

Competing Technologies, Increasing Returns, and Lock-In by Historical Events



W. Brian Arthur

The Economic Journal, Vol. 99, No. 394. (Mar., 1989), pp. 116-131.

Stable URL:

<http://links.jstor.org/sici?sici=0013-0133%28198903%2999%3A394%3C116%3ACTIRAL%3E2.0.CO%3B2-R>

The Economic Journal is currently published by Royal Economic Society.

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/about/terms.html>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/res.html>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is an independent not-for-profit organization dedicated to creating and preserving a digital archive of scholarly journals. For more information regarding JSTOR, please contact support@jstor.org.

COMPETING TECHNOLOGIES, INCREASING RETURNS, AND LOCK-IN BY HISTORICAL EVENTS*

W. Brian Arthur

This paper explores the dynamics of allocation under increasing returns in a context where increasing returns arise naturally: agents choosing between technologies competing for adoption.

Modern, complex technologies often display increasing returns to adoption in that the more they are adopted, the more experience is gained with them, and the more they are improved.¹ When two or more increasing-return technologies ‘compete’ then, for a ‘market’ of potential adopters, insignificant events may by chance give one of them an initial advantage in adoptions. This technology may then improve more than the others, so it may appeal to a wider proportion of potential adopters. It may therefore become further adopted and further improved. Thus a technology that by chance gains an early lead in adoption may eventually ‘corner the market’ of potential adopters, with the other technologies becoming locked out. Of course, under different ‘insignificant events’ – unexpected successes in the performance of prototypes, whims of early developers, political circumstances – a different technology might achieve sufficient adoption and improvement to come to dominate. Competitions between technologies may have multiple potential outcomes.

It is well known that allocation problems with increasing returns tend to exhibit multiple equilibria, and so it is not surprising that multiple outcomes should appear here. Static analysis can typically locate these multiple equilibria, but usually it cannot tell us *which* one will be ‘selected’. A dynamic approach might be able to say more. By allowing the possibility of ‘random events’ occurring during adoption, it might examine how these influence ‘selection’ of the outcome – how some sets of random ‘historical events’ might cumulate to drive the process towards one market-share outcome, others to drive it towards another. It might also reveal how the two familiar increasing-returns properties of *non-predictability* and *potential inefficiency* come about: how increasing returns act to magnify chance events as adoptions take place, so that *ex-ante* knowledge of adopters’ preferences and the technologies’ possibilities may not suffice to predict the ‘market outcome’; and how increasing returns

* I thank Robin Cowan, Paul David, Joseph Farrell, Ward Hanson, Charles Kindleberger, Richard Nelson, Nathan Rosenberg, Paul Samuelson, Martin Shubik, and Gavin Wright for useful suggestions and criticisms. An earlier version of part of this paper appeared in 1983 as Working Paper 83–90 at the International Institute for Applied Systems Analysis, Laxenburg, Austria. Support from the Centre for Economic Policy Research, Stanford, and from the Guggenheim Foundation is acknowledged.

¹ Rosenberg (1982) calls this ‘Learning by Using’ (see also Atkinson and Stiglitz, 1969). Jet aircraft designs like the Boeing 727, for example, undergo constant modification and they improve significantly in structural soundness, wing design, payload capacity and engine efficiency as they accumulate actual airline adoption and use.

might drive the adoption process into developing a technology that has inferior long-run potential. A dynamic approach might also point up two new properties: *inflexibility* in that once an outcome (a dominant technology) begins to emerge it becomes progressively more 'locked in'; and *non-ergodicity* in that historical 'small events' are not averaged away and 'forgotten' by the dynamics – they may decide the outcome.

This paper contrasts the dynamics of technologies' 'market shares' under conditions of increasing, diminishing and constant returns. It pays special attention to how returns affect predictability, efficiency, flexibility, and ergodicity; and to the circumstances under which the economy might become locked-in by 'historical events' to the monopoly of an inferior technology.

I. A SIMPLE MODEL

Nuclear power can be generated by light-water, or gas-cooled, or heavy-water, or sodium-cooled reactors. Solar energy can be generated by crystalline-silicon or amorphous-silicon technologies. I abstract from cases like this and assume in an initial, simple model that two new technologies, *A* and *B*, 'compete' for adoption by a large number of economic agents. The technologies are not *sponsored*² or strategically manipulated by any firm; they are open to all. Agents are simple consumers of the technologies who act directly or indirectly as developers of them.

Agent *i* comes into the market at time t_i ; at this time he chooses the latest version of either technology *A* or technology *B*; and he uses this version thereafter.³ Agents are of two types, *R* and *S*, with equal numbers in each, the two types independent of the times of choice but differing in their preferences, perhaps because of the use to which they will put their choice. The version of *A* or *B* each agent chooses is fixed or frozen in design at his time of choice, so that his payoff is affected only by past adoptions of his chosen technology. (Later I examine the expectations case where payoffs are also affected by future adoptions.)

Not all technologies enjoy increasing returns with adoption. Sometimes factor inputs are bid upward in price so that diminishing returns accompany adoption. Hydro-electric power, for example, becomes more costly as dam sites become scarcer and less suitable. And some technologies are unaffected by adoption – their returns are constant. I include these cases by assuming that the returns to choosing *A* or *B* realised by any agent (the net present value of the version of the technology available to him) depend upon the number of previous adopters, n_A and n_B , at the time of his choice (as in Table 1⁴) with

² Following terminology introduced in Arthur (1983), *sponsored* technologies are proprietary and capable of being priced and strategically manipulated; *unsponsored* technologies are generic and not open to manipulation or pricing.

³ Where technologies are improving, it may pay adopters under certain conditions to wait; so that no adoptions take place (Balcer and Lippman, 1984; Mamer and McCardle, 1987). We can avoid this problem by assuming adopters need to replace an obsolete technology that breaks down at times $\{t_i\}$.

⁴ More realistically, where the technologies have uncertain monetary returns we can assume von Neumann-Morgenstern agents, with Table 1 interpreted as the resulting determinate expected-utility payoffs.

Table 1
Returns to Choosing A or B given Previous Adoptions

| | Technology A | Technology B |
|---------|--------------|--------------|
| R-agent | $a_R + m_A$ | $b_R + m_B$ |
| S-agent | $a_S + sm_A$ | $b_S + sm_B$ |

increasing, diminishing, or constant returns to adoption given by r and s simultaneously positive, negative, or zero. I also assume $a_R > b_R$ and $a_S < b_S$ so that R -agents have a natural preference for A , and S -agents have a natural preference for B .

To complete this model, I want to define carefully what I mean by ‘chance’ or ‘historical events’. Were we to have infinitely detailed prior knowledge of events and circumstances that might affect technology choices – political interests, the prior experience of developers, timing of contracts, decisions at key meetings – the outcome or adoption market-share gained by each technology would presumably be determinable in advance. We can conclude that our limited discerning power, or more precisely the limited discerning power of an implicit *observer*, may cause indeterminacy of outcome. I therefore define ‘historical small events’ to be those events or conditions that are outside the *ex-ante* knowledge of the observer – beyond the resolving power of his ‘model’ or abstraction of the situation.

To return to *our* model, let us assume an observer who has full knowledge of all the conditions and returns functions, except the set of events that determines the times of entry and choice $\{t_i\}$ of the agents. The observer thus ‘sees’ the choice order as a binary sequence of R and S types with the property that an R or an S comes n th in the adoption line with equal likelihood, that is, with probability one half.

We now have a simple neoclassical allocation model where two types of agents choose between A and B , each agent choosing his preferred alternative when his time comes. The supply (or returns) functions are known, as is the demand (each agent demands one unit inelastically). Only one small element is left open, and that is the set of historical events that determine the sequence in which the agents make their choice. Of interest is the adoption-share outcome in the different cases of constant, diminishing, and increasing returns, and whether the fluctuations in the order of choices these small events introduce make a difference to adoption shares.

We will need some properties. I will say that the process is: *predictable* if the small degree of uncertainty built in ‘averages away’ so that the observer has enough information to pre-determine market shares accurately in the long-run; *flexible* if a subsidy or tax adjustment to one of the technologies’ returns can always influence future market choices; *ergodic* (not path-dependent) if different sequences of historical events lead to the same market outcome with probability one. In this allocation problem choices define a ‘path’ or sequence of A - and B -technology versions that become adopted or ‘developed’, with early adopters

possibly steering the process onto a development path that is right for them, but one that may be regretted by later adopters. Accordingly, and in line with other sequential-choice problems, I will adopt a 'no-regret' criterion and say that the process is *path-efficient* if at all times equal development (equal adoption) of the technology that is behind in adoption would not have paid off better.⁵ (These informal definitions are made precise in the Appendix.)

Allocation in the Three Regimes

Before examining the outcome of choices in our R and S agent model, it is instructive to look at how the dynamics would run in a trivial example with increasing-returns where agents are of one type only (Table 2). Here choice order does not matter; agents are all the same; and unknown events can make no difference so that ergodicity is not an issue. The first agent chooses the more favourable technology, A say. This enhances the returns to adopting A . The next agent *a-fortiori* chooses A too. This continues, with A chosen each time, and B incapable of 'getting started'. The end result is that A 'corners the market' and B is excluded. This outcome is trivially predictable, and path-efficient if returns rise at the same rate. Notice though that if returns increase at different rates, the adoption process may easily become path-inefficient, as Table 2

Table 2
An Example: Adoption Payoffs for Homogeneous Agents

| Number of previous adoptions | 0 | 10 | 20 | 30 | 40 | 50 | 60 | 70 | 80 | 90 | 100 |
|------------------------------|----|----|----|----|----|----|----|----|----|----|-----|
| Technology A | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 |
| Technology B | 4 | 7 | 10 | 13 | 16 | 19 | 22 | 25 | 28 | 31 | 34 |

shows. In this case after thirty choices in the adoption process, all of which are A , equivalent adoption of B would have delivered higher returns. But if the process has gone far enough, a given subsidy-adjustment g to B can no longer close the gap between the returns to A and the returns to B at the starting point. Flexibility is not present here; the market becomes increasingly 'locked-in' to an inferior choice.

Now let us return to the case of interest, where the unknown choice-sequence of two types of agents allows us to include some notion of historical 'small events'. Begin with the constant-returns case, and let $n_A(n)$ and $n_B(n)$ be the number of choices of A and B respectively, when n choices in total have been made. We can describe the process by x_n , the market share of A at stage n , when

⁵ An alternative efficiency criterion might be total or aggregate payoff (after n choices). But in this problem we have two agent types with different preferences operating under the 'greedy algorithm' of each agent taking the best choice at hand for himself; it is easy to show that under any returns regime maximisation of total payoffs is never guaranteed.

n choices in total have been made. We will write the difference in adoption, $n_A(n) - n_B(n)$ as d_n . The market share of A is then expressible as

$$x_n = 0.5 + d_n/2n. \quad (1)$$

Note that through the variables d_n and n – the difference and total – we can fully describe the dynamics of adoption of A versus B . In this constant-returns situation R -agents always choose A and S -agents always choose B , regardless of the number of adopters of either technology. Thus the way in which adoption of A and B cumulates is determined simply by the sequence in which R - and S -agents ‘line up’ to make their choice, $n_A(n)$ increasing by one unit if the next agent in line is an R , with $n_B(n)$ increasing by one unit if the next agent in line is an S , and with the difference in adoption, d_n , moving upward by one unit or downward one unit accordingly. To our observer, the choice-order is random, with agent types equally likely. Hence to him, the state d_n appears to perform a simple coin-toss gambler’s random walk with each ‘move’ having equal probability 0.5.




Fig. 1. Increasing returns adoption: a random walk with absorbing barriers

In the increasing-returns case, these simple dynamics are modified. New R -agents, who have a natural preference for A , will switch allegiance if by chance adoption pushes B far enough *ahead* of A in numbers and in payoff. That is, new R -agents will ‘switch’ if

$$d_n = n_A(n) - n_B(n) < \Delta_R = \frac{(b_R - a_R)}{r}. \quad (2)$$

Similarly new S -agents will switch preference to A if numbers adopting A become sufficiently ahead of the numbers adopting B , that is, if

$$d_n = n_A(n) - n_B(n) > \Delta_S = \frac{(b_S - a_S)}{s}. \quad (3)$$

Regions of choice now appear in the d_n, n plane (see Fig. 1), with boundaries between them given by (2) and (3). Once one of the outer regions is entered, both agent types choose the same technology, with the result that this technology further increases its lead. Thus in the d_n, n plane (2) and (3)

describe barriers that 'absorb' the process. Once either is reached by random movement of d_n , the process ceases to involve both technologies – it is 'locked-in' to one technology only. Under increasing returns then, the adoption process becomes a random walk with absorbing barriers. I leave it to the reader to show that the allocation process with diminishing returns appears to our observer as a random walk with *reflecting* barriers given by expressions similar to (2) and (3).

Properties of the Three Regimes

We can now use the elementary theory of random walks to derive the properties of this choice process under the different linear returns regimes. For convenient reference the results are summarised in Table 3. To prove these properties, we

Table 3
Properties of the Three Regimes

| | Predictable | Flexible | Ergodic | Necessarily path-efficient |
|---------------------|-------------|----------|---------|----------------------------|
| Constant returns | Yes | No | Yes | Yes |
| Diminishing returns | Yes | Yes | Yes | Yes |
| Increasing returns | No | No | No | No |

need first to examine long-term adoption shares. Under constant returns, the market is shared. In this case the random walk ranges free, but we know from random walk theory that the standard deviation of d_n increases with \sqrt{n} . It follows that the $d_n/2n$ term in equation (1) disappears and that x_n tends to 0.5 (with probability one), so that the market is split 50-50. In the diminishing returns case, again the adoption market is shared. The difference-in-adoption, d_n , is trapped between finite constants; hence $d_n/2n$ tends to zero as n goes to infinity, and x_n must approach 0.5. (Here the 50-50 market split results from the returns falling at the same rate.) In the increasing-returns-absorbing-barrier case, by contrast, the adoption share of A must eventually become zero or one. This is because in an absorbing random walk d_n eventually crosses a barrier with probability one. Therefore the two technologies cannot coexist indefinitely: one *must* exclude the other.

Predictability is therefore guaranteed where the returns are constant, or diminishing: in both cases a forecast that the market will settle to 50-50 will be correct, with probability one. In the increasing returns case, however, for accuracy the observer must predict A 's eventual share either as 0 or 100%. But either choice will be wrong with probability one-half. Predictability is lost. Notice though that the observer *can* predict that one technology will take the market; theoretically he can also predict that it will be A with probability $s(a_R - b_R) / [s(a_R - b_R) + r(b_S - a_S)]$; but he cannot predict the actual market-share outcome with any accuracy – in spite of his knowledge of supply and demand conditions.

Flexibility in the constant-returns case is at best partial. Policy adjustments

to the returns can affect choices at all times, but only if they are large enough to bridge the gap in preferences between technologies. In the two other regimes adjustments correspond to a shift of one or both of the barriers. In the diminishing-returns case, an adjustment g can always affect future choices (in absolute numbers, if not in market shares), because reflecting barriers continue to influence the process (with probability one) at times in the future. Therefore diminishing returns are flexible. Under increasing returns however, once the process is absorbed into A and B , the subsidy or tax adjustment necessary to shift the barriers enough to influence choices (a precise index of the degree to which the system is 'locked-in') increases without bound. Flexibility does not hold.

Ergodicity can be shown easily in the constant and diminishing returns cases. With constant returns only extraordinary line-ups (for example, twice as many R -agents as S -agents appearing indefinitely) with associated probability zero can cause deviation from fifty-fifty. With diminishing returns, any sequence of historical events – any line-up of the agents – must still cause the process to remain between the reflecting barriers and drive the market to fifty-fifty. Both cases forget their small-event history. In the increasing returns case the situation is quite different. Some proportion of agent sequences causes the market outcome to 'tip' towards A , the remaining proportion causes it to 'tip' towards B . (Extraordinary line-ups – say S followed by R followed by S followed by R and so on indefinitely – that could cause market sharing, have probability or measure zero.) Thus, the small events that determine $\{t_i\}$ decide the path of market shares; the process is non-ergodic or path-dependent – it is determined by its small-event history.

Path-efficiency is easy to prove in the constant- and diminishing-returns cases. Under constant-returns, previous adoptions do not affect pay-off. Each agent-type chooses its preferred technology and there is no gain foregone by the failure of the lagging technology to receive further development (further adoption). Under diminishing returns, if an agent chooses the technology that is ahead, he must prefer it to the available version of the lagging one. But further adoption of the lagging technology by definition lowers its payoff. Therefore there is no possibility of choices leading the adoption process down an inferior development path. Under increasing returns, by contrast, development of an inferior option can result. Suppose the market locks in to technology A . R -agents do not lose; but S -agents would each gain ($b_S - a_S$) if their favoured technology B had been equally developed and available for choice. There is regret, at least for one agent type. Inefficiency can be exacerbated if the technologies improve at different rates. An early run of agent-types who prefer an initially attractive but slow-to-improve technology can lock the market in to this inferior option; equal development of the excluded technology in the long run would pay off better to both types.

Extensions, and the Rational Expectations Case

It is not difficult to extend this basic model in various directions. The same qualitative results hold for M technologies in competition, and for agent types

in unequal proportions (here the random walk 'drifts'). And if the technologies arrive in the market at different times, once again the dynamics go through as before, with the process now starting with initial n_A or n_B not at zero. Thus in practice an early-start technology may already be locked in, so that a new potentially-superior arrival cannot gain a footing.

Where agent numbers are finite, and not expanding indefinitely, absorption or reflection and the properties that depend on them still assert themselves providing agent numbers are large relative to the numerical width of the gap between switching barriers.

For technologies *sponsored* by firms, would the possibility of strategic action alter the outcomes just described? A complete answer is not yet known. Hanson (1985) shows in a model based on the one above that again market exclusion goes through: firms engage in penetration pricing, taking losses early on in exchange for potential monopoly profits later, and all but one firm exit with probability one. Under strong discounting, however, firms may be more interested in immediate sales than in shutting rivals out, and market sharing can reappear.⁶

Perhaps the most interesting extension is the expectations case where agents' returns are affected by the choices of future agents. This happens for example with *standards*, where it matters greatly whether later users fall in with one's own choice. Katz and Shapiro (1985, 1986) have shown, in a two-period case with strategic interaction, that agents' expectations about these future choices act to destabilise the market. We can extend their findings to our stochastic-dynamic model. Assume agents form expectations in the shape of beliefs about the type of stochastic process they find themselves in. When the *actual* stochastic process that results from these beliefs is identical with the *believed* stochastic process, we have a rational-expectations fulfilled-equilibrium process. In the Appendix, I show that under increasing returns, rational expectations also yield an absorbing random walk, but one where expectations of lock-in hasten lock-in, narrowing the absorption barriers and worsening the fundamental market instability.

II. A GENERAL FRAMEWORK

It would be useful to have an analytical framework that could accommodate sequential-choice problems with more general assumptions and returns mechanisms than the basic model above. In particular it would be useful to know under what circumstances a competing-technologies adoption market must end up dominated by a single technology.

In designing a general framework it seems important to preserve two properties: (i) That choices between alternative technologies may be affected by the numbers of each adopted at the time of choice; (ii) That small events 'outside the model' may influence adoptions, so that randomness must be allowed for. Thus adoption market shares may determine not the next

⁶ For similar findings see the literature on the dynamics of commodity competition under increasing returns (e.g. Spence, 1981; Fudenberg and Tirole, 1983).

technology chosen directly but rather the *probability* of each technology's being chosen.

Consider then a dynamical system where one of K technologies is adopted each time an adoption choice is made, with probabilities $p_1(\mathbf{x}), p_2(\mathbf{x}), \dots, p_K(\mathbf{x})$, respectively. This vector of probabilities \mathbf{p} is a function of the vector \mathbf{x} , the adoption-shares of technologies 1 to K , out of the total number n of adoptions so far. The initial vector of proportions is given as \mathbf{x}_0 . I will call $\mathbf{p}(\mathbf{x})$ the *adoption function*.

We may now ask what happens to the long run proportions or adoption shares in such a dynamical system. Consider the two different adoption functions in Fig. 2, where $K = 2$. Now, where the probability of adoption of A is higher than its market share, in the adoption process A tends to increase in proportion; and where it is lower, A tends to decrease. If the proportions or adoption-shares settle down as total adoptions increase, we would conjecture that they settle down at a fixed point of the adoption function.

In 1983 Arthur, Ermoliev, and Kaniovski proved that under certain

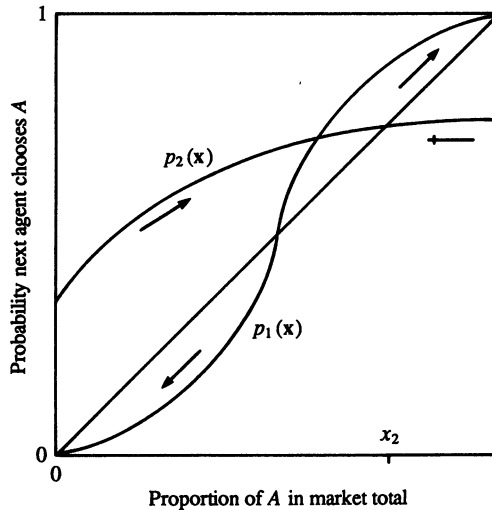


Fig. 2. Two illustrative adoption functions.

technical conditions (see the Appendix) this conjecture is true. A stochastic process of this type converges with probability one to one of the fixed points of the mapping from proportions (adoption shares) to the probability of adoption. Not all fixed points are eligible. Only 'attracting' or stable fixed points (ones that expected motions of the process lead towards) can emerge as the long run outcomes. And where the adoption function varies with time n , but tends to a limiting function \mathbf{p} , the process converges to an attracting fixed point of \mathbf{p} .

Thus in Fig. 2 the possible long-run shares are 0 and 1 for the function p_1 and x_2 for the function p_2 . Of course, where there are multiple fixed points, which one is chosen depends on the path taken by the process: it depends on the cumulation of random events that occur as the process unfolds.

We now have a general framework that immediately yields two useful theorems on path-dependence and single-technology dominance.

THEOREM I. *An adoption process is non-ergodic and non-predictable if and only if its adoption function \mathbf{p} possesses multiple stable fixed points.*

THEOREM II. *An adoption process converges with probability one to the dominance of a single technology if and only if its adoption function \mathbf{p} possesses stable fixed points only where \mathbf{x} is a unit vector.*

These theorems follow as simple corollaries of the basic theorem above. Thus where two technologies compete, the adoption process will be path-dependent (multiple fixed points must exist) as long as there exists at least one unstable 'watershed' point in adoption shares, above which adoption of the technology with this share becomes self-reinforcing in that it tends to increase its share, below which it is self-negating in that it tends to lose its share. It is therefore not sufficient that a technology gain advantage with adoption; the advantage must (at some market share) be self-reinforcing (see Arthur, 1988).

Non-Linear Increasing Returns with a Continuum of Adopter Types

Consider, as an example, a more general version of the basic model above, with a continuum of adopter types rather than just two, choosing between K technologies, with possibly non-linear improvements in payoffs. Assume that if n_j previous adopters have chosen technology j previously, the next agent's payoff to adopting j is $\Pi_j(n_j) = a_j + r(n_j)$ where a_j represents the agent's 'natural preference' for technology j and the monotonically increasing function r represents the technological improvement that comes with previous adoptions. Each adopter has a vector of natural preferences $\mathbf{a} = (a_1, a_2, \dots, a_K)$ for the K alternatives, and we can think of the continuum of agents as a distribution of points \mathbf{a} (with bounded support) on the positive orthant. We assume an adopter is drawn at random from this probability distribution each time a choice occurs. Dominance of a single technology j corresponds to positive probability of the distribution of payoffs Π being driven by adoptions to a point where Π_j exceeds Π_i for all $i \neq j$.

The Arthur-Ermoliev-Kaniovski theorem above allows us to derive:

THEOREM III. *If the improvement function r increases at least at rate ϵ as n_j increases, the adoption process converges to the dominance of a single technology, with probability one.*

Proof. In this case, the adoption function varies with total adoptions n . (We do not need to derive it explicitly however.) It is not difficult to establish that as n becomes large: (i) At any point in the neighbourhood of any unit vector of adoption shares, unbounded increasing returns cause the corresponding technology to dominate *all* choices; therefore the unit-vector shares are stable fixed points. (ii) The equal-share point is also a fixed point, but unstable. (iii) No other point is a fixed point. Therefore, by the general theorem, since the limiting adoption function has stable fixed points only at unit vectors the

process converges to one of these with probability one. Long-run dominance by a single technology is assured. ■

Dominance by a single technology is no longer inevitable, however, if the improvement function r is bounded, as when learning effects become exhausted. This is because certain sequences of adopter types could bid improvements for two or more technologies upward more or less in concert. These technologies could then reach the upper bound of r together, so that none of these would dominate and the market would remain shared from then on. Under other adopter sequences, by contrast, one of the technologies may reach the upper bound sufficiently fast to shut the others out. Thus, in the bounded case, some event histories dynamically lead to a shared market; other event histories lead to dominance. Increasing returns, if they are bounded, are in general *not* sufficient to guarantee eventual monopoly by a single technology.

III. REMARKS

(1) To what degree might the actual economy be locked-in to inferior technology paths? As yet we do not know. Certainly it is easy to find cases where an early-established technology becomes dominant, so that later, superior alternatives cannot gain a footing.⁷ Two important studies of historical events leading to lock-ins have now been carried out: on the QWERTY typewriter keyboard (David, 1985); and on alternating current (David and Bunn, 1987). (In both cases increasing returns arise mainly from coordination externalities.)

Promising empirical cases that may reflect lock-in through learning are the nuclear-reactor technology competition of the 1950s and 1960s and the US steam-versus-petrol car competition in the 1890s. The US nuclear industry is practically 100% dominated by light-water reactors. These reactors were originally adapted from a highly compact unit designed to propel the first nuclear submarine, the U.S.S. *Nautilus*, launched in 1954. A series of circumstances – among them the Navy's role in early construction contracts, political expediency, the Euratom programme, and the behaviour of key personages – acted to favour light water. Learning and construction experience gained early on appear to have locked the industry in to dominance of light water and shut other reactor types out (Bupp and Darian, 1978; Cowan, 1987). Yet much of the engineering literature contends that, given equal development, the gas-cooled reactor would have been superior (see Agnew, 1981). In the petrol-versus-steam car case, two different developer types with predilections toward steam or petrol depending on their previous mechanical experience, entered the industry at varying times and built upon on the best available versions of each technology. Initially petrol was held to be the less

⁷ Examples might be the narrow gauge of British railways (Kindleberger, 1983); the US colour television system; the 1950s programming language FORTRAN; and of course the QWERTY keyboard (Arthur, 1984; David, 1985; Hartwick, 1985). In these particular cases the source of increasing returns is network externalities however rather than learning effects. Breaking out of locked-in technological standards has been investigated by Farrell and Saloner (1985, 1986).

promising option: it was explosive, noisy, hard to obtain in the right grade, and it required complicated new parts.⁸ But in the United States a series of trivial circumstances (McLaughlin, 1954; Arthur, 1984) pushed several key developers into petrol just before the turn of the century and by 1920 had acted to shut steam out. Whether steam might have been superior given equal development is still in dispute among engineers (see Burton, 1976; Strack, 1970).

(2) The argument of this paper suggests that the interpretation of economic history should be different in different returns regimes. Under constant and diminishing returns, the evolution of the market reflects only *a-priori* endowments, preferences, and transformation possibilities; small events cannot sway the outcome. But while this is comforting, it reduces history to the status of mere carrier – the deliverer of the inevitable. Under increasing returns, by contrast many outcomes are possible. Insignificant circumstances become magnified by positive feedbacks to ‘tip’ the system into the actual outcome ‘selected’. The small events of history become important.⁹ Where we observe the predominance of one technology or one economic outcome over its competitors we should thus be cautious of any exercise that seeks the means by which the winner’s innate ‘superiority’ came to be translated into adoption.

(3) The usual policy of letting the superior technology reveal itself in the outcome that dominates is appropriate in the constant and diminishing-returns cases. But in the increasing returns case *laissez-faire* gives no guarantee that the ‘superior’ technology (in the long-run sense) will be the one that survives. Effective policy in the (un-sponsored) increasing-returns case would be predicated on the nature of the market breakdown: in our model early adopters impose externalities on later ones by rationally choosing technologies to suit only themselves; missing is an inter-agent market to induce them to explore promising but costly infant technologies that might pay off handsomely to later adopters.¹⁰ The standard remedy of assigning to early developers (patent) rights of compensation by later users would be effective here only to the degree that early developers can appropriate later payoffs. As an alternative, a central authority could underwrite adoption and exploration along promising but less popular technological paths. But where eventual returns to a technology are hard to ascertain – as in the U.S. Strategic Defence Initiative case for example – the authority then faces a classic multi-arm bandit problem of choosing which technologies to bet on. An early run of disappointing results (low ‘jackpots’) from a potentially superior technology may cause it

⁸ Amusingly, Fletcher (1904) writes: ‘...unless the objectionable features of the petrol carriage can be removed, it is bound to be driven from the road by its less objectionable rival, the steam-driven vehicle of the day.’

⁹ For earlier recognition of the significance of both non-convexity and path-dependence for economic history see David (1975).

¹⁰ Competition between *sponsored* technologies suffers less from this missing market. Sponsoring firms can more easily appropriate later payoffs, so they have an incentive to develop initially costly, but promising technologies. And financial markets for sponsoring investors together with insurance markets for adopters who may make the ‘wrong’ choice, mitigate losses for the risk-averse. Of course, if a product succeeds and locks-in the market, monopoly-pricing problems may arise. For further remarks on policy see David (1987).

perfectly rationally to abandon this technology in favour of other possibilities. The fundamental problem of possibly locking-in a regrettable course of development remains (Cowan, 1987).

IV. CONCLUSION

This paper has attempted to go beyond the usual static analysis of increasing-returns problems by examining the dynamical process that 'selects' an equilibrium from multiple candidates, by the interaction of economic forces and random 'historical events'. It shows how dynamically, increasing returns can cause the economy gradually to lock itself in to an outcome not necessarily superior to alternatives, not easily altered, and not entirely predictable in advance.

Under increasing returns, competition between economic objects – in this case technologies – takes on an evolutionary character, with a 'founder effect' mechanism akin to that in genetics.¹¹ 'History' becomes important. To the degree that the technological development of the economy depends upon small events beneath the resolution of an observer's model, it may become impossible to predict market shares with any degree of certainty. This suggests that there may be theoretical limits, as well as practical ones, to the predictability of the economic future.

Stanford University

Date of receipt of final typescript: May 1988

APPENDIX

A. Definitions of the Properties

Here I define precisely the properties used above. Denote the market share of A after n choices as x_n . The allocation process is:

(i) *predictable* if the observer can *ex-ante* construct a forecasting sequence $\{x_n^*\}$ with the property that $|x_n - x_n^*| \rightarrow 0$, with probability one, as $n \rightarrow \infty$;

(ii) *flexible* if a given marginal adjustment g to the technologies' returns can alter future choices;

(iii) *ergodic* if, given two samples from the observer's set of possible historical events, $\{t_i\}$ and $\{t'_i\}$, with corresponding time-paths $\{x_n\}$ and $\{x'_n\}$, then $|x'_n - x_n| \rightarrow 0$, with probability one, as $n \rightarrow \infty$;

(iv) *path-efficient* if, whenever an agent chooses the more-adopted technology α , versions of the lagging technology β would not have delivered more had they been developed and available for adoption. That is, path-efficiency holds if returns Π remain such that $\Pi_\alpha(m) \geq \text{Max}_j \{\Pi_\beta(j)\}$ for $k \leq j \leq m$, where there have been m previous choices of the leading technology and k of the lagging one.

¹¹ For other selection mechanisms affecting technologies see Dosi (1988), Dosi *et al.* (1988), and Metcalfe (1985).

B. The Expectations Case

Consider here the competing standards case where adopters are affected by *future* choices as well as past choices. Assume in our earlier model that R -agents receive additional net benefits of Π_A^R, Π_B^R , if the process locks-in to their choice, A or B respectively; similarly S -agents receive Π_A^S, Π_B^S . (Technologies improve with adoption as before.) Assume that agents know the state of the market (n_A, n_B) when choosing and that they have expectations or beliefs that adoptions follow a stochastic process Ω . They choose rationally under these expectations, so that actual adoptions follow the process $\Gamma(\Omega)$. This actual process is a *rational expectations equilibrium process* when it bears out the expected process, that is, when $\Gamma(\Omega) \equiv \Omega$.

We can distinguish two cases, corresponding to the degree of heterogeneity of preferences in the market.

Case (i). Suppose initially that $a_R - b_R > \Pi_B^R$ and $b_S - a_S > \Pi_A^S$ and that R and S -types have beliefs that the adoption process is a random walk Ω with absorption barriers at Δ'_R, Δ'_S , with associated probabilities of lock-in to A , $P(n_A, n_B)$ and lock-in to B , $1 - P(n_A, n_B)$. Under these beliefs, R -type expected payoffs for choosing A or B are, respectively:

$$a_R + m_A + P(n_A, n_B)\Pi_A^R \quad (4)$$

$$b_R + m_B + [1 - P(n_A, n_B)]\Pi_B^R. \quad (5)$$

S -type payoffs may be written similarly. In the actual process R -types will switch to B when n_A and n_B are such that these two expressions become equal. Both types choose B from then on. The actual probability of lock-in to A is zero here; so that if the expected process is fulfilled, P is also zero here and we have n_A and n_B such that

$$a_R + m_A = b_R + m_B + \Pi_B^R$$

with associated barrier given by

$$\Delta_R = n_A - n_B = -(a_R - b_R - \Pi_B^R)/r. \quad (6)$$

Similarly S -types switch to A at boundary position given by

$$\Delta_S = n_A - n_B = (b_S - a_S - \Pi_A^S)/s. \quad (7)$$

It is easy to confirm that beyond these barriers the actual process is indeed locked in to A or to B and that within them R -agents prefer A , and S -agents prefer B . Thus if agents believe the adoption process is a random walk with absorbing barriers Δ'_R, Δ'_S given by (6) and (7), these beliefs will be fulfilled, and this random walk will be a rational expectations equilibrium.

Case (ii). Suppose now that $a_R - b_R < \Pi_B^R$ and $b_S - a_S < \Pi_A^S$. Then (4) and (5) show that switching will occur immediately if agents hold expectations that the system will definitely lock-in to A or to B . These expectations become self-fulfilling and the absorbing barriers narrow to zero. Similarly, when non-improving standards compete, so that r and s are zero, in this case again beliefs that A or B will definitely lock-in become self-fulfilling.

Taking cases (i) and (ii) together, expectations either narrow or collapse the switching boundaries. They exacerbate the fundamental market instability.

C. The Path-Dependent Strong-Law Theorem

Consider a dependent-increment stochastic process that starts with an initial vector of units \mathbf{b}_0 , in the K categories, 1 through K . At each event-time a unit is added to one of the categories 1 through K , with probabilities $\mathbf{p} = [p_1(x), p_2(x), \dots, p_K(x)]$, respectively. (The Borel function \mathbf{p} maps the unit simplex of proportions S^K into the unit simplex of probabilities S^K .) The process is iterated to yield the vectors of proportions $\mathbf{X}_1, \mathbf{X}_2, \mathbf{X}_3, \dots$.

THEOREM. *Arthur, Ermoliev, and Kaniovski (1983, 1986)*

(i) Suppose $p: S^K \rightarrow S^K$ is continuous, and suppose the function $p(x) - x$ possesses a Lyapunov function (that is, a positive, twice-differentiable function V with inner product $\{[p(x) - x], V_x\}$ negative). Suppose also that the set of fixed points of p , $B = \{x: p(x) = x\}$ has a finite number of connected components. Then the vector of proportions $\{X_n\}$ converges, with probability one, to a point z in the set of fixed points B , or to the border of a connected component.

(ii) Suppose p maps the interior of the unit simplex into itself, and that z is a stable point (as defined in the conventional way). Then the process has limit point z with positive probability.

(iii) Suppose z is a non-vertex unstable point of p . Then the process cannot converge to z with positive probability.

(iv) Suppose probabilities of addition vary with time n , and the sequence $\{p_n\}$ converges to a limiting function p faster than $1/n$ converges to zero. Then the above statements hold for the limiting function p . That is, if the above conditions are fulfilled, the process converges with probability one to one of the stable fixed points of the limiting function p .

The theorem is extended to non-continuous functions p and to non-unit and random increments in Arthur, Ermoliev and Kaniovski (1987*b*). For the case $K = 2$ with p stationary see the elegant analysis of Hill *et al.* (1980).

REFERENCES

- Agnew, H. (1981). 'Gas-cooled nuclear power reactors.' *Scientific American*, vol. 244, pp. 55-63.
- Arthur, W. B. (1983). 'Competing technologies and lock-in by historical small events: the dynamics of allocation under increasing returns.' International Institute for Applied Systems Analysis Paper WP-83-92, Laxenburg, Austria. (Center for Economic Policy Research, Paper 43, Stanford).
- (1984). 'Competing technologies and economic prediction.' *Options*, International Institute for Applied Systems Analysis, Laxenburg, Austria. No. 1984/2, pp. 10-3.
- Arthur, W. B., Ermoliev, Yu. and Kaniovski, Yu. (1983). 'On generalized urn schemes of the polya kind.' *Kibernetika*, vol. 19, pp. 49-56. English translation in *Cybernetics*, vol. 19, pp. 61-71.
- and — (1986). 'Strong laws for a class of path-dependent urn processes.' In *Proceedings of the International Conference on Stochastic Optimization, Kiev 1984*, Springer Lecture Notes Control and Information Sciences, vol. 81, pp. 187-300.
- and — (1987*a*). 'Path-dependent processes and the emergence of macro-structure.' *European Journal of Operational Research*, vol. 30, pp. 294-303.
- and — (1987*b*). 'Non-linear urn processes: asymptotic behavior and applications.' International Institute for Applied Systems Analysis. Paper WP-87-85, Laxenburg, Austria.
- (1988). 'Self-reinforcing mechanisms in economics.' In *The Economy as an Evolving Complex System* (P. Anderson, K. Arrow and D. Pines (eds)). Reading, Massachusetts: Addison-Wesley.

- Atkinson, A. and Stiglitz, J. (1969). 'A new view of technical change.' *Economic Journal*, vol. 79, pp. 573-80.
- Balcer, Y. and Lippman, S. (1984). 'Technological expectations and the adoption of improved technology,' *Journal of Economic Theory*, vol. 34, pp. 292-318.
- Bupp, I. and Derian, J. (1978). *Light Water: How The Nuclear Dream Dissolved*. New York: Basic.
- Burton, R. (1976). 'Recent advances in vehicular steam engine efficiency.' Society of Automotive Engineers, Preprint 760340.
- Cowan, R. (1987). 'Backing the wrong horse: sequential choice among technologies of unknown merit.' Ph.D. Dissertation, Stanford.
- David, P. (1975). *Technical Choice, Innovation, and Economic Growth*. Cambridge: Cambridge University Press.
- (1985). 'Clio and the economics of QWERTY.' *American Economic Review Proceedings*, vol. 75, pp. 332-7.
- (1987). 'Some new standards for the economics of standardization in the information age.' In *Economic Policy and Technological Performance*, (P. Dasgupta and P. Stoneman, eds.) Cambridge: Cambridge University Press.
- and Bunn, J. (1987). 'The economics of gateway technologies and network evolution: lessons from electricity supply history.' Centre for Economic Policy Research, Paper 119, Stanford.
- Dosi, G. (1988). 'Sources, procedures and microeconomic effects of innovation.' *Journal of Economic Literature* vol. 26, pp. 1120-71.
- , Freeman, C., Nelson, R., Silverberg, G. and Soete, L. (eds.) (1988). *Technical Change and Economic Theory*, London: Pinter.
- Farrell, J. and Saloner, G. (1985). 'Standardization, compatibility, and innovation.' *Rand Journal of Economics*, vol. 16, pp. 70-83.
- (1986). 'Installed base and compatibility: innovation, product preannouncements and predation.' *American Economic Review*, vol. 76, pp. 940-55.
- Fletcher, W. (1904). *English and American Steam Carriages and Traction Engines* (reprinted 1973). Newton Abbot: David and Charles.
- Fudenberg, D. and Tirole, J. (1983). 'Learning by doing and market performance.' *Bell Journal of Economics*, vol. 14, pp. 522-30.
- Hanson, W. (1985). 'Bandwagons and orphans: dynamic pricing of competing systems subject to decreasing costs.' Ph.D. Dissertation, Stanford.
- Hartwick, J. (1985). 'The persistence of QWERTY and analogous suboptimal standards.' Mimeo, Queen's University, Kingston, Ontario.
- Hill, B., Lane, D. and Sudderth, W. (1980). 'A strong law for some generalized urn processes.' *Annals of Probability*, vol. 8, pp. 214-26.
- Katz, M. and Shapiro, C. (1985). 'Network externalities, competition, and compatibility.' *American Economic Review*, vol. 75, pp. 424-40.
- , and — (1986). 'Technology adoption in the presence of network externalities.' *Journal of Political Economy*, vol. 94, pp. 822-41.
- Kindleberger, C. (1983). 'Standards as public, collective and private goods.' *Kyklos*, vol. 36, pp. 377-96.
- Mamer, J. and McCardle, K. (1987). 'Uncertainty, competition and the adoption of new technology.' *Management Science*, vol. 33, pp. 161-77.
- McLaughlin, C. (1954). 'The Stanley Steamer: a study in unsuccessful innovation.' *Explorations in Entrepreneurial History*, vol. 7, pp. 37-47.
- Metcalf, J. S. (1985). 'On technological competition.' Mimeo, University of Manchester.
- Rosenberg, N. (1982). *Inside the Black Box: Technology and Economics*. Cambridge: Cambridge University Press.
- Strack, W. (1970). 'Condensers and Boilers for Steam-powered Cars.' NASA Technical Note, TN D-5813, Washington, D.C.
- Spence, A. M. (1981). 'The learning curve and competition.' *Bell Journal of Economics*, vol. 12, pp. 49-70.