ELSEVIER

Research Policy 31 (2002) 1275–1289

research policy

www.elsevier.com/locate/econbase

The discrete charm of the bourgeoisie: quantum and continuous perspectives on innovation and growth

Gerald Silverberg

International Institute of Infonomics, Maastricht Economic Research Institute on Innovation and Technology, Maastricht University, Postbus 616, NL-6200 Maastricht, MD, The Netherlands

Received 13 September 2001; received in revised form 19 October 2001; accepted 7 March 2002

Abstract

This paper discusses the relative merits of discrete versus continuous perspectives on innovation, technical change, and economic growth. It discusses the innovation time series literature in some detail to extract the continuous and clustering properties of the historical record on innovation. It then proposes a mosaic/avalanche model based on percolation theory and self-organized criticality to address this question.

JEL classification: C32; O31; O4

Keywords: Innovation; Percolation; Economic growth; Technological change; R&D

1. Introduction

Is the process of technological innovation and economic growth better characterized as one of continuous and incremental change, or rather as marked by discrete events and quantum jumps? Different scholars have advocated each point of view, sometimes simultaneously. Is there a mechanism whereby a continuous stream of activity is converted into an intermittent or pulsating process, or conversely, one in which discrete events induce smooth unfoldings? To some extent such questions are a matter of levels, with aggregation leading to the appearance of smoothness and disaggregation revealing the graininess of the underlying system. Yet in other cases small microscopic events do seem to have significant macroscopic implications.

In this paper, I want to address the following questions. First, is the innovation process better regarded as a discrete one producing distinct and qualitatively different entities at identifiable points in time? If so, what is the time pattern of their production? Second, what is the relationship between the innovation process and the sectoral and aggregative state of an economy? Does the possible discreteness of the innovative process translate into particular characteristics of the meso- and macroeconomy?

The paper proceeds as follows. In the next section, I briefly review the literature on the history of innovations and their statistical properties. This is followed by a section reviewing some statistical studies I have been involved with on this issue that challenge a number of published results. The methodological basis for this critique is discussed and some implications of the analyses are drawn for the characterization of the innovation process.

E-mail address: gerald.silverberg@merit.unimaas.nl (G. Silverberg).

0048-7333/02/$ – see front matter © 2002 Elsevier Science B.V. All rights reserved.
PII: S0048-7333(02)00063-X

The following section examines the nexus between innovation and sectoral and aggregate productivity growth and macroconomic dynamics. I then briefly discuss the historical record of sectoral change and aggregate growth and the ways in which researchers have attempted to come to terms with it statistically and mathematically.

The final section sketches a framework for dealing with discontinuous change, still at a very preliminary stage, which I call the mosaic and avalanche theory. Mosaic refers to the fact that the innovation process consists of many small steps that need to be assembled to create a final 'Gestalt' before an operational innovation is achieved. The dynamics of this process is modelled on the lines of percolation and thus involve a *phase transition*, or, in the words of Hegel, a transformation of quantity into quality. This allows us to address the question of how a more or less constant stream of innovation activity can trigger clusters of innovation, as well as how innovation activity can focus itself or be induced as the outlines of the Gestalt emerge. Furthermore, it provides a natural framework for different *trajectories* of technological development to emerge, converge or cross-fertilize, and for science to act as an attractor for subsequent technological development.

The avalanche aspect addresses the unpredictable character of the extent of sectoral diffusion of an innovation and thus ultimately of the magnitude of its effects on the overall economy. Following the literature on self-organized criticality (SOC), these diffusion effects results from chain-reactions of local interaction subject to non-linear thresholds and can be used to endogenize (at least in a statistical sense) the growth waves initiated by an innovation.

1.1. Incrementalism versus radicalism in innovation studies

Traditional neoclassical theory of growth and technical change is most closely associated with a thoroughgoing incrementalist approach. Technical change is envisioned as the smooth shift in time (at an exogenously given rate) of a smooth, substitutable production function. There is no room in fact for distinguishable, discrete technologies (does one identify a technology with a point on the production function, or the entire curve?). In some sense they have been aggregated into one macrotechnology.[1] This has changed to some extent with the advent of Schumpeterian endogenous growth models (initiated by Aghion and Howitt, 1992), which share many, if not all, elements with an evolutionary perspective. Technologies, as in Silverberg and Lehnert (1993), are nodes on a directed graph, here the graph being the very simplest case of a linear array.[2] We do indeed perceive technologies as discrete entities in this framework, although by assumption in the first instance (as in Silverberg and Lehnert) and as derived from an intertemporal perfect foresight equilibrium argument based on patent-race reasoning (in Aghion and Howitt), the stochastic arrival rate of these discrete innovations is constant. In stochastic terms, this represents a steady flow, but the realization is a Poisson jumping process. In AG, in fact, the whole economy performs these jumps, since technologies are disembodied and there are no diffusion lags. In SL, due to diffusion and intertechnological competition, the aggregate effects demonstrate complex dynamics as innovations propagate through the economy.

This highly simplified picture concentrates on the salient aspects of the Schumpeterian vision: radical, discrete innovations drive the economy. However, we know from innovation studies that a considerable part of technical change is due to so-called learning effects, i.e. the accumulation of almost imperceptible small increments of skill, design, and organizational improvements. While the mechanism behind the various forms of technological learning (learning by doing, learning by using) is not well understood, phenomenologically it can be characterized by power-law relationships, known as learning curves, of the form:

$$y = x^{\alpha}$$

where y is some measure of performance (e.g. productivity, unit costs) and x is some measure of experience (e.g. cumulative production). On this basis, Arrow (1962) produced the first endogenous growth model, and most of the endogenous growth literature à la

[1] Indeed this was the conclusion of the famous Cambridge controversies of the 1960s—an aggregate production function only made sense when all sectoral technologies had the same technical coefficients.

[2] For a model also based on a directed graph of technologies but allowing for branch points, see Vega-Redondo (1994).

Romer, Lucas, etc. relies on some similar mechanism (Solow, 2000). Once again, as in the 'exogenous' growth neoclassical literature, there are no identifiable 'discrete' technologies, and growth inevitably takes the form of an exponential steady state.

Finally, the stochastic continuity of the Poisson model can be ruptured by some notion of clustering or punctuation of innovation at large time scales. What meaning these should have statistically has never been explicitly formalized. One interpretation is that clusters of innovations occur more or less periodically, in synchrony with purported 50 years Kondratieff cycles of economic activity. A slightly variant interpretation is that a particular phase of long-term economic fluctuations triggers higher (radical) innovation activity due to a higher propensity of risk taking (the 'depression-trigger hypothesis'). A more agnostic interpretation is that the underlying stochastic process has a natural tendency to random clustering at different time scales, independent of any economic influences, much like earthquakes.

Economic history and the history of technology of course have always suggested a close relationship between emergent and dominant technologies and the salient economic character of different epochs. Whether these were formulated as ages of steam, steel, electricity, the automobile, or the computer, on the one hand, or the first, second, third industrial revolutions on the other, this quasi-Marxian technological determinism has implicitly influenced much historical thinking. Attempts to capture these effects quantitatively have relied, e.g., on the methodology of technological substitution (cf. Nakicenovic, 1987; Grübler, 1990, 1998). In particular, large-scale infrastructural technologies (primary energy sources, transport infrastructure such as canals, railroads, highways, airways) display clear evidence of 50 years replacement cycles and logistic diffusion.

Recently, the concept of *general purpose technologies* (GPT) has been coined to represent those innovations whose influence becomes so pervasive that they become significant inputs and sources of productivity growth in a large number of sectors and final consumption goods (cf. Helpman, 1998). Although these authors often fail to highlight the connection of this concept with Schumpeter's idea of radical innovations, it is clear that this is a closely allied notion.

The neo-Schumpeterian evolutionary literature has placed particular emphasis on discontinuities in economic development. However, in its formal implementations, this is not quite so obvious. The 'standard' model of evolutionary economic growth due to Nelson and Winter (1982) considers technologies as equivalent to behavioural routines. Innovation takes place in a characteristics space of technologies (labour and capital coefficients) by jumping with a certain probability determined by R&D expenditures between points (technologies) generated at random. The only structure imposed on the space is the topology induced by a technological distance metric. Since the technologies are capital-disembodied, no investment going beyond the original R&D is necessary, the productivity effects of innovation are immediate and no diffusion takes place, although imitation effects can lead to the diffusion of the routine from one firm to another.

In the appreciative and applied evolutionary literature much has been made of the concepts of technological paradigm (Dosi, 1982) and natural trajectories (Nelson and Winter, 1977). This is indeed an attempt to impose additional structure on technology and differentiate discrete interrelationships in technological space from one another, if only ex post (empirical evidence for such trajectories has been advanced by Sahal, Saviotti and Leonard, e.g.). This should be contrasted with the smooth, substitutable, unbounded production possibility sets of neoclassical theory (but is possibly related to the notion of localized learning introduced by Atkinson and Stiglitz, 1969). Very little in the way of insightful modelling has been done in this regard without assuming the exogenous existence of parameterized trajectories (as in Silverberg et al., 1988). A recent attempt in this direction is the so-called history-friendly modelling approach (Malerba et al., 1999), but again, one has the impression that too much explicit historical structure has to be hardwired into the model first to get anything plausible out of it.

1.2. The time pattern of radical innovations and the size distribution of innovations in general

Schumpeter is the author most responsible for highlighting the role of radical innovations in economic change. His theory of economic development posited the existence of innovation clustering and bandwagon effects in investment to drive waves of creative destruction, i.e. periods of intense structural change and productivity growth attendant on the introduction of new products, processes, and forms of business organization. But already in 1940, Kuznets (1940) in his review of Schumpeter's *Business Cycles*, raised serious objections to the logical consistency and historical plausibility of this perspective.

It was not until the 1970s, that any serious attempt was made to examine the time pattern of radical innovations empirically. Mensch (1975) drawing on the innovation data contained in Jewkes et al. (1958) undertook the first influential statistical analysis of innovation time series. While the test he used (a runs test) was appropriate for testing his null hypothesis of iid of the yearly count data, his treatment of the data was challenged by Clark et al. (1981). Indeed, all attempts at compiling data on radical innovations have been marred by definitional and dating problems. Is there an objective definition of what constitutes a radical innovation? Is there an unambiguous way of dating its time of introduction? Almost all major innovations have gone through stages of development in different times and places which make it difficult to determine exactly when they have reached maturity and have attained a recognizable and viable form. Their impact on the economy also varies considerably as they pass through more advanced stages of development. While these ambiguities cast doubt on the reliability of any of the innovation time series employed in this literature, as we shall later argue, this state of affairs can also be made a rich source of dynamic structure which can be used to elucidate the relationship between minor and major innovations and advance our understanding of the temporal unfolding of the innovation process.

Additional innovation time series have been compiled by Clark et al. (1981); Haustein and Neuwirth (1982); van Duijn (1983); Kleinknecht (1987, 1990a) on the basis of UK patent data published by Baker (1976). An attempt to reconcile the disparities in these series and compile a consistent super series is presented in Silverberg and Verspagen (2000). Statistical analysis of these time series has been undertaken by Sahal (1974), Kleinknecht (1981, 1987, 1990a,b); Solomou (1986); Silverberg and Lehnert (1993) and most recently Silverberg and Verspagen (2000).

It has not always been clear in this literature what the appropriate null hypothesis should be. While Mensch simply tested for iid in the yearly count data, Kleinknecht and Solomou applied z- and t-tests to binary comparisons of adjacent sub-periods (chosen with reference to their datings of long cycles in aggregate growth). As pointed out by Silverberg and Lehnert (1993), z- and t-tests are only appropriate for normally distributed random variables. However, if the innovation process is thought of as a stochastic but non-clustering point process, then the time-homogeneous Poisson process is the obvious

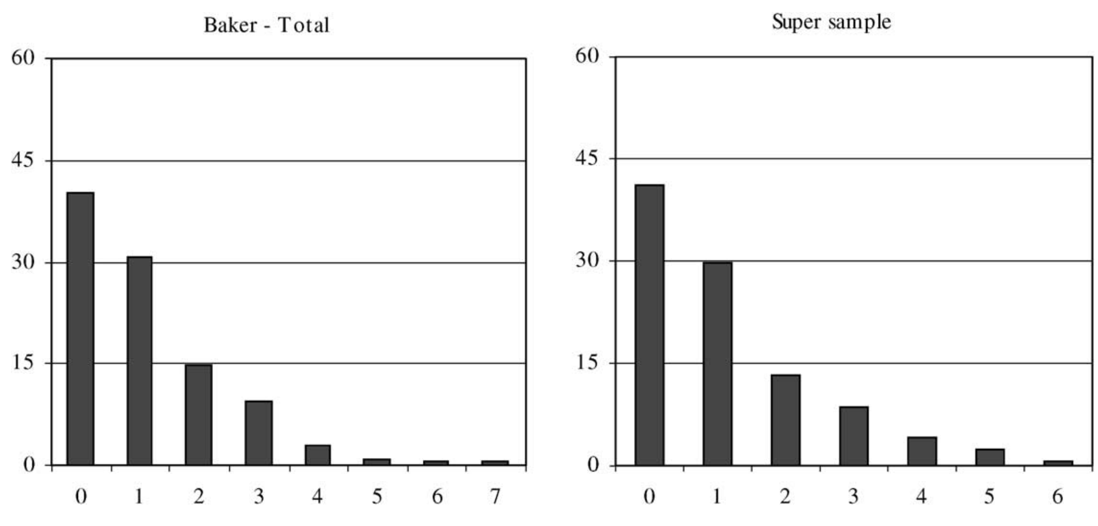

Fig. 1. Histogram of innovation time series for Baker and supersample data (x-axis: number of innovations per year, y-axis: percent of years).

null. In this case, the count data for various sub-periods will certainly not be normally distributed, as histograms readily reveal (Fig. 1). Furthermore, the use of sub-periods is open to the objection that it does not exclude the possibility of a selection bias of the samples. If the criterion used to select the sub-periods correlates somehow with the innovation data (even though they themselves may be completely random), the sub-periods will extract periods of higher and lower mean realized activity and thus invalidate the tests. And as Silverberg and Lehnert demonstrate using a dynamic model even on the assumption of independent homogeneous Poisson innovation data, such correlations with aggregate macroeconomic data (with causality running from innovations to macrovariables) are to be expected. Thus, ideally one should test the null hypothesis of a homogeneous Poisson process or other point process on the entire data set without invoking sub-periods. This is what Sahal and Silverberg and Lehnert have done using non-parametric methods, and what Silverberg and Verspagen have done using parametric Poisson regression techniques.

All of these tests concur in rejecting the time-homogeneous Poisson process (see top panel of Table 1). A visual inspection of any of the time series, however, immediately reveals that this could also be due to the presence of an underlying trend (Fig. 2). The second panel of Table 1 presents the results of detrending the original series using an estimator of the exponential growth rate of a non-time-homogeneous Poisson process (see Silverberg and Lehnert, 1993 for details on the techniques employed). Regardless of one's position with respect to the clustering issue, this exercise has succeeded in uncovering one fact that had been buried in the time series and obscured by the clustering controversy—namely, that the rate at which innovations have been made has itself been increasing. Its own exponential rate of increase has been in the range of 1/2 to 1% p.a., depending on the series and the time period examined, with product innovations increasing faster than process ones. This is obviously one example of continuity of the innovation process that also reinforces the notion of a robust evolutionary time trend.

Kleinknecht, under the influence of the 'depression-trigger' hypothesis of Mensch, has claimed that clustering of innovations exists and occurs in the depression or early recovery phases of 50 years long

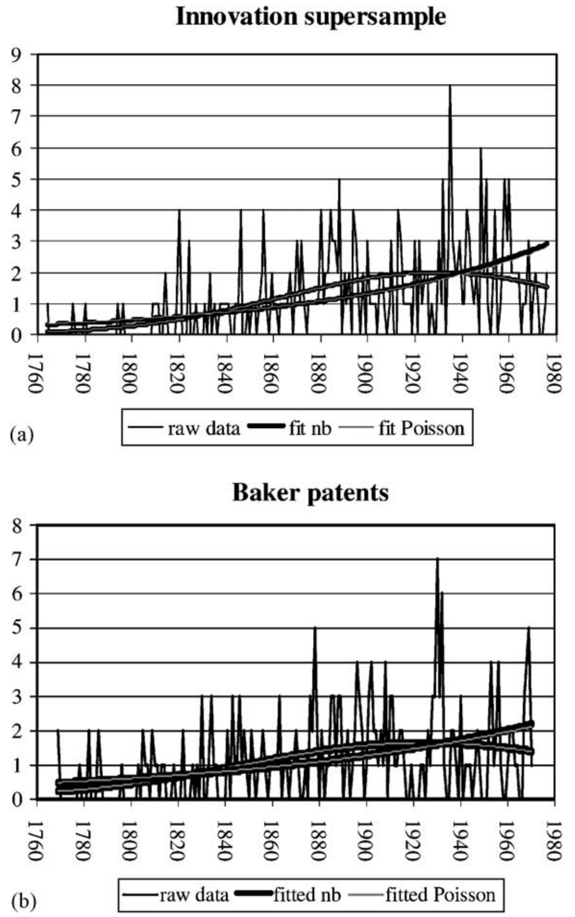

Fig. 2. Raw data and four fitted regression models for supersample and Baker data.

waves of general economic activity. This is a conception of causality running from general economic conditions to the propensity to innovate, to which one can immediately take exception, e.g. by pointing to the general decline in R&D activity in depression periods.[3] In any event this does not result from any inherent tendency of innovations to cluster by themselves.[4] Silverberg and Verspagen (2000) attempt to test variants of a Schumpeterian clustering of innovation hypothesis using Poisson regression tech-

[3] To which Kleinknecht has replied that it is the relative propensity to embrace radical versus incremental innovations that is at issue here.

[4] In contrast, Clark et al. (1981) have pointed to the common origins of many innovations of the 1930s and 1940s in advances in polymer chemistry as a reason for this particular cluster.

Table 1
Non-parametric tests of Poisson distribution of original and detrended innovation time series

Time series	Dispersion (d)	H-test	d.f.	Growth rate (%)	z-Trend
Original innovation time series					
Baker (1976) all years					
Product	741.08[b]	731.59[b]	273	0	18.11[b]
Process	465.67[b]	442.84[b]	273	0	7.93[b]
All	775.17[b]	815.21[b]	273	0	19.19[b]
Baker (1976) from 1769 on					
Product	438.66[b]	465.69[b]	202	0	11.88[b]
Process	304.17[b]	308.17[b]	202	0	2.66[b]
All	422.05[b]	443.79[b]	202	0	11.09[b]
Kleinknecht (1990a); based on Baker (1976)					
Product	336.27[b]	287.75[b]	201	0	7.11[b]
Process	236.94[a]	214.91	201	0	1.11
All	295.98[b]	305.15[b]	201	0	6.15[b]
Clark et al. (1981)	76.39	86.99[a]	64	0	0.77
Haustein and Neuwirth (1982)	336.20[b]	310.24[b]	211	0	6.68[b]
van Duijn (1983)	195.58[b]	171.79[b]	115	0	2.29[a]
Detrended innovation time series					
Baker (1976) all years					
Product	381.01[b]	380.79[b]	273	1.24	2.43[a]
Process	374.41[b]	378.44[b]	273	0.63	0.80
All	402.24[b]	426.68[b]	273	0.99	2.06[a]
Baker (1976) from 1769 on					
Product	288.78[b]	320.09[b]	202	1.00	1.20
Process	288.14[b]	301.09[b]	202	0.28	0.05
All	298.49[b]	318.64[b]	202	0.72	0.66
Kleinknecht (1990a); based on Baker (1976)					
Product	245.79[a]	235.12[a]	201	1.15	1.16
Process	234.56[a]	213.68	201	0.20	0.07
All	246.90[a]	266.72[b]	201	0.73	0.58
Clark et al. (1981)	76.43	86.41[a]	64	0.52	0.13
Haustein and Neuwirth (1982)	300.00[b]	264.48[b]	211	0.86	0.83
van Duijn (1983)	179.26[b]	166.48[b]	115	0.78	0.28

[a] Significant at 5% level.
[b] Significant at 1% level.

niques. By examining residuals of trends of different orders, one can test for autoregressive ‘knock-on’ effects or periodicity. The main conclusion is that an overdispersed model such as the negative binomial[5] with a polynomial trend of third order is preferred, and that no obvious periodicity is present. While autoregressive elements remain, they show no tendency for innovation shocks to persist.

There is another stream of research, which has attempted to measure the size or significance distribution of innovations and scientific publications. This has been approached using scientometric methods such as citation analysis to measure the importance of a patent or scientific publication by the number of times it is

[5] In the negative binomial model the Poisson arrival rate itself become a random variable and is drawn independently from a gamma distribution. While the mean remains the same, the variance of the process will exceed the mean, i.e. the process becomes overdispersed in comparison with a pure Poisson process.

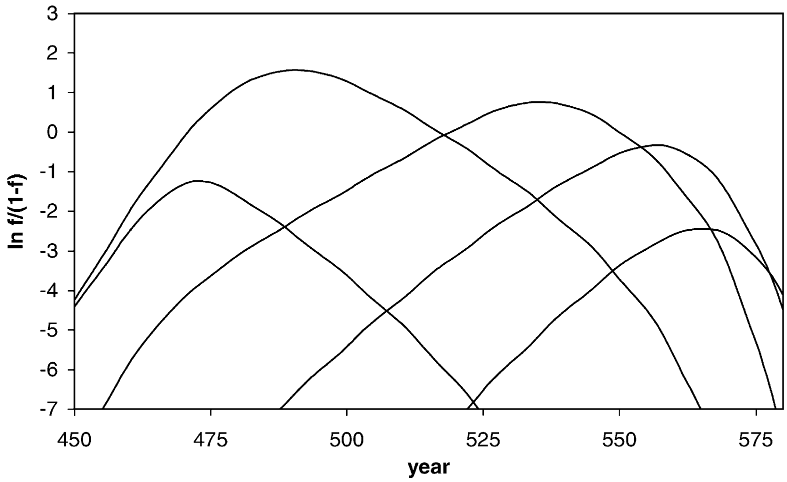

Fig. 3. Diffusion curves for a succession of innovations: slopes and saturations levels are endogeneous (the Fisher-Pry transformation $\ln f/(1-f)$ is plotted on the y-axis, where f is the share in total output of each technology). Straight segments correspond to logistic diffusion.

subsequently cited. The work of Trajtenberg (1990) and van Raan (1990) has demonstrated that these distributions are highly skewed and can be interpreted as indicative of a fractal mechanism underlying intellectual search and discovery.

1.3. Innovation and productivity growth

There is a vast literature discussing the relationship between innovation and productivity growth. There are three points, I want to emphasize in this connection:

1. Radical innovations can be seen as initiating growth 'pulses' with random amplitudes and at random times (Fig. 3). These pulses only translate into macroeconomic productivity gains with considerable delay due to diffusion lags. Even in such a simple model as Silverberg and Lehnert (1993), the pulses are realized with different speeds and different amplitudes although the underlying innovations each represent the same quantum improvement over their predecessors. The net result of this random (but steady) flow of pulses is complex dynamics in aggregate variables, with fluctuations over a range of time scales but more structure than a random walk.
2. What a model like Silverberg and Lehnert (as well as most other models) does not capture are the cyclical productivity effects of major innovations that are not directly derived from the innovations themselves. In that radical innovations usually initiate a surge in investment (justified or not) to build up the necessary productive capacity and infrastructure associated with them, due to backward linkages, multipliers and the inflation of general animal spirits and effective demand, the rates of utilization of other (traditional) sectors will go up. This may lead to a cyclical boost in productivity (Okum's law) even before the productivity gains of the new technologies can be realized in any significant way. Thus, Gordon's (2000) scepticism about whether the 'new economy' represents more than just a cyclical effect may in one sense be justified. In another sense, however, the cyclical effect may well be a non-trivial result of the investment boom unleashed by that very 'new economy' which would otherwise not have happened. The true dividends of the latter may not be realized for decades, if at all, and may or may not be comparable to those of other radical innovations like indoor plumbing and electricity.
3. These backward linkages in turn can induce real technological productivity gains in other sectors not directly related to the new technologies by accelerating the rate in which these sectors go down their learning curves (Verdoorn effect). Whether

these gains should be attributed to the radical innovations is a moot point, but they will certainly not be picked up by growth accounting exercises as such.

1.4. The mosaic/avalanche perspective on the innovation process

As we have seen, even the identification and dating of radical innovations are fraught with difficulty. Kleinknecht (1987, p. 61) illustrates this problem in a revealing way with a quote from Brockhoff (1972, p. 283; Kleinknecht's translation with my improvements):

> In 1818, K.V. Drais de Sauerborn presented his Draisine, a kind of walk-drive bicycle (Laufrad). In 1839, Mannilau demonstrated how wheels can be driven by pedals, and in 1861 at the latest pedals were built into the Draisine. In 1867, they were used on the front wheel by Michaux, and during the next few years the bicycle industry in France grew rapidly. A model of the bicycle approaching the one we are accustomed to today was constructed by Lawson in 1879, but a commercially successful 'safety bike' was not introduced by Starley until 1885. If we take 1818, 1839 or 1861 alternatively as years of invention, and 1867, 1879 or 1885 alternatively as years of basic innovation, we can obtain nine different results for the time-span between invention and innovation.

Undoubtedly, numerous other examples could be found in the history of technology to reinforce this point. What we normally perceive as a unitary entity, a radical innovation, in reality is usually composed of a number of smaller steps dispersed in time, often involving borrowing from other fields or dependent on specific unrelated advances in order to make the final step possible. In the bicycle case, we could add the availability of pneumatic tires and ball bearings (and thus precision machining, the precision grinding machine ...) as essential complementary innovations without which the bicycle boom of the 1890s would have been unthinkable. The bicycle is not one innovation but a succession of several smaller ones. In fact, our problem is not reducible à la Schumpeter to just radical versus incremental innovations; rather innovations come in all sizes, suggesting a fractal structure to the process of innovation.

Two lessons can be drawn from this example. First, all major innovations are decomposable into smaller, almost microscopic elemental steps, many of which if not the majority being missteps lost to history. It is these elemental steps that, I would posit, arrive independently and stochastically at an almost uniform rate until such time as the final form or 'Gestalt' of the innovation becomes increasingly visible from the random pieces of the mosaic. It is only then that innovative activity accelerates and becomes more purposeful and focused as the pack closes in for the kill. Second, associated with each such innovative trajectory is a threshold of performance, which, when first surpassed, unleashes significant diffusion into specific areas of the economy. Before that point is reached, the innovation remains a prototype, an object of tinkering and the consumption of hobbyists. As each threshold of performance is passed, diffusion into ever-wider reaches of the economy (differentiated e.g. by sector or income category) is triggered.

This mechanism would thus appear to be a highly non-linear phenomenon: first, the coalescence of the innovation only when the mosaic is completed, and second, the surge of diffusion as performance (or better, price-performance) thresholds are surpassed. In the following, I will sketch a modelling framework for addressing these phenomena in an admittedly highly abstract but, I would suggest, empirically non-trivial manner.

Imagine a space of discrete technologies, with nearness corresponding to some measure of technological proximity (Fig. 4). The bottom line represents such a space in one dimension. In reality, of course, technological characteristics are so multitudinous that a much higher dimensional space might be more appropriate. And other topologies might be more suitable than the lattice we will impose on it, such as some sort of network structure. Periodic boundary conditions (pasting the left and right edges together) are also sensible. The vertical dimension represents the performance level of each technological category (also seen as one dimensional for simplicity). On this half plane, we regard the elemental innovation process as the filling of a lattice site with a certain probability P, i.e. as what is

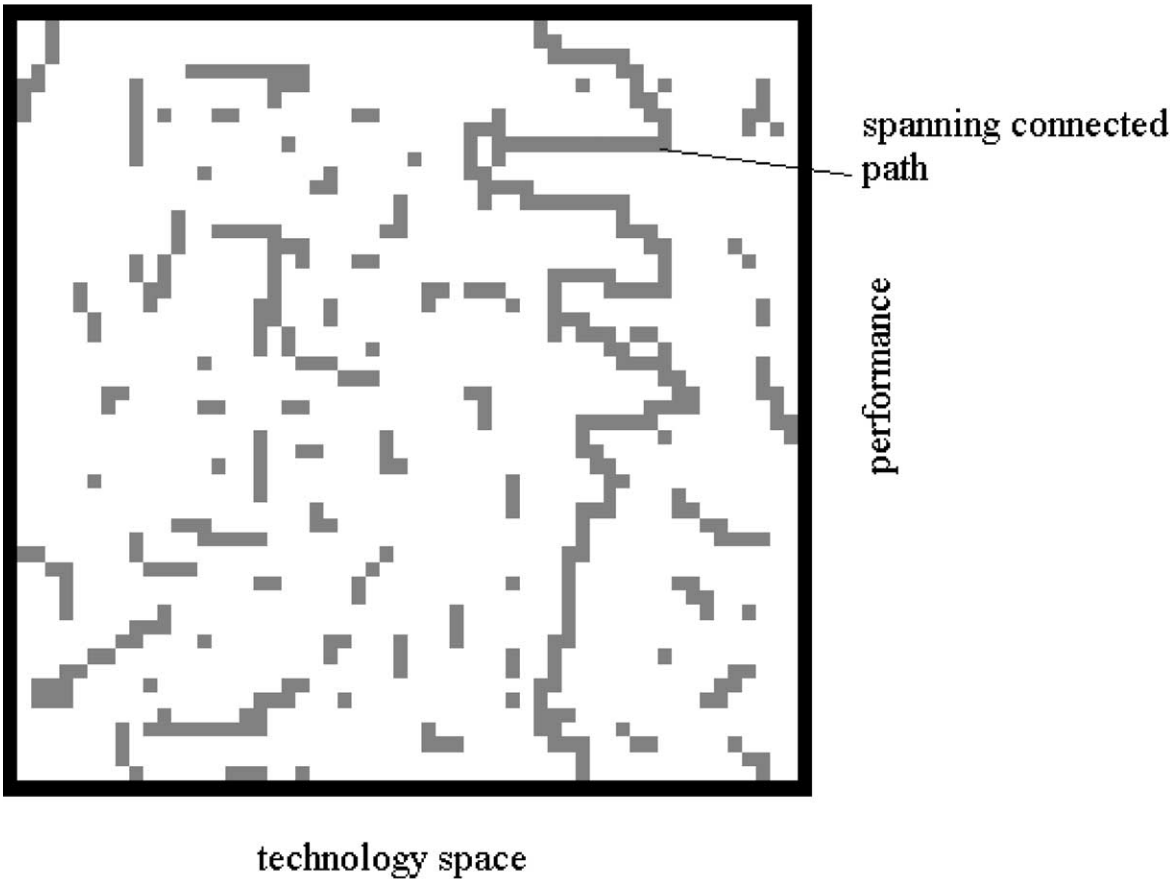

Fig. 4. Percolation diagram in technology-performance space. Lattice sites are filled at random. A site is viable when it connects to the baseline.

known as a *percolation* problem.[6] We will consider a filled site to represent an operational discovered technology if it is connected to the bottom of the diagram by a filled path of such sites (using nearest-neighbour connections, i.e. up, down, left, right).

The essential property of percolation is the behaviour of connected sets as a function of the (uniform and independent) probability of occupation of sites. On an infinite lattice (including the half plane) there exists a threshold probability, P_c, below which there is no infinite connected set and above which with probability 1 there is one (and only one) infinite connected set. The probability that any site will belong to the infinite connected set is obviously zero below P_c and increases continuously and monotonically above P_c.[7] For bounded lattices such as in Fig. 3, the interesting question is the probability of finding a connected path

[6] In this case, we speak of *site* percolation, as opposed to working with the lines connecting nodes, known as *bond* percolation (Grimmett, 1989; Stauffer and Aharony, 1994). For the purposes of this paper, there is no obvious preference for one or the other (and bond percolation can always be reformulated as a site model). An early application of percolation theory to technological change can be found in Cohendet and Zuscovitch (1982). David and Foray (1994) applied a hybrid site and bond percolation model to the standardization and diffusion problem in electronic data interchange networks. Some recent applications of percolation theory to social science problems include Solomon et al. (1999), Goldenberg et al. (2000), Gupta and Stauffer (2000) and Huang (2000).

[7] For bond percolation it can be proven that P_c is exactly 1/2. For site percolation it has been numerically established to be around 0.59.

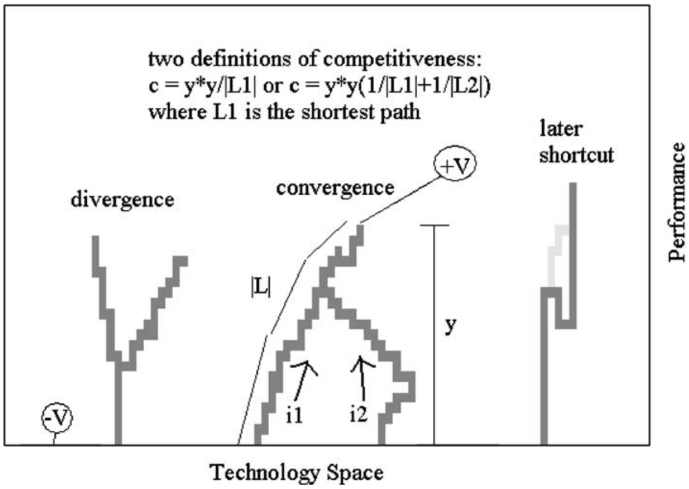

Fig. 5. Convergence, divergence and shortcuts, and two methods of defining a technology's competitiveness.

spanning the lattice from the bottom edge to the top one. This will increase rapidly and non-linearly in the neighbourhood of P_c. A metaphor that may help to sharpen intuition is to regard rain falling on a yard as a percolation problem. After only a bit of rain the yard consists of islands of wetness surrounded by dry pavement. After more rain has fallen the yard suddenly flips to being islands of dryness surrounded by wetness.

Suppose R&D is being undertaken in a region of our space such as Fig. 4. We imagine the search for new technologies to be akin to firing a *blunderbuss*, i.e. it cannot be targeted too closely to where we want to go, but instead its shot is smeared out over a finite region such as this one. Then there will be a probability P, which will be a function of the intensity of search, that any site in the region will be hit. At this point, there are two different perspectives we can adopt about whether a given site represents successfully uncovered technological knowledge (it does not represent an operational technology until it is connected by a path of such sites to the bottom line):

1. The *social construction of technology (SCT)* perspective says that any site we try is valid technological knowledge that can potentially be incorporated into a viable technology. Thus, in this case, a tried site will immediately become occupied and coloured red. The paths that result from innovative search will be pure accidents of history.
2. The alternative *technological determinism (TD)* perspective says that a tested site only represents true technological knowledge if it accords with the a priori underlying laws of nature. Thus, when we try a site, we must then test whether it is technologically valid. If it is, we fill it, if not, we leave it blank. This is a bit like playing the game minesweeper. The paths that result will be a selection from the technologically possible ones.

The laws of nature can be represented by creating a prepercolation on the lattice with some probability, q. If $q < P_c$, then there will only be finite connected sets (clusters) and technological change will eventually come to an end. If, however, nature is so bountiful that $q > P_c$, then there will potentially be infinite unbounded paths of innovation.[8] And the larger q, the denser the network of potentially viable technologies will be. Social construction of technology results from technological determinism in the limiting case $q = 1$.

Fig. 5 shows how connected paths may represent some relevant technological phenomena. First, any connected path beginning on the bottom line can be thought of as a natural trajectory. On the left, we see two trajectories diverging from a common origin. In the middle, we see technological convergence (e.g. the convergence of mechanical and electronic

[8] If our technological space (the horizontal axis) is bounded, this will not be true with probability 1, however.

technologies to mechatronics, or optical and mechanical technologies to optronics). While the purely technological performance characteristics of an operational site are measured by its height above the baseline, its economically relevant *technological competitiveness* can be measured in different ways. The point of introducing a separate technological competitiveness is to reflect the ease of realization (related to cost) of a given level of technological performance and allow subsequent incremental innovations to operate. Additionally, we may want the extent of parallelism in the realization of a technology to be counted as an advantage. Thus, I propose two separate measures of competitiveness, both based on path length (if L is a path then let $|L|$ be its length. The first measure is:

$$c_1 = \frac{y^2}{|L_s|}$$

where L_s is the shortest path connecting the site to the baseline. If this path is simply a straight vertical line, then $c_1 = y$. The more indirect the path, the more the competitiveness is diminished. The second measure corresponds to the current that would be extracted at the site if we apply a 1 V potential difference between the site and the baseline and set the resistance of a single lattice nearest-neighbour link to one. If two paths L_1 and L_2 converge at a site, then:

$$c_2 = y^2 \left(\frac{1}{|L_1|} + \frac{1}{|L_2|} \right)$$

For more complicated connections Kirchhoff's laws have to be applied.

A relevant technological analogy would be the different generations of microprocessors. While each generation represents a certain gain in performance, it usually comes at a certain price. However, over time that price declines as learning takes place in the production and design of the product. This can be captured in a natural way in our framework by allowing subsequent shortcuts (which we identify with *incremental* innovation) to reduce the length of the connecting base of a site (rightmost in Fig. 5). Thus, we will allow innovation to take place both ahead and behind the current best practice frontier, so that radical and incremental innovation take place simultaneously. The best practice frontier at any given time, shown in Fig. 6, consists of those operational (i.e. discovered and connected) sites lying highest above each point of the baseline. Also shown are *inventions*, which are viable discovered but not yet operational (i.e. connected) sites lying just above the frontier, and *scientific discoveries and fictions* (the latter possibly

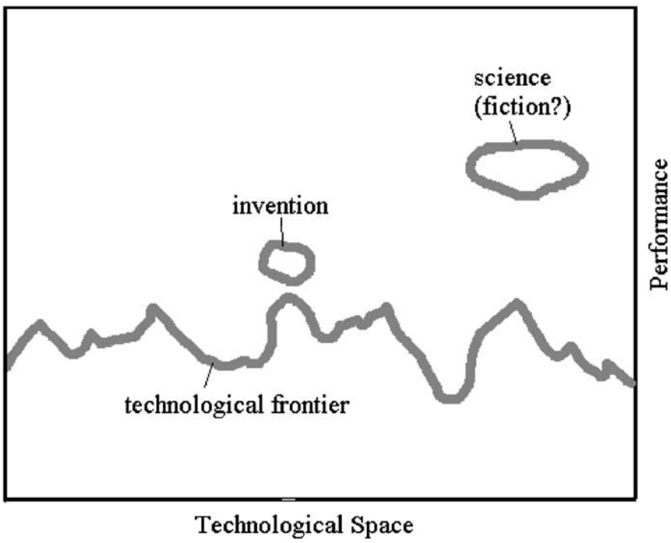

Fig. 6. Near disjoint regions represent inventions, far off discoveries science, and clusters that can never be connected to the baseline science "fictions".

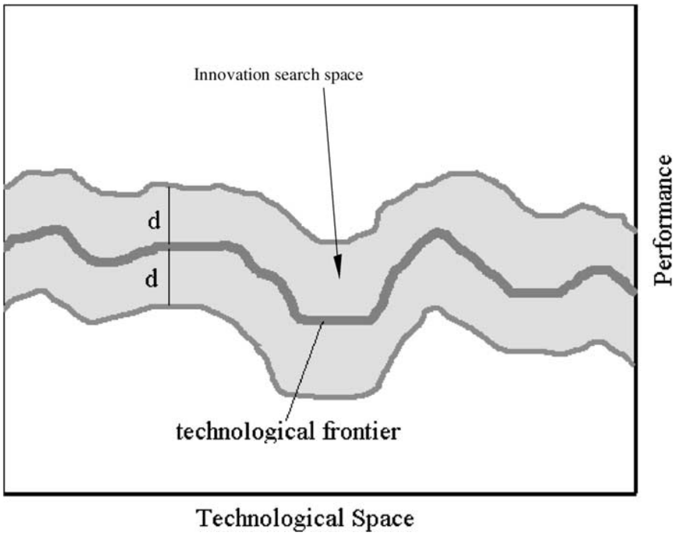

Fig. 7. New innovations are generated with probability P in a region d units above and below the technological frontier.

never being connectable to the operational network), which lie considerably above the frontier.

Consistent with our blunderbuss vision of the search process, we allow innovation to take place in a square region of side length d_0 centred around each point on the frontier. The union of these regions creates a band of innovative percolation extending ahead and behind of the frontier (Fig. 7). Within this region, new sites will be tested at random with some probability P. A discovered site of course need not connect immediately with the operational network. It is this fact that permits innovations of variable length (as measured by the jump in y they entail) to occur spontaneously. Thus, we obtain a natural explanation of innovation clustering (but of the random kind), as shown in Fig. 8. This happens when a disjoint extended network of

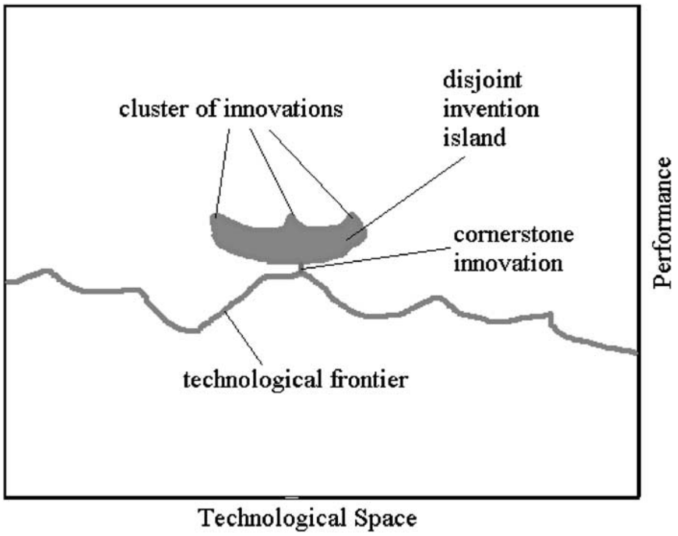

Fig. 8. A cluster of simultaneous invention occurs when a disjoint island of invention is suddenly joined to the frontier by a single ‘cornerstone’ innovation.

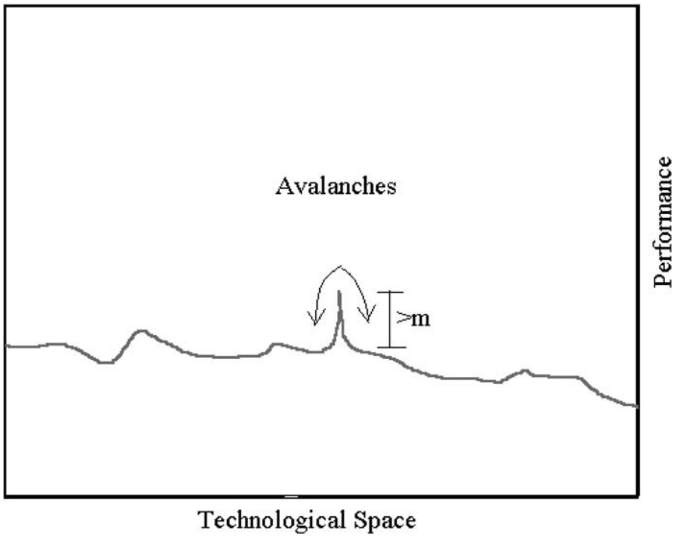

Fig. 9. Peaks above a certain threshold height trigger innovation attempts in neighbouring columns without themselves losing height. SOC?

discovered but not yet operational sites is finally connected to the technological frontier.

The final element in our vision consists of an *avalanche* mechanism. The intuition behind this is that extreme advances in a technology can spill over to adjacent technologies, which in turn can spill over to their neighbours. This differs from normal technological interrelatedness, whereby advances at one point in the frontier induce a probability of making discoveries in neighbouring sites. We posit that when a technological promontory such as shown in Fig. 9 becomes so extreme (e.g. sticks out above its neighbours on the frontier by more than m), a breakthrough mechanism is triggered, so that neighbouring sites just above the frontier are also immediately tried. In contrast to the analogous mechanism in sandpile models of *self-organized criticality*,[9] however, the performance of the original site is not diminished when it initiates spillovers, reflecting the fact that knowledge is not lost in one area when it is acquired in a related one. In the SCT case avalanches will always propagate at least one site, since every site tested is accepted. In the TD case, an avalanche may be stopped in its tracks because there is always a chance that a tested site will be rejected.

1.5. Conclusions and prospects for further research

The mosaic/avalanche model, as sketched above, incorporates the following virtues into one relatively homogeneous framework while making a minimal number of assumptions regarding the structure of the innovation process:

1. By invoking elemental innovation steps on a lattice or other graph, complex structures of technological advance and technological interrelatedness can be generated endogenously, dispensing with ad hoc assumptions about technological paradigms or production possibility frontiers. The ambiguity of invention and innovation datings is a natural feature of the mosaic perspective, but one amenable to statistical analysis.
2. The fractal structure of the innovation record (i.e. a distribution of innovation types from small, incremental innovations to radical or paradigmatic ones) results naturally from the percolation-induced clustering and the manner in which nature's prepercolation is uncovered by the R&D search process. I conjecture that the resulting distribution of

[9] Cf. Bak et al. (1987); Bak (1996). For the relationship between SOC and percolation see Grassberger and Zhang (1996).

innovation sizes and the time pattern of innovation activity will be highly skewed. A complete numerical implementation of this model will have to be undertaken to verify this conjecture.
3. Incremental innovation emerges as a by-product of the R&D process (given that this is necessarily also backward-looking from the perspective of the technological frontier) as paths to operational lattice sites shorten over time or alternative routes are created.
4. The invention/innovation distinction emerges naturally based on whether a discovered lattice site is disconnected from the technological frontier or not.
5. Science and science fiction can be conceived as extensions of the R&D process resulting from search with higher values of d. Science fictions are clusters of technologically imaginable sites not connectable in principle to the baseline, science ones that are (Fig. 6).
6. The avalanche mechanism, making use as it does of non-linear thresholds, translates the uneven advance of the technological frontier into bursts of technological spillovers into related areas, and into waves of diffusion.

What do we gain by invoking this framework? First, it allows macroinnovations to emerge spontaneously from the steady stream of small, elemental innovations and discoveries. We do not need a separate theory of each. And the macroinnovations will naturally fall into some distribution of sizes, highly skewed and unbounded from above.

How this mechanism proceeds in time is something that can only be understood by actually simulating such a percolation structure and relating it to existing results in the literature. This is something, I hope to do in the near future. To obtain an overall picture of economic growth and sectoral change, it will be necessary to round out the model with a component relating innovations in each segment of the technology space to diffusion pulses in investment and utilization of the corresponding goods and services in the real economy. I suspect that an appropriate model of such sectoral diffusion pulses would also benefit from the avalanche perspective. This would allow us to endogenize the range of intersectoral spillovers in economic activity initiated by breakthroughs in a few specific sectors. While such an approach might not provide a detailed picture of the emergence of the current 'new economy,' it might give us a perspective on when we are entitled to speak of such things, how often they occur, and how long they last.

References

Aghion, P., Howitt, P., 1992. A model of growth through creative destruction. Econometrica 60, 323–351.

Arrow, K., 1962. The economic implications of learning by doing. Review of Economic Studies 29, 155–173.

Atkinson, B., Stiglitz, J.E., 1969. A new view of technological change. Economic Journal 79, 573–578.

Bak, P., 1996. How Nature Works: The Science of Self-Organized Criticality. Springer, New York.

Bak, P., Tang, C., Wiesenfeld, K., 1987. Self-organized criticality: an explanation of $1/f$ noise. Physical Review A 38, 364–374.

Baker, R., 1976. New and Improved: Inventors and Inventions That Have Changed the Modern World. British Museum Publications, London.

Brockhoff, K., 1972. Zur Dynamik des Technischen Fortschritts. Zeitschrift für Betriebswirtschaftslehre, 42, 283.

Cohendet, P., Zuscovitch, E., 1982. Progrès Technique et Percolation. BETA, Universite de Strasbourg Working Paper.

Clark, J., Freeman, C., Soete, L., 1981. Long waves, inventions, and innovations. Futures 13, 308–322.

David, P.A., Foray, D., 1994. Percolation structures: markov random fields and the economics of EDI standard diffusion. In: Pogorel (Ed.), Global Telecommunications Strategies and Technological Changes. North Holland, Amsterdam.

Dosi, G., 1982. Technological paradigms and technological trajectories. Research Policy 11, 147–162.

Goldenberg, J., Libari, B., Solomon, S., Stauffer, D., 2000. Marketing percolation. Physica A 284, 335–347.

Gordon, R.J., 2000. Does the new economy measure up to the great inventions of the past? Journal of Economic Perspectives 14 (4), 49–74.

Grassberger, P., Zhang, Y., 1996. Self-organized formulation of standard percolation phenomena. Physica A 224, 169–179.

Grimmett, G., 1989. Percolation. Springer, New York.

Grübler, A., 1990. The Rise and Decline of Infrastructures. Dynamics of Evolution and Technological Change in Transport. Physical-Verlag, Heidelberg.

Grübler, A., 1998. Technology and Global Change. Cambridge University Press, Cambridge.

Gupta, A.K., Stauffer, D., 2000. Social percolation on inhomogeneous networks. International Journal of Modern Physics C 11 (4), 695–705.

Haustein, H.D., Neuwirth, E., 1982. Long waves in world industrial production, energy consumption, innovations, inventions, and patents and their identification by spectral analysis. Technological Forecasting and Social Change 22, 53–89.

Helpman, E. (Ed.), 1998. General Purpose Technologies. MIT Press, Cambridge, MA.

Huang, Z.F., 2000. Self-organized model of information spread in financial markets. European Physical Review B 16 (2), 379–385.

Jewkes, J., Sawers, D., Stillermann, R., 1958. The Sources of Invention. Macmillan, London.

Kleinknecht, A., 1981. Observations on the Schumpeterian swarming of innovations. Futures 13, 293–307.

Kleinknecht, A., 1987. Innovation Patterns in Crisis and Prosperity: Schumpeter's Long Cycle Reconsidered. Macmillan, London.

Kleinknecht, A., 1990a. Are there Schumpeterian waves of innovations? Cambridge Journal of Economics 14, 81–92.

Kleinknecht, A., 1990b. Schumpeterian waves of innovation? Summarizing the evidence. In: Vasko, T., Ayres, R., Fontvieille, L. (Eds.), Life Cycles and Long Waves. Springer, Berlin.

Kuznets, S., 1940. Schumpeter's business cycles. American Economic Review 30, 257–271.

Malerba, F., Nelson, R., Orsenigo, L., Winter, S., 1999. A history-friendly model of the evolution of the computer industry. Industrial and Corporate Change 8, 3–40.

Mensch, G.O., 1975. Das Technologische Pat. Umschau, Frankfurt.

Nakicenovic, N., 1987. Technological substitution and long waves in the USA. In: Vasko, T. (Ed.), The Long-Wave Debate. Springer, Berlin.

Nelson, R.R., Winter, S.G., 1977. In search of a useful theory of innovation. Research Policy 6, 36–76.

Nelson, R.R., Winter, S.G., 1982. An Evolutionary Theory of Economic Change. Harvard University Press, Cambridge, MA.

Sahal, D., 1974. Generalized Poisson and related models of technological innovation. Technological Forecasting and Social Change 6, 403–436.

Silverberg, G., Dosi, G., Orsenigo, L., 1988. Innovation, diversity and diffusion: a self-organisation model. Economic Journal 98, 1032–1054.

Silverberg, G., Lehnert, D., 1993. Long waves and evolutionary chaos in a simple Schumpeterian model of embodied technical change. Structural Change and Economic Dynamics 4, 9–37.

Silverberg, G., Verspagen, B., 2000. Breaking the Waves: A Poisson Regression Approach to Schumpeterian Clustering of Innovations, ECIS and MERIT Working Paper, available from http://meritbbs.unimaas.nl. Forthcoming in Cambridge Journal of Economics.

Solomon, S., Weisbuch, G., de Arcangelis, L., Jan, N., Stauffer, D., 1999. Social percolation models. Physica A 277, 239–247.

Solomou, S., 1986. Innovation clusters and Kondratieff long waves in economic growth. Cambridge Journal of Economics 10, 101–112.

Solow, R., 2000. Growth Theory: An Exposition. Oxford University Press, Oxford.

Stauffer, D., Aharony, A., 1994. Introduction to Percolation Theory. Taylor & Francis, London.

Trajtenberg, M., 1990. A penny for your quotes: patent citations and the value of innovations. Rand Journal of Economics 21, 172–187.

Vega-Redondo, F., 1994. Technological change and path dependence: a co-evolutionary model on a directed graph. Journal of Evolutionary Economics 4, 59–80.

van Duijn, J.J., 1983. The Long Wave in Economic Life. Allen & Unwin, Boston.

van Raan, A.F.J., 1990. Fractal dimension of co-citations. Nature 347, 626.