Market Structure and the Direction of Technological Change*

Matthew Mitchell[†] and Andrzej Skrzypacz[‡]

February 13, 2006

Abstract

We study a model where innovation comes in two varieties: improvements on existing products, and new products that expand the scope of a technology. We make this distinction in order to highlight how market structure can determine not only the quantity of innovation but also its direction. We study two market structures. The first is the canonical one from the endogenous growth literature, where innovations can be developed by anyone, and developers market their own innovations. We then consider a more concentrated industry, where all innovation and pricing for a given technology is monopolized. We study the implications of the different market structures for both types of innovation, focusing on differences they induce in the direction of technological change. We apply our model model to the case of a hardware/software technology and analyze which market structure offers greater profits to a monopolist who can monopolize either hardware or software. We compare social welfare across the market structures, and discuss whether one type of innovation should be subsidized over another.

*We thank Philippe Aghion, Paul Romer, and participants and the 4th Duke/Texas/Northwestern IO theory conference for suggestions.

[†] University of Iowa.

[‡] Stanford GSB.

1 Introduction

In this paper we analyze the relationship between market structure and incentives to innovate. We argue that market structure is an important determinant of not only the quantity but also the direction of technological change. We present a framework to study the decisions of researchers and firms about both the intensity of research and the allocation of research efforts. This multi-dimensionality of research decisions is the new tradeoff we focus on. We argue that market power weakens incentives to generate innovations that increase the efficiency of existing uses of the technology and strengthens the incentives to innovate by generating new uses for the technology. This tradeoff creates different innovation dynamics across different market structures, and leads to insights for firm strategy and public policy.

In Section 2 we introduce a continuous time model of innovation. We consider innovations of two types: efficiency improvements to existing uses of the technology (efficiency innovations) and developments of new uses for the technology (scope innovations). The main differentiating feature is that improvements are characterized by larger "business stealing" that is, the profits accrued to new, improved versions come with decreased profits of old products. On the other hand, new uses have relatively little impact on the profits of existing lines. We assume that developing a new use is more costly than developing an improvement to an existing use and that there are physical barriers on the number of possible uses (while the scope for improvements is unbounded). The last assumption is sufficient for us to characterize contrast cases where potential scope is relatively important or unimportant.

The main economic force is that firms with market power over all efficiency levels of a given use internalize the "business stealing" effects of efficiency improvements. As a result, compared with a competitive market (with firms that each own only one efficiency level for one use) firms with market power have weaker incentives to develop improvements but stronger incentives to develop new uses. These differences influence innovation dynamics, long-run product structure and welfare.

There are many ways in which innovations differ in practice. The differences we focus on (we refer to them as efficiency vs. scope dimensions or improvements vs. new use dimensions) can be illustrated via the following examples. In the market for computer operating systems, some innovations are improvements to existing functions, which allow a more efficient execution of some tasks, while other innovations are new functions that allow the computer to do new tasks. Similar distinction can be made for business applications or video games. In the market for pharmaceuticals some innovation comes in the form of an improvement to an existing cure, but also often researchers come up with a cure for a previously uncured ailment. In general, the relevant distinction comes from the demand side rather than from the technological side.

For simplicity we study two polar market structures. On the one hand, in Section 4 we consider a monopolist who controls all possible innovation, both in terms of efficiency levels and scope of the technology. Such a situation leads to expansion in scope for the reason described above: the monopolist seeks to avoid displacing his own profits on existing products and does not have to worry about having its profits taken over by a competitor. Over time, the amount of research done by the monopolist decreases, as the physical limits to the number of possible product lines are approached.

On the other hand, we consider in Section 3 a competitive innovation case that mirrors the common setup of the endogenous growth literature. Each efficiency level is monopolized, but all innovations compete with one another, and there is free entry into innovation. Innovation often starts slower than in the monopoly case, but as there are more and more products developed the total intensity of research increases. At the same time, there is a shift of research from new uses to improvements and eventually the market reaches a steady-state in which further research is done solely on improvements. In the long run the competition leads to a more narrow scope than the monopoly achieves, but innovation never dries out.

Comparisons across these two polar market structures depend crucially on the potential scope of technology – i.e. on how large is the limit on potential product lines. The larger it is, the stronger are the monopolist incentives to innovate (while the competitive market is usually unaffected), which makes this market structure more efficient.

We apply this logic to a classic IO question of the optimal pricing of hardware and software. In particular, in Section 5 we consider a hardware monopolist who cannot commit to future innovation levels in software, but can commit to a market structure for software. By giving up market power in the software market, he might enhance the value of the hardware he sells by increasing the amount of software that results.

As our prior intuition suggests, the choice hinges on the degree to which scope innovations are possible. When the hardware promises a wide potential for software to undertake different jobs, the monopolist is best off providing the software himself. However, if the potential for scope is not so great, the monopolist may be better off giving up market power in software to enhance the value of hardware, since consumers realize that business stealing concerns will limit the rate of innovation for a software monopolist.

Our logic extends naturally to the question of social welfare and market structure, which we touch upon in Section 6. Similar intuition suggests that the socially optimal market structure depends on the potential a technology has for being applied to a wide variety of areas. More generally, this points out that considerations of the potential directions of technological change might be important in answering questions about which market structure most favors innovation and for informing intellectual property rights.

We also analyze how (and if) the government should subsidize innovation. To answer this question, in Section 7 we compare a steady-state of a competitive market and calculate the benefits to research on the two types of innovation. An efficiency innovation creates a permanent increase in the level of welfare. In contrast, a scope innovation, even if it has a smaller immediate impact, it increases private returns to research and leads to a higher rate of innovation (reflected in the growth rate of welfare).

This result has several implications. First, a society might want to think about not only how much innovation will result from a particular stimulus, but also might care about the type. For example, the government might prefer that the NSF focuses on basic research to develop new research areas, rather than on applications that increase the efficiency of the earlier innovations.

Furthermore, the government might care about the side of the research market in which it intervenes. Typically the economic incidence of a tax or subsidy is independent of its statutory incidence, so this concern does not arise. Here, though, as in Romer (supply versus demand), the government might be better off subsidizing innovations in particular areas on the demand side, rather than giving a broad subsidy to the factors of production engaged in research.

The relationship between market structure and incentives to innovate is a fundamental topic in Industrial Organization. Since innovation and productivity growth are so tightly connected, the topic is also relevant in macroeconomics. Most of the existing research focuses attention on the impact of market structure on the amount of innovation. However, recent research such as Acemoglu (2002) stresses that the direction of technological change.

In this paper we argue that market structure is an important determinant of the direction of technological change.

Moreover, the model that we develop provides insight into the classic literature on the measurement of the relationship between scale and innovation (see, for instance, Scherer (1980), and, more recently, Aghion, et al. (2005)). One way to interpret the two market structures we study is that it entails a single large firm, with more market power, versus an industry with many smaller firms. If one measured the relationship between firm size and innovation using data from our two polar market structures, there would be a variety of pitfalls. First, the industry life-cycle would matter; initially the closed standard (large firms) do more innovating, but later the open standard overtakes it. Moreover, the closed standard does different sorts of innovation than the open standard, so comparing conventional measures of innovation like patent counts might not be a consistent comparison across market structures.

Our model has the potential to be developed in ways that have implications for other important policy issues. In discussing antitrust issues in industries like computer software, Schmalensee (2000) suggests the importance of the industry's Schumpeterian character. We further this line by using modern Schumpeterian models to frame both positive and normative discussion of such industries. We choose a structure that resembles most closely that of Grossman and Helpman (1991), but of course the same ideas could be embedded in a model along the lines of Aghion and Howitt (1992). Our model adds endogenous variety, and in that sense is similar to a long line of growth theory papers such as Romer (1987).

Our model provides a rationale for paying attention to whether support for innovation takes place on the supply or demand side, as Romer (2000) stresses. Since different types of innovations have different social benefits, it may be that the government has an interest in guiding through targeted support (such as NSF support for "basic" research) a particular type of innovation

The commitment benefit of opening markets that we study has been developed in other contexts. For example, Shepard (1987) showed that licensing technologies to multiple competing firms can serve as a commitment device for a monopolist seeking to deliver high quality. Similarly, Economides (1996) discusses how a firm selling a product with network benefits might allow competition as a way to commit to a large customer base.

Our model has implications relative to the long literature on durable goods monopolies (including Coase (1972), Stokey (1981), and Bulow (1982)) where a monopolist can sell an object up-front, but faces commitment problems in future actions. Our model has two dimensions for future actions, efficiency improvements and increases in scope. The former involves the classic commitment problem for the monopolist; the latter, however, brings a comparative advantage for the monopolist by being able to coordinate innovative activity in a way that internalizes the effects on existing applications. In our model, when the number of applications is small, quality improvements are relatively important, and solving commitment problems with software competition is relatively attractive. On the other hand, when the number of potential applications is large, the software monopolist benefits from internalizing business stealing.

2 The Model

2.1 Standards and Applications

We study a continuous time, infinite horizon model. In each instant, consumers derive utility from $\bar{N}$ *functions* that are related under a common *standard.* The standard allows a function to be accomplished via a specialized *application.* A given application j for function i has quality $q_i^j \geq 1$ per physical unit. Without an application, the consumer can accomplish the function directly at a fixed quality $q_i^0 = \phi \leq 1$ per physical unit. If the standard were a particular type of computer, examples of directly accomplishing a function, instead of using a specialized application, would be using pencil and paper instead of a spreadsheet or a typewriter instead of a word processor. We will consider two ways in which applications are consumed. Applications may be consumed directly (which we call *stand alone applications*), or consumers might have to purchase first a piece of hardware which allows the purchase of applications as software (termed *software applications*).

We define the standard to be *open* if each application is owned by a different firm. The standard is *closed* if all applications are owned by the same firm and only this firm can innovate to obtain new innovations.[1]

[1]In the future we hope to consider a mixed setup with both multi-application and single-application firms competing in a market.

2.2 Preferences

For a given function i, the representative consumer consumes d_i^j physical units of application j. This leads to e_i efficiency units of the function, where[2]

$$e_i = \sum_j q_i^j d_i^j$$

The representative consumer's instantaneous utility from a bundle of efficiency units $\{e_i\}$ of the various functions is

$$u(\{e_i\}) - \sum_i \sum_j p_i^j d_i^j - h$$

where p_i^j is the price paid for application j on function i and h is the amortized cost of the hardware.

In equilibrium consumers will choose only one quality level, the highest quality, denoted simply q_i, per function. Therefore, given consumption of d_i units of that quality level of application i, utility is

$$u(\{q_i d_i\}) - \sum_i p_i d_i - h$$

To simplify the analysis, for most of the paper we parameterize the utility function to have a simple form:

Assumption 1: $u(\{q_i d_i\}) = \sum_i \ln(q_i d_i)$.

This utility implies that the demands are independent across different applications and the representative consumer spends a constant share of his income on every application. In particular, if the representative consumer buys applications q_i^j at price p, his demand is $d_i^j = 1/p$.

2.3 Output Production: Firms and Competition

Qualities fall on a ladder with rungs of size $\lambda > 1$; i.e., for the jth quality level on ladder i, $q_i^j = \lambda q_i^{j-1}$. The first application has quality $q_i^1 = \lambda$. Each physical unit requires one unit of labor to be produced, so the marginal cost of production, per physical unit, is normalized to 1.

[2] The preference structure for a given application follows Grossman and Helpman (1991).

Competition Between Qualities in an Open Standard First, following the endogenous growth literature such as Aghion and Howitt (1992) and Grossman and Helpman (1991), suppose that each quality application is monopolized, perhaps due to a patent or a trade secret. There is a Bertrand competition between qualities within a ladder. Non-leading-edge qualities price at marginal cost; to match this price per efficiency unit, the leading edge quality charges $p_i^j = \lambda$ for $j > 1$ and $p_i^1 = \lambda/\phi$.

Given demand x_i on ladder i, profit flows for the leader are $d_i(p_i - 1)$, where under Assumption 1, $d_i = 1/p_i$. For instance, if $\phi = 1$, $d_i = 1/p_i = 1/\lambda$ and profits are $\frac{\lambda-1}{\lambda} \equiv \pi$. Although our assumption of no dependence of profits on the number of applications is an extreme one, it is simply an expedient abstraction to the idea that there is less business stealing on the creation of new applications than on the creation of quality improvements.

Pricing in a Closed Standard If a single firm controls pricing on all applications, it faces only the limit price from consumers' direct accomplishment of the task. Therefore it can set price at most $p_i^j = \lambda^j/\phi$. Note that with the Assumption 1, since unit elastic demand implies an infinite monopoly price, the limit price always binds.

2.4 Pricing Hardware

When hardware is sold, it is by a monopolist. He charges all the expected surplus from the applications in the standard. He cannot commit to a future stream of innovations or prices for applications. Hardware is produced at zero marginal cost.

2.5 Innovation

Innovation comes through research. Research can be done on either developing new applications (i.e. applications for functions with no applications yet) or on improvements to existing applications (i.e. creating applications of higher quality for functions with existing applications). In both cases, research takes place continuously and innovations arrive according to a Poisson process. The arrival rate is proportional to the amount of research intensity, denoted x_e for existing applications and x_f for new (frontier) applications.

We assume that research intensity comes from one input, researchers, and that the pool of this input is heterogenous in their skills. A researcher of type θ can provide one unit of research intensity, at a cost flow of θ for an existing application, and a cost flow of $\theta + \eta$ for a frontier application. The inclusion of $\eta > 0$ means that new applications are more costly to research. Researchers' types are distributed on $[\theta_l, \infty)$ according to the cumulative distribution function $F(\theta)$, with $\theta_l > 0$. We normalize the outside option of researchers to zero. In the open-standard case researchers may enter freely into research at any instant. In the closed-standard case the researchers are hired by the firm controlling all applications at a uniform wage.[3] In order to keep the model from trivially having no innovation, we assume that there are types θ smaller than the value of profits of the leading application, or $F(\pi/r - \eta) > 0$.

The critical feature of our model that ties different applications together is the fact that they draw researchers from a common pool of scarce talent. As a result, innovative effort on one application has an impact on the marginal cost of innovation for all of the applications using the common factor. This equilibrium effect is what gives rise to all of the results about the dynamics across different applications that we develop below, and the differences that develop between open and closed standards.

In the open-standard regime, when a researcher finds an innovation he forms a firm and markets it (becoming one of the producers in the open standard). In the closed-standard regime, workers are hired by the monopolist at a fixed research wage w. Researchers (and firms hiring them) maximize the expected sum of discounted profits/wages net of research costs and use a common discount rate r.

3 Innovation Dynamics in Open Standard

In this section we consider the equilibrium of the economy when the standard is open. We will keep Assumption 1 and assume that applications are stand alone, so $h = 0$ (i.e. there is no hardware to buy). We will also disregard the assumption that $\bar{N}$ is bounded and instead will establish an endogenous bound on the number of functions for which under the open standard applications will ever be developed.

[3] In the open-standard case, WLOG the researchers can also be hired by a continuum of firms that have a free entry to the innovation process and compete for the researchers.

3.1 Analytic Results

Given the profits π from the Bertrand pricing game, the innovation takes place competitively. We will characterize a recursive equilibrium, in which the decision rules of the researchers depend only on the current state of the industry. Since obtaining an improvement over an existing application yields the same profit flows regardless of the identity of the function or the current quality level of the application, in such an equilibrium the strategies of the researchers (and expected profits of the producers) depend only on the current number of functions with existing applications, which we denote by N.

For simplicity we take $\phi = 1$, which allows us to reduce the payoff flows for the current highest-quality application for any function to be $\pi = \frac{\lambda-1}{\lambda}$. Once we deliver the results we will discuss the case of $\phi < 1$.

Assumption 2: $\phi = 1$.

Denote total research by $x(N) \equiv x_e(N) + x_f(N)$. Since the benefits of research are independent of type, research follows a cutoff rule: if type θ does research on one of the applications, all types $\theta' < \theta$ do as well (although they can do it on some other application). The cutoff $\bar{\theta}(N)$ must satisfy:

$$F(\bar{\theta}(N)) = x(N) \tag{1}$$

The cutoff $\bar{\theta}(N)$, as well as the allocation across the two activities is determined by free-entry conditions: the cutoff type must be indifferent between researching any of the existing applications (unless $x_e(N) < 0$), researching a frontier application (unless $x_f(N) < 0$) and opting out of research.

Define $c(x) = F^{-1}(x)$, the cost of the marginal researcher in the existing application search, so that $\bar{\theta}(N) = c(x(N))$. It is increasing from (1). Let $\rho(N) = \frac{x_e(N)}{x(N)}$ be the probability that the new innovation is for an existing application. Let the random variable $\tau(N)$ be the arrival time of an innovation given the aggregate research intensity $x(N)$. The expected discount factor can be calculated using the Poisson distribution as:

$$\delta(N) \equiv E\left[e^{-r\tau(N)}|x(N)\right] = \frac{x(N)}{x(N)+r}$$

Note that the expected time it takes for an application to be replaced by an improved one depends on the search intensities, hence the expected discount factor depends on the current state of the industry, summarized by N.

To calculate the expected benefit of doing research, let first $V(N)$ be the value of an incumbent firm just as a new innovation arrives, but without knowing what application it is for. It is described recursively by

$$\begin{aligned} V(N) &= \rho(N)\left(1-\frac{1}{N}\right)\left((1-\delta(N))\pi+\delta(N)V(N)\right) \\ &\quad +(1-\rho(N))\left((1-\delta(N+1))\pi+\delta(N+1)V(N+1)\right) \end{aligned} \tag{2}$$

This allows us to define the flow of expected benefit to the two research activities:

$$\begin{aligned} V_e(N) &= (1-\delta(N))\pi+\delta(N)V(N) \\ V_f(N) &= (1-\delta(N+1))\pi+\delta(N+1)V(N+1) \end{aligned} \tag{3}$$

$V_e(N)/r$ and $V_f(N)/r$ are the expected total profits of a researcher conditional on achieving one of the corresponding innovations when the current state is N. By the properties of Poisson distribution they also represent the flow of expected profits from innovation. Therefore, the free-entry conditions for the researchers are:

$$\begin{aligned} V_e(N)/r &\leq \bar{\theta}(N) \\ V_f(N)/r &\leq \bar{\theta}(N)+\eta \end{aligned} \tag{4}$$

with equality whenever the corresponding task is undertaken by a positive mass of researchers.

Formally, the recursive equilibrium requires that agents optimize according to equation (4) given (2) and (3), and these individual decisions agree with aggregate variables in (1).

Next we summarize the results of this section in a proposition. We follow the proposition with a discussion, including a series of lemmas, that establish the result, as well as establishing a sufficient condition for uniqueness of the equilibrium.

Proposition 1 *There exists some N^* such that $x_f(N) > 0$ and $x_e(N) > 0$ for $N < N^*$. Further, $x_f(N^*) = 0$ and $x_e(N^*) > 0$. Finally, $x(N)$ is increasing in N.*

Despite the fact that the per period profits are the same for new applications and improvements to existing applications (π), competitive innovators may be willing to pay the additional cost to develop a new application. The reason is that, due to the increasing cost of researchers, research intensity on existing applications rises less than proportionally to the number of existing applications. As a result, the amount of research on improvements, per application, is declining in the number of applications; this implies that the value of having a marketable application rises as the total number of applications increases. If that value is rising fast enough from N to $N+1$, the extra value from developing a new application makes the extra cost worth paying. Eventually this increased value gets small and may be insufficient to draw research in new applications. We then reach a steady state number of applications, N^*, at which point there is only research in existing applications.

On the other hand, anytime there is research on new applications, there must also be research on improving existing applications. If there were only research on new applications for a given N, then developing an improvement would make more profits than a new application: it would earn profits until the next new application, at which point the continuation value would be as much as the new application would have made. This is of course impossible since improvements are less expensive. We get the following picture of innovation. For N below the steady state, there are both types of research, with total research intensity increasing due to the rising value of a leading edge application.

In order to see the results in Proposition 1, suppose first that there are some states N and $N+1$ in which both research dimensions are active in equilibrium. Using the free entry condition we can now characterize the equilibrium aggregate research intensity for times when both activities are undertaken.

Note that (3) implies $V_e(N+1) = V_f(N)$. Combining it with the free entry conditions we get:

$$c(x(N+1)) = c(x(N)) + \eta \tag{5}$$

For example, if $F(\theta) = (\theta - \theta_l)/a$, then $c(x) = ax + \theta_l$ and:

$$x(N) = x(N+1) - \eta/a$$

This condition allows us to show that aggregate research is increasing in N :

Lemma 1 *As long as both research tasks are active, aggregate research effort $x(N)$, and value $V(N)$ are increasing in N.*

Proof. Monotonicity of $x(N)$ follows directly from (5) and monotonicity of $c(x)$.

Regarding $V(N)$, from the free-entry conditions we have

$$\begin{aligned} V_e(N+1)/r - V_e(N)/r &= \eta \\ &\Downarrow \\ (\delta(N) - \delta(N+1))\pi + \delta(N+1)V(N+1) - \delta(N)V(N) &= r\eta \\ &\Downarrow \\ \underbrace{(\delta(N) - \delta(N+1))}_{<0}\underbrace{(\pi - V(N))}_{>0} + \delta(N+1)(V(N+1) - V(N)) &= r\eta \end{aligned}$$

where $(\delta(N) - \delta(N+1)) < 0$ because we have proven that $x(N)$ is increasing. As the first element on the LHS is negative, we must have $V(N+1) > V(N)$ for the equality to hold. ■

Next, we establish the existence of a steady-state and that before stead-state both research tasks are indeed active:

Lemma 2 *For any N, $x(N) > 0$. For any $N > 0$ such that $x_f(N) > 0$, it must be that $x_e(N) > 0$. Finally, there exists N^* such that for all $N \geq N^*$ $x_f(N) = 0$.*

Proof. We start with the second claim. Suppose $x_e(N) = 0$ and $x_f(N) > 0$. Then, combining (2) and (3) we get $V_e(N) = V_f(N)$, which contradicts the free entry conditions (4).

Now, suppose that there exists an N such that $x(N) = 0$. Then $\delta(N) = 0$ and $V_e(N) = \pi$. Since there are researchers with $\theta < \pi/r$ (which we assumed to get any innovation in equilibrium), that violates free entry condition.

Finally, suppose that for all N, $x_e(N)$ and $x_f(N)$ are positive. From the analysis before we know that this would imply $V_e(N+1) = V_e(N) + r\eta$ for all N. But that is not possible as $V_e(N) \in (0, \pi)$. ■

We will assume that this $N^* < \bar{N}$ (the total number of feasible functions). In a steady-state N^*, the value and research intensity can be easily determined as they satisfy the Bellman equation and the free-entry condition:

$$V(N^*) = \left(1 - \frac{1}{N^*}\right)((1 - \delta(N^*))\pi + \delta(N^*)V(N^*)) \tag{6}$$

$$\underbrace{(1-\delta(N^*))\,\pi+\delta(N^*)\,V(N^*)}_{V_e(N^*)}=rc\,(x_e\,(N^*)) \tag{7}$$

Call the solution to these two equations, for arbitrary N, $\hat{V}(N)$ and $\hat{x}_e(N)$, with the associated $\hat{\delta}(N)$.

Lemma 3 *$\hat{V}(N)$ is increasing in N; $\hat{x}_e(N)$ is increasing in N.*

Proof. Pick any N and consider $N'=N+1$. $\hat{V}(N)$ increasing: suppose not. Then $\hat{x}_e(N)$ would weakly decrease to satisfy (7) But that implies that $\hat{\delta}\,(N)$ would weakly decrease; since $V(N)<\pi$ (the maximum flow payoff), (6).implies that $V(N^*)$ is increasing.

Finally, suppose $\hat{x}_e(N)$ is weakly decreasing. That would require that $\delta\,(N^*)$ is weakly decreasing and would violate (7) since we have already established that $\hat{V}(N)$ is increasing. ■

In order for N^* to be a steady state, it must be the case that it is not profitable to search for a new application at N^*, even if no further effort on new applications were researched. In other words, at the steady state N^*, the following inequality holds:

$$(1-\widehat{\delta}\,(N+1))\pi+\widehat{\delta}\,(N+1)\,\hat{V}(N+1)\leq r\,(c\,(\hat{x}_e\,(N))+\eta) \tag{8}$$

where

$$\widehat{\delta}\,(N+1)=\frac{\hat{x}_e(N+1)}{r+\hat{x}_e(N+1)}$$

is the expected discount factor if at state $N+1$ the research is $x_e=\hat{x}_e(N+1)$ and $x_f=0$. Condition (8) can be simplified to

$$c\,(\hat{x}_e\,(N+1))-c\,(\hat{x}_e\,(N))\leq\eta$$

If $c\,(\hat{x}_e\,(N))$ is concave, there clearly exists exactly one "crossing point" which provides a sufficient condition for uniqueness of the steady-state and equilibrium:

Lemma 4 *Suppose $c\,(\hat{x}_e\,(N))$ is concave in N. Then the steady state N^* is unique.*

It can be verified directly that $c\,(\hat{x}_e\,(N))$ is in fact concave for many distribution functions F, for example a linear one (see also the numerical example in the next section). The intuition why we should expect it to hold is as follows: suppose first that $\delta(N)$ is constant. Then the solution to (6) $V(N)$ is concave because $\left(1-\frac{1}{N}\right)$ is concave. Now, for (7) to hold, $c(\hat{x}_e(N))$ and $\delta(N)$ have to increase. This adjustment has to be larger the more $V(N)$ increases, so it is smaller for larger N. Therefore, for a lot of shapes of $c(x)$ we would obtain that the $c(\hat{x}_e(N))$ would be concave.

Once we find N^* and $x_e(N^*)$, we can solve the open-standard model by working from the eventual steady state. In particular, iterating on (5) we get:

Lemma 5 *Aggregate research effort $x(N)$ is increasing in N according to $c(x(N+1)) = c(x(N)) + \eta$.*

Given $x(N)$ for all $N \leq N^*$ and $V(N^*)$, we can use equation (4) to calculate $V(N)$. Finally, to compute the individual values $x_e(N)$ and $x_f(N)$, given $x(N)$ and $V(N)$, we can use equation (2). Note that this construction is unique for a given N^*, so that $c(\hat{x}_e(N))$ concave delivers a unique equilibrium.

3.2 Numerical Example

We assume that F is linear, $F(\theta) = (\theta - \theta_l)/a$, and so $c(x) = ax + \theta_l$.

We consider the following parameters: $r = 5\%$ (annual interest rate) $a = 0.01$, $\theta_l = 0.2$, $\eta = \theta_l/25$, $\lambda = 1.5$ (so quality increases by 50%). Then $N^* = 9$. In the steady-state $x(N^*) \approx 9.66$, which is also the average number of improvements per year. The research intensities are shown in the figure as a function of N: the top line is total investment, the decreasing line is the investment in new applications, and the third line is investment in existing applications.

Note that the difference between the extra cost η of a new application is at most four percent of the cost of researching an improvement (for $\theta = \theta_l$), but yet differences between intensities research intensities new and existing applications are large.

The expected time to the next frontier application is drawn in the next figure:

It takes less than 5 months for the first application to be invented. The time till next frontier application stays below one year only until 3 ladders

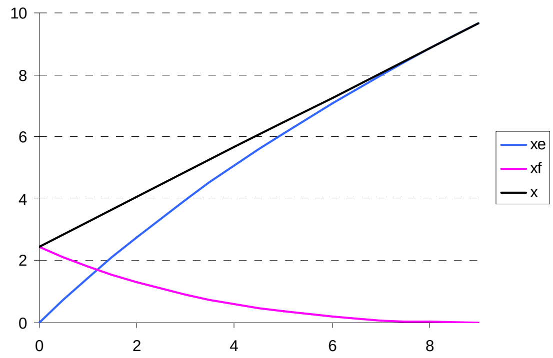

Figure 1: Innovation in an Open Standard

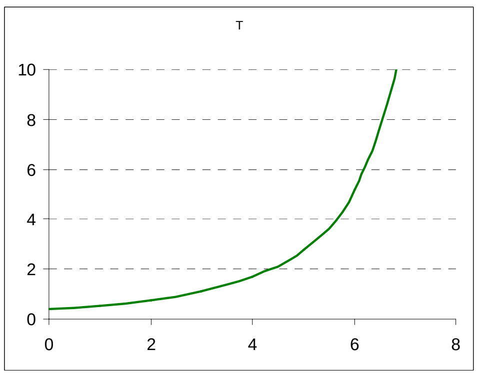

Figure 2: Time to Next Application

exist and finally reaches over 49 years (not shown on the figure) for the move from $N = 8$ to N^*.Note that the difference between the extra cost η of a new application is at most 4% of the cost of researching an improvement (for $\theta = \theta_l$), but yet differences between intensities research intensities new and existing applications are great.

4 Innovation Dynamics in a Closed Standard

4.1 Analytic Results

As we described above, with a closed standard pricing of the applications will be different as the monopolist owns all applications. The optimal strategy of the monopolist will depend now not only on N but on the whole state of the applications, i.e. not only on how many functions have application but also on the quality of existing applications.

Recall that quality improves at rungs of $\lambda > 1$ and let j_i denote the number of rungs that the best application has over the default level $\phi = 1$, so that $q_i^j = \lambda^j$. If there is no application for a given function, then let $j_i = 0$. Let $Q = \{j_1, ..., j_{\overline{N}}\}$ denote the state of the applications. The monopolist chooses research activities $x_e(Q), x_f(Q)$ to maximize expected profits. As we argued before, the optimal price of the best application for function i is λ^{j_i}. At that price demand is $x_i = 1/p_i = 1/\lambda^{j_i}$ and hence the monopolist obtains a flow of profits

$$\pi_j = 1 - \frac{1}{\lambda^j}$$

which is increasing and concave in j. The profit gain from a unit increase in quality is

$$\Delta_j = \pi_{j+1} - \pi_j = \frac{1}{\lambda^j}\left(1 - \frac{1}{\lambda}\right)$$

which is decreasing in j and $\Delta_j \to 0$. For now, assume $\Delta_0 - \eta > \Delta_1$, so that the monopolist prefers to invest in frontier applications before improving existing ones. We will later discuss how the optimal innovation strategy differs if this condition is not satisfied.

We begin by summarizing the results as a proposition.

Proposition 2 *If $j_i = 0$ for any i, then $x_f > 0$ and $x_e = 0$. Existing applications are improved only if $j_i = \min Q$. Total innovation x is strictly* *decreasing in the number of innovations achieved. There exists a level j^* such that all research stops once all applications reach j^*.*

We will prove this proposition through a sequence of lemmas and at the same time we will provide a more detailed characterization. We first argue that the optimal strategy is to continue research only up to a level j^* and to put research activity only at the current lowest-level application, i.e. one that maximizes Δ_j. As a result, research is done in layers: at any time the monopolist works on the lowest-level applications until he brings all of them to the next level. This continues until he reaches level j^* with all of them.

Lemma 6 *Optimal research strategy in the closed standard satisfies:*
a) In the long run research stops once all applications reach a level j^* which is the smallest integer s.t. $\Delta_{j^*} < r\theta_l$
b) In the long run all functions have applications.
c) In any time the research is done on a product with the highest Δ_j (i.e. lowest j_i).

Proof. See appendix. ■

This partial characterization allows us to reduce the state space to relevant points with applications on two levels only. Abusing slightly notation, let Q now denote a pair (j, k) where j is the lowest level of applications across i and k is the number of functions on this level. We will use a convention that $Q + 1 \equiv (j, k - 1)$ and $(j, 0) \equiv \left(j + 1, \overline{N}\right)$. Finally, let $C(x) = xc(x)$ denote the total cost of hiring x mass of researchers, each at wage $c(x)$, where $c(x) = F^{-1}(x)$ if $x_f = 0$ and $c(x) = F^{-1}(x) + \eta$ if $x_f > 0$.

We can prove the following monotonicity result.

Lemma 7 *$x(Q)$ is strictly decreasing.*

Proof. See appendix. ■

The intuition is that the rewards to research decrease for two reasons. First, the immediate rewards, Δ_j, are weakly decreasing. Second, as there is only a finite number of rewards, the increase in continuation-payoffs are strictly decreasing: finding the first innovation brings closer the profits from all the subsequent innovations, an effect missing for the last innovation.

We will now determine recursively the levels $x(Q)$ of optimal research. The value function of the monopolist is defined recursively through the optimization problem:

$$V_M(Q) = \max_x \frac{r}{x+r}(\pi(Q) - C(x)) + \frac{x}{x+r} V_M(Q+1) \tag{9}$$

The first order condition is

$$\frac{V_M(Q+1) - (\pi(Q) - C(x))}{(x+r)} = C'(x) \tag{10}$$

For example, if $F(\theta) = (\theta - \theta_l)/a$, then $C(x) = ax^2 + \theta_l x$, the FOC becomes

$$V_M(Q+1) - \pi(Q) = ax^2 + 2rax + r\theta_l$$

and the optimal choice is

$$x = \sqrt{r^2 + \frac{V_M(Q+1) - \pi(Q) - r\theta_l}{a}} - r$$

which is positive as long as $V_M(Q+1) - \pi(Q) > r\theta_l$, i.e. the gain is worth hiring some researchers.[4]

From lemma 6 we can find j^* as the closest integer higher than the solution to $\Delta_j = r\theta_l$. This can be simplified rewritten as $j = \ln((\lambda - 1)/(r\lambda\theta_l))/\ln\lambda$. Given this j^*, let $Q^* = (j^*, \overline{N})$. The value in this steady-state is:

$$V_M(Q^*) = \overline{N}\pi_{j^*}$$

That allows us to solve recursively for $x(Q)$ and $V_M(Q)$ by iterating on (10) and (9) and finishes the description of the equilibrium.

Remark 1 *We have assumed that $\Delta_0 - \eta > \Delta_1$. If that does not hold, the optimal research policy varies only slightly. Let j' be the smallest j such that $\Delta_0 - \eta > \Delta_1$. If $j' > 1$, then the optimal strategy is to develop first the j' applications for a given function before developing a frontier application for a new function.*

Remark 2 *Given the symmetry across applications, all strategies with a given level of aggregate $x(Q)$ that put all the research activity at the lowest-quality applications, but differ in the division of $x(Q)$ among them are payoff equivalent and lead to the same evolution of applications.*

[4] For $x_f > 0$ these formulas have to be modified to replace θ_l with $\theta_l + \eta$.

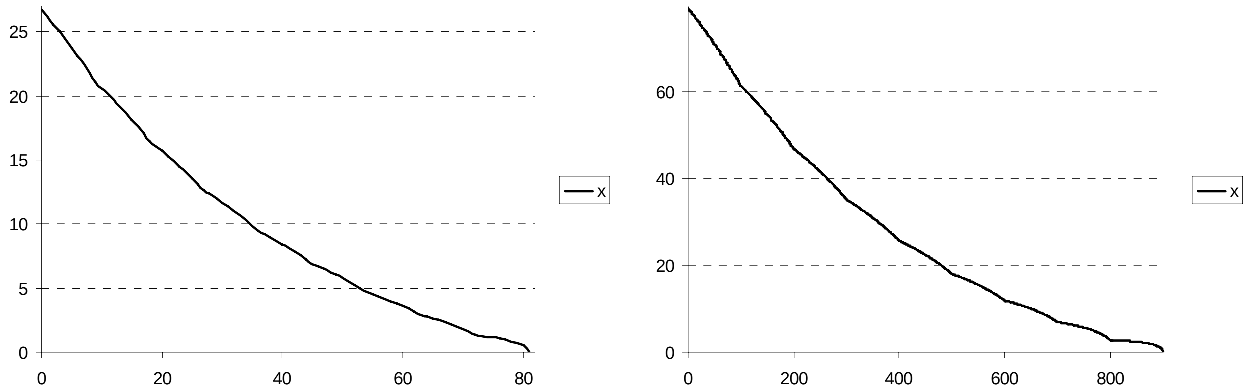

Figure 3: Innovation in a Closed Standard

4.2 Numerical Example

We keep the parameters from the open-standard example: $r = 5\%$ (annual interest rate), $a = 0.01$, $\theta_l = 0.2$, $\eta = \theta_l/25$, $\lambda = 1.5$. The new parameter we need to specify is $\overline{N}$. We will compare two values: $\overline{N} = 9$ and $\overline{N} = 100$. Given these parameter values, $j^* = 8$ so the monopolist will research up to 8 applications per ladder. The figures show the dynamics of innovation as a function of $J = \sum j_i$ for $\overline{N} = 9$ on the left and for $\overline{N} = 100$ on the right.

The intensity of innovation is huge in the beginning (as compared with the open standard): for $\overline{N} = 9$ it takes on average less than a month in between the first 30 innovations (in the open standard in the steady-state the innovations come on average 1.2 months apart and the speed of innovation is much slower in the beginning - the first innovation arrives an order of magnitude faster in the closed standard). Note also that for $\overline{N} = 100$. the innovation starts almost 3 times as large!

Note further that one can interpret these differences as differences between innovation by large firms (the monopolized closed standard) and small firms (the open standard). However, in this case, the model suggests that measurement of these differences is difficult. First, whether you observe more innovation by large or small firms depends on the point in the product cycle; in the beginning, it appears that large firms do more innovating, but later on the pattern reverses. Moreover, our model suggests that different market structures lead to different types of innovation, so simply measuring a one dimensional innovation variable for each is not an apples-to-apples comparison.

5 Software Applications with Open and Closed Standards

The analysis so far focused on *stand-alone applications* and took the market structure as given. We now use the model to consider *software applications* and to better understand the tradeoffs in choosing a closed or open standard regime. In this section we consider the choice by the owner of hardware and in the next section we discuss the welfare consequences.

In particular, we allow for the possibility that hardware can be sold by the monopolist. The hardware in our model has no use in itself, but is necessary to use any application. We assumed that hardware is infinitely durable and bought in the beginning of the game, in anticipation of future applications (that assumption is unrealistic, as hardware always comes with some applications immediately available, but that simplification is not affecting the economics we describe).

The consumers will be willing to pay up to their expected surplus for the hardware, which equals their total utility minus the utility that they could obtain from directly performing the function. In the closed standard, the monopolist always follows limit pricing of applications to the direct option. Therefore, there is zero expected surplus from the hardware, all profits come from sales of the applications. In other words, if the software market is monopolized, the hardware can be sold by competitive firms.

5.1 Hardware Pricing in Open Standard

If the standard is open, then prices are $p_i = 1/\lambda$, quantity is $1/\lambda$ and consumer surplus (over the outside option) from good i with quality $q_i^j = \lambda^j$ is

$$CS(j) = (j-1)\ln\lambda$$

Therefore the hardware seller can extract up to the expected present value of $CS(j)$ summed over all applications.

Let $J = \sum j_i$ be the number of quality improvements summed over all applications. Then the current flow of surplus given J is

$$CS\left(J,N\right)=\left(J-N\right)\ln\left(\lambda\right)$$

In the steady-state the expected total consumer surplus for a current J is:

$$\begin{aligned} U\left(J,N^{*}\right) &= \left(1-\delta^{k}\left(N^{*}\right)\right)\sum_{k=0}^{\infty}\delta^{k}\left(N^{*}\right)CS\left(J+k,N^{*}\right) \\ &= \ln\left(\lambda\right)\left(J-N^{*}+\frac{\delta\left(N^{*}\right)}{1-\delta\left(N^{*}\right)}\right) \end{aligned}$$

In a state $N<N^{*}$, the total surplus is:

$$U\left(J,N\right)=\left(1-\delta\right)CS\left(J,N\right)+\delta\left(\rho U\left(J+1,N\right)+\left(1-\rho\right)U\left(J+1,N+1\right)\right)$$

(note that $\delta=\delta\left(N\right)$ and $\rho=\rho\left(N\right)$). Since research intensity for the open standard depends only on N, we can expand this as

$$\begin{aligned} U\left(J,N\right) &= \left(1-\delta\right)\left(CS\left(J,N\right)+\delta\rho CS\left(J+1,N\right)+\delta^{2}\rho^{2}CS\left(J+2,N\right)...\right) \\ &\quad +\delta\left(1-\rho\right)U\left(J+1,N+1\right)+\delta^{2}\rho\left(1-\rho\right)U\left(J+2,N+1\right)+... \end{aligned}$$

and so

$$U\left(J,N\right)=\ln\left(\lambda\right)\left(1-\delta\right)\sum_{k=0}^{\infty}\left(\delta\rho\right)^{k}\left(J-N+k\right)+\left(1-\rho\right)\delta\sum_{k=1}^{\infty}\left(\delta\rho\right)^{k-1}U\left(J+k,N+1\right)$$

We guess that

$$U\left(J,N\right)=\ln\left(\lambda\right)\left(J-N+a_{N}\right)$$

and plug this back in for $U(J+k,N+1)$ to get :

$$U\left(J,N\right)=\ln\left(\lambda\right)\left(J+\frac{a\delta\left(1-\rho\right)+\delta\rho}{1-\delta\rho}\right)$$

This confirms the guess; we get the following recursive equation for a :

$$a_{N}=\frac{a_{N+1}\delta\left(N\right)\left(1-\rho\left(N\right)\right)+\delta\left(N\right)\rho\left(N\right)}{1-\delta\left(N\right)\rho\left(N\right)} \tag{11}$$

Since we have found before $a_{N^*}=\frac{\delta(N^*)}{1-\delta(N^*)}$, we can use (11) to calculate a_N for all $N < N^*$.Then we can calculate

$$U(0,0) = \ln(\lambda) a_0$$

and the optimal hardware price is

$$\ln(\lambda) a_0 / r$$

The analysis so far allows us to calculate and compare the profits in the two regimes numerically, but direct analytic comparison is difficult. Therefore in the next two sections we develop further results to highlight the role scope of feasible innovation plays in the choice of the hardware monopolist. For simplicity we will focus on the linear case $F(\theta) = (\theta - \theta_l)/a$.

5.2 Large $\bar{N}$

First we develop a result for the case where the future is relatively important (low r) $\bar{N}$ is high and η is small. In this case, it is better to extract surplus through monopolizing applications, since there is a great deal of benefit in internalizing the benefits of making the scope as large as possible. Of course one must be sufficiently patient to make this plan profitable.

We first establish two lemmas, which serve to compute the payoff from a closed standard and bound above the payoff for an open standard

Lemma 8 *In the linear case, $F(\theta) = (\theta - \theta_l)/a$, as $\overline{N} \to \infty$, the value from the closed standard converges to $\frac{1}{4}\frac{(\pi - r(\theta_l+\eta))^2}{ar^2}$ (in flow terms)*

Proof. See appendix. ■

Lemma 9 *The hardware price for the open standard is bounded above by $\ln(\lambda)\frac{\pi - r\theta_l}{r}$ (in flow terms).*

Proof. See Appendix. ■

Proposition 3 *If $r < \frac{\pi}{\theta_l + 4a\ln\lambda}$ then there exists $\eta^* > 0$ such that for all $\eta \leq \eta^*$ if $\overline{N}$ is sufficiently large, the monopolist prefers to have a closed standard rather than an open one.*

Proof. From the above calculations for open and closed standards, we will calculate a ratio

$$\frac{1}{4}\frac{\left(\pi - r\left(\theta_l + \eta\right)\right)^2}{ar^2} / \left(\ln\left(\lambda\right)\frac{\pi - r\theta_l}{r}\right)$$

which is continuously decreasing in η. At $\eta = 0$ the ratio is:

$$\frac{1}{4ar\ln\lambda}\left(\pi - r\theta_l\right)$$

This expression is positive and decreasing in r. Solving for it to be equal 1 yields the bound. Given r less than this bound we can find an $\eta^* > 0$ such that ratio is still less than 1. Then, for any $\eta \leq \eta^*$ we can find $\overline{N}$ large enough that the payoffs under the closed are arbitrarily close to the calculated bound (for this and all larger $\overline{N}$). ■

Remark 3 *For a given η we can calculate a bound on the interest rates:*

$$r^* = \pi\frac{\theta_l + \eta + 2a\ln\lambda - 2\sqrt{\eta a\ln\lambda + a^2\ln^2\lambda}}{\left(\theta_l + \eta\right)^2 + 4a\theta_l\ln\lambda}$$

Then, if $r < r^*$ we can find $\overline{N}^*$ such that for all $\overline{N} \geq \overline{N}^*$ the closed standard is more profitable.
This bound converges to the one in proposition as $\eta \to 0$. For example, if $\theta_l = a = 0.1$, $\eta = \theta_l/10$ and $\lambda = 1.5$ (so that $\pi = \frac{1}{3}$) then the bound is $r^* = 118\%$. Any discount rate smaller than that makes the monopolist prefer a closed standard for sufficiently high $\overline{N}$.

5.3 Small $\bar{N}$

We will now argue that if $\overline{N}$ is small and r is small, then the open standard choice will be more profitable (we keep $F(\theta)$ linear). We will focus on the case $\overline{N} = 1$ as it illustrates the intuition best.

For the open standard, once the first application is developed, the free-entry condition is:

$$\left(1 - \delta\right)\pi/r = c\left(x\right)$$

which in the linear case yields that steady-state x^* solves:

$$\frac{\pi}{r + x^*} = ax^* + \theta_l \tag{12}$$

Direct calculation shows x^* is decreasing in r.

Following the calculations in Section 5.1, consumer surplus for $j > 0$ is $CS(j) = (j-1)\ln\lambda$, and $CS(0) = 0$. Let $\delta = \frac{x^*}{x^*+r}$. The expected surplus after the first application is developed is:

$$\begin{aligned} U(1,1) &= (1-\delta) CS(1,1) + \delta(1-\delta) CS(2,1) + \delta^2 (1-\delta) CS(3,1) ... \\ &= (1-\delta)\ln(\lambda) \sum_{k=0}^{\infty} k\delta^k = \ln(\lambda) \frac{x^*}{r} \end{aligned}$$

Using condition (12) we can find that $\lim_{r\downarrow 0} x^* > 0$. so $U(1,1)$ is on the order $O\left(\frac{1}{r}\right)$ for small r.

We can now bound $U(0,0) = h$. The free entry condition at $N = 0$ is

$$(1-\delta)\pi + \delta \underbrace{V(1)}_{=0} = r(ax_f + \theta_l + \eta)$$

which yields:

$$x_f = \left(\frac{\pi}{r+x^*} - (\theta_l + \eta)\right)/a \tag{13}$$

Remark 4 *We have assumed that η is small enough so that $x_f > 0$. This assumption corresponds in this case to:*

$$\eta < \frac{\pi}{r+x^*} - \theta_l$$

the above inequality holds for small r if

$$\pi > \eta \frac{\theta_l + \eta}{a}$$

If η is too large compared with π, then the open standard will fail to develop even the first application. Then the closed standard is necessary to overcome the public-good aspect of the higher cost of developing a new application and clearly the closed standard will be more profitable.
On the other hand, if the hardware comes with the first application, then this additional restriction is not necessary.

Assuming $\pi > \eta\frac{\theta_l+\eta}{a}$, x_f converges to a strictly positive number as $r \to 0$. Therefore the price of hardware (in payoff flows) is

$$h = U(0) = \frac{x_f}{x_f + r} U(1) \to O\left(\frac{1}{r}\right)$$

In the closed standard the profit flow is at most $\pi_{j\to\infty} = 1$. Therefore:

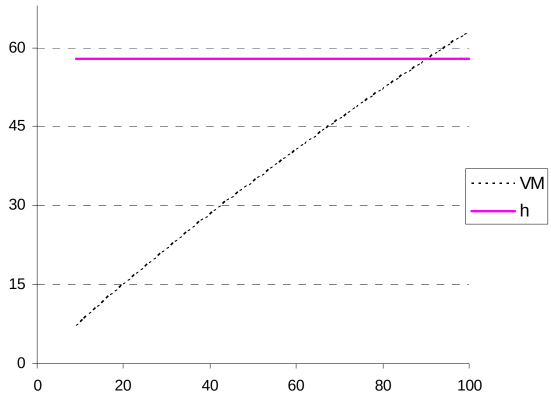

Figure 4: Surplus Extraction

Proposition 4 *If $\pi > \eta \frac{\theta_l + \eta}{a}$, then there exists $r^* > 0$ such that if $r < r^*$ and $\overline{N}$ is sufficiently small, the open standard is more profitable to the hardware owner than the closed standard.*

Proof. Follows directly from the calculations above: for $\overline{N} = 1$ and small r, the total profits from closed standard are at most $1/r$, while the price of hardware in the open standard is on the order $O\left(1/r^2\right)$. ■

5.4 Numerical example (effects of $\overline{N}$)

Return to our two numerical examples. As $N^* = 9$, then as long as $\overline{N} \geq 9$, the price of hardware is independent of $\overline{N}$. Contrary, the profits in the closed standard depend crucially on $\overline{N}$. In the figure below we compare the profits from hardware in the open standard (the horizontal line) with the profits from software in the closed standard regime (the increasing curve). As the discussion in this section points out, for small $\overline{N}$ the profits are higher from an open standard and when $\overline{N}$ is large they are higher from the closed standard case (in this example the cutoff is $\overline{N} = 91$).

6 Social Welfare

In the previous section we calculated the surplus the hardware owner can expect from different market structures. For policy questions it is useful to calculate social welfare under various production structures.[5] The only portion of surplus that we have not calculated so far is the surplus of researchers.

6.1 Surplus of the researchers

The supply function of the researchers is $c(x)$. Therefore the flow of producer surplus is simply:

$$PS(x) = \int_0^x (c(x) - c(u))\, du$$

In the linear case, $c(x) = ax + \theta_l$, we have

$$PS(x) = \frac{1}{2}ax^2$$

(Note that it does not depend on θ_l, so the surplus calculation when the researchers work on existing applications or frontier ones is the same).

In the closed standard case, the total expected discounted researcher surplus is:

$$V_R(Q) = (1 - \delta(Q))\, PS(x(Q)) + \delta(Q)\, V_R(Q+1)$$

and we have $V_R(Q^*) = 0$. So we can find the $V_R(0)$ iteratively.

In the open standard case, given an N, the total expected discounted researcher surplus is:

$$V_R(N) = (1 - \delta(N))\, PS(x(N)) + \delta(N)\,(\rho V_R(N) + (1-\rho)\, V_R(N+1))$$

Hence

$$V_R(N) = ((1 - \delta(N))\, PS(x(N)) + \delta(N)(1-\rho)\, V_R(N+1)) / (1 - \delta(N)\rho)$$

Also, at N^* the surplus flow becomes constant:

$$V_R(N^*) = PS(x(N^*))$$

[5] Note that we do not study a planner with complete control, as he would trivially do only quality improvements, and sell at marginal cost, given our preference structure.

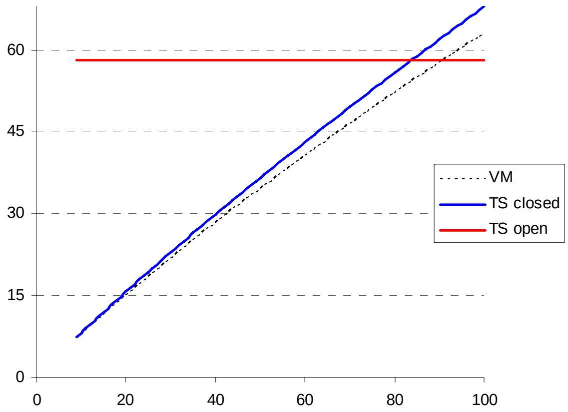

Figure 5: Social Welfare

so again it is easy to find the time-zero expected discounted producer surplus of the researchers.

Note that in our model the consumers are left with no surplus due to their homogeneity and unit-elastic demand. As a result, the difference between total surplus and the surplus of the owner of hardware is small unless research is done on a very large scale. Therefore the inefficiency of the decision of the monopolist will be in general small (see the following example). A better model to study policy questions would expand the model and introduce heterogenous consumers to allow for consumer surplus in either regime.

6.2 Numerical Example

After adding the welfare of researchers to obtain total welfare calculation for the numerical examples we have presented before, the picture as a function of $\overline{N}$ looks like this:

For comparison we have also plotted the value in the closed standard (the difference between the total surplus in the open standard and the hardware price is negligible as the research intensity is small in the open standard).

As we see, for a wide range on $\overline{N}$ the hardware owner's preferences agree with maximization of total surplus. Only in the range $\overline{N} \in \{84, ...90\}$ the hardware monopolist would choose to open the standard while the total surplus would be maximized by keeping it closed. Still, the potential social inefficiency is small in this example.

7 Subsidizing innovation

The equilibria we have characterized above do not achieve social first-best along many directions: the prices of the final products are above marginal cost of production, too little innovation takes place and in the closed standard case, the monopsony power in the market for researchers creates inefficiency there. Hence there are many ways a government can intervene to improve efficiency. We now focus on one possibility: subsidizing innovation in an open standard and ask how the social returns compare between subsidizing frontier or existing applications.

To keep things simple, we focus on one-time unexpected subsidy for the marginal researcher in the steady state of the open standard to avoid crowding-out of private innovation caused by the anticipated future government subsidy. In equilibrium the steady state innovation level x^* is socially inefficient, in the sense that the total surplus would increase if additional researchers joined the innovation effort. The reason is that the private return is equal to $\frac{r}{x^*+r}\pi$ while the social return is equal to CS_Δ/r: the second number is higher because the social benefits accrue forever, while the private returns occur only until the firm gets replaced by an improvement (plus, given our demand structure, the increase in flow of consumer surplus is higher than the profit flow: $CS_\Delta = \ln\lambda > \frac{\lambda-1}{\lambda} = \pi$).[6] Therefore a subsidy to research increases total welfare. Any such policy would have to decide whether to favor improvements of existing applications or development of frontier applications.

To model this, we simply ask what the planner's payoff would be to one arrival of each type of innovation. For a quality improvement, the benefit is just $CS_\Delta/r = (\ln\lambda)/r$. Innovation in a frontier application would have two effects. First, it creates additional profit flow, without business stealing, hence a return to firms of π/r (assuming that the innovation is sold to a firm that then sells it at profit maximizing price; if instead the price is set at marginal cost, then the return is larger, but we want to focus on intervention in innovation alone). Second, a frontier innovation increases N and hence the steady state intensity of research from $\hat{x}_e(N^*)$ to $\hat{x}_e(N^*+1)$ (recall from Sec-

[6] CS_Δ denotes an increase of equilibrium consumer surplus flow of having one additional application: $CS_\Delta = CS(J+1,N) - CS(J,N)$.

tion 3 that $\hat{x}_e(N)$ is defined as the equilibrium research intensity if no further frontier applications are expected in the future).[7] This second effect is going to dominate in the long run, even though flow of benefit from inventing an improvement to existing applications is higher than from inventing a frontier application, because increase in x leads to an increase in the *growth* rate of future welfare. Therefore a sufficiently patient planner will prefer subsidizing frontier applications.

Formally, the steady-state free-entry condition is:

$$\frac{\pi}{r + \hat{x}_e(N)/N} - c(\hat{x}_e(N)) = 0 \tag{14}$$

This expression is increasing in N and decreasing in $\hat{x}_e$, hence unless $c(x)$ is vertical at $\hat{x}_e(N^*)$, it must be the case that $\hat{x}_e(N^* + 1) > \hat{x}_e(N^*)$. Furthermore, notice that this expression is decreasing in $\hat{x}_e$ faster if $c(x)$ is increasing faster (to the right of $\hat{x}_e(N)$).

One might be concerned about the behavior or N^* as r gets small; the following lemma shows that N^* converges to a finite number, and so, for large enough $\bar{N}$, an additional ladder is feasible even for small r.

Lemma 10 *$lim_{r \to 0} N^* < \infty$.*

Proof. See Appendix. ■

Given that N^* is well defined for small r, we can state formally the comparative static in r :

Proposition 5 *Suppose $c(x)$ is finite for all x. For sufficiently low r, the planner obtains a higher social return from (one-time, unexpected) investment in frontier applications than in existing applications.*

Proof. As we argued above, the total social return the one-time creation of one additional existing application is simply a constant flow of one additional "step" consumer surplus: CS_Δ / r. When instead a new frontier application is invented, customers do not gain any additional surplus immediately (as we assumed that the product will be sold by a monopolist) but they will enjoy a

[7] We assume that $c(\hat{x}_e(N))$ is concave in N to make sure that $N^* + 1$ will be indeed a steady state, see Lemma 4. There is also a third effect that the frontier application costs more, but compared to all future benefits it is likely to be small and hence we ignore η in this section.

faster rate of arrival of future innovations: $\hat{x}_e(N^*+1)$ instead of $x(N^*)$. (In terms of total surplus there is also an increase of profits for the industry, but it is sufficient to compare the consumers' gain). That leads to a total gain of:

$$\frac{CS_\Delta}{r^2}\left(\hat{x}_e(N^*+1) - x(N^*)\right)$$

The ratio of the total returns to customers from the two types of innovation is hence

$$\frac{\hat{x}_e(N^*+1) - x(N^*)}{r} \tag{15}$$

As $r \to 0$, condition (14) converges to:

$$N\pi = \hat{x}_e c(\hat{x}_e(N))$$

hence even in the limit $\hat{x}_e(N)$ is strictly increasing in N. Therefore, for sufficiently small r the return to customers (and total social return) is higher if the planner subsidizes frontier applications. ■

To understand the role of the shape of $c(x)$ on the planner's preference, consider the following comparative static: fix $c(x)$ for $x \leq x(N^*)$, but, for $x > x(N^*)$, let $\tilde{c}(x) < c(x)$, so that $\tilde{c}(x)$ is flatter (more elastic) than $c(x)$. Under $\tilde{c}(x)$, an additional ladder has more social benefit (because it will increase x by more than if the costs were given by $c(x)$,according to (15)). But, by contrast, an additional ladder is *less* attractive to the innovator under $\tilde{c}(x)$, for the same reason; the high rate of future arrivals discourages frontier research. As a result, for both cost functions, the steady state is N^* with research $x(N^*)$, but the social benefit of an additional ladder is greater under $\tilde{c}(x)$.

This logic is summarized in the following proposition.

Proposition 6 *Let $\tilde{c}(x) = c(x)$ for $x \leq x(N^*)$, but let $\tilde{c}(x) < c(x)$ for $x > x(N^*)$. Then the planner has a greater preference for frontier research under $\tilde{c}(x)$*

One interpretation of new ladders versus improvements is the contrast between basic research and applied developments. The model suggests both a rationale for governmental support for basic research (the greater research intensity that they foster), and an intuition for when this impact is likely to justify government involvement. For flat $c(x)$, the private benefit to frontier research is small, but the social benefit from an additional ladder is large. Both are driven by the fact that flat $c(x)$ leads to a big impact of a new ladder on equilibrium research intensity.

8 Conclusion

We have shown that analyzing two margins for innovation (improvements versus new products), the tradeoff between different market structures is not so much about *quantity* of innovation as about *allocation* of innovative efforts. Closed standards/firms with market dominance are good at coming up with new products while open standards/markets with many small firms are better at improving products. The welfare consequences depend crucially on the fundamentals of the technology, like possible scope of applications.

The models suggests that market structure impacts not only the amount but also the direction of technological change. Considering other ways that market structure affects the direction of technological change, as introduced by Acemoglu (2002), seems to be an interesting way to extend the basic structure.

The model could be applied in other ways. Klepper and Thompson (2005) show that a variety of empirical facts about firm dynamics can be explained through a model of submarkets, where submarkets arrive exogenously. One can view the arrival of new applications as a sort of new submarket; one would only have to add obsolescence for it to match Klepper and Thompson's notion more directly. We have focused on two market structures which correspond to some real world situations. Considering other market structures, including an optimal one for extraction of surplus for the monopolist, is a topic for future research.

9 Appendix A: Proofs

Proof of lemma 6. a) the benefit to innovation at a given product is at most $(1-\pi_j)/r = \frac{1}{r\lambda^j}$. As the cost of hiring researchers is at least $c(0) = \theta_l$, for high enough j the expected benefit becomes larger than marginal cost, making further research unprofitable. Now, suppose that the research stops at some j. If $\Delta_j > r\theta_l$, then a profitable deviation is to hire researchers with $\theta \in [\theta_l, \theta_l + \varepsilon]$ for some small ε. If $\Delta_{j-1} < r\theta_l$ then a profitable deviation is to reduce research activity to zero once the application reaches level $j-1$.

b) That follows directly from the assumption that $F(\pi/r - \eta) > 0$ (or equivalently, $\pi_1/r > \theta_l + \eta$), so the expected benefit of the first application is higher than the cost of hiring the most efficient researchers.

c) Given a researcher is hired, it is optimal to put his activity in a product with the highest expected return. The immediate return, Δ_j, is the highest for the goods with the lowest j and the continuation return (from subsequent innovations, $\Delta_{j+1}...\Delta_{j^*-1}$) is also highest for those goods. Hence such strategy is optimal (see the proof of lemma 7 for a more detailed calculation). ■

Proof of lemma 7. Rewrite the closed standard problem in the following way. Let $p_Q = (\pi(Q) - \pi(Q-1))/r$, the incremental benefit of the Qth innovation, where there are $Q^* = j^*\bar{N}$ innovations. Note that p_Q is weakly decreasing. Let $C(x) = xc(x)$.

The closed standard problem is choosing x_n, $n \in \{1, 2, ...Q^*)$ to maximize

$$\begin{aligned}
&\frac{x_1}{r+x_1}p_1 - \frac{r}{r+x_1}C(x_1) + \\
&\frac{x_1}{r+x_1}\left(\frac{x_2}{r+x_2}p_2 - \frac{r}{r+x_2}C(x_2)\right) + \\
&\frac{x_1}{r+x_1}\frac{x_2}{r+x_2}\left(\frac{x_3}{r+x_3}p_3 - \frac{r}{r+x_2}C(x_3)\right) + \\
&... \\
&\left(\prod_{i=1}^{N-1}\frac{x_i}{r+x_i}\right)\left(\frac{x_N}{r+x_N}p_N - \frac{r}{r+x_N}C(x_N)\right)
\end{aligned}$$

Let

$$G(Q) = \max_{\{x_Q,...x_{Q^*}\}} \sum_{i=Q}^{Q^*}\prod_{j=Q}^{i-1}\frac{x_j}{r+x_j}\left(\frac{x_i}{r+x_i}p_i - \frac{r}{r+x_i}C(x_i)\right)$$

Note that the solution to the whole problem is $G(1)$.Denote the solution to this problem for any Q by $\{x_n^Q\}$. Note further that $G(Q)$ is strictly decreasing, since, for $G(Q)$, choosing $\{x_Q^Q, x_{Q+1}^Q..., x_{Q^*-1}^Q, x_{Q^*}^Q\} = \{x_{Q+1}^{Q+1}, x_{Q+2}^{Q+1}..., x_{Q^*}^{Q+1}, 0\}$ gives at least as high a payoff as $G(Q+1)$, and increasing the last term from zero makes the payoff strictly higher.

The first order condition for x_n in the full problem is

$$p_n + G(n+1) = C'(x_n)(r+x_n) - C(x_n)$$

If $C'(x_n)(r+x_n) - C(x_n)$ is increasing, then it is immediate that, since the right hand side is strictly decreasing, the left hand side must be strictly increasing. Extending the proof to arbitrary C is a direct application of simple ironing techniques; letting $R(x) = C'(x_n)(r+x_n) - C(x_n)$, and defining the

$$\begin{aligned}\kappa(x) &= \int_0^x R(x)dx \\ \bar{\kappa}(x) &= conv(\kappa(x))\end{aligned}$$

the virtual value for the right hand side is defined almost everywhere by $\frac{d\bar{\kappa}}{dx}$, and extended by right limits everywhere else. ■

Proof of lemma 8. Note that, for the closed standard,

$$V_M(Q) = \max_x (1-\delta)(N\pi - xc(x)) + \frac{x}{x+r}(V_M(Q+1))$$

As $\overline{N} \to \infty$, x converges to a constant that maximizes

$$x\pi/r - xc(x)$$

The FOC is:

$$\pi/r = xc'(x) + c(x)$$

in the linear case, $F(\theta) = (\theta - \theta_l)/a \Rightarrow c(x) = ax + \theta_l + \eta$, the optimal choice is:

$$x = \frac{1}{2}\frac{\pi - (\theta_l + \eta) r}{ar} \tag{16}$$

and the sum of expected discounted profits:

$$\begin{aligned}V_M(0) &= (1-\delta)(0 - xc(x)) + \delta(1-\delta)(\pi - xc(x))... \\ &= (1-\delta)\sum_{k=0}^{\infty}\left(\delta^k k\pi\right) - xc(x) \\ &= \pi\frac{x}{r} - xc(x)\end{aligned}$$

(the last equality uses $\frac{\delta}{1-\delta} = x/r$).

Substituting the optimal x from (16):

$$V_M(0) = \frac{1}{4}\frac{(\pi - r(\theta_l + \eta))^2}{ar^2}$$

■

Proof of lemma 9. The hardware price (in terms of payoff flow) for an open standard is bounded by the steady state discounted sum of consumer surplus, which is

$$\ln(\lambda)\max\left\{1, \frac{\delta(N^*)}{1-\delta(N^*)}\right\}$$

In order to calculate $\frac{\delta(N^*)}{1-\delta(N^*)} = \frac{x(N^*)}{r}$, calculate:

$$\begin{aligned} V(N^*) &= \left(1-\frac{1}{N^*}\right)\left(\left(1-\delta\left(N^*\right)\right)\pi + \delta\left(N^*\right)V(N^*)\right) \\ &\Downarrow \\ V(N^*) &= \pi r\frac{N^*-1}{N^*r + x\left(N^*\right)} \end{aligned}$$

Now, the free-entry condition (7) is:

$$\frac{r}{x\left(N^*\right)+r}\pi + \frac{x\left(N^*\right)}{x\left(N^*\right)+r}V(N^*) = r\left(ax(N^*)+\theta_l\right)$$

Combining with $V\left(N^*\right)$ yields:

$$x\left(N^*\right) = \frac{1}{2a}\left(\sqrt{\left(arN^*+\theta_l\right)^2 + 4aN^*\left(\pi - r\theta_l\right)} - \left(arN^*+\theta_l\right)\right)$$

For any r, $x\left(N^*\right)$ is increasing in N^* and

$$\lim_{N^*\to\infty} x\left(N^*\right) = \pi - r\theta_l$$

which is a uniform bound on $x^*\left(N^*\right)$, that allows us to bound the price of hardware without finding the steady state. The hardware price (in terms of payoff flow) is at most:

$$h \leq \ln\left(\lambda\right)\frac{\pi - r\theta_l}{r}$$

Proof of lemma 10. Recall that $\hat{x}_e(N)$ is defined by the solution to (6) and (7) which yields

$$c\left(\hat{x}_e(N)\right) = \frac{\pi}{r + \hat{x}_e(N)/N}$$

So $\hat{x}_e(N)$ is decreasing in r and increasing in N. As $r \to 0$, this condition becomes:

$$c\left(\hat{x}_e(N)\right)\hat{x}_e(N) = N\pi \tag{17}$$

so $\hat{x}_e(N)$ converges to a number. Now, to see that the steady-state N^* is bounded away from ∞ as $r \to 0$,suppose $c'\left(x\right)$ is bounded from above. Suppose that for every N,

$$c\left(\hat{x}_e(N+1)\right) - c\left(\hat{x}_e(N)\right) > \eta$$

(the opposite inequality is our condition for N^*, see discussion before Lemma 4). Then, in (17) the LHS is growing at least linearly, while the RHS is growing much slower than linearly. So it cannot work. Therefore, even as $r \to 0$, N^* is bounded. ■ ■

10 Appendix B: $\phi < 1$

Throughout the paper we have maintained Assumption 2 that $\phi = 1$. How would the analysis change if instead:

Assumption 2B: $\phi < 1$.

The fundamental change is that profits earned by a frontier application will be higher than profits earned by subsequent improvements, that is $\pi_1 > \pi_{j>1} = \pi$. How is it going to affect the dynamics of innovation?

In the open standard case, if $\pi_1 - \eta < \pi$, then the equilibrium changes only slightly. In particular, equation (2) for non-frontier ($j \geq 2$) applications has to be supplemented by a value of frontier applications ($j = 1$) :

$$\begin{aligned} V_1(N) &= \rho(N)\left(1 - \frac{1}{N}\right)\left((1 - \delta(N))\pi_1 + \delta(N) V_1(N)\right) \\ &\quad + (1 - \rho(N))\left((1 - \delta(N+1))\pi_1 + \delta(N+1) V_1(N+1)\right) \end{aligned} \tag{18}$$

Also, we need to modify the benefit to x_f :

$$V_f(N) = (1 - \delta(N+1))\pi_1 + \delta(N+1) V_1(N+1)$$

and the free-entry condition changes accordingly. The N^* has to be calculated with the new value of $V_f(N)$ and a lower ϕ clearly leads to a (weakly) higher N^*. If π_1 is very large (in particular, sufficiently larger than $\pi - \eta$), then the equilibrium can change qualitatively as well: it is possible that for some N all research activity will be in the frontier applications, that is $x_f > 0 = x_e$. The reason Lemma 2 can be overturned is that when $\pi_1 - \pi$ is larger than η, paying more for a frontier application can be more than compensated by the higher profits (for this to happen it is not sufficient that $\pi_1 - \eta > \pi$, since due to business stealing returns to innovation are smaller than π_j / r).

In the closed standard the changes are less dramatic: we have already taken into account that the $C(x)$ function differs for frontier and non-frontier applications by ηx. Now we have to allow for $\pi(Q)$ to vary as well and increase Δ_0 accordingly. A higher π_0 will clearly imply a more intensive research on the frontier applications, but once $j = 1$ is reached for all $\overline{N}$, it will not have any further effects.

Finally, in terms of the choice between closed and open standard, the impact of a higher π_1 is ambiguous. On one hand, if $\overline{N} > N^*$ the closed standard will enjoy the higher π_1 from a larger number of products. On the other hand, a larger π_1 may tip an increase in N^* (even from 0 to 1 - see Remark 4) and hence to a large increase in the price of hardware.

References

[1] Acemoglu, Daron (2002). "Directed Technical Change," *Review of Economic Studies*, vol. 69. pp. 781-810.

[2] Aghion, Philippe and Peter Howitt (1992), "A Model of Growth through Creative Destruction," *Econometrica*, vol. 60(2), pages 323-51.

[3] Aghion, Philippe, Nick Bloom, Richard Blundell, Rachel Griffith and Peter Howitt (2005), "Competition and Innovation: An Inverted-U Relationship," *The Quarterly Journal of Economics,* vol. 102(2), pp. 701 - 728

[4] Bulow, Jeremy I. (1982), "Durable-Goods Monopolists," *Journal of Political Economy*, vol. 90(2), pages 314-32.

[5] Coase, Ronald H. (1972), "Durability and Monopoly," *Journal of Law & Economics*, vol. 15(1), pages 143-49.

[6] Economides, Nicholas (1996). "Network Externalities, Complementarities, and Invitations to Enter," *European Journal of Political Economy* vol. 12, pp. 211-232.

[7] Grossman, Gene M. and Elhanan Helpman (1991), "Quality Ladders in the Theory of Economic Growth," *Review of Economic Studies* 58, 43-61.

[8] Romer, Paul M (1987), "Growth Based on Increasing Returns Due to Specialization," *American Economic Review*, vol. 77(2), pages 56-62.

[9] Romer, Paul M. (2000), "Should the Government Subsidize Supply or Demand in the Market for Scientists and Engineers?," NBER Working Papers 7723, National Bureau of Economic Research, Inc.

[10] Scherer, Frederic M. (1980), *Industrial Market Structure and Performance,* Boston: Houghton Mifflin.

[11] Schmalensee, Richard (2000), "Antitrust Issues in Schumpeterian Industries," *American Economic Review* 90 (2), pp. 192-196.

[12] Shepard, Andrea (1987), "Licensing to Enhance Demand for New Technologies" *Rand Journal of Economics*18 (3), pp. 360-368.

[13] Stokey, Nancy L. (1981). "Rational Expectations and Durable Goods Pricing," *Bell Journal of Economics*, vol. 12(1), pages 112-128.