

ECONOMIC RESEARCH REPORTS

TECHNOLOGICAL VARIETY AND COMPETITION:
ISSUES OF DIFFUSION AND INTERVENTION

by

Robin Cowan

R.R. #89-23

September 1989

**C. V. STARR CENTER
FOR APPLIED ECONOMICS**



**NEW YORK UNIVERSITY
FACULTY OF ARTS AND SCIENCE
DEPARTMENT OF ECONOMICS
WASHINGTON SQUARE
NEW YORK, N.Y. 10003**

**TECHNOLOGICAL VARIETY AND COMPETITION:
ISSUES OF DIFFUSION AND INTERVENTION**

Robin Cowan
New York University

March 1989
revised September 1989

Recent work in the area of competing technologies has stressed the result that when several technologies compete for the same niche in the market, under a wide variety of circumstances, all but one of the technologies will be driven from the market place. Further, it need not be the best technology which remains. This paper argues that there is a place for intervention when technologies are being chosen or standards are being set. Intervention in the market can optimally balance the tradeoff between the desire for immediate payoffs and the desire for more information about different technologies—it can raise the expected present value of the adoption process. It remains the case, however, that even with optimal intervention, the market will lock in to a single technology, and not necessarily the superior one. The paper discusses these results with reference to technological changes in high tech and information technology

Prepared for the International Seminar on Science, Technology and Economic Growth. OECD, Paris, June 1989. The author thanks Terry Sicular for many helpful comments, and Yvonne Manu for very able research assistance. The author also acknowledges the technical assistance of the C.V. Starr Center for Applied Economics.

TECHNOLOGICAL VARIETY AND COMPETITION:
ISSUES OF DIFFUSION AND INTERVENTION

Introduction

It has been observed that the history of science contains many examples of multiple, independent discoveries of the same idea. Indeed, Merton (1973) claims that this is the rule rather than the exception regarding advances of scientific knowledge. The same observation must be made about the history of technology. Many technological advances appear in multiples (Merton's term referring to the independent discovery of the same idea by more than one scientist). It is also the case, however, that in the history of technology the multiples tend to consist of different technologies. Several innovators may simultaneously discover ways of overcoming a particular bottleneck, or practical difficulty. Though these discoveries, in their fully developed forms, will all be means to the same end, they may differ considerably in the way this end is achieved. Though they will perform the same function, the discoveries are technically distinct. Gas and electric refrigerators, heavy water, light water and gas-graphite nuclear reactors, wet and dry smoke-stack scrubbers are all instances of technological multiples. (See Arthur (1983), Cowan (1988b), David (1985), David & Bunn, Karlson, Besen for other examples.) Economists now refer to these situations as cases of competing technologies.

One of the interesting features of competing technologies is that often, as history unfolds, some of the technologies disappear as agents stop adopting them. It is also often the case that when several technologies are competing all but one of them disappear and this single technology comes to dominate or monopolize the market.

Standardization *per se*, that is to say having a single technology dominate a market, has both costs and benefits. (For a more complete discussion of the costs and benefits, see Braunstein & White.) Many of the benefits can be seen as arising from increasing returns.

If the market has standardized on a single technology, the scale of production and the level of cumulated past production will be greater than the case in which several technologies are competing. Scale economies can be exploited, and there will be greater benefits from learning by doing, learning by using and network externalities. Further, consumers gain flexibility, as when they are able to mix and match stereo components. Users of the technology are also able to avoid translating costs, and are able to fully utilize all aspects of a system, network or repertory. This consideration is relevant particularly when the technology has a hardware/software form.

The costs of standardization can be seen as being of two types. First, there are market structure effects. In any process of standardization which occurs over time, there will be users who have adopted the non-standard technology. There will be orphans. The orphans must either suffer the problems of having a non-standard technology, or bear the costs of standardizing. Second, standardization entails a loss of variety. Technologies having unique properties may be lost. This has two effects: The short run effect is that consumers who have strong preferences for those properties will find their preferences frustrated. The long run effect is that the scope for future developments will be curtailed. Technological advances that depend on the prior development of these unique, and now lost, properties are put in jeopardy.

Agricultural seed breeders are very concerned with this second issue. At the beginning of a search for plants with particular properties they explore many types of germplasm. As the process continues, however, some varieties are put aside as not promising, and the breeders breed down particular lines. The setting aside of this germplasm can result in a decreased ability to focus on characteristics which are its strengths. A more important problem, though, follows from the success of the plant breeder. The variety he has developed to have certain characteristics will be released and adopted by farmers. It will often replace indigenous varieties. The germplasm put aside by the breeder in his breeding process is preserved for possible later use. The indigenous

varieties, however, can disappear forever. This can remove from future possibilities characteristics peculiar to them. If the indigenous varieties have adapted to features of that agrosystem, there may be real costs involved.

The possibility that the market fails to choose the right standard makes the costs of standardization potentially greater. Locking in to an inferior technology involves immediate costs, in that users are forced to use an inferior technology. There are also future costs. To lose the prospect of the future development of a bad technological path will not impose large costs. On the other hand, to lose the ability to pursue a good technological path may be very costly.

“Right” and “wrong” are not precisely defined here. Below I define the superior (right) technology as that technology which, as the standard, would yield the highest benefits. This is clearly a simplification. One of the lessons which is clear from the historical work on competing technologies (see especially David & Bunn, Cowan (1988b)) is that which technology is “right” and which technology is “wrong” changes with time and circumstances. This makes the task of even the most foresighted decision maker very difficult indeed.

One explanation for the disappearance of a technology is simply that it was not a very good one. If a technology is clearly more costly than its rival, if the benefits from using it are lower, it would be surprising if it did not disappear. A different explanation, however, is that the technology that remains does so because it somehow gained an advantage in the competition for market share, and the others fell further behind until they simply disappeared. In contrast to the previous explanation, this explanation makes no claims about which of the technologies is inherently superior.

If the second explanation is to be of use it must be that there is a way in which one of the technologies can “gain an advantage” in the competition. One mechanism often discussed is increasing returns. Any technology which exhibits learning by doing, learning by using, or network externalities will benefit from having been adopted in the past. The

degree to which each of the competing technologies has penetrated the market will be an important consideration for an agent faced with a choice among them. If a technology is such that as more people adopt it, it becomes more valuable to an adopter, it is clear that it can build up momentum at the expense of other technologies.

There has recently been increased interest in the economics of markets in which several increasing returns technologies are competing. (See Arthur (1988) for a survey of this literature.) Most of the theoretical work, however, shares two features: it has been concerned solely with market forces; and the future payoff streams of the technologies are well known from the start of the process. This literature has produced two striking results. Under a wide variety of circumstances these markets will standardize, or lock in to a single technology. And, under a similarly wide variety of circumstances, the technology to which the market locks in need not be inherently superior. “Superior” is defined in an *ex post* sense: Given that the market locks in to one technology, the superior technology is that technology which, if dominant, provides the highest net benefits. Note that there is an *ex ante* sense of superior: In this sense the superior technology at time t is that technology which an economically rational agent would adopt at time t . One of the interesting features of this work is that even though all agents adopt the *ex ante* superior technology, the market can lock in to an *ex post* inferior technology. For the remainder of this paper “superior” and “inferior” will refer to the *ex post* sense.

This paper is concerned with questions of sequential technology choice, and the process of standardization. If there are costs to standardization, the issue of intervention to reduce them arises. If market forces are enough to ensure that one technology dominates, then intervention is unlikely to be able to eliminate the costs entirely, particularly costs due to lost variety. It may be possible, though, to manage the process of standardization in such a way as to reduce the costs and increase the overall net benefits, and to reduce the probability of standardizing on the wrong technology.

This paper presents two simple models of choice among competing technologies. In both models agents act sequentially to adopt technologies. The technologies are embodied in capital goods. In the first model the development paths of the technologies are well known, in the second there is uncertainty regarding these paths. Both models are developed in the context of a market with agents who make only a single adoption. The question is then asked of each of the models: What would be the effect of a policy maker who acts to steer the adoption process from one technology to another according to an optimal policy? The paper concludes with a discussion of different factors that affect the outcome of the adoption process and the efficacy of intervention.

Technology Choice with no Uncertainty

Following Arthur 1988, consider a market having two technologies, A and B, and homogeneous agents. The payoffs to both technologies increase with the number of adopters and are well known. Thus it is known at the outset that $P_A^1 \leq P_A^2 \leq P_A^3 \dots$ and $P_B^1 \leq P_B^2 \leq P_B^3 \dots$ where P_j^i is the payoff to the i^{th} user of technology j . Each period one agent adopts one or another of the technologies. He adopts that technology which will maximize his benefits. Suppose now that $P_A^1 > P_B^1$ (See Figure 1). The first adopter will choose technology A—it maximizes his benefits. The second adopter is either the second adopter of technology A, or the first of technology B, whichever maximizes his benefits. To maximize his benefits he will be the former. Clearly, the same will be true of all who follow. The market locks in to the technology that has the highest initial payoff, in this case A. This is so regardless of the fact that the technology having the highest initial payoff may, as the only technology adopted, yield a lower present value of the adoption stream than other technologies.¹

¹ This model is essentially the simplified version of Arthur's 1988 model. (See Arthur (1988) p.119.) Arthur's full model is a generalization of this one in that he has two types of adopters, each type having a preference for one of the two technologies. In his model the market locks in, not necessarily to the first technology chosen, but possibly to the inferior technology. The idea underlying this model is also present in the models of Farrell and Saloner .




FIGURE 1
Payoffs to two technologies, A and B, which change with the degree of use.

Intervention can clearly be efficacious here. At the beginning of the adoption process the policy maker can simply do present value calculations for two possible futures—one in which only A is adopted, and one in which only B is adopted. By forcing the adoption of the technology with the higher present value, he will be doing at least as well as the market. It is clear that to force the adoption of a single technology is the optimal policy. This also follows from a theorem known as the Doctor/Lawyer Theorem, which states that if annual incomes from the two professions are non-decreasing with experience, if a person at the outset of his career chooses law, say, it will never pay him to try his hand at doctoring.²

The simplicity of these results stems from the fact that the development paths of the two technologies are known with certainty. The policy maker can judge at the outset which is the optimal path and ensure that it is followed.

This certainty regarding payoff streams, however, is not a typical feature of new technologies. In general, how well a technology will perform, and precisely what it will do are not known when it is first introduced. (See Rosenberg, 1982, especially Part II.) This type of uncertainty is central to the next model of sequential technology choice. The dynamics of the adoption process differ from the case of certainty, both with and without

² For a proof of this intuitively obvious, yet technically difficult theorem, see Cowan (1988) Appendix One.

intervention in the market. When this uncertainty is present one would expect to see both technologies being tried early in the process, as information about how they perform is being gathered. We might expect however, that eventually enough information will be gathered that the market (or policy maker) makes a choice, in effect, and fixes on a single technology. This is in fact the case.

Technology Choice with Uncertainty

Assume two technologies³ both of which are subject to dynamic increasing returns through learning by using. Wet scrubber and dry scrubber technologies, for example, are available to reduce sulphur dioxide emissions from coal-fired electricity generating stations. The technologies are unsponsored,⁴ and are competitively supplied. The technologies are embodied in capital goods—particular scrubbers in particular smoke-stacks. The capital goods have known lifetimes and operate unaltered throughout those lifetimes. Each period a consumer adopts one of the technologies. He installs a wet scrubber, say. The information generated from his adoption and use generates improvements to that technology, so the next wet scrubber built is slightly better than the last. We can think of each consumer as adopting a new variant of the technology. Each successive variant is better than its predecessors, but once adopted, that variant (that particular wet scrubber) is unchanged throughout its lifetime. In other words, the learning, and improvements to the technologies, cannot be applied retroactively.

Dynamic increasing returns imply that the payoff to the next adopter of a technology is an increasing function of the number of agents who have previously adopted that technology. This function is unknown, but its family is known. Specifically, assume that the learning function for a technology is one of the family of functions described by $Qf(n)$, where Q is a parameter and $f(n)$ is a known, deterministic function, increasing in n . So

³ These results can be easily generalized to any finite number of technologies.

⁴ See Arthur (1983) for a discussion of sponsored and unsponsored technologies.

stated, a single observation is enough to determine the precise curve (assuming $f(1) \neq 0$) and so after each technology has been adopted once, we have returned to the world of certainty. This argument assumes that observations can be made with precision and that the first observation is representative of all that follow. This is clearly false. There is measurement error, and a single technology will operate differently under different conditions and with different users. The exact degree of emission reduction is affected by precisely the type of coal used, which changes both geographically and over time, as well as by various details of the physical design of the plant and smoke-stack. To capture these features, assume that Q is a random variable. For technology A then, the true learning curve is given by $Q_A f(n_A)$ where n_A is the number of agents who have adopted technology A. A similar relation holds for technology B: $Q_B f(n_B)$. We can see the payoffs as being distributed around a learning curve, $\pi f(n)$, where π is defined as the mean of Q .⁵ (See Figure 2)




FIGURE 2

In this figure, $f(n)$ is the deterministic function which determines the shape of the learning curve. For the technology pictured here, the true curve is centered around its mean, $\pi f(n)$, and π is less than 1. The distributions sketched around $\pi f(n)$ indicate that the payoffs do not fall exactly on this curve, but are distributed around it.

Since Q is a random variable, the payoff to each adopter is also a random variable. Once an agent adopts a technology, though, the payoff is observed and the value of that

⁵ This particular form, $Qf(n)$, is not crucial, the random component could as well be additive: $Q + f(n)$. The same results would follow.

random variable is known. A scrubber is installed, and the degree of emission reduction is observed. Each adoption can be seen as an experiment which has a random payoff. It is known, however, that there is a tendency for payoffs to increase, due to the existence of learning. What is not known is how strong this tendency is.

Agents are expected-profit maximisers. Thus their concern is with the means of Q_A and Q_B , π_A and π_B . Estimates of these values, μ_A and μ_B , are updated every period. Knowledge is public, and all agents begin with the same priors. If, for example, Q is distributed normally, with a known variance, σ , and unknown mean, π , then the posterior distribution on the value of π would be normal with mean \bar{x} , and variance $\frac{\sigma^2}{n}$, where \bar{x} is the mean of the n previous observations of Q . In effect then, there is a single, public estimate for each of the two means.

Each period, on the basis of expected profit, an agent adopts a technology. He adopts A if $\mu_A f(n_A) \geq \mu_B f(n_B)$; he adopts B otherwise. His net payoff is observed and this (public) information is used to update the estimate of Q for the technology he adopted. Next period, a new agent uses the new priors to choose between the technologies.

There are three important aspects to every adoption. The first is the immediate payoff to that adopter: sulphur dioxide emissions from that smoke-stack are reduced. The second is the learning which occurs and enables improvements to be made to the technology: the next wet (dry) scrubber installed will, in general, be better. Third is the observation on the payoff, which allows the estimate of the true value of the technology to be updated. We can better estimate whether wet or dry is the better technology.

In this model one of the technologies is used only a finite number of times. This implies that the market share of one of the technologies goes to one, the other to zero. The argument behind the proof of this proposition is very straightforward for the case in which the true means of Q_A and Q_B differ. Suppose both technologies were used an infinite number of times. Eventually it would be known what the true means are. Then the functions $Q_A f(n_A)$ and $Q_B f(n_B)$ would be known with certainty. Given that f is non-

decreasing, it will be clear, and remain so, which technology is superior. Only that technology will be used. This follows again from the Doctor/Lawyer Theorem. Thus it cannot be that both technologies are used an infinite number of times. This is the thrust of the formal proof—as the technologies are used, agents get closer and closer to perfect knowledge, and eventually close is close enough. A very similar argument exists for the case where $Q_A = Q_B$.

It can also be shown that in this model the inferior technology can dominate. Here “inferior” is well defined: A technology is inferior if it has a lower mean of Q . (A is inferior if $\pi_A < \pi_B$.) If an inferior technology is ever used, it can get locked in by the following process. Initial good luck can cause it to be used early on. As it is being used, it is building up improvements through learning. It can be used enough so that the improvement compensates for the true low value of the mean of Q . That is, when the truth is finally learned, it will have advanced so far along its learning curve that it remains, in the sense of immediate, one period payoffs, better than the other, inherently superior technology.

It is also the case that an inferior technology can dominate the market even if there are no increasing returns. The presence of increasing returns facilitates the proof, in that more use *per se* gives a technology an advantage over other technologies. It facilitates the actual process as well, however, in the following sense. After lock-in has occurred, the estimate of the “value” of the locked in technology, A say, converges to the true value: approaches π_A . In the absence of increasing returns, it must be the case that the most recent estimate of the value of B is less than the true value of A: $\mu_B < \pi_A$. If B is the better technology, it must have experienced bad luck such that the estimate of its worth has been driven down. The same need not be true under increasing returns. Because the level of payoffs increase with use, μ_B need not have been driven below π_A . “Knowing” that $\pi_A < \mu_B$ is consistent with further use of A providing greater payoffs than switching to B.

We see that, not surprisingly, the addition of uncertainty about the “true value” of the competing technologies and their development paths does not change the results: One technology will dominate the market, and with positive probability it will be an inferior technology.

Again, this suggests that there may be a role for intervention. One might think that a policy maker could simply subsidize both technologies until he knows which one is superior and then force its adoption. This is only partly so. We will argue below that even with complete control, when technologies improve with use but the policy maker cannot precisely predict the paths of development, it is possible to lock in to an inferior technology. Discounting, and a concern for social profitability can cause the policy maker to pursue a technology which he believes to be best, at the cost of further learning about, and development of, other technologies. This allows erroneous estimates of the relative goodness of technologies to be perpetuated.

To the model just discussed we now add a policy maker has complete control of the adoption process.

Again, there are two technologies A and B, both subject to learning by using. Their learning curves are unknown but are known to belong to the family of $Qf(n)$ where Q is a random variable, and $f(n)$ is a deterministic, increasing function of the number of times the technology has been adopted. For each technology, i , the mean of Q_i , π_i , is unknown, but estimated by μ_i . Technologies are unsponsored and supplied competitively.

Each period, on the basis of the prior beliefs about π_A and π_B , an agent adopts a technology. The capital good in place, his net payoff is observed, and this information is used to update beliefs about that technology. As in the previous model there are three aspects to every adoption: There is an immediate payoff to the adopter; there is learning which generates improvements; and an observation is made which is used to improve the estimate of the true worth of that technology. In this model, in contrast to the previous one, the policy maker balances these considerations optimally. By imposing taxes and subsidies

every period, he is able to force the adoption of the optimal technology, where “optimal” is defined relative to current information. This is the *ex ante* notion referred to above. His objective is to maximize the expected value of the discounted benefits of the entire adoption process.

What the policy maker faces is a difficult dynamic programming problem. He must balance the desire for present payoffs with that, not only with the desire for future payoffs, but also with the desire for knowledge with which to make better future decisions. Cowan (1988a) presents a full specification of this problem. The dynamic programming problem can be finessed, however, by applying the two-armed bandit model from probability theory.

A two-armed bandit is a slot machine that has two arms. The objective is to play the arms, one at a time in any order, so as to maximize the expected value of one’s winnings. The probabilities of getting jackpots from either arm is unknown. The payoff probabilities are independent, and each play generates one of two results: failure, in which case the player gets nothing; or success, in which case the payoff is of some known and fixed size. This model of technology choice differs from the two-armed bandit in two respects. The possible payoffs from adopting a technology are not simply success or failure. There is a payoff to each adoption, which can be seen as having been drawn from some, possibly continuous, distribution. The second difference is that with each adoption there is a deterministic change—a movement along the learning curve. Regardless of the size of the payoff to the current adopter, there is a tendency for the payoffs to increase, whereas for the two-armed bandit the size of jackpots is unchanging.

The work of Gittins and Jones has produced an optimal policy for a large class of models that share the objective of maximizing the expected discounted value of the stream of earnings. This is the Gittins Index Policy. The Gittins index is a scalar assigned to each of the arms of a multi-armed bandit, which measures the value of a try on that arm. The value of the index for a particular arm is independent of what is known about all the other arms. The policy consists in playing, each period, the arm with the highest index.

The model of technology choice described above is a member of the class of models for which the Gittins index policy is optimal; optimal, that is, with respect to the goal of maximizing the expected discounted value of payoffs to future adopters. That the Gittins index policy is optimal implies that a policy maker will “do better” than will an unguided market. Intuitively, he will do better because of the potential trade-off between immediate payoff and information. A policy maker, in deciding which technology to subsidize, takes into account what will be learned about the merits of a technology. If there is a large amount to be learned (the variance of his prior is large) he may be willing to forego some immediate benefits to gain that knowledge. It may be optimal to adopt a technology that has a lower expected payoff. A single adopter, however, would be unwilling to make this decision. His goal is to maximize his immediate benefits. He gets no benefits from the fact that future adopters are able to make better informed decisions. In effect, a market is equivalent to the one-step-look-ahead policy for the multi-armed bandit. Gittins and Jones have shown this to be inferior to the Gittins index policy.

The optimality of the Gittins index policy shows that intervention can indeed generate a better outcome than will the market. This is only with respect to the assumed objective, however, and is not to say that intervention will either prevent or encourage standardization, or that it can prevent the dominance of an inferior technology. Both of these issues have to do with the long run dynamics of the system, with which Gittins and Jones were not concerned.

Cowan (1988a) is concerned with the long run dynamics and proves two results analogous to those that have emerged in the work on market systems. The first is that the policy maker cannot prevent one technology from dominating the market. Lock-in is inevitable. The intuition is straightforward. (For a formal proof see Cowan (1988a).) The argument is by contradiction. Suppose that market dominance does not take place; both technologies are used an infinite number of times. In this case, after some finite length of time the policy maker has beliefs about the distributions of Q_A and Q_B which are very close

to the true distributions. It will also be true that these beliefs change very little after this point. Call this time t^* . It will also be true at t^* that the expected value from using technology B, say, exclusively thereafter is greater than that from using A exclusively. Is there a time after t^* at which it is better to use A than B? The answer is no: If it is better to use A at $t^* + n$, it will be better to do so at $t^* + n + 1$ as well. This is so because beliefs about the two distributions (of Q_A and Q_B) effectively do not change, but increasing returns are operating. Clearly if A should be used at $t^* + n$, it should be used exclusively thereafter. Between t^* and $t^* + n$, however, technology B has been subject to increasing returns, so it looks even better, in terms of future payoffs, than it did at t^* . Thus if, at $t^* + n$, using A exclusively is preferable to using B exclusively, this will also be true at t^* . But this, by hypothesis, is not true, and we have a contradiction.⁶ Eventually, but in finite time (possibly before t^*), one of the technologies will disappear from the market.

The second result in Cowan (1988a) is that the remaining technology need not be the superior one. Assuming that neither technology is believed to be so bad that it is written off from the very start (technically, the intersection of the support of the priors is not empty), then with positive probability the inferior technology will come to dominate the market.⁷

Again the intuition is very simple: By assumption it is possible to use the worse technology at least once. Suppose this happens, and that a good result is obtained. The policy maker will be inclined to use it again. Good readings, so to speak, can be received for many successive adoptions. But each time the technology is used it moves along its learning curve, so the expected value of its *true* net benefits increases with each adoption. If this goes on long enough, this technology will have advanced sufficiently along its learning

⁶ This result is similar in spirit to the Doctor/Lawyer Theorem. Rothschild proves that the proposition holds for a constant returns case. The proposition is true for all returns regimes, with the exception of unbounded decreasing returns.

⁷ Another technical assumption is that $f(n)$ not increase too rapidly. If $f(n)$ increases very rapidly, only the first technology used will ever be used.

curve that discounting ensures that the other technology, no matter how good, is left behind.⁸

The analysis developed states that a policy maker who maximizes social profitability cannot prevent the market from standardizing on one technology, nor can he prevent it from standardizing on an inferior technology. The same is true of a policy maker who is constrained not to ignore social profitability. Preventing standardization will involve intervening in the market forever (since if he stopped after finite time market forces will take over, and they are enough to ensure lock-in). Similarly, it will also force agents to make ex ante suboptimal decisions forever. To be able to prevent standardizing on the wrong technology would involve knowing which was the right technology. This cannot be known (with probability one) in finite time, it would take an infinite amount of time and so an infinite expenditure.

The optimality of the Gittins index policy, however, does imply that intervention will lower the probability of locking in to the inferior technology. To see that this is so, consider the following question: “Could it be that policy P^* gives a higher expected payoff than another policy, P , yet is more likely to lead to standardization on the inferior technology?” For the answer to be “Yes”, the following two statements would have to be true of P^* relative to P :

- 1) Early in the process, P^* dictates that the superior technology be used more frequently than does P .
- 2) P^* generates a higher probability of switching away from the superior technology than does P .

1) must be true to ensure that the expected value of the total payoff is higher under P^* than it is under P . By hypothesis, P^* is more likely to lead to the inferior technology being adopted later in the process. Thus the payoffs to P^* will, in expected value, be less

⁸ Again, increasing returns are not necessary for this proposition to be true. Their presence does facilitate both the process and the proof however.

than those to P for the later part of the process. Therefore for P* to yield a higher total expected present value, it must give higher expected payoffs early in the process. Given 1), 2) must be true since under P* the probability of locking in to the inferior technology is higher.

There are, however, two effects which follow from 1). The superior technology, early on, will make greater advances along its learning curve, than will the inferior technology under P* than under P. This suggests that, contrary to statement 2), the probability of switching away from the superior technology is greater under P*. To switch away would involve suffering through the early, low payoff portion of the learning curve of the other technology. The second effect is that information about the superior arm is gathered faster. In general, then, under P* at any time in the early part of the process the policy maker will have more information on which to base his estimate of the distribution for the superior technology than he will under policy P. Since the true mean of the superior arm is higher than that of the inferior arm, this implies that under P* the policy maker is less likely to switch away from the superior arm. (*Very loosely, the more technology A is used, the closer is μ_A to π_A . Thus to switch away from A, it must be that μ_B , the estimate of π_B , is erroneously high. But the probability of having an erroneous estimate of π_B is not a function of the policy used, rather it is a function of the initial priors and of the realizations of the payoffs to the two technologies.*) These two effects suggest that statements 1) and 2) are not compatible, and so P*, the policy which generates a higher expected discounted payoff, yet a lower probability of standardizing on the good technology cannot exist. Thus, a policy maker, by employing the Gittins index policy not only maximizes the expected value of the payoff stream, but also generates, relative to a market, a higher probability of standardizing on the superior technology.

Discussion

The previous section has shown that even with the presence of a powerful policy maker the system of sequential technology choice retains two features of the market. The first is that it is certain that one of the technologies eventually disappears from the market. It is not difficult to see the mechanism whereby this happens: At the beginning of the process, the policy maker, by forcing the use of both technologies, is trying to form an accurate estimate of their relative merits. At some point one of the technologies has a string of good luck. Several consecutive installations of wet scrubbers are very effective in reducing sulphur dioxide. This has two effects. The policy maker will revise upward his beliefs about the size of payoffs that technology will provide; and that technology moves along its learning curve. Both effects, which occur while the other technology sits idle, give this technology an advantage in the competition for market share. If the technology continues to have reasonably good luck, the estimate of its inherent worth will not fall too fast, and it will continue to be used. The other technology will get left behind. It will become too costly to switch to the other technology, since in switching, one would be forced to suffer through the early, low benefit, segment of the learning curve before payoffs rise to match those of the much-used technology.

Lock in will occur even in the absence of increasing returns (see Rothschild), in this case, though, the mechanism is slightly different. A technology gets locked in by the fact that with experience the estimate of the distribution of Q (which determines the level of payoffs) converges to the true distribution of Q . As a technology is used more and more, beliefs about its value stabilize. Eventually they change so little that we effectively know how good a technology is and that it is better than what we believe the other technology to be. Lock in will occur under increasing returns, constant returns, or bounded decreasing returns. If returns are decreasing and unbounded, it will always be the case that after a finite number of adoptions of technology A, say, for the next adopter, the payoff from A will be lower than that from technology B, and so the system will switch technologies. If

decreasing returns are bounded, as the bounds are approached, the system approximates a constant returns system, and so lock in will occur.

In both regimes, increasing and constant returns, early in the adoption process the optimal policy can dictate that we use that technology which has lower expected immediate (one-period) payoffs. This will be so if the prior distribution for that technology has, relative to that of its competitor, close to the same mean but larger higher moments. While we would lose in immediate payoffs, we would gain in decreasing uncertainty. As the technologies are used, the higher order moments of the priors fall, and this possibility (of sacrificing immediate payoffs for reduced uncertainty) disappears. In a constant returns regime this means that the technology with the higher μ will be used. Late in the process then, the policy maker will always believe he is using the better technology, where “better” means “has higher π ”. This is not so under increasing returns. Late in the process the technology with higher $\mu f(n)$ will be used. Because $f(n)$ is an increasing function this is consistent with using the technology with lower μ . So the policy maker can believe he has locked in to the wrong technology. In a constant returns regime there is no room for this Strotz-type regret. We always believe (sometimes wrongly) that changing past experiences with the technologies would have no effect on the optimality of our decision to use A today (though it may well have an effect on our beliefs about it). We can lock into a wrong technology but we would never think so. Strotz-type regret can exist if there are increasing returns—we can believe we are trapped on the wrong technology.

High technology and lock in

High technology industries are now often seen as the new engine of growth. As smoke-stack industries were seen as the source of an economy’s dynamism earlier this century, high technology is now pointed to as that which will spur the next economic advance. Referring to high technology as the new engine of growth may be somewhat hyperbolic, (though it is hyperbola used both by politicians and the popular press) but it is

certainly the case that high technology represents far-reaching potential for technological change and advance. This change is manifested in both product and process innovation, and seems to be occurring throughout the economy. The pervasiveness with which high tech has been accepted and used, and what appears to be the considerable potential for further technical change does suggest that high technology possesses at least some of the features of an engine of growth.

This shift in the nature of the engine of growth is significant in two ways. First it is significant simply as a shift. Second, it is significant that the new engine is high technology.

When there is a large scale shift from one type of technology to another, it is accompanied by a large increase in technological uncertainty. It is well-understood in the research and development literature that there is relatively little uncertainty involved with making modifications to existing products or processes. Developing an entirely new product or process, by contrast, is subject to considerable uncertainty, both about how to do it, and whether it can in fact be done. (See DeLeon and Jovanovic and Rob.) A switch from an old type of technology to a new one will be accompanied by a similar increase in technological uncertainty, though perhaps on a different scale. We have, now, very good knowledge about smoke-stack type technologies—what can be done with them, and how to maximize payoffs from them. The same is not true about high technology. For particular technologies perhaps we do, but as a technology-type, we have little idea how it will progress and evolve: In 1981, the Commodore PET was introduced. It had 8 kilobytes of random access memory, and a half-sized keyboard. One of the very exciting features was that it could be programmed in BASIC! Looking at the Commodore PET, it would clearly not have been possible to make a sensible prediction about the uses to which micro-computers would be put in 8 years.

A second effect of a wholesale technological shift is that there will be a considerable amount of learning about the new technology-type. There will be learning about what it will

do. This will be true both of the technology-type and particular instances of it. High technology is continually being put to use in new ways, often in the form of new, discreet technologies. It is also true that micro-computers, an example of a high-tech technology, are being put to new uses. This type of learning is, in part, a reduction of the uncertainty discussed above. There is also, however, learning about how to make the technology provide higher payoffs—how to make it perform particular tasks more efficiently. This will occur through both learning by doing and learning by using.⁹ Rosenberg (1982) provides an extensive discussion and documentation of this feature of new technologies. This effect clearly diminishes as technology-types and technologies mature, but to the extent that high technology is still a new technology-type, learning will remain important.

The learning that accompanies new and strange technologies has two effects. There is the direct learning by using effect (see Rosenberg, 1982) which is itself a source of dynamic increasing returns. There is also an indirect effect, though, in that uncertainty about the performance features of a technology will generate information networks. Users of a technology will find great benefits in being able to share information about problems and solutions they have encountered. One sees this much less with familiar technologies. Having bought an automobile, I know very well what I can do with it, where to take it to be fixed and so on. The scope for learning from other owners exists, but is limited. Furthermore, much of the knowledge is documented, and readily available in print. This was not true, of course, when the automobile was first invented, nor is it true of microcomputers now. For microcomputers, information networks exist to share information both about how to solve problems, and about just what can be done with the machine. The desire to be part of a large network, to gain precisely this type of information,

⁹ In this discussion there have been three levels of generality: technology types, for example high technology; technologies themselves—the Macintosh computer; and particular variants or instances of a technology—the computer in my office. The statements I have made about uncertainty and learning are likely to be true at all levels. That is, if they are true at the most general level, they will be true as well at subordinate levels.

is often given as a reason for the adoption of light water nuclear reactor technology rather than one of its competitors. (See Cowan (1988b).)

The second significant feature of this technological shift is that the new technology-type is high technology. Certainly, the introduction of very innovative technologies and technology types in the past—the steam engine; the automobile; nuclear power, for example—has been accompanied by learning by using and the creation of information networks. These effects are likely to be common to all such inventions. It may be the case, though, that they are more pronounced for high technology. Relative to manufacturing technologies, high technology differs in the degree of increasing returns, both static and dynamic. “[E]conomies based on high-technology products operate subject to very large design and set-up costs; to significant learning effects in manufacturing; and to possibilities of standardization due to “network externalities.” These act to give market advantage to products that have achieved high volume or high visibility. Therefore, modern, high-technology economies contain several sources of positive feedback and market self-reinforcement.” (Arthur, *Emergent Structures*, p. 2)

This discussion suggests that there has been an increase in the degree of dynamic increasing returns present in the economy. This is due partly to the fact that there has been a technological shift to a new engine of growth, and also due to the fact that the shift has been to high technology. Technology competitions of the type modelled above, then, will be taking place under stronger increasing returns and greater uncertainty than they have in the past. In terms of the model, stronger increasing returns means that $f(n)$ increases more quickly. Increased uncertainty means that the original priors on π have a larger variance.

Increasing returns are not necessary to force the market to lock in, but their presence and degree, or strength, will have an effect on the speed with which the lock-in occurs. The greater the increasing returns, the faster lock-in will occur. Simplifying the second model to make Q a Bernoulli random variable allows us to cast the process as a random walk. If Q is Bernoulli, a sensible prior will be a Beta distribution: Beta (α, n), to

estimate the probability of success π . Here α can be seen as the number of past successes with the technology, and n the number of past trials. (For ease of exposition assume the original prior was Beta (0,0).) As a technology is used it can be seen as a random walk through (α, n) space. Each point in this space corresponds to a value of the Gittins index, $G(\alpha, n)$. Clearly a technology, A, will continue to be used provided it does not fall below the boundary formed by the locus of points $\{ (\alpha, n) \mid G(\alpha, n) = G_B \}$, where G_B is the current index value for technology B. It can be shown that this locus, or iso-index line, shifts down as $f(n)$ shifts up. Thus the stronger are the increasing returns, the higher is the propensity for a technology to continue being used, and so the sooner the system will lock in. (For a more complete presentation of this argument see Cowan (1987), Chapter 4.)

Three potential pitfalls of intervention in the technology choice process have been pointed out by David (1987). The Narrow Windows Paradox refers to the problem that the time at which intervention can be effective may be very restricted. The Blind Giants Quandary encapsulates the notion that the windows tend to be open early in the process, at which point the least is known about any of the developing technologies. At the time at which a policy maker is most able to steer the process, he least knows what to do. Angry Orphans are early adopters who, perhaps due to efforts by the policy maker, adopted a technology which did not become the standard.

One effect of the faster lock-in time is that Narrow Windows become more narrow still. They also become more difficult to keep open. The windows close, in part, because increasing returns make it too costly to shift from one technology to the other. To hold the windows open, a policy maker must provide large enough subsidies to overcome this cost and to force a shift to a technology that has not advanced along its learning curve. Clearly the stronger the increasing returns the faster these costs increase, and the larger the subsidies must be.

On the other hand, though, a shorter lock-in time will make for fewer Angry Orphans. The time to lock in is extended if there is extensive experimentation with both technologies. The more experimentation there is, the more will the non-standard technology have been used, and so the more orphans. In contrast, it could happen that B, say, is used many times, then A is used only a few times, and then B is used again, this time exclusively thereafter. This would involve a long lock-in time, yet few orphans. This could only occur, though, under very particular sequences of results from using the technologies. The strong correlation between lock-in time and number of orphans seen in simulation data indicates that this is not a common sequence of events. Clearly a reduction in the time to lock-in will reduce the costs of orphans and orphanages. (See David (1987) for a discussion of these costs.)

It was suggested above that a shift in technology and technology type will involve an increase in technological uncertainty. New technologies within a new technology type are likely to be surrounded with considerable uncertainty regarding how they work and what they will do. In the early days of nuclear power, it was the consensus that the fast breeder reactor would quickly become the technology of choice. Indeed the first nuclear power was generated using this technology. The history of nuclear power, however, has shown the early projections about the relative merits of various nuclear technologies to be far from borne out. Relative to new technologies within an *old* technology type, there are very few examples on which to base extrapolations about how the technology will develop. A change in the degree of uncertainty—the blindness of Paul David's Giants, or the variance and higher moments of the initial priors in the model above—will also have an effect on the time to lock-in and on the number of orphans.

This uncertainty arises, and increases, due to the novelty of the technologies (Rosenberg, 1982). It can also arise from a different source. (This point is due to DeLeon.) What is important for the purposes of technology choice are the relative merits of the two technologies. If the two technologies being compared are very different, this can

lead to problems on that score. This is not necessarily an issue of recognizing the merits of the particular technologies, rather it is one of devising a scale on which to compare them. The first best would be some scalar measure. In principle this is not difficult, we simply need two functions which map from the properties of a technology into a scalar. The problem is that there is uncertainty about exactly what these functions should be in order to make the right comparison. Thus, not only will there be learning about the precise properties of the technology, there will be learning about the functions used to compare them. This translates into initial priors with large variances.

In general, one would expect that as the uncertainty increases, the time to lock-in will also increase. If the planner is very confident that A is superior to B, he may try B a few times “just to make sure”, but will quickly turn to A, and stick with it in the absence of disastrous results. On the other hand, if it is not at all clear which is superior, we would expect to see considerable experimentation before one technology emerges as dominant. (The intuitive appeal of this story may be driven in part by risk aversion, but I think it is correct even absent risk aversion.)

There is a counteracting force however. This arises from the fact that if there is considerable uncertainty, a small number of observations can have a large effect on perceptions of relative merits of the technologies. Thus the higher the variances of the priors, the faster the evaluations can diverge, and so the faster can a winner emerge. This effect is more important in the absence of a policy maker: Myopic agents are concerned solely with the mean of the posterior distribution, to which this effect refers specifically. While the mean of the prior can change quickly, if the prior has a large variance, the variance itself will also fall relatively quickly. This is of no concern to a myopic agent, but, all else equal, the Gittins Index Policy favours the prior with the higher variance. The rapidity with which the mean will change if the variance is large is offset, under the Gittins Index Policy, by the speed with which the variance falls, and so the effect of the fast diverging means is somewhat reduced.

Simulation results indicate that the former is the dominant effect: An increase in uncertainty will increase the expected time to lock-in. This result is more pronounced if there is a policy maker acting than if market forces alone operate.

Conclusion

This paper has modelled an adoption process wherein adopters are faced with a choice among several technologies. It has argued that in the case where there is uncertainty about the relative merits of the competing technologies there is scope for intervention in the market. Intervention can raise the expected value of the adoption process by balancing the three effects from an adoption: the immediate payoff from the adoption; the improvement to the adopted technology; and the added information about the relative merits of the different technologies. Even optimal balancing of these features, however, cannot eliminate the possibility of locking in to an inferior technology, though it will reduce the probability of doing so.

Major innovations, which can sometimes be referred to as the introduction of a new type of technology, bring with them considerable potential for further innovation and economic growth. The advance in knowledge creates opportunities for new economic activities, and as has been argued above, major innovations are often subject to strong increasing returns. This potential comes at the cost of greater uncertainty however. The uncertainty is about how to proceed—which of many paths to pursue, in order to maximize the newly created potential.

An implication of the analysis developed above is that the market undersupplies experimentation. The undersupply will be more significant when there is more uncertainty about which is the better path to follow, which suggests that there is most scope for effective intervention in technology competitions immediately following major innovations. This is the time when the reduction of uncertainty is most valuable. Reducing this uncertainty is a costly activity, however, and involves the use of resources on activities that

may not be directly productive. It may be rational, though, to bear these costs, in order to reap larger benefits from standardization further down the road.

BIBLIOGRAPHY

- Arthur, W. Brian "On Competing Technologies and Historical Small Events: The Dynamics of Choice Under Increasing Returns", International Institute for Applied Systems Analysis, Working Paper WP-83-90, 1983.
- Arthur, W. Brian "Competing Technologies: An Overview" in *Technical Change and Economic Theory* G. Dosi, C. Freeman, R. Nelson, G. Silverberg and L. Soete eds. London, Pinter 1988.
- Arthur W. Brian, "Competing Technologies, Increasing Returns, and Lock-In by Historical Events" *Economic Journal*, Vol. 99, Mar. 1989.
- Besen, S. "Compatibility Standards, Competition and Innovation in the Broadcasting Industry", Mimeo, 1986.
- Braunstein, Y. M. and L.J. White "Setting Technical Compatibility Standards: An Economic Analysis" *Antitrust Bulletin*, Summer 1985.
- Callon, M. "The State and Technical Innovation: A Case Study of the Electrical Vehicle in France", *Research Policy*, Vol. 9, 1980.
- Cowan, R. "Backing the Wrong Horse: Sequential Choice Among Technologies of Unknown Merit", New York University, 1988.
- Cowan, R. "Nuclear Power Reactors: A Study in Technological Lock In" C.V. Starr Center for Applied Economics, Research Report R.R. #88-33, 1988.
- Cowan, R. *Backing the Wrong Horse: Sequential Technology Choice Under Increasing Returns* Ph.D. Dissertation, Stanford University, 1987.
- Cowan, Ruth S. *More Work for Mother: The ironies of household technology from the open hearth to the microwave* New York, Basic Books, 1983.
- David, Paul A. "CLIO and the Economics of QWERTY" *American Economic Review*, Vol. 75, 1985.
- David, Paul A. with Julie Bunn, " 'The Battle of the Systems' and the Evolutionary Dynamics of Network Technology Rivalries", mimeo, 1987.
- David, Paul A. "Some New Standards for the Economics of Standardization in the Information Age" in Dasgupta, P. and P. Stoneman eds. *Economic Policy and Technological Performance*. Cambridge University Press, 1987.
- DeLeon, P. "Evaluation of Technology R&D: A Continuing Dilemma" *Research Policy* Vol. 11, 1982.
- Eads, George and Richard R. Nelson "Governmental Support of Advanced Civilian Technology: Power Reactors and the Supersonic Transport", *Public Policy*, Vol. 19, 1971.
- Farrell, J.R. and G. Saloner "Standardization, Compatibility and Innovation", *Rand Journal of Economics*, Vol.16, 1985.
- Farrell, J.R. and G. Saloner "Installed Base and Compatibility", *American Economic Review*, Vol. 76, 1986.
- Freeman, Christopher *The Economics of Industrial Innovation*, MIT Press, Cambridge, Mass. 1982.

- Gittins J.C. and D.M. Jones "A Dynamic allocation Index for the Sequential Design of Experiments", in *Progress in Statistics*, J.Gani, K.Sarkadi, and I.Vincze eds., North-Holland, 1974.
- Jovanovic, B. And R. Rob, "Long and Short Waves: Growth through intensive and extensive search" New York University, 1989.
- Karlson, S.H. "Adoption of Competing Inventions by United States Steel Producers", *Review of Economics and Statistics*, Vol. 1986.
- Katz, Michael L. and Carl Shapiro, "Network Externalities, Competition and Compatibility", *American Economic Review*, Vol. 75, 1985b.
- Katz, Michael L. and Carl Shapiro, "Technology Adoption in the Presence of Network Externalities", *Journal of Political Economy*, Vol.94, 1986.
- Keck, O. "Government Policy and Technical Choice in the West German Reactor Programme", *Research Policy*, Vol 9, 1980.
- Mansfield, Edwin *The Economics of Technical Change*, W.W. Norton and Company, New York, 1968.
- Merton, R.K. *The Sociology of Science: Theoretical and Empirical Investigations*. N.W. Storer, ed. University of Chicago Press, 1973.
- Nelson, Richard R. "The Allocation of Research and Development Resources: Some Problems of Public Policy", in R.A. Tybout, ed. *Economics of Research and Development*, Ohio State University Press, 1965.
- Nelson, Richard R. "Uncertainty, Learning, and the Economics of Parallel Research and Development Efforts", *Review of Economics and Statistics* Vol. 63, 1961.
- Rosenberg, Nathan *Inside the Black Box: Technology and Economics*, Cambridge University Press, Cambridge, 1982.
- Rosenberg, Nathan "The Direction of Technical Change: Inducement mechanisms and focusing devices" *Economic Development and Cultural Change* 1969.
- Rothschild, M. "A Two-Armed Bandit Theory of Pricing", *Journal of Economic Theory*, Vol. 9, 1974.
- Santa Fe Institute *Emergent Structures* March 1989.
- Teubal M. and E. Steinmueller, "Government Policy, Innovation and Economic Growth: Lessons From a Study of Satellite Communications.", *Research Policy*, Vol. 11, 1981.