TEXTOS PARA DISCUSSÃO

nº 313

Are Technology Diffusion Processes Inherently Historical?

J. L. de Araújo

Instituto de Economia Industrial
UFRJ

Universidade Federal do Rio de Janeiro
Instituto de Economia Industrial

Textos para Discussão

Are Technology Diffusion Processes Inherently Historical?

J. L. de Araújo

junho 1994

Instituto de Economia Industrial
U F R J

Diretor
José Ricardo Tauile
Coordenador de Ensino
Carlos Aguiar de Medeiros
Coordenador de Pesquisa
João Luis Maurity Sabóia
Gerente Administrativa
Sebastiana de Sousa Barros
Supervisão
Maria Lúcia Ferro
Projeto Gráfico
José Antonio de Oliveira
Editoração Eletrônica
Jorge Amaro
Impressão
Célio de Almeida Mentor e Ronei José Gomes

Ficha Catalográfica

ARAÚJO, João Lisardo de

Are Technology Diffusion Processes Inherently Historical? / J. L. de Araújo. -- Rio de Janeiro: UFRJ/IEI 1994.

27 p., 21 cm. --(Texto para Discussão IEI/UFRJ; nº 313)

Inclui gráficos, tabelas

1. Difusão de Tecnologia. 2. Inovações Tecnológicas. I. Título. II. Série

Universidade Federal do Rio de Janeiro
Instituto de Economia Industrial
Pálacio Universidade do Brasil
Av. Pasteur, 250 - Praia Vermelha
CEP 22290 - Rio de Janeiro - RJ
295 1447 e 541 8148 (fax)

Are Technology Diffusion Processes Inherently Historical?

1. *Introduction*

Since the first efforts to model technological diffusion processes (Ayres, 1969), a basic tool has been the logistic curve deriving from the diffusion equation:

$$(1)\ \dot{n}(t) = \alpha \cdot n(t) \cdot \left(M - n(t)\right)$$

In this simple "contagion model" a is a diffusion coefficient, M is the population of "susceptibles" and n(t) is the number of "infected" at time t. This equation has been adapted to more complex models to try and account for observed complexities. It has two strong points: it is simple and has shown a good fit to a variety of phenomena. There is however a fundamental drawback in it for economists: it is essentially deterministic and its parameters have to be empirically determined without reference to economic theory. Unsurprisingly, efforts were made during the eighties to overcome both faults. Such efforts may be divided into three general categories: efforts embedding technology diffusion within the neo-classical framework, efforts modelling diffusion as a random process, and efforts using an evolutionary approach.

The first school of research (exemplified by Cameron and Metcalfe, 1987) analyses technology diffusion from the point of view of economic theory. It starts from supply and demand schedules for both a established technology and for an innovation. From this, niches are deduced for both old and new technologies under neo-classical assumptions. Furthermore, with the aid of an empirically determinable contagion parameter, it may be shown that a logistic equilibrium path exists under these conditions.

The second and third schools are both concerned with the determinism of the diffusion equation. Interesting results were obtained by Arthur (Arthur, 1989; Arthur et al., 1987), treating the competition between two innovations A and B as a random walk with absorbing barriers and with transition probabilities dependent on the relative returns of each. This model departs from the contagion approach; instead, a stream of agents enters the market for technologies, and examines the two technologies available using both their basic parameters (two types of agents are assumed, each armed with a set of parameters favouring one technology) and the number of choices already made in favour of each (n_A and n_B). The sequence n_A-n_B is a random walk with absorbing barriers.

When returns increase with accumulated adoption the process is unstable, ending in one technology completely expelling the other. Furthermore, in this case the outcome is unpredictable: the economy may become locked-in by an inferior (considering long-term returns) technology. Even worse, any attempt by agents to forecast future evolution (e.g. by using rational expectations) intensifies unpredictability. Conversely, if returns are constant or decreasing with accumulated adoption the process is ergodic and predictable.

The third school avoids both neo-classical and probabilistic modelling. Instead, it tries to describe agents' behaviour in an evolutionary model, following a Schumpeterian tradition. A good example is given by Silverberg, Dosi and Orsenigo (1988). In this, the degrees of freedom present in the behavioural equations for distinct types of agents make for complex diffusion trajectories. This is a theoretically attractive approach and may achieve high explanatory power. However, the latter depends critically upon the confidence in the behavioural equations and on assumed values for parameters. Such confidence may decrease as the complexity of the model increases.

The result obtained by Arthur is interesting and suggests a line of investigation by focusing on random mechanisms. However, it also means a significant break with the diffusion model by doing away with the time dimension. Given the empirical evidence in favour of logistic diffusion paths, as well as theoretical results such as those by Cameron and Metcalfe, it seems worthwhile to explore stochastic approaches to the contagion models. Likewise, the complexity of evolutionary models may hide some basic features, that a simpler model could make visible.

The present paper treats non-deterministic diffusion trajectories for two competing technologies as a birth-and-death process, using the basic contagion equations as a starting point. This has the advantage of being comparable to earlier efforts and of addressing the question: how does the introduction of randomness affect the diffusion trajectories for two competing technologies? A probabilistic treatment of simple contagion is of course well known in the literature of stochastic processes and epidemiology, using however diffusion processes, rather than birth-and-death processes. For our purposes, birth-and-death processes have the attractive feature of modelling individual decisions which look random to an observer, as well as being able to deal with finite markets.

The structure of the paper is as follows. The second section goes from the basic diffusion process as a pure birth process to introduce the competing diffusion process as a two-dimensional, constrained birth process. This is discussed with a view to its properties. Section three generalizes the model to a birth-and-death process. Finally, the last section discusses implications and questions for investigations. Throughout the paper, theoretical results are illustrated with computer simulations.

2. Technology diffusion as a pure birth process

This section introduces the basic models, in their simplest form. The simplicity highlights some essential features, at the cost of realism.

2.1. The basic contagion process

Assume that a new technology is introduced at time t=0 by an innovator, and that the potential number of adopters totals M. The adoption process is assumed to follow a pure birth process with diffusion coefficient a (dependent on product and market features). That is to say, time intervals between adoptions are independent random variables, with exponential distribution with average time interval $1/\lambda_n$ between the nth and the (n+1)th adoptions. The parameter λ_n satisfies.

(2) $\lambda_n = \alpha \cdot n \cdot (M - n)$

This model may be interpreted in the following way. After n adoptions, each remaining potential adopter independently opts for the technology in the time interval Δt with the probability $\alpha \times n \times \Delta t + o(\Delta t)$, where the second term goes to zero faster than Δt; furthermore, the probability of simultaneous adoptions is zero. In other words, the rate of adoption is proportional to the number of adoptions already made times the number of remaining potential adopters. Note that probability in this context refers to the point of view of an external observer; it makes little economic sense to assume that an individual will throw dice to make a technological choice, although this might occur in other contexts.

The process $\{N(t), t \geq 0\}$, with N(0) = 1, is then a pure birth process with a single absorbing state (n=M); all the other states are transient. It is easy to show that N(t) converges to M with probability 1. Furthermore, realizations of the process show the familiar logistic pattern (see Figure 1).

Figure 1

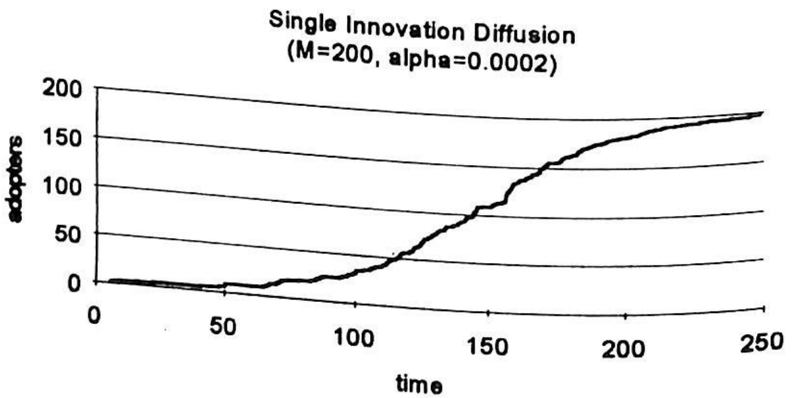

2.2. Competitive diffusion as a pure birth process

The simple diffusion process is trivial. The introduction of competition, however, makes things far more interesting. If two technologies, 1 and 2, are introduced at time 0 with diffusion coefficients α_1 and α_2, we have a two-dimensional process. This may be noted as $\{(N_1(t), N_2(t)), t \geq 0\}$, $N_1(0) = N_2(0) = 1$, where transitions in $N_1(t)$ and $N_2(t)$ are independent save for the 'birth rates' λ_1 and λ_2. These are now linked through the equation

(3) $\lambda_i = \alpha_i \cdot n_i \cdot (M - n_1 - n_2), i = 1,2$.

This may be interpreted as follows. During the time interval $(t, t + \Delta t)$ each remaining potential adopter has three choices: wait and see, choose 1, or choose 2. If Δt is small enough, Pr{choosing i} $\cong \alpha_i \cdot N_i(t) \cdot \Delta t$, i = 1,2. Furthermore, each agent chooses independently from the others, and the probability of two simultaneous choices is zero.

It is not difficult to show that all states of the form $n_1 + n_2 = M$ (numbering a total M-1 states, since the states (0, M) and (M, 0) are excluded in the pure birth model) are absorbing, and all others are transient. Therefore, the process is not ergodic; once it enters one absorbing state it cannot leave. Even though the distribution of final states may be connected with a_1 and a_2, the outcome is unpredictable. An illustration is given in Figure 2, which presents a histogram of one hundred simulation runs for two equally attractive technologies ($a_1 = a_2 = 0.0002$) competing for a market of 200 potential adopters.

Figure 2

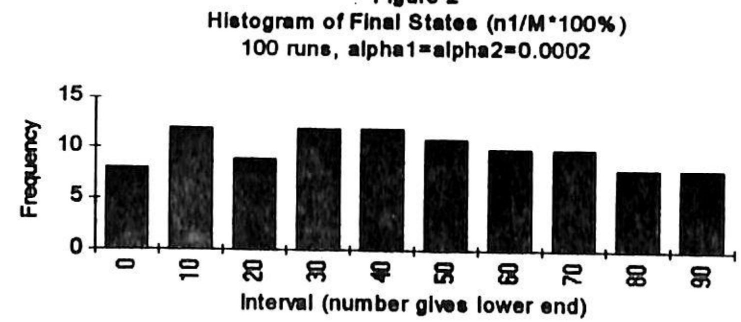

The flat distribution of final states may come as a surprise to those expecting a convergence to a 50% market share with probability one (this would be the result with a deterministic model). In fact, it is a direct consequence of a finite market. Finiteness is the important factor, not market size. One may easily show that, when $\alpha_1 = \alpha_2$, for any finite market size M and any final market split-up (n, M-n) with n between 1 and M-1, the probability that this split-up will be reached from the initial state (1,1) is 1/(M-1). In other words, the distribution of final states is uniform. See Appendix A.1 for a proof.

On the other hand, unequal diffusion rates lead to lopsided distributions of final states, as was to be expected. Neither in this case is there convergence to a state with probability one, although now size has an effect: larger markets tend to make the distribution more lopsided in favour of the technology with higher diffusion rate. For instance, if $\alpha_1 = 2\alpha_2$, when M = 10 the ratio between probabilities of final states (9,1) and (1,9) is 13.5 to 1; when M = 100, the ratio $P_{(1,1)}(99,1):P_{(1,1)}(1,99)$ is 441 to 1. In other words, if the market is sufficiently large a technology with higher diffusion rate tends to crowd out the other[1], provided they start diffusion at the same time and conditions. Figure 3 illustrates the distribution when one diffusion rate is the double of the other, both technologies start at time zero with one adopter each and the market size equals 100 potential adopters.

It is well to discuss a little the reasons why this should be so.

Figure 3
Probability Distribution of Final States
starting from (1,1)
M=100, alpha1=0.0002, alpha2=0.0001

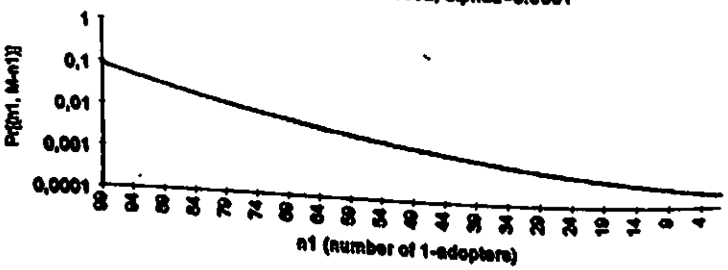

To begin with, the contagion process described by equation (3) intrinsically refers to a finite population of potential adopters. Letting M tend to infinity will make λ_1 and λ_2 infinite, if $n_1 + n_2 < \infty$, and undefined otherwise; in the infinite market, the logistic should be replaced by the exponential. For M finite, the process stops after a finite number of transitions and asymptotic results are useless. On the other hand, for any given n_1, n_2 and any M $> n_1 + n_2$, we have

(4) $\frac{\lambda_1(t)}{\lambda_2(t)} = \frac{\alpha_1 \cdot n_1}{\alpha_2 \cdot n_2}$, independently of M.

Thus, the size of the potential market affects only the rate at which adoptions are taking place. It does not affect at all the share of each technology in these adoptions. In contrast, this share is directly affected by the ratio n_1/n_2. Initial deviations are thus reinforced through it, with the result that each realization of the process is a pair of logistic curves, noisy but apparently well-behaved and arriving at an equilibrium.

Ex-post analysis will lead to estimates of diffusion rates with little or nor relation to parameters. First, it is not possible to identify α_1 and α_2 from observation, since this would require a reliable estimate of their potential markets; all one observes, however, are "equilibrium niches". Second, equilibrium points are unpredictable a priori, especially if diffusion rates are equal. For an illustration, see Figure 4. The outcome shown has precisely the same a priori probability as any other. The smoothness of the diffusion path is remarkable, despite its totally random origin.

Figure 4
Diffusion of two competing innovations
(simulation run)
M=1000, alpha1=alpha2= 0.0002

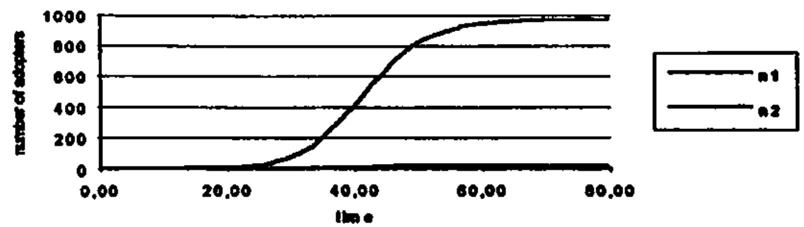

On the other hand, for unequal diffusion rates, in large markets a less "attractive" technology will tend to be crowded out if it starts in the same conditions as its competitor. In these cases, therefore, there would seem to be a degree of predictability.

However, if a technology lags the other it may lose out even if it has a higher diffusion rate (i.e., a higher intrinsic utility for buyers). Thus, if in the example of Figure 3 technology 1 (with diffusion rate double that of technology 2) enters when technology 2 already has 6 adopters (6% of the market), the final state distribution will be biased in favour of technology 2 (see Figures 5 and 6).

Figure 5
E[n1] given that initial state is (1,b)
for M=100, alpha1=0.0002 = 2*alpha2

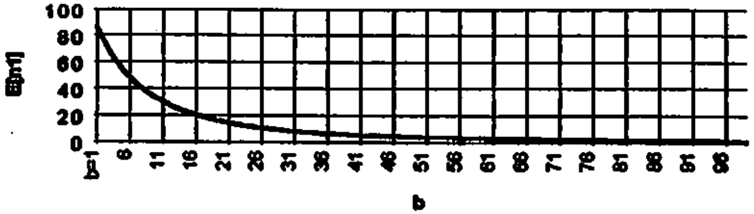

Figure 6
Distribution of final states (n1, M-n1) from (1, b) for several values of b.
M=100 and alpha1 = 0.002 = 2* alpha2

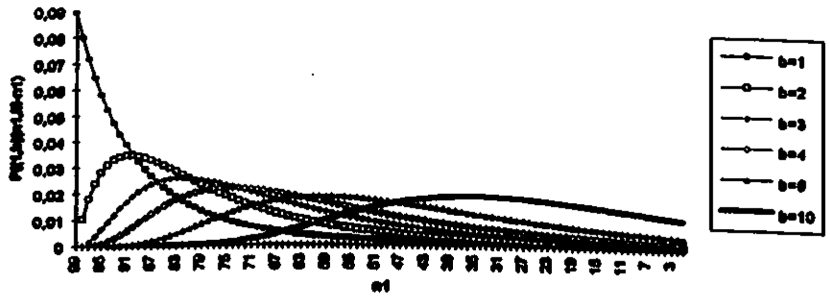

This suggests that historical circumstances are very important in defining the actual diffusion paths and saturation levels of competing technologies. The necessary lead to reverse the bias from a higher diffusion rate will tend to increase with market size in absolute terms, but not necessarily in relative terms. A closed expression is out of question, but a few results may be shown.

For small values of M, it is possible to evaluate $E[n_1 | b]$ for various b. Experimenting with different values of M between 10 and 120 and $\alpha_1 = 2\alpha_2$, we find that bmin, the initial advantage that allows technology 2 to offset its smaller diffusion rate in terms of expected market share, grows approximately as $O(M^{0.45})$. In other words, the relative offsetting lead decreases with growing market size. This experimentation is however limited by the fact that, for large M, the computational effort involved in calculating the distribution of final states, or even $E[n_1]$, may be excessive. Instead, a simple measure of the impact of b may be given by the ratio P2P1(M, b, b) = Pr{all adopt 2 | b} . Pr{all adopt 1 | b} for several values of M and $\beta = \alpha_2 / \alpha_1$, where the initial state is (1, b) and b < 1 (that is, technology 1 is more attractive than technology 2 but this latter starts with b adopters against 1).

We are interested in $b_{min}(M, b)$, the minimum b for which this ratio is greater than 1, for a given M and b. For all b at or above this value, the initial stock more than compensates for the disadvantage in diffusion rates. Even though this does not imply $E[n_1] < M/2$ (compare Figures 5 and 7), it is a meaningful measure. The figure below gives the result for a range of M and b. (See also Appendix A.2).

Figure 7
Minimum b for compensating beta in P2P1

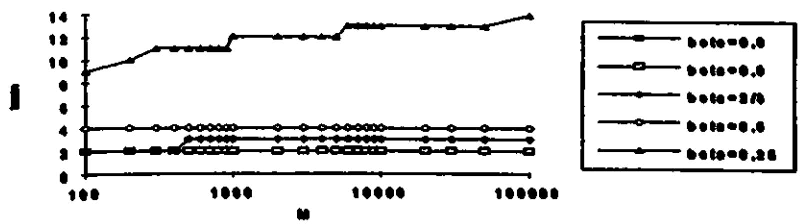

One striking feature is the insensitivity of $b_{min}(M, b)$ to market size. Only for large inequalities between diffusion rates is there a noticeable increase for the range of market sizes shown, and even then by several orders of magnitude less than market increase: When one diffusion rate is four times larger than the other, a thousandfold increase in market size leads to a 55% increase in b_{min}. In other words, for a market of a hundred thousand potential adopters a 4:1 disadvantage in diffusion rates may be offset by a lead of less than 0.002% of the market. Even more, the offsetting lead decreases in relative terms with market size. This confirms that competitive technology diffusion following a pure birth model is very sensitive to historical factors, even with relative intensities proportional to adoption.

3. Competitive diffusion as a birth-and-death process

We shall now extend the model to the case where regrets may occur, so that an adopter of technology i may abandon it. Contrary to the pure birth case, there is no natural model to follow. We shall therefore conduct most of the discussion using a general functional form for the rate at which adopters leave a technology.

To be more precise, we shall make the following assumption: given that at time t technology i has $n_i(t)$ adopters, the probability that one adopter will relinquish it and join the "potential adopters" crowd during the interval [t, t+h) tends to $\mu_i(n_1(t), n_2(t), M) \cdot h$ as $h \to 0^+$. That is to say, the instantaneous 'death rate' for each technology depends on the number of adopters of each technology and on the size of the market. The complete model is then

$$(5) \quad \begin{cases} \text{brith rate } \lambda_i = \alpha_i \cdot n_i(t) \cdot (M - n_1(t) - n_2(t)),\ i = 1,2 \\ \text{death rate } \theta_i = \mu_i(n_1(t), n_2(t), M),\ i = 1,2 \end{cases}$$

with the usual constraint that simultaneous events have zero probability. We further assume that q_i is positive for all $n_i > 0$.

With this formulation, the model undergoes a radical change. All states satisfying $n_1 > 0$ or $n_2 > 0$ are now transient. In effect, they cannot be reached from states having zero adopters for either technology, and the 'no-adopters' state ($n_1 = n_2 = 0$) behaves like an absorbing barrier. With (5), in the very long term neither technology will have adopters.

Curiously enough, even in this case realizations of the process will behave approximately as in the pure birth case, reaching apparently stable (on a finite observation time window) market shares. The figure below is an example, with the death rates linear in the number of adopters; in other words, $\theta_i = \mu_i \cdot n_i(t)$.

Figure 8

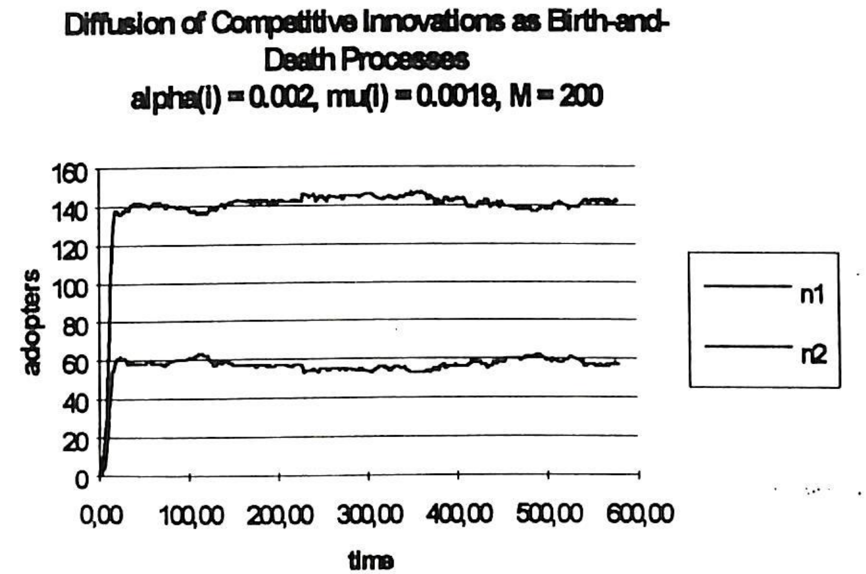

The aspect of the curve might be explained by the fact that, in the beginning, the process behaves like a pure birth process. Deaths only become relevant when market approaches saturation; and the oscillations around the equilibrium point suggest that for most practical time horizons the fact that (0,0) is an absorbing state is irrelevant. What is important is the distribution of the "transient equilibria"[2] such as the one shown in Figure 8 above. From the argument above, these should be distributed roughly as equilibria under a pure birth process.

One may also observe that.

$\lambda_i/\theta_i = \alpha_i/\mu_i \cdot (M - n_1 - n_2),\ i = 1,2.$ That is to say, as long as death coefficients are not much larger than diffusion coefficients births tend to dominate over deaths whenever the state drifts away from market saturation. Taken together with $\lambda_1/\lambda_2 = (n_1\alpha_1)/(n_2\alpha_2)$, this tends to lead to stability of the equilibrium, even though with infinite time there is extinction with probability 1.

Another question of interest is the limiting behaviour of the process if there is a renewal at extinction, i.e. whenever the process attains the state (0,0) it starts anew at a given state (1, b). With this change the chain becomes irreducible, allowing evaluation of limit probabilities and investigation of the consequences of initial advantages. Note, however, that for even b the chain is periodic with period 2. Moreover, differently from the remaining specifications, renewal at extinction does not allow a natural economic interpretation. Rather, it serves as a means to study what would happen in the very very long term if extinctions were excluded.

It is out of question to obtain limit probabilities for large markets, or even for moderate-sized ones. For instance, for M=100 we would have to find 5151 steady state probabilities, components of the eigenvector associated with the unit eigenvalue of a (M+1)(M+2)/2 by (M+1)(M+2)/2 matrix for the embedded Markov chain of the process. Nevertheless, it is possible to obtain the flavour of the solution through a small-scale instance.

Figure 9 shows the result for M=10; $a_i = 0.02$, $i = 1,2$; death rates are assumed linear in adoption as in Figure 8, with m_i = 0.02. It may be noticed that the least likely states are those with both n_1 and n_2 positive: all have limit probabilities between $1 \cdot 10^{-7}$ and $2 \cdot 10^{-6}$. Thus, the intrinsic dynamics of the competitive birth-and-death process with renewal tends to exclusive dominance of either technology, in an unpredictable way. Decreasing the death rate coefficients will only accentuate this feature; the probabilities of positive (n_1, n_2) drop by three orders of magnitude when death coefficients are halved.

Figure 9

Limit Probabilities for the competitive diffusion process with linear death rate, m(i)=0.02=α(i), and renewal at extinction to (1,1), and M= 10

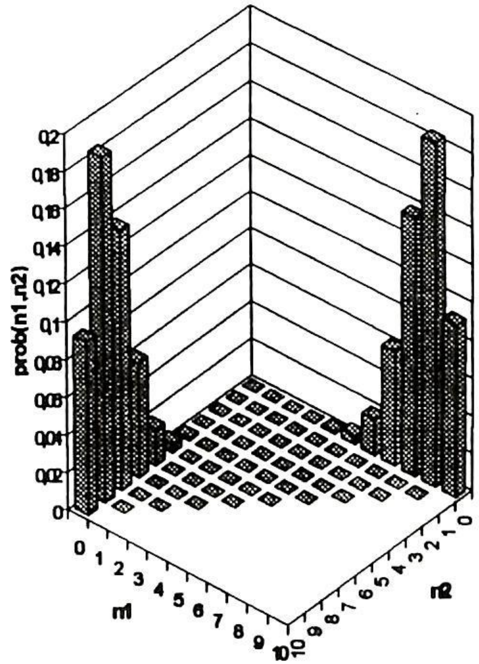

Since the model allows the investigation of initial advantages, let the renewal be to the state (1,3) instead of (1,1). Figure 10, below, shows the results.

Figure 10

Limit Probabilities for the competitive diffusion process with linear death rate, m(i)=0.02=a(i), renewal at extinction to (1, 3), and M = 10

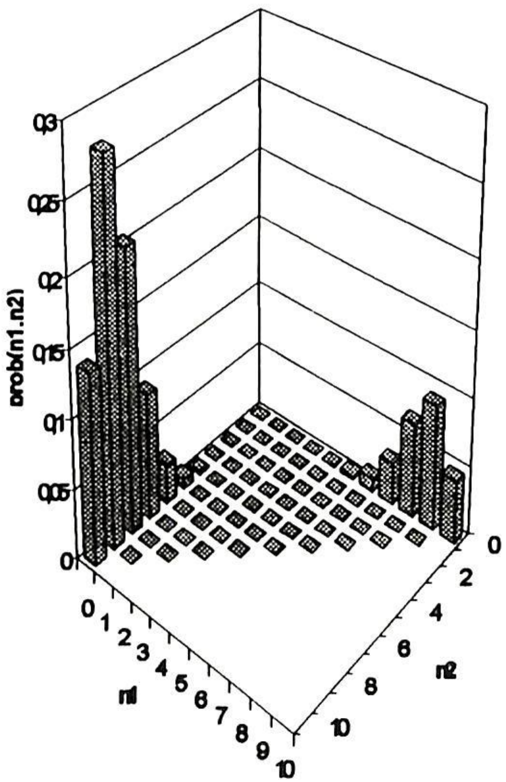

It may be observed that P(0,i) is approximately the triple of P(i,0), as was to be expected. It requires an decrease of a_2 to 0.0151 to compensate the initial advantage.

A question may be raised about the assumption of a linear death rate in this example; other functional dependencies might be proposed, such as a constant death rate (this would mean that the likelihood of one adopter abandoning a given technology varies inversely with the number of co-adopters of that technology). The effect of this modification will depend on the value chosen for the death rate θ, relative to M*α. Using θ_i ° 0.01 or 0.02 in the above example, the result is to decrease even more the probability of states with both technologies. On the other hand, large θ_i (so that at (M-1,0) and (0,M-1) birth and death rates are the same) brings about a distribution centred on small adoption values for both technologies. In this case, the probability of both technologies having positive numbers of adopters is not negligible. The same happens when in the linear death rate very large m's are used (in the above example, $m_i = 0.4 = 2 \times M \times \alpha_1$ — which implies a death rate larger than the birth rate for $n_i \geq 1$ leads to $P(n_1 \times n_2 > 0) \cong 0.33$; see Figure 11).

Figure 11

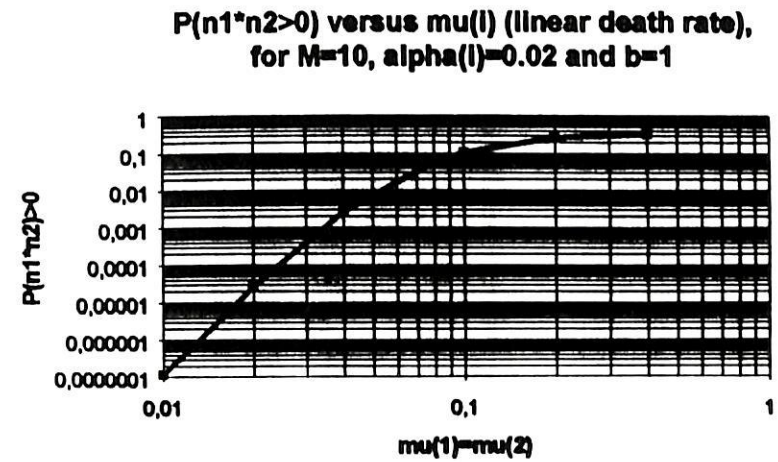

However, such large death rates would mean massive user dissatisfaction with a technology upon experiment; this would imply the need for a more sophisticated model, dividing the population between susceptibles and immunes as in epidemiological studies (Bailey, 1975). Since we are not concerned with this aspect, but only with the occasional regret of a technology by an adopter through accidental causes, it seems more appropriate to consider only small death rates in relation to $M^*\alpha$. Within this range, the qualitative results for a birth-and-death competitive diffusion process seem independent of the precise death rate definition. The main conclusions to draw from the above results are, one, that the limit distribution is centred on states where only one technology is present; and two, that initial advantages may be significant if both technologies are equally attractive.

Whether initial advantages are significant when diffusion parameters are not the same for both technologies is another question. To investigate it, we shall use the probability ratio PR = $P(n_1>n_2)/P(n_2>n_1)$. Figure 12 shows the behaviour of PR when α_2 varies, with constant death rates, for distinct values of b (initial advantage). Qualitatively similar results obtain with linear death rates (not shown).

Figure 12

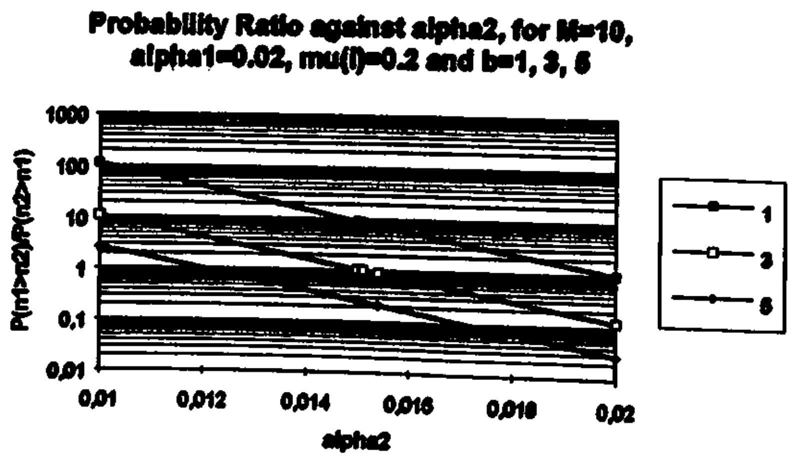

This figure allows two conclusions. First, PR is highly sensitive to differences in diffusion coefficients: halving α_2 increases PR a hundredfold, cæteris paribus. Second, initial advantages significantly change PR from the "no advantage" case by a roughly invariant amount (one order of magnitude from b=1 to b=3), over the range analysed for α_2. These gains are only marginally sensitive to m_i, decreasing slightly with larger death coefficients (not shown). A similar conclusion may be drawn keeping birth coefficients fixed and varying death coefficients (figure not shown). The main differences are that doubling μ_2 increases PR a thousandfold for our example, and larger birth coefficients slightly increase the sensitivity.

A final question remains: It may be argued that the extremely low values for $P(n_1n_2>0)$ reflect the linear birth rate used. Arguably, if there were "diseconomies of adoption", limit probabilities would be more evenly distributed, or possibly even centred upon equitable market shares. To test this hypothesis, we modified the functional form of birth rates to

$$(6)\ \lambda_i = \alpha_i \cdot n_i^{\theta} \cdot (M - n_1 - n_2),\ \theta \in [0,1],\ i = 1,2$$

Figure 13, below, shows the result for M = 10 and linear death rates. There is a discontinuity at q = 0, where $P(n_1n_2>0)$ @ 0.96; for all other values, $P(n_1n_2>0)$ only attains very low values. (The behaviour with constant death rates is similar)

Figure 13

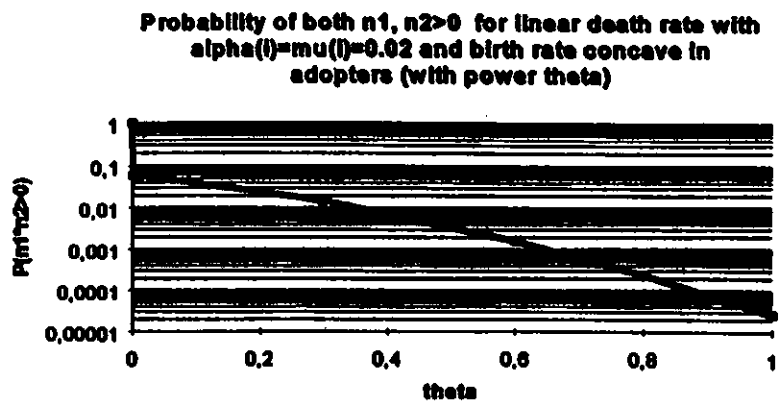

These conclusions need qualifiying, since they arise from limited experimentation on a small-scale case. However, they confirm the results obtained in the pure birth case as long as the death rate does not dominate the behaviour. That is to say, technology diffusion appears to be a priori uncertain and highly sensitive to historical events.

4. Implications and avenues for investigation

The model investigated suggests an intriguing possibility. That is to say, in market phenomena for which birth-and-death processes are an adequate model, apparent equilibria[3] may occur with only a tenuous relationship to the intrinsic advantages of competing alternatives. In other words, apparently stable competitive equilibria need not be efficient in such cases. For these classes of phenomena, relationships based on Pareto's theorems are not valid, not even in a probabilistic sense. Nothing (or very little) can be said regarding social optima in these cases by looking at market outcomes, with or without redistribution of utilities among agents. In effect, this is an almost perfect opposite of chaotic behaviour from deterministic mechanisms: instead, we are talking about unpredictable "equilibria" (that is, apparently orderly states) generated by random mechanisms. In such cases, public regulation would be legitimate in the search for a social optimum.

Given this possibility, two questions we must ask are, first, if birth-and-death processes are a plausible model for technology diffusion processes; second, if other economic phenomena are amenable to a similar description.

The answer to the first question is a qualified yes. Adoption and rejection in a market is a series of individual decisions, taking into account features of technologies as well as their established bases. Besides, given individual differences in tastes and expectations, modelling an individual decision as a stochastic phenomenon from the point of view of an external observer is a valid option, particularly if the number of individuals is large, loosely organized, without perfect information and mainly undifferentiated. In cases where individuals are differentiated

into a small number of coherent groups, evolutionary models may be a better description.

Accepting the description of technology diffusion as a stochastic process, birth-and-death processes are a good candidate to model it for their simplicity and power. As a matter of fact, the random walk employed by Arthur (1989) may be viewed as a sequence of competitive births only, with some special features: composition of birth rates between two types of individual decision-makers, absorption at the (X,0) and (0,Y) states (with X and Y > 0), no absorption at (0,0), unlimited market size. The last feature prevents it from being embedded into a (composite) pure birth process for finite markets, and in fact the key difference between Arthur's results and ours is that pure birth processes in a finite market reach equilibria covering the whole range of market partititions (and the same happens with the "transient equilibria" in birth-and-death processes). All the other features might conceivably be reproduced by a suitably adapted birth-and-death process. In fact, the introduction of renewal creates an unlimited market; it seems no coincidence that in this case our results are similar to Arthur's.

However, the case without renewal presents the most remarkable result, whether in the pure birth form or in the transient equilibria of the birth-and-death model. That is to say, observed equilibria span the whole range and their distribution is affected by initial advantages. This has far-reaching implications; thus, ex-post diffusion rates and niches tell little or nothing regarding ex-ante parameters.

Birth-and-death rates could be made explicitly dependent on perceived economics of competing technologies. However, whether those rates should model decisions based on supply and demand schedules is not crucial to the argument. What matters is that qualitative conclusions drawn from the birth-and-death model do not depend on the precise form of birth-and-death rates. They are in fact remarkably robust, as shown in the preceding sections. There is moreover some empirical justification for them, as Arthur (1989) argues[4]; and we agree with him that there seems to be a clear case for testing these conclusions against more systematic empirical evidence.

The second question requires more careful examination. In principle, one may model a market as an auction where successive bidders (buyers and sellers) raise or lower prices as successive units are sold. Nevertheless, it is not clear that a birth-and-death process could be profitably used to model this auction in most cases. An investigation of market phenomena for which this approach is adequate might be an interesting and potentially relevant endeavour.

5. Bibliography

ARTHUR, W.B. (1989), Competing Technologies, Increasing Returns, and Lock-In by Historical Events, The Economic Journal, 99 , pp 116-131

--, ERMOLIEV, Yu. and KANIOVSKI, Yu. (1987), Path-Dependent Processes and the Emergence of Macro-Structure; European Journal of Operational Research, 30, pp 294-303

AYRES, R. U., Technological Forecasting and Long Range Planning, McGraw-Hill, 1969, New York

BAILEY, NORMAN T. J., The Mathematical Theory of Infectious Diseases and its Applications, Charles Griffin & Company Ltd., 2nd Edition, 1975, London and High Wycombe

CAMERON, H.M. & METCALFE, J.S. (1987), On the Economics of Technological Substitution, Technological Forecasting and Social Change, 32, pp 147-162

ROSS, S. M., Stochastic Processes, John Wiley & Sons, 1983, New York

SILVERBERG, G.; DOSI, G.; ORSENIGO, L., "Innovation, diversity and diffusion: a self-organisation model"; The Economic Journal, 98, pp. 1032-1054, December 1988

Appendix

A.1 Distribution of final states in the competitive pure birth model

In order to determine the distribution of final states, we may omit the time dimension and limit ourselves to the embedded Markov chain of the process. That is, if we note the state (a, b) as meaning that a have adopted technology 1 and b have adopted technology 2, from this only two transitions are possible: to (a+1, b) with probability $\alpha_1 \cdot a/(\alpha_1 \cdot a + \alpha_2 \cdot b)$, or to (a, b+1) with probability $\alpha_2 \cdot b/(\alpha_1 \cdot a + \alpha_2 \cdot b)$. If we note $P^{M-a-b}_{(a,b)(n_1,M-n_1)}$ the probability of going from state (a, b) to state (n_1, M-n_1) in M-a-b steps, we have from the Chapman-Kolmogorov equations (noting b = a_2/a_1):

$$P^{M-a-b}_{(a,b)(n_1,M-n_1)} = \frac{\alpha_1 \cdot a}{\alpha_1 \cdot a + \alpha_2 \cdot b} P^{M-a-b-1}_{(a+1,b)(n_1,M-n_1)} + \frac{\alpha_2 \cdot b}{\alpha_1 \cdot a + \alpha_2 \cdot b} P^{M-a-b-1}_{(a,b+1)(n_1,M-n_1)} =$$

$$= \frac{a}{a+b\beta} \cdot P^{M-a-b-1}_{(a+1,b)(n_1,M-n_1)} + \frac{b\beta}{a+b\beta} \cdot P^{M-a-b-1}_{(a,b+1)(n_1,M-n_1)}$$

From this recurrence $P^{M-b-1}_{(1,b)(n_1,M-n_1)}$ may be computed for all n_1, given M and b.

Observe that these are not limiting probabilities, since the chain is not ergodic. The probability of reaching a given final state depends upon the initial state. This is in stark contrast to many equilibrium models in mathematical economics.

In the special case where $\alpha_1 = \alpha_2$, we obtain $P^{M-2}_{(1,1)(n_1,M-n_1)} = \frac{1}{M-1}, \forall n_1$. The simplest way to prove this is to note that in this case any sequence of transitions that leads from (1,1) to (n_1, M-n_1) has probability

$$\frac{1}{2} \cdot \frac{2}{3} \cdots \frac{n_1-1}{n_1} \cdot \frac{1}{n_1+1} \cdot \frac{2}{n_1+2} \cdots \frac{M-n_1-1}{M-1} = \frac{(n_1-1)! \cdot (M-n_1-1)!}{(M-1)!}$$

Since there are $C^{n_1-1}_{M-2} = \frac{(M-2)!}{(n_1-1)! \cdot (M-n_1-1)!}$ such sequences, the result follows Q.E.D.

A.2 Derivation of P2P1(M, b, b)

Starting from (1, b), the probability that all remaining M-b-1 potential adopters choose technology 1 is

$$\frac{\alpha_1}{\alpha_1 + b \cdot \alpha_2} \cdot \frac{2\alpha_1}{2\alpha_1 + b \cdot \alpha_2} \cdots \frac{k\alpha_1}{k\alpha_1 + b \cdot \alpha_2} \cdots \frac{(M-b-1)\alpha_1}{(M-b-1)\alpha_1 + b \cdot \alpha_2} =$$

$$= \frac{1}{1+b\beta} \cdot \frac{2}{2+b\beta} \cdots \frac{k}{k+b\beta} \cdots \frac{M-b-1}{M-b-1+b\beta}$$

Similarly, the probability that all (M-b-1) remaining potential adopters choose technology 2 is

$$\frac{b\alpha_2}{\alpha_1 + b\alpha_2} \cdot \frac{(b+1)\alpha_2}{\alpha_1 + (b+1)\alpha_2} \cdots \frac{(b+k-1)\alpha_2}{\alpha_1 + (b+k-1)\alpha_2} \cdots \frac{(M-2)\alpha_2}{\alpha_1 + (M-2)\alpha_2} =$$

$$= \frac{b\beta}{1+b\beta} \cdot \frac{(b+1)\beta}{1+(b+1)\beta} \cdots \frac{(b+k-1)\beta}{1+(b+k-1)\beta} \cdots \frac{(M-2)\beta}{1+(M-2)\beta}$$

Therefore,

$$P2P1(M,\beta,b)=b\beta\cdot\frac{(b+1)\beta(2+b\beta)}{2(1+(b+1)\beta)}\cdot\ldots\cdot\frac{(b+k-1)\beta(k+b\beta)}{k(1+(b+k-1)\beta)}\cdot\ldots\cdot\frac{(M-2)\beta(M-b-1+b\beta)}{(M-b-1)(1+(M-2)\beta)}=$$

$$=\prod_{k=1}^{M-b-1}\frac{\beta\left(1+\frac{b\beta}{k}\right)}{\beta+\frac{1}{b+k-1}}$$

Notas

[1] For a deterministic model, final shares are linked through the equation $x_1^{1/a_1} = \left[x_1(0)/x_2(0)\right]\cdot x_2^{1/a_2}$, as it may be easily seen by solving the equations.

[2] The probability law of the duration of these "pseudo-equilibria" is itself a non-trivial research object. It may however be noted that, when market is near saturation, birth rates are of the same order as death rates, i.e. times between transitions are longer than in the intermediate stages. Moreover, birth rates are proportional to market shares, which reinforces stability.

[3] Even though in the very long run all competitors are extinct in birth-and-death processes without renewal, actual realizations spend a significant amount of time in apparent equilibrium. The analysis of the probability laws of this 'equilibrium' might be rewarding, as noted above.

[4] The best-known instance cited by Arthur is the QWERTY keyboard.