tau\right] \\ &= \mathbb{E}_t\left[\int_t^\infty e^{-\rho(\tau-t)}f(X_\tau)g(Y_\tau)d\tau\right] \\ &= \int_t^\infty e^{-\rho(\tau-t)}\mathbb{E}_t[f(X_\tau)g(Y_\tau)]d\tau \\ &= \int_t^\infty e^{-\rho(\tau-t)}\mathbb{E}_t[f(X_\tau)]\mathbb{E}_t[g(Y_\tau)]d\tau \\ &= yw^{-\lambda}\int_t^\infty e^{-\rho(\tau-t)}\mathbb{E}_t[e^{\lambda X_\tau}]\mathbb{E}_t[g(Y_\tau)]d\tau \end{aligned}$$

By Lemma 6,

$$\mathbb{E}_t[e^{\lambda X_\tau}] = e^{\lambda[X_t + \frac{1}{2}(\lambda\sigma^2+2\gamma)(\tau-t)]},$$

therefore

$$V(X_t, Y_t) = yw^{-\lambda}e^{\lambda X_t}\int_t^\infty e^{-\beta(\tau-t)}\mathbb{E}_t[g(Y_\tau)]d\tau,$$

where $\beta = \rho - \frac{1}{2}\lambda(\lambda\sigma^2 + 2\gamma) = \rho - \lambda\gamma - \frac{1}{2}\lambda^2\sigma^2$. □

We will now solve for $G(Y_t)$ from the above lemma and its proof. There are three regimes for firm 1, as before: when it sets a monopoly price, when it charges firm 2's marginal cost, and when it is priced out by firm 2. Formally,

$$G(Y) = \begin{cases} G_m(Y) & \text{if } Y > \bar{y}, \\ G_d(Y) & \text{if } 0 < Y < \bar{y}, \\ G_0(Y) & \text{if } Y < 0, \end{cases}$$

We will now derive each part of G by solving its corresponding Fokker-Planck equation, and then glue them together via value matching and smooth pasting. Define the following roots: $\kappa_1 = (-\delta + \sqrt{\delta^2 + 2\sigma^2\beta})/\sigma^2 > 0$ and $\kappa_2 = -(\delta + \sqrt{\delta^2 + 2\sigma^2\beta})/\sigma^2 < 0$.

Lemma 5. *The function $G_0(Y)$ satisfies $\beta G_0 = \delta G_0' + \frac{1}{2}\sigma^2 G_0''$ and $G_0(Y) \to 0$ as $Y \to -\infty$ if and only if $G_0(Y) = A_0 e^{\kappa_1 Y}$ for some $A_0 \in \mathbb{R}$.*

Proof. The argument here is a standard derivation of the unique solution of the ODE above. Clearly, $\beta A_0 e^{\varepsilon Y} = \delta\kappa_1 A_0 e^{\varepsilon Y} + \frac{1}{2}\sigma^2\varepsilon^2 A_0 e^{\varepsilon Y}$ implies $\beta = \delta\varepsilon + \frac{1}{2}\sigma^2\varepsilon^2$, which holds for $\varepsilon = \kappa_1, \kappa_2$. Therefore, $G_0(Y) = A_0 e^{\kappa_1 Y} + B_0 e^{\kappa_2 Y}$, but the limiting condition $G_0(Y) \to 0$ as $Y \to -\infty$ implies that $B_0 = 0$, which gives the solution $G_0(Y) = A_0 e^{\kappa_1 Y}$ for some A_0. □

Lemma 6. *Let $-\kappa_1 < \omega < -\kappa_2$. The function $H(Y)$ satisfies*

$$\beta H = C e^{\omega Y} - \delta H' + \tfrac{1}{2}\sigma^2 H''$$

if and only if

$$H(Y) = A e^{-\kappa_1 Y} + B e^{-\kappa_2 Y} + D e^{\omega Y}$$

for some coefficients A and B, where $D = C/(\beta + \delta\omega - \frac{1}{2}\sigma^2\omega^2)$.

Proof. The linear ODE above has a unique solution given boundary conditions. Assume that $-\kappa_1 < \omega < -\kappa_2$. Guessing a form $A e^{\varepsilon Y} + D e^{\omega Y}$ implies that

$$r(A e^{\varepsilon Y} + D e^{\omega Y}) = C e^{\omega Y} - \delta(\varepsilon A e^{\varepsilon Y} + \omega D e^{\omega Y}) + \tfrac{1}{2}\sigma^2(\varepsilon^2 A e^{\varepsilon Y} + \omega^2 D e^{\omega Y}).$$

Matching coefficients multiplying $e^{\varepsilon Y}$ and $e^{\omega Y}$ implies that (i) $r = -\delta\varepsilon + \frac{1}{2}\sigma^2\varepsilon^2$, with roots $\varepsilon = -\kappa_1, -\kappa_2$, and (ii) $D = C/(r + \delta\omega - \frac{1}{2}\sigma^2\omega^2)$. Therefore,

$$H(Y) = A e^{-\kappa_1 Y} + B e^{-\kappa_2 Y} + D e^{\omega Y}$$

for some coefficients A and B to be determined by boundary conditions, as claimed. □

We seek G_k for $k \in \{m, d\}$ satisfying the ODE in Lemma 9 above such that

$$\beta G_m = g_m - \delta G_m' + \tfrac{1}{2}\sigma^2 G_m'' \quad \text{and} \quad \beta G_d = g_d - \delta G_d' + \tfrac{1}{2}\sigma^2 G_d'',$$

where g_m and g_d are defined in Lemma 7 above.

Applying Lemma 9, we obtain the following general solution for G_m:

$$G_m(Y) = A_m e^{-\kappa_1 Y} + B_m e^{-\kappa_2 Y} + D_m e^{\lambda Y}.$$

The coefficient D_m, found by substitution and coefficient matching in the ODE above, equals

$$D_m = \frac{(1-\alpha)\alpha^\lambda}{\beta + \lambda\delta - \frac{1}{2}\lambda^2\sigma^2}.$$

For this to make sense as a present discounted value, we need $\beta + \lambda\delta - \frac{1}{2}\lambda^2\sigma^2 > 0$. Since $\beta = \rho - \lambda\gamma - \frac{1}{2}\lambda^2\sigma^2$, we require $\rho > \lambda\mu + \lambda^2\sigma^2$ for D_m to be interpreted as a present value. The coefficients A_m and B_m reflect changes in the value associated with switching from regime m to regime d. However, as $Y \to \infty$, a switch from m to d becomes arbitrarily far away in time, with present value asymptotically negligible. As such the value must converge to $D_m e^{\lambda Y}$. This implies that $B_m = 0$, since otherwise $B_m e^{-\kappa_2 Y}$ would explode as $Y \to \infty$ because $\kappa_2 < 0$.

We now turn to G_d. Again by Lemma 9,

$$G_d(Y) = A_d e^{-\kappa_1 Y} + B_d e^{-\kappa_2 Y} + C_d e^{-\lambda Y} - D_d e^{-(\lambda+2)Y},$$

where

$$C_d = \frac{1}{\beta - \lambda\delta - \frac{1}{2}\lambda^2\sigma^2} \quad \text{and} \quad D_d = \frac{1}{\beta - (\lambda+2)\delta - \frac{1}{2}(\lambda+2)^2\sigma^2}.$$

Again, for these present value coefficients to be well defined, we require $\beta > \lambda\delta + \frac{1}{2}\lambda^2\sigma^2$ and $\beta > (\lambda+2)\delta + \frac{1}{2}(\lambda+2)^2\sigma^2$. Of course, the latter implies the former, which in turn implies the restriction for D_m above. We have thus established the following result.

Lemma 7. *If $\beta > (\lambda+2)\delta + \frac{1}{2}(\lambda+2)^2\sigma^2$ then G satisfies*

$$G(Y) = \begin{cases} A_m e^{-\kappa_1 Y} + D_m e^{\lambda Y} & \text{if } Y > \bar{y}, \\ A_d e^{-\kappa_1 Y} + B_d e^{-\kappa_2 Y} + C_d e^{-\lambda Y} - D_d e^{-(\lambda+2)Y} & \text{if } 0 < Y < \bar{y}, \\ A_0 e^{\kappa_1 Y} & \text{if } Y < 0, \end{cases}$$

for some A_0, A_d, B_d, A_m to be determined below, with C_d, D_d and D_m defined above.

Our last task is to find the coefficients A_0, A_d, B_d and A_m. These are found by connecting G_m, G_d and G_0 via smooth pasting and value matching:

$$G_m(\bar{y}) = G_d(\bar{y}), \quad G_d(0) = G_0(0),$$

$$G'_m(\bar{y}) = G'_d(\bar{y}), \quad G'_d(0) = G'_0(0).$$

These are four equations in four unknowns.

Proposition 9. *Let* $\beta > (\lambda+2)\delta + \frac{1}{2}(\lambda+2)^2\sigma^2$ *and define the following constants:*

$$\begin{aligned}
R_1 &= C_d - D_d,\\
R_2 &= \frac{(\kappa_1+\lambda+2)D_d - (\kappa_1+\lambda)C_d}{2\kappa_1},\\
R_3 &= e^{(\kappa_1-\lambda)\bar{y}}C_d - e^{(\kappa_1-\lambda-2)\bar{y}}D_d - e^{(\kappa_1+\lambda)\bar{y}}D_m,\\
R_4 &= (\lambda+2)e^{-(\lambda+2)\bar{y}}D_d - \lambda e^{-\lambda\bar{y}}C_d - \lambda e^{\lambda\bar{y}}D_m.
\end{aligned}$$

The coefficients of the previous lemma are equal to

$$\begin{aligned}
A_0 &= R_1 + R_2 - \frac{\kappa_1 e^{-\kappa_1\bar{y}}R_3 + R_4}{2\kappa_1 e^{-\kappa_2\bar{y}}},\\
A_d &= R_2 - \frac{1}{2}\left(1+\frac{\kappa_2}{\kappa_1}\right)\frac{\kappa_1 e^{-\kappa_1\bar{y}}R_3 + R_4}{(\kappa_2-\kappa_1)e^{-\kappa_2\bar{y}}},\\
A_m &= R_2 - \frac{1}{2}\left(1+\frac{\kappa_2}{\kappa_1}\right)\frac{\kappa_1 e^{-\kappa_1\bar{y}}R_3 + R_4}{(\kappa_2-\kappa_1)e^{-\kappa_2\bar{y}}} + R_3 + \frac{\kappa_1 R_3 + e^{\kappa_1\bar{y}}R_4}{(\kappa_2-\kappa_1)}, \qquad \textit{and}\\
B_d &= \frac{\kappa_1 e^{-\kappa_1\bar{y}}R_3 + R_4}{(\kappa_2-\kappa_1)e^{-\kappa_2\bar{y}}}.
\end{aligned}$$

This completes the definition of firm 1's present value profit. Firm 2's present value profit is obtained by antisymmetry about $Y=0$*: if* $\tilde{G}$ *is the corresponding* G*-function for firm 2 then*

$$\tilde{G}(Y) = G(-Y).$$

Proof. By the previous lemmas,

$$\begin{aligned}
G_0(Y) &= A_0 e^{\kappa_1 Y},\\
G_d(Y) &= A_d e^{-\kappa_1 Y} + B_d e^{-\kappa_2 Y} + C_d e^{-\lambda Y} - D_d e^{-(\lambda+2)Y}, \qquad \text{and}\\
G_m(Y) &= A_m e^{-\kappa_1 Y} + D_m e^{\lambda Y}.
\end{aligned}$$

Plugging in the value matching and smooth pasting conditions gives

$$\begin{aligned}
A_0 &= A_d + B_d + C_d - D_d,\\
\kappa_1 A_0 &= -\kappa_1 A_d - \kappa_2 B_d - \lambda C_d + (\lambda+2)D_d,\\
A_m e^{-\kappa_1\bar{y}} + D_m e^{\lambda\bar{y}} &= A_d e^{-\kappa_1\bar{y}} + B_d e^{-\kappa_2\bar{y}} + C_d e^{-\lambda\bar{y}} - D_d e^{-(\lambda+2)\bar{y}},\\
-\kappa_1 A_m e^{-\kappa_1\bar{y}} + \lambda D_m e^{\lambda\bar{y}} &= -\kappa_1 A_d e^{-\kappa_1\bar{y}} - \kappa_2 B_d e^{-\kappa_2\bar{y}} - \lambda C_d e^{-\lambda\bar{y}} + (\lambda+2)D_d e^{-(\lambda+2)\bar{y}}.
\end{aligned}$$

We now solve these four linear equations in four unknowns. Define the following constants:

$$\begin{aligned}
R_1 &= C_d - D_d, \\
R_2 &= \frac{(\kappa_1 + \lambda + 2)D_d - (\kappa_1 + \lambda)C_d}{2\kappa_1}, \\
R_3 &= e^{(\kappa_1 - \lambda)\bar{y}} C_d - e^{(\kappa_1 - \lambda - 2)\bar{y}} D_d - e^{(\kappa_1 + \lambda)\bar{y}} D_m, \\
R_4 &= (\lambda + 2)e^{-(\lambda+2)\bar{y}} D_d - \lambda e^{-\lambda\bar{y}} C_d - \lambda e^{\lambda\bar{y}} D_m.
\end{aligned}$$

Symbolic arithmetic yields

$$\begin{aligned}
A_0 &= A_d + B_d + R_1, \\
A_d &= R_2 - \frac{1}{2}\left(1 + \frac{\kappa_2}{\kappa_1}\right) B_d, \\
A_m &= A_d + e^{(\kappa_1 - \kappa_2)\bar{y}} B_d + R_3, \\
B_d &= \frac{\kappa_1 e^{-\kappa_1\bar{y}} R_3 + R_4}{(\kappa_2 - \kappa_1)e^{-\kappa_2\bar{y}}}.
\end{aligned}$$

Therefore,

$$A_d = R_2 - \frac{1}{2}\left(1 + \frac{\kappa_2}{\kappa_1}\right) \frac{\kappa_1 e^{-\kappa_1\bar{y}} R_3 + R_4}{(\kappa_2 - \kappa_1)e^{-\kappa_2\bar{y}}}.$$

This implies that

$$\begin{aligned}
A_0 &= R_1 + R_2 - \frac{1}{2}\left(1 + \frac{\kappa_2}{\kappa_1}\right) \frac{\kappa_1 e^{-\kappa_1\bar{y}} R_3 + R_4}{(\kappa_2 - \kappa_1)e^{-\kappa_2\bar{y}}} + \frac{\kappa_1 e^{-\kappa_1\bar{y}} R_3 + R_4}{(\kappa_2 - \kappa_1)e^{-\kappa_2\bar{y}}} \\
&= R_1 + R_2 - \frac{1}{2}\left(\frac{\kappa_2}{\kappa_1} - 1\right) \frac{\kappa_1 e^{-\kappa_1\bar{y}} R_3 + R_4}{(\kappa_2 - \kappa_1)e^{-\kappa_2\bar{y}}} \\
&= R_1 + R_2 - \frac{\kappa_1 e^{-\kappa_1\bar{y}} R_3 + R_4}{2\kappa_1 e^{-\kappa_2\bar{y}}}.
\end{aligned}$$

Finally,

$$\begin{aligned}
A_m &= R_2 - \frac{1}{2}\left(1 + \frac{\kappa_2}{\kappa_1}\right) \frac{\kappa_1 e^{-\kappa_1\bar{y}} R_3 + R_4}{(\kappa_2 - \kappa_1)e^{-\kappa_2\bar{y}}} + R_3 + e^{(\kappa_1 - \kappa_2)\bar{y}} \frac{\kappa_1 e^{-\kappa_1\bar{y}} R_3 + R_4}{(\kappa_2 - \kappa_1)e^{-\kappa_2\bar{y}}} \\
&= R_2 - \frac{1}{2}\left(1 + \frac{\kappa_2}{\kappa_1}\right) \frac{\kappa_1 e^{-\kappa_1\bar{y}} R_3 + R_4}{(\kappa_2 - \kappa_1)e^{-\kappa_2\bar{y}}} + R_3 + \frac{\kappa_1 R_3 + e^{\kappa_1\bar{y}} R_4}{(\kappa_2 - \kappa_1)}.
\end{aligned}$$

This completes the proof. □

C Proofs

C.1 Proof of Proposition 1

[insert proof]

C.2 Proof of Proposition 2

[insert proof]

C.3 Proof of Lemma 1

We will show that X and Y are distributed independently, as the rest of the Lemma follows by direct calculation. Indeed, clearly W^+ and W^- are uncorrelated:

$$\begin{aligned}
\mathbb{E}[(W_t^+ - W_s^+)(W_t^- - W_s^-)] &= \tfrac{1}{4}\mathbb{E}[(W_{jt} + W_{kt} - W_{js} - W_{ks})(W_{jt} - W_{kt} - W_{js} + W_{ks})] \\
&= \mathbb{E}[(W_{jt} - W_{js} + (W_{kt} - W_{ks}))(W_{jt} - W_{js} - (W_{kt} - W_{ks}))] \\
&= \mathbb{E}[(W_{jt} - W_{js})^2] - \mathbb{E}[(W_{kt} - W_{ks})^2] = 0.
\end{aligned}$$

Independence now follows by joint normality of W_j and W_k.

C.4 Proof of Proposition 3

[insert proof]

C.5 Proof of Proposition 4

[insert proof]

References

Philippe Aghion, Nick Bloom, Richard Blundell, Rachel Griffith, and Peter Howitt. Competition and innovation: An inverted-u relationship. *The Quarterly Journal of Economics*, 120(2):701–728, 2005. 3

Kenneth J. Arrow. The economic implications of learning by doing. *The Review of Economic Studies*, 29:155–173, 1962. 3

David Besanko, Ulrich Doraszelki, Yaroslav Kryukov, and Mark Satterthwaite. Learning-by-doing, organizational forgetting, and industry dynamics. *American Economic Review*, 78: 453–508, 2010. 3, 4

David Besanko, Ulrich Doraszelki, and Yaroslav Kryukov. The economics of predation: What drives pricing when there is learning-by-doing? *American Economic Review*, 104:868–897, 2014. 3, 4

David Besanko, Ulrich Doraszelski, and Yaroslav Kryukov. How efficient is dynamic competition? the case of price as investment. *American Economic Review*, 109(9):3339–3364, 2019. 3, 4

Michele Boldrin and David K. Levine. The case against patents. *Journal of Economic Perspectives*, 27:3–22, 2013. 1, 4, 8, 17, 20

Kevin A. Bryan and Heidi L. Williams. Innovation: Market failures and public policies. In Kate Ho, Ali Hortacsu, and Alessandro Lizzeri, editors, *Handbook of Industrial Organization 5(1)*, pages 281–388. Elsevier, Amsterdam, Netherlands, 2021. 1

Kenneth Burdett and Kenneth L. Judd. Equilibrium price dispersion. *Econometrica*, 51(4): 955–969, 1983. 28

Kenneth Burdett and Guido Menzio. The (q,s,s) pricing rule. *Review of Economic Studies*, 85:892–928, 2018. 28

Gerard R. Butters. Equilibrium distributions of sales and advertising prices. *Review of Economic Studies*, 44(3):465–491, 1977. 28

F. A. Hayek. Der wettbewerb als entdeckungsverfahren. *Kieler Vorträge*, 56:1–20, 1968. 2

Robert Lucas and Benjamin Moll. Knowledge growth and the allocation of time. *Journal of Political Economy*, 122(1):1–51, 2014. 8

Robert E. Lucas. On the mechanics of economic development. *Journal of Monetary Economics*, 22:3–42, 1988. 3

Eric Maskin and Jean Tirole. A theory of dynamic oligopoly, i: Overview and quantity competition with large fixed costs. *Econometrica*, 56:549–569, 1988a. 3

Eric Maskin and Jean Tirole. A theory of dynamic oligopoly, ii: Price competition, kinked demand curves, and edgeworth cycles. *Econometrica*, 56:571–599, 1988b. 3

Petra Moser. Patents and innovation: Evidence from economic history. *Journal of Economic Perspectives*, 27(1):23–44, 2013. 19

William D. Nordhaus. *Invention, Growth, and Welfare: A Theoretical Treatment of Technological Change*. MIT Press, Cambridge, MA, United States, 1969. 1

Paul M. Romer. Increasing returns and long-run growth. *Journal of Political Economy*, 94: 1002–1037, 1986. 3

Joseph Schumpeter. *The Theory of Economic Development*. Routledge, Oxfordshire, United Kingdom, 1980. 3

Nancy L. Stokey. Learning by doing and the introduction of new goods. *Journal of Political Economy*, 96:701–717, 1988. 3

Hal R. Varian. A model of sales. *American Economic Review*, 70(4):651–659, 1980. 28

Heidi L. Williams. How do patents affect research investments? *Annual Reviews of Economics*, 9:441–469, 2017. 1