The Direction of Innovation over the Product Life Cycle*

Matthew Mitchell† and Andrzej Skrzypacz‡

January 14, 2007

Abstract

Papers such as Gort and Klepper (1982) and Utterback and Abernathy (1975) describe the evolution of products from their birth through maturity, termed the product life cycle. Both the number of firms and the number of patents grows as the product reaches its peak. There is also a clear shift in the composition of innovation. Utterback and Abernathy (1975) have pointed out the change from product innovations to process innovation. Gort and Klepper (1982) document a shift from what they call major to minor innovations.

In this paper we introduce a model of directed technological change that can account for these changes. Innovative effort can be used either for process or product innovation. We argue that, in the model, product innovations are the ones more likely to be classified as major. Our model predicts a slowdown in these innovations in absolute terms, along with a rise in process innovations. We investigate the product cycle data and show other important features that are consistent with the model.

*We thank Philippe Aghion, Paul Romer, and participants at seminars at the Minnesota Fed, Duke, Iowa State and the 4th Duke/Texas/Northwestern IO theory conference for suggestions.

†University of Iowa.

‡Stanford GSB.

1 Introduction

In this paper we develop a model that captures some of the salient features of how the quantity and direction of innovation changes over a product's rise to maturity. We then investigate an updated version of the Gort and Klepper data to show that the new model predictions are broadly consistent with other key features in the data.

The empirical literature on innovation stresses some important features of innovation dynamics of a product category/industry. In particular, both the quantity and nature/direction of innovation change over the product life cycle. These changes have been documented in papers such as Gort and Klepper (1982) and Utterback and Abernathy (1975). In particular, innovations are not only measured in terms of quantity, but are also categorized by type which describes the direction of innovation. Gort and Klepper (1982) show that major innovations fall as a share of total innovations as a product matures. Utterback and Abernathy (1975) show that innovations move from product to process innovations, with product innovations steadily falling and process innovations rising to the peak of the industry life cycle. Utterback and Abernathy's depiction of innovation over the product cycle is often described by the following picture:[1]

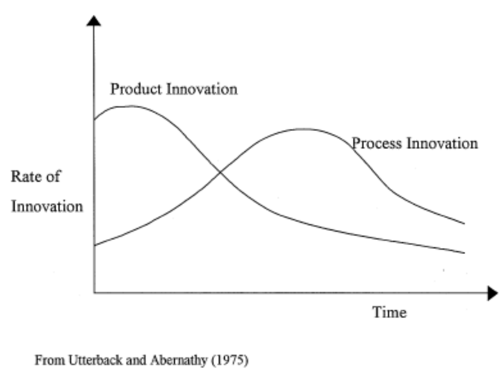

Figure 1: Product and Process Innovation

We focus our attention on the period of increasing innovation in the product cycle. Our first goal is to construct a model of innovations of different types. Some innovations are cost reductions to existing technologies, and

[1] Here taken from Adner and Levinthal (2001).

transfer market share from one firm to another. We associate these with process improvements. Other innovations increase the scope of a technology into a fundamentally new area. These innovations are major in the model, in the sense that they impact industry wide innovation and firm counts, and are clearly product innovations, as they bring about a new stream of profits (we refer to them as inventions of new versions of a product while the minor innovations are simply production cost reductions for one of the existing products).

We study our model from birth to steady state. We show that research intensity follows a simple implicit equation, which implies rising levels of arrival of innovations, as the data suggests. However, the composition is changing. Over time major innovations (new products) decline, and minor (process) innovations rise. In steady state, there are only process innovations. This is a stark depiction of an industry where innovations are increasingly minor or process, as the evidence suggests.[2]

A key feature of the model is that the two types of innovations 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 since 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, consistent with the evidence. Instead, we focus on the fact that product innovations tend to increase the total market more so than process innovations (as in Cohen and Klepper (1996)).

Our model echoes the ideas 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. We show how one can relable our model by referring to "versions" as actually different components of the design. When viewed this way, our steady state describes an endogenously determined "dominant design," as Utterback and Abernathy (1975) describe. Once the industry reaches the steady state, no more “new designs” arrive (though firms con-

[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.

tinue getting more and more cost efficient in producing the existing design). Our model endogenizes the timing of the switch from product to process improvements, and shows that the decline in product innovation may not be tied to a technological change, but rather to a gradual decline in the incentives of firms to generate product improvements.

Another interpretation of the model is that major innovations can be thought of as the development of new submarkets. Klepper and Thompson (2006) study how *exogenous* submarket development can explain some important facts about firm dynamics. Our model is as a step towards the development of a model of industry dynamics with *endogenous* submarket formation. Our model is simplified on the dimension of competition within a submarket (as one firm takes all of a given submarket), but the model could be extended to allow for multiple firms to make profits from a submarket along the lines of Klepper and Thompson (2006).

As mentioned above, our model has the feature that major innovations raise the number of firms in the industry. In response, the arrival rate of innovations increase, but at a slower rate than the increase in firms. We show that this is a clear feature of the product cycle data. Moreover, the data on entry and exit is broadly consistent with the (new) 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 (1996). As the product develops, market shares become more stable.

We study primarily the case where all submarkets are equally profitable. Firms research in order to become a leader in these symmetric submarkets. Since the reward (per unit of time) is the same across submarkets, and new submarkets are more expensive to research, one might think that there is no incentive for competitive researchers to develop new submarkets. However, there are two important features that arise when developing a new submarket. Total research intensity rises. However, because of the increasing marginal cost of research, research per submarket decreases. This implies that developing a new submarket slows arrival of market-stealing improvements; this is where the incentive to develop new submarket arises.[3]

Other models have addressed the life cycle impact on product and process innovation. Filson (2002) studies a version of the Jovanovic and MacDonald's

[3] We also consider the case that initial innovations (the first in a submarket) are more profitable than subsequent innovations and get similar results.

(1994) model augmented to include product and process innovation. Rather than replicating facts, he uses the model to back out implied product and process innovation shares for the computer industry. Saha (2006) studies the role of heterogenous consumer tastes in driving the relative shares of product and process innovation.

Our approach to modeling innovation draws heavily on the endogenous growth literature, such as Grossman and Helpman (1991) and Aghion and Howitt (1992) (see for example Kortum and Klette (2004)). 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 product and process innovation, and studies out of steady state dynamics, in order to address the product life cycle evidence.

Other papers in the growth literature such as Peretto (1999) have included endogenous variety and innovation. There is only process innovation and the incentive to engage in innovation depends on the endogenous number of firms that are in competition. One can think of the number of firms as a measure of product diversity. Paretto focuses on the role of competition on the rate of aggregate innovation, and shows that more firms may lower the rate of process innovation.

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 that new empirical results from the product cycle data.

Although our model has the feature (consistent with the product life cycle evidence) that more firms are associated with greater levels of innovation, the model does not necessarily have the feature that more competition is necessarily better for innovation. To show this, in section **??** we briefly describe how a monopolist would operate in this world, and show that the rate of innovation for a monopolist can be either greater or less than the case of research competition that we study. A more general analysis of market structure and innovation in an environment with two directions of innovation, like this one, is contained in Mitchell and Skrzypacz (2006).

Our general equilibrium model allows policy experiments. In 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 additional research into product components over process improvements.

2 The Model

2.1 Products

At any given instant there are N versions of the product for sale. The lowest cost producer for a given version captures the entire market for that version and makes profits π per instant (so that total industry profits are $N\pi$).

This simple model of profits by version or submarket is analogous to assumptions in Kortum and Klette (2005) and Klepper and Thompson(2006). 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 is also motivates Kortum and Klette (2005)). We also discuss some relaxations of this assumption, many of which are straightforward. The important feature of the assumption is that it is meant to not bias research toward one type of research or another (i.e. profits from new versions and improvements to current versions have identical profit flows), or to 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 simple assumed differences in profitability.

2.2 Innovation

Innovation comes through research. Research can be done on either developing new versions (product innovation) or on improvements to production costs of existing versions (process innovations). 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 process improvements to existing versions and x_f for new (frontier) versions.

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 version, and a cost flow of $\theta + \eta$ for a frontier version. The inclusion of $\eta > 0$ means that new versions 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. Researchers may enter freely into research at any instant and choose which product to research on. 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 version, or $F(\pi/r - \eta) > 0$.

The critical feature of our model that ties different versions together is that they draw researchers from a common pool of scarce talent. As a result, innovative effort on one version has an impact on the marginal cost of innovation for all of the versions using the common factor. This equilibrium effect is what gives rise to all of the results about the dynamics across different versions that we develop below. While we use the heterogeneous types of researchers to generate the scarce resource, anything that generated cost increasing in total research intensity would generate the same results.

We assume that when a researcher finds an innovation he forms a firm and markets it (becoming a producer). Researchers (and firms hiring them) maximize the expected sum of discounted profits net of research costs and use a common discount rate r.

3 Innovation Dynamics over the Product Cycle

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

3.1 Analytic Results

Given the profits π from the Bertrand pricing game, consider the reduced-form game with competition among researchers. We focus on 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 product version yields the same profit flows regardless of the identity of the version or the current cost level of the version, in such an equilibrium the strategies of the researchers (and expected profits of the producers) depend only on the current number of product versions, N.

We start by introducing some notation that allows 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 follows a cutoff rule:

if type θ does research, all types $\theta' < \theta$ do as well (although they can do it on some other version). 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 versions (unless $x_e(N) = 0$), researching a frontier version (unless $x_f(N) = 0$) and opting out of research.[4]

Define $c(x) = F^{-1}(x)$, the cost of the marginal researcher in the existing version search, so that $\bar{\theta}(N) = c(x(N))$. By (1), $c(x)$ is increasing. Let $\rho(N) = \frac{x_e(N)}{x(N)}$ be the probability that conditional on new innovation arriving, it is for an existing version. 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 it depends on N because so does $x(N)$.

Let $V(N)$ be the value of an incumbent firm just as a new innovation arrives, but without knowing what version it is for. It is defined recursively by

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

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

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

[4] I the whole paper the equilibria we construct pin down only the share of active researchers working on the particular tasks but not the allocation of particular types. If the difference in costs of the two types of research depended on θ we would obtain this additional characterization.

$V_e(N)/r$ 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).

The following proposition summarizes the characterization of the equilibrium.

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 it 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 innovators can ever be willing to pay the additional cost to develop a new version. 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 $N^* > 1$. The reason is that, due to the increasing cost of researchers, research intensity on existing versions rises less than proportionally to the number of existing versions. Therefore $x_e(N)/N$ - the amount of research on improvements per version - tends to decline in N, and the value of having a marketable version 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, $x_e(N)/N$ drops by very little for large N. Eventually the increase in value of an version gets small and is insufficient to draw research in new versions. The steady state N^* is reached, at which point there is only process innovation.

Second, anytime there is research on new versions, there must also be research on improving existing versions. If there were only research on new versions for a given N, then developing an improvement would make more profits than a new version: it would earn profits until the next new version, at which point the continuation value would be as much as the new version 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 expected value of a leading edge version.

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 products (the first version) 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).

We now turn to a construction of the equilibrium. 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 (3) implies $V_e(N+1) = V_f(N)$. Combining it with the free entry conditions we get:

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

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 (5) and monotonicity of $c(x)$.

Regarding $V(N)$, from the free-entry conditions (4) 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 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, 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 $V_e(N) \in (0, \pi)$. ■

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:

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

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

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 not profitable to search for a new version at N^*, even if no new versions were researched afterwards. 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 the whole 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. 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 (6) 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 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 (5) we get:

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

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

Finally, to show $N^* = \arg\min_{N\in\{1,2,\ldots N^*\}} x(N)/N$, first consider $x_e(N)$. If $x_e(N)/N$ were minimized for some $N < N^*$, then $V(N)$ would be maximized at N, and therefore could not rise from N to N+1,so $N^* = \arg\min_{N\in\{1,2,\ldots N^*\}} x_e(N)/N$. Since $x_f(N^*) = 0$, clearly $N^* = \arg\min_{N\in\{1,2,\ldots N^*\}} x_f(N)/N$, and as a result $N^* = \arg\min_{N\in\{1,2,\ldots N^*\}} x(N)/N$. That establishes all the claims in the Proposition 1.

3.2 Numerical Example

In order to see how the competitive equilibrium evolves in a numerical example, we assume that F is linear, $F(\theta) = \left(\theta - \theta_l\right)/a$, and so $c\left(x\right) = ax + \theta_l$.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% - see Section 5 for details). Then $N^* = 9$. In the steady-state $x\left(N^*\right) \approx 9.66$, which is also the average number of innovations 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 versions, and the third line is investment in existing versions.

Note that the difference between the extra cost η of a new version 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 versions are large.

3.3 The Evolution of a Dominant Design

Utterback and Abernathy (1975) argue that the shift from product to process innovation is driven by the development of a "dominant design." This is consistent with the model we present. To see this, reinterpret "versions" i

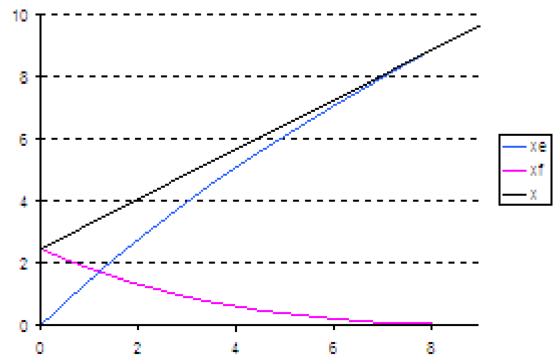

Figure 2: Innovation Dynamics

as "components"[5] of a common design. For simplicity, the model has the interpretation that components are sold separately and aggregated by the consumer; it would be identical to have a constant returns competitive firm that aggregated various components into the final product.

Then, in a sense, the model gives and endogenous dominant design; namely 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. In other words, dominant design is being importantly impacted by the resources the industry has to draw on.

4 Evidence

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 1 (birth) to stage 3 (the peak level of firms).

Gort and Klepper (1982) document that the rise in firms is met with a

[5] A term from Utterback and Abernathy (1975).

rise in patenting. Interpreting innovations in the model as patents, one can see that pattern emerge, since each version is sold by one firm. Utterback 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.[6] 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 versions 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 versions are more researched early on, and less researched later on, and therefore innovation switches from product to process and major to minor.

4.2 Additional Evidence

One key feature of the model is that innovation is tightly linked to the number of firms in the industry. To address this issue empirically, 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) accord-

[6] Note that existing models such as Klepper (1996) focus on the relative amounts of the two types, rather than the absolute quantities.

ing 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)

The data is consistent with a key feature of the model. Expanding the total number of firms (N) does increase innovation ($x(N)$), but lowers innovation per firm ($x(N)/N$, which is minimized at the maximum number of firms N^*), since the coefficient on ln(firms) is between zero and one.[7]

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 or done by new ones. To take one tangible example, assume that inventors market their own inventions, so every innovation leads to a new firm. In other words, the model associates the entry with x. The imposition of this strong assumption implies that the model may have a much harder time explaining the entry and exit data.

Since x increases less than proportionally to firm numbers, entry should, too, under this interpretation. We duplicate, then, the prior estimation, estimating that the log arrival rate of entrants is a function of ln(firms) and the fixed effects.

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

This is once again consistent with less-than-proportional growth that the model predicts.

Exit occurs when a firm is overtaken on a given version; therefore exit is associated with x_e. This too increases less than proportionally than 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

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

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. We expect, then, that replacing exits with entrants in the previous regression should yield a higher coefficient, but still less than one.

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

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 zero. 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.

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 version 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.

5 Policy Implication: Subsidizing Innovation

5.1 A Model with Preferences

In order to study policy implications, it is necessary to fully describe the demand side of the economy. We assume that different firms j produce version i at costs $c_i^j \leq 1$. Firms are ordered in decreasing order of quality. For a given version i, the representative consumer consumes d_i^j units of cost level j. This leads to e_i efficiency units of the version, where[8]

$$d_i = \sum_j d_i^j$$

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

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

where p_i^j is the price paid for cost level j on version i.

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

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

To deliver the constant profits per ladder assumed above, we parameterize the utility function to be logarithmic:

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

Log utility is common in the growth literature because it generates stationarity with constant percentage-decreases in cost levels. Here it also delivers us, in a stark way, what we want to assume about new versions versus cost reductions. Demands are independent across different versions and the representative consumer spends a constant share of his income on every version i. In particular, if the representative consumer buys version c_i^j at price p, his demand is $d_i^j = 1/p$.

Costs fall on a ladder with rungs of size $\lambda > 1$; i.e., for the jth> 0 cost level on ladder i, $c_i^j = \lambda c_i^{j+1}$. The first version has cost $q_i^1 = 1/\lambda$. For each version, if nothing has been invented, an outside option is provided competitively at marginal cost 1. One can interpret this as the next best alternative product that might substitute for version i.

Following the endogenous growth literature such as Aghion and Howitt (1992) and Grossman and Helpman (1991), suppose that each cost level is

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

monopolized, perhaps due to a patent or a trade secret. There is a Bertrand competition between firms within a ladder. Non-leading-edge 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 $p_i^1 = 1$.[9] Given demand x_i on ladder i, profit flows for the leader are $d_i(p_i - c_i)$, where under Assumption 1, $d_i = 1/p_i$ and profits are $\frac{\lambda-1}{\lambda} \equiv \pi$.

5.2 Social Benefit of Different Types of Innovation

We ask whether the competitive market structure provides sufficient incentives to innovate. Here we add to the standard discussion a new dimension: if the government should subsidize innovation, should it subsidize any particular direction? 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 now ask how the social returns compare between subsidizing research on new or existing versions.

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 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, where CS_Δ is the extra flow consumer surplus generated by the extra quality improvement. Typically, 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; additionally, given our demand structure, the increase in flow of consumer surplus is higher than the profit flow: $CS_\Delta = \ln\lambda > \frac{\lambda-1}{\lambda} = \pi$.[10] Therefore a subsidy to research increases total welfare. Any such policy would have to decide whether to favor process improvements of existing versions or development of new versions.

To model this, we simply ask what the planner's payoff would be to one arrival of each type of innovation. For a process improvement, the benefit is just $CS_\Delta/r = (\ln\lambda)/r$. Innovation in a new version would have two effects. First, it creates additional profit flow, hence a return to firms of π/r

[9] As is typical, this equilibrium requires assuming that consumers choose the higher quality product when indifferent between two quality levels.

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

(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 Section **??** that $\hat{x}_e(N)$ is defined as the equilibrium research intensity if no further frontier applications are expected in the future).[11] Since 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 process improvement is higher than from inventing a new version. Therefore a sufficiently patient planner would prefer subsidizing new versions.

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

$$\frac{\pi}{r + \hat{x}_e(N)/N} - c(\hat{x}_e(N)) = 0 \tag{9}$$

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 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 frontier application is feasible even for small r.

Lemma 6 $\lim_{r\to 0} N^* < \infty$.

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

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

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

$$c(\hat{x}_e(N))\,\hat{x}_e(N) = N\pi \tag{10}$$

[11] 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.

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 $c'(x)$ is bounded from above. Suppose that for every N,

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

(the opposite inequality is our condition for N^*, see discussion before Lemma 4). Then, in (10) 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. ■

With N^* well-defined for small r, we can state formally the comparative static in r :

Proposition 2 *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 a new version than in a cost reduction.*

Proof. As we argued above, the total social return to one-time creation of one additional existing application is simply a constant flow of one additional "step" in 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 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

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

As $r \to 0$, condition (9) 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 $\tilde{c}(x)$, a new version 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 version 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 version 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 planner has a greater preference for product innovations under $\tilde{c}(x)$*

One interpretation of new versions versus process 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 frontier invention is large. Both are driven by the fact that a flat $c(x)$ leads to a big impact of a new version on equilibrium research intensity.

Another interpretation is related to the discussion of Romer (2000), who asks whether the government should subsidize supply or demand in researchers. Here, subsidizing the supply of researchers is undirected; it may be better to choose a particular type of innovation (for instance new versions) and subsidize hiring researchers working on that dimension.

6 Conclusion

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. The model relies on innovation rising with firm numbers, but less than proportionally. We show that this is indeed true in a well known product cycle data set. Moreover, the behavior of entry and exit is consistent with the model in several ways.

The model relies on relatively few key ingredients. Product innovations are more costly than process 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 product innovations.

References

[1] Adner, Ron and Daniel Levinthal (2001), "Demand Heterogeneity and Technology Evolution: Implications for Product and Process Innovation," *Management Science* Vol. 47, No. 5, May 2001 pp. 611–628

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

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

[4] 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.

[5] Filson, Darren (2002). "Product and process innovations in the life cycle of an industry," *Journal of Economic Behavior & Organization*, vol. 49(1), pp. 97-112.

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

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

[8] Jovanovic, Boyan and Glenn M. MacDonald (1994), "The Life Cycle of a Competitive Industry", *Journal of Political Economy*, volume 102, pp. 322.

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

[10] Klepper, Steven and Peter Thompson (2006), "Submarkets and the Evolution of Market Structure," *RAND Journal of Economics*, forthcoming.

[11] Milling, Peter and Joachim Stumpfe (2000), "Product and Process Innovation: A System Dynamics-Based Analysis of the Interdependencies," Proceedings of the 18th International Conference of the System Dynamics Society..

[12] Mitchell, Matthew and Andrzej Skrzypacz (2006), "Market Structure and the Direction of Technological Change," mimeo.

[13] Peretto, Pietro (1996). "Sunk Costs, Market Structure, and Growth," *International Economic Review*, Vol. 37(4), pp. 895-923.

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

[15] 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.

[16] Saha, Souresh (2006). "Consumer Preferences and Product and Process R&D," forthcoming *RAND Journal of Economics*.

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