Collective Learning, Innovation and Growth in a Boundedly Rational, Evolutionary World

Gerald Silverberg & Bart Verspagen **93-019**

September 1993

Mailing address:

Gerald Silverberg / Bart Verspagen
MERIT/Faculty of Economics
University of Limburg
P.O. Box 616
6200 MD Maastricht
The Netherlands

Phone:
(31) (0)43-883875

Collective Learning, Innovation and Growth in a Boundedly Rational, Evolutionary World

Gerald Silverberg and Bart Verspagen

We thank participants at seminars at the Universities of Alicante, Tilburg and Limburg (Maastricht) for helpful comments. The usual disclaimer applies. The research of Bart Verspagen has been made possible by a fellowship of the Royal Netherlands Academy of Arts and Sciences.

September 1993

Maastricht Economic Research Institute on Innovation and Technology
University of Limburg
P.O. Box 616
NL-6200 MD Maastricht
The Netherlands
Tel. +31 43 883869, Fax +31 43 216518
Email: Gerald.Silverberg@Merit.Rulimburg.NL and Bart.Verspagen@Merit.Rulimburg.NL

1. Introduction

The recognition that economic growth and market dynamics are inextricably intertwined with technical change is nothing new. In recent years, however, attention has focused on both a more precise representation of technology and on the economics of innovation itself, in particular, its economic incentives, returns, and determinants (for a recent survey see Silverberg and Soete 1993). Under the name endogenous growth theory, for example, (cf Romer 1986, Romer 1990, Lucas 1988, Aghion and Howitt 1992, Helpman 1992), mainstream economists have attempted to reconcile a general equilibrium approach to growth with some of the recognized features of technical change: increasing returns, increasing product variety, imperfect competition and 'creative destruction'. While recognizing that innovation is associated with many imponderables and uncertainties, these models have still hinged on the standard tools of market clearing, classical rationality, intertemporal optimization, rational (technological) expectations and the identification of steady-state equilibria.

It is doubtful, however, if the applicability of these tools has not been overreached in the interests of maintaining the general equilibrium framework. First, although many authors recognize the uncertainty surrounding the innovation process and thus represent it stochastically, there is still no reason to assume agents have common knowledge of the underlying stochastic process. Rather, these agents seem to be confronted with something much more akin to a multiagent, multiarm bandit problem, with uncertainty ever present with respect to the prospect of making a viable innovation, competitors' actions, and the ability to appropriate returns from the invested effort. The assumption of a rational expectations equilibrium based on the decision process of a representative agent seems to be even more inappropriate here than in more traditional areas of economic analysis, where it has also been intensely debated.

Moreover, the representation of technology has often been oversimplified to correspond to known and tractable special cases. Thus, in the model most closely related to the structure of our own, Aghion and Howitt (1992) assume that at any time only one technology prevails in the economy. When a new innovation is made, it instantly and costlessly (except for the sunk cost of the R&D which went to invent it) replaces its predecessor. The monopoly returns associated with it, for as long as it prevails, are completely recouped by the single innovator, with no imitation or spillovers. Thus the appropriate stochastic decision problem to determine the level of R&D spending, they argue, can be represented as a patent race with a commonly known probability of making an innovation as a function of R&D effort.

As any student of the history of technology knows, however, this is a far cry from anything that has ever prevailed economy-wide. And even in some particular industries where a monopoly of a truly key innovation has been defended for as long was possible (one thinks of the Boulton and Watt patent on the condensing steam engine), earlier technologies (such as the Newcomen engine and water power) retained a significant place in production for a considerable time (partly because of the monopoly, no doubt).

These cursory reflections suggest a number of requirements for an "endogenous" growth model regarding both a model's representation of technology and of the decision problem confronting the innovator:

1) At any given time a number of technologies will be concurrently in use, particularly if these technologies are capital embodied (the vintage effect);
2) Even at the investment "frontier", different technologies may be adopted simultaneously;

3) The aggregate rate of technical change will be a function of the rate of diffusion of new technologies and not of the rate of instantaneous innovation;
4) Even if one may accept that the return to innovative effort is representable as a draw from a stationary random process, the parameters of this process are not *a prior* known to agents, and their subjective priors may differ widely. Moreover, even if they are Bayesians, they may not live long enough to draw more precise conclusions. Their problem is not dissimilar to that of the multiarmed bandit;
5) Technological knowledge, like information in general, can have a public and codifiable, a private, and a tacit character. It can only be imperfectly protected as private property (the mere knowledge that something can be done, which a patent discloses, can already be very useful information). This state of affairs will certainly influence innovative activity, but in ways that are difficult to anticipate.[1]

In the model we present in the following, firms must determine what proportion of their profits to invest in R&D as an operating routine. In contrast to the more neoclassical approaches, we consider this to be an archetypical example of a problem of *bounded rationality*, where agents can have only vague ideas about the relationship between their actions and outcomes. To provide some anecdotal evidence, we recall an interview with the director of R&D of the Japanese firm Canon published in *The Financial Times* a number of years ago. The director reported that the firm had some time before raised its R&D/turnover ratio from 11% to 11½%. This appeared to have been beneficial to firm, so that the directors were now debating whether to cautiously raise it even further. There in fact did not seem to be any way to determine where an upper limit might lie, and what the optimum policy might be, short of actually trying it out, but the firm was set to continue in this direction.

This sort of reasoning appears to be much more typical of what actually occurs in business firms than the sophisticated dynamic optimization under uncertainty so beloved by mathematical economists. The standard 'as if' argument offered in defence of optimization (Friedman 1953), based on market selection of the optimum strategy, is something that should be explicitly modeled, not assumed. In recent years this seems to have become more widely accepted (see for instance Selten 1989 for a provocative contrast of theoretical perspectives). Yet economists have been reluctant to embrace bounded rationality as an alternative program, perhaps less because of a belief in the realism of traditional methods than in a fear of the Pandora's box of unfettered behavioral possibilities and *'ad hocery'* that might thereby be opened.

2. Selection, Learning and the Artificial Worlds Modeling Philosophy

Nevertheless, the selectionist viewpoint has been reintroduced formally into economics from biology in the form of evolutionary game theory (cf. Friedman 1991 for a survey). The static notion of *evolutionary stable strategy* is the biologist's refinement of a Nash equilibrium,

[1]In classical articles Nelson 1959 and Arrow 1962 underscored the disparity which may exist between the social and the private rates of return and incentives to innovation. Recent literature has uncovered the possibility of both insufficient as well as socially excessive, redundant R&D, depending on the precise assumptions made. Nelson 1990 brings the empirical literature to bear on these issues, and show in particular that not insubstantial own R&D efforts are necessary even to imitate.

while dynamic versions have usually been based on replicator dynamics (cf. Hofbauer and Sigmund 1988 for the biological applications). The key biological stability concept is *invadability*, the inability of a small number of mutants to invade an equilibrium population. As in standard game theory, the interesting results concern possible existence of multiple equilibria, but in contrast to game theory the dynamic theory provides a means of equilibrium selection based on initial conditions. Most of this work has been in a deterministic framework in economics, although this has not been the case in biology. But evolutionary game theory reflects just one half of the evolutionary process, namely selection, and has focused on the equilibria as once-and-for-all asymptotic states.

Recently, attention has shifted to the role of mutation and stochasticity in further refining the notion of stability in evolutionary games (Foster and Young 1990, Young 1993, Kandori, Mailath, and Rob 1993, Binmore and Samuelson 1993). In contrast to conventional evolutionary game theory, agents or species are allowed to mutate continually over time before an equilibrium is attained. In general, the equilibria that remain (usually in the limit as the mutation probability goes to zero) will be a significant further refinement of the limiting states. This framework has been proposed as a model of learning by ongoing trial and error with selection. Once again, though, the focus has been on asymptotic states and the refinement of equilibrium concepts.

Approaching learning from another extreme is work based on artificial intelligence and computer science such as neural nets, genetic algorithms, and classifier systems. While these approaches are also based on complex interdependence of interacting subunits as in game theory, selection, and stochastic perturbations, the complexity of the problems workers have addressed has made analytical results difficult to come by. Whereas in the beginning, most work focused on computing standard optimization problems and thus converging to a hopefully unique point, more recent efforts have been directed at using these tools open-endedly to simulate the self-organization and evolution of life-like systems, whether they be abstract ecosystems, microorganisms, individual behavior or human economies (cf. the work in Langton 1989, Langton, Taylor, Farmer and Rasmussen 1992, Arthur 1991, and Lane 1993).

For systems exceeding a certain complexity in their organization, if not in the nature of their constituent parts and their interactions, the existence and uniqueness of an asymptotic steady state may be less interesting than two heuristic phenomena that have often been commented on without precise definition. The first is *emergent properties*, i.e., complex but identifiable patterns of behavior that emerge spontaneously at some point in the history of the system and are not in any obvious way inherent in the constitution of its parts. The second is *punctuated equilibria*, periods of quasi-stable behavior separated by usually very rapid periods of transition and disorder. Evolutionary game theory, by focusing on very low-dimensional game structures (usually only two, in fact) may well provide an explanation of the emergence and stability of a single such behavior, but has deliberately excluded the complexity apparently necessary to generate an ongoing sequence of such states. There is obviously a tradeoff here between analytic tractability and behavioral richness as well as a philosophical difference of opinion on how science should proceed. Lane (1993) argues that the artificial world framework differs from conventional approaches in one or more of the following ways:

1) Transients are at least as important as steady states;
2) The concept of stability must be relativized to some notion of metastability, i.e., one which may contain the seeds of its own destruction;

3) New statistical methods will have to be developed to enable metastable states and emergent properties to be identified and characterized.
4) Computational methods will attain a scientific status equal to that reserved until now for analytical ones.

Most of the work in artificial worlds has been based on discrete genetic codings of strategies such as are called for by genetic algorithms or classifier systems. One may ask whether this perspective, borrowed from biological genetics and eminently suitable for the digital computer, should not be made more congruent with the way agents may be interpreted to formulate and modify their strategies in reality. Thus, although a genetic algorithm may be used to solve an optimization problem in a continuous space by representing real values in binary form, in some cases it may be more natural to represent for example mutation as a local operation rather than the discrete and possibly very large jumps implicated by a binary coding in a genetic algorithm. In our case at hand—the determination of firms' R&D investment—it may make more sense to mutate locally around existing values than to allow for jumps across the entire parameter space. This "realistic" approach to evolutionary modeling based on stylized behavior is characteristic of the work initiated by Nelson and Winter (1982) and in the computer science realm to the real-valued evolutionary algorithms of Schwefel (1981). Thus mutation or trial and error will be represented in our model by a draw from a normal distribution centered around the current value of the R&D to profit ratio. Imitation, in contrast, does permit large discrete jumps in parameter space but will be modified here to reflect a satisficing principle—only firms with low profits will go out and imitate.

2. The Model

The model consists of three basic blocks. The first block describes how the artificial economy evolves with a given set of technologies and firms. This block consists of equations for the rate of capital accumulation, the diffusion of new technologies in the total capital stock of the firms, and the real wage rate. The second block describes a set of rules that is used to introduce new technologies and firms into the economy. This block takes the innovative behavior of firms (to be explained below) as given, and then describes the probability that individual firms will make an innovation, as well as how this innovation is introduced. The third block describes how innovative behavior changes under the influence of the evolution of the economy and firm learning. This block, in other words, describes a feedback from performance to innovative behavior and thus a form of collective learning. The parameters of the model and the values used in the simulations are summarized in the Appendix.

a. The evolution of the artificial economy with a given set of firms and technologies

The basic framework of the model is taken from Silverberg and Lehnert (1993), which in turn draws on Silverberg (1984) and Goodwin (1967). Let hats above variables denote proportional growth rates, w be the (real) wage rate, v the employment rate (persons employed as a fraction of the labor force), and m and n parameters (both positive). Then the growth rate of real wages is described by the following linear Phillips curve:

$$\hat{w} = -m + nv. \tag{1}$$

It is assumed that there are q firms in the economy, while each of these firms has a number p_q of different types of capital goods that it utilizes to produce a homogeneous product. New capital arises from the accumulation of profits, a process described by the following equation:

$$\hat{k}_{ij} = (1-\gamma_i)r_{ij} + \alpha(r_{ij} - r_i). \tag{2}$$

The capital stock is denoted by k, and r stands for the profit rate. The subscript i ($1..q$) denotes a firm, and j ($1..p_q$) the type of capital (absence of any these indices indicates an aggregation over this particular dimension). Equation (2) assumes that the principal source for type ij-capital accumulation is profits generated by ij-capital. This is modelled by the first term on the rhs of (2), i.e., $(1-\gamma)_i r_{ij}$. A firm specific portion of profits (denoted by γ) is used for the development of knowledge (R&D). R&D is restricted to cases where there are positive firm-profits (i.e., when $r_i<0$, γ_i is set to zero).

However, to a certain extent, profits are redistributed such that more profitable types of capital accumulate even faster. The mechanism used to model this was first proposed by Soete and Turner (1984), and is represented by the second term on the rhs of eq. (2). By changing the value of α, redistribution of profits takes place faster (larger α) or slower (smaller α).

It is assumed that each type of capital is characterized by fixed technical coefficients, c and a (for capital coefficient and labor productivity, respectively). The capital coefficient is assumed to be fixed throughout the economy (and time), while labor productivity is assumed to change under the influence of technical progress. The profit rate of ij-capital is then given by $(1-w/a_{ij})/c$.

The principal variable used to describe firm dynamics is the share of the labor force employed on each capital stock. Production is assumed to be always equal to production capacity (the influence of effective demand is absent), so that the amount of labor employed by each capital stock is equal to $k_{ij}/(a_{ij}c)$. Dividing this by the labor force (assumed to grow at a fixed rate β) gives the share of labor employed, l_{ij} (called employment share hereafter). The expression for the growth rate of this variable is

$$\hat{l}_{ij} = \hat{k}_{ij} - \beta. \tag{3}$$

R&D also has an employment effect. We assume that the ratio between R&D expenditures and R&D labor input is equal to a fraction δ of economy-wide aggregate labor productivity. The employment rate v_q resulting from production is then found by summing l over i and j. Under these assumptions, it can then be shown that the overall employment rate v is equal to $(1+\delta\gamma(1-w/a))v_q$.

Equations (1)—(3) constitute a system of differential equations that, for given initial conditions, technologies and firms, describe how the economy evolves over time. However, for the given set of technologies, long-run (per capita) growth is not possible beyond a situation where all firms only apply the newest technology. Therefore, the next section outlines how new technologies can enter the system, and thereby open up the possibility for long-run growth.

b. The introduction of new technologies and firms into the economy

The main idea that distinguishes this model from the ones in Silverberg (1984) and Silverberg and Lehnert (1993), is the way in which innovation is "endogenized". It is assumed that in each time period, firms devote resources (R&D) to the systematic search for new production possibilities (i.e., new types of capital). The outcome of this search process is assumed to be stochastic.

Each time an innovation occurs, the firm creates a new type of capital. The labor productivity of this type of capital is given by the following process (τ is the fixed proportional increase in labor productivity between innovations, $a_{i,t}^*$ the firm-specific best practice labor productivity):

$$a_{i,t}^* = (1+\tau)a_{i,t-1}^*. \tag{4}$$

The new type of capital is seeded with a small employment share (say 0.0001). In order to keep the total employment rate constant, this seed value is (proportionally) removed from the other types of capital of the innovating firm. The number of technologies employed by any given firm will then not be fixed but may vary in time.

In this setup, if real wages rise over time, every technology will generate negative profits at some stage (because of its fixed labor productivity). It is assumed that these losses are financed by an equivalent decrease of the capital stock. In other words, losses imply that capital will be scrapped, and the scrapped capital can be transformed to cover for the losses. Note that for individual capital stocks, the point at which scrapping occurs lies *prior* to the point where profits are negative, due to the α-related diffusion term in equation (2). When a technology employs a labor share smaller than a specified (very small) value, it is scrapped completely.

A firm's R&D activities as well as possibly those of its rivals enter an innovation potential function T_i. This in turn determines the firm's probability of making an innovation according to a Poisson process with arrival rate ρ_i. The simplest relation is simply linear:

$$\rho_i = A\,T_i + \rho_{\min}, \tag{5}$$

where ρ_{min} is the (small) autonomous probability of making a fortuitous innovation without doing formal R&D, and A is the innovation function slope. One can also posit a nonlinear relationship with both increasing and decreasing returns to R&D, such as a logistic:

$$\rho_i = \frac{\rho_{\min}\rho_{\max}}{\rho_{\min}+(\rho_{\max}-\rho_{\min})\,e^{-AT_i}}. \tag{6}$$

This logistic function has intercept ρ_{min} and (asymptotic) saturation level ρ_{max}. In this case, the parameter A determines the speed at which the saturation level is approached.

T_i, the innovation potential, is determined both by the firm's own R&D level (h, to be defined below) and its ability to profit from other firms' R&D (technological spillovers):

$$T_i = h_i + \phi_1 h + \phi_2 h h_i. \tag{7}$$

These spillovers can take two forms. First, there is a term related to the economy-wide value of h (written without subscript). The economy-wide R&D level h is defined to be the market share weighted average of firm-specific R&D levels. Second, there is a term related to the product of the economy-wide and firm-specific values of h. This latter term takes into account the argument that in order to assimilate spillovers, a firm has to have some technology-generating proficiency itself (see Cohen and Levinthal 1989, Nelson 1990). The parameters ϕ_1 and ϕ_2 determine the importance of each spillover mode.

The firm-specific R&D level h_i is defined to be the ratio of a moving average of firm R&D investment to its total physical capital stock. A ratio is used to normalize for firm size, since otherwise such a strong positive feedback between R&D and firm growth exists that monopoly becomes inevitable. While *a priori* it is by no means clear why the size of individual R&D effort should not directly relate to innovative success, a pure scale effect must be ruled out by the continuing existence of competition and the ability of small countries to remain or even advance in the technology race. The exponential moving average $<RD>$ on R&D for a lag of L (or a depreciation rate of $1/L$) is given by the following differential equation:

$$\frac{d}{dt}<RD> = (\gamma r_i k_i - <RD>)/L. \tag{8}$$

Hence the firm-specific R&D level is

$$h_i = <RD>/k_i. \tag{9}$$

An innovation can be defined in a narrow or a wide sense. In the wide sense, the adoption of any technology not yet employed by a firm (or a country) is an innovation to that unit. In the narrow sense, only technologies that have never been employed before anywhere are considered innovations at their time of introduction. If firms innovate according to the above Poisson arrival rates in the narrow sense, however, a very considerable intertemporal externality is created, because firms' innovations always build on each other. Thus there can be no duplication of effort and, as long as firms maintain a minimal level of R&D, no cumulative falling behind. On the other hand, once an innovation has been introduced somewhere into the economy, it should be progressively easier for other firms to imitate or duplicate it; it should not be necessary to reinvent the wheel. We capture this by introducing a catch-up effect. Let the labor productivity of the economy-wide best practice technology be a^*, and the best practice technology of firm i a^*_i. Then firm i's innovation potential T_i is augmented by a measure of its distance from the best practice frontier:

$$T_i' = T_i(1 + \kappa \ln(a^*/a_i^*)). \tag{10}$$

Thus adopting an old innovation is facilitated for backward firms, but they are still required to invest in their technological capacity to reap these catchup benefits. Here, however, R&D efforts should be interpreted in the larger sense of technological training and licensing, reverse engineering, or even industrial espionage (all costly activities, if not as costly as doing state-of-the-art R&D).

We have also experimented with innovations in the narrow sense, but the results on strategic selection are rather ambiguous. This is not surprising, since the import of the intertemporal externality is indeed quite large. We consider the Ansatz in eq. 10 therefore to be a justifiable first formulation, since technology adoption decisions are never passive, but rather require technological efforts of the adopting firm. However, it does place too much of the burden of catching up onto R&D, which is probably misplaced and should eventually be replaced by a more appropriate formulation.

In the artificial economy modelled here, entry of a new firm occurs only as a result of exit of an incumbent firm. Exit occurs whenever a firm's employment share (excluding its R&D employment) falls below a fixed level E. While exit of incumbent firms is completely endogenous, entry only occurs in case of exit, so that the total number of firms is constant. Naturally, this feature of the model is not very realistic, as in reality entry may be independent of exit and the total population of firms may vary. However, it is not the aim of this model to describe the phenomena of entry and exit as such. Instead, the main function of entry and exit is to maintain potential variety in the population of firms while providing for firm elimination.

Whenever entry occurs, the entrant is assigned a single technology with an amount of capital corresponding to an employment share of $2E$ (the remaining employment is proportionally removed from other firms so that total employment remains constant). The labor productivity of this technology is drawn uniformly from the range $[(1-b)A,(1+b)A]$, where A is the unweighted mean value of labor productivity of all the firms in the economy, and b is a parameter. The values for h and γ are (uniformly) drawn from the range existing in the economy at the time of entry.

c. Firm-strategies for innovation

In Sections (a) and (b) we have outlined the system whereby innovating firms generate technical change and undergo selection in a closed economy model as a function of their R&D strategy parameter γ. Learning now enters the picture in the form of two "genetic" operators: mutation and imitation. These are summarized below:

$$\begin{aligned} &Prob\ \Pi:\ \gamma_{it}=Min(1,Max(\gamma_{it-1}+\varepsilon,0)),\quad \varepsilon\sim N(0,s_i),\\ &Prob\ \Pi_i^c:\ \gamma_{it}=\gamma_{jt-1},\ j(\neq i)\in[1..q]_i,\\ &Prob\ 1-\Pi_i-\Pi_i^c:\ \gamma_{it}=\gamma_{it-1}. \end{aligned} \tag{11}$$

With probability Π each "year", which is set exogenously and equal for all firms, a firm will draw from a normal distribution and alter its strategy within the admissible range [0,1] (mutation). With variable probability Π^c_i the firm simply imitates a strategy of another firm. The imitation probability is partly endogenous to reflect satisficing behavior. Only firms with unsatisfactory rates of profit with respect to economy leaders will choose or be forced (for example by their stockholders or by hostile takeovers) to adopt the strategy of a competitor:

$$\Pi_i^c = \mu\left(1 - \frac{y_i - y_{\min}}{y_{\max} - y_{\min}}\right) \qquad (12)$$

y_i is the firm's rate of expansion of physical capital (defined as $min(r_i,(1-\gamma_i)r_i)$), $y_{\max}$ and y_{min} are the maximum and minimum values of y in the sample, and μ is the (exogenously determined) maximum imitation probability, which is associated with the least profitable firm in the sample. Thus, the more profitable a firm is, the less the chance it will change its strategy by imitating another firm. The most profitable firm has an imitation probability equal to zero. Once a firm has decided to imitate, it selects a firm to imitate randomly from the industry with weight equal to the target firm's market share in output.

3. Identification of Steady States by Random Generation

The model has been implemented to run on both MS-DOS and UNIX computers. To make the solution as time-step invariant as possible, the selection mechanism, which is basically a system of differential equations, is solved using a fixed-step, fourth order Runge-Kutta algorithm. The innovation decisions are executed during each computational step (using Poisson arrival rates scaled by the computation time step), and when an innovation is made, the corresponding changes in initial conditions, number of equations, and coefficients are made for the next step. Mutation and imitation are only performed at fixed intervals of one "year", which may be many times the step employed in the Runge-Kutta algorithm.

We have concentrated up to now on investigating the linear innovation function case (eq. 5) in order to demonstrate as unequivocally as possible the existence in some sense of an "evolutionary attractor" to the dynamics. To zero in on this attractor as quickly as possible, we have initialized the system in a "grapeshot" mode we term random generation in which the initial γ's are drawn from a uniform distribution over [0,1]. As Figure 1 demonstrates, the system does converge (up to stochastic fluctuations) to a single value for γ consistently for the five runs shown, each generated with a different random seed. In Figure 2 we have plotted the means of the market-share-weighted γ's for the last 4000 years of 5000-year runs (in order to allow transients to die out) against stepped values of the innovation function slope A. Again, five runs have been performed per value to allow for stochastic variance. For low values of A this is indeed essential, since outliers demonstrate that a high growth/R&D regime need not always come about, or only after consider delay. At higher values of A this is no longer a problem, and convergence is apparent. Even including the outliers, we obtain a significant relationship, with a negative coefficient, between equilibrium R&D and innovation slope:

$$\begin{aligned} RD = \; & \underset{(0.0032)}{-0.0066} \; A + \underset{(0.073)}{0.544} \\ R^2 = \; & 0.076, \end{aligned}$$

where the standard errors are given in parentheses below the values. This is significant at the 5% level. Omitting the outliers, of course, increases the significance and R^2 considerably. The rate of technical change demonstrates a highly significant positive relationship to the innovation slope (Figure 3). The regression results are

$$\begin{array}{rcl} TC = & \underset{(0.00035)}{0.00459}\ A - & \underset{(0.0082)}{0.00086} \\ & R^2 = 0.76 & \end{array}$$

Evidently, with increasing A the "optimal" share of R&D in investment can decline in order to provide the additional investment resources necessary to finance the associated higher rates of technical change and capital turnover.

Even within this tight convergence, however, a relationship between technical change and concentration can be detected. Figure 4 plots time averages over the last 4000 years of 5000 year runs of the economy-wide rate of technical change vs. the Herfindahl concentration index[2], for 20 runs all at an innovation slope of 10. The highly significant inverse relationship is characterized by the following regression:

$$\begin{array}{rcl} TC = & \underset{(0.0093)}{-0.0999}\ CON + & \underset{(0.003)}{0.063} \\ & R^2 = 0.86 & \end{array}$$

A similar pattern is observed by plotting the mean γ against concentration (Figure 5). The regression here is

$$\begin{array}{rcl} RD = & \underset{(0.0226)}{-0.3027}\ CON + & \underset{(0.007)}{0.544} \\ & R^2 = 0.91 & \end{array}$$

again, highly significant. The mechanism underlying this relationship will be discussed in the next section, where we shall also see that the concentration index is indicative of a change of growth regime.

4. Spontaneous Generation: The Takeoff to a High Growth Regime

Instead of searching as rapidly as possible for evolutionary steady states, we can also set up runs to recreate the emergence of R&D activity from "medieval" initial conditions, i.e., with all γ's and all h's set to zero. The rate at which aggregate γ can grow is bounded by the mutation rate and the mean size of the mutation step. This development, should it take place, need not be monotonic, however. Periods of stagnation and even regression are quite possible.

Figure 6 shows the market-share-weighted mean and variance of γ for a run with $A = 10$, over 8000 years. After a period of stagnation near zero, the series begins to rise, oscillating

[2]This is defined as $H = \Sigma f_i^2$, where f_i is the market share of the ith firm. It ranges from $1/n$, for n equally sized firms, to 1, for complete monopoly.

for some 1000 years just below 0.1, until it begins a long, almost linear march up to its equilibrium value of around 0.5. The heavy line, tracing the variance of the population, shows that short periods of high variance are interspersed between relatively long periods of rather large market uniformity.

Figure 7 plots the rate of aggregate technical change (thin line) and the concentration index (heavy line) for the same run. Here the relationship between endogenous market structure and technical change hinted at in the random generation analysis becomes evident in an extreme form. In the initial stagnation phase as well as long into the climb, the rate of technical change remains low, with marked long-wave fluctuation (cf. Silverberg and Lehnert 1993). Most remarkable, however, is the dynamics of market structure, which we have termed the "guild pattern". Most of the time a single firm dominates, relegating the rest of the economy to a competitive fringe subject to entry and exit. At nearly periodic intervals of up to several hundred years, a small firm encounters a run of innovative success and rather rapidly displaces the incumbent. Then the market reverts again to a monopolistic pattern with low rates of technical change until the next challenger comes along. This pattern eventually breaks down in stages, however, as aggregate γ passes through thresholds. First an intermediate period ensues, in this run between about year 4500 and 5500, with moderate levels of concentration around 0.3 and a plateau of technical change oscillating around 3%. Finally, as the economy attains "equilibrium" a very sudden transition occurs, concentration declines and remains almost stable near the minimum value of 0.1 while technical change shoots up to a new plateau in the range 4-7%. This pattern is found, independent of A, for all runs that eventually attain respectable rates of technical change. Over the 8000 years we have followed the spontaneous generation runs there are cases, particularly for low values of A, in which no takeoff seems to occur. One may conjecture, though, that in the limit a takeoff is inevitable, with probability one, although one may have to wait eons for it to happen.

This is borne out in Figure 8, where the time averages of mean γ over the last 6000 years of 8000 year runs are plotted against innovation slope. The dispersion is quite large, with many runs remaining at low values. The dashed line superimposed on the plot is the regression line from the random generation runs (Figure 2) for "equilibrium" γ. The solid line is the regression for these data points. Apparently, with increasing A the probability of a takeoff increases and their average time at incidence declines, so that the time averages go up. But they are constrained from above by the "equilibrium" values represented by the dashed line, which is a *declining* function of A. The regression results here are

$$\begin{array}{lcccc} RD = & \underset{(0.0039)}{0.0074} & A & + & \underset{(0.0912)}{0.2295} \end{array}$$

$$R^2 = 0.06$$

still significant at the 10% level.

The relationship between the rate of technical change, again averaged over the last 6000 years of the run, and innovation slope is shown in Figure 9. This is markedly positive, again a result of the earlier onset of takeoff with increasing slope. Since these are time averages over disparate regimes, it is not surprising that they generally fall short of the values attained in the random generation runs (Figure 3).

By pooling the data over all slope values we again recover a practically linear relationship between the rate of technical change and the concentration index (Figure 10). Numbers over the data points indicate the corresponding value of A. There is a clustering of runs in the lower right corner, associated with late or no takeoff, and a smaller cluster toward the upper left corner, also associated with higher values of A, deriving from the genuine takeoff histories. The regression results are

$$TC = \underset{(0.0042)}{-0.1084} \; CON + \underset{(0.0044)}{0.0719}$$

$$R^2 = 0.93$$

What is the mechanism underlying this strong relationship between market concentration and growth? Recall that in the existing theoretical and empirical work on market structure and innovation (for an overview, see Scherer and Ross 1990 and Kamien and Schwartz 1982), two different conclusions dominate. One strand of the literature (starting from the so-called Schumpeterian hypotheses) argues that monopoly power is conducive to innovative efforts (R&D). Another part of the literature maintains that there is some optimal combination of competition and market power which maximizes innovative effort (thus presupposing an inverted U-curve relationship).

The results of the present model point in a different direction: a highly competitive regime and high R&D activity emerge jointly. The causal relationship between the two is far from evident, however. The interpretation we offer for the "guild pattern" is that in situations in which the overall rate of innovation is low (i.e., the arrival rate of new innovations is low), a firm which has a run of statistical luck immediately acquires a large dynamic advantage over its competitors. This leads to the monopoly. However, because the monopolist does not do a lot of R&D, a new entrant eventually also experiences a run and occasionally takes over. When firms, in the course of history, discover that higher R&D activity may lead to better performance, runs no longer produce large permanent advantages. Hence, a more competitive market structure arises.

5. R&D Spillovers and Evolutionary Steady States

What now happens if the spillover terms are turned on? We have examined this question for an innovation slope A of 10 and random generation (i.e., initial γ's drawn uniformly from [0,1]) as a function of ϕ_1 and ϕ_2 (eq. 7) separately. Figure 11 shows the time averaged mean R&D/profit ratios (last 4000 years of 5000 year runs) for five different values of ϕ_1, with five runs per value. Not unexpectedly, the presence of spillovers from the general level of R&D to individual firm innovativeness reduces the incentive to perform R&D in evolutionary "equilibrium". With increasing values of the spillover term the dispersion of the results also increases, with outliers indicating a number of runs that never or only tardily take off into the high growth regime. A regression on these data yields the following results:

$$RD = \underset{(0.0541)}{-0.2922} \; \phi_1 + \underset{(0.0506)}{0.5514}$$

$$R^2 = 0.51$$

The same phenomenon is observed for the realized rates of technical change (Figure 12). A regression yields

$$TC = \underset{(0.0083)}{-0.0211}\ \phi_1 + \underset{(0.0077)}{0.0375}$$

$$R^2 = 0.19$$

The situation is quite similar for type 2 spillovers, as shown in Figure 13. Evidently the dispersion of the results increases even more markedly, with a tendency for runs at the same value of ϕ_2 to bifurcate between a high and a low growth regime. The regression results are correspondingly weaker:

$$RD = \underset{(0.055)}{-0.101}\ \phi_2 + \underset{(0.0605)}{0.5446}$$

$$R^2 = 0.093$$

A similar pattern holds for the rate of technical change (Figure 14). The regression results are

$$TC = \underset{(0.0084)}{-0.0158}\ \phi_2 + \underset{(0.0099)}{0.0355}$$

$$R^2 = 0.097$$

Thus the net result of the addition of spillover terms to the innovation function is to depress R&D and technical change somewhat in the majority of cases, and raise the probability of retarding or precluding a takeoff into a high growth regime in others. This breakup of the histories into two structures can be clearly seen once again by plotting the rate of technical change against the concentration index (both time averaged) as in Figure 15 for the pooled date over the range of ϕ_2. The runs tend to cluster at either end of the regression line, with a few intermediate cases spread out in between.

6. Conclusions and Directions for Further Research

In this paper we have developed an evolutionary model describing the relation between endogenous technological change and economic growth along the lines of an "artificial world" modeling philosophy. By this we mean that the economy is disaggregated into diverse individual behavioral subunits (instead of the representative agent so prevalent in most macroeconomic modeling) connected by nontrivial nonlinear dynamic interactions based on plausible notions of disequilibrium competition and investment. Rather than search for a strategic equilibrium based on a concept of rationality, we have assumed that these agents use boundedly rational behavioral procedures. In the present case this is an extremely simple rule for R&D investment as a share of profits (or gross investment), which is parameterized by a single real number between 0 and 1. Learning is modelled by allowing for mutation and imitation rules operating on the agents' strategy parameters. An element of behavioral realism is injected into the model by insisting that mutations are local in the strategic "phenotype" space, and that imitation is only prompted by less than satisfactory performance.

Using a linear innovation function we were able to show with the random generation initialization that a unique evolutionary "equilibrium" exists that is attractive in the behavioral space and tracks variations in the innovation opportunity and spillover parameters in plausible ways. Thus the model does establish a case for endogenous growth in the sense of demonstrating that economic competition, even with very relaxed assumptions about individual goal-seeking behavior and profit maximization, leads to an approximately steady-state growth path with a positive rate of technical change and R&D investment.

However, the spontaneous generation experiments underline the fact that the mere existence of such a steady state does not mean that history does not matter. Quite the contrary. A society starting with no or low rates of R&D will stagnate for long periods in a sporadic low-growth trap and a "guild" market structure with an extreme degree of concentration, but with periodic upheavals or "palace revolts" of market leadership on a time scale of centuries. Eventually such an economy will "bootstrap" itself to higher rates of R&D and technical change, while preserving the "guild" market pattern (Mercantilism?). At some point a threshold may be passed which triggers a rapid transition to a first "industrial revolution" with moderate rates of concentration and a respectable rate of technical change. This may give way to yet another, very rapid, phase transition—a second "industrial revolution"—characterized by substantial levels of R&D[3] and a shifting market structure with low levels of concentration.

While we do not wish to overburden such a simple model with historical interpretations, the point must still be made that it would be unfortunate to restrict the concept of endogenous growth to steady-state growth paths with no real structural development, social learning, and historical contingency. For this reason an evolutionary approach appears to offer an attractive alternative explanation of how an economy can "bootstrap" itself in historical time through a succession of growth phases and market structures.[4]

A number of extreme simplifications have been made in the interests of a clearer preliminary analysis. First, the assumption of a linear innovation function not surprisingly leads to the uniqueness of the evolutionary attractor. A nonlinear function such as a logistic (eq. 6) would possibly result in distinct competing attractors and hence path dependence even in long-term growth trajectories. This is a question we intend to investigate further in the near future. Second, we have represented technologies as points on a linear directed graph. The succession is clear and the proportional improvements are fixed. One could also assume more complex technology "spaces" such as the branching and reticular digraphs analyzed by Vega-Redondo (1993), where technological *trajectories* and not just individual technologies compete with each other. Finally, from the point of view of obtaining analytic results the model may not yet be simple enough. It may prove worthwhile to make judicious simplifications in the

[3]The rates observed—on the order of 50%—are not all that implausible if interpreted as the share of R&D expenditures in gross investment. The most progressive firms in technologically dynamic industries in the USA, Japan and Europe have already attained comparable rates.

[4]The concepts of growth or development stages, takeoffs and changes of regime were prominent in a classical line of thought associated with Marx, Schumpeter and later Rostow. The criticism that these theories smacked of rigid mechanistic determinism is neatly sidestepped by the artificial worlds methodology, which demonstrates that reversions, variable delays, and path dependence resulting from the underlying stochasticity cannot be excluded. Needless to say, such a broad and "nonlinear" perspective on economic growth has mostly fallen by the wayside in the postwar literature on growth and development.

dynamical structure in order to derive asymptotic results along the lines of the new stochastic evolutionary game theory, while maintaining at least some of the essential features of the present framework.

Appendix

A summary of the parameters and the values employed in the runs analyzed in the paper is presented below.

$q = 10$	number of firms
$m = 0.9$	parameters of the Phillips curve eq. 1
$n = 1$	
$\alpha = 1$	Soete-Turner coefficient eq. 2
γ (endogenous)	R&D/profit ratio in eq. 2
$c = 3$	capital-output ratio
$\beta = 0.01$	rate of growth of labor force eq. 3
$\delta = 1$	ratio of productivity in goods and R&D sectors
$\tau = 0.06$	proportional jump in labor productivity eq. 4
A (variable)	innovation slope in eq. 5
$\rho_{min} = 0.01$	autonomous rate of innovation eq. 5
ϕ_1 (variable)	type 1 spillover coefficient in eq. 7
ϕ_2 (variable)	type 2 spillover coefficient in eq. 7
$L = 5$	lag for R&D moving average eq. 8
$\kappa = 4$	catch-up parameter eq. 10
$\Pi = 0.02$	mutation probability eq. 11
$s = 0.02$	standard deviation of mutation step size eq. 11
$\mu = 0.02$	maximum imitation probability eq. 12
$E = 0.005$	exit level in employment share
$b = 0.1$	labor productivity bandwidth for entrants

References

Aghion, P. and Howitt, P., 1992, "A Model of Growth through Creative Destruction", *Econometrica,* **60:** 323-351.

Arrow, K., 1962, "Economic welfare and allocation of resources for invention", in Nelson, R., (ed), *The Rate and Direction of Inventive Activity,* Princeton: NBER/Princeton University Press.

Arthur, W.B., 1991, "Designing Economic Agents that Act Like Human Agents: A Behavioral Approach to Bounded Rationality", *American Economic Review. Papers and Proceedings,* **81:** 353-359.

Binmore, K.G. and Samuelson, L., 1993, "Muddling Through: Noisy Equilibrium Selection", Alicante: paper presented at the workshop "Learning and Industrial Dynamics", May 31-June 2, 1993.

Cohen, W.M. and Levinthal, D.A., 1989, "Innovation and Learning: The Two Faces of R&D", *Economic Journal,* **99:** 569-596.

Foster, D. and Young, P., 1990, "Stochastic Evolutionary Game Dynamics", *Theoretical Population Biology,* **38:** 219-232.

Friedman, M., 1953, "The Methodology of Positive Economics", in *Essays in Positive Economics,* Chicago: Chicago University Press.

Friedman, D., 1991, "Evolutionary Games in Economics", *Econometrica,* **59:** 637-666.

Goodwin, R.M., 1967, "A Growth Cycle", in C.H. Feinstein, (ed), *Socialism, Capitalism and Economic Growth,* London: Macmillan.

Helpman, E., 1992, "Endogenous Macroeconomic Growth Theory", *European Economic Review,* **36:** 237-267.

Hofbauer, J. and Sigmund, K., 1988, *The Theory of Evolution and Dynamical Systems,* Cambridge: Cambridge University Press.

Kamien, M.I. and Schwartz, N.L., 1982, *Market Structure and Innovation,* Cambridge: Cambridge University Press.

Kandori, M., Mailath, G.J. and Rob, R., 1993, "Learning, Mutations, and Long Run Equilibrium in Games", *Econometrica,* **61(1):** 29-56.

Lane, D.A., 1993, "Artificial Worlds and Economics. Parts 1 and 2", *Journal of Evolutionary Economics,* **3(2&3):** 89-108, 177-197.

Langton, C.G., (ed), 1989, *Artificial Life,* Redwood City, CA: Addison-Wesley.

Langton, C.G., Taylor, C., Farmer, J. D. and Rasmussen, S., (eds), 1992, *Artificial Life II,* Redwood City, CA: Addison-Wesley.

Lucas, R., 1988, "On the Mechanisms of Economic Development", *Journal of Monetary Economics,* **22:** 3-42.

Nelson, R., 1959, "The Simple Economics of Basic Scientific Research", *Journal of Political Economy,* **67:** 297-306.

Nelson, R., 1990, "What is Public and What is Private about Technology?", Berkeley: CCC Working Paper No. 90-9.

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

Romer, P.M., 1986, "Increasing Returns and Long-Run Growth", *Journal of Political Economy,* **94:** 1002-1037.

Romer, P.M., 1990, "Endogenous Technological Change", *Journal of Political Economy,* **98:** 71-102.

Scherer, F.M and Ross, D., 1990, *Industrial Market Structure and Economic Performance,* Third Edition, Boston: Houghton Mifflin Company.

Schwefel, H.P., 1981, *Numerical Optimization of Computer Models,* Chichester: John Wiley.

Selten, R., 1989, "Evolution, Learning, and Economic Behavior", *Games and Economic Behavior,* **3:** 3-24.

Silverberg, G., 1984, "Embodied Technical Progress in a Dynamic Economic Model: the Self-Organization Paradigm", in Goodwin, R., Krüger, M. and Vercelli, A., (eds), *Nonlinear Models of Fluctuating Growth,* Berlin-Heidelberg-New York: Springer Verlag.

Silverberg, G. and Lehnert, D., 1993, "Long Waves and 'Evolutionary Chaos' in a Simple Schumpeterian Model of Embodied Technical Change", *Structural Change and Economic Dynamics,* **4** (in press).

Silverberg, G. and Soete, L., (eds), 1993, *The Economics of Growth and Technological Change. Technologies, Nations, Agents,* Cheltenham: Elgar (forthcoming).

Vega-Redondo, F., 1993, "Industrial Dynamics, Path-Dependence, and Technological Change", Alicante: paper presented at the workshop "Learning and Industrial Dynamics", May 31-June 2, 1993.

Young, H.P., 1993, "The Evolution of Conventions", *Econometrica,* **61(1):** 57-84.

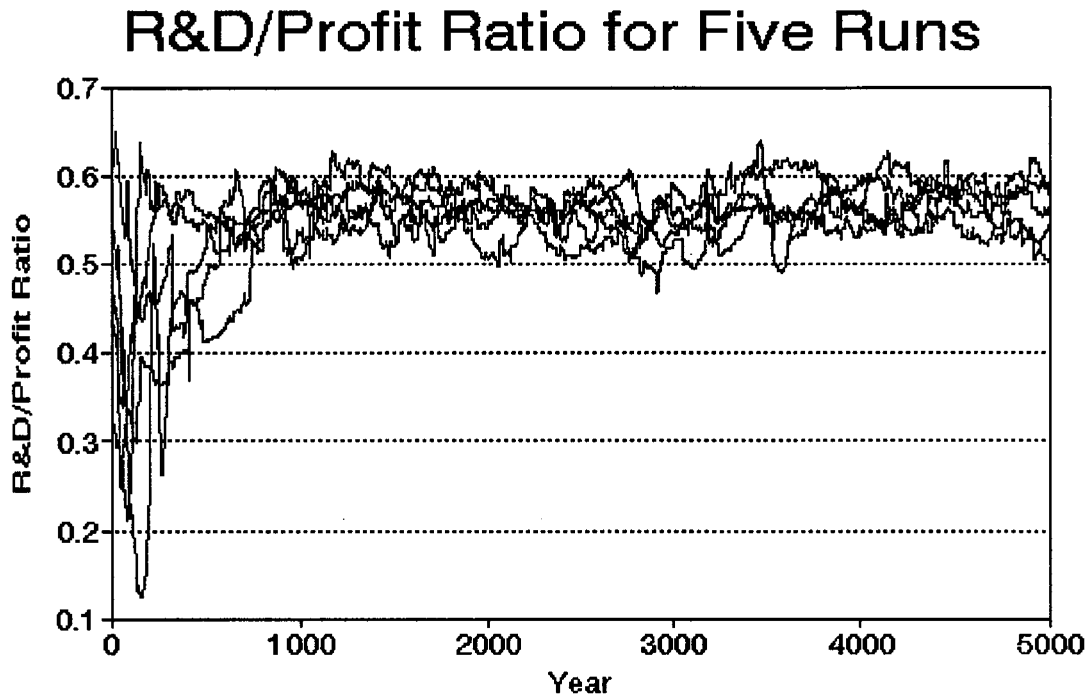

Figure 1. Five runs with $A = 7$ but using different random seeds and the "random generation" initial conditions.

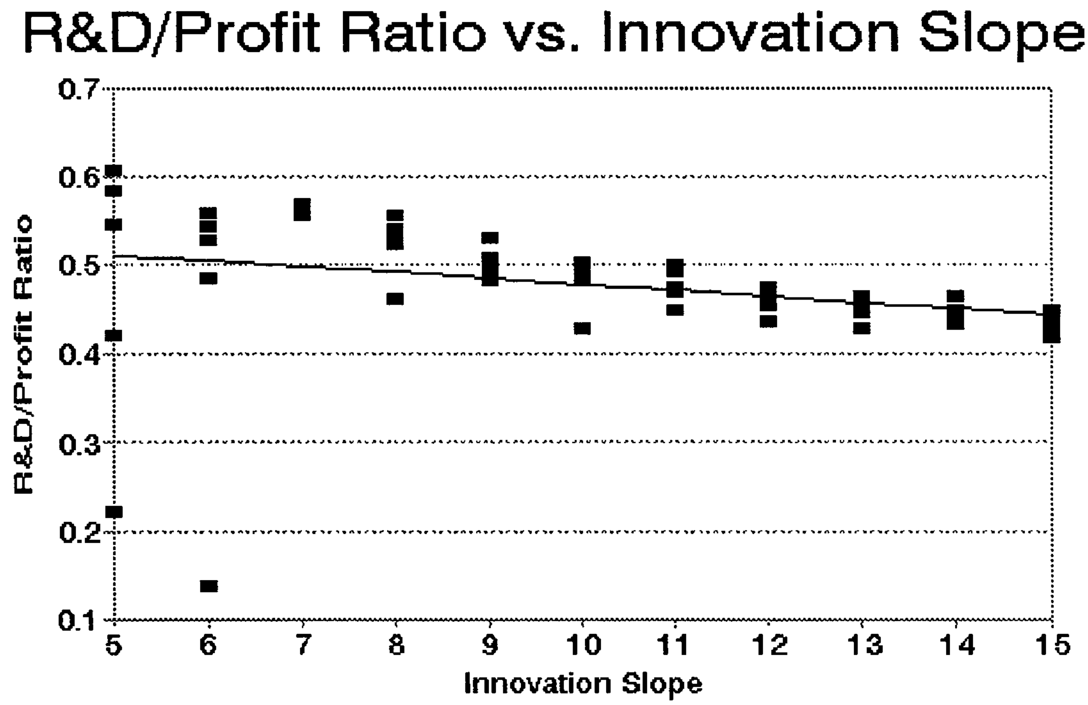

Figure 2. Random generaton with five runs per value.

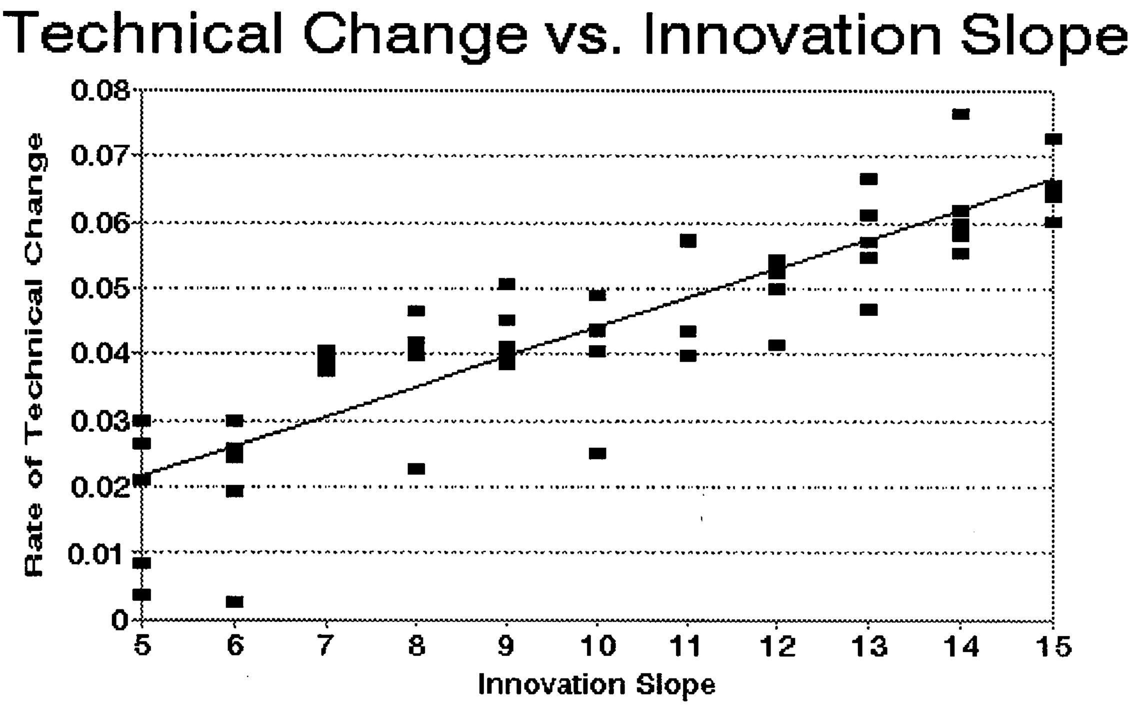

Figure 3. Random generation with five runs per value.

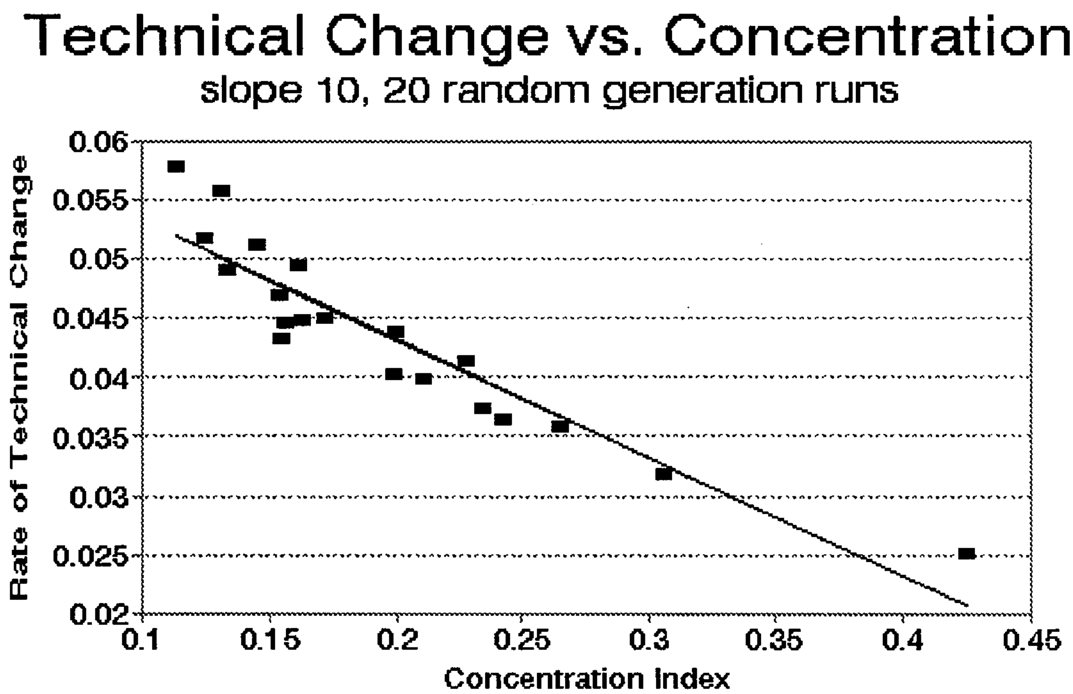

Figure 4. Random generation, 20 runs with fixed $A = 10$.

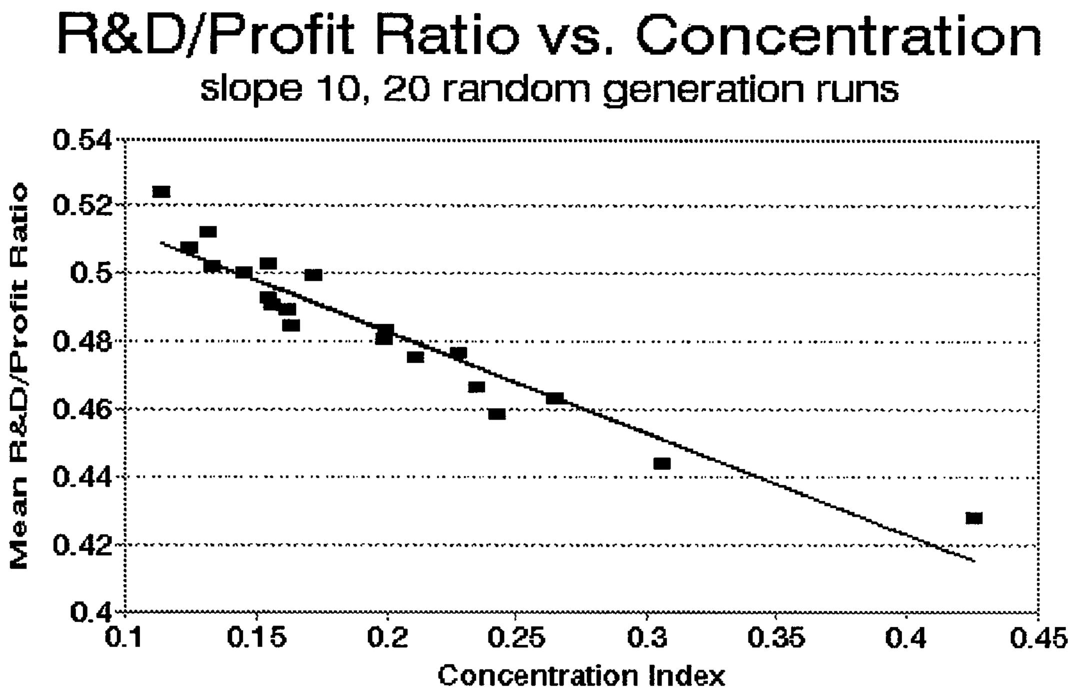

Figure 5. Random generation, 20 runs with fixed A = 10.

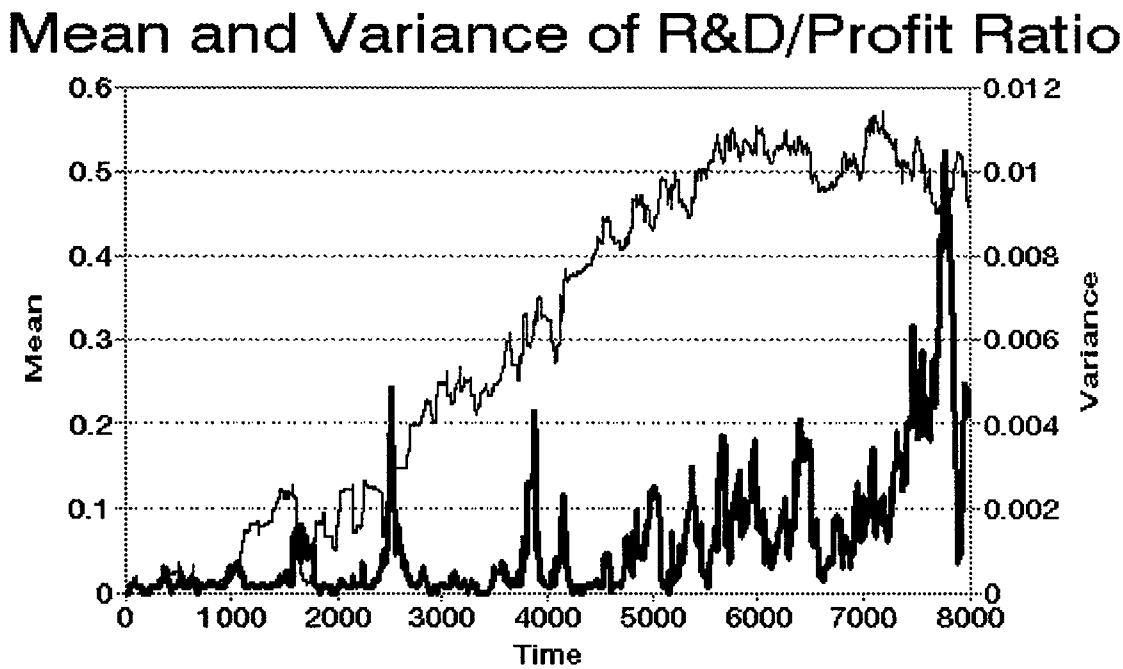

Figure 6. Light line is the mean of the aggregate R&D/Profit ratio. The heavy line is its variance.

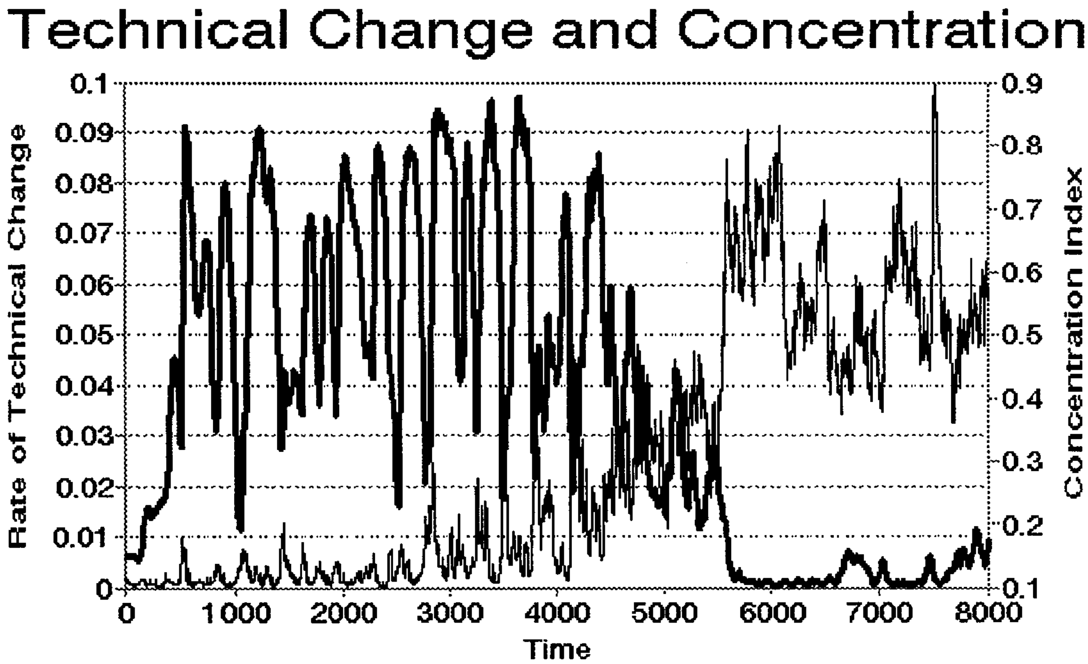

Figure 7. Light line is the rate of technical change (in % p.a., left scale), the heavy line is the Herfindahl concentration index (right scale).

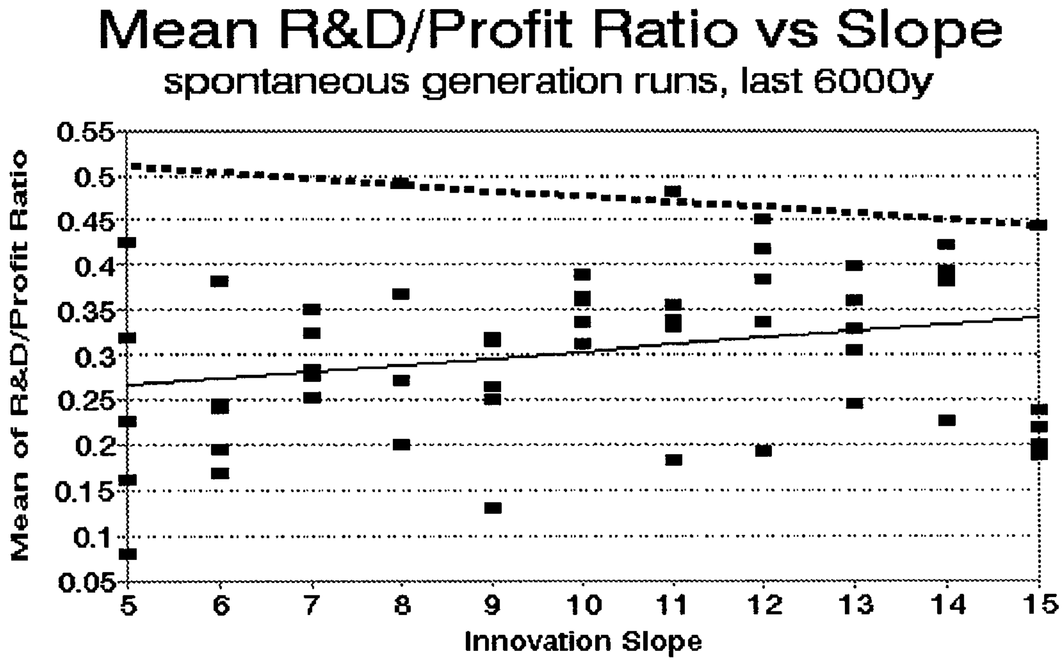

Figure 8. Dashed line is the regression obtained from the comparable random generation runs. Solid line is the regression for these data.

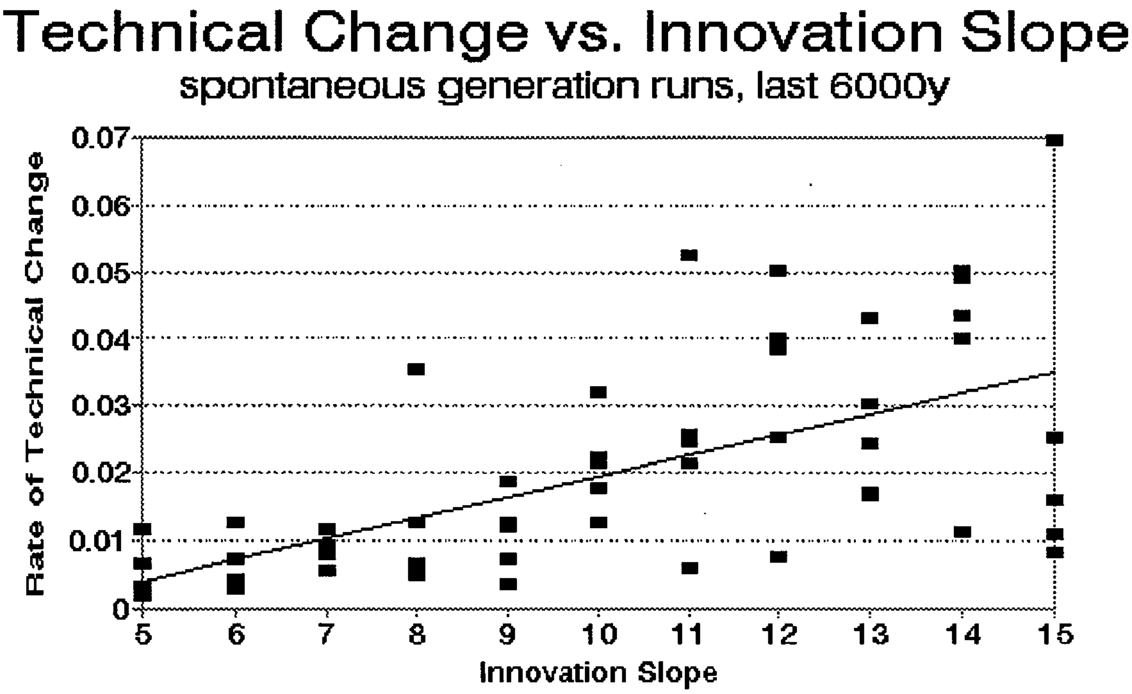

Figure 9. Five spontaneous generation runs per value.

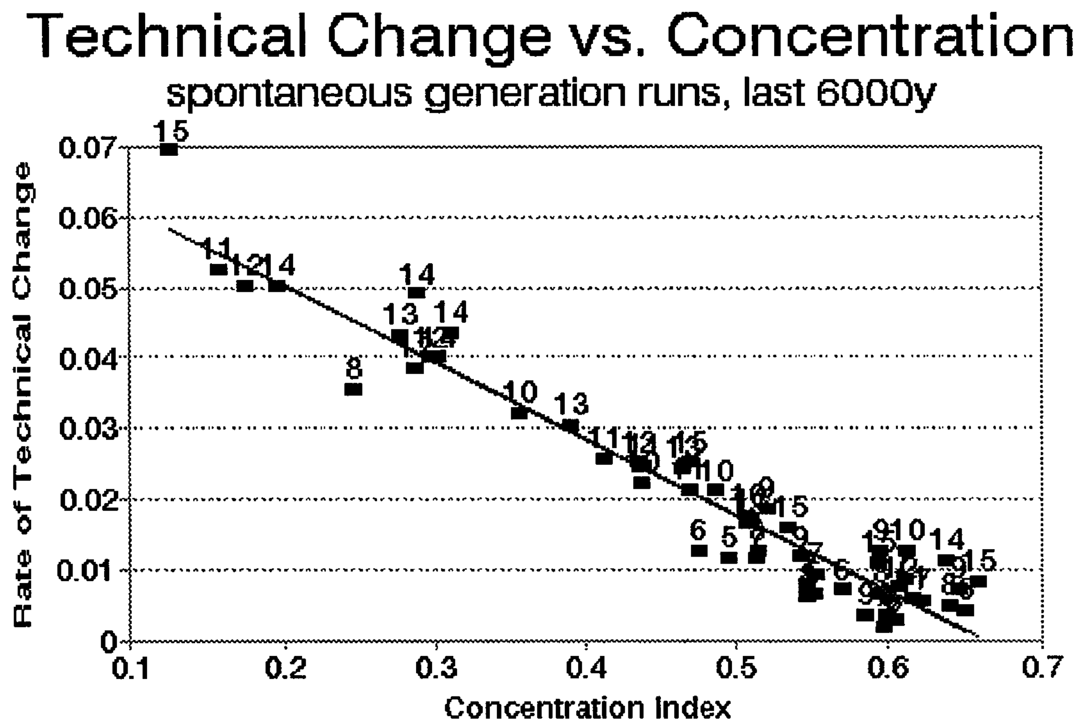

Figure 10. 55 pooled runs over all innovation slope values demonstrate clustering regimes.

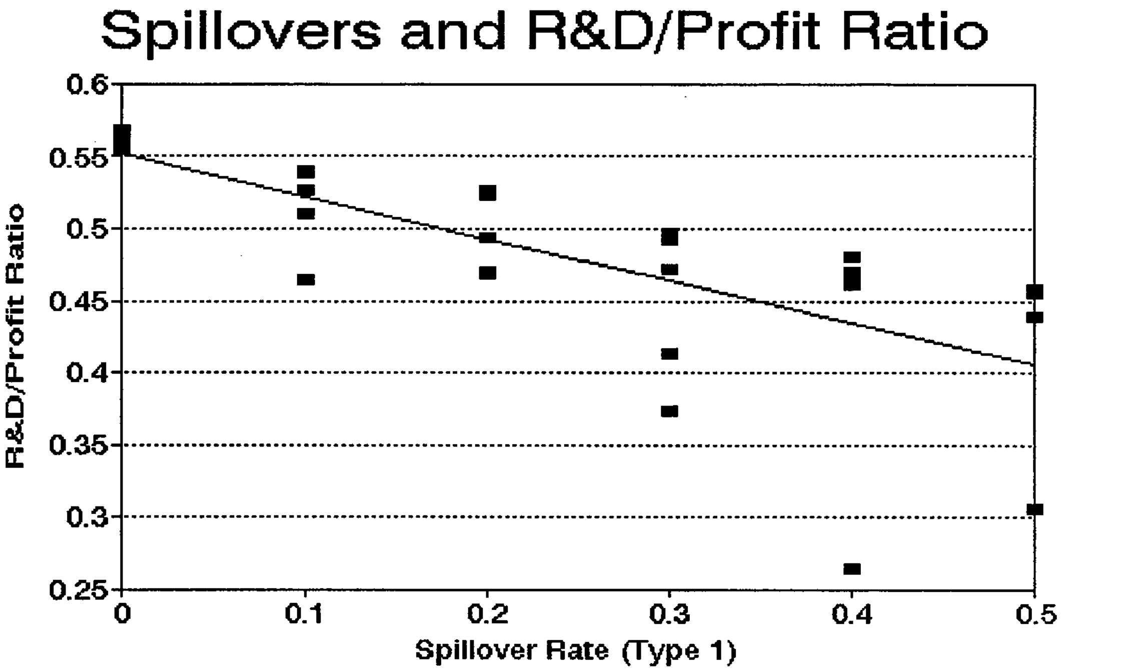

Figure 11. Five random generation runs per value with $A = 10$.

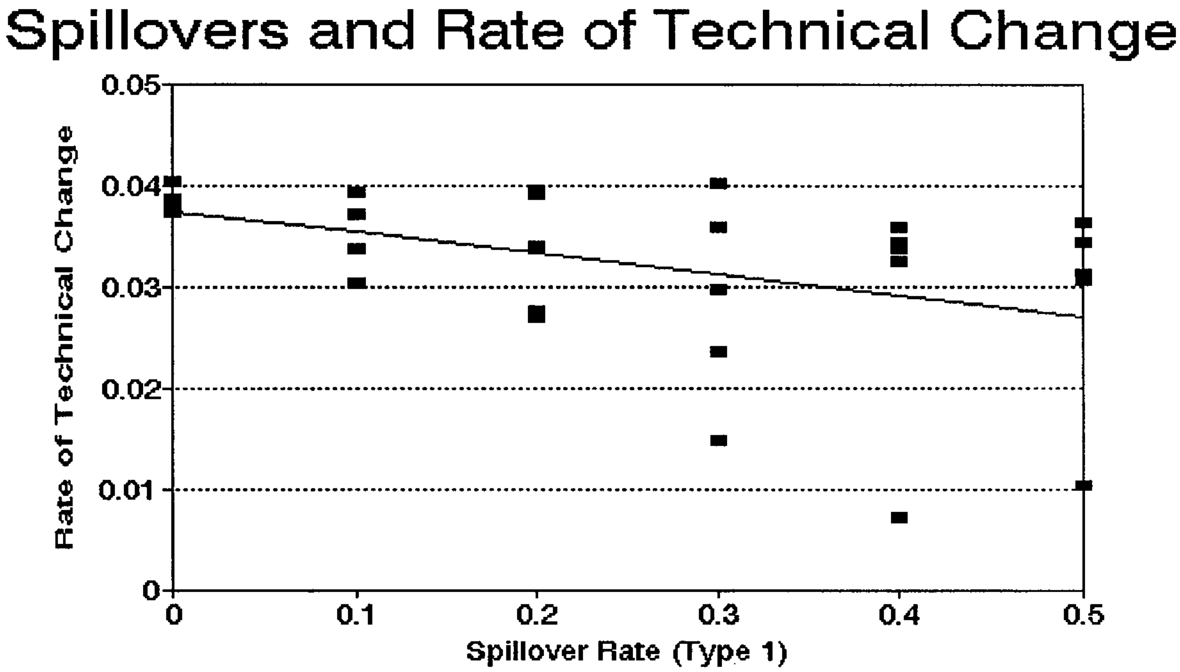

Figure 12. Five random generation runs per value with $A = 10$.

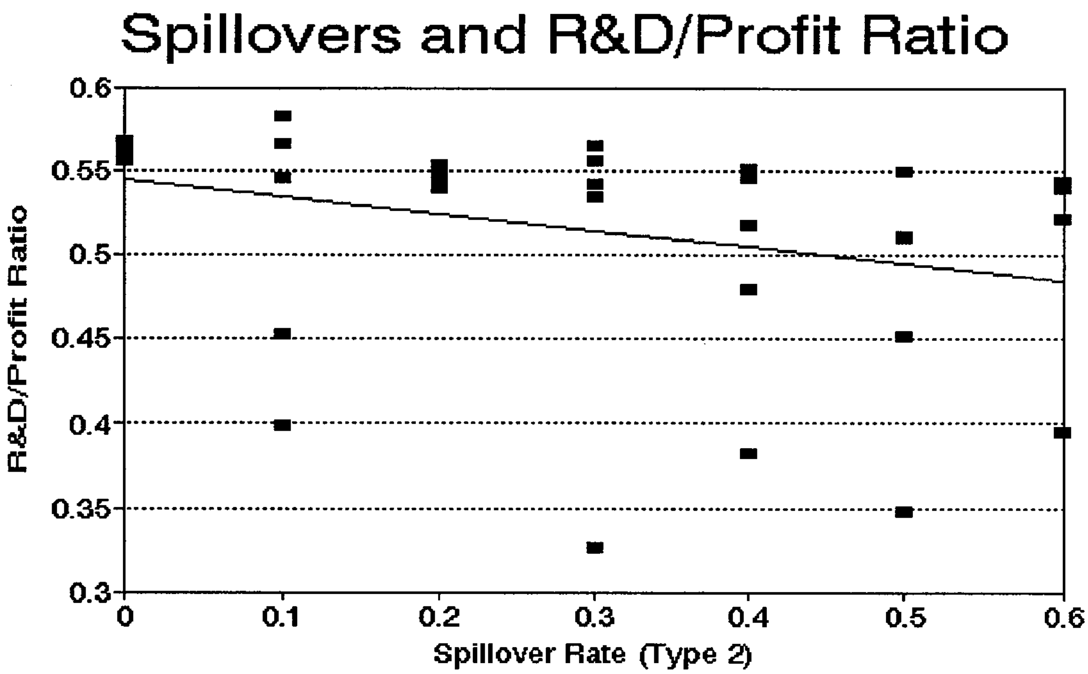

Figure 13. Five random generation runs per value with $A = 10$.

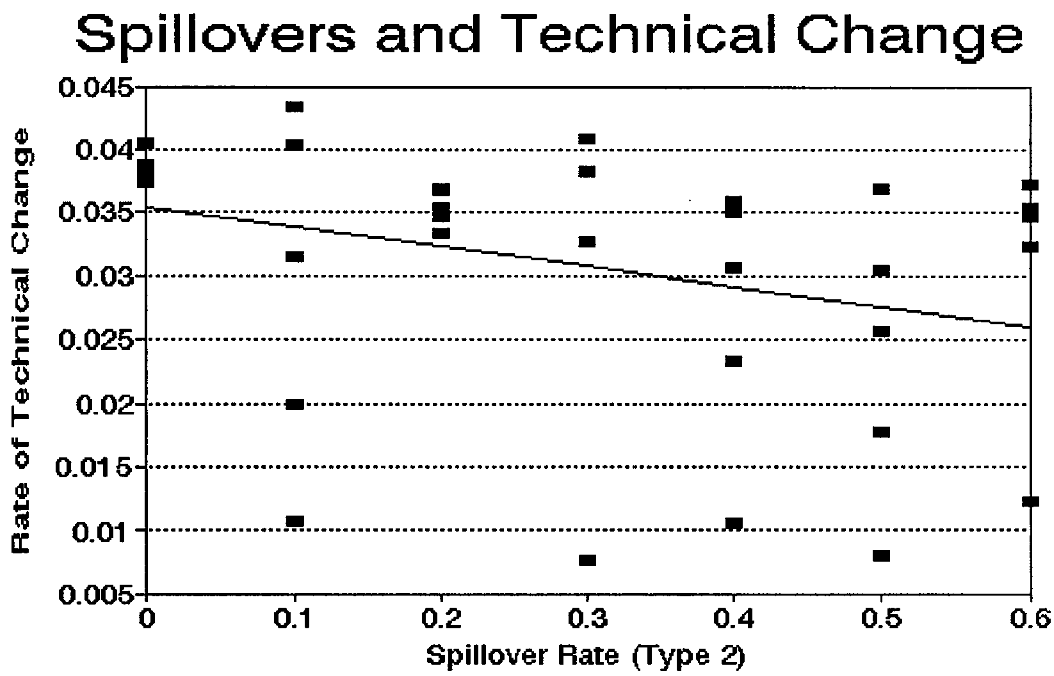

Figure 14. Five random generation runs per value with $A = 10$.

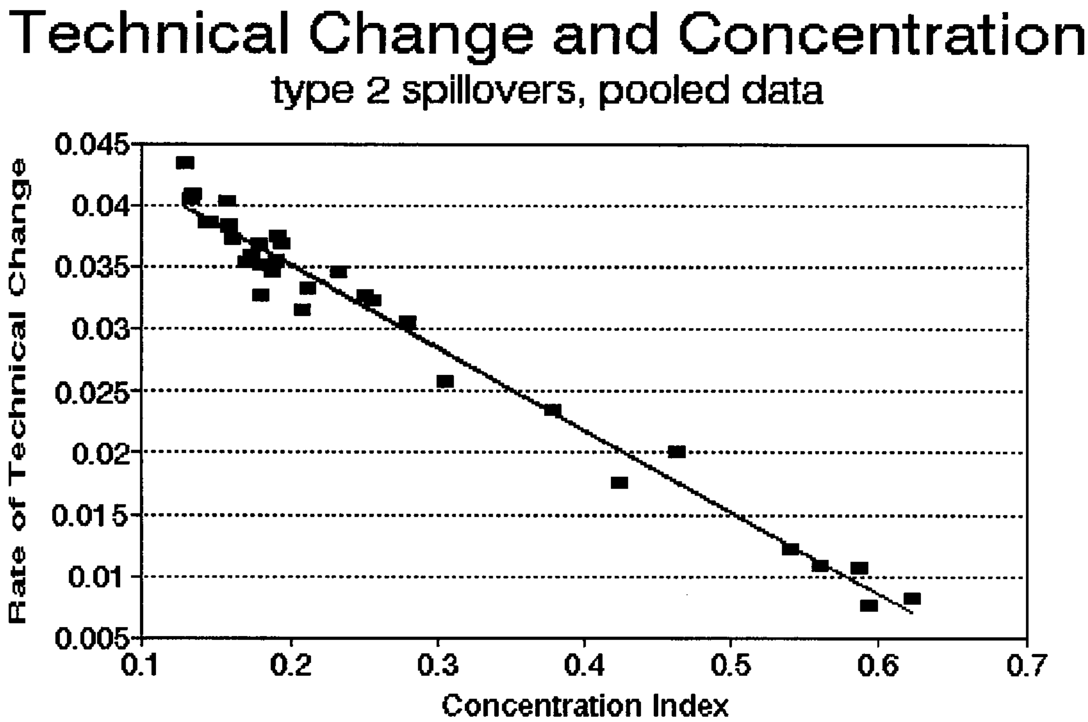

Figure 15. Pooled data display clustering in high and low growth regimes.