A Theory of Market Pioneers*

Matthew Mitchell† and Andrzej Skrzypacz‡

May 24, 2011

Abstract

An important type of innovation is one that pioneers a new submarket. Klepper and Thompson (2006) and Sutton (1998) show that innovation driven by the scope of the market can explain a variety of empirical facts. We introduce a model where innovators must decide whether to pioneer a new submarket or compete in an existing one. Unlike the prior research on submarkets and industry evolution, we endogenize this decision and show how the model generates an equilibrium long run scope of the market. We show that the model can explain some of the important existing facts about the product life cycle. We investigate the product cycle data and show other important features that are consistent with the model. Finally, we show that the model has implications for policy directed at particular kinds of innovations.

*We thank Philippe Aghion, Serguey Braguinsky, Alex Frankel, Paul Romer and participants at seminars at the Carnegie Mellon, Duke, the Federal Reserve Bank of Minneapolis, Iowa State, and Toronto as well as at the 4th Duke/Texas/Northwestern IO theory conference for suggestions.

†Rotman School of Management, University of Toronto.

‡Stanford GSB.

1 Introduction

Firms are driven by a constant process of finding new profit opportunities. Sometimes those profit opportunities involve overtaking an existing firm's place at the top of one activity. Other times the opportunities involve pioneering new activities, such as entering into a new submarket, which can be called "market pioneering." Klepper and Thompson (2006) and Sutton (1998) show that the arrival of new submarkets is an important driver of industry evolution. While those papers treat arrivals as exogenous, this paper introduces a theory of the trade-off between entering existing activities versus market pioneering, and shows how the model can endogenize an important margin for explaining the evolution of industries, the scope of an industry across distinct submarkets. To the extent that the scope of the market is an important determinant for industry evolution, understanding the forces that shape the equilibrium scope of the market is important to understanding industry dynamics. We further show how the model generates predictions that are familiar from some well known facts in industry dynamics, particularly about innovation over product life cycles. We also study several policy implications.

There are two main forces in our model. The force behind market pioneering in our model is one that we term an "escape competition" motive. Casual empiricism suggests that the introduction of new submarkets increases total industry profits. For instance, a new variety of laser studied in Klepper and Thomson (2006) might open up new applications of laser technology. Consistently with this view, in our model, when a firm pioneers a new submarket, industry resources are spread across a wider set of activities This makes the total value of the industry greater by making industry resources more scarce. As a result, pioneering a new submarket not only increases total profits but also increases profits per submarket, making pioneering attractive.

Countering this force, we assume that market pioneering is more expensive than overtaking existing submarkets.[1] We model market pioneering as this trade-off between escaping competition and higher costs. We describe the dynamics of the market innovation from a start of an industry to a steady state. We show how, even with symmetric submarkets to pioneer, these forces determine the steady state of innovation and the long-run breadth of the mar-

[1]Robinson, et al (1994) document the higher cost of innovation faced by market pioneers.

ket (i.e. the number of submarkets in a market). Certainly other forces, such as natural technological limitations, are at the heart of determining what is done by a technology such as the laser industry studied by Klepper and Thompson. The model developed here stresses an economic rationale behind the eventual scope of a technology. One can think of this idea as similar to the role of economic obsolescence in determining depreciation; there, both physical characteristics and economic motivations matter for determining the value of old versus new capital. Here the set of existing submarkets determines the relative value of innovations that generate new submarkets relative to innovations that improve existing ones.

Klette and Kortum (2004) study a model similar in spirit to ours, but where firms find new profit opportunities exclusively by overtaking other firms. They show that such a model can explain facts about the stationary distribution of firms. Our model enriches this idea by studying the breadth of the market. By adding the flavor of submarkets from Klepper and Thompson (2006) and Sutton (1998) to a model of innovation in the spirit of Klette and Kortum (2004), we can study facts concerning the innovation over the life-cycle of an industry that are fundamentally non-stationary.

We focus our attention on the period of increasing innovation in the product cycle, since our model is about the determination of the peak scope of the industry. We compare it to our model from birth to steady state. We show that in equilibrium research intensity follows a simple implicit equation, which implies rising levels of arrival of innovations. However, the composition of innovation is changing. Over time, the development of new submarkets declines, and overtaking innovations rise. In steady state, there are only overtaking innovations; all market pioneering is complete.

It is natural to interpret process innovations as disproportionately non-pioneering, while product innovations represent, at least partially, pioneering of new submarkets. With that interpretation, we can compare the model's predictions on pioneering to well known evidence on product innovations over the course of an industry life cycle. This evidence was documented first by Utterback and Abernathy (1975), and has been further discussed in papers including Cohen and Klepper (1996) and Klepper (1996). Innovations move from product to process innovations, with product innovations steadily falling and process innovations rising until the peak of the industry life cycle. Total innovation rises from birth to peak. Our stark depiction of an industry where total innovation is increasing while innovations are increasingly non-pioneering while introduction of new submarkets is dying out naturally follows this stylized fact. It generates both the changeover from product to process innovation, as well as the rise in total innovation, purely as a response to the number of submarkets that result.[2]

A key feature of our model is that the two types of innovation (pioneering and non-pioneering) draw on a common pool of a scarce resource, and therefore are jointly determined. This is in contrast to standard models of the relative fall in product innovation, such as Cohen and Klepper (1996) and Klepper (1996). They take the view that process innovation rises as firms get larger because it is more scale dependent than product innovation; the cost of each type of innovation is independent of the other. Our paper relies on no scale effects associated with one particular type of innovation, and delivers an absolute fall in the amount of product innovation over the life cycle, consistent with the evidence and contrary to the previous contributions. Instead, the driving force is that product innovations tend to increase the total market more so than process innovations. Our model shows that the decline in product innovation may not be tied to a change in the ability to find new product areas, but rather to a gradual decline in the incentives of firms to generate them.[3]

A fundamental and previously unexplored implication of our model is that the arrival rate of innovations increases over the product cycle, but at a slower rate than the increase in the number of firms. We show that this is a clear feature of the Gort-Klepper product cycle data. Moreover, the data on entry and exit is broadly consistent with the model's predictions. While our model is mostly silent on cross-sectional facts about firms of different sizes, it is consistent with data on movements in market share summarized in Klepper

[2]Our model does not include the final phase of the cycle, the eventual decline that leads to decreased innovation of all types. Extending the model to incorporate this would be straightforward. In our opinion, it is an attractive feature of the model that even without assuming declines in the opportunities to innovate, we get convergence to the steady state with pioneering decreasing over the life-cycle.

[3]A long line of papers, reviewed by Lieberman and Montgomery (1988), have discussed the relative merits of entering a market early (a first mover advantage) versus later. In those models, the market to enter is taken as given. Our model focuses on the arrival of the new opportunities for entry. Much of the literature on market pioneering in the management and marketing literature's focuses on the differences in outcomes between pioneers and subsequent entrants. Our model can be made consistent with this evidence, as we discuss, but explaining the possible advantage that pioneering firms enjoy is not the purpose of this paper. Rather, we focus on the less studied question of the trade-offs of pioneering versus "normal" innovation.

(1996). As the market matures, market shares become more stable.[4]

We show that modeling these different types of innovation is relevant to policy questions (see section 5). We consider the possibility that the government might want to subsidize particular types of innovation. We find that patient planners might prefer to subsidize pioneering research over non-pioneering research even when the profit and consumer surplus from each type of innovation is identical. The reason is that pioneering innovations have an additional impact on innovation, through the increase in profits per submarket that they generate.

The idea that some innovations change the set of available opportunities is essential to the model of long waves and short waves of Jovanovic and Rob (1990). Their model, however, focuses on contrasting epochs, one where only innovations in new areas occur, and another where only improvements to existing areas are developed. For the purposes of understanding an industry life cycle, our model has the useful feature that the two types of innovation coexist. This is because the innovation incentives are different in the two models; we stress the pioneer's dynamic motivation for inventing new submarkets, whereas in their model new areas are explored only when existing areas become too infertile. In our model both directions of innovation are always equally productive.

Finally, as in Klette and Kortum (2004), our approach to modeling innovation draws heavily on the endogenous growth literature, such as Grossman and Helpman (1991) and 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 departs from these growth papers in studying a finite number of products, focusing on the trade-off between pioneering and non-pioneering innovation, where the decision has a non-marginal impact on the evolution. We study transitional dynamics in order to address the product life cycle evidence.

[4]Our model also echoes the ideas of the original contribution of Utterback and Abernathy (1975), who argue that the data can be explained by imagining that product innovation proceeds until a dominant design emerges, at which point firms focus on process improvements. Similarly, our model predicts that falling product innovation is linked to stabilization in the breadth of the product: as the number of active submarkets increases, the pioneering benefits decrease while pioneering remains more costly than regular innovation. Once the industry reaches the steady state, one can think of the scope of the industry as the "design" that remains constant, although firms continue getting more and more cost efficient in producing the existing design in the steady state.

In the next section we introduce the model. Section 3 characterizes the equilibrium. In Section 4 we compare the results to both existing facts and some new empirical results from the product cycle data. We then discuss policy implications in section 5.

2 The Model

2.1 Products

At any given instant there is an industry made up of N submarkets (where N can change over time). For simplicity, we assume a very stark structure within a given submarket: the lowest cost producer for a given submarket captures it in its entirety and makes profits π per instant (so that total industry profits are $N\pi$). All agents use the common discount rate r.

This simple model of profits by submarket is analogous to assumptions in Klette and Kortum (2004) and Klepper and Thompson (2006); Klette and Kortum use the term "goods" and Klepper and Thompson use "submarkets.". In Section 5 we introduce an explicit preference structure that delivers this as the outcome of a Bertrand price setting game (along the lines of Grossman and Helpman (1991), which also motivates Klette and Kortum (2004)).

The important feature of the assumption is that submarket profit flows are constant across submarkets and time. It is meant to not bias innovation toward one type of research or another (i.e. profits from new submarkets and improvements to current submarkets have identical profit flows), or to drive the amount of research that takes place over time (as profits per instant are constant in time). Therefore the life-cycle changes in total innovation and the shares of different types of innovation are not driven by *assumed* differences in profitability. Appending such differences is straightforward, and does not change the basic message about the endogenous determination of industry scope.

2.2 Innovation

Innovation comes through research. Research can be done on either developing new submarkets or on improvements in existing submarkets. Because we think of product innovation as being necessary for pioneering new submarkets, we will later use the term product innovation to describe those pioneering innovations, and term improvements (say in cost) of producing existing products as process innovation.[5]

In both cases, research takes place continuously and innovations arrive according to a Poisson process. The arrival rate is equal to the amount of research intensity, denoted x_e for improvements to existing submarkets and x_f for new (frontier) submarkets. We assume that research intensity comes from one input, researchers, and that the pool of this input is heterogeneous in their skills. A researcher of type θ can provide one unit of research intensity, at a cost flow of θ for an existing submarket, and a cost flow of $\theta + \eta$ for a frontier submarket. The inclusion of $\eta > 0$ means that new submarkets are more costly to research (pioneering is expensive), as has been documented. 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. Researchers may enter freely into research at any instant and choose which submarket to research on. In order to keep the model from trivially having no innovation, we assume that θ_l is small enough so that in equilibrium at least one innovation takes place. We also assume that F is atomless and strictly increasing with derivatives bounded away from zero and infinity.

The critical feature of our model that ties different submarkets together is that they draw researchers from a common pool of scarce talent. As a result, innovative effort on one submarket has an impact on the marginal cost of innovation for all of the submarkets that use the common factor. This equilibrium effect plays an important role in the industry dynamics that we develop below.[6]

We assume that when a researcher generates an innovation he forms a firm and markets it, becoming the new low cost producer and earning profits per instant π. This mirrors the market structure commonly used in the endogenous growth literature, and in applications such as Klette and Kortum (2004). In the case where an innovation arrives in a submarket that has been pioneered before, the new firm replaces the previous leader as the new low-cost producer, leaving the whole industry profits unchanged, while in case of a pioneering innovation, the new firm is the first firm in the new submarket,

[5]Throughout we consider non-pioneering innovations to be cost reductions; equally valid are quality improvements that lower cost per efficiency unit provided.

[6]While we use the heterogeneous types of researchers to generate the scarce resource and hence an upward-sloping supply of research, anything that generated cost increasing in total research intensity would generate the same results.

increasing the total industry profits by π. Researchers (and the firms they form) maximize the expected sum of discounted profits net of research costs. We define the Markov equilibrium of this model after we introduce some more notation in the next section.

3 Innovation Dynamics

In this section we define and characterize the equilibrium of the industry. We establish an endogenous bound on the number of submarkets that will ever be developed in equilibrium and describe equilibrium dynamics of innovation.

3.1 Equilibrium

Keeping fixed the profit flows π (that can be endogenized as in Section 5), consider the reduced-form model of competition among researchers. We focus on symmetric Markov equilibria, in which the decision rules of the researchers depend only on the current state of the industry which is summarized by N, the number of submarkets pioneered so far. Since obtaining an improvement over an existing submarket yields the same profit flows regardless of the identity of the submarket or the current cost level of the submarket, we focus on equilibria where the strategies of the researchers (and expected profits of the producers) are symmetric: all submarkets that have been pioneered are researched equally.

We start by introducing some notation that allows us to make the model more tractable. Denote total research by $x(N) \equiv x_e(N) + x_f(N)$. Since the benefits of research are independent of type, researchers follow a cutoff rule: if type θ chooses to research, all types $\theta' < \theta$ do as well (although they can do it on some other submarket). By definition, the cutoff $\bar{\theta}(N)$ solves:

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

Total intensity $x(N)$ (and hence 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 submarkets (unless $x_e(N) = 0$), researching a frontier submarket (unless $x_f(N) = 0$) and opting out of research.[7]

[7]The equilibria we construct pin down only the share of active researchers working on

Define $c(x) \equiv F^{-1}(x)$, the cost of the marginal researcher in an existing submarket given the total research intensity x, so that $\bar{\theta}(N) = c(x(N))$. Note that $c(x)$, the industry supply of researchers, is increasing. Let $\rho(N) = \frac{x_e(N)}{x(N)}$ be the probability that conditional on a new innovation arriving, it is for an existing submarket. Let the random variable $\tau(N)$ be the arrival time of an innovation given the aggregate research intensity, $x(N)$.

Let $V(N)$ be the (flow) value of an incumbent firm just as a new innovation arrives, but without knowing what submarket it is for, including the possibility that it pioneers a new submarket. Let $V_e(N)$ be the flow value of a firm that has just developed a non-pioneering innovation and $V_f(N)$ be the flow value of a firm that has just developed a pioneering innovation, in each case where previously there were N submarkets. These functions are defined recursively by:

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

In these recursions, the expected discount factor (to the time of arrival of a new innovation), incorporating both the arrival rate and the rate of discount, 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 it depends on N through $x(N)$.

$V_e(N)/r$ 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:

$$\begin{aligned} V_e(N)/r &\leq \bar{\theta}(N) \\ V_f(N)/r &\leq \bar{\theta}(N) + \eta \end{aligned} \tag{3}$$

the particular tasks but not the allocation of particular types. This additional characterization can be achieved by introducing heterogeneity in research costs across different submarkets. For example, allowing η differ with θ one can pin down which types will be pioneers.

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

Formally, the (symmetric, Markov) equilibrium is an aggregate strategy profile $(x_e(N), x_f(N))$ such that all agents optimize. Optimization is equivalent to strategies such that the free entry equations (3) hold, with entry values defined by (2), and individual decisions agree with aggregate variables in (1).[8]

The following proposition summarizes the characterization of the equilibrium. Informally, the industry starts with $N = 0$ and all researchers are involved in pioneering. After the first submarket is created, research activities change: some researchers work on improvements in the first submarket and some try to pioneer a new submarket. As new submarkets are introduced, the total research activity grows in absolute terms but tends to decrease in per-submarket terms. After several pioneering inventions the market reaches a steady state, $N = N^*$, and all pioneering ceases: from that moment on all research is directed at existing submarkets.

Proposition 1 *In equilibrium 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, with $N^* = \arg\min_{N\in\{1,2,\ldots N^*\}} x(N)/N$.*

Before we prove this characterization of equilibrium with a series of lemmas (and provide a sufficient condition for uniqueness of the equilibrium), we discuss the economic intuition behind the results.

First, it might seem surprising that in equilibrium competitive researchers can ever be willing to pay the additional cost to develop a new submarket. After all, we have assumed that the profit flows, π, are the same from both types of innovation! Still, in equilibrium if η is not too large (and the supply of researchers, $c(x)$, is not too elastic) the steady-state will have $N^* > 1$. The reason is that pioneering generates a non-marginal impact on profits per submarket. Due to the increasing cost of researchers, research intensity on existing submarkets rises less than proportionally to the number of existing submarkets. Therefore $x_e(N)/N$ – the amount of research on improvements per submarket – tends to decline in N, and the value of having a marketable submarket rises in N. If it rises enough from N to $N+1$, it is worth paying the extra cost η. However, since $1/N$ is convex and $x_e(N)$ is weakly increasing,

[8]In the proof of the proposition characterizing equilibria, we provide sufficient conditions for the equilibrium to be unique.

$x_e(N)/N$ drops by very little for large N. Eventually the increase in the profit per submarket, $V(N)$, becomes small and is insufficient to draw research into pioneering new submarkets.

Second, anytime there is pioneering research, there must also be research on improving existing submarkets. If there were only research on new submarkets for a given N, then developing an improvement would yield more profits than a new submarket: it would earn profits until the next new submarket, at which point the continuation value would be as much as the new submarket would have made. This is of course impossible since improvements are less expensive.

We now turn to a construction of equilibria and hence a proof of the proposition. 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 characterize the equilibrium aggregate research intensity for these states. Note that (2) implies $V_e(N+1) = V_f(N)$. Combining it with the free entry conditions we get:

$$c(x(N+1)) = c(x(N)) + \eta \tag{4}$$

This 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 (4) and monotonicity of $c(x)$.

Regarding $V(N)$, from the free-entry conditions (3) we have

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

where $(\delta(N) - \delta(N+1)) < 0$ because we have proved 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 the steady state both research dimensions 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, by (2) we get $V_e(N) = (1-\delta)\pi + \delta V_f(N) \geq V_f(N)$ which contradicts the free entry conditions (3).

Now, suppose that there exists an $N > 0$ such that $x(N) = 0$. Then $\delta(N) = 0$ and $V_e(N) = \pi$. Then any researcher who found it profitable to do pioneering research at $N = 0$ finds research on existing submarkets even more profitable at N, which contradicts $x_e(N) = 0$.

Finally, suppose that $x_f(N) > 0$ for all N. Then $x_e(N) > 0$ and therefore we know that $V_e(N+1) = V_e(N) + r\eta$ for all N. But that is not possible as $V_e(N) \in (0, \pi)$. ■

In a steady state N^*, the values and research intensity can be easily determined as they satisfy the Bellman equation and the free-entry condition:

$$V(N^*) = \left(1 - \frac{1}{N^*}\right)\left[(1-\delta(N^*))\pi + \delta(N^*)V(N^*)\right] \tag{5}$$

$$\underbrace{(1-\delta(N^*))\pi + \delta(N^*)V(N^*)}_{V_e(N^*)} = rc(x_e(N^*)) \tag{6}$$

Call the solution to these two equations, for arbitrary $N \geq 1$, $\hat{V}(N)$ and $\hat{x}_e(N)$, with the associated $\hat{\delta}(N) = \frac{\hat{x}_e(N)}{\hat{x}_e(N)+r}$.

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

Proof. First, condition (6) implies that $\hat{V}(N)$ is increasing if and only if $\hat{x}_e(N)$. To see this, note that

$$\left(1 - \frac{x}{x+r}\right)\pi + \frac{x}{x+r}V - rc(x)$$

is increasing in V and decreasing in x when $V < \pi$ (which is the case since $\hat{V}(N) < \pi$), and hence for (6) to hold for all N, $\hat{V}(N)$ and $\hat{x}_e(N)$ have to change in the same direction in N.

Second, suppose by contradiction that $\hat{V}(N)$ is decreasing and hence $\hat{\delta}(N)$ is decreasing. However, since (5) can be re-written as:

$$\hat{V}(N) = \pi\left(1 - \frac{1}{N - \hat{\delta}(N)(N-1)}\right)$$

if $\hat{\delta}(N)$ is decreasing in N, (5) would imply that $\hat{V}(N)$ is increasing, a contradiction. ■

In order for N^* to be a steady state, it must not be profitable to search for a new submarket at N^*, even if no new submarkets were researched after. In other words, at the steady state N^*, the following inequality holds:

$$(1-\widehat{\delta}\,(N+1))\pi + \widehat{\delta}\,(N+1)\,\hat{V}(N+1) \leq r\,(c\,(\hat{x}_e\,(N)) + \eta) \tag{7}$$

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 (7) 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 the whole equilibrium:

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

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. The intuition why we should expect $c\,(\hat{x}_e\,(N))$ to be concave is as follows: start with a constant $\delta\,(N)$. Then the solution $V\,(N)$ to (5) is concave because $\left(1-\frac{1}{N}\right)$ is concave. Now, for (6) to hold, $c\,(\hat{x}_e\,(N))$ and $\delta\,(N)$ have to increase in N. 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 a concave $c\,(\hat{x}_e\,(N))$.

Once we find N^* and $x_e\,(N^*)$, we can solve for the rest of the equilibrium by working from the eventual steady state. In particular, iterating on (4) 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 the first two lines of (3) 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 the last line of (3). Note that this construction is unique for any given N^*, so that a concave $c\,(\hat{x}_e\,(N))$ implies a unique equilibrium (as we claimed in Lemma 4). Finally, we claim:

Lemma 6 $N^* = \arg\min_{N\in\{1,2,\ldots N^*\}} x(N)/N$.

Proof. First consider $x_e(N)/N$. If $N^* \notin \arg\min_{N\in\{1,2,\ldots N^*\}} x_e(N)/N$, let $\hat{N} < N^*$ be the largest N such that $x(N)/N$ is minimized. When the current state is N,define by $p_N(\tau)$ the probability of any current firm being the submarket leader τ periods forward. Since $x_f(\hat{N}) > 0$ and $x_e(N)/N > x_e(\hat{N})/\hat{N}$ for $N > \hat{N}$, $p_N(\tau) \in [0, e^{-(x_e(\hat{N})/\hat{N})\tau})$. We can write $p_N(\tau)$ recursively as

$$p_N(\tau) = \int_0^\tau e^{-(x_e(N)/N)t} x_f(N) e^{-x_f(N)t} p_{N+1}(\tau - t)dt + e^{-x_f(N)\tau} e^{-(x_e(N)/N)\tau}$$

The integral adds up possible arrival dates of the $N+1st$ market pioneer; the last term adds the chance of being in business with no pioneer arriving before τ.

Further, define the operator on the space of bounded, continuous functions, (with the sup norm)

$$T(p(\tau)) = \int_0^\tau e^{-(x_e(\hat{N})/\hat{N})t} x_f(\hat{N}) e^{-x_f(\hat{N})t} p(\tau - t)dt + e^{-x_f(\hat{N})\tau} e^{-(x_e(\hat{N})/\hat{N})\tau}$$

Since this operator satisfies Blackwell's sufficient conditions for a contraction mapping, it is a contraction and has a unique fixed point. The fixed point can be calculated to be $e^{-(x_e(\hat{N})/\hat{N})\tau}$. Since $p_{\hat{N}+1}(\tau) <. e^{-(x_e(\hat{N})/\hat{N})\tau}$, and iterations on T starting from any initial function are converging monotonically to the fixed point $e^{-(x_e(\hat{N})/\hat{N})\tau}$ of T, it must be the case that $p_{\hat{N}}(\tau) = T(p_{\hat{N}+1}(\tau)) > p_{\hat{N}+1}(\tau)$ for all τ. The chance of being in business in τ periods is higher at $\hat{N}$ than $\hat{N}+1$. Therefore $V(\hat{N}+1) < V(\hat{N})$, and therefore the return to frontier research at $\hat{N}$ is no greater than existing submarket research at $\hat{N}$, but is more costly. This contradicts the in equilibrium both $x_f(\hat{N})$ and $x_e(\hat{N})$ are strictly positive.

Since $x_f(N^*) = 0$, clearly $N^* = \arg\min_{N\in\{1,2,\ldots N^*\}} x_f(N)/N$, and as a result $N^* = \arg\min_{N\in\{1,2,\ldots N^*\}} x(N)/N$. ■

Intuitively, if $x_e(N)/N$ were minimized somewhere before N^*, then the innovators would always be facing worse prospects for overtaking if they raise the state past the point where $x_e(N)/N$ is minimized. They would therefore do no frontier research. Since $x_f(N)$ is minimized by definition at N^*, it must be the case that $x(N)/N$ is minimized at N^*. That finishes the proof of Proposition 1.

Although the model is admittedly stylized and we have isolated only one particular reason for competitive research leading to increases in scope (the reduction in $x_e(N)/N$), other reasons could be easily incorporated as well. For instance, new submarkets (the pioneering innovation) might initially be harder for others to improve on, or initial innovations might be more valuable. Initial innovation might lower future research costs or provide a private benefit in research for the pioneering firm. Whatever the other forces, as long as the supply of researchers, $c(x)$ is increasing, by the same economic intuition a higher N is most likely to yield a higher $x(N)$ (so that innovation is accelerating over time).

3.2 Numerical Example

In order to see how the equilibrium evolves in a numerical example, we assume that F is linear, $F(\theta) = (\theta - \theta_l)/a$, and so $c(x) = ax + \theta_l$. Further, we consider the following parameters: $r = 5\%$ (annual interest rate), $a = 0.01$, $\theta_l = 0.2$, $\eta = \theta_l/25$, $\pi = 1/3$ (which corresponds to $\lambda = 1.5$ so that costs decrease by 50% with every innovation - see Section 5 for details). Then the steady state number of submarkets is $N^* = 9$. In the steady-state $x(N^*) \approx 9.66$, which is also the average number of innovations per year (across the whole industry). 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 submarkets (pioneering), and the third line is investment in existing submarkets.

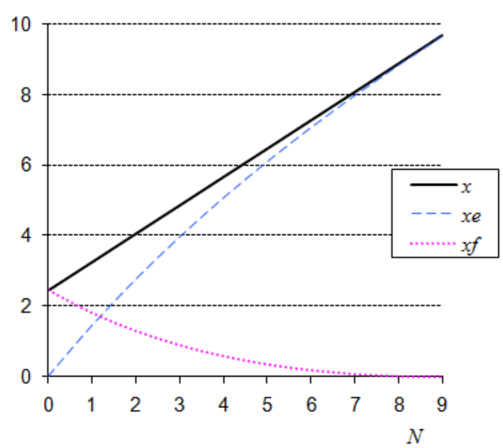

Note that the difference between the extra cost η of a new submarket is at most four percent of the cost of researching an improvement (for $\theta = \theta_l$), yet differences between intensities research intensities new and existing submarkets are large.

4 Evidence

Does our model help us understand empirical patterns of innovation? Admittedly, it is very difficult to match a stylized model like ours to data, especially that we do not have good measures of research intensities and innovation (in terms of measurement, we follow the literature and use patent counts as a proxy for innovation rates). Moreover, unfortunately, there is little empirical work on the rate of pioneering vs. non-pioneering innovation over the industry life cycle.

Instead, most of the existing evidence on direction of innovation is focused on the patterns of product vs. process innovation. Yet, we think that there is a fundamental connection between these two ways of categorizing innovation because pioneering is overwhelmingly a product innovation while improvements to existing products can take the form of either product or process innovation. Suppose that α share of the improvements are product innovations and $(1-\alpha)$ are process innovations. Then, letting x_{Pd} be the amount of product innovation and x_{Pc} be the amount of process innovation we can impute $x_e = \frac{x_{Pc}}{1-\alpha}$ and $x_f = x_{Pd} - \frac{\alpha}{1-\alpha} x_{Pc}$. As long as α is not too large (an assumption we will maintain), even though the levels of process and product innovation are going to be different than the levels of pioneering and non-pioneering innovations, the relative patterns will be very similar.

4.1 Gort and Klepper (1982) and Utterback and Abernathy (1975)

Our model reproduces some key features of the evolution of innovation over the product cycle, from what has commonly been termed stage I (birth) to stage III (the peak level of firms).

Gort and Klepper (1982) document that the rise in firms is met with a rise in patenting. Interpreting innovations in the model as patents, one can see that pattern emerge, since each submarket is sold by one firm.[9] Utterback

[9] Clearly, the same pattern emerges if the number of firms is proportional to the number

and Abernathy (1975) stress that product innovation declines, and process innovation increases over this period of the product cycle moves forward. This is clearly a feature of the model. Note that existing models such as Klepper (1996) focus on the relative amounts of the two types, rather than the absolute quantities and do not deliver an absolute decrease in product innovation. Utterback and Abernathy (1975) also document a change from innovations that require original components, to ones that focus on adopted components and products. That change can again be interpreted as the move from x_f to x_e; research into new submarkets requires the extra cost η because it does not build on prior art in the same way a process innovation does.

Gort and Klepper (1982) also document a shift from major to minor innovations. There is a sense that product innovations in the model are major: they increase the total number of firms and the amount of research in the industry. Process innovations merely replace one firm with a lower cost firm, increasing consumer surplus, but leaving industry variables unchanged.

The tight connection between product innovations and major ones might seem excessive. However, it can easily be loosened. Our process innovations are analogous to quality improvements (in the standard language of Grossman and Helpman (1991)) increasing efficiency units per physical unit by a factor of λ, and leaving cost unchanged. Our main argument, however, is that while x_e might include some product innovation, x_f is by its nature all product innovation. Our basic story is that new submarkets are more researched early on, and less researched later on, and therefore innovation switches from product to process and major to minor.

4.2 Prices and Innovation

Gort and Klepper (1982) point out that price changes do not reflect changes in innovation rates, as measured by patents; as patents rise with the industry moving to maturity, rates of price decline fall. One explanation is that patents are not a good proxy for innovation; the model suggests another. As the number of submarkets rises, there is more innovation in total, but it is spread across more submarkets, and therefore each innovation has less impact on market-wide prices.

Suppose each innovation within that submarket lowers the price in that

of submarkets, for example having two firms active per submarket.

submarket by a fixed percentage.[10] If the rate of price decline is measured at the market level by averaging across the submarkets, prices are falling at the rate of $x(N)/N$, which is minimized at maturity (N^*); in other words, innovation rates are rising but prices are falling less and less rapidly.

4.3 Number of Firms and the Rate of Innovation

One key feature of the model is that innovation is tightly linked to the number of firms in the industry. To see if that prediction fits the data, we study the expanded version of the Gort and Klepper data introduced by Agarwal (1998). It includes data on firm numbers and patents (our proxy for innovations) for 31 products. We focus our attention on the data for each product cycle up to the maximum number of firms is achieved, which corresponds to N^* in the model and Stage III in the language of Gort and Klepper (1982).

We estimate the arrival rate of innovations (measured by patents) according to a negative binomial specification; that is, the log of the arrival rate is linear in a set of regressors. Our specification allows for random effects in the arrival rate across product-year pairs. Our model suggests that arrival rates are increasing in the number of firms, but less than proportionally on average (as $x(N)/N$ is minimized at N^*). We therefore estimate the log of the arrival rate as a function of the log of the number of firms, as well as year and product fixed effects:

variable	coefficient	standard error
ln(firms)	.16	(.03)

Dependent variable - ln(arrival rate of patents)

From the regression results we see that the data is consistent with a key feature of the model. Since the coefficient on ln(firms) is between zero and one, total number of firms, N, is positively correlated with an increase of total innovation, $x(N)$, but negatively with innovation per firm, $x(N)/N$.[11]

The data also includes entry and exit counts. The model has a less clear implication for entry and exit. Inventions could be purchased by existing firms, achieved in-house or done by new ones. For a tangible example we can make the extreme assumption that all inventors market their own inventions,

[10] See Section 5 for a specific model that delivers such price dynamics.

[11] The results survive the inclusion of various forms of product-age trends, suggesting that firm counts are not simply a proxy for age.

so every innovation leads to a new firm. That is, in the model arrival rate of innovations and entry are one-to-one. Of course, the imposition of this strong assumption implies that the model may have a much harder time explaining the entry and exit data.

Since in the model x increases less than proportionally to firm numbers, entry should, too, under this interpretation. To test this, we estimated the relationship between the log arrival rate of entrants and the log of the number of firms (with year and product fixed effects):

variable	coefficient	standard error
ln(firms)	.75	(.05)

Dependent variable - ln(rate of entry)

This is once again consistent with less-than-proportional growth that the model predicts (but since the coefficient is different than in the previous regression, the relationship between innovation and entry is more complicated than the extreme assumption we suggested).

In our model exit occurs only when a firm is overtaken on a given submarket; therefore exit is associated with x_e. This too increases less than proportionally with firm numbers; note, however, that the model suggests a stronger positive relationship for x_e (which is increasing but less than proportionally) than x, which the sum of x_e and the declining x_f. For instance, for the numerical example calculated in Section 3.2, a simple regression of $\ln(x_e)$ on $\ln(N)$ gives a slope coefficient of .21, but the same regression of $\ln(x)$ on $\ln(N)$ gives a coefficient of .13. Therefore, we expect that replacing the entry rate with exit rate in the the previous regression should yield a higher coefficient, but still less than one. The regression results are:

variable	coefficient	standard error
ln(firms)	1.00	(.06)

Dependent variable - ln(rate of exit)

The results do not contradict the model. In particular, the parameter is significantly greater than the one from the prior estimation; it cannot be statistically distinguished whether the parameter is greater or less than one In summary, the entry data is in accord with the theory of less than proportional growth in x with firms. The exit data is consistent with being more positively related to firm numbers than entry, and is not inconsistent with less than proportional growth.

Finally, Klepper (1996) points out that market shares become more stable as the product cycle develops. If each innovation is marketed by a new firm, the model predicts a per-instant expected fall in market share for an incumbent of

$$\frac{1}{N}\frac{x_e(N)}{N} + x_f(N)\left(\frac{1}{N} - \frac{1}{N+1}\right)$$

The first term is the probability of being replaced by an improvement (and hence losing all market share); the second term is the probability that a new submarket arrives ($x_f(N)$) times the change in market share (market share declines from $1/N$ to $1/(N+1)$). Since $x_e(N)/N$ and $x_f(N)$ are both declining on average (since they are minimized at N^*), this is on average decreasing in N, and therefore market shares become more stable as N rises. So the model is consistent also with this empirical fact.[12]

5 Policy Implication: Subsidizing Innovation

We now ask how the model can inform us about a typical policy question in the innovation literature, the appropriate subsidy for innovative activity. We focus on the new question raised in our model: the direction of innovation and ask about the relative social returns to a subsidy to a pioneering vs. non-pioneering innovation.

5.1 A Model of Preferences

In order to study policy implications, it is necessary to fully describe the demand side of the economy (to be able to talk about welfare).

In a given moment of time with N submarkets, in submarket i there is a set of firms J_i. Firm $j \in J_i$ can produce the good at a constant marginal cost

[12] Utterback and Abernathy (1975) argue that the shift from product to process innovation is driven by the development of a "dominant design." It is possible to interpret our model to be consistent with this pattern. To see this, reinterpret "submarkets" as "components" (a term from Utterback and Abernathy) of a common design. Then the model endogenizes dominant design, namely to be one with N^* components. When the dominant design arrives, process innovation takes its maximum value. Notice that the time to the arrival of the dominant design is endogenous, and not determined by a technological constraint, but rather by the shape of the cost of innovation (and stochastic results of research). In other words, our model adds the idea that the dominant design is being importantly impacted by the resources the industry has to draw on.

$c_i^j \leq 1$. Firms within a submarket are ordered in a decreasing order of costs. For a given submarket i, the representative consumer consumes d_i^j units of products from firm j. This leads to d_i units from the submarket, where[13]

$$d_i = \sum_j d_i^j$$

The representative consumer's instantaneous utility from a bundle of units $\{d_i\}$ of the various submarkets is

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

where p_i^j is the price paid for good j in submarket i.

In equilibrium consumers will all consume the lowest cost product, denoted simply c_i, for each submarket. Therefore, given consumption of d_i units of that cost level of submarket i, utility is

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

To achieve stationarity, we parameterize the utility function to be logarithmic: [14]

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

With this utility, demands are independent across submarkets and the representative consumer spends a constant share of his income on every submarket i. In particular, if the representative consumer buys in submarket i

[13] The preference structure for a given sub-market follows Grossman and Helpman (1991). Although we are focusing on cost reductions, an analogous model can be written with all firms having the same costs but inventions increasing quality of the product. These two models are mathematically equivalent if consumers care equally about quality. The mapping is achieved by measuring the cost of providing quality-adjusted units. With that accounting an improvement in product quality while cost stays constant is equivalent a reduction of cost while quality stays constant.

[14] Additive and logarithmic utility is common in the growth literature because it generates stationarity of profits with constant percentage-decreases in cost levels. Here it also delivers us, in a stark way, that profit flows are the same for all leading firms in each submarket and do not change with the introduction of new submarkets (no business stealing across submarkets).

at price p, his demand is $d_i^j = 1/p$ (where we normalize the total spending per submarket to 1).

On the supply side, we assume that innovation reduces costs by a factor $\lambda > 1$. That is, if in a given submarket the lowest cost firm j has a cost c_i^j, if a new researcher invents an improvement, his firm will have costs $c_i^{j+1} = c_i^j/\lambda$. The first firm to operate in a submarket has cost $c_i^1 = 1/\lambda$. For each submarket, if nothing has been invented, the consumers have an outside option that is provided competitively at marginal cost 1. One can interpret this as the next best alternative product that might substitute for submarket i.

Following the endogenous growth literature such as Aghion and Howitt (1992) and Grossman and Helpman (1991), we suppose that each cost level is monopolized, perhaps due to a patent or a trade secret. There is a Bertrand competition between firms within a ladder. Non-lowest-cost firms price at marginal cost; to match this price, the lowest cost producer charges $p_i^j = 1/\lambda^{j-1}$ for $j > 1$ and $p_i^1 = 1$. Given demand x_i on ladder i, profit flows for the leader are $d_i(p_i - c_i)$, where under Assumption 1, $d_i = 1/p_i$. As a result, profits are $\frac{\lambda-1}{\lambda} \equiv \pi$.

5.2 Social Benefit of Different Types of Innovation

The equilibria we have characterized above do not achieve social first-best along many directions. Hence there are many ways a government can intervene to improve efficiency. We focus on how the social returns compare between subsidizing pioneering or non-pioneering research.

For clarity, we focus on a one-time unexpected subsidy for the marginal researcher in the steady state to avoid crowding-out of private innovation caused by an anticipated future government subsidy.

In equilibrium the steady state innovation level x^* is typically 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 $\frac{\pi}{r+x^*/N}$ while the social return is equal to CS_Δ/r, where CS_Δ is the extra flow consumer surplus generated by the extra reduction in costs and prices.[15] Typically, the second number is higher because the social ben-

[15] The private benefit is the profits stolen from the previous lowest-firm firm that accrue until a new firm enters the submarket. The social benefits depend only on the change of consumer surplus, since on the producer side there is simply a transfer of profits from the previously-leading firm to the new entrant. The new entrant lowers price, which results

efits accrue forever, while the private returns occur only until the firm gets replaced by an improvement. Additionally, given our demand structure, the increase in flow of consumer surplus is higher than the profit flow: $CS_{\Delta} = \ln\lambda > \frac{\lambda-1}{\lambda} = \pi$.

Therefore, for the utility functions assumed above, a subsidy to non-pioneering research in the steady-state unambiguously increases total welfare. But is it larger or smaller than the social return to subsidizing a pioneering research? To model this, we simply ask what the planner's payoff would be to one arrival of each type of innovation. As we argued, for a non-pioneering innovation, the benefit is just $CS_{\Delta}/r = (\ln\lambda)/r$. A pioneering innovation creates a new submarket and has two main effects.[16] First, it creates additional profit flow, 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 pioneering innovation increases N and hence the steady state intensity of research from $\hat{x}_e(N^*)$ to $\hat{x}_e(N^*+1)$ (recall from Section 3 that $\hat{x}_e(N)$ is defined as the equilibrium research intensity if no further frontier applications are expected in the future).[17] Since an increase in x increases the growth rate of future welfare, the second effect is going to dominate in the long run, even though flow of benefit from inventing a cost improvement is higher than from inventing a new submarket. Therefore a sufficiently patient planner would prefer subsidizing new submarkets.

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

$$\frac{\pi}{r + \hat{x}_e(N)/N} - c(\hat{x}_e(N)) = 0 \tag{8}$$

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 $\hat{x}_e(N^*+1) > \hat{x}_e(N^*)$.[18] Furthermore, notice that this expression is decreasing in $\hat{x}_e$ faster if $c(x)$ is increasing faster (to the right of $\hat{x}_e(N)$).

in gain in consumer surplus that accrues forever.

[16]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.

[17]For this section we assume that $c(\hat{x}_e(N))$ is concave in N to make sure that N^*+1 will be indeed a new steady state, see Lemma 4.

[18]We have assumed that $F()$ is strictly increasing with bounded derivatives, which implies that $c'(x)$ exists and is bounded from above. It guarantees that a higher N leads to a strictly higher research intensity.

Since our results are "for sufficiently low r", one might be concerned about the behavior of 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 submarket is feasible even for small r.

Lemma 7 $\lim_{r\to 0} N^* < \infty$.

Proof. Recall that $\hat{x}_e(N)$ is defined by the solution to (5) and (6) 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:

$$c\left(\hat{x}_e(N)\right)\hat{x}_e(N) = N\pi \tag{9}$$

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 to the contrary 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). Evaluating (9) at N and $N+1$ and subtracting, we get that for all N :

$$\underbrace{\left(c\left(\hat{x}_e(N+1)\right) - c\left(\hat{x}_e(N)\right)\right)\hat{x}_e(N)}_{>0} + \underbrace{c\left(\hat{x}_e(N+1)\right)}_{>N\eta}\underbrace{\left(\hat{x}_e(N+1) - \hat{x}_e(N)\right)}_{>0} = \pi \tag{10}$$

Note that since we assumed F is atomless with bounded derivatives, c also has bounded derivatives. Therefore, $c\left(\hat{x}_e(N+1)\right) - c\left(\hat{x}_e(N)\right) > \eta$ implies that $\left(\hat{x}_e(N+1) - \hat{x}_e(N)\right)$ is uniformly bounded away from zero and hence it is not possible that (10) holds for all N, a contradiction. ■

With N^* well-defined for small r, we can state formally the relative benefits of the two types of subsidy as a function of r :

Proposition 2 *For sufficiently low r, the planner obtains a higher social return from (one-time, unexpected) subsidy of pioneering innovation than a subsidy of a research into one of the existing submarkets.*

Proof. As we argued above, the total social return to one-time creation of one cost improvement is simply a constant flow of one additional "step" in consumer surplus: CS_Δ/r. When instead a new submarket is created, customers do not gain any additional surplus immediately (as we assumed that the product will be sold by a monopolist while the outside option is sold competitively) but they will enjoy a 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 from introduction of a new submarket, 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

$$\frac{\hat{x}_e(N^*+1) - x(N^*)}{r} \tag{11}$$

As $r \to 0$, condition (8) converges to:

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

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 $\tilde{c}(x)$, a new submarket creates a higher social benefit (because it induces a larger increase of x than if the costs were given by $c(x)$,according to (11)). But, by contrast, an additional submarket 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 submarket is greater under $\tilde{c}(x)$.

This logic is summarized in the following proposition.

Proposition 3 *Let $\tilde{c}(x) = c(x)$ for $x \leq x(N^*)$, but let $\tilde{c}(x) < c(x)$ for $x > x(N^*)$. Then the relative social return to a pioneering research vs. non-pioneering research is greater under $\tilde{c}(x)$.*

One interpretation of new submarkets/pioneering versus cost reductions/product improvements in existing submarkets 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 pioneering research is small, but the social benefit from inventing additional submarkets is large. Both effects are driven by the fact that a flat $c(x)$ leads to a big impact of a new submarket on equilibrium research intensity.

One way to understand further the comparative static is to think of a flat $c(x)$ function as akin to innovations being easily imitated. From the perspective of the innovator, this is the worst case for innovation; for the planner, it is the best case, since it generates both innovations and potentially valuable imitations. The private and social benefits go in opposite directions as $c(x)$ flattens.

Another interpretation is related to the discussion of Romer (2000), who asks whether the government should subsidize supply or demand in researchers. If we interpret subsidizing the supply of researchers as undirected, then it may be better to choose a particular type of innovation (for instance new submarkets) and subsidize hiring researchers working on that dimension, or indirectly subsidize the demand side of that market to increase the private returns.

6 Conclusion

We have introduced a model of the trade-off faced by an innovator who might pioneer a new submarket or focus on an existing one. The model relies on relatively few key ingredients. Pioneering innovations are more costly than non-pioneering innovations. Innovation requires a scarce resource. The model not only delivers facts about innovation dynamics, but also suggests a sense in which a planner might favor one over another, with a patient planner favoring encouragement of pioneering innovations.

The study of innovation over the product cycle has brought to light a variety of regularities. We have introduced a model capable of explaining the regularities related to the rise of the product cycle. Although we had to make some particular modeling choices, the results rely on the supply of research being upward-sloping which implies that innovation should be rising with the number of submarkets, but less than proportionally. We showed that this key prediction of our model is consistent with a well-known product cycle data set. We believe that one can expand our model in several dimensions and keep the force behind market pioneering that we highlight.

References

[1] Agarwal, Rajshree (1998). "Evolutionary trends of industry variables," *International Journal of Industrial Organization*, Vol.16 (4), pp. 511-525.

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

[3] Cohen, Wesley M. and Steven Klepper (1996). "Firm Size and the Nature of Innovation within Industries: The Case of Process and Product R&D," *The Review of Economics and Statistics*, vol. 78(2), pp. 232-243.

[4] Gort, Michael and Steven Klepper (1982). "Time Paths in the Diffusion of Product Innovations," *The Economic Journal,* vol. 92(367), pp.630-653.

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

[6] Jovanovic, Boyan and Rafael Rob (1990). "Long Waves and Short Waves: Growth Through Intensive and Extensive Search," *Econometrica*, Vol. 58(6), pp. 1391-1409.

[7] Klepper, Steven (1996). "Entry, Exit, Growth, and Innovation over the Product Life Cycle," *American Economic Review*, vol. 86(3), pages 562-583.

[8] Klepper, Steven and Peter Thompson (2006), "Submarkets and the Evolution of Market Structure," *RAND Journal of Economics*,vol. 37(4) pp. 861-886

[9] Klette, Tor Jakob and Samuel Kortum (2004). "Innovating Firms and Aggregate Innovation," *Journal of Political Economy*, vol. 112(5), pp. 986-1018.

[10] Lieberman, M. and D. Montgomery (1988). First-Mover Advantages. *Strategic Management Journal*. Volume 9, Issue S1 , Pages 41 - 58.

[11] Robinson, WT, G Kalyanaram, GL Urban (1994) "First-mover advantages from pioneering new markets: A survey of empirical evidence" *Review of Industrial Organization,* Vol 9(1), pp. 1-23.

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

[13] 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.

[14] Sutton, John (1998). *Technology and Market Structure*. Cambridge, Mass.: MIT Press.

[15] Utterback, James and William Abernathy (1975). "A dynamic model of process and product innovation," *Omega*, vol. 3(6), pp.639.