A Punctuated-Equilibrium Model of Technology Diffusion

Christoph H. Loch
INSEAD
Fontainebleau, France

Bernardo A. Huberman
Xerox PARC
Palo Alto, CA 94304

Abstract

We present an evolutionary model of technology diffusion in which an old and a new technology are available, both of which improve their performance incrementally over time. Technology adopters make repeated choices between the established and the new technology based on their perceived performance, which is subject to uncertainty. Both technologies exhibit positive externalities, or performance benefits from others using the same technology.

We find that the superior technology will not necessarily be broadly adopted by the population. Externalities cause two stable usage equilibria to exist, one with the old technology being the standard and the other with the new technology the standard. Punctuations, or sudden shifts, in these equilibria determine the patterns of technology diffusion. The time for an equilibrium punctuation depends on the rate of incremental improvement of both technologies, and on the system's resistance to switching between equilibria. If the new technology has a higher rate of incremental improvement, it is adopted faster, and adoption *precedes* performance parity if the system's resistance to switching is low. Adoption of the new technology *trails* performance parity if the system's resistance to switching is high.

1. Introduction

In the age of gene technology, superconductors and supercomputers, it is often claimed that innovations are mainly breakthroughs from existing technologies. For example, it has been stated that new audio technologies may have failed to replace the established audio CD because none has been able to offer the "necessary ten-fold performance improvement" over the CD technology (Economist 1996).

On the other hand, it has been established that much technological progress stems from streams of incremental improvements over time. In a number of industries it is also observed that long periods of incremental improvement tend to be interrupted by short periods of radical innovation (Abernathy and Utterback 1978, Utterback and Suarez 1993). This pattern has been called "punctuated equilibrium" (Tushman and Anderson 1986, Mokyr 1990).

The question is what drives the interplay between the two: When is incremental innovation dominant, and when (if at all) do radical changes in technology appear? The literature presents a number of obstacles to switching between technologies. For example, organizational inertia and stable industry constellations may prohibit significant innovations and the associated change for long periods until change is forced by a crisis. Cultural "openness" may also foster or inhibit significant techno-logical changes, and interdependencies among multiple component technologies in complex products can prevent radical innovations for compatibility reasons.

In this paper, we develop a model that offers a different explanation for punctuated technology diffusion. In this theory, punctuated equilibrium may occur among rational adopters who at any point in time choose the technology with the better performance. This can happen if positive externalities as well as uncertainty in the evaluation of a new technology are present in the system.

After reviewing the relevant literature we present the basic model in Section 3. We then demonstrate the full dynamic adoption behavior in the punctuated equilibrium model in Section 4 through the use of computer simulations. Section 5 develops two extensions: one with users differing in their attitude toward the new technology (innovators and followers), and the other with users basing their adoption decision not on current technology performance, but on expectations about future performance evolution.

2. Literature Review

The acceptance and spread of new technology in a market or user community is commonly referred to as *diffusion*. It is an important topic of research in several disciplines, such as marketing, strategy, organizational behavior, economics, and the

history of technology. They offer several possible explanations for a punctuated-equilibrium-like adoption of radically new technologies.
The classical diffusion model in the marketing literature is the S-curve model of spreading innovations. S-curve growth (logistical growth) results when growth is proportional to the established base (contagion) and to the remaining untapped potential. This model has successfully been fitted to new product innovations in many industries (in particular, consumer industries). It was found that the diffusion speed is highly industry-specific. The typical application of diffusion S-curves is to new products or product categories opening up a new market potential, but not to the competition between an established and a new technology (Kumar and Kumar 1992). The present paper examines how a newly available technology diffuses in competition with an established technology.

Abernathy and Utterback (1978) first pointed out that industries often go through cycles of incremental innovations, punctuated by short periods of radical change. One reason for this is the necessary initial uncertainty reduction about a new technology, both for the providers and for the users, leading to instability and experimentation for some time. After the technology is better understood, an incremental era of refinement sets in (Tushman and Anderson 1986). Clark (1985) and Henderson and Clark (1990) argue that the distinction between "incremental" and "radical" innovation is not clearly defined; they add as an intermediate type *architectural* innovations, in which the components of a product stay (largely) unchanged, but the architecture of how they are combined and interact changes. This intermediate type explains why sometimes seemingly small "conceptual" changes in technology are in fact radical in the sense that they require a new way of thinking about the whole.

Organizational theorists show that another reason for long periods of incremental change is "inertia," or the inability of organizations to change their procedures and routines unless forced to do so by a crisis (e.g., Tushman and Romanelli 1985, Sastry 1994). A separate reason for an organization's inability to deal with new technologies is closely linked to Clark's (1985) architectural uncertainty, which may negate the usefulness of a firm's current set of procedures and thus prove "competence-destroying," creating resistance to adoption. A third source of inertia is the external environment in which organizations operate. The structure of complex products is mirrored in alliances, associations and coop-erations with interlocking interests, and a new technology, independent of its merits, may be blocked because it upsets this system of interests (Anderson and Tushman 1990). Specifically, firms that dominate their market with an established technology may suppress new technology to defend their market position and to prevent cannibalization of their dominant products (Ghemawat 1991).

Similarly, historians of technology have demonstrated the role of the external environment in the interplay between incremental and radical innovations over time. Mokyr (1990) offers two possible circumstances under which radical innovations may occur often occur.

First, social, economic, political, and religious factors in the environment greatly influence how “receptive” society is for innovations, or whether they tend to be suppressed.

In addition to the environment, however, the characteristics of the innovations themselves play a role: First, if the innovation affects a large, complex system of interdependent parts, a significant change may make a component incompatible with the rest of the system and thus render the innovation infeasible. This is demonstrated by the sailing ship, which changed unrecognizably between the 16th and 19th century with only incremental innovations (Mokyr 1990, p. 294 f., see also Rosenkopf and Tushman 1994). This argument refers not to the compatibility of group interests (which Anderson and Tushman 1990 emphasize), but to compatibility of the technology components themselves.

Second, an innovation that is location-specific may not become widely accepted without further adaptation. As an example, agricultural innovations took place very slowly over the centuries, partly because each area had to adapt technology advances to their soil and climatic conditions (Mokyr 1990, p. 297).

To summarize, the following reasons have been put forward to explain the existence of punctuated equilibria in technology diffusion: 1. A radical innovation creates uncertainty (for producers as well as users), which needs to be resolved before wide-spread adoption can occur. 2. The new characteristics of the technology may destroy existing firm competences, which contributes to inertia within firms. 3. In addition, a new technology may be incompatible technically with other components of complex systems of which it is a part. 4. It may also upset the balance of cooperations and interests in the business network that has evolved around the old technology and its complements. 5. Finally, it may encounter resistance in society at large.

In this paper, we develop an analytical model of technology diffusion that offers a different explanation for the existence of punctuated equilibrium behavior. Our model argues that even in the absence of inertia[1] or compatibility issues (i.e. for a single-component technology) punctuated equilibrium-type diffusion can happen, provided that two factors are present: some positive network externalities or switching costs, and (initial) uncertainty about the performance of the new technology. Network externalities may cause multiple stable equilibria to exist, and performance uncertainty may trigger a switch of the population between equilibria. The equilibrium switch, or punctuation, may occur before the new technology has caught up in performance, or the incumbent equilibrium may persist long beyond the time when the new technology has become superior. We derive within the model under which circumstances the expected time of the equilibrium switch is earlier or later.

[1] Here, inertia is used in the sense of “organizational resistance to change.” Sometimes, inertia refers in a much wider sense to any outcome of a superior technology not being adopted. In this sense, our model offers an alternative explanation of inertia.

Our model follows the tradition of evolutionary models of technological change introduced by Nelson and Winter (1982).[2] We follow their approach by first assuming that actors are profit driven but unable to optimize because of bounded rationality. Actors simply choose the "best" out of the currently available technolo-gies, without being capable of perfect evaluation or of anticipating the system equilibrium (the latter assumption will be relaxed in Section 5). Second, our model focuses on the off-equilibrium behavior of the system, or the system dynamics, until an equilibrium is possibly reached.

Our model is also related to path dependence models with positive externalities, or increasing returns (Arthur 1983 and 1994, David 1994). Path dependence models emphasize that small initial advantages may determine which one of possibly several competing technologies is chosen by the user community. Path dependence may have several roots (David 1994): historically grown mutually consistent expectations, from "sunk" investmenets in channnels and coordination protocols, or from positive externalities (benefits resulting from other users adopting the same technology, see Granovetter 1978, Cusumano *et al.* 1992). Positive externalities are also referred to as "bandwagon" adoptions (e.g., Abrahamson and Rosenkopf 1997). Such "critical-mass" or "path-dependent" models have recently been attacked because of a lack of empirical evidence and also because they imply that once the initial "perturbation" has happened, the fate of the technology can no longer be influenced (Rogers 1991, Liebowitz and Margolis 1995).

Positive externalities also play an important role in the present article, as wider acceptance of the technology increases its usefulness (as in the case of the telephone) or encourages the provision of complementary assets (such as application programs for a PC platform). The difference is that we are focusing on when and how fast technology adoption occurs, and how adoption is related to the underlying technology performance. We do not focus on the path-dependent choice of which of several competing new technology comes to dominate the market.[3]

---

[2] There also exists a literature on technology races in economics (e.g. Reinganum 1985), which applies more to planned innovations in anticipation of competitive moves, but less to unplanned "radical" changes.

[3] In fact, the "new" technology in Sections 4 and 5 will be chosen eventually with probability one, since we assume that the "old"technology can no longer improve its basic performance. This assumption reflects our interest in the explanation of equilibrium punctuations. Our model can also be used to shed light on competing new technologies, in which case it may exhibit path-depent technology choice.

3. The Basic Model

Consider an industry with N actors, for example, firms. The actors use a certain technology to conduct their business, but at some point in time, a new technology becomes available. Each firm can choose repeatedly to stay with the old technology or to switch to the new one. We assume that there are no switching costs between the technologies.[4] A firm does not evaluate technologies continuously, but according to a Poisson process with rate α. The Poisson assumption is realistic and well-established for situations where many events (here the firms' evaluations) have independent sources. The assumption is not restrictive, since α can be made very large. It ensures that firms do not evaluate technologies in lockstep, following a common clock, but perform evaluations asynchronously, rather.[5]

Firms choose technologies based on performance. For now, we assume that all firms are identical in the sense that they measure performance in the same way. This assumption will be relaxed later. Call the current performance of technology i P_i, where $i = 1$ refers to the "old" and $i = 2$ to the "new" technology. Thus, the firms are assumed to be capable of a performance-based choice at each evaluation, but not of further optimization. Neither do firms anticipate the system equilibrium or act accordingly; we will introduce expectations into the model in Section 5.

This basic performance is subject to several additional influences. First, both technologies exhibit some positive externalities, or benefits from other users sharing the same technology. We assume the externalities to be linear.[6] We let $f_i = n_i / N$ denote the fraction of firms that are currently using technology i (note that $f_2 = 1 - f_1$). That is, $f_2 = 0$ corresponds to all actors using the old technology, $f_1 = f_2 = 0.5$ to each technology being used by half the population, and $f_2 = 1$ means that all users have adopted the new technology. We use f_2 as the *state variable* of the system. Positive externalities add an additional performance benefit of b_i for each additional fraction of technology i users, expressed by f_i. b_i is a technology-specific externality coefficient.

Second, the performance of the new technology is uncertain and can only be evaluated up to a random component ξ. For concreteness, we assume that ξ has a symmetric exponential distribution with parameter β and density

---

[4] This assumption can be relaxed by incorporating a shift in the performance difference between the technologies, which is described in the conclusion.

[5] Synchronicity of evaluations, or the existence of a common "clock" may significantly influence the dynamic behavior of the system, leading to complicated patterns and chaotic time paths that are not realistic. For a discussion see Huberman and Glance 1993a.

[6] This is notationally convenient, and there is also some empirical evidence that externalities are linear at least over part of their range, see Arthur (1983 and 1994).

$$f_\xi(x) = 1/2\ \beta\, e^{-\beta x} \text{ for } x \geq 0, \text{ and } f_\xi(x) = 1/2\ \beta\, e^{\beta x} \text{ for } x < 0 \qquad (1)$$

Thus, ξ has zero mean and variance $1/\beta^2$. Each actor evaluates the technology separately and independently, so the random components across actors are iid. random variables. The third additional influence is an incremental performance increase of both technologies over time. For the analysis of the simplest analytical model in this section, we take performance as fixed in time for now. Sections 4 and 5 will incorporate performance improvement rates over time.

Using the approach of Ceccato and Huberman (1989) for punctuated processes, we can now specify the performance of the two technologies as follows. Each user, when evaluating technology i, measures the total performance G_i, which results as the sum of the basic technology performance P_i and the additional externality benefit from every additional user of the same technology.

$$\begin{aligned} G_1 &= P_1 + f_1\, b_1 \\ G_2 &= P_2 + f_2\, b_2 + \xi \end{aligned} \qquad (2)$$

Let us first compare the performance of the two technologies, as shown in Figure 1. The horizontal axis is the state variable f_2, where we drop the subscript for the remainder of this Section to simplify notation. G_2 increases over f, while the old technology's performance decreases over f. The slopes of the two curves in each graph are the externality coefficients b_i.

Two cases are shown. On the left-hand side of Figure 1, the new technology is so inferior that even with all possible externality benefit it still cannot match its counterpart's performance. In this case, there is exactly one equilibrium, with all actors choosing the old technology (see, for example, Schelling 1978, p.). If the new technology's uncertainty $1/\beta^2$ is large enough, there will be a few actors "trying out" the new technology from time to time, but these trials will die out. The equilibrium is stable, and the new technology is not adopted. If the relative performance of the two technologies is reversed, then by the same argument, there will be one equilibrium with all actors using the new technology.

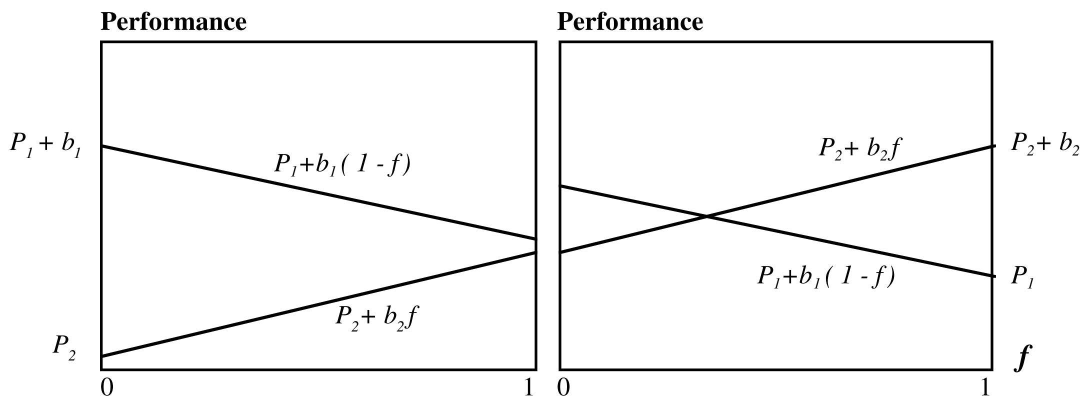

Figure 1: Technology Performance at a Fixed Point in Time

On the right hand side of Figure 1, the situation is different. The new technology is inferior when all actors use the old technology, but superior when adopted by all actors. The performance crossover point is often referred to as "critical mass" (e.g., Schelling 1978, Rogers 1985). In our model, it is given by

$$\hat{f} = \frac{P_1 + b_1 - P_2}{b_1 + b_2} \tag{3}$$

Whenever the crossover $\hat{f}$ is interior, there are two stable equilibria, one with all actors in one camp and one with all actors in the other camp. Both equilibria are stable in the sense that every actor will find it more beneficial to stay with the equilibrium than to deviate from it.

Consider a situation where the system starts out with all actors using the old technology. The question is whether the system will ever cross over to adopting the new technology. In the absence of uncertainty (β is very large), the answer is no, since no actor wants to deviate from the initial equilibrium. With increasing uncertainty about the new technology's performance, however, the chance that enough actors may find the new technology appealing also increases, to bring the system over the critical fraction of adopters where the performance curves cross. In that case, the system switches to the other equilibrium, with the whole population adopting the new technology. For any finite value of β, the crossover will occur with probability *1* at *some* point as *t* approaches infinity. But how long will it take in expectation?

We can use the probability distribution of ξ to characterize for any individual actor the probability of choosing technology *i* at time *t* via:

$$\eta_2(f) = \begin{cases} \frac{1}{2} e^{-\beta D(f)} & \textbf{\textit{if technology 1 is better,}} \quad f \le \hat{f} \\ 1 - \frac{1}{2} e^{\beta D(f)} & \textbf{\textit{if technology 2 is better,}} \quad f > \hat{f} \end{cases} \tag{4}$$

where $D(f) = P_1 + b_1 - P_2 - (b_1 + b_2) f$ is the "performance distance" between the two technologies. Furthermore, $\eta_1(f) = 1 - \eta_2(f)$. Equation (4) implies that the technology with a performance disadvantage is chosen with a probability that decreases exponentially with the size of the handicap. The choice probabilities depend on the uncertainty parameter β because high uncertainty (low β) makes the performance handicap less imposing. The probabilities also depend on the system state f because, with increasing penetration of the new technology, network externalities reduce the established technology's performance advantage.

In order to characterize the dynamic behavior of the system, we now introduce a system state evolution function g in the following way:

$$g(f) = (1-f)\eta_2(f) - f\eta_1(f). \tag{5}$$

The first summand in the brackets corresponds to the expected fraction of actors switching from the old to the new technology, and the second one to the expected fraction switching away from the new back to the old technology. Thus $g(f)$ is the net growth of f over one evaluation cycle. Multiplying g with the rate at which actors evaluate technology performance, we have thus heuristically argued that the expected change rate of the system state is described by

$$\frac{df}{dt} = \alpha g(f) = \alpha\left[(1-f)\eta_2(f) - f\eta_1(f)\right]. \tag{6}$$

We are now in a position to define a potential function $\Omega(f)$ as the integral over g:

$$\Omega(f) = \int_f^1 g(x)dx. \tag{7}$$

We show in the Appendix that the minima of this potential function, or the zeroes of the state evolution function g, indicate the most likely system states and thus characterize system equilibria. This can also be seen intuitively by noting that a large positive value of $\Omega(f)$ indicates high values of $g(x)$ for $x \geq f$, implying that the system state tends to move upward and has a low likelihood of reaching f. Thus, the system equilibria along with the minima of Ω are characterized by $g(f) = 0$ and $g'(f) < 0$.

We can write g and Ω in terms of the parameters in the following way, observing that $D(0) = P_1 - P_2 + b_1$.

$$g(f) = \begin{cases} e^{-\beta D(f)} - 2\mathrm{f} & f \leq \hat{f} \\ -e^{\beta D(f)} + 2(1-f); & f > \hat{f} \end{cases} \quad ; \quad \Omega(f) = \begin{cases} \dfrac{e^{\beta D(0)} - e^{-\beta D(f)}}{\beta(b_1+b_2)} + 2 - \frac{(1+f)^2}{2} + \frac{f^2 - \hat{f}^2}{2} & f \leq \hat{f} \\ \dfrac{e^{\beta D(0)} - e^{\beta D(f)}}{\beta(b_1+b_2)} + 2 - \frac{(1+f)^2}{2} & f > \hat{f} \end{cases} \tag{8}$$

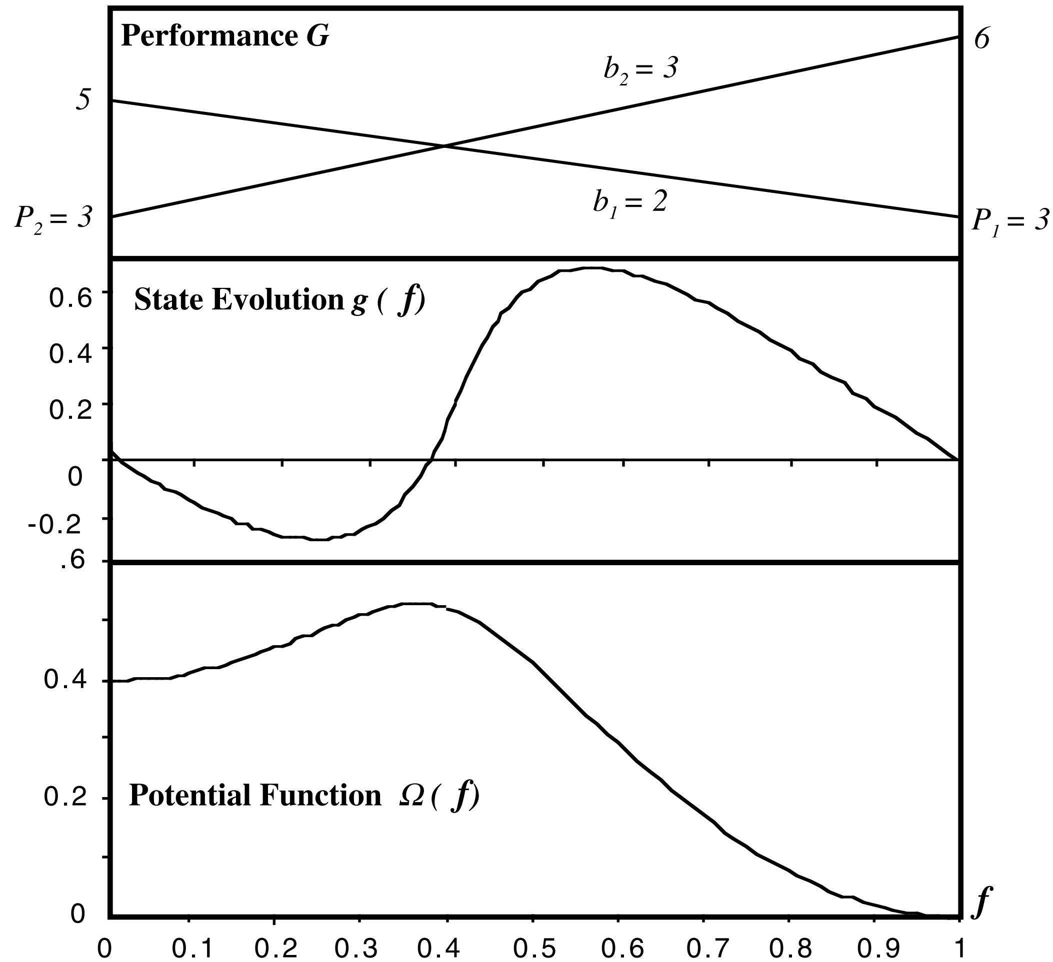

Figure 2: Example of Evolution and Potential Functions

We have plotted a numerical example of the payoff functions, the system state evolution function g and the potential function Ω in Figure 2. We see the two local minima of Ω, corresponding to the stable system equilibria, to the left and right, as we argued qualitatively at the beginning of this section. These equilibria are not exactly at the endpoints, since (depending on the uncertainty of β) there is always a small positive expected number of trials of the respectively inferior technology.

We are now in a position to describe the dynamic adoption behavior in our system. First, if the system is not in equilibrium, it moves (in expectation) toward the next (in a "downhill" direction in Ω) local minimum of Ω, governed by Equation (6). The solution of (6) is an exponential, so adoption settles down at this equilibrium exponentially fast. Here, $g = 0$, so no further movement occurs except random fluctuations from which the system relaxes back to the equilibrium. However, with very low probability a random fluctuation may be sufficient to push adoption to the other side of the maximum of Ω in Figure 2. If that occurs, the system will flip, and adoption will settle at the other equilibrium. This is the punctuated equilibrium behavior of sudden technology adoption.

In the Appendix we show that the steady state probability of a certain system state is proportional to *exp(-N Ω(f))*. Therefore, the expected time until a random fluctuation pushes the system to the other equilibrium is proportional[7] to *exp(β N ΔΩ)*, where *ΔΩ* denotes the "barrier" that needs to be crossed, that is, the difference in height between the maximum of *Ω* and *Ω* in the equilibrium from which the system starts.[8]

Thus, the system may flip from the metastable equilibrium $f = 0$ in Figure 2 after a long time, but it will take exponentially longer to flip from the global minimum (because of the higher barrier). Thus, the global minimum of Ω may be viewed as the final settling point of system adoption. Call f^* the value where $\Omega(f)$ achieves its maximum. Then we can write the expected time to switch from the left equilibrium as follows:

$$\textbf{\textit{Expected Time to Switch}} \cong e^{-\beta N[\Omega(f^*)-\Omega(0)]} = \exp\left\{\frac{\eta_2(f^*) - \eta_2(0)}{b_1 + b_2} N + \frac{N\beta f^{*2}}{2}\right\} \qquad (9)$$

Equation (9) tells us what drives the expected time to switch from the initial to the new equilibrium and how it may be influenced.[9] A large performance advantage (P_1 - P_2) of the established technology makes the "barrier" higher by requiring more trial uses to trigger a switch. Increasing ($b_1 + b_2$) simultaneously increases the barrier, but increasing b_2 alone makes the new technology more attractive and reduces the expected switching time. Low performance uncertainty of the new technology (a large value of β) makes it less likely that enough actors try out the new technology at any given time, increasing the expected time until the triggering level of trials occurs. Finally, a large population size N reduces the impact of any individual actor, and thus requires more trials at the same time. All of these parameters influence the expected switching time exponentially.

Thus, we have seen in the simplest version of our model how the system first moves to the closest stable point exponentially fast and then performs small random perturbations around it. However, the random fluctuations cause the system to switch between equilibria after a "long" period of time whose expectation is described by (9). Once the switch occurs, it is again exponentially fast and looks almost immediate in comparison to the time that has elapsed until the switch occurs.

In this simple version, we have held the performance of the two technologies constant over time. What happens when both technologies improve through incremental

---

[7] Equivalently, the probability that the system is pushed to the other equilibrium at a given point in time is *inversely proportional*

[8] A rigorous derivation of this relationship can be found in van Kampen 1981. The flip is formally equivalent to particle decay in a bistable potential.

[9] By the envelope theorem, we can analyze the marginal influence of the parameters on the expected switching time without considering the change of f^*. If uncertainty is very high or externalities very low, Ω may become flat, without a maximum. This corresponds to a system where the state f "wanders" randomly because users switch back and forth continuously. We do not further analyze this case, where behavior is dominated by uncertainty.

innovations? In order to answer this question, suppose the system starts out with only the old technology in use, as in Figure 2. For simplicity, suppose there is a constant difference of d in improvement rates between the two technologies, that is, $D(0)$ changes linearly over time. If d is positive, the old technology will pull ahead, and the likelihood of an equilibrium switch will continuously shrink over time. If, on the other hand, d is negative, the new technology will reduce the performance gap over time. In this case, adoption behavior depends on which of the following two things happens first: One possibility is that the new technology will completely catch up in performance before being fully adopted. However, externality benefits may keep the incumbent equilibrium in place, possibly even to the point where the new technology has become superior across the whole range of f, when the incumbent equilibrium in fact disappears. In this case, adoption of the new technology *trails* its performance evolution. A second possibility is that the performance catchup is very slow and that the equilibrium punctuation may occur first and *precede* the catching up of the new technology. In both cases, we may see punctuated equilibrium behavior, that is, the system switching from one stable equilibrium to another because of performance uncertainty.

Thus, we find two possible regimes of technology adoption, depending on how fast one technology catches up with the other, and on how much the system parameters support the possibility of equilibrium punctuation. In the next section, we present a more realistic representation of performance improvement over time and demonstrate the different adoption regimes with the help of simulations.

4. A Model With Usage Dependent Technology Improvement

In this section, we introduce a more realistic representation of incremental technology improvements over time. Suppose that actors evaluate technologies according to a Poisson process, as in Section 3. In the realistic case, perceived technology performance changes between evaluations not only because of positive externalities, but also because of incremental improvements along an S curve. The upper limit of the S curve implies an inherent performance ceiling that a technology cannot surpass, denoted by $\overline{P}_i$.

The improvement rate depends on three influences, which together determine how quickly the technology "learns." The first influence is the installed base of current users, since a larger market encourages improvement efforts. This source of improvement can be interpreted as "Learning by Doing," or the classical learning curve effect (Lévy 1965): it depends on the cumulative volume of technology use only and happens "automatically."

The second influence is current performance, as increasing knowledge about the technology opens up additional improvement opportunities. And third, improvement is proportional to the remaining gap to the technology's inherent performance ceiling, since achievement of the final portion of possible performance becomes increasingly difficult to attain. A small constant c_i determines how fast a technology advances towards its performance ceiling. c_i indicates the learning characteristics of the technology, such as the competence of its developers, or complexity. This is summarized as follows:

$$P_i(t) := P_i(t) + c_i f_i(t) P_i(t) \frac{\overline{P}_i - P_i(t)}{\overline{P}_i}. \tag{10}$$

Both technologies are endowed with a beginning performance value $P_i(0)$. At time 0, all actors use the "old" technology $(i = 1)$. At time t, an actor who chooses to perform a technology evaluation compares the alternatives according to:

$$G_i(t) = P_i(t) + f_i(t) b_i + \xi_i(P_i), \tag{11}$$

where $\xi_i(P_i(t)) = \xi_i * (\overline{P}_i - P_i(t)) / \overline{P}_i$, and ξ_i is the same random variable as in Section 3, with technology specific variance parameter β_i. Thus, technology performance uncertainty is high when the technology is still far from its performance limit and thus "new," and the uncertainty diminishes as the technology becomes well understood and advances toward its performance limit. Whenever an actor performs a technology evaluation, it adopts the technology with the higher perceived performance, increasing or decreasing $f(t)$ by $1/N$. The next time technologies are compared, the current performance levels are updated by the incremental improvements.

Below, we simulate two examples. Both exhibit a user population of $N = 20$ actors, who choose between the two competing technologies. Technology 1, the "old" technology, is initially used by the entire population, and it has an initial performance advantage of 2 units vs. 1 for its competitor, not counting network externalities. Each actor performs a technolgy evaluation at a rate of 0.05, or on average every 20 time units. Thus, over the whole population, evaluations happen at a rate of 1 per time unit. The old technology has perfectly known performance (zero uncertainty), and the new technology has an initial uncertainty of $\beta = 0.5$, corresponding to a performance variance of 1. Performance limits are (4, 5). Over time, the uncertainty of the new technology is reduced as it approaches its performance limit.

Figures 3 and 4 show the behavior of this system under two different parameter constellations, simulated over 750 time units.[10] The base performance of each technology (before network externalities) over time is shown (corresponding to the right-hand scale), as well as the evolution of the system state f (corresponding to the left-hand scale).

[10] Each Figure represents one sample path of the simulations. Although not formally describing average system behavior, these sample paths are typical. Descriptions of simulation averages and variances can be obtained from the authors.

Technology adoption in both examples is initially dominated by incremental improvements of the old technology, which enjoys the full user base and thus advances quickly toward its performance limit. Once it has reached the limit, it stagnates. The new technology is handicapped by a base performance shortfall as well as the network externalities benefits captured by the incumbent. It is therefore not capable of attracting a permanent user group. The stable system equilibrium is at $f_2 = 0$.

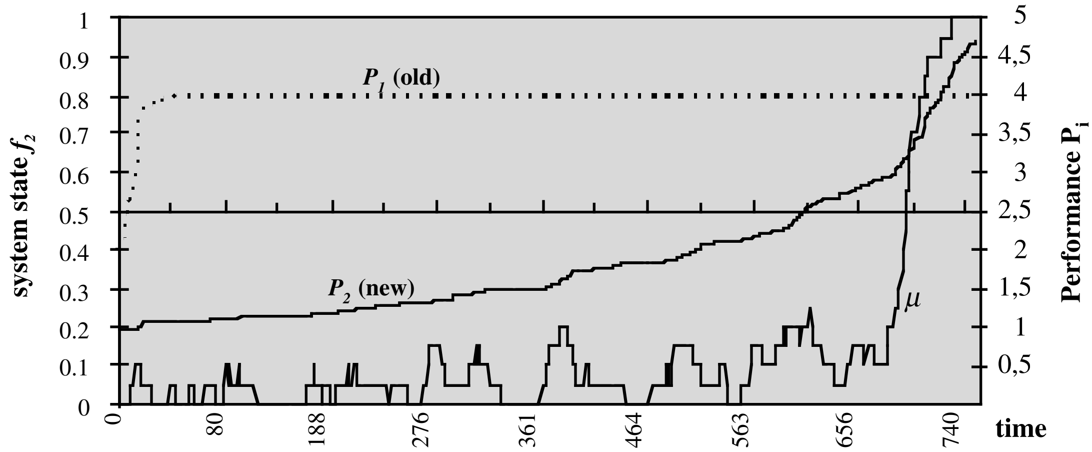

Example of a punctuated equilibrium, where the new technology gains acceptance before it has caught up in performance. The parameters are: P_i = (4, 5); c_i = (0.1, 0.04); b_i = (0.8, 5); $\beta = 0.5$; $\alpha = 0.05$; $N = 20$. The initial values of P_i are (2,1).

Figure 3: Punctuated Equilibrium Example

The system experiences random perturbations around the equilibrium, because users, from time to time, try out the new technology. They are prompted by a positive evaluation, which is made possible because of the new technology's performance uncertainty. However, all of these users eventually revert back to the incumbent technology and its superior performance.

Each adoption, even if temporary, helps the new technology to incrementally improve its performance. Thus, it slowly progresses over time, narrowing the gap to the incumbent. In Figure 3, the incremental improvement rate of the new technology is relatively slow ($c_2 = 0.04$). However, network benefits are greater for the new, then for the old technology (*5* vs. *0.8*). Thus, once the basic performance gap has narrowed sufficiently, externalities make full adoption of the new technology an equilibrium. In this case, the system may switch from one equilibrium to the other, even before the new technology has caught up in performance. This is what we see in Figure 3: Adoption precedes performance catch-up. After *680* time units, usage switches very quickly (within *50* time units) from almost full dominance of the incumbent to full adoption of the new technology. Only afterwards does the new technology catch up in base performance, driven by its large user base. This equilibrium is now permanently stable, since effective uncertainty disappears as the new technology reaches its performance limit and becomes incumbent. The equilibrium will persist until another new technology appears.

The situation in Figure 4 differs from that in Figure 3 by the presence of faster incremental improvements (c_2 = *0.15*), a higher performance limit of *6*, but weaker network externalities (b_2 = *1*) of the new technology. Now the new technology incrementally improves to the point of clear superiority, but the network externality benefit of the old technology maintains the incumbent equilibrium for another *150* time units before the system finally switches to the new permanently stable equilibrium. Thus, complete adoption of the new technology trails performance catch-up. Once the switch occurs, it happens very fast, within *35* time units.

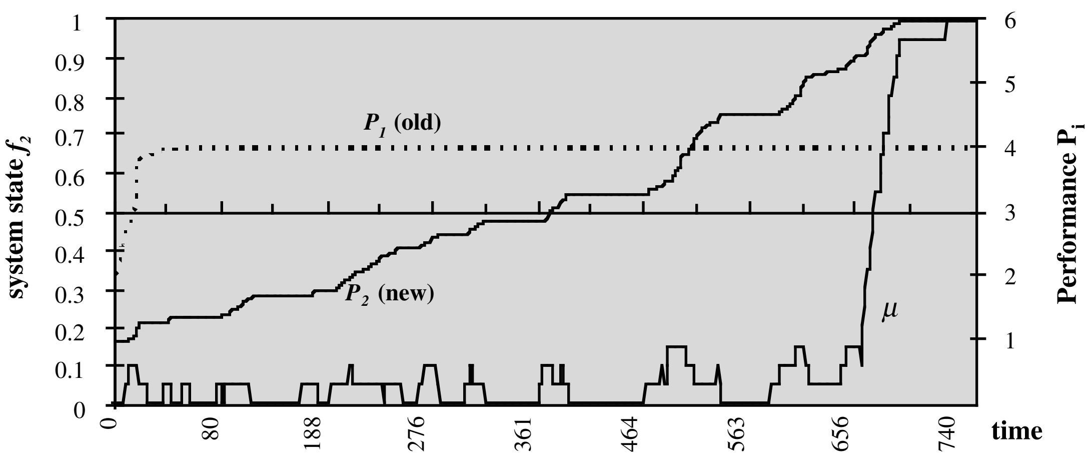

Example of a punctuated equilibrium, where the new technology first catches up in performance and only then gains acceptance. The parameters have been chosen such that the time horizon is the same as in Figure 3: P_i = (4, 6); c_i = (0.1, 0.15); b_i = (2, 1); β = 0.5; α = 0.05; N = 20.

Figure 4: Technology Catch-up Example

The two examples in this Section show how the analysis of the simple model in Section 3 illuminates system behavior in a more complicated model. Technology performance and adoption are governed by incremental improvements as well as equilibrium switches. Only incremental improvements are seen as long as one technology dominates. In addition, incremental innovations help a new, rising technology to narrow the relevant performance gap. Positive externalities may *bring about* a switch between equilibria, preceding dominance of the new technology, as seen in Figure 3. But externalities may also *delay* an equilibrium switch, as seen in Figure 4. Here, the new technology catches up before the residual performance uncertainty may allow a switch between equilibria. Adoption occurs only when the incumbent equilibrium has in fact disappeared, the only equilibrium remaining at full dominance of the new technology.

The more complete model in this section adds additional driving parameters, namely the learning rates c_i and the performance limits $\overline{P}_i$, to the previous parameters, $(P_1 - P_2)$, $(b_1 + b_2)$, β, and N. The former two parameters describe the relative performance progress (or learning) rates, and the latter the resistance of the system to switching. If the new technology learns fast and the system exhibits strong resistance, it is likely that the initial equilibrium will persist beyond the time when the new technology has achieved superiority (externalities delay the switch). If, on the other hand, the new technology learns more slowly and the system exhibits weak resistance to a switch, adoption may precede performance parity.

5. Diverse Technology Preferences and Performance Expectations

In this Section, we present two extensions of the model. The first extension intro-duces users that are not homogeneous in their performance evaluations.[11] The diffusion literature, for example, speaks of "innovators," with a positive bias toward a new technology, and "laggards," adopting only after everyone else, as specific groups in the general population.

One possibility of distinguishing between the user groups is to vary their externality coefficients. We introduce three different groups into the example from Figure 3. In our population of twenty users, one group of five innovators has a doubled externality benefit from the new technology, which makes it more attractive to them. The mass of the population (ten users) has the same externality benefits as the overall average population (and as in Figure 3). There are five conservative users with a bias against new technologies, who perceive no network benefit from the new technology, $b_2 = 0$. Thus, the average externalities over the whole population are unchanged.

[11] See also Huberman and Glance (1993b) who examine the effect of diverse users in a static framework.

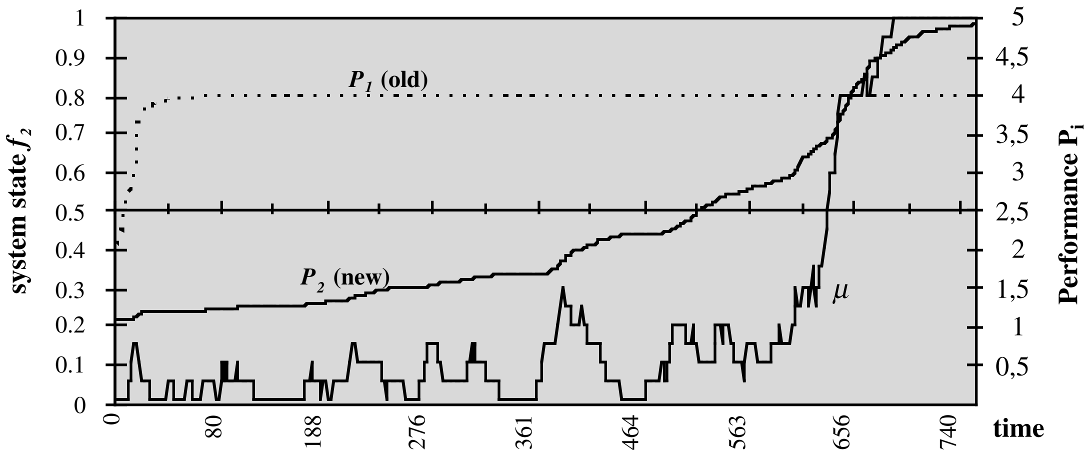

Example of a punctuated equilibrium with diverse actors. The parameters are the same as in Figure 3, with the exception of externalities: For one group of 5 innovators, b_i = (0.8, 10); for the 15 average actors in the middle, as before b_i = (0.8, 5), and for a group of 5 technology conservatives, b_i = (0.8,0).

Figure 5: Adoption Behavior With Diverse Actors

The resulting diffusion pattern is shown in Figure 5. Diffusion occurs significantly faster than in Figure 3: In a simulation of 100 runs, the average time point of punctuation time (measured as the time when $\mu=0$ is passed) is 506, with a standard deviation of 168. In Figure 3, in contrast, the average time of punctuation is 563 (with standard deviation 154). The two distributions are very close to normal, and the mean in Figure 5 is smaller at a statistical significance level of less than 1%.

The technology enthusiasts ensure a higher "base level" of usage early on for the technology (clearly apparent from inspecting the two Figures), and thus its performance can improve faster, which leads to earlier full adoption. Both the point of equal usage of both technologies and the time of full adoption are reached roughly *100* time units earlier. In the end, even the conservative group switches, because the old technology can no longer provide externality benefits, and the new technology is now superior. This behavior is consistent with previous models of diffusion, which also observe that diverse groups with differing adoption biases may speed up technology diffusion.

Note that, as in Figure 3, the adoption pattern represents a punctuated equilibrium with the equilibrium switch preceding performance parity. The system jumps over the barrier of the potential function Ω before the locally stable incumbent equilibrium (usage of the old technology) disappears.

The second extension that we present in this section introduces *expectations* by the users about the future performance evolution of the two technologies. Users often base their decision of whether or not to adopt a new technology, not on the current performance, but on the performance they expect in the foreseeable future. This is rational because they will use the technology for a period of time over which the actual performance changes.

However, the future performance is stochastic and unforeseeable. A simple and reasonable rule that adopters may use is *linear extrapolation* of the performance evolution observed over the recent past. That is, a user compares the observed technology performance today, at the current adoption decision, with the performance observed the last time it performed an evaluation. It then linearly extrapolates the trend for the same length of time into the future. The same extrapolation period is used since users evaluate technologies according to regular intervals of rate α.

The reader may note that this extrapolated trend includes learning as well as the effects of externality changes (caused by changed adoption numbers) and the effect of performance uncertainty (the random variable ξ). This reflects a situation where users compare the benefits they actually perceive from either technology, without analyzing in detail how much of the improvement may stem from learning and how much from having fellow users of the same technology. This "total perceived" performance is what is compared and extrapolated into the future.

Figure 6 shows the profound impact expectations may have on the diffusion pattern. All parameters are the same as in Figure 3, yet diffusion occurs radically faster. Expectations act as self-fulfilling prophecies by sustaining trends in adoption: If a technology happens to have shown an improvement, this will be extrapolated, causing more usage and thus more network externality benefits as well as faster incremental performance improvements, further enhancing the extrapolated trend.

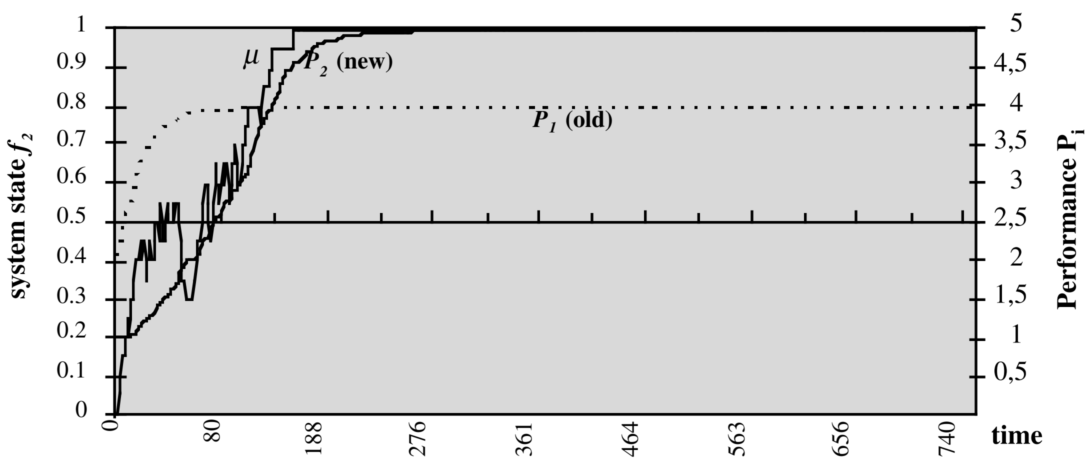

Example of a punctuated equilibrium with expectations. The parameters are the same as in Figure 3. Actors extrapolate the performance improvement since the last evaluation into the future and base their current decision on the extrapolation.

Figure 6: Adoption Behavior With Linear Expectations

The trend brings the new technology up to 50% adoption after only 60 time units, but then it reverses and pushes usage down again to 30% ($f_2 = 0.3$). Full adoption is reached after only 180 time units. Because of the large user base, the new techno-logy's base-performance P_2 also grows fast and overtakes the incumbent after 140 time units.[12]

The two examples in this section show how the system's resistance to switching equilibria may be dramatically lowered by diversity of the user population or by users basing their adoption decisions not on current performance, but on expectations about future performance.

6. Outlook and Conclusion

In this article, we presented an evolutionary model of technology diffusion. An old and a new technology are available, which both improve their performance incrementally over time. At the outset, the old technology is superior and used by the whole population in a stable equilibrium. The performance of the new technology cannot be perfectly evaluated and thus contains a random component. Both technologies exhibit positive externalities, or performance benefits that result from others using the same technology.

The new technology may gradually close its performance gap through trial usages and incremental improvements. However, it will not simply be adopted broadly when it achieves performance superiority. Network externalities may cause two stable usage equilibria to exist, one with the old technology being the standard and one with the new technology the standard. In this case, the existence of performance uncertainty may cause the system to switch between the equilibria.

Our results show that the expected time to adoption of the new technology is governed by two key characteristics: the rate of incremental improvement, or learning, of the new technology, and the system's resistance to switching between equilibria. Learning increases with the installed base of current users, current performance, the remaining gap to the technology's inherent performance ceiling, and the learning capability of the technology itself. The learning capability of the technology (and its owners) and its inherent performance limit are taken as exogenous parameters.

The system's resistance to switching equilibria increases exponentially if the size of the user population is large, if the network externality benefits are small, if the performance advantage of the established technology is large, and if the uncertainty about the performance of the new technology is low.

[12] As before, the pattern is quite stable; over 100 simulation runs, the average time of full adoption of the new technology was 180, with a standard deviation of 48.

Technology diffusion is determined by punctuations, or sudden shifts, between the two stable usage equilibria. We show with simulation examples that adoption may *precede* performance parity of the new technology if its learning rate is slow and the system's resistance to equilibrium switching is low. On the other hand, adoption of the new technology may *trail* performance parity, with the incumbent equilibrium being kept in place by its externality benefits, if the new technology learns fast, but the system's resistance to switching is high.

In two model extensions, we demonstrate two effects that may reduce the system's resistance to switching equilibria. One is diversity of users with respect to their appreciation of the new technology. The second is the existence of expectations about future performance, on which individual adoption decisions are based. Both may drastically speed up adoption behavior.

A number of additional features may be included in our model. One example is the existence of switching costs. If users pay a penalty every time they switch between technologies, adoption decisions, once made, will persist longer. A "hysteresis effect" will result which we expect to again make the overall adoption behavior of the system more conservative. A second example is adoption in growing markets: when the user population N grows, an equilibrium punctuation becomes more difficult over time, similar to the case of an uncertainty reduction. This may lead to an adoption occurring early or never.

Other interesting effects may be caused by users influencing one another through their networks of contacts. This makes externalities dependent on who the other users of the same technology are. Adoption may be influenced by these network structures in unexpected ways (see, for example, Abrahamson and Rosenkopf 1997). Also, expectations may be more sophisticated than the ones assumed in Figure 6. Finally, one may include a more highly developed model of each technology's learning behavior in the model.

We conclude that sudden technology adoption by systemwide punctuation of stable equilibria may happen without the presence of any resistance to change, technical compability issues, or constituency interests. The simple model presented in this paper offers a new mechanism which, through the presence of positive externalities and performance uncertainty alone, may lead to equilibrium punctuation. We offer several parameters that may allow management of an organization to influence whether the punctuation will precede or trail performance parity of the new technology, and thus suggest managerial actions to speed (or delay) adoption. The simple model may be extended to include the influence of learning, diversity, expectations, or organizational networks, on innovation diffusion.

References

Abernathy, W. J., and J. Utterback. Patterns of Industrial Innovation. *Technology Review* June-July 1978, 40 - 47.

Abrahamson, E., and L. Rosenkopf. Social Network Effects on the Extent of Innova-tion Diffusion: A Computer Simulation. Forthcoming in *Organization Science* 1997.

Arthur, W. B. Competing Technologies and Lock-In by Historical Small Events: The Dynamics of Allocation Under Increasing Returns. Center for Economic Policy Research, Stanford University 1983.

Arthur, W. B. Competing Technologies, Increasing Returns and Lock-In by Historical Small Events. Chapter 2 in: *Increasing Returns and Path Dependence in the Economy*, Ann Arbor: University of Michigan Press 1994, 13 - 32.

Anderson, P., and M. L. Tushman. Technological Discontinuities and Dominant Designs. *Administrative Science Quarterly* 35, 1990, 604 - 633.

Ceccato, H. A., and B. A. Huberman. Persistence of Nonoptimal Strategies. *Procedures of the National Academy of Sciences USA* 86, 1989, 3443 - 3446.

Clark, K. B. The Interaction of Design Hierarchies and Market Concepts in Technology Evolution. *Research Policy* 14, 1985, 235 - 251.

Cusumano, M.A., Y. Mylonadis, and R. S. Rosenbloom. Strategic Maneuvering and Mass-Market Dynamics: The Triumph of VHS over Beta. *Business History Review* 66, 1992, 51 - 94.

David, P. Why Are Institutions the Carriers of History? *Structural Change and Economic Dynamics* 5, 1994, 205 - 220.

Economist. Technology Brief: Being Digital is not Enough. *The Economist*, September 28, 1996, 108.

Ghemawat, P. Market Incumbency and Technological Inertia. *Marketing Science* 10, 1991, 161 - 171.

Granovetter, M. Threshhold Models of of Collective Behavior. *American Journal of Sociology* 83, 1978, 1420 - 1443.

Henderson, R., and K. B. Clark. Architectural Innovation: The Reconfiguration of Existing Product Technologies and the Failure of Established Firms. *Administrative Science Quarterly* 35, 1990, 9 - 30.

Huberman, B. A., and N. S. Glance. Evolutionary Games and Computer Simulations. *Proceedings of the national Academy of Sciences USA* 90, 1993a, 7716 - 7718.

Huberman, B. A., and N. S. Glance. Diversity and Collective Action. In: H. Haken and A. Mikhailov (eds.): *Interdisciplinary Approaches to Nonlinear Complex Systems*, Springer Series in Synergetics 62, 1993b, 44 - 64.

Kumar, U. and V. Kumar. Technological Innovation Diffusion: The Proliferation of Substitution Models and Easing the User's Dilemma. *IEEE Transactions on Engineering Management* 39, 1992, 158-168.

Lévy, F. K. Adaptation in the Production Process. *Management Science* 11*,* 1965, B136 - B154.

Liebowitz, S., and S. E. Margolis. Policy and Path Dependence -- From QWERTY to Windows 95. *Regulation,* 1995, 33 - 41.

Mokyr, J. *The Lever of Riches.* Oxford: Oxford University Press 1990.

Nelson, R., and S. Winter. *An Evolutionary Theory of Economic Change.* Cambridge: Harvard University Press 1982.

Reinganum, J. F. Innovation and Industry Evolution. *The Quarterly Journal of Economics,* 1985.

Rogers, E. M. The "Critical Mass" in the Diffusion of Interactive Technologies in Organiztions. Chapter 8 in: K. L. Kraemer (ed.): *The Information Systems Research Challenge: Survey Research Methods.* Boston: Harvard Business School Press 1985, 245 - 271.

Rosenkopf, L., and M. L. Tushman. The Coevolution of Technology and Organizations, in: J. A. C. Baum and J. V. Singh (eds.): *Evolutionary Dynamics of Organizations.* Oxford: Oxford University Press 1994, 403 - 424.

Sastry, M. A. Problems and Paradoxes in a Theoretical Model of Punctuated Organizational Change. MIT Working Paper 1994.

Schelling, T. C. *Micromotives and Macrobehavior.* New York: Norton 1978.

Tushman, M. L., and E. Romanelli. Organizational Evolution: A Metamorphosis Model of Convergence and Reorientation. *Research in Organizational Behavior* 7, 1985, 171 - 222.

Tushman, M. L., and P. Anderson. Technological Discontinuities and Organizational Environments. *Adminsitrative Science Quarterly* 31, 1986, 439 - 465.

Utterback, J. M. and F. F. Suárez. Innovation, Competition, and Industry Structure. *Research Policy* 22, 1993, 1-21.

van Kampen, N. G. *Stochastic Processes in Physics and Chemistry*. Amsterdam: North Holland, 1981.

Appendix

Theorem 1. The maximum of the equilibrium distribution $\Pi(f_2)$ of the system is reached at the minimum of $\Omega\,(f_2)$, which is characterized by the first-order condition $g(f_2) = 0$.

Theorem 1 is a special case of Ceccato and Huberman (1989). Take the system state as $n = f_2 N$, the number of adopters of the new technology. By the Poisson character of evaluations, no more than one actor switches technologies within a small time interval. Then the probability that a user defects is $\alpha\, n\, \eta_1(n/N)$ (since currently n actors use technology 2), and the probability that a new adopter is added is $\alpha\,(N - n)\,\eta_2(n/N)$. Thus, we can write the stochastic flow balance equations for the stationary distribution $P_e(n)$ as

$$P_e(n+1)\,\alpha\,(n+1)\,\eta_1(\tfrac{n+1}{n}) = P_e(n)\,\alpha\,(N-n)\,\eta_2(\tfrac{n}{N}) \quad \forall\, n \in (0, \cdots, N).$$

The solution of this system of $(N+1)$ equations is given by

$$P_e(n) = \binom{N}{n} \prod_{r=1}^{n} \left[\eta_2(\tfrac{r-1}{N}) \big/ \eta_1(\tfrac{r}{N})\right] P_e(0) \qquad \text{for } P_e(0) > 0\ .$$

We want to determine the maximum of this probability distribution. In order to write it in an appropriate form, we take logarithms on both sides:

$$\ln P_e(n) = \ln\binom{N}{n} + \sum_{r=1}^{n}\left[\ln\eta_2(\tfrac{r-1}{N}) - \ln\eta_1(\tfrac{r}{N})\right] + \ln P_e(0)$$

We approximate the right-hand side of this equation by applying Stirling's formula ($\ln x! \approx x\ln x - x$) and by approximating the sum in r by an integral in f_2. After some algebraic manipulations, we get

$$\ln P_e(n) \approx -N\left[\tfrac{n}{N}\ln\left(\tfrac{n}{N}\right) + (1-\tfrac{n}{N})\ln\left(1-\tfrac{n}{N}\right)\right] + N\int_0^{f_2} \ln\eta_2(x) - \ln\eta_1(x)\,dx \ + \ln P_e(0)$$

$$\approx N\left[-f_2\ln(f_2) - (1-f_2)\ln(1-f_2) + \int_0^{f_2} \ln\eta_2(x) - \ln\eta_1(x)\,dx\right] + \ln P_e(0)$$

With this approximation, we take the derivative:

$$\frac{d\ln P_e(n)}{df_2} \approx \frac{N}{2}\left[-\ln(f_2) + \ln(1-f_2) + \ln\eta_2(f_2) - \ln\eta_1(f_2)\right] = N\ln\left(\frac{(1-f_2)\,\eta_2(f_2)}{f_2\,\eta_1(f_2)}\right)$$

Setting the derivative to zero yields the first order condition (FOC): $f_1\,\eta_2(f_2) = f_2\eta_1(f_2)$. Thus, the zeroes of the function $g(f_2) = \alpha\,[f_1\,\eta_2(f_2) - f_2\,\eta_1(f_2)\,]$ describe the local extrema of the steady state distribution function $P_e(n)$.

The second order condition for a maximum of P_e requires that $g'(x)<0$ at x solving the FOC. This corresponds to a minimum of $\Omega(f_2)$ by Equation (7).