

NBER WORKING PAPER SERIES

WAVES OF CREATIVE DESTRUCTION:
CUSTOMER BASES AND THE
DYNAMICS OF INNOVATION

Jeremy C. Stein

Working Paper No. 4782

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 1994

This research is supported by MIT's International Financial Services Research Center and the National Science Foundation. Thanks to Maureen O'Donnell for help in preparing the manuscript, and to Owen Lamont for research assistance. I am also grateful for the comments and suggestions of Daron Acemoglu, Ricardo Caballero, Ken Froot, Rebecca Henderson, Julio Rotemberg, and Alwyn Young. This paper is part of NBER's research program in Corporate Finance. Any opinions expressed are those of the author and not those of the National Bureau of Economic Research.

WAVES OF CREATIVE DESTRUCTION:
CUSTOMER BASES AND THE
DYNAMICS OF INNOVATION

ABSTRACT

This paper develops a model of repeated innovation with knowledge spillovers. The model's novel feature is that firms compete on two dimensions: 1) product quality or cost, where one firm's innovation ultimately spills over to other firms; and 2) distribution costs, where there are no spillovers across firms and where incumbent firms' existing customer bases give them a competitive advantage over would-be entrants. Customer bases have two important consequences: 1) they can in some circumstances dramatically reduce the long-run average level of innovation; 2) they lead to endogenous bunching, or waves, in innovative activity.

Jeremy C. Stein
E52-448, Sloan School of Management
Massachusetts Institute of Technology
50 Memorial Drive
Cambridge, MA 02142-1347
and NBER

1. Introduction

Recent work in the literature on innovation and growth has emphasized the process that Schumpeter (1942) labeled "creative destruction." In the models of Aghion and Howitt (1992), Grossman and Helpman (1991a, 1991b), Segerstrom (1991), Segerstrom, Anant and Dinopoulos (1990), and Caballero and Jaffee (1993), new, higher-quality products are introduced by new firms, thereby displacing incumbent firms. The introduction of these new products in turn sets the stage for yet another round of innovation, entry, and displacement, because of the existence of knowledge spillovers--once a new product is introduced, future generations of innovators can learn from it and improve upon it.

The image that comes through in these papers is of a very fluid corporate sector, with new companies continually pushing aside existing ones. Indeed, this fluidity is seen as an essential element in the growth process: it is only through the destruction of existing firms' market shares and profits that new, better products--and the embedded knowledge that accompanies them--come into being.

Yet casual empiricism suggests that the corporate landscape may be much less fluid than these models envision. Many companies maintain their market shares in given product areas for a very long time. Often they manage to do so even as the technological environment in which they operate changes dramatically, and the nature of the products they sell change along with it.

There are two broad reasons why established firms might be expected to survive in a rapidly changing environment. First, and most simply, established firms may just have a comparative advantage in innovation. This would imply that established firms naturally tend to

stay "on the cutting edge" of new products and new technology, since they actually pioneer the majority of improvements. If this is the case, the survival of such firms would be no surprise--they survive because they are either the highest quality, or most efficient producers at any point in time.¹

A second possibility is that established firms are not always on the cutting edge of new products and technology, but rather, they have a number of other competitive weapons at their disposal that allow them to fend off more innovative newcomers and to thereby forestall the process of creative destruction. For example, long-established firms are likely to have a base of loyal customers and well-developed distribution networks that allow them to market their products more efficiently. Or they may have long-standing relationships with capital suppliers that allow them to finance their plant and inventories more cheaply. In either case, the net result is that the established firms may be able to compete successfully against newcomers, even if the newcomers bring with them substantial technological improvements.

The latter explanation of corporate survival in the face of outside product innovation--which I will call the "customer base" story--fits well with the literature on competitive strategy (e.g., Porter, 1980). Work in this area repeatedly emphasizes that innovative, high-quality

¹The literature has identified a number of factors that may help to foster innovation inside already-established firms. Gilbert and Newbery (1982) argue that incumbents are likely to have a greater strategic incentive to invest in innovative activity than potential entrants when the innovations in question are incremental in nature. Moreover, large established firms may have certain organizational advantages in doing incremental innovation, as stressed by Schumpeter (1942) -- e.g., preferential access to information or trained scientific personnel. Combining these strategic and organizational arguments, Henderson (1993) concludes that: "established firms are likely to dominate incremental innovation, while entrants are likely to dominate radical innovation" (p. 252). Thus if most innovation proceeds at an incremental pace, one might expect that established firms would naturally be able to stay on the cutting edge.

products are but one route to "competitive advantage", and that loyal customers and a strong distribution network can also be crucial links in the "value chain". The customer base story also seems to jibe with how some managers of established firms describe their own competitive strengths. For example, Frank Perna, CEO of Magna Tek, a \$1.3 billion manufacturer of electrical equipment, argues that: "I think the core competence of our company is really its distribution channels--far more so than its products." (emphasis added)²

This paper takes the customer-base story of corporate survival as a starting point, and investigates its implications for the dynamics of innovation. That is, I adopt an entrepreneurial slant, taking it as given that many good ideas for product improvement are naturally generated outside the established firm sector.³ I then ask: if established firms have a countervailing advantage over newcomers in the form of an existing customer base, how is innovative activity by these newcomers affected?

More specifically, I construct a dynamic model with the following features: 1) incumbent firms' customer bases give them a "distribution cost" advantage over potential entrants, an advantage which gets more pronounced the longer the incumbents are able to survive; 2) potential entrants on the other hand, may have access to a new superior production technology; 3) technological innovations have spillovers--once an entrant develops a new technology and introduces its product to the marketplace, future generations of innovators can learn from it, and improve upon it. The model has implications for how customer bases affect both the long-run

²This quote is taken from "Continental Bank Roundtable on Global Competition in the '90's", Journal of Applied Corporate Finance, Spring 1993 (p. 40).

³As the references in footnote 1 suggest, this entrepreneurial premise may be most relevant when the ideas in question represent radical, rather than incremental innovations.

average level of innovative activity, as well as the temporal pattern of innovation.

In its focus on the long-run average level of innovation, the model is similar in spirit to the recent work in the growth literature that also analyzes repeated innovations with intertemporal knowledge spillovers.⁴ A common theme of this work is that, due to the spillovers, the resources devoted to innovation, and hence the rate of growth of the economy, can in some circumstances be less than socially optimal. That basic under-innovation effect is at work in this model too, although it can be dramatically amplified by the presence of customer bases. Loosely speaking, in my model firms compete on two dimensions: a "high spillover" dimension--i.e., product quality or cost, where one firm's innovation is ultimately passed on to other firms--and a "zero spillover" dimension--distribution costs, where an incumbent's advantage never spills over to other firms. In some circumstances, an incumbent can survive a very long time (or even forever) by being strong in the zero spillover dimension. While this is privately optimal for the incumbent, it can be very costly for society, as generation after generation of potentially high-spillover innovators are warded off.

Where the model departs more sharply from the papers mentioned above is in its emphasis on the timing of innovation. Most of these other works have the feature that innovation proceeds at a steady rate.⁵ To take just one example, in Aghion and Howitt (1992)--which apart from the customer bases resembles this model quite closely--there always exists a

⁴In addition to the papers cited above, see also Romer (1990).

⁵See, however, Shleifer (1986) for a model where innovations are implemented in bunches. The mechanism in Shleifer's model is quite different from that emphasized here--bunching is driven by aggregate demand spillovers that lead all firms to coordinate their implementation in periods when demand is high. In contrast, demand plays absolutely no role in my model; everything is driven by the relative competitive strengths of the rival firms.

steady-state equilibrium in which innovative activity is constant over time. In contrast, a central feature of the model in this paper is that when a new firm successfully displaces an incumbent in any given period, this has a positive externality on future generations of potential entrants--it makes it easier for them to gain access to the market. This externality, which I call the "shakeup" externality, arises because when a new firm succeeds, it breaks the incumbent's stranglehold on the customer base. Thus the market is now "up for grabs", which tilts the playing field more toward technologically strong newcomers, and away from established firms.

The consequence of this shakeup externality is that even if the underlying research and development technology is stable over time, innovations will tend to occur in waves. That is, if there is an innovation today, the odds of another innovation tomorrow may be substantially higher. Thus on the one hand, there may be long periods of stagnation, in which no newcomers enter the market. But these periods of stagnation can suddenly give way to rapid bursts of innovative activity. This wave-like aspect of innovation was also stressed by Schumpeter (1936):

"...new combinations are not, as one would expect according to general principles of probability, evenly distributed through time...but appear, if at all, discontinuously in groups or swarms....Why do entrepreneurs appear, not continuously, that is singly in every appropriately chosen interval, but in clusters? Exclusively because the appearance of one or a few entrepreneurs facilitates the appearance of others, and these the appearance of more, in ever-increasing numbers." (pp. 223, 228)

As this passage indicates, Schumpeter also believed that the entry of any one entrepreneur imparted a positive externality to future would-be entrepreneurs, and that this was the source of waves, or swarms, in innovative activity. However, he was somewhat less clear as to exactly what the sources of this externality were. One major contribution of this paper lies in

delineating more precisely the mechanisms that give rise to this positive externality.

The remainder of the paper is organized as follows. The basic model is developed in Section 2. For the purposes of this section, it is just assumed that inventions arrive exogenously to potential entrants, who must then decide whether or not to spend the money to further develop these inventions and enter into product market competition against incumbent firms. In Section 3, I extend the model to allow the probability of invention to be an endogenous function of the resources devoted to research. As will be seen, this extension further strengthens the positive "shakeup" externality effect, and leads to an even more pronounced bunching of innovative activity. Section 4 discusses some further extensions of the basic framework, and Section 5 concludes.

2. The Model

2.1 Types of Firms

The model is an infinite horizon one. In each period t , there is a single "incumbent" firm, defined as the firm that was active in period $t-1$. With probability p , there is also a "potential entrant" firm, that invents a superior new technology, and that may choose to further develop the technology and challenge the incumbent for the market in period t . For the time being, p is taken as exogenous; later it will be made an endogenous function of the resources devoted to research.

If the potential entrant is successful in its challenge, it will become the new incumbent in period $t+1$. Only one firm is ever active in any given period--an incumbent always has 100% market share. As will be seen, this feature emerges from the assumptions that are made below

about firms' cost structures and the nature of product market competition.

In addition to incumbents and potential entrants, there is a third class of firms, called "copycats". As will be seen below, copycats are always strictly less efficient than incumbents. Thus in equilibrium, they never capture any market share. Nonetheless, they play an important rôle: their costs serve to tie down market prices in periods when the incumbent firm is not challenged by a potential entrant, and otherwise would be an uncontested monopolist.

2.1.A Incumbents' cost structures

The incumbent, the potential entrant and the copycats all have different cost structures; I begin by describing that of the incumbent. The incumbent has no fixed costs. Its marginal costs are of two kinds: costs of "production" and costs of "distribution". The marginal cost of production for an incumbent at time t is denoted by C_t^i . For any given incumbent, production costs do not change over time. That is, if the same firm is the incumbent in periods $t-1$ and t , then $C_t^i = C_{t-1}^i$. Thus incumbents make no technological progress in terms of their production costs.⁶

The marginal cost of distribution for an incumbent at time t is denoted by D_t^i . The costs of distribution are assumed to take the following form:

⁶The assumption that incumbents make no technological progress is obviously extreme and unrealistic. However, it is made for expositional convenience only. It would be trivial to extend the model to consider the case where incumbents make steady technological progress on their own. In either case, the principal question would remain: when outsiders develop better technologies than incumbents, what is the likelihood that these better technologies are implemented?

$$\begin{aligned}
 D_t^i &= dC_t^i \beta^{A_t^i} \text{ for } A_t^i \leq \bar{A}; \\
 &= dC_t^i \beta^{\bar{A}} \text{ for } A_t^i > \bar{A}
 \end{aligned}
 \tag{1}$$

where $\beta < 1$ and A_t^i is the "age" of the incumbent at time t --that is, A_t^i represents the number of consecutive periods (prior to t) over which the incumbent has been active.

There are two things to note about the form of the incumbent's distribution costs in (1). The first is that they are proportional to the incumbent's production costs. This assumption is not really critical for the basic point to be made. However, it makes it possible to couch the analysis in a steady-state equilibrium framework. Indeed, in order to facilitate the steady-state approach, all costs incurred by any firm in the model will be proportional to that firm's contemporaneous production costs. As will become clear shortly, this simplification allows one to derive a set of time-invariant decision rules for firms--rules that do not depend on, for example, a firm's current ratio of production costs to distribution costs.

The second, more important assumption embedded in (1) is that an incumbent firm's distribution costs are a decreasing function of the length of time it has been an incumbent. In particular, distribution costs decline geometrically over the first \bar{A} periods of incumbency and then remain flat after that. While the specific functional form is not critical, it is crucial that these costs do decline over some range. It is not enough simply to assume--as would be implied by a switching cost model--that an incumbent has lower distribution costs than an entrant, but that all incumbents have the same costs regardless of age. It must be the case that the longer an incumbent has been around, the stronger their competitive advantage along this dimension.

How should this assumption--that distribution costs decline with an incumbent firm's

age--be interpreted? One possibility is to think of falling distribution costs as reflecting the accumulated knowledge about customers that is a byproduct of an ongoing firm-customer relationship.⁷ For example, the better a computer manufacturer knows its corporate customers, the more efficiently it can market to them, customize the computers it sells them, provide continuing service that meets their specific needs, etc.

At a general level, what is important about the distribution costs is that they represent a competitive attribute where there is effectively the equivalent of learning-by-doing at the firm level, but no spillovers across firms. As will become clear, it is this lack of spillovers across firms that differentiates distribution costs from production costs in the model. When a new production technology is pioneered by one firm, other firms are eventually able to learn from it and improve upon it. In contrast, firms are unable to inherit the distribution cost advantages created by their predecessors--if a new firm enters the market, it must effectively start from scratch in building an efficient distribution system.

As this discussion suggests, the words "distribution costs" are really intended as a metaphor for something much broader--any source of competitive advantage that tends to become more pronounced with an incumbent firm's age and that does not spill over to other firms. For example, the work of Diamond (1989) implies that access to financing may be another important competitive attribute which meets this description. In Diamond's model, the longer a firm survives, the stronger is its "reputation" in the debt market, and the more efficiently it can raise

⁷In this sense, the value of long-term relationships is similar to that which has been emphasized in the literature on banking. Bank relationships are valuable, this literature argues, because of the accumulated knowledge about borrowers that they generate. See, e.g., Fama (1985), Sharpe (1990), Rajan (1992), and Petersen and Rajan (1994), among others.

external finance. Thus long-established incumbents are likely to have a distinct edge over newcomers in this dimension. To the extent that this financing advantage allows incumbents to, say, carry inventories more cheaply, they will have lower marginal costs of "distribution", in the sense of the model.

2.1.B Potential entrants' cost structures

It is assumed that the potential entrant has invented a new technology, which if "developed", will allow it to produce at lower cost than the incumbent. That is, the potential entrant's post-development production costs in period t , C_t^e , are given by:

$$C_t^e = \lambda_t C_t^i \quad (2)$$

where $\lambda_t < 1$ represents the magnitude of the innovation. λ_t is a random variable distributed according to the time-invariant cumulative density function $G(\lambda)$ --the smaller is the realization of λ_t , the more significant is the period- t innovation.

Although I use the terminology "production costs" throughout, the innovations need not be literally thought of as only reducing costs. Equivalently, they can be thought of as allowing the innovating firm to produce a higher quality product than the incumbent for the same cost; this is the "quality ladder" interpretation adopted by Aghion and Howitt (1992), Grossman and Helpman (1991a, b) and others. All that really matters is that innovations lead to reductions in quality-adjusted costs for the firms that develop them.

A critical feature of the production cost innovations is that they exhibit spillovers across firms. In particular, if the potential entrant does indeed develop its innovation--i.e, it decides

to go ahead with production in period t and takes over the market--others can learn about the new technology, beginning in period $t+1$. This will allow future generations of innovators to stand on the shoulders of the entrant, so that their innovations will further reduce costs from a new base level of C_t^e .

Note that this assumption is already built into the notation: if a potential entrant decides to go ahead with production in period t , it will become the incumbent at time $t+1$. Therefore $C_{t+1}^i = C_t^e$ if there is entry in period t . If another innovator enters in period $t+1$, that second innovator will improve on the costs of the period $t+1$ incumbent, so that $C_{t+1}^e = \lambda_{t+1} C_{t+1}^i = \lambda_{t+1} \lambda_t C_t^i$. Thus if there are two consecutive rounds of innovation followed by development, production costs will fall by a factor of $\lambda_{t+1} \lambda_t$. More generally, if there are n rounds of innovation and development in any given interval, production costs will fall by a factor equal to the product of all the λ 's that were developed.

While the potential entrant has an edge over the incumbent in terms of production costs, it is at a disadvantage in terms of distribution costs. The potential entrant's distribution costs, denoted by D_t^e , satisfy:

$$D_t^e = dC_t^e \quad (3)$$

The form of (3) is similar to that of (1)--distribution costs are proportional to production costs for both incumbents and entrants--but the ratio of distribution costs to production costs is higher for entrants. Essentially, an entrant has the same distribution costs that would be associated with an "age zero" incumbent with the same production technology. Clearly, if $\lambda_t = 1$, and the new entrant does not have any production-cost advantage, it will be strictly less

efficient overall than the incumbent.

In addition to the marginal production and distribution costs, an entrant must pay a one-time fixed development cost to begin production with the new technology. This development cost can be thought of as the amount that must be spent to turn the research discovery into a commercially viable technology. It is denoted F_t^e , and satisfies:

$$F_t^e = fC_t^e \quad (4)$$

Like the distribution costs, the development cost is proportional to the firm's current production costs. As noted above, this is done to allow for a steady-state analysis of the model.

2.1.C Copycats' cost structures

In any period t , there is also a competitive fringe of copycat firms that can mimic the production technology of the current incumbent. That is, a copycat's cost of production, denoted by C_t^c , is given by:

$$C_t^c = C_t^i \quad (5)$$

However, since the copycats have no existing customer base, their distribution cost, D_t^c , is given by:

$$D_t^c = dC_t^c \quad (6)$$

Since they use only existing production technology, copycat firms do not have to pay any fixed development costs. Simply put, a copycat is like an incumbent, but without the

distribution cost advantage.⁸ Thus its overall marginal costs are always strictly higher than those of the current incumbent.

2.2 Product Market Competition

I now turn to the competition in the product market that determines both: 1) who will control the market in any period t ; and 2) the associated profits. To keep things simple, the demand side of the market is modelled in such a way as to be essentially irrelevant. This is accomplished by assuming that there are N consumers, each of whom wishes to purchase exactly $1/N$ units of the good, up to a reservation price of X_t . X_t is sufficiently large that it always exceeds the copycats' marginal costs, $C_t^c + D_t^c$. This implies that prices and profits in the model will be determined solely by the relative cost structures of the competing firms.

There are two cases to distinguish. In the first case, which occurs with probability $(1-p)$, there is no new invention in period t , and hence no potential entrant. In this case, the incumbent competes against just the fringe of copycat firms. This implies that the incumbent, by virtue of its lower marginal cost, captures the entire market, at a price equal to the copycats' marginal cost, $C_t^c + D_t^c$. Thus the incumbent's profits in this case are given by $(C_t^c + D_t^c) - (C_t^i + D_t^i) = dC_t^i(1-\beta^{\hat{A}_t^i})$; where \hat{A}_t^i is defined as $\min(A_t^i, \bar{A})$.

In the second case, which occurs with probability p , there is a potential entrant in period t . This case is a bit more complicated, and the timing of events is as follows. First, the

⁸The underlying assumption here is that the production technology does not have patent protection, and hence can be costlessly imitated after a one-period lag. Although this assumption simplifies the analysis somewhat, it is of no significant consequence. Similar results would follow if one were to grant incumbents patent protection and use another device to tie down prices in those periods in which an incumbent is not challenged by a new entrant.

potential entrant must decide whether or not to sink the development cost F^e . If it does not, then the incumbent is left alone with the copycats, and the outcome is exactly the same as described just above. If, on the other hand, the entrant does sink the development cost, then the entrant and the incumbent compete à la Bertrand. The copycats will be irrelevant to the outcome of this Bertrand competition, as their costs are now strictly higher than those of the other two types of firms.

To simplify the nature of the Bertrand competition between the incumbent and the potential entrant, I make the following strong assumption:

$$C_1^e + D_1^e < C_1^i + D_1^i \text{ for any } A_t^i \text{ and all } t; \text{ or equivalently,} \quad (7)$$

$$\lambda_1 < (1 + d\beta^{\bar{A}})/(1 + d) \text{ for all } t.$$

In words, (7) says that once development costs are sunk, the entrant's innovation is sufficiently valuable that it always has lower overall marginal costs than an incumbent of any age. The value of this assumption is established in the following lemma:

Lemma: If (7) holds, then the Bertrand equilibrium always involves the potential entrant gaining 100% market share and charging a price equal to the combined production and distribution costs of the incumbent firm, $C_1^i + D_1^i$.

Without the assumption that (7) holds, things might be quite different, and substantially more complicated. For example, suppose that $\bar{A} = 5$, and that there is a 3-period-old incumbent who is currently at a cost disadvantage relative to the entrant--i.e., $\lambda_1 < (1 + d\beta^3)/(1 + d)$.

However, suppose further that (7) is violated, such that $\lambda_t > (1 + d\beta^t)/(1 + d)$. The incumbent might reason as follows: "Suppose I price below cost today, in an effort to hang onto my customer base. Next period the potential entrant will surely still be around--it has already sunk the development cost--but I will be in a stronger position, because my distribution costs will have fallen as I have aged. Indeed, if I can hang on until next period I will have lower marginal costs than the entrant and will be able to earn a profit." Thus the incumbent might engage in a complex intertemporal pricing strategy designed to trade off current profits in order to maintain its customer base.

However, if (7) holds, this sort of strategy is useless. No matter how long the incumbent holds off the potential entrant, it will never have the lower costs. Thus there is no gain to the incumbent from pricing below current cost, and it will never do so. At the same time, the Bertrand equilibrium can also never involve a price higher than the incumbent's current cost, for the usual reason--the incumbent would deviate to a lower price, capture the entire market, and make a positive profit.

Thus (7) ensures that, once the development cost F_t^e is sunk, the entrant will always take over the market completely. However, it does not guarantee that the entrant will find it worthwhile to spend the F_t^e in the first place. For that to happen, it must be that expected profits net of the development cost--denoted by π_t^e --exceed zero.

These profits can be broken into two components. First, there is the net profit (after development cost) earned by the entrant in the period it enters--i.e., in period t . This can be calculated directly from the results of the lemma as: $C_t^i + D_t^i - C_t^e - D_t^e - F_t^e =$
 $C_t^e [((1 + d\beta^{t+1})/\lambda_t) - 1 - d - f]$

In addition to these immediate profits, a firm that enters in period t can also earn profits

in later periods if there is no subsequent entry by later innovators, because then it only has to contend with the weaker copycat firms. Of course, if subsequent entry by an innovating firm does occur in some future period $t+k$, the period- t entrant earns nothing from $t+k$ onward.

We have already calculated the profits that a period- t entrant will earn in period $t+1$, conditional on no entry by a new innovator at this time. This profit is simply that earned by a 1-period old firm competing only with copycats, or $dC^c(1 - \beta)$. Similarly, the profits that a period- t entrant will earn in period $t+2$, conditional on no entry by an innovator by this time, are given by $dC^c(1 - \beta^2)$.

Thus overall, the expected present value of net profits to a potential entrant in period t is given by:

$$\begin{aligned} \pi_t^e = & C_t^c \{ [(1+d\beta^{\hat{A}})/\lambda_t] - 1 - d - f \} \\ & + [(d-d\beta)/(1+r)] * (\text{Prob of no entry at } t+1) \\ & + [(d-d\beta^2)/(1+r)^2] * (\text{Prob of no entry through } t+2) \\ & + \dots + [(d-d\beta^{\hat{A}})/(1+r)^{\hat{A}}] * (\text{Prob of no entry through } t+\hat{A}) \\ & + \dots + [(d-d\beta^{\hat{A}})/(1+r)^{\hat{A}+1}] * (\text{Prob of no entry through } t+\hat{A}+1) \dots \} \end{aligned} \quad (8)$$

where r is the per-period discount rate. The first term inside the curly brackets represents profits earned immediately upon entry in period t ; the second term represents expected profits in period $t+1$, the third term represents expected profits in period $t+2$, etc.

2.3 Steady-State Equilibrium Development Rules

The appropriate equilibrium concept for this model is a set of equilibrium development rules. In particular, for any given age of the incumbent firm A , we want to calculate a "threshold" $\lambda^*(A)$ such that a potential entrant will choose to develop its invention if and only if $\lambda_t < \lambda^*(\hat{A})$. Intuitively, in this equilibrium, a potential entrant will weigh both the magnitude

of its own innovation, and the age of the current incumbent, in making the development decision. At the same time, it will take as given the set of equilibrium development rules when it attempts to assess the likelihood of entry by future generations of innovators.

For example, if the potential entrant does decide to go forward in period t , the probability that there will be another round of innovation and entry in period $t+1$ is given by: $pG(\lambda^*(1))$. Similarly, the probability of another round of innovation and entry in period $t+2$ (conditional on there being no entry in period $t+1$) is $pG(\lambda^*(2))$.

One can re-write equation (8) in more compact form as:

$$\pi_t^e = C_t^e \{ [((1+d\beta^{\hat{A}_t})/\lambda_t) - 1 - d - f] + \text{FUTURE} \} \quad (8')$$

where the important new observation is that FUTURE depends on the equilibrium rules $\lambda^*(1), \dots, \lambda^*(\bar{A})$, but is independent of the period- t data, λ_t and \hat{A}_t . Moreover, it is easy to show that FUTURE is a monotonically decreasing function of each of the λ^* 's. This makes intuitive sense: the greater is any λ^* , the greater is the probability that a future innovator will decide to enter, and hence the lower are the expected profits associated with entry today. This is similar to the "creative destruction" effect noted by Aghion and Howitt (1992) and others.

The λ^* 's are determined by setting the appropriate variants of (8') equal to zero. For example, $\lambda^*(1)$ satisfies:

$$[((1+d\beta)/\lambda^*(1)) - 1 - d - f] + \text{FUTURE} = 0 \quad (9)$$

Analogously, $\lambda^*(2)$ satisfies:

$$[((1+d\beta^2)/\lambda^*(2)) - 1 - d - f] + \text{FUTURE} = 0 \quad (10)$$

Given that FUTURE is independent of λ_i and \hat{A}_i , we can use equations like (9) and (10) to derive the following recursive relationship among the λ^* 's:

$$\lambda^*(A)/\lambda^*(1) = (1+d\beta^A); \text{ for all } A \leq \bar{A} \quad (11)$$

In order to solve explicitly for the individual λ^* 's, we need to pin down a "boundary condition"--i.e., we need to establish the value of, say $\lambda^*(1)$. While this is difficult to do in closed form, it is easy to establish the existence and uniqueness of a solution, and to do some simple comparative statistics. From (9), $\lambda^*(1)$ is determined by the following equality:

$$1 + d + f - (1+d\beta)/\lambda^*(1) = \text{FUTURE} \quad (9')$$

We have already seen that the term on the right-hand side of (9'), FUTURE, is a decreasing function of $\lambda^*(1)$. It is also easy to see that the term on the left-hand side of (9') is an increasing function of $\lambda^*(1)$. Thus the unique equilibrium value of $\lambda^*(1)$ will be determined as in Figure 1. In the figure, the increasing function $1 + d + f - (1+d\beta)/\lambda^*(1)$ is denoted by "CURRENT", since this function represents the net current (i.e., period t) losses associated with entry when $\lambda_t = \lambda^*(1)$. Thus the equilibrium point has the simple interpretation of being that

value of λ , where the current costs of entry are just equal to the expected future profits.

Example 1: In this example, the parameter values are chosen as follows: $p = 1$; $d = 1$; $f = 5$; $\beta = 0.8$; $\bar{A} = 10$; and $r = .10$. In addition, the distribution $G(\lambda)$ is assumed to be uniform over the interval $(0, .5)$. It is easily checked that this distribution satisfies the condition in (7).

These parameters lead to a value of $\lambda^*(1) = .268$. From (11), it follows that the λ^* 's then decline monotonically to a value of .165 for $\lambda^*(10)$. The probability of a developed innovation when the incumbent is only one period old is $.268/.5 = .536$. However, when the incumbent is ten or more periods old, this probability falls to $.165/.5 = .330$. It is also possible to calculate the unconditional, steady-state probability of a developed innovation, which in this example is equal to .491.

2.4 Implications for the Pace of Innovation

One result that emerges immediately from Figure 1 has to do with the effect of development costs on the pace of innovation:

Result 1: An increase in the ratio f of development costs to production costs reduces the threshold λ^* 's--thereby deterring the development of new innovations--but on less than a one-for-one basis.

The impact of increased development costs on innovation is dampened somewhat because

there are two competing effects at work. On the one hand, an increase in f has a direct negative impact on the profits earned by an entrant in the period these costs are paid—i.e., in the period that entry occurs. This is manifested as an inward shift of the upwards-sloping CURRENT curve in the figure. On the other hand, an increase in development costs raises the expected profits that an entrant will earn in the periods after entry, because it lowers the probability of further rounds of entry by subsequent innovators. This is reflected in the fact that the CURRENT curve shifts along the downwards-sloping FUTURE curve, rather than along a horizontal line.

Example 2: Maintain all the same parameter values as in Example 1, but double f , from 5 to 10. $\lambda^*(1)$ falls from its previous value of .268 to .161. Thus the probability of a developed innovation when the incumbent is one period old falls from .536 to .322. The unconditional probability of a developed innovation falls from .491 to .269.

More central to this paper are the effects of customer bases on innovative activity. A first natural question to ask is: as customer bases become relatively more important in inter-firm competition, does the pace of innovation tend to accelerate or slow down? That is, does the unconditional probability of a developed innovation rise or fall?

There are two ways to parametrize the importance of customer bases: either an increase in d or a decrease in β can be thought of as raising the competitive advantage of incumbent firms along this dimension. (In the extreme cases where either $d = 0$ or $\beta = 1$, incumbents never have any advantage.) Whichever measure is used, the following result obtains:

Result 2: The effect of customer bases on the pace of innovation is in general ambiguous. In some cases, an increase in the importance of customer bases (i.e., an increase in d , or a reduction in β) can dramatically stifle innovation. However, there are other cases in which an increase in the importance of customer bases can actually promote the development of innovative new technologies. The latter outcome is more likely when p is close to zero.

The source of the ambiguity can be heuristically understood by reference to Figure 1. For concreteness, consider the impact of an decrease in β . On the one hand, this reduces the current-period appeal of entry, thereby leading to an inward shift of the CURRENT curve. At the same time, once entry has occurred, a lower value of β makes future profits higher in each period in which the new entrant remains in control of the market. This causes an outward shift of the FUTURE curve. The net effect of these two shifts on $\lambda^*(1)$ cannot in general be signed. However, the latter effect is more significant when p is low, because then new inventions arise infrequently and it is more likely that the higher post-entry profits can be sustained for a longer period of time. This is why customer bases are relatively more favorable to innovation when p is close to zero.

On the one hand, strong customer bases create a type of entry barrier which can make it initially very costly for a new innovating firm to enter a market. However, the flip side is that if and when that firm does get established and has its own customer base, it will be able to earn higher profits, because its competitive advantage over other firms (i.e., the copycats) will be more pronounced.

Example 3: The potential for customer bases to have a negative impact on innovation can be illustrated most starkly by considering a case where: 1) $(1+d)/(1+d+f) > \lambda^*(\bar{A})$ (this can always be accomplished by making d and \bar{A} large enough); 2) the distribution $G(\lambda)$ is such that λ_1 always satisfies $\lambda^*(\bar{A}) < \lambda_1 < (1+d)/(1+d+f)$; and 3) p is strictly less than one. In this case, the economy eventually must get "stuck" in a situation where the incumbent firm is at least \bar{A} periods old, and where there are never any more new innovations developed. All it takes to get an incumbent to be this old is a run of \bar{A} consecutive periods with no new inventions; with $p < 1$, such a run will occur in finite time with probability one. And once the incumbent reaches this age, there is no innovation that is of sufficient magnitude to displace it, since it is always the case that $\lambda^*(\bar{A}) < \lambda_1$.

In contrast, if we maintained the same $G(\lambda)$, but set $\beta = 1$, so that customer bases did not matter, the economy could never get stuck in this way. Rather, every innovation that arose would always be developed. This is guaranteed by the assumption that $\lambda_1 < (1+d)/(1+d+f)$, which makes entry immediately profitable against an incumbent with no distribution cost advantage. The net result is a new developed innovation in a fraction p of the periods.

The observation that customer bases can be damaging to innovation may sound intuitive and not all that surprising. However, recall that the words "customer base" are just a euphemism for firm-specific learning-by-doing. So what the model is really saying is that the potential for firm-specific learning-by-doing can actually be very harmful to innovation and therefore to long-run growth. Phrased this way, the result sounds much more striking.

Example 4: To see how customer bases can actually help foster innovation when p is

low, consider a case where the parameter values are all the same as in Example 1, except that $p = .10$ instead of 1. In this case, the unconditional probability of a developed innovation is .064.

If, however, we take these same parameters but set $\beta = 1$ (keeping $p = .10$) the unconditional probability falls somewhat, to .057. Thus there is actually more innovation when customer bases matter than when they do not.

While Examples 3 and 4 illustrate the potentially dramatic consequences of customer bases for the long-run average level of innovation, perhaps the most novel aspect of the model has to do with its implications for the timing of innovations:

Result 3: Customer bases lead to an endogenous "bunching" of innovative activity. In particular, whenever there is an invention of sufficient merit to be developed in period t , this raises the likelihood of further development in future periods.

This result follows from the recursive formula for the λ 's in equation (11), which states that, the older is the current incumbent firm, the less attractive it is for a potential entrant to develop a given invention. This observation in turn implies that when a potential entrant does go forward in period t , this has a positive externality on future generations of potential entrants—it raises their prospective returns to development, since they will be facing a younger incumbent, on average. This positive externality might be termed a "shakeup" externality, since its essence is that entry in period t breaks an existing incumbent's hold on customers, thereby shaking up

the market and facilitating the entry of younger firms whose competitive advantage lies more in technological prowess and less in having an established customer base.

Example 5: Maintain all the same parameter values as in Example 1. Suppose there is an incumbent firm in period t that is over \bar{A} periods old, and that a potential entrant is deciding whether or not to proceed with development. Let us ask what effect the potential entrant's decision has on the probability of development in future periods, $t+1$, $t+2$, etc.

If the potential entrant does not develop its invention in period t , the new potential entrant in period $t+1$ will still be facing an incumbent over \bar{A} periods old. Thus the probability of a developed innovation in period $t+1$ is $pG(\lambda^*(\bar{A})) = .330$. In contrast, if the potential entrant does proceed with development in period t , the new potential entrant in period $t+1$ will have an easier task, since it will be facing an incumbent only 1 period old. Thus the probability of development in period $t+1$ rises to $pG(\lambda^*(1)) = .536$.

The potential entrant's decision in period t exerts a similar, though less pronounced influence, on the conditional probabilities for period $t+2$. If there is no development in period t , the probability of development in period $t+2$ is given by: $(1 - pG(\lambda^*(\bar{A})))pG(\lambda^*(\bar{A})) + (pG(\lambda^*(\bar{A})))pG(\lambda^*(1)) = .398$. If, on the other hand, there is development in period t , the probability of development in period $t+2$ rises to: $(1 - pG(\lambda^*(1)))pG(\lambda^*(2)) + (pG(\lambda^*(1)))^2 = .514$.

More generally, given that there is development in period t , the conditional probabilities of development in further-out periods decay slowly back to the unconditional value of .491. For example, the conditional probability at $t+3$ is .504; at $t+4$ it is .499, etc. Figure 2 illustrates

this time path of conditional probabilities.

3. Endogenous Research Expenditures

Thus far, the probability p of an invention arriving in any given period has been taken as exogenous. More realistically, this probability will be a function of the resources that potential entrants devote to research activities. If this is the case, there is an additional channel that leads to bunching of innovations: a shakeup of the market in period t not only makes it more attractive to develop an existing invention in period $t+1$, it also will encourage more research in period $t+1$, and hence raise the likelihood of there being an invention in the first place.

To capture this idea in the simplest possible way, assume as before, that there is just one potential entrant in each period. The potential entrant now begins the period by deciding how much effort to expend on research. The probability that the research will be successful and yield an invention in period t is then given by $p(e_t)$, where e_t is the level of research effort, and $p(\cdot)$ is an increasing, concave function. Regardless of whether or not the research is successful, research effort is costly to the potential entrant; in particular, the cost of an effort level e_t is given by R_t^e , which satisfies:

$$R_t^e = e_t C_t^i \quad (12)$$

As with all the other costs in the model, research costs are proportional to contemporaneous production costs. The only slight twist is that the potential entrant's research costs are proportional to the incumbent's current production costs. This simplifies the interpretation, since one can think of the incumbent's production costs as already known at the

time that the potential entrant begins research.⁹

In this extended version of the model, an equilibrium will consist of two sets of decision rules. First, as before, there will be the threshold λ^* 's, that tell us when development is optimal as a function of the incumbent's age. Second, there will be a set of optimal research levels, denoted by $e^*(1), \dots, e^*(\bar{A})$, which tell us how much research will be done as a function of the incumbent's age.

In order to begin thinking about the properties of this equilibrium, let us begin by focusing on equilibrium at the development stage. That is, take the research rules--the e^* 's-- as given, and solve for the λ^* 's. To do so, note that equation (8'), which describes an entrant's profits from development, is still valid. The only difference is that FUTURE is now a more complicated function since it depends also on the e^* 's. It follows immediately that the recursive equation (11) for the λ^* 's is exactly the same as before. In other words, once we pin down $\lambda^*(1)$, the ratios of the λ^* 's will be the same as they were earlier.

The computation of the $\lambda^*(1)$ itself is slightly more complicated, but the logic seen earlier still applies exactly. That is, for any given set of e^* 's, one can still show that there is a unique equilibrium value of $\lambda^*(1)$ as in Figure 1. The only change is that the FUTURE curve in this figure, while still downwards sloping, may have a slightly different shape, due to its more complicated functional form.

Having solved for the equilibrium development rules, one can now fold backwards and

⁹In contrast, it is a little more awkward to assume that the entrant's research costs are proportional to its own production costs, since these will not be known until the research is completed. Nonetheless, using this alternative assumption would have no effect on the results below.

solve for the equilibrium research levels. Expected profits at the research stage, denoted by π^r , satisfy:

$$\pi^r = C_i^i \left\{ p(e_i) \left[\int_{\lambda < \lambda^*(A)} \frac{\pi_i^d(\hat{A}_i^i, \lambda) g(\lambda) d\lambda}{C_i^i} \right] - e_i \right\} \quad (13)$$

Here $\pi_i^d(\hat{A}_i^i, \lambda)$ is the profit at the development stage for a given value of \hat{A}_i^i and λ . It is calculated using equation (8).

From equation (13), one can immediately derive the first order conditions that must be satisfied by each of the e^r 's:

$$p'(e^*(A)) = \left[\int_{\lambda < \lambda^*(A)} \frac{\pi_i^d(A, \lambda) g(\lambda) d\lambda}{C_i^i} \right]^{-1} = \left[\int_{\lambda < \lambda^*(A)} (1 + d\beta^A - \lambda(1 + d + f) + \lambda \text{FUTURE}) g(\lambda) d\lambda \right]^{-1} \quad (14)$$

Equation (14) makes it clear that, in any equilibrium, $e^*(A)$ must be a decreasing function of A . This follows from the fact that both the integrand in (14), as well as the limit of integration, $\lambda^*(A)$, are decreasing functions of A . Intuitively, the older the incumbent, the lower is the marginal productivity of investment in research by a potential entrant. This is true for two related reasons. First, development of any given invention is less profitable with an older

incumbent. Second, fewer inventions are worth developing at all with an older incumbent. The bottom line is that less research is done when the incumbent is older.

This leads to the principal result of this section:

Result 4: Endogenizing the level of research makes the bunching of innovations more pronounced. In particular, the positive effect of an innovation in period t on the conditional probability of innovation in period $t+1$ is stronger when research is endogenous.

The reasoning for this result is straightforward. The probability of a developed innovation in period $t+1$ is $p(e^{*(\hat{A}_{t+1}^i)})G(\lambda^{*(\hat{A}_{t+1}^i)})$. When there is an innovation in period t , this has the effect of setting $\hat{A}_{t+1}^i = 1$. With endogenous research, this reduction in the age of the incumbent has two distinct beneficial effects: it increases both the probability of an invention, and the likelihood that the invention will subsequently be developed. In contrast, in the simpler model of Section 2, only the latter effect was at work.

Example 6: Maintain all the same parameter values as Example 1, except: the probability p of an invention can be either 1 or .5, depending on the level of research effort. In any period t , it costs an amount equal to $.25C_t^i$ to set $p = 1$; it is costless to set $p = .5$.

It can be shown that these parameters lead to $\lambda^*(1) = .28$. Moreover, when $A_t^i = 1$, the expected profits conditional on obtaining an invention in period t are equal to $.505C_t^i$. Thus it is (just barely) worth it to spend the money on research to raise the probability of an invention from .5 to 1. In contrast, when A_t^i exceeds 1, it is never worth it to invest in research.

We can now revisit the question asked in Example 5, namely how does development in period t affect the conditional probability of development in period $t+1$? If we begin period t with an age-10 or older incumbent and there is no development, the conditional probability for $t+1$ is only $.5G(\lambda^*(10)) = .17$. In contrast, if there is development in period t , the conditional probability for $t+1$ rises to $G(\lambda^*(1)) = .56$.

4. Further Extensions

4.1 Gradual Displacement of Incumbents: Implications of the Theory for the Number of Firms in an Industry

The current version of the model has the unrealistic feature that when a stronger entrant arrives on the scene, it immediately and completely displaces the incumbent. Thus in equilibrium, there is always just one firm with 100% market share active at any point in time. While this feature makes it considerably easier to derive the sort of steady-state research and development rules seen above, it is by no means central to the key economic intuition of the paper, which is this: because of the shakeup externality, entry in period $t+1$ is more attractive when there has been entry in period t . This intuition survives in a more realistic setting in which new entry at time t only partially displaces the existing incumbent.

The basic idea can be illustrated as follows. Suppose that we have the same set-up as before, with one modification. Conditional on development, new entrants have lower marginal costs than incumbents for low levels of production, but these marginal costs are increasing. That is, new entrants have limited capacity and cannot immediately start producing cost-effectively on the same scale as incumbents.

Now suppose we start from a situation where there is a single incumbent firm that is over \bar{A} periods old, and that has 100% market share. Imagine that a new entrant has an invention of sufficient importance that it decides to proceed with development in period t . Because of scale constraints, the new entrant takes over some of the market--say 25% of it--but not all of it. The incumbent retains its distribution cost advantage with the 75% of the market that it continues to serve in period t . However, this advantage is erased for the 25% of the market that switches over to the new entrant, since the incumbent's ties to these customers have been severed.

Next, suppose that another potential entrant arrives in period $t+1$. Entry will look more attractive than it did in period t . Although 75% of the market is still being served by an "old" firm (more than \bar{A} periods old) the other 25% is now being served by a firm that is only one period old and therefore has less of a distribution cost advantage with its customers. Thus the basic logic set out above still holds--entry in period t still dislodges some customers from the grip of old incumbents, and therefore makes it easier for period $t+1$ entrants to attract these customers.

One attractive feature of this version of the model is that it yields implications for the number of firms active in the market at any point in time. In other words, because there is not complete and immediate displacement of incumbents, entry in any period t now corresponds to an increase in the number of firms in the market. Thus the model would now predict that a major innovation that shakes up a market would typically be followed by a period of supra-normal growth in the number of firms active in the market.

This sort of prediction can be compared to the stylized facts on industry evolution that

emerge from the work of Gort and Klepper (1982) and Klepper and Graddy (1990). This work documents a tendency for the number of firms in a typical industry to follow a distinctive, non-monotonic pattern. In particular, Gort and Klepper (1982) identify five stages of industry evolution, based on their study of 46 products:

"Stage I encompasses the interval in which the number of producers in the market remains relatively small (usually between one and three). Stage II is the interval from the "take-off" point of net entry to the time that net entry decelerates drastically. Stage III is the ensuing period of low or zero net entry, and Stage IV is the subsequent period of negative net entry. Stage V represents the new equilibrium in the number of producers that coincides with the maturity of the product market and continues until some new fundamental disturbance generates a change in market structure."

The arguments made above suggest that the notion of a shakeup externality can be especially helpful in understanding the "take-off" phase of an industry—i.e., the phase during which the industry makes the transition from the Stage I of no growth to the Stage II of rapid growth in the number of producers. Think of Stage I as a time during which there is, in the language of the model, a single incumbent more than \bar{A} periods old. During this time there is a relatively low probability of entry. However, if and when an innovation large enough to prompt entry is drawn from the $G(\lambda)$ distribution, the market will be shaken up and there will be a higher incidence of further entry leading to a period of rapid growth in the number of active firms—this

is the take-off phase.¹⁰

It is also important to note that in the Gort and Klepper (1982) data, the rapid growth in the number of firms that occurs in Stage II is accompanied by dramatic increases in productivity. The mean duration of Stage II in their sample is 9.7 years, and the mean annual price decrease over this time is 13%. In contrast, price decreases are significantly smaller in magnitude during the later stages when there is reduced entry. These findings fit broadly with the primary implications of the model, namely: 1) technological improvements in an industry will tend to be "bunched", rather than occurring at a steady pace; and 2) these improvements will be associated with the entry of new firms into the industry.

4.2 Marketable Innovations

Up to this point, it has been assumed that a newcomer with a developed innovation always enters the market and displaces the incumbent. When this happens, the incumbent's distribution cost advantage is not exploited. However, if the incumbent and the potential entrant could somehow join forces, it would seem that their total joint surplus could be raised—i.e., by

¹⁰Other recent papers have also advanced theoretical arguments that can rationalize some or all of the Gort and Klepper (1982) findings. However, the story given here to explain the take-off phase differs from those in other papers. For example, in Hopenhayn (1993), growth in the number of firms is driven by exogenous growth in demand. And in Jovanovic and MacDonald (1993) there is an exogenous change in the technological environment—an outside "invention"—that makes it suddenly possible for firms to pursue innovative activity. In contrast with the latter paper, in my model firms always have the ability to pursue innovation—it is just that they are endogenously more inclined to do so after a shakeup of the market makes innovation more attractive.

taking advantage of both the newcomer's better technology and the incumbent's loyal customer base, they could together make more profit than the newcomer could by itself.

One way that this could happen is if the incumbent were able to buy the innovation from the potential entrant, and then sell it through its own existing distribution channels. Casual observation suggests that such acquisition of new technologies by established firms happens with some frequency. Of course, this is not to suggest that it is always possible for inventors to sell or license their innovations. Clearly, there will be cases in which information or moral hazard problems are so pronounced that it becomes infeasible to do so. Nonetheless, it does seem worth at least considering the possibility that some outside innovations will be acquired by incumbent firms.

To do so in the simplest possible way, assume there are two types of innovations. "Non-marketable" innovations, as before, simply cannot be transferred from one firm to the other. Thus the only way that a non-marketable innovation can ever be introduced is by a new entrant. In contrast, "marketable" innovations can, once they are developed by the potential entrant, be transferred costlessly to the incumbent firm. The interpretation is that marketable innovations are ones for which post-development information and agency problems are relatively small, so that there are no significant frictions in having the original owner sell them.

The first observation to make about this extended version of the model is that marketable innovations will, if developed, always be acquired and distributed by the incumbent firm. This is just the Coase theorem at work--total profits are always higher when the incumbent's distribution cost advantage is exploited. Therefore, the introduction of a marketable innovation always leaves the incumbent in place. The immediate corollary is that only non-marketable

innovations have the externality effect of "shaking up" the market--i.e., of displacing the older incumbent firm.

This suggests that if we took the polar extreme where all innovations were marketable, the model would be reduced back to something very similar to that in Aghion and Howitt (1992). The age of the incumbent would cease to be a relevant state variable, as the incumbent would always be greater than \bar{A} periods old. In equilibrium, the probability of innovation would be constant from one period to the next, and there would be no tendency for bunching of innovations.

Given that marketable innovations have no shakeup externalities, the natural issue to focus on in this setting is the externality effects associated with non-marketable innovations. Now, we already know from above that a non-marketable innovation in period t has a positive effect on future non-marketable innovations. But what is the effect of a non-marketable innovation in period t on the probability of future marketable innovations? In other words, if a new entrant shakes up the market by entering in period t , is it more or less likely that marketable innovations will occur in future periods?

The answer to this question depends crucially on the relative bargaining power of the incumbent and the potential entrant when they negotiate the price at which the marketable innovation will be transferred. The extreme case where the incumbent has all the bargaining power is easiest to understand. In this case, the incumbent just has to compensate the potential entrant for what it would earn were it to bring the innovation to market itself. Thus the potential entrant earns exactly the same post-development profit it would were the innovation non-marketable. Consequently, the potential entrant's optimal research and development strategies

will be identical whether it is considering a marketable or a non-marketable innovation.

The bottom line is that if the incumbent has all the bargaining power, a shakeup of the market in period t has exactly the same beneficial impact on marketable innovations in period $t+1$ as it does on non-marketable innovations. The intuition is simple. As before, a shakeup of the market in period t leads to a younger incumbent in period $t+1$ and therefore makes entry more attractive for a newcomer at this time. Even when the innovation in question is marketable and the newcomer will not literally enter the fray, it is in a stronger position when negotiating with a relatively young incumbent. Knowing this, the newcomer will be more willing to expend resources on both research and development when the incumbent firm is younger.

A similar logic applies when the incumbent has "most" but not all of the bargaining power. However, when the potential entrant has a great deal of bargaining power, and is able to capture most of the total surplus in the negotiating process, the results may be reversed. This is because an older, more profitable incumbent has more to lose by being forced out of the market. Thus if the entrant has a lot of bargaining power, it may be able to extract more from an old incumbent than from a younger incumbent. In this case, a shakeup of the market in period t would stimulate future non-marketable innovations, but might actually discourage (at least in relative terms) future marketable innovations.

Figure 3 summarizes the discussion, illustrating the very different spillover properties of marketable and non-marketable innovations. While both add to society's stock of knowledge, a marketable innovation in period t does not do anything to encourage innovation of any sort in future periods, since it does not shake up the market. In contrast, we have just seen that a non-marketable innovation in period t can stimulate future innovation of both types, particularly if

incumbents have most of the bargaining power. Thus in such circumstances, it can be said that non-marketable innovations tend to impart larger positive externalities, all else equal.

5. Conclusions

The premise of this paper is a simple one: even in industries in which innovation and knowledge spillovers are critically important, firms are unlikely to compete solely on the basis of such innovation-driven variables as product quality or cost. Rather, they will also compete in part on the basis of variables such as the strength of their respective customer bases. With customer bases there is the equivalent of learning-by-doing at the firm level, but there are no spillovers across firms.

Introducing customer bases into a dynamic model of repeated innovation significantly alters the model's predictions. First, customer bases can in some circumstances dramatically reduce the long-run average level of innovation. Or stated somewhat more generally, the potential for firm-specific learning-by-doing can, ironically, be quite harmful to long-run growth.

Second, customer bases tend to generate endogenous waves in innovative activity of the sort described by Schumpeter (1936). The key to these waves is what I have termed a shakeup externality. When a new firm successfully enters the market, it breaks the incumbent's stranglehold on the customer base. This in turn makes it more attractive for the next generation of innovators to enter.

The model developed above is extremely stripped-down and stylized. This was done so as to illustrate the important consequences of customer bases in the simplest possible way. Unfortunately, this simplicity also makes the model less empirically realistic, and thus less

appropriate for directly confronting the data. One noteworthy deficiency of the model is the very artificial industry structure in which only one firm at a time is ever active. Given that some of the most interesting evidence regarding innovative industries (e.g., Gort and Klepper (1982), Klepper and Graddy (1990) centers on the number of firms present at any point in time, this deficiency is particularly striking. Thus, if one is interested in generating a more empirically relevant model, it seems that enriching the industry structure along the lines discussed in Section 4.1 would be a useful first step.

References

- Aghion, Philippe and Peter Howitt, 1992. "A Model of Growth through Creative Destruction," *Econometrica* 60, 323-351.
- Caballero, Ricardo J. and Adam B. Jaffee, 1993. "How High Are the Giants' Shoulders: An Empirical Assessment of Knowledge Spillovers and Creative Destruction in a Model of Economic Growth," *NBER Macroeconomics Annual 1993*, 15-74.
- Continental Bank Roundtable on "Global Competition in the '90's", *Journal of Applied Corporate Finance*, Spring 1993, 38-59.
- Diamond, Douglas 1989. "Reputation Acquisition in Debt Markets," *Journal of Political Economy*, 97, 828-862.
- Fama, Eugene 1985. "What's Different About Banks?" *Journal of Monetary Economics*, 15, 29-36.
- Gilbert, R.J. and Newbery, D.M. 1982. "Preemptive Patenting and the Persistence of Monopoly," *American Economic Review*, 72, 514-526.
- Gort, M. and S. Klepper 1982. "Time Paths in the Diffusion of Product Innovations," *Economic Journal*, 92, 630-53.
- Grossman, Gene M. and Elhanan Helpman 1991 a. "Quality Ladders in the Theory of Growth," *Review of Economic Studies*, 58, 43-61.
- _____ 1991 b. "Quality Ladders and Product Cycles", *Quarterly Journal of Economics*, 106, 557-586.
- Henderson, Rebecca. 1993. "Underinvestment and Incompetence as Responses to Radical Innovation: Evidence from the Photolithographic Alignment Equipment Industry," *RAND Journal of Economics*, 24, 248-270.
- Hopenhayn, Hugo A. 1993. "The Shakeout," Working Paper, Universitat Pompeu-Fabra.
- Jovanovic, Boyan and Glenn M. MacDonald 1993. "The Life-Cycle of a Competitive Industry," Working Paper, National Bureau of Economic Research.
- Klepper, Steven and Elizabeth Graddy 1990. "The Evolution of New Industries and the Determinants of Market Structure," *RAND Journal of Economics*, 21, 27-44.

- Petersen, Mitchell and Raghuram Rajan 1994. "The Benefits of Lending Relationships: Evidence from Small Business Data," *The Journal of Finance*, 49, 3-37.
- Porter, Michael E. 1980. *Competitive Strategy*, New York: The Free Press.
- Rajan, Raghuram 1992. "Insiders and Outsiders: The Choice between Informed and Arm's Length Debt," *Journal of Finance*, 47, 1367-1400.
- Romer, Paul M. 1990. "Endogenous Technological Change," *Journal of Political Economy*, 98, S71-S102.
- Schumpeter, Joseph A. 1936. *Theory of Economic Development: an Inquiry into Profits, Capital, Credit, Interest, and the Business Cycle*, Harvard University Press.
- Schumpeter, Joseph A. 1942. *Capitalism, Socialism and Democracy*, New York, Harper and Brothers.
- Segerstrom, Paul S., T.C.A. Anant, and Elias Dinopoulos 1990. "A Schumpeterian Model of the Product Life Cycle," *American Economic Review*, 80, 1077-1091.
- Segerstrom, Paul S. 1991. "Innovation, Imitation and Economic Growth," *Journal of Political Economy*, 99, 807-827.
- Sharpe, Steven 1990. "Asymmetric Information, Bank Lending, and Implicit Contracts: A Stylized Model of Customer Relationships," *Journal of Finance*, 47, 1343-1366.
- Shleifer, Andrei 1986. "Implementation Cycles," *Journal of Political Economy*, 94, 1163-1190.

FIGURE 1

Cost and benefits of entering market




FIGURE 2

Probability of innovation at time $t+k$
Conditional on innovation at time t




Figure 3

Spillover Effect of Period-t
Innovation on Period-t+1 Innovation

		Type of innovation in period t	
		Non-marketable	Marketable
Effect on probability of innovation of this type in period t+1	Non- marketable	(+)	0
	Marketable	(+) if incumbent's bargaining power is high (-) if entrant's bargaining power is high	0