No. 2146

MODELS OF TECHNOLOGY DIFFUSION

Paul A Geroski

INDUSTRIAL ORGANIZATION

Centre for Economic Policy Research

ISSN 0265-8003

MODELS OF TECHNOLOGY DIFFUSION

Paul A Geroski

Discussion Paper No. 2146
May 1999

Centre for Economic Policy Research
90–98 Goswell Rd
London EC1V 7RR
Tel: (44 171) 878 2900
Fax: (44 171) 878 2999
Email: cepr@cepr.org

CEPR Discussion Paper No. 2146

May 1999

ABSTRACT

Models of Technology Diffusion*

The literature on new technology diffusion is vast, and it spills over many conventional disciplinary boundaries. This paper surveys this literature by focusing on alternative explanations of the dominant stylized fact in this area: namely, that the usage of new technologies over time typically follows an S-curve. The most commonly found model which is used to account for this model is the so-called epidemic model, which builds on the premise that what limits the speed of usage is the lack of information available about the new technology, how to use it and what it does. The leading alternate model is often called the probit model, which follows from the premise that different firms, with different goals and abilities, are likely to want to adopt the new technology at different times. In this model, diffusion occurs as firms of different types gradually adopt it. There are actually many ways to generate an S-curve, and the third class of models which we examine are models of density dependence popularized by population ecologists. In these models, the twin forces of legitimation and competition help to establish new technologies and then ultimately limit their take-up. Finally, we look at models in which the initial choice between different variants of the new technology affects the subsequent diffusion speed of the chosen technology. Such models often rely on information cascades, which drive herd like adoption behaviour when a particular variant is finally selected.

JEL Classification: L00, L60
Keywords: technology diffusion, epidemics, probit models,
density dependence, information cascades

Paul A Geroski
Department of Economics
London Business School
Sussex Place
Regents Park
London, NW1 4SA
UK
Tel: (44 171) 262 5050 x 3477
Fax: (44 171) 402 0718
Email: pgeroski@lbs.ac.uk

*I am obliged to Luis Cabral, Al Link, Paul Stoneman, Mariana Mazzucato, Bill Putsis, Stan Metcalfe, Paul David and a referee for help in various forms. The usual disclaimer applies. This paper will be published in the forthcoming special issue of *Research Policy* focusing on 'The Economics of Technology Policy'.

Submitted 12 April 1999

NON-TECHNICAL SUMMARY

It sometimes takes a long time for things to happen, and this is particularly the case in the area of technology diffusion. Many empirical studies of diffusion have observed a time path of adoption which resembles an S-curve: a slow period of early take-up is followed by a phase of rapid adoption and then a gradual approach to satiation (i.e. the rate of diffusion first rises and then falls over time). Much of the literature on technology diffusion has been built up around accounting for this empirical observation, and this paper surveys four leading models of this phenomenon.

By far and away the most popular account of the S-curve is based on the premise that the adoption of a new technology is limited by the diffusion of information about it. The diffusion process is, in this view, analogous to the process by which epidemics spread: each user of the new technology passes information on to one or more non-users who, in turn, adopt the technology and also spread the word. In the early phases of diffusion, most of the population are non-users, which means that passing information is easy but take up is slow because few users exist to pass the word. In the later stages of diffusion, many users exist to pass on the information, but their chances of meeting one of the few remaining non-users is low, and hence the rate of adoption is also low. In between, adoption rates are much higher since the many users are quite likely to meet one or more of the many non-users and convert them.

One of the main problems with the epidemic model is that information typically diffuses much faster than the use of new technology does; another is that the analogy with epidemics is misleading – potential users need to be persuaded and not just informed about the new technology. Possibly most damning of all, there are plenty of other reasons why particular firms might be faster or slower than other firms in adopting a new technology. Considerations such as these have led some scholars to apply probit models to the explanation of diffusion. The idea here is that each potential user has its own valuation of the new technology, giving some an incentive to adopt before others. As the costs of the new technology gradually come down, more and more potential users become actual users. The kinds of factors which create such differences between firms include: firm size, various types of switching costs, firm capabilities, etc.

The hallmark of an S-curve is an initial period in which the rate of adoption rises, followed by a period in which it falls. A third type of model focuses on accounting for what happens during these two phases. The first phase, sometimes called 'legitimation', describes the process by which a new technology becomes established, and its features become well known. Until this happens, take up rates will be low (because the new technology is perceived as quite risky); when it happens, however, take up rates rise. The second phase describes the effect that ‘competition’ has on take up. Early adopters of a new technology realize gains from being in a privileged position in product markets, but, as more and more firms adopt the new technology, the rents from early adoption become dissipated by the competition which occurs between using firms. This, in turn, tends to inhibit diffusion, lowering the rate of take up.

The final model that we consider is a model of the process by which a technology first arrives in a market. The point here is that new technologies generally come in a number of variants, and early adopters effectively choose between different variants. This choice process is both costly and risky and, as a consequence, it inhibits adoption, particularly when network externalities are present. However, when a choice has been made between the several variants present in the market, adoption rates rise (this is often propelled by a kind of bandwagon effect). This model has the great virtue of accounting for one important stylized fact: while successful inventions or innovations typically display an S-curve, most inventions or innovations are not successful. Further, a fair amount of casual evidence suggests that the process by which initial choices are made can have a big effect on subsequent diffusion rates. This is also a feature of this model.

I. INTRODUCTION

It is not easy to understand why things sometimes take a long time to happen, particularly when one views events with the benefit of 20:20 hindsight. In part, this lack of understanding is a reflection of how we think about social phenomena. For economists and others who use comparative statics (or equilibrium) analysis to answer questions about what will happen in given circumstances and why, the question of when that thing will occur is often not even regarded as an interesting question to ask. The problem of understanding how long things take to happen also reflects the inherent difficulty of the question: social phenomena involve many people making choices, often in an interdependent manner, and there are no basic reference points (like the speed of light) which can be used as a metric to measure the passage of time in such processes. Unlike molecules which act and react mechanically, people try to think before they act and this can be a very slow and unpredictable business for some of them.

The diffusion of new technology is a good example of this problem. Sometimes it seems to take an amazingly long period of time for new technologies to be adopted by those who seem most likely to benefit from their use. The literature which tries to explain why this happens is enormous, and it sprawls over several disciplinary boundaries. For many, the question of why things diffuse slowly has become very focused on a single stylized fact about that slowness, namely that the time path of usage usually follows an S-curve: diffusion rates first rise and then fall over time, leading to a period of relatively rapid adoption sandwiched between an early period of slow take up and a late period of slow approach to satiation. My goal is in this paper is examine how we typically think about what gives rise to S-curve diffusion patterns.[1] Mental models often have an amazingly powerful effect on how people think about particular phenomena, an effect that is sometimes stimulating and sometimes limiting. The premise behind this particular survey is the thought that if we are going to think creatively about public policies toward diffusion, we may need to think reflectively about how we think about technology diffusion.

The plan is as follows. Probably the most popular explanation of S-curve is an **epidemic model** of information diffusion, while the leading alternative is a **probit model** which argues that differences in adoption time reflect differences in the goals, needs and abilities of firms. I will explore these two ways of thinking about diffusion in Sections II and III below. I will also explore two other ways of thinking about diffusion. The first is drawn from the literature on organizational ecology, and argues that the primary drivers of S-curves are the processes of **legitimation** and **competition**. The second is as much a model of technology choice as it is one of technology diffusion, and it is based on the phenomena of **information cascades**, aided and abetted by network externalities. These last two models will be explored in Sections IV and V. I will close with some final reflections on what all of this might mean for technology policy in Section VI.

II. EPIDEMIC MODELS

The central feature of most discussions of technology diffusion is the apparently slow speed at which firms adopt new technologies.[2] If a new technology really is a significant improvement over existing technologies, it is important to ask why some firms shift over to it more slowly than other firms. Possibly the most obvious explanation is that they just find out about the new technology later than other firms do. If this is truly the case, one is likely to learn a lot about the time path of technology diffusion by studying the spread of information about it.

Suppose that there are N potential users of a new technology, and that each adopts the technology when s/he hears about it.[3] At time t, y(t) firms have adopted and $\{N - y(t)\}$ have not. Suppose further that information is transmitted from some **central source**, reaching α% of the population each period. If $\alpha = 1$, then the source contacts all N potential users in the first period, and diffusion is instantaneous. If, on the other hand, $\alpha < 1$, then information spreads gradually and so, therefore, does usage of the new technology. A transmitter that contacts α% of the current population of non-users, $\{N - y(t)\}$, at time t over the time interval Δt increases awareness (or usage) by an amount $\Delta y(t) = \alpha\{N - y(t)\}\Delta t$, and, taking the limit as $\Delta t \rightarrow 0$ and solving for the time path of usage,

$$(1) \qquad y(t) = N\{1 - \exp[-\alpha t]\}.$$

Equation (1) is a modified exponential function and Figure I plots it's time path (see the curve labelled A). Clearly, the smaller is α, the slower is diffusion and the smaller the number of users at any time (given that there are y(0) initial users). What is equally clear is that this particular information diffusion process does not produce the S-curve we expected to observe: it lacks an initial convex segment (curve B in Figure I is S-shaped).

This kind of model of information diffusion is not an implausible story of how people become aware of a new yoghurt product or news about the fall of the Berlin Wall. However, technology adoption often takes an order of magnitude longer than it takes for information to spread. To understand what lies behind this difference, it is useful to draw a distinction between the "*hardware*" and the "*software*" aspects of new technology (Rogers, 1995, pp.12). The hardware is the tool, machine or physical object that embodies the technology, while the software is the information base needed to use it effectively. Although some of the software can be transmitted impersonally through a users manual, much of the software of a particular technology is built up from the experience of using it, and at least some of that valuable knowledge will be tacit. As a consequence, it must be transmitted from person to person, and cannot effectively be broadcast from a common source. Thus, while the common source model embodied in (1) may usefully describe the transmission of information about the existence of a new hardware, it may not accurately trace flows of information about the associated software. And, without good software knowledge, many potential users will not adopt the new technology, however aware they are of it's existence.

To pass on software knowledge, potential users need to be able to communicate directly with current users who have accumulated experience with the new technology. This suggests that software knowledge may often follow a **word of mouth** information diffusion process in which the main source of information is previous users. Suppose that each existing user independently contacts a non-user with probability β. If there are $y(t)$ current users, then the probability that contact will be made with one of the $\{N - y(t)\}$ current non-users is $\beta y(t)$, meaning that usage will increase over the interval Δt by an amount $\Delta y(t)$

$= \beta y(t)\{N - y(t)\}\Delta t$. Assuming that there are $y(0) > 0$ initial users, taking the limit as $\Delta t \rightarrow 0$ and solving for the time path of usage yields

$$(2) \qquad y(t) = N\{1 + \phi \exp[-\kappa t]\}^{-1},$$

where $\kappa \equiv \beta N$ and $\phi \equiv (N - y(0))/y(0)$. Equation (2) is a logistic function and Figure I plots it's time path (see curve B). Smaller values of β mean smaller values of κ (for a given population, N) and, therefore, slower diffusion. It is clear that, unlike the common source model discussed above, this model traces out an S-curve over time: the rate of infection gradually rises (as the population of users gradually rises, increasing the aggregate stock of software information that can be passed on) until it hits a maximum at N/2, and then it declines (as non-users get increasingly hard to find and, therefore, to infect). As Figure I shows, this means that while usage increases year by year over time, it does so more rapidly in the early years after the introduction of a new technology than it does after the technology has become fairly well established.

Although the word of mouth model generates the kind of S-shaped diffusion curve that we are looking for, it has a serious weakness: it cannot explain the diffusion of an innovation from the date it is invented, but only from the date when some number, $y(0) > 0$, of early users have begun using it. Word of mouth diffusion processes can only begin to happen after an initial base of users has been built up, and, needless to say, the larger is this initial base of users, the faster is diffusion. Since early adopting individuals (or firms) have evidently chosen to use the technology despite not having had access to the experience of a previous user, it seems clear that they are somehow different from subsequent users.[4] This suggests that a more satisfactory model should distinguish between (at least) two different types of agents, a suggestion that we will explore shortly. Alternatively, one might say that these initial users are much more susceptible to common source information than subsequent users, who insist on receiving word of mouth information before they adopt. This second suggestion means that the right model of information diffusion might actually be a mix of (1) and (2).

Putting (1) and (2) together is pretty straightforward, although the resulting mathematics are not pretty. Over the time interval Δt, existing non-users are subject to two sources of information, and the probability that one of them will be informed (or infected) is $\{\alpha + \beta y(t)\}$. The first term in the brackets is the common source information which reaches users at a rate which is constant over time (for at least as long as the common source is transmitting). The second term reflects word of mouth diffusion in which the contact rate is dependent on the current size of the population of existing users. Performing the usual manipulations, the time path of adoption in this **mixed information source model** is

(3) $$y(t) = N\{1 - \exp[-(\alpha/\sigma)t]\}\{1 + \psi\exp[-(\beta/\sigma)t]\}^{-1},$$

where $\sigma \equiv \alpha/(\alpha + \kappa)$ measures the relative strength of the common source: if $\kappa = 0$, then no word of mouth transmission occurs and $\sigma = 1$, while if $\alpha = 0$ then the common source does not broadcast and $\sigma = 0$. Note that when $\sigma = 0$, $y(t) = 0$ for all t since no common source of information exists to create the initial user base that is needed to start a word of mouth process. When σ is "small" then the time path of y(t) will resemble the logistic curve shown as B in Figure I, but with an inflection point at $(N/2)\{(1-2\sigma)/(1-\sigma)\} < N/2$. As σ rises, the inflection point falls and the logistic curve becomes increasingly asymmetric, meaning that the upper, concave segment of the curve "stretches out"(we shall return to this point below), and the lower, convex segment shrinks. In the limit as $\sigma \to 1$, the time path of y(t) resembles the modified exponential function (curve A on Figure 1).[5]

In principle, it is not difficult to develop hypotheses about the potential determinants of β. Diffusion is likely to be faster for simpler technologies where software knowledge is easily learned and transmitted, for populations which are densely packed and where mixing is easily, where early users spread the word with enthusiasm (and do not die or forget what they have learned), and in situations where the new technology is clearly superior to the old one and no major switching costs arise when moving from one to the other. For many economists, these diverse factors often boil down to expected profits, learning and risk. The problem in empirical work is not with theory but with practice: most of these factors are difficult to observe and measure.[6] As a consequence, it is very difficult to see much in the way of really persuasive formal tests of the hypothesis that expected profits (or risk) drive diffusion rates (much less an assessment of just how much they really matter) in this literature, which is probably not a serious problem since no one really doubts the importance of expected profits and risk in principle. More to the point, most studies use a measure of learning which reflects the passage of time since the innovation was first introduced, and, at least in the early years of diffusion, this almost always has the kind of positive effect on diffusion rates or adoption decisions which one expects from epidemic models.

In many ways, it seems more natural to apply (1), (2) or (3) to the diffusion of information flows rather than to the diffusion of artefacts. For economists interested in the diffusion of new technology, the information flows of most interest have typically been technology spillovers; i.e. involuntary flows of information between rivals in the same market. A large empirical literature has pursued the question of whether such spillovers exist, and, if so, how large they are. In general, few doubt that spillovers occur, although it is not clear how fast the process takes place and what path the information flows take through the economy. A more recent literature which has tried to track flows of information by looking at patent citations is a little more illuminating. The focus in this work has typically been on trying to build up a map which identifies the main “information superhighways” in the economy, and it suggests that such knowledge flows are often localized. This is consistent with a word of mouth process, and suggests that the speed of diffusion might depend on how fast knowledge flows between different geographical regions. Further, most patents receive the majority of their citations shortly after they are issued, but some (particularly those in Drugs and Medicine) are sometimes cited for long periods of time (i.e 15-20 and more years). This suggests that the total number of citations of any particular patent is likely to follow an asymmetric S-curve over time: a very rapid initial rise in citations (for those which are cited at all) is followed by a slowing rate of citation over time which may only tail off after 25 or 30 years.[7] One way or the other, none of this is inconsistent with an epidemic model of information diffusion.

The basic hypothesis that we have been exploring is that it takes time for information about new technology to reach all potential users, and we have observed that different mechanisms of knowledge transfer – common source and word of mouth - affect the pattern of diffusion over time. However, these two models are rather simple, and the relatively even flow of information between individuals which they are built on is plausible only when applied to *homophilic* populations.[8] When populations are heterophilic, differences between individuals can impede the process of communication or, more likely, the process of persuasion. To understand diffusion in this context, one needs to understand which individuals are particularly influential and how they meet other individuals over whom their influence is decisive. It is hard to see much in the way of a general model (or class of models) emerging from the analysis of network structures[9], but one can get a sense of what might happen by considering a simple extension of the word of mouth model to two populations.

Suppose that there are **two populations**, N_1 and N_2, which do not interact with each other. Each has an initial number of users, $y_1(0)$ and $y_2(0)$, who initiate word of mouth diffusion processes with speeds β_1 and β_2 respectively. Following exactly the same argument as earlier, the increase in the total number of users, $y(t) \equiv y_1(t) + y_2(t)$, over the interval Δt is

(4) $$\Delta y(t) = \{\beta_1 y_1(t)\{N_1 - y_1(t)\} + \beta_2 y_2(t)\{N_2 - y_2(t)\}\}\Delta t.$$

It is relatively easy to extend (4) to the case where the two groups interact. Suppose, for example, that users in population 1 contact non-users in population 2 at a rate η_{12} while users in 2 contact non-users in 1 at a rate η_{21}. Then (4) can be written as:

(5) $$\Delta y(t) = \{[\beta_1 y_1(t) + \eta_{12} y_2(t)]\{N_1 - y_1(t)\} + [\beta_2 y_2(t) + \eta_{21} y_1(t)]\{N_2 - y_2(t)\}\}\Delta t,$$

which is very similar to the Lotka-Volterra model of competitive exclusion that is often discussed by population ecologists (e.g. see Roughgarden, 1996). It is relatively easy to imagine what (4) might look like. Suppose that the first population adopts first and at higher speed. Amongst other things, it then acts as a source of information for the second population which begins word of mouth diffusion at some time t*>0. The aggregate diffusion path is the vertical sum of the S-curves of these two populations, and it is likely to be an asymmetric S-curve with what looks like a long upper tail: aggregate usage will display a relatively quick initial diffusion (mainly driven by what happens within the first population), followed by a relatively slow subsequent approach to satiation (caused by the gradual increase in usage by the slower adopting second population). The aggregate speed of diffusion will depend on a weighted average of the two βs and ηs, while the overall limit to diffusion is the sum of the two populations.

One of the big problems with the epidemic model is that it takes N and β as fixed, and the two population model that we have just discussed is useful because it is an easy way to get round these drawbacks. In particular, it can be used to mimic a process in which β declines over time. There are any number of reasons why this might happen. One obvious possibility is that users become increasingly resistant to word of mouth communication (i.e. resistance to the disease increases and infection rates fall off); another is that late adopters may simply be less able to understand the new technology than early adopters. If, for example, the total population of potential users is composed of "*more able*" and "*less able*" firms, and the former start earlier and diffuse information faster (i.e. have a large β), then the diffusion of the new technology is likely to follow the time path of (4) or (5) with β falling over time as users from the second population become more numerous. The two population model can also be used to mimic a situation in which the total pool of potential users, N, is not fixed but increases over time. Suppose, for example, that there are two groups in the population: those for whom the new technology is ideally suited, and those for whom it initially does not work as well as existing alternatives. Further, imagine that the new technology gradually improves in a way which makes it increasingly suited to the needs of the second group. Then, diffusion is likely to proceed as follows: the process will start with the new technology diffusion amongst the first group, and then, when the technology has been modified enough to appeal to the second group, they will join in the process. This can carry on indefinitely: as diffusion spreads and the new technology develops or matures, N will gradually increase as usage spreads to other, more remote or peripheral populations.

Finally, it is worth emphasising that one rarely encounters symmetric S-curves in the actual diffusion of new technology. In almost all cases, the later stages of diffusion occur much more slowly than would be predicted by a symmetric S-curve. For example, Dixon, 1980, extended and reworked Griliches original data on the diffusion of hybrid seed corn and found that an asymmetric model like the Gompertz fitted the data better than the logistic model, equation (2), in 27 of the 31 US states examined. Similarly, Davies, 1979, observed that 8 of the diffusion processes in his sample of 22 displayed a positive skew, while 7 were symmetric (7 could not be easily classified). Asymmetry is a property of several of the models which we have examined thus far (and it will feature in many of those that we will be considering below), and it is worth listing some of it's more important causes. Asymmetry arises when: populations are heterogeneous and diffusion involves the introduction of progressively slower diffusion population groups; information processes involve the addition of common source information diffusion to word of mouth processes; the "infection rate", β, declines over time as knowledge depreciates, early converts lose some of their evangelical zeal or non-users increase their resistance to the new technology; and when the population of potential users increases over time, forcing diffusion to pursue a moving target that gradually drifts upward. And, finally, an information diffusion process in which $\Delta y(t) = \beta\{\log N - \log[y(t)]\}$ generates a Gompertz diffusion curve which also has positive skew.

So, where does all of this leave us? The basic premise of the epidemic model is that information diffusion drives technology diffusion. In a sense, it is hard to dispute this. However, the word of mouth diffusion process is clearly inconsistent with the data in many (sometimes not terribly important) ways, and one might, therefore, be tempted to reject it. At one level, fixing the simple word of mouth model is not all that difficult. It seems clear that a satisfactory epidemic model needs a good story of how the process of infection starts, and it probably must allow for imperfect mixing of some sort, endogenous decreases in β and increases in N. The net result is almost certain to be an asymmetric S-curve. However, a deeper problem with the model is that it is hard to believe that information diffuses as slowly as new technologies typically do. Evidently what matters is the type of information which users need before they are willing to adopt the new technology. There is an important distinction to be draw between understanding something and being persuaded, between hearing and acting on what you have heard. We have tried to capture this in drawing a distinction between "*hardware*" and "*software*", and there are doubtless other useful distinctions which can be drawn. However, the important point is that once one begins to think seriously about diffusion as a process of persuasion rather than simply as a process of spreading news, the analogy with epidemics begins to break down.

Some scholars have followed this line, and have tried to rework this model by focussing on risk or uncertainty rather than on the process by which people become informed about something. This usually results in a model with the property that some measure of uncertainty declines at increasing and then at decreasing rates over time, leading a pattern of technology adoption in which risk taking users rapidly, and then risk averse users more gradually, climb on board (i.e. leading to an S-curve). In fact, there is a basic equivalence between learning, expectations and uncertainty, one which becomes transparent when one thinks about the costs of knowledge transfer (including both transmission and receiving costs). If there are no knowledge transfer costs, learning will be instantaneous and knowledge will be complete. If, however, knowledge transfer is costly, agents will not acquire complete information about the new technology and they will be uncertain about just what new the technology does and how best to use it. Further, the ability of individuals to learn has an effect on learning costs, and their degree of risk aversion affects how they react to the uncertainty which they experience when their learning is incomplete. This observation suggests that it might be more useful to concentrate on the individual adoption decisions made by particular firms, examining (*inter alia*) the effects that information transfer costs, risk aversion and other firms specific factors have on the decision made by particular firms. Probit models are a natural way to do this.

III. PROBIT MODELS

Epidemic models abstract from differences in the goals, capabilities or actions of individual members of the population in order to focus on the diffusion of information in a simple, tractable, non-strategic setting. This is a particularly useful simplification whenever social structures affect information transmission, or when externalities and competitive effects of one type or another are created by the density of usage. However, it is a simplification. It is important not to lose track of the fact that the decision to adopt is a choice made by a particular individual (or firm), and that agents frequently make different choices for the best of reasons. It follows that differences between individuals may have a potentially important role to play in explaining patterns of diffusion. One natural way to think about this is by using a **probit model** to analyse individual adoption decisions.

Possibly the simplest way to think about an individual choice based model of diffusion is to suppose that individuals differ in some characteristic, x_i, which affects the profitability of adopting the new technology. Further, suppose that they will adopt if x_i exceeds some threshold level, x^*. Individuals differ in their characteristics, and we will suppose that x_i is distributed across the population according to some function f(x). Figure II shows two possibilities. In the top panel, the distribution of abilities across the population is normal. Those agents with levels of x_i larger than x^* choose to adopt (the shaded area), while the rest don't. Clearly, if x^* falls (i.e. shifts left) at a constant rate over time, the rate of adoption will gradually rise (as we climb up the right hand side of the distribution function) and then fall (as we go down the left side), generating an S-shaped diffusion curve. The lower panel of Figure II shows a situation where the distribution of abilities across the population is rectangular (or uniform). In this case, an S-shaped diffusion curve will come about if x^* falls at a rate which first rises and then falls over time (see David, 1969, for a classic discussion of this model).

To get a feel for how probit models work, it is worth dwelling on one particular model in some depth. Davies, 1979, argues that a firm will choose to adopt at time t if it's expected return π_i exceeds a threshold π^*; if, on the other hand, $\pi_i < \pi^*$, then adoption will not occur.[10] Expected returns and thresholds are not observable, but Davies supposes that both are simple functions of firm size, S; i.e. that $\pi_i / \pi^* = \theta S^{\upsilon}$. It follows that there is a critical size of firm $S^* \equiv (1/\theta)^{1/\upsilon}$ at which $\pi_i = \pi^*$. Hence, if $S_{it} > S^*$, then firm i will adopt the new technology at time t, while if $S_{it} < S^*$, it will wait until either it gets larger or S^* gets smaller before it adopts. This is the story told in Figure II when $x_i \equiv S_i$ and $x^* \equiv S^*$. The model has a very simple (but quite strong) prediction: when $\upsilon > 0$, then firms will adopt over time in order from largest to smallest, and conversely when $\upsilon < 0$.

As we saw earlier, the precise shape of the diffusion curve depends on how x_i is distributed across the population, and on how x^* changes over time. In the Davies model, firm size is assumed to be distributed lognormally (an assertion which is roughly consistent with the facts) and θ to follow a time trend. If $\theta = ct^{\psi}$ for a particular type of innovation (he calls these "*group A innovations*"), then it's diffusion curve will be cumulative log normal; if, on the other hand, $\theta = ce^{\psi t}$ ("*group B innovations*"), then it will be cumulative normal. Both are S-curves, but, as we saw on Figure I, they have slightly different shapes. Group A innovations are taken to be relatively transparent and their software can diffuse easily. As a consequence, θ falls quite rapidly. Group B innovations, on the other hand, require more search and may also experience regular post-innovation improvements which delay adoption, causing the fall in θ to occur much later after the innovation is first introduced than would be the case with a more transparent technology. The bottom line, then, is that group A innovations diffuse much more rapidly (the cumulative log normal diffusion curve is asymmetric, showing a rapid early rise and a point of inflection well below N/2) than group B innovations, which take-off more slowly but reach their plateau more quickly when they are established (the cumulative normal is a symmetric, S-shaped diffusion curve).

The trick with probit models is to identify interesting and relevant characteristics x_i. **Firm size** turns out to be a very commonly explored variable in the empirical literature on diffusion. This is partly because it is relatively easy to observe, and partly because it is typically taken as a proxy for all kinds of things: large firms are sometimes thought to be more capable (they may have higher quality or more technically able people on their staff), and, for this reason, they may be less likely to need word of mouth persuasion to adopt; they may use process innovations more intensively (e.g. on a larger scale) and so earn more profits from adopting than smaller firms would[11]; they might be less (or, for that matter, more) risk averse; they may be freer from financial constraints; they might have market power or be more inclined to strategically pre-empt smaller rivals; the new innovation might be complementary with other activities they undertake or be capable of being applied to a wider range of activities than would be the case if the adopting firm were specialised; and so on. Needless to say, these different interpretations of what firm size might mean are not always mutually consistent, and consequently it is hard to unambiguously interpret the empirical results which have been reported in the empirical literature. What is clear, however, is that the preponderance of the evidence suggests that, for one reason or another, large firms are, by and large, quicker imitators than small firms.[12]

Firm size is not the only interesting characteristic of firms which might be thought to drive decisions to adopt new technology. Anything which induces x^* to shift to the left, either at the same time as $f(x)$ shifts right or in the absence of such a shift, will make adoption more attractive for a firm. Here it might be useful to think of x_i as the net benefit of adoption, and x^* as the cost of acquiring the new equipment in which it is embodied. One set of important agents who will affect these cost-benefit calculations are **suppliers**. They are frequently responsible for facilitating the flow of information about the new technology, and, more generally, for marketing it. Their pricing and servicing policies have a direct bearing on the cost of new technology acquisition, and their success at designing a new technology which exactly meets the needs of the using population can often be the deciding factor between successful, rapid diffusion and outright failure. Finally, whatever technology they have designed and however they have chosen to market it, the learning process which suppliers undergo is likely to lead to a downward trajectory in prices which will push x* to the left at an initially high but subsequently declining rate. For a fixed distribution f(x), this is likely to induce an asymmetric S-curve describing downstream adoption. One way or the other, the important point is that conditions of competition upstream will affect diffusion downstream [13]

Suppliers are also interesting because they are often key players in the competition between the new technology and the older, existing technology which it displaces: in some cases, they control both technologies, while in others different groups of suppliers champion the different technologies. One way or the other, it is rarely the case that the existing technology which is being threatened remains static, and any sort of incremental innovations of the existing technology will clearly slow the diffusion of the new technology.[14] Similarly, the new technology is unlikely to arrive on the market in it's final form, and, in both cases, **technological expectations** are likely to have a major impact on diffusion: current or expect near future improvements in either the old or the new technology are likely to inhibit the diffusion of the new technology. Again, if x is the net benefit and x* the cost of acquiring the new technology, then any technological progress which makes the old technology more attractive, or lowers the benefits of adopting the new technology now (as opposed to adopting it in the near future) effectively raises the opportunity cost of adopting now; i.e. it shifts x* to the right.

One final class of exogeneous drivers of diffusion worth considering are costs. Those who come to probit models from the literature on epidemics will naturally focus in the first instance on **learning** and **search costs**. When they are first introduced, the benefits of adopting new technologies are often are often difficult to gauge with certainty, and they may seem too risky to be worth it.

However, as time passes (and usage increases), more information becomes available which enables firms to reassess the expected returns and risk involved. How fast this occurs depends, *inter alia*, on how firms learn (i.e. on how they update their prior information). Many models of this phenomena use statistical updating rules (such as Bayesian learning) which have the property that early bits of acquired information have a much bigger impact in changing prior views than information bits acquired when the firm has already undergone substantial learning. If x_i is a measure of the amount of risk a firm is willing to tolerate and x^* is the current estimate of risk, then learning models typically describe a process in which x^* shifts to the left at initially high but subsequently declining speeds.[15] Alternately, one might interpret x_i as the expected value of acquiring software information on a new technology and x^* as the search costs of doing so. Firms who initially have high expectations about the new technology (those in the shaded area of Figure II) are willing to make the investment, while others are unwilling to invest in search. However, as more firms become familiar with the new technology, search costs fall and x^* shifts to the left.

A variety of factors lock firms into existing technologies, raising **switching costs** and slowing the diffusion of new technologies. In terms of Figure II, think of x_i as a measure of the net benefits of adopting the new technology: firms with higher switching costs will have lower values of x_i (*ceteris paribus)* and are, therefore, less likely to pass the threshold x^* and adopt. These costs include the direct cost of software acquisition, something which ought to depend on a firm's ability to learn (sometimes called it's "*absorptive capacity*").[16] They also depend on the often very long learning process which a firm must go through in order to use the new technology to it's fullest. Since new process technologies often create new products or services (or, at least provide the means to differentiate existing goods and services more fully) and new products mean developing new markets, the costs of coming to terms with a new technology often include marketing and other costs incurred downstream. Finally, the more fundamental the break with previous activities caused by a new technology, the greater are switching costs. Firms that find it easier to spot costs than new sources of revenues may well may well be more reluctant to adopt a new technology than others.[17] Similarly, some new technologies will augment the competencies of a firm, strengthening them and making it more competitive, while others will disrupt existing competencies. In this later case, the costs of adopting a new technology include those associated with developing the new competencies needed to make the most of the new technology.[18] Switching costs can also be affected by government regulations.[19]

Opportunity costs are also important, and can be created by previous investments in machinery which have not fully depreciated. In particular, firms with new vintages of capital stock are less likely to switch to a new technology than firms with older, less valuable vintages, since the net benefits will be lower and the capital costs of change will be greater. This is particularly true for capital equipment that is so specialised that the costs incurred on installing it have been sunk, since, in this case, there are no second hand markets which the old equipment might be sold on to. Let x_i measure the age of a firm's capital goods, and suppose that when it's equipment is older than x^*, the firm switches to the new technology. The distribution of equipment age in the industry – shown as f(x) in Figure II – will gradually shift to the right, and the speed at which it does so will determine the speed of diffusion.[20]

Where does all of this leave us? One clear gain from thinking about diffusion using the probit model is that it enables one to generate a long and fairly impressive list of firm specific potential determinants of diffusion speeds. What is more, the link to decision making puts a certain precision on these arguments, and makes it possible (in principle) to identify a "*who*" and a "*why*" for each point on the diffusion curve. Finally, the fact that it identifies observable factors which, in certain circumstances, will trigger an adoption decision makes it possible (in principle at least) to identify a number of levers which policy makers can use to speed up (or slow down) the diffusion of particular technologies.

To put these gains into some perspective, it is worth distinguishing the **probit model of diffusion** from what one might call **population models of diffusion** (of which the epidemic model is a classic example). This latter type of model focuses on explaining the percentage of the population of firms who have adopted the new technology at any point in time. Population models have a natural appeal if one is primarily interested in the gradually unfolding impact that a new innovation has on markets, since the size of this impact depends (at least in the first instance) on aggregate usage, and not on which firms in particular are using it. Further, while probit models seem more natural and more attractive to economists (because they focus on individual decision making), they are less transparent than population models in describing phenomena which occur between individuals. The gradual increase of information available to potential users (or the decline in the risk which they perceive) appears to be exogenously driven in probit models; epidemic models at least have the virtue of making the true endogeneity of this phenomena absolutely plain to see. The real question here is whether diffusion is a social process that is something other than the sum of it's parts. Anyone who thinks that this might be the case will find the focus on (apparently exogenously determined) differences in firm characteristics in the probit model a little unsatisfactory.

All of this said, there is not much in the choice between population and probit models. There are no drivers of diffusion which feature in population models which cannot be expressed one way or the other in probit form. And, as we have seen, population models can be extended to allow for heterogeneous populations defined by differences in some characteristic x_i.[21] The really interesting choice, I think, is between different types of population models, and our next task is to explore some of these.

IV. LEGITIMATION AND COMPETITION

The analogy between word of mouth epidemic processes and the S-curve typically observed in the diffusion of new technologies is so well established in the literature that it is probably worth pointing out what should in any case be rather obvious: there are many different models which have nothing to do with information diffusion that can be used to generate an S-curve. Probably the leading alternative is the probit model which we have just discussed; others include the so-called "stock adjustment" model which has featured in several studies, and evolutionary models of diffusion.[22] There is, however, a fourth alternative derivation of the S-curve which derives from the population ecology literature.

Population ecologists use **density dependent** growth models to account for the systematic increases and decreases in net birth rates which they observe in natural settings. Suppose that y(t) is a count of the members of some particular population which inhabits a particular environmental niche, and that it increases at constant rate g. Then: $dy(t)/dt = gy(t)$, and, in principle, the population will eventually explode. This will never happen of course. Constraints imposed by the limitations of the niche as population density rises will depress birth rates, r_b, and raise death rates, r_d. If, for example, $r_b = b - k_b y(t)$ and $r_d = d - k_d y(t)$, then the net rate of increase in the population is $g \equiv r_b - r_d = r - rKy(t)$, where $r \equiv b - d$, and $K \equiv r/(k_b + k_d)$. As a consequence, population growth is given by: $dy(t)/dt = ry(t)\{1 - (y(t)/K\}$, which yields a logistic time path for y(t) of

(6) $$y(t) = K\{1 + \eta\exp[-rt]\}^{-1},$$

where $\eta \equiv \{[K/y(0)] - 1\}$. This is, of course, exactly the form of (2), but in this case the characteristic rise and fall of growth rates is caused by the effects of density on birth and/or death rates. The two parameters r and K in this model have natural interpretations: r is sometimes called "*the natural rate of increase*" of the population (i.e. that which would occur if there were no constraints on birth or death processes), while K is the "*carrying capacity*" of the niche (and gives an upper bound on the population size which can be supported by the niche). Both are directly analogous to the parameters β and N in the epidemic model (i.e. equation (2)).

An interesting extension of the density dependent population growth model has recently been advanced by sociologists studying the growth (and decline) in the populations of different types of organizations (Hannan and Freeman, 1989, Hannan and Carroll, 1992, and others). This model posits the existence of two forces affecting the birth and death rate of organizations over time: "***competition***" and "***legitimation***". Competition arises whenever resource constraints limit the number of organizations which can survive in a particular market (or social setting), and depends mainly on population density in these models. In the context of organizations, legitimation is the process by which a new type of organization becomes accepted, institutionalized or simply just taken for granted, and it clearly depends (amongst other things) on the number of such organizations already in existence. Plausible assumptions about the effects of competition and legitimation on birth and death rates produces a non-linear relationship between the net birth rate of the population and population density (rather than the linear dependence which underlies (6)), but, with some additional assumptions, the model can be made to generate an S-shaped or logistic curve tracing the number of firms alive in the market. Intuitively, the argument goes as follows. Legitimation gradually erodes barriers facing a new type organization, raising it's birth rate and increasing it's survival prospects. This drives a gradual increase in the (net) birth rate of the new organization. However, as the population of the new type of organization increases, a competition for resources sets in. This crowding is likely to have the effect of lowering birth rates and raising death rates, and, as a consequence, it lowers the rate of organizational expansion. This rise and fall in the (net) birth rate is, of course, just what underlies the S-curve.

This sort of story translates very naturally into the context of new technology diffusion. Consider the S-shaped curve labelled B on Figure I once again, and, for simplicity, divide the elapsed time of diffusion into two periods: the "*early period*" up until time $t = \tau$, and the *"late period"* which takes place after $t = \tau$ (there is, of course, no reason to focus only on the time it takes to reach the inflection point as the dividing point). In the early period, what matters to everyone involved with the new technology is whether it will work, whether it is superior to any other new technologies which might possibly arrive in the near future, whether there is a supply infra-structure available to support adopters, whether buyers will resist products made from the new technology, and so on. This legitimation process is clearly analogous to a standards setting processes, and that means that it's length is likely to depend on switching costs between the old standard and the new standard, the size of the installed base of new users and expectations about market growth and the future evolution of technology.[23] By time $t= \tau$ (or some such time), this legitimation process will be complete; i.e. the new technology will have become established. As it continues to be adopted, however, a second set of forces begins to limit it's diffusion in the market. As more and more firms begin to use the new technology, competition in the market for the goods or services which use this technology begins to lower the returns earned by early adopters, and this, of course, reduces the returns that non-users can expect if they adopt. This slows diffusion rates and ultimately brings the whole process to an end.

In simple ecology models, density dependence is the main driver of population dynamics. These models have the great virtue of providing a simple and tractable account of market dynamics, and the two forces which we have examined – legitimation and competition – help to account for the distinctive feature of S-curves: their initial convexity and subsequent concavity. However, density dependence is just too simple a story. Economic agents are not ants: their incentive structures are more complicated, and they often behave strategically in response to environmental pressures. This, unfortunately, means that the effects of legitimation and competition on technology diffusion can cut both ways.

The simplest stories about competition which are told by economists are density dependent in nature. In simple Cournot models of competition for example, profits per firm almost always declines as the number of firms (using a particular technology and selling similar goods) operating in the market rises. This kind of argument suggests that competition will slow diffusion rates in just the same manner as discussed above. There is, however, more to the story than this. If agents are at all foresighted, they will realise that the market will eventually get crowded, and they will wish to adopt the new technology before the returns from using it are competed away, pre-empting as many rivals as possible. This type of strategic behaviour may actually speed up diffusion, at least initially. Further, early users will try to create barriers to the entry of later adopters, and this may gradually slow the rate of adoption. [24] In fact economists have identified two other competitive forces that are likely to affect the timing decisions of firms: the "*pre-emption effect*" and "*rent displacement*". The first arises when the new technology complements the existing activities of some firms more than it does others, giving rise to an incentive for these more favoured firms to move first and adopt it before their rivals do. Rent displacement arises when the new technology cannibalizes some of a firm's existing activities, making adoption more costly than it would be in the absence of such activities. This argument is often used to explain why incumbents can be slower to adopt new technologies than new entrants (who have nothing to cannibalize), and it is likely to be part of any story about why market leaders who are champions of old technologies are often slower than others to adopt new competence displacing technologies. [25]

One might deduce from all of this that the net effect of competition on the rate of diffusion is ambiguous, and this is certainly a feature of the empirical literature which has looked at the effects of competition on diffusion rates. Not all of the statistical results point in the same direction, and most suggest that measures of competition like concentration ratios of counts of the number of firms in an industry experiencing diffusion are not all that important as a driver of diffusion.[26] It would, however, be imprudent to conclude from this work that competition has no effect on diffusion: the real issue is "*competition from who, or what?*". There is an extensive case study literature which suggests that incumbent firms are often very slow to adopt new technologies when entry barriers are high, and this suggests that it may be that it is competition from entrants (or threats of entry) which matters most in stimulating diffusion. This, of course, only makes plain what common sense suggests, namely that the degree of competition is likely to be endogenous to the process of diffusion. What is more, competition changes both it's character and it's intensity as diffusion proceeds. Initially competition is between the old technology and different variants of the new technology; when "*the*" new technology has been legitimized, competition is between the various firms who use the new technology to serve the market. In short, competition probably does speed diffusion rates, but the degree of competition felt by adopters and non-adopters at any time probably depends on the rate and extent of diffusion which has occurred up to that time.

Density dependence is also too simple to provide a really satisfactory account of the process of legitimation or standardization in markets, and for the same reason: economic agents often behave strategically, anticipating the effects of increased density and, as a consequence, altering the evolution of market structures. Standardization processes also generate externalities which can complicate market processes and either hasten or delay the development of a standard What is more – and what is more important – standardization processes involve making choices between alternatives, meaning that some technologies will fail while others succeed. More generally, it seems reasonable to believe the process of making choices between alternatives ought to have a profound affect on the time path of adoption of the technology which is ultimately selected. To make any kind of progress on this issue, however, means pushing well beyond the model of density dependence and looking at diffusion models which encompass both the choice between alternative new technologies, and the time path of imitation which follows that choice.

V. INFORMATION CASCADES

The literature on new technology diffusion is really a literature about S-curves, and in many ways this is rather limiting. S-curves have the virtue of being plausible (which is more than can be said for alternatives like instantaneous diffusion or linear adoption time paths), they can be a nice way to parameterize the diffusion process (making it straightforward to do empirical work on the determinants of diffusion speeds) and, last but by no means least, they are roughly consistent with the facts. However, they are only roughly consistent with the facts. As we have already seen, diffusion curves tend to be asymmetric in practice. More fundamentally, the fact is that most innovations fail (i.e. they do not diffuse at all), and it seems reasonable to insist that any serious model of diffusion ought to include failure as a possible outcome.

There are several ways forward from this observation, and, in what follows, I would like to focus on one particular type of model. My starting point is the observation that new technologies come to the market in a variety of forms, often leading to a small explosion in new products (or new product variants) based on that technology.[27] These new products are either sold directly to consumers or to downstream firms (where they appear as process innovations). "Adopting" a new technology in these circumstances involves choosing between these variants in the first instance (this is similar to the "legitimation" process discussed earlier), and then tracing the time path of imitation that results when one particular variant has been adopted (this is the conventional "diffusion" process). An extension of diffusion models to include an explanation of the initial choice between alternate variants of the new technology is worth considering for three reasons: it may be more descriptively accurate than (at least) the epidemic model; it will turn out to encompass several of the models discussed earlier, providing what might be a useful framework on which to hang a great number of more specific arguments about the determinants of diffusion; and, last but by no means least, it will identify when, and how, initial choices made in the market amongst the possible variants of the new technology have an effect on the diffusion of whatever variant is chosen (if, that is, any become established on the market).

Suppose that two variants of a new technology, A and B, simultaneously appear on the market and threaten to displace an existing technology. No one really knows for sure whether A is better than B or B is better than A, much less whether either is better than the existing technology. If, for some reason, early users are willing to experiment with the new technology and prefer A to B, then early trials with the new technology are likely to generate more information about A than B. If A turns out to be better than the existing technology, then it will gradually become more commonly used. These early adoption decisions are investment decisions, but as more and more information becomes available about A in particular, later adopters will be less and less willing to invest in making a serious choice between A and B: after all, if A seems to work better than the existing technology, why invest in B and take the risk that it will be worse? It follows that something of a bandwagon is likely to develop, with later adopters making the same choices as early adopters without having gone through the same investment in learning by experience. This process is sometimes referred to as an **information cascade**.[28]

When network externalities are present, they can strengthen these effects. After some point, A is likely to become more attractive than B regardless of it's intrinsic merits simply because A has a larger installed base. Moreover, the possibility of creating a similar sized installed base for B in the future will be smaller than it once was, since there are fewer potential adopters left and they are even less likely to choose B than A (in the future, A's installed base will be even larger and there will be even fewer potential adopters available to create a comparably sized installed base for B than there are now). For both of these reasons, the incentive to try B falls as the diffusion of A proceeds, and this effect will be stronger the more important are network externalities and the more marginal is the technical difference between A and B. In fact, network externalities can have two effects on diffusion: the lock-in effect just discussed, and a risk creating effect which can delay diffusion. When network externalities exist, early users risk making the “wrong” choice and becoming stranded with a technology which has failed to generate the network externalities it is potentially capable of. This may make early users reluctant to move first, and may delay the adoption bandwagon (this phenomena is sometimes called “*excess inertia*” in the standards literature).[29] The consequence will be an initial convexity in the time path of diffusion.

In fact, one might identify three phases in a diffusion process driven by information cascades: **the initial choice** between A and B, the **lock-in** to A, and then the **bandwagon** induced by imitation. As we have just seen, incentives to invest in information and network externalities help to explain the lock-in to A, but this is obviously not the whole story. One obvious reason why A might be chosen is that it appears on the market before B, and the expectations of early users at that time are that B is just not worth waiting for (indeed, the arrival of B might be a surprise). A might also be more effectively championed by it’s suppliers, who may even design A jointly with some of their major customers (typically large or symbolically important early users). For these or other reasons, A’s characteristics may suit the needs of early users better than B’s characteristics. Finally, suppliers may also play a strategically important role even when A and B arrive simultaneously on the market. A may be priced more economically, software information about A may be diffused more effectively or the support infra-structure (or other sources of network externalities) may be more effectively organized by A’s suppliers. All of this is to say that the initial choice between A and B may be hard to predict, and it may appear as if A were chosen “*by accident*”. As a consequence, the early time path of diffusion may be largely stochastic. However, once a choice is made and lock-in occurs, this early uncertainty is likely to die away, and the subsequent dynamics of the system (driven by an information cascade) will look much more deterministic. If it turns out that there are lots of firms waiting for others to make the initial choice between A and B, then, when it becomes clear that A has been chosen, a sudden burst of adoption will occur. As time passes and early users are succeeded by the imitative behaviour of the herd, the rate of adoption will begin to rise sharply and then, after the rush has passed, tail off. This, of course, can easily look a lot like an S-curve (although of course it doesn't have to).[30]

The important point is that choices which occur early on in the process may have an extremely powerful effect on the time path of diffusion. When the initial choice between A and B is made quickly and clearly and when A is clearly superior to the existing technology, then diffusion is likely to be rapid (quick and decisive decision making will quickly stampede the herd into action). If, however, these early choices are muddled, then the processes which generate and swell an information cascade are likely to be fragmented and weak. In these circumstances, one might observe only a very flat diffusion path (because it takes a while for a "winner" to be established"), incomplete diffusion (both A and B come to share the market) or no diffusion at all (A and B kill each other off amid confusion or general indifference by potential users). This point is, of course, not new. In path dependent processes, initial conditions matter. If diffusion is a path dependent process, then clearly it is hard to conceptually separate the process by which new technology spreads from the process by which economic agents make choices between the different variants of that new technology.

These arguments prompt a second general observation. In this model, S-curves are not the starting point of analysis, but they are one of several possible outcomes. Diffusion stories which are designed to explain the S-curve usually take the appearance of "*the*" new technology for granted, and focus on the question of why it takes so long for it diffuse, However, it is rarely clear to anyone at the time that "*the*" new technology has arrived, or which of several variants it is: it is only with the benefit of hindsight that "*the*" new technology stands out. Further, this clarity (or lack of it) is bound to have a profound effect in shaping (at least) the initial convexity which we observe early on in the diffusion of "*the*" new technology. In fact, most of the competing variants of "*the*" new technology are likely to fail, meaning that they have no S-curve. What this means that our observations of S-curves are coloured by sample selection bias: only successful innovations have an S-curve, and it is, therefore, by no means clear that the typical new innovation actually generates an S-curve. This, in turn, prompts the question: if our observations of the S-curve are the result of sample selection bias, should the S-curve be the centre-piece of our models?

Finally, notice that the model also says something about the extent of diffusion in the long run: if there are many more potential users of A than B, then the initial choice of A will also determine the eventual extent of diffusion: had early users opted for B, then the new technology would, in these circumstances, have diffused much less widely. Actually, market size is likely to endogenous to diffusion in a deeper sense than this. New products are typically targeted on specific users (e.g. those who prefer variant A), and then gradually adopted to other users (e.g. those who prefer variant B or C or D) as time passes. These adaptations look like post-innovation technological progress (which they are), but arguably the real action comes from the demand side and not the supply side: that is, they are market widening innovations. To put the point a slightly different way, the user population, N, is likely to gradually increase over time and more and more marginal agents gradually become users of the new technology. This is a model of diffusion which (like the probit model) says that diffusion is a phenomena which largely arises from the heterogeneity of user populations. However (unlike the probit model), the nature of the externalities which drive diffusion when information cascades are present mean that the most interesting and important users are the first users. Without them there to start the bandwagon, not much happens at all.

The focus on early events as prime determinants of subsequent diffusion patterns is a virtue of the kind of model which we have been examining in this section. However, there is a sense in which the initial choice between technologies A and B is not well explained in this model: we can describe the process by which choices are made, but it is very hard to predict which alternative will be selected. This observation completes a circle which brings us back to the epidemic model. In that model, the early pool of users – y(0) – who drive the model is exogenously given. These pioneering users have also made a choice, and their choice builds up a base of software information and reduces risks or increases the benefits perceived by subsequent users. In a sense, they "legitimize the innovation", and once that happens, an epidemic or information cascade drives subsequent adoption. This observation raises two questions: *should the "arrival" of "the" new innovation be dated from it's first appearance, or from the time it has gained legitimacy?*, and: *is the legitimation of a new innovation any less fundamental than the act of invention which brings it about?*. These two questions suggest that the real problem may not be understanding how the process of diffusion unfolds, but understanding how it starts.

VI. SOME REFLECTIONS ON TECHNOLOGY POLICY

We use models to help illuminate phenomena that we find difficult to understand, or to solve problems which are too difficult to think through. These benefits come because models simplify reality, and make it tractable enough for our limited powers of understanding to grasp. These benefits, however, also bring costs. Models can easily become prisons. They can severely limit the way in which we think about things, and so limit the range of actions which we might choose to take when we have completed our analysis. It is arguably the case that this has happened in the area of technology policy. The dominance of the epidemic model in particular seems to have created a set of policy presumptions which are surprisingly limited. Thinking about diffusion in terms of legitimation, competition and information cascades raises questions and issues which go well beyond the standard policy stance.[31]

The epidemic model is built on the presumption that diffusion happens too slowly, mainly because information does not diffuse fast enough amongst potential users. Anyone who really believes this will become interested in the question of whether public policy makers can do anything to improve the mechanism by which information spreads through the economy. Policy makers might become the common source, they might promote worth of mouth communication or subsidize the externalities involved with it, and they may try to identify key actors (those who are particularly persuasive) and try to motivate them or at least support their evangelical activities. If the key actors turn out to be users or suppliers, subsidies may rain down on them or policy makers may try to put together forums in which all parties can get together and communicate with each other. The bottom line seems to be that diffusion is a problem which public policy can ameliorate with a judicious mix of information provision and subsidies.[32]

If the probit model broadens the range of percieved policy options, it does so because it points to firms themselves as the source of the problem (if there is a problem). Firms often need to acquire special skills and they may lack enough incentives to move quickly. Aside from policies which speed up information diffusion or help suppliers to fly down their learning curves, this insight suggests that one might contemplate subsidies which encourage the building up of various types of human capital, and policies which stimulate competition (particularly by new entrants). In a sense, however, the probit model shuts down many policy options. If the problem really lies within firms, then there are real limits to what public policy makers can do short of running the firms themselves. Policies can be devised which make firms more aware of their opportunities and more able to exploit them, but it is hard to think of a policy which actually forces them to act when they don't wish to. One can subsizie all kinds of things, but that may not be enough.

The list of "exogeneous" drivers of diffusion which often feature in probit models also frequently includes variables which purport to reflect the effect that competition has on diffusion. At a common sense level, very few people doubt that a little bit of competition stimulates diffusion, but beyond that things become very murky. It is altogether possible that "*too much*" competition slows diffusion, either because it lowers the returns to adoption (the population of users becomes too dense) or because it muddles the initial choice between alternatives. Further, the nature of the firms who are the source of competitive pressures may matter: domestic rivals who are located close to other users may have much more effect than foreign rivals who compete through imports.[33] Finally, as we noted earlier, the nature of "competition" changes over the product life cycle, shifting from competition between alternative variants of the new technology to competition from a user population which gradually becomes more and more heterogeneous over time. All of these puzzles mean that competition policy (or strategic trade policy) are possible policy tools which can be used to stimulate diffusion, but they are possibly too blunt and indirectly associated with the diffusion process to be of much practical use.

Models of diffusion which focus on legitimation or information cascades open up several new insights into what the public policy problem associated with technology diffusion might be, and how it might be ameliorated. First and foremost, they destroy any clear presumption that diffusion is "*too slow*". The process of making choices between alternative technologies is a complicated one, and there are plenty of examples where market processes made choices too quickly and set in motion adoption processes which led agents to adopt second best technologies. Second, these models suggest that there is only a limited window in which policy can have important effects, and that is during the choice process. What speeds diffusion is that choices are made cleanly and clearly, and that the process of choice throws up enough information to create a strong bandwagon. However, once a choice has been made and the bandwagon has started, there are probably only limited effects which policy makers can have on what happens next. Third, although the policy window is small, the effects of policy are potentially very large. It is in the nature of information cascades (and many other externality driven processes) that small initial effects can have very large ultimate consequences. Since efficient and effective policy making should focus on situations where increasing returns (in this sense) exist, this observation reinforces the last: namely, that the timing of policy intervention may be at least as important as it's substance. Finally, since the choice process is inherently market specific, these observations suggest that technology policy must necessarily be selective if it is to have any substantive effects. Non-selective policies like subsidies or running technology fairs and forums are administratively convenient and they are consistent with the popular (but sometimes grossly misinformed) view that civil servants are indecisive, bureaucratic and totally ignorant of market realities. They are, however, very blunt policy tools, and it is hard to believe that one cannot do better.

Epidemic and probit models point to information provision and subsidies as the major tools of policy, and these alternative models add at least three further tools to the public policy portfolio. All of them concentrate on the process of choosing between alternative variants of new technologies. Standards setting processes are sometimes a an important way to resolve the many externalities which surround choice, and administrative processes can be very appealing when it is important not to choose too quickly (i.e. when the basic underlying technology is still evolving rapidly and unpredictably). Publicly sponsored (or at least sanctioned) pre-competitive consortia and R&D joint ventures are usually thought to be policy tools which stimulate the production of new technologies, but a moments reflect suggests that production choices have a profound effect on the nature and number of new variants which appear. Further, new technology production processes which operate in close tandem with users are likely to lead to a much quicker take up of new technologies than those which happen in isolated R&D labs. Finally, public procurement is, in principle, a powerful tool of technology policy. Governments are very heavy users of new technologies, and they are often well large, informed and rather insensitive to price. This makes them important potential agents in the bandwagon process, particularly when they insist that technology developments associated with their purchases are put in the public domain.

Perhaps more fundamentally, legitimation and information cascade models of diffusion challenge the basis of the commonly made distinction between technology policies which are oriented towards the generate of new technology and those which are oriented towards diffusion.[34] This distinction is sometimes used to make the helpful point that technology policy can be more than just a matter of shovelling funds into big science projects or uncritically supporting high tech defence spending. Further, this distinction directs attention to the sometimes distressingly large hinterland of firms who operate well away from the best practice frontier. It might well be much more sensible to spend money trying to move these firms to the frontier than trying to push the frontier further away from them. However, all of this said, it remains the case the problem of diffusion is not just one of slow imitation. Diffusion is about matching new technology to what is usually a wide range of different user needs, and this issue is as important on the day when a scientist or engineer first starts speculating about what might be as it is on the day when the last potential user finally catches up with what is. Diffusion is as much a process by which new technologies are developed as it is a process by which usage spreads, and this means that there is probably not a hard and fast distinction to be drawn between technology policies design to generate new technology and those designed to increase the usage of existing technologies.

REFERENCES

Arthur, B. (1989) “Competing Technologies, Increasing returns and Lock-in by Historical Small Events”, Economic Journal, 99, 116-131.

Arthur, B., Y. Ermoliev and Y. Kaniovski (1987) “Path Dependent Processes and the Emergence of Macrostructure”, European Journal of Operations research294 - 303.

Balcer, Y. and S. Lippman (1984) “Technological Expectations and Adoption of Improved Technology”, Journal of Economic Theory, 34, 292-318.

Bannerjee, A. (1992) “A Simple Model of Herd Behavior”, Quarterly Journal of Economics, 107, 797-817.

Baptista, R. (1999) The Diffusion of Process Innovations: A Selective Survey”, International Journal of the Economics of Business, 6, 107-130.

Bartholomew, D. (1973) Stochastic Models for Social Processes, 2nd ed., John Wiley, New York.

Bass, F. (1980) “The Relationaship between Diffusion Raytes, Experience Curves and Demand Elasticities for Consumer Durables Technological Innovations”, Jhournal of Business, 53, 551-567.

Beath, J., Y. Katsoulacos and D. Ulph (1995) “Game Theoretic Approaches to the Modelling of Technological Change”, in P. Stoneman, ed., Handbook of the Economics of Innovation and Technical Change, Basil Blackwell, Oxford.

Bikhchandani, S., D. Hirschleifer and I Welch (1998) “Learning from the Berhaviour of Others: Conformity, Fads and Informational Cascades”, Journal of Economic Perspectives, 12, 151-170.

Bikhchandani, S., D. Hirschleifer and I. Welch (1992) “A Theory of Fads, Fashion, Custom and Cultural Change as Informational Cascades”, Journal of Political Economy, 100, 992-1026.

Cabellaro, R. and M. Hammour (1994) “The Cleansing Effects of Recession” American Economic Review, 84, 1350-1368.

Cabral, L. (1990) “On the Adoption of Innovations with Network Externalities”, Mathematical Social Sciences, 19, 299-308.

Chatterjee, R., J. Eliashberg and V. Rao (1998) “Competitive Diffusion Models”, mimeo, University of Pittsburgh.

Choi, J. (1997) “Herd Behaviour, the Penguin Effect and the Suppression of Informational Diffusion”, Rand Journal of Economics, 28, 407-425.

Chow, G. (1967) “Technological Change and the Demand for Computers”, American Economic Review, 57, 1117-1130.

Christensen, C. (1997) The Innovator’s Dilemma, Harvard Business School Press, Boston.
Cohen, W. and D. Levinthal (1989) “innovation and Learning: The Two Faces of R&D”, Economic Journal, 99, 569-596.

David, P. (1969) “A Contribution to the Theory of Diffusion”, mimeo, Stanford University and forthcoming as Chapters 1-3 in his Behind the Diffusion Curve, Westview Press, Boulder, Colorado.

David, P. (1975) “The Mechanization of reaping in the Ante-Bellum Midwest”, in his Technical Change, Innovation and Economic Growth, Cambridge University Press, Cambridge.

David, P. (1986) “Technology Diffusion, Public Policy and Industrial Competitiveness”, in R. Landon and N. Rosenberg (eds), The Positive Sun Strategy, National Academy Press, Washington DC.

David, P. and S. Greenstein (1990) “The Economics of Compatibility Standards”, Economics of Innovation and New Technology, 1, 3-42.

David, P. and T. Olsen (1986) “Equilibrium Dynamics of Diffusion when Incremental Technological Innovations are Foreseen”, Richerche Economiche, 40, 738-770.

David, P. and T. Olsen (1992) “Technology Adoption, learning Spillovers and the Optimal Duration of Patent Based Monopolies”, International Journal ofd Industrial Organization, 10, 517-544.

Davies, S. (1979) The Diffusion of Process Innovations, Cambridge University Press, Cambridge.

Debresson, C. and F. Amesse (1991) “Networks of Innovations”, Research Policy, 20, 363-380.

De Vany and Walls, (1996) “Bose-Einstein Dynamics and Adaptive Contracting in the Motion Picture Industry”, Economic Journal, 106, 1493-1514.

Dixon, R. (1980) “Hybrid Corn Revisited”, Econometrica, 48, 1451-146.

Ergas, H. (1987) "The Importance of Technology Policy", in P. Dasgupta and P. Stoneman (eds), Economic Policy and Technological Performance, Cambridge University Press, Cambridge

Farrell, J. and G. Saloner (1985) "Standardization, Compatibility and Innovation", Rand Journal of Economics, 16, 70-83.

Farrell, J. and G. Saloner (1986) "Installed Base and Compatibility", American Economic Review, 76. 940-955.

Freeman, C. (1986) The Economics of Industrial Innovation, 2nd ed., MIT press, Cambridge Mass.

Gandal, N., M. Kende and R. Rob (1999) The Dynamics of Technological Adoption in Hardware/Software Systems: the Case of Compact Disk Players", CEPR Working Paper, London..

Geroski, P. (1995) "Do Spillovers Undermine the Incentive to Innovate?", in S. Dowrick, ed, Economic Approaches to Innovation, Edward Elghar, Aldershot.

Geroski, P. (1995) "Markets for Technology: Knowledge, Innovation and Appropriability", in P. Stoneman, ed, Handbook of the Economics of Innovation and Technical Change, Basil Blackwell, Oxford.

Geroski, P. and P. Gregg (1998) Coping With Recession, Cambridge University Press, Cambridge.

Griliches, Z. (1957) "Hybrid Corn: An Exploration in the Economics of Technical Change", Econometrica, 48, 501-522.

Griliches, Z. (1992) "The Search for R&D Spillovers", Scandinavian Journal of Economics, 94, Supplement, 29-47.

Grubver, H. and F. Verboven (1999) "The Diffusion of Mobile Telecommunications Services in the European Union"" CEPR Working Paper, London.

Hannan, M. and J. Freeman (1989) Organizational Ecology, Harvard University Press, Cambridge, Mass.

Hannan, M. and G. Carroll (1992) Dynamics of Organizational Populations, Oxford University Press, Oxford.

Hannan, T. and J. McDowell (1984) "The Determinents of Technology Adoption: The Case of the Banking Firm", Rand Journal of Economics, 15, 328-335.

Harley, C. (1973) "On the Persistence of Old Techniques: The Case of North American Wooden Shipbuilding", Journal of Economic History, 33, 372-398.

Henderson, R. (1994) “The Evolution of Integrative Capability: Innvation in Cardovascular Drug Discovery”, Industrial and Corporate Change, 3, 607-630.

Iansiti, M. and K. Clark (1994) “Integration and Dynamic Capability: Evidence from product Development in Automobiles and Mainframe Computers, Industrial and Corporate Change, 3, 557-606.

Ingham, H. and S. Thompson (1993) “The Adoption of New Technology in Financial Services: The Case of Building Societies"” Economnics of Innovation and New Technology, 2, 263-274.

Ireland, N. and P. Stoneman (1982) “The Role of Supply Factors in the Diffusion of New Process Technology”, Economic Journal, Conference Paper, 65-77.

Jaffee, A. and M. Trajtenberg (1996) “Modelling the Flows of Knowledge Spillovers”, NBER Working paper, Cambridge Mass.

Jaffee, A., M. Trajtenberg and R. Henderson (1993) “Geographical Localization of Knowledge Spillovers as Evidenced by Patent Citations”, Quarterly Journal of Economics, 108, 577-598.

Jensen, R. (1982) “Adoption and Diffusion of an Innovation of Uncertain Profitability”, Journal of Economic Theory, 27, 182-193.

Jensen, R. (1983) “Innovation Adoption and Diffusion when there are Competing Innovations”, Journal of Economic Theory, 29, 161-171.

Jovanovic, B. and G. MacDonald (1994) “Competitive Diffusion”, Journal of Political Economy, 102, 24-52.

Jovanovic, B. and S. Lach (1989) “Entry, Exit anmd Diffusion with Learning by Doing”, American Economic Review, 79, 690-699.

Kapur, S. (1995) “Technology Diffusion with Social Learning” Journal of Industrial Economics, 43, 173-195.

Karshenas, M. and P. Stoneman (1993) “Rank, Stock, Order and Epidemic Effects in the Diffusion of New Process Technologies”, Rand Journal of Economics, 24, 503-528.

Karshensas, M. and P. Stoneman (1995) “technological Diffusion”, in P. Stoneman (ed) Handbook of the Economics of Innovation and Technical Change, Basil Blackwell, Oxford.

Katz, M. and C. Shapiro (1985) “Network Externalities, Competition and Compatibility”, American Economic Review, 73, 424-440.

Lekvall, P. and C. Wahlbin (1973) “A Study of Some Assumptions Underlying Innovation Diffusion Functions”, Swedish Journal of Economics, 75, 362-377.

Levin, S., S. Levin and J. Meisel (1987) “A Dynamic Analysis of the Adoption of a New Technology: the case of Optical Scanners”, Review of Economics and Statistics, 69, 12-17.

Link, A. and P. Kapur (1994) “A Note on the Diffusion of Flexible Manufacturing Systems Technology”, Economics Letters, 46, 369-374.

Macleod, C. (1992) “Strategies for Innovation: The Diffusion of New Tecxhnology in Nineteenth century British Industry”, Economic History Review, 45, 285-307.

Mahajan, V., E. Muller and F. Bass (1990) “New Product Diffusion Models in Marketing: A Review and Directions for Research”, Journal of Marketing, 54, 1-26.

Mansfield, E. (1961) “Technical Change and the Rate of Imitation”, Econometrica, 29, 741-766.

Mansfield, E. (1963) “The Sped of Response of Firms to New Technologies”, Quarterly Journal of Economics, 77, 290-311.

Mansfield, E. (1989) “Industrial Robots in Japan and the USA”, Research Policy, 18, 183-192.

McWilliams, B. and D. Zilberman (1996)”Time of Technology Adaption and Learning by Using”, Economics of Innovation and New Technology, 4, 139-154.

Metcalfe, S. (1970) “Diffusion of Innovation in the Lancashire Textile Industry”, Manchester School, 38, 145-162.

Metcalfe, S. (1981) “Imp[ulse and Diffusion in the Process of Technological Change” Futures, B, 347-359.

Metcalfe, S. (1988) “The Diffusion of Innovation: An Interpretative Survey”, in G. Dosi et al (eds) Technical Change and Econbomic Theory, Frances Pinter, London.

Metcalfe, S. (1995) The Economic Foundation of Technology Policy: Equilibrium and Evolutionary Perspectives”, in P. Stoneman (ed) Handbook of the Economic of Innovation and Technical Change, Basil Blackwell, Oxford.

Midgley, D., P. Morrison and J. Roberts (1992) “The Effect of Network Structure in Industrial Diffusion Processes”, Research Policy, 21, 533-552.

Mokyr, J. (1990) The Lever of Riches, Oxford University Press, Oxford.

Oster, S. (1982) “The Diffusion of Innovations Among Steel Firms: The Basic Oxygen Furnace”, Bell Journal of Economics, 13, 45-56.

Oster, S. and J. Quigley (1977) “Regulatory Barriers to the Diffusion of Innovation: Some Evidence from Building Codes”, Bell Journal of Economics, 8, 361-377.

Oulton, N. (1989) “Productivity Growth in Manufacturing, 1963-85: the Roles of New Investment and Scrapping”, National Institute Economic Review132, 71-91.

Pennings, J. and F. Harianto (1992) “The Diffusion of Technological Innovation in the Commerical Banking Industry”, Strategic Management Journal, 13, 29-46.

Porter, M. (1990) The Competitive Advantage of Nations, Macmillan Press, London.

Postell, S. (1990) “Competing Networks and Proprietary Standards: The Case of Quadraphonic Sound”, Journal of Industrial Economics, 39, 169-185.

Quirmbach, H (1986) “The Diffusion of New Technology and the Market for Information”, Rand Journal of Economics, 17-33-47.

Reinganum, J. (1981a) “Market Structure and the Diffusion of New Technology”, Bell Journal of Economics, 12, 618-624.

Reinganum, J. (1981b) “On the Diffusion of New Technology: A Game Theoretic Approach”, Review of Economic Studies, 48, 395-405.

Reinganum, J. (1989) “The Timing of Innovation: Research, Development and Diffusion”, in R. Schmalensee and R. Willing, eds., Handbook of Industrial Organization, Volume I, Elsevier, Amsterdam.

Rogers, E. (1995) Diffusion of Innovations, 4th ed., The Free Press, New York.

Romeo, A. (1975) “Inter-industry and Inter-firm Differences in the Rate of Diffusion of an Innovation”, Review of Economics and Statistics, 57, 311-319.

Romeo, A. (1977) “The Rate of Imitation of a Capital Embodied Process Innovation”, Economica, 44, 63-69.

Rose, N. and P. Joskow (1990) “The Diffusion of New Technologies: Evidence from the Electric Utility Industry”, Rand Journal of Economics, 21, 354-373.

Rosenberg, N. (1976) “On Technological Expectations”, Economic Journal, 86, 223-235.

Roughgarden, J. (1996) The Theory of Population genetic and Evolutionary Ecology, Prentice Hall, Upper Saddle River, New Jersey.

Schmalensee, R. (1982) “Product Differentiation Advantages of Pioneering Brands”, American Economic review, 72, 349-365.

Silverberg, G., G. Dosi and L. Orsenigo (1988) “Innovation, Diversity and Diffusion: A Self-Organizing Model”, Economic Journal, 98, 1032-1054.

Stoneman, P. (1976) Technological Diffusion and the Computer Revolution, Cambridge University Press, Cambridge.

Stoneman, P. (1981) “Intra-Firm Diffusion, Bayesian Learning and Profitability”, Economic Journal, 91, 375-388.

Stoneman, P. (1983) The Economic Analysis of Technological Change, Oxford University press, Oxford.

Stoneman, P. (1987) the Economic Analysis of Technology Policy, Oxford University Press, Oxford.

Stoneman, P. and J. Vickers (1988) “The Assessment: The Economics of Technology Policy”, Oxford Review of Economic Policy, 4, i–xvi.

Stoneman, P. and M.-J. Kwon (1994) “The Diffusion of Multiple Process Technologies”, Economic Journal, 104, 420-431.

Stoneman, P. and P. David (1986) “Adoption Subsidies vs Information Provision as Instruments of Technology Policy”, Economic Journal, Supplement, 96, 142-150.

Stoneman, P. and P. Diederen (1994) “Technology Diffusion and Technology Policy”, Economic Journal, 104, 918-930.

Teece, D. and G. Pisano (1994) “The Dyunamic Capabilities of Firms: An Introduction”, Industrial and Corporate Change, 3537,556.

Thirtle, C. and VC. Ruttan (1987) the Role of Demand and Supply in the Generation and Diffusion of Technical Change, Harwood Academic Publishers, London.

Tirole, J. (1988) The Theory of Industrial Organization, MIT Press, Cambridge Mass.

Tonks, I. (1986) “The Demand for Innovation and the Diffusion of New Product”, International Journal of Industrial Organization, 4, 397-408.

Trajtenberg, M., R. Henderson and A. Jaffee (1997) “University Versus Corporate Patents: A Window on the Basicness of Invention”, Economics of Innovation and New Technology, 5, 19-50.

Utterbeck, J. (1994) Mastering the Dynamics of Innovation, HBS Press, Cambridge Mass.,

Vettas, N. (1998a) “Supply and Demand in new Markets: Diffusion with Bilateral Learning”, Rand Journal of Economics, 29, 215-233.

Vettas, N. (1998b) “Investment Dynbamics in Markets with Endogenous Demand”, CEPR Working Paper, London.

Vickery and Northcott (1995)

Wozniak, G. (1987) “Human Capital, Information and the Early Adoption of New Technology”, Journal of Human Resources, 22, 101-112.

NOTES

---

[1] For recent surveys of this literature, see Vickery and Northcott, 1995, Stoneman, 1983 and 1987, Karshensas and Stoneman, 1995, Thirtle and Ruttan, 1987, Metcalfe, 1981, Baptista, 1999, and others.

[2] Mansfield, 1989, for example, observed the following times for half the population of potential users to adopt new technologies: by-product coke ovens, 15 years; centralized traffic control, 14 years; car retarders and continuous annealing, 13 years; industrial robots, 12 years; and diesel locomotives, 9 years. At the other end of the spectrum are: tin containers, 1 year; continuous mining machines, 3 years; and numerically controlled machine tools, shuttle cars and pallet loading machines, 5 years.

[3] I have in mind that the unit of adoption is the firm, so that y(t) measures the number of firms using the new technology. However, firms rarely fully adopt a new technology (meaning that inter-firm diffusion often happens simultaneously with intra-firm diffusion). In the literature, some models focus on the output produced by firms using the new technology, or the market share that output represents. I am also going to talk as if it were a single artefact that was diffusing. In fact, subsequent developments of the original artefact often generate a sequence of artefacts which diffuse over time.

[4] There is no doubt that the psychological and sociological characteristics, educational achievements, outlook and attitudes to risk of early users are likely to differ from later users in a number of ways. This is a subject which has loomed much larger in the marketing and sociology literature than in the economics literature; see Rogers, 1995, for a survey of some of it.

[5] The exposition in the text follows Lekvall and Wahlbin, 1973, who call these two models "*external*" and "*internal*" influence models respectively. The mixed model is frequently used in the marketing literature; see Mahajan et al, 1990, for a survey. For some attempts to empirically discriminate between (2) and (3), see Karshensas and Stoneman, 1992, and Zettelmeyer and Stoneman, 1993, who observe that although the epidemic model implicitly presumes that all users participate in the process of informing others, epidemic models seem to work better empirically when one allows for the possibility that less than the full stock of users actively contribute to learning.

[6] Equation (2) is the workhorse of much of the diffusion literature in economics (equation (3) seems to play much the same role in the marketing literature). The pioneering studies of diffusion speed effectively imposed (2) on the data, modelling β as a function of various observables of interest; see Griliches, 1957, Mansfield, 1961 and 1963, Romeo, 1975 and 1977, Link and Kapur, 1994, and others (for informal applications of epidemic models in case study work, see Nabseth and Ray, 1974). However, recent work by Oster, l982, Hannan and McDowell, 1984, Levin et al, 1987, Rose and Joskow, 1990, Karshenas and Stoneman, 1993, and others has focused on modelling hazard rates or the time taken for particular firms to adopt a particular technology.

[7] The maximum citation frequency is 5 years in the study by Jaffee and Trajtenberg, 1996; see also Jaffee et al, 1993, Trajtenberg et al, 1997, and others. See Griliches, 1992, or Geroski, 1995, for two (amongst many) recent surveys of the empirical literature on spillovers.

[8] Rogers defines "*homophily*" as: "*…the degree to which two or more individuals who interact are similar in two or more attributes, such as beliefs, education, social status and the like"*, observing that: "*…homophilous communication …(*is) *… more likely …*(and) *…also more likely to be effective*" (1995, pp. 19; see also his Chapter 8).

[9] See Bartholomew, 1973, for a review of several more complex epidemic models which involve imperfect mixing; Rogers, 1995, Midgeley et al, 1992, and Debresson and Amesse, 1991, contain stimulating discussions of the effect of network structure on diffusion (and references to the relevant literature). For an introduction into the work ont eh geographical diffusion of innovations, see the survey by Baptista, 1999.

[10] The Davies model is actually expressed in terms of pay-back period: firm i adopts if it can pay-back the money invested in the new technology is less than some threshold time. I have also simplified the model somewhat by eliminating a discussion of the separate determinants of π_i and π^*.

[11] David, 1975, provides a nice illustration of this kind of argument: the adoption of mechanical reapers only makes sense if the savings in wage costs exceeds the cost of the machine, and this, of course, depends on how many workers are on the payroll (and on the number of acres over which the fixed costs of the machine will be spread). As the price of reapers drops, smaller and smaller farms will find adoption to be economic.

[12] David, 1969, Romeo, 1975, Davies, 1979, Hannan and McDowell, 1984, Ingham and Thompson, 1993, Levin et al, 1987, Metcalfe, 1970, Rose and Joskow, 1990, Karshenas and Stoneman, 1993, Pennings and Harianto, 1992, and many of the studies reported in Nasbeth and Ray, 1974, all report positive correlations between firm size and the speed of adoption. However, Mansfield, 1963, found insignificant effects, and there have been several negative correlations reported. Possibly the most famous of these was reported by Oster, 1982, in her study of the diffusion of the basic oxygen furnace in the US steel industry.

[13] For some theoretical work that explores the effects that the market structure (and other features) of supplying industries may have on diffusion rates in using industries, see Stoneman and Ireland, 1982, David and Olsen, 1984, 1986 and 1992, Metcalfe, 1981, Bass, 1980, and others.

[14] Diffusion of a new technology is also slowed whenever it stimulates technical progress in the older, established technology: for some examples, see Mokyr, 1990, pp. 90 and pp. 129 , Macleod, 1992, Harley, 1973, and others. On the other hand, the early introduction of relatively inexpensive complementary goods ought to speed diffusion: for some work on the effects of CD software on the diffusion of CD players, see Gandal et al, 1999.

[15] See Stoneman, 1981, Jensen, 1982 and 1983, Balcer and Lippman, 1984, Tonks, 1986, and others for a variety of learning models. Rosenberg, 1976, is a stimulating discussion of the role that expectations can play more generally in affecting diffusion.

[16] See Cohen and Levinthal, 1989, for a discussion and some evidence. Wozniac, 1987, provides some evidence on the effect of education and human capital formation has on diffusion. Case study and statistical evidence are united in suggesting that firms with higher absorptive capacity (however it is measured) adopt new technologies sooner than others. For example, higher educational attainment seems to have been associated with more rapid adoption of computer technology by California farmers (McWilliams and Zilberman, 1996), while the adoption of video banking by US commercial firms was fastest among those who has previous experience with IT and with the number of inter-firm arrangements it had (Pennings and Harianto, 1992).

[17] One particular example of this which has attracted recent attention are firms who resist new technologies for fear of causing too much disruption for their customers, or fail to adopt them because they are not well suited to their current customers current needs. Such firms are said to be "*too close to their customers*", and are generally thought likely to adopt "*sustaining*" new technologies but not "*disruptive*" ones; see Christensen, 1997, for some vivid case studies of this phenomena.

[18] There is a literature which argues that firms who possess "*dynamic competencies*" will be more able to adopt new technologies, and it seems plausible to believe that new technologies will diffuse more rapidly in industries where these skills abound; see Teece and Pisano, 1994, and the suggestive case studies by Insiti and Clark, 1994, Henderson, 1994, and others.

[19] Much of the recent econometric work suggests that governments rarely speed things up, and government owned enterprises rarely move faster than private owned ones; see Hannan and McDowell, 1984, Oster and Quigley, 1977, Rose and Joskow, 1990, and others. On the other hand, countries that were quick to grant licenses to mobile phone operators seem to have stimulated much higher rates of diffusion than others (particularly when the also have introduced competition and facilitated the switch from analogue to digital technology); see Gruber and Verboven, 1999.

[20] There is no reason to think that equipment ages at the same rate over time, and many scholars believe that scrapping rates rise in recessions (this is sometimes called "the pit stop theory of recession; see Cabellaro and Hammour, 1994, and, for some evidence, Oulton, 1989, Geroski and Gregg, 1998, and others). If this is true, then one expects to see variations in diffusion rates which can be associated with macroeconomic fluctuations.

[21] The equivalence between population models which count the number of users and probit models which describe the choice by a firm of whether or not to use a new technology is exact. However, firms can use a new technology more or less intensively (in practice, they rarely switch over wholesale to a new technology), and this can drive a wedge between population models which track usage and probit models which describe the adoption decision. Note that intra-firm diffusion rates can

sometimes be very slow. For example, it took both Ford and GM more than 20 years to reach 50% of their 1989 usage of robots (Nissan took 13 years to accomplish the same thing); see Mansfield, 1989, and references cited therein.

[22] For work using the stock adjustment model, see Chow, 1967, and Stoneman, 1976. These models typically just posit an adjustment rule, such as: $dS/dt = \alpha S_t\{S^* - S_t\}$, where S* is the equilibrium stock of usage, S, towards which the system gradually adjusts. In certain circumstances, this yields a logistic curve tracking usage over time (e.g. if adjustment were thought to be proportional to the difference between logs* and log S_t, then the diffusion curve would be Gompertz). Evolutionary models share with probit models the presumption that users are heterogeneous, and then examine the effect that selection mechanisms have on technology adoption choices: see Silverberg et al, 1988, Metcalfe, 1995, and others.

[23] See, for example, Katz and Shapiro, 1985, Farrell and Saloner, 1985 and 1986, and others; David and Greenstein, 1990, surveys much of the early literature on standards.

[24] See, Reinganum, 1981a and 1981b, Quirmbach, 1986, and others for models of sequential adoption which look at the effect of market power on diffusion; Reinganun, 1989, and Beath et al, 1995, survey the game theoretic literature, while Chatterjee et al, 1998, survey marketing models of competitive diffusion.

[25] See Tirole, 1988, Chapter 10, for an exposition of these forces; Stoneman and Kwon, 1994, document *"significant cross technology effects*" between different types of machine tools which seem similar to those which underlie the pre-emption effect; the oxygen steel case discussed by Oster, 1982, is a good illustration of rent displacement (there are many others). The distinction between "*competence enhancing*" and "*competence destroying*" technologies points to very similar effects on behaviour; see also the very similar distinction between "*sustaining"* and "*disruptive"* technologies made by Christensen, 1997.

[26] For what it is worth, Hannan and McDowell, 1984, found a positive association between adoption speeds and market concentration, while Levin et al, 1987 found a negative effect; Romeo, 1975, found that diffusion speeds increased with the number of firms in the market and fell with increases in the variance of the firm sizes (another measure of concentration), while Davies, 1979, found that diffusion fell with increases in firm numbers (and also with increases in the variance of firm sizes). Rather more satisfying is the approach of Karshenas and Stoneman, 1993. They distinguish "*rank effects*" (which reflect differences in the characteristics of adopters), "*order effects"* (which reflect returns to adoption associated with pre-emption of late movers by early movers) and "*stock effects*" (which reflect the decline in benefits which arise over time as more and more firms adopt the new technology). The latter two capture different types of competitive effects which might affect diffusion, but neither seemed to play much role in explaining the diffusion of computer numerically controlled machine tools in the UK.

[27] There is a large literature associated with this sylized fact, much of it associated with the "dominant design hypothesis"; for a recent exposition, see Utterbeck, 1994.

[28] Information cascades are defined as situations in which "*…it is optimal for an individual, having observed the actions of those ahead of him, to follow the behaviour of the preceding individual without regard to his own information"*, and are often used to explain "*herd*" behaviour*;* see Bikhchaandani, Hirschleifer and Welch, 1992, pp. 994 and 1998, De Vany and Walls, 1996, Bannerjee, 1992, and others for models of information cascades. The model of technology choice outlined in the text is based on Arthur, 1989; see also Arthur et al, 1987, who examine the underlying polyna urn process of the model. The argument that early investment choices may give pioneering brands long lasting advantages is discussed in Schmalensee, 1982, and others.

[29] On the other hand, competitive pressures may encourage firms to adopt new technologies "too soon" in order to pre-empt their rivals. For work on the effects of network externalities on diffusion, see Farrell and Saloner, 1986, Cabral, 1990, Choi, 1997, and others.

[30] For similar work in this spirit, see Vettas, 1998a and 1998b, Kapur, 1995, Jovanovic and MacDonald, 1994, Jovanovic and Lach, 1989, and others who identify conditions in which an S-curve will to emerge from a "*social learning*" process.

[31] What follows is not designed to be a general survey of the literature on technology policy (much less that associated with diffusion); for broad overviews of the area, see Metcalfe, 1995, Mowery, 1995, Stoneman and Vickers, 1988, Stoneman, 1987, David, 1988, Stoneman and Diederen, 1994, and others.

[32] See Stoneman and David, 1988, for a good discussion of the trade-off between information provision policies and subsidies.

[33] One classic example of a diffusion process which was killed by too many initial variants (and the behaviour of their sponsors) is quadrophonic sound (see Postrell, 1990); for some arguments about the effects of different types of competitors, see Porter, 1990.

[34] For stimulating discussions, see Ergas, 1987, Freeman, 1986, and others.

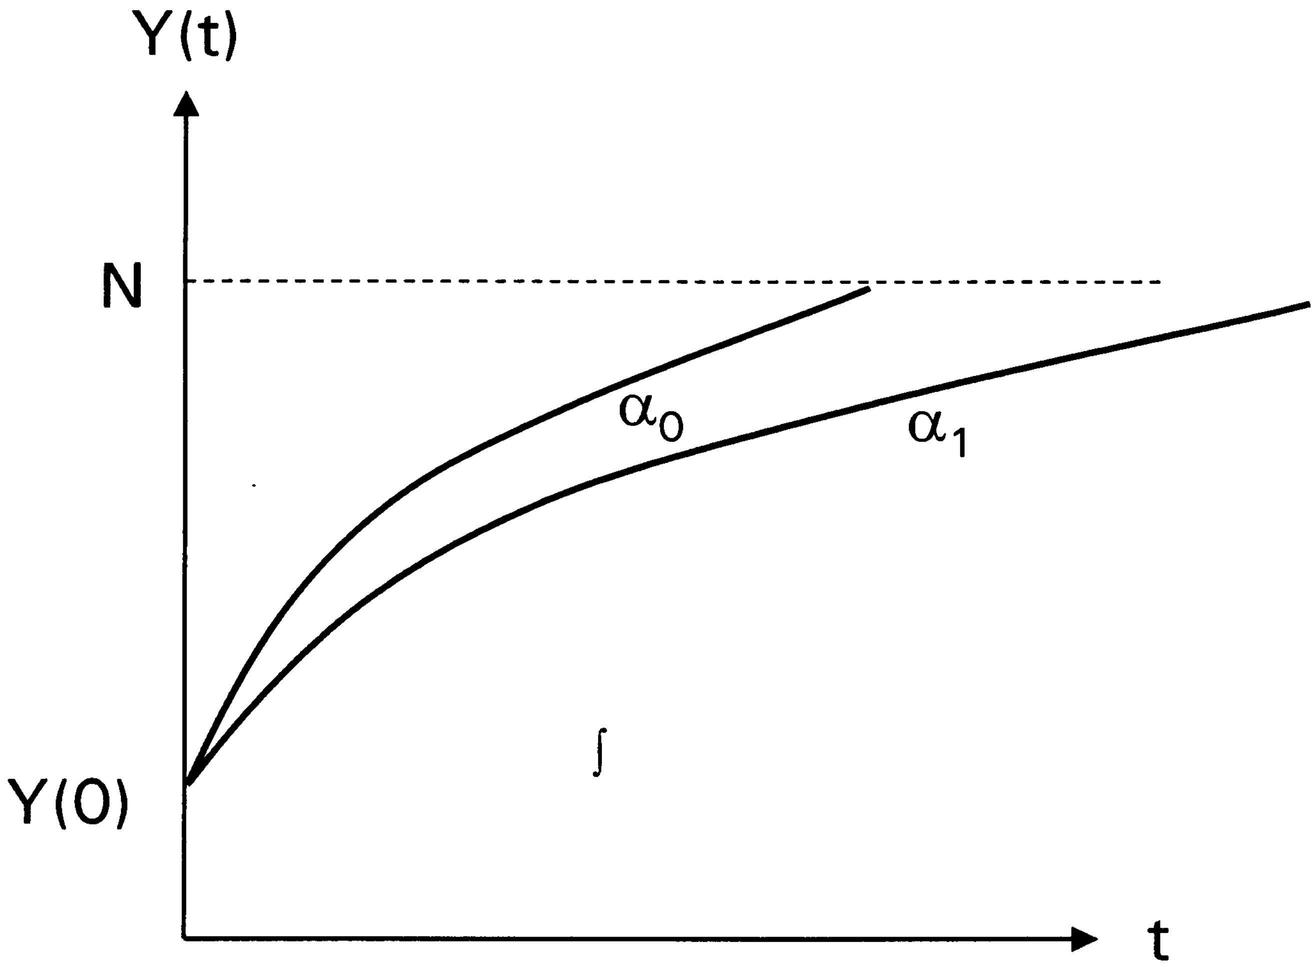

FIGURE I : A MODIFIED EXPONGITIAL FUNCTION diffusion time path generated by the "common source" model, equatial (1), for two values of $\alpha : \alpha = \alpha_0$ **and** $\alpha = \alpha_1 \angle \alpha_0$

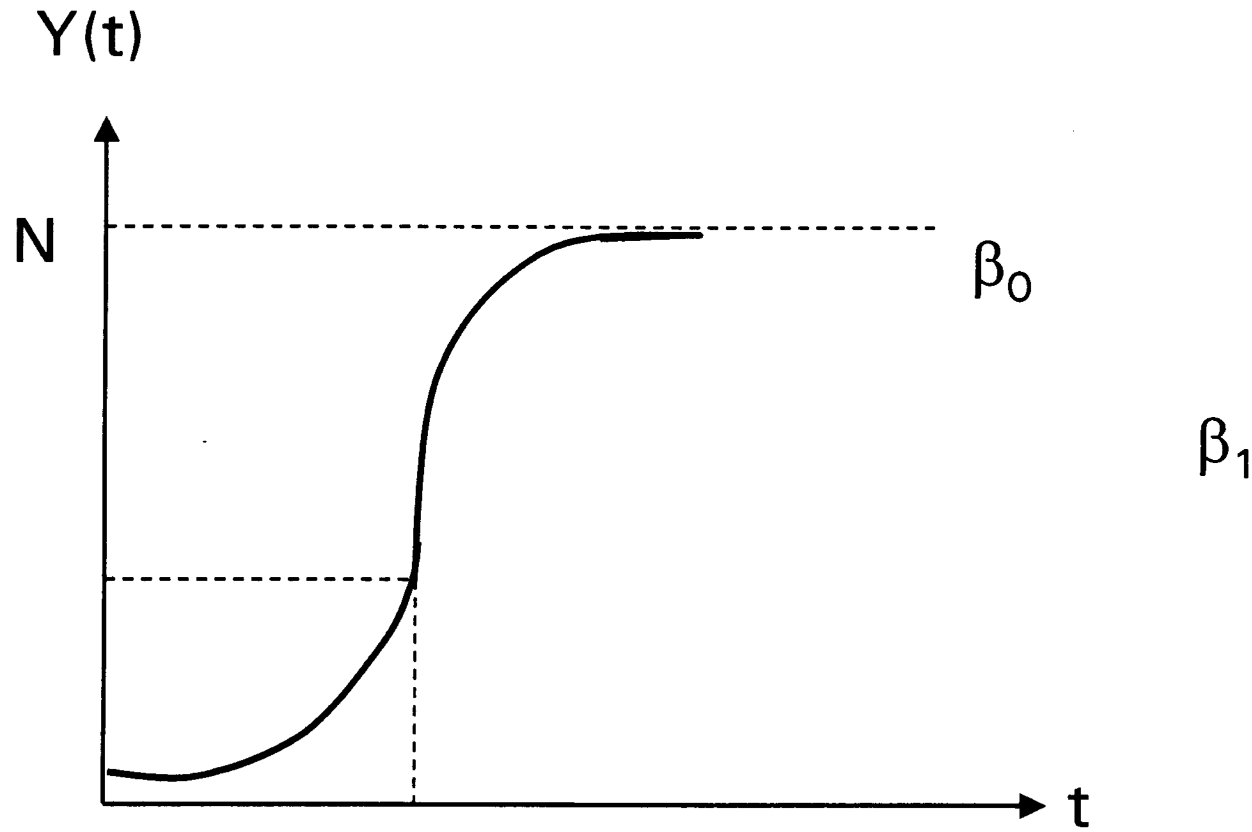

FIGURE II: A LOGISTIC FUNCTION
diffusion time path generated by the "word of mouth" model, equation (2) for two values of β: $\beta = \beta_0$ and $\beta = \beta_1 \angle \beta_0$. Note that it is symmetric around the inflection point $N/_2$ (assuming that $y/_0$) is not too different from zero), which is reached at twice $t = \uparrow$

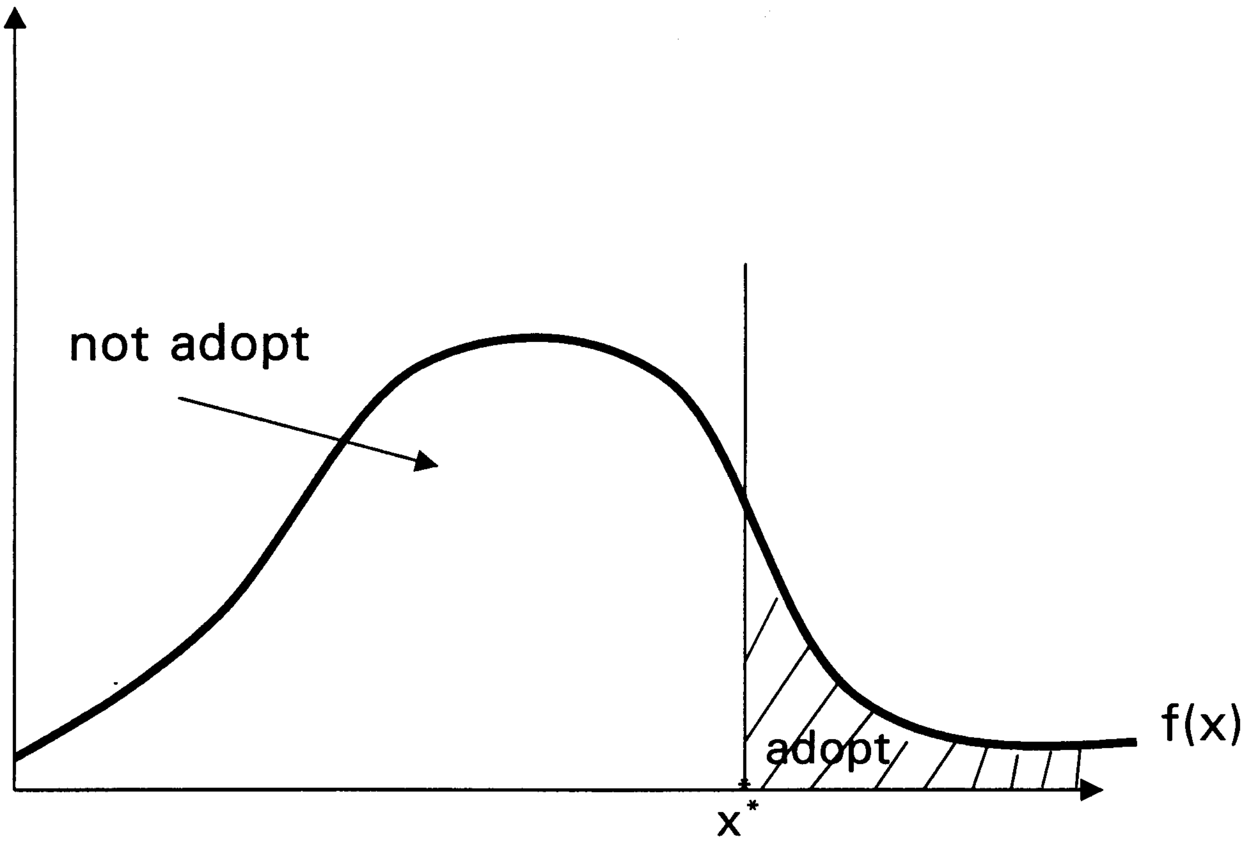

(a) Normal distribution of x_i

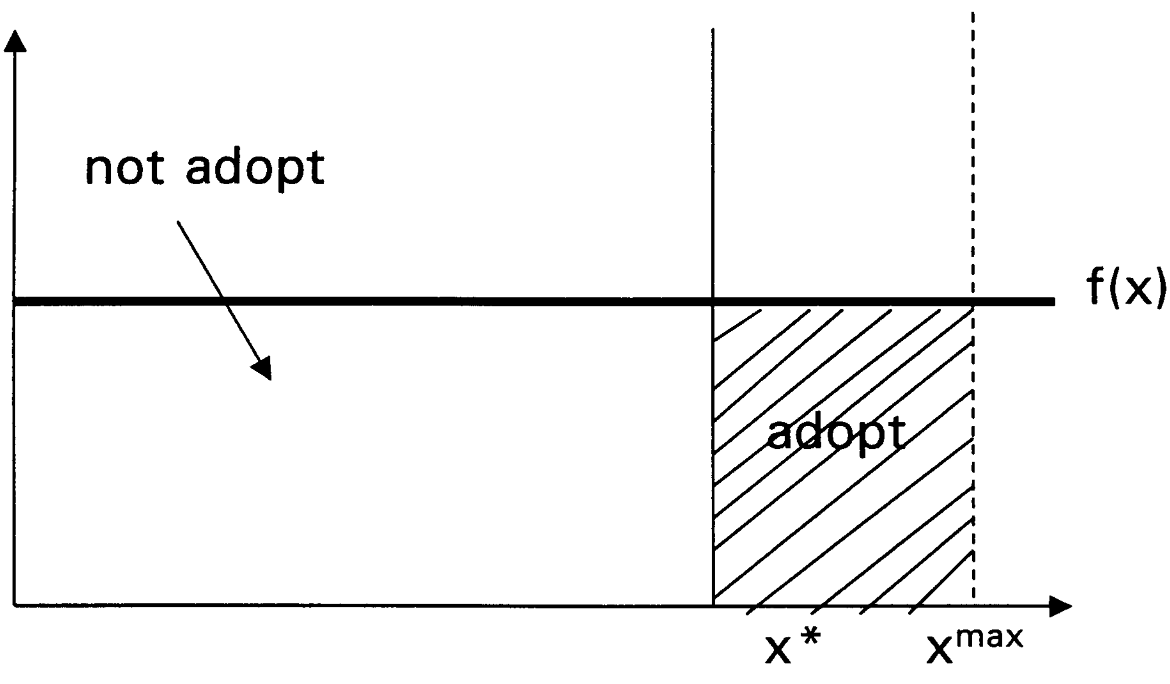

(b) rectangular distribution of x_i

FIGURE II: Two distributions of f(x) with thresholds separating adopters from non - adopters

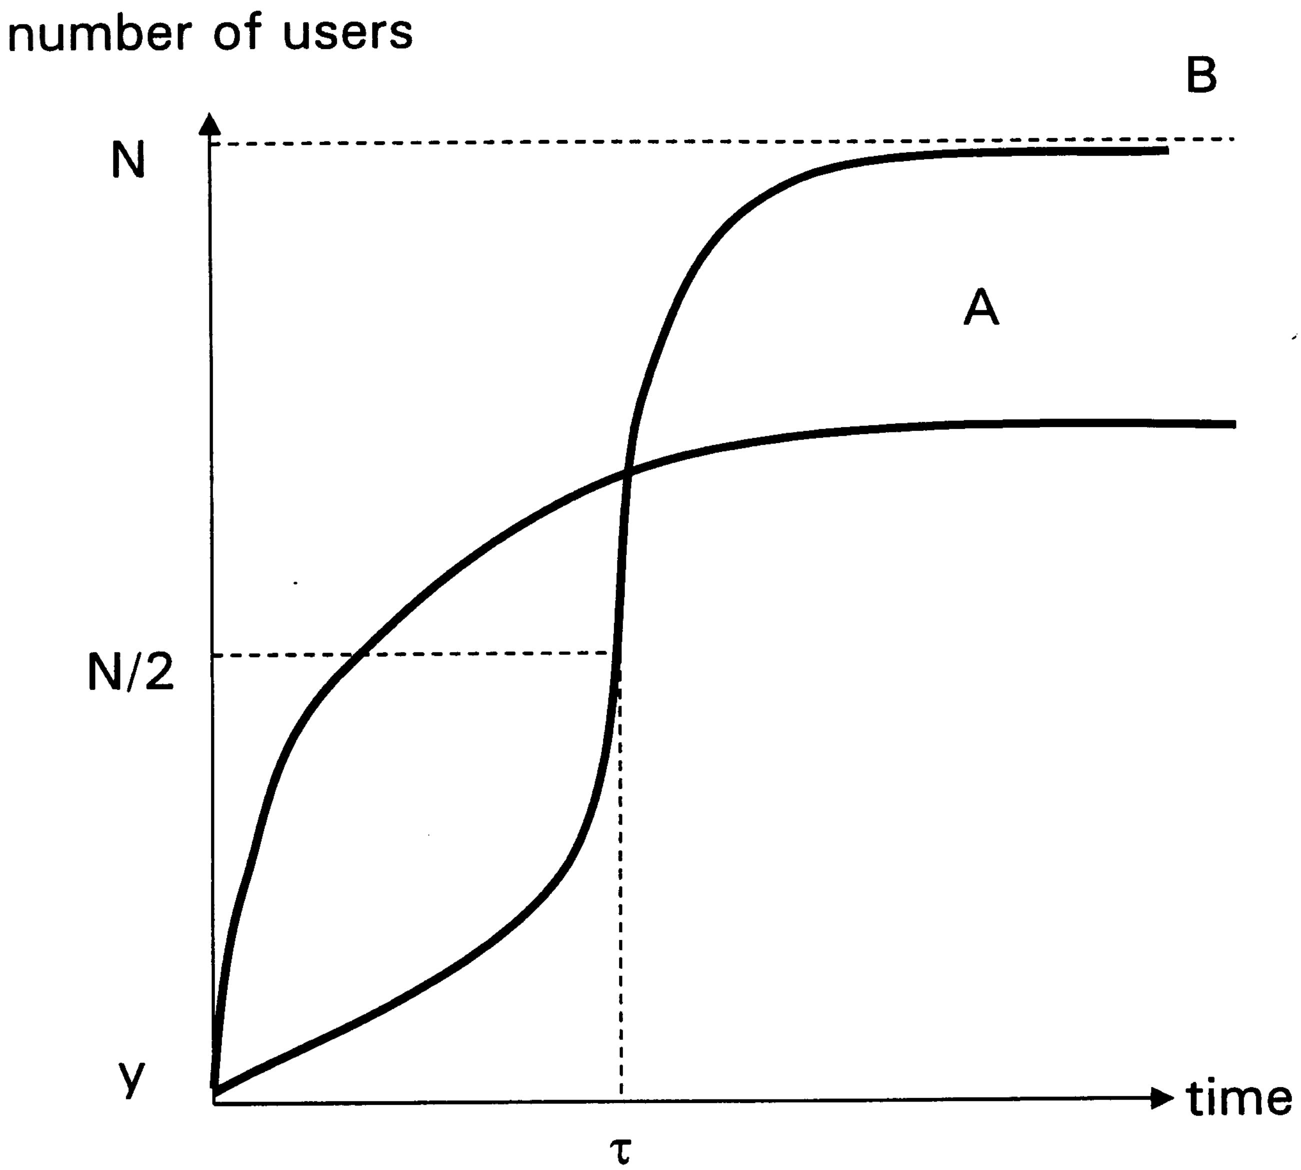

FIGURE I: Plots of the modifed exponential (A) and logistic (B) diffusion functions.