

MERIT-Infonomics Research Memorandum series

*Self-organization of R&D search in
complex technology spaces*

Gerald Silverberg & Bart Verspagen

2005-017



*MERIT – Maastricht Economic Research
Institute on Innovation and Technology*

PO Box 616
6200 MD Maastricht
The Netherlands
T: +31 43 3883875
F: +31 43 3884905

<http://www.merit.unimaas.nl>
e-mail: secr-merit@merit.unimaas.nl



International Institute of Infonomics

c/o Maastricht University
PO Box 616
6200 MD Maastricht
The Netherlands
T: +31 43 388 3875
F: +31 45 388 4905

<http://www.infonomics.nl>
e-mail: secr@infonomics.nl

Self-organization of R&D search in complex technology spaces

Gerald Silverberg^{*} and Bart Verspagen^{**}

Paper prepared for the WEHIA 2005 Conference, Exeter University

May 2005

Abstract

We extend an earlier model of innovation dynamics based on invasive percolation by adding endogenous R&D search by economically motivated firms. The $\{0,1\}$ seeding of the technology lattice is now replaced by draws from a lognormal distribution for technology ‘difficulty’. Firms are rewarded for successful innovations by increases in their R&D budget. We compare two regimes. In the first, firms are fixed in a region of technology space. In the second, they can change their location by myopically comparing progress in their local neighborhoods and probabilistically moving to the region with the highest recent progress. We call this the moving or self-organizational regime. We find that as the mean and standard deviation of the lognormal distribution are varied, the relative rates of aggregate innovation switches between the two regimes. The SO regime has higher innovation rates, other things being equal, for lower means or higher standard deviations of the lognormal distribution. This results holds for increasing size of the search radius. The clustering of firms in the SO regime grows rapidly and fluctuates in a complex way around a high value which increases with the search radius. We also investigate the size distributions of the innovations generated in each regime. In the fixed one, the distribution is approximately lognormal and certainly not fat tailed. In the SO regime, the distributions are radically different. They are much more highly right skewed and show scaling over at least two decades with a slope of almost exactly one, independently of parameter settings. Thus we argue that firm self-organization leads to self-organized criticality.

Keywords: innovation, percolation, search, technological change, R&D, clustering, self-organized criticality

JEL codes: C15, C63, D83, O31

^{*} MERIT, Maastricht University, P.O. Box 616, NL-6200 MD Maastricht, The Netherlands, Tel. +31-43-3883868, Fax +31-43-3216518. Email: gerald.silverberg@merit.unimaas.nl.

^{**} ECIS, Eindhoven University of Technology, P.O. Box 513, NL-5600 MB Eindhoven, The Netherlands, Tel. +31-40-2475613, Fax +31-40-2474646. Email: b.verspagen@tm.tue.nl.

1. Introduction

The paradoxical characteristic of innovations is that their nature, significance and date of arrival are intrinsically unknowable in advance – if we knew them in detail, then the innovation would in effect have already happened. It is this property of intrinsic uncertainty that makes innovations so central and different from other factors in the theory of long-term economic evolution. On the assumption of total ignorance and independence, the natural approach would be to regard discrete innovations as generated by a simple stochastic point process such as the time-homogeneous Poisson. At the same time we know that technologies are not picked out of a hat at random times in random orders – to some extent there is a logical order in which they can be discovered, and they build on each other. Modern computers could not exist without a mastery of electronics (although Babbage tried and failed to make a purely mechanical one in the 19th century), electronics without a mastery of electricity, and electricity without the metallurgical skills necessary to make wires. Thus we shall argue in the following, based both on empirical evidence and a theoretical model, that the innovation process, while highly uncertain and stochastic, is still more structured in important respects than such a null hypothesis would suggest.

While the study of the statistical properties of the innovation process is scientifically interesting in its own right, it also has important implications for economic theory and innovation management. If innovations are drawn from a highly skewed and even infinite variance process (Pareto), then economic growth may be even more erratic than if they are of constant ‘size’ but generated by a Poisson process (see Sornette and Zajdenweber 1999 on the former case, Silverberg and Lehnert 1996 on the latter). And if they are drawn from an infinite variance and even infinite mean process, than R&D risk management and portfolio policy are confronted with such high risk that the standard tools of capital asset management theory are inapplicable (Scherer and Harhoff 2000).

Three earlier papers (Silverberg, 2002; Silverberg and Verspagen, 2005 and 2003a) examined a model based on percolation theory that takes several stylized facts about innovation into account. In Silverberg and Verspagen (2003a), we summarized the stylized facts about innovation under three types of ‘clustering’. First (major) innovations tend to be clustered in time: they “are not evenly distributed in time, but ... on the contrary they tend to cluster, to come about in bunches, simply because first some, and then most firms follow in the wake of successful innovation” (Schumpeter, 1939, p. 75, see also, Silverberg and Verspagen 2003b). Second, innovations are clustered in ‘technology space’ (a concept that will be operationalized in terms of our model). In the economic literature analyzing the development of technological change there are numerous suggestions that the innovative process follows relatively ordered pathways that can be measured ex post in technology characteristics space. Examples of propositions in this direction are Nelson and Winter’s (1977) *natural trajectories*, Sahal’s (1981) *technological guideposts*, and Dosi’s (1982) *technological paradigms*. Empirically oriented contributions that illustrate the point are, e.g., Foray and Grübler (1990), Saviotti (1996) and Frenken and Leydesdorf (2000). Third, recent literature such as Scherer (1998), Harhoff, Narin, Scherer and Vopel (1999), and Scherer, Harhoff and Kukies (2000),

suggests that the distribution of innovation sizes, as captured by some measure of economic returns to R&D investment, is highly skewed, with most innovations having low or negative returns but with a highly skewed tail extending into the region of extremely high rates of return. The same tendency can be observed using the data compiled by Trajtenberg (1990) for the ‘value’ of patents proxied by the number of patent citations. These data suggest that the distribution of innovations may follow a power law, at least in the tails (cf. Silverberg and Verspagen 2004).

Our earlier model was aimed at explaining these stylized facts from the simplest possible assumptions regarding the nature of the innovative process. Thus we abstracted from any economically motivated, active search process. The earlier model did have ‘research’ in the form of search in technology space, but the efforts put into this process were completely exogenous, both in terms of their size (“amount of expenditures”) and direction (selection of promising avenues for research). It is the aim of this paper to introduce a more elaborate, economically motivated basis for technological search, and to investigate its implications for the stylized facts of innovation that our model addresses.

The process of technological search will be motivated by two crucial economic factors. The first concerns the way in which firms¹ select the parts of technological space they want to search. Each of the firms in the model can only address a (small) part of the technological space (we exogenously set a search radius for all firms), but we now allow the firms a certain latitude in determining the search region themselves. This implies that, contrary to our earlier model, search may now be concentrated in selected neighbourhoods, when firms collectively decide to locate their R&D activities there. Other parts of technological space may conversely be abandoned. Although firms can, in principle, move freely through technology space, the model does assume that there are costs associated with this. Although we do not model such costs explicitly, we do assume that relocation of search in technology space is the result of two counteracting tendencies. On the one hand, firms want to move to the places where technological opportunities seem to be largest, but, on the other hand, they also tend to stick to locations that they know from previous experience (something Nelson and Winter 1977 termed ‘local search’).

The second assumption consists of introducing a positive feedback resulting from successful innovations. Firms that realize innovations will generate resources to invest in new R&D efforts in the next period proportional to their success. Although in the real world financial markets may also reward (economically) unproven innovations (venture capital), it is fair to say that in many cases ‘success breeds success’ in innovative activity. This is what Winter (1984) has called a “routinized” regime of innovation

In Section 2 we formally describe this new part of the model, as well as the basic structure retained from the previous version. This section also discusses our approach to modeling technology space, based on percolation theory. Section 3 describes some of the results of the model. Although the model has relatively few parameters, the total parameter space is

¹ We use this term to describe the abstract agents that operate in our technological space, but our model only addresses the R&D function of these agents. We ignore various aspects of firm behaviour that would normally be the subject of economic models, such as production, sales, investment and firm growth.

large, and we have only just begun to analyze it in a systematic way. We focus here on comparing the fixed firm regime with the self-organizational one with moving firms with respect to several indicators. In particular the size distribution of innovations is shown to change radically between the two regimes. The results are summarized and directions for further analysis are outlined in the concluding section 4.

2. The model

2.1. Technology space

Our probabilistic model of innovation is an elaboration of the model in Silverberg and Verspagen (2005). As in the original model, the present model hinges on two essential properties. First, technologies constitute a discrete topological space with a neighborhood structure reflecting their technological interrelatedness, and second, over time technologies can only come ‘online’ by becoming contiguous to previously operational technologies, even if R&D search takes place in a more ‘leapfrogging’ or farsighted manner.

For simplicity, consider a lattice, unbounded in the vertical dimension, anchored on a baseline (or space), with periodic boundary conditions in the horizontal dimension. The horizontal space represents the universe of technological niches, with neighboring sites being closely related. While the technology space is represented here and in the following as one-dimensional, it can easily be generalized to higher dimensions or different topologies. The vertical axis measures an indicator of performance intrinsic to that technology and could also be conceived as multidimensional. For simplicity we will restrict ourselves to a two-dimensional lattice in the following.

A lattice site \mathbf{a}_{ij} can be in one of three states: 0 or not yet discovered, 1 discovered but not yet viable, and 2, discovered and viable. Compared to the original version of the model, the present model has one fewer state. Whereas Silverberg and Verspagen (2005) distinguished between sites that were potentially technologically possible and those that were excluded by the laws of nature, the present model does not make this distinction. No sites are excluded by the laws of nature, but sites do differ with regard to the difficulty of discovering them. Some sites are easy to discover, others more difficult if not nearly impossible.

A site may become discovered, i.e., move from state 0 to 1, by means of repeated effort by the agents searching in its region of the technology lattice. Agents invest R&D with the aim of discovering the site. Each site on the lattice is randomly initialized with a ‘resistance’ value, which we denote by \mathbf{q}_{ij} . This value is drawn from a lognormal distribution with mean $\langle q \rangle$ and standard deviation σ . When an agent invests \mathbf{b} units of R&D with the aim of discovering the site, the resistance value is diminished according to the following process:

$$\mathbf{q}_{ij,t+1} = \mathbf{q}_{ij,t} - \mathbf{b} \omega,$$

where ω is a random variable drawn from a uniform distribution on $[0, 1)$ (this represents the stochastic nature of the R&D process), and the subscripts t and $t+1$ denote the value of the

resistance factor before and after the agent's R&D project. A site becomes an invention with state 1 when \mathbf{q}_{ij} becomes zero or negative.

A site moves from state 1 to 2, i.e., from discovered to viable (invented to innovated), when there exists a contiguous path of discovered or viable sites connecting it to the baseline. The neighborhood we shall use is the von Neumann one of the four sites top, bottom, right and left $\{\mathbf{a}_{i\pm 1,j}, \mathbf{a}_{i,j\pm 1}\}$, with periodic boundary conditions horizontally. The intuition here is that a discovered technology only becomes viable or operational when it can draw on an unbroken chain of supporting technologies already in use. Until such a chain is completed, the technology is still considered to be under development – it is still an invention, not an innovation.

At any point in time t a best-practice frontier can be defined consisting of the highest sites in state 2 for each baseline column (of which there are N_c):

$$BPF(t) = \{(i, j(i)), i = 1, N_c\}, \text{ where } j(i) = (\max j \mid a_{i,j} = 2).$$

If there is no viable site in column i we set $j(i) = -1$.

2.2. The firm-based R&D process

An innovation is defined as a jump in the BPF in the vertical dimension in a single time period. The size of the innovation, denoted by s , is defined as the number of levels (rows) that the frontier has moved upward in an individual column. The payoff of an innovation will be assumed proportional to s . The firm's R&D budget consists of a fixed part, which is equal for all firms and all periods, and a part deriving from the payoffs to previous innovations. This is formulated as follows:

$$\mathbf{B}_t = \beta + \sum_k s_{k,t-1} \pi,$$

where \mathbf{B}_t is the total R&D budget that the firm spends in period t , β is the base part of the R&D budget, $s_{k,t-1}$ is the size of the firm's innovation (if any) in column k that the firm made when it last had a turn at performing R&D, and π is the payoff per 'unit' of innovation, and k is summed over all columns in the firm's search neighborhood. If the firm was unsuccessful in its previous R&D round (no innovations were realized), its R&D budget falls to β .

The firm operates from a single position (site) in the lattice, which is updated periodically. Its search neighborhood consists of a (diamond-shaped) neighborhood of radius m in the 'Manhattan' metric induced by the neighborhood relation centered around the firm's present site. This neighborhood contains $2m(m+1)$ points. The R&D budget \mathbf{B} is distributed equally over all sites in the neighborhood, irrespective of whether or not they have already been discovered. Thus, the R&D budget available for a single site is $\mathbf{b} = \mathbf{B} / 2m(m+1)$.

Since R&D is aimed at the local environment of the firm, an important element of the model is how firms determine their position in technology space from which R&D is undertaken. Just before a firm starts an R&D cycle to search its local neighborhood, it may move its position on the lattice. We now differentiate between two distinct regimes of firm behavior. If we are in the *fixed firm* regime, the firm simply moves vertically in its present column to the present level of the BPF that it inherits from the last innovator. In the *moving firm* regime, the

firm also first moved vertically to the BPF, but then also examines the heights of the points on the BPF in the columns within a radius of m from its current column. It then decides probabilistically whether to move to one of these columns and then to the associated point on the BPF according to the following calculation. For each column j in its m -radius column neighborhood, we calculate a value $u_j = e^{h_j - h_i}$, where h_j is equal to column j 's BPF technological level (i.e., height on the lattice) and h_i is the current level of the firm. We include the possibility that the firm stays where it is. We calculate the probability p_{ij} that the firm will move to column j on the BPF from its current column i (or stay where it is, with probability p_{ii}) as

$$p_{ij} = \frac{u_j}{U}, \quad U = \sum_j u_j.$$

In this scenario, firms may wind up concentrated in a small number of regions and leave other regions of technology space completely unpopulated.

In both scenarios, we initially populate every column on the lattice with exactly one firm. We also assume that before a firm comes up 'to bat' in an innovation round, it first advances to the BPF of its current technology column. Thus firms benefit from an interfirm technology externality after one innovation period.

After the firm moves, it performs R&D in its (new) local neighborhood. Payoffs are awarded (added to the R&D budget of next period) after this R&D process, and the global BPF is updated. The fact that search continues to take place below the BPF means that the path connecting sites on the frontier to the baseline may shorten over time as 'shortcuts' and missing links are discovered. We regard this as one way of representing *incremental* innovation, but we will not deal with this aspect here.

2.3. Innovation dynamics

A discovered site need not initially connect up 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 Figure 1. This happens when a disjoint extended network of discovered but not yet operational sites is finally connected to the technological frontier, and/or when an 'overhanging cliff' advances laterally, pulling up the BPF at neighboring sites by increments that can be much larger than m , the search radius, and are in fact unbounded from above.

The basic unit of time in the model is one R&D cycle by one firm. At the beginning of each cycle, a single firm is drawn randomly from the population. By convention, we set the number of firms equal to the number of columns in the lattice, and at the beginning of each run, one firm is placed at the baseline of each column of the lattice.

The computer implementation of the model is illustrated in Figure 2, which is a screen shot of the user interface in interactive mode. The rectangle on the upper right shows the state of the lattice at this point of time. Grey dots represent undiscovered lattice sites (state 0), with

darker colors indicating higher values of q . Green sites represent discovered but not yet viable sites (state 1), and yellow sites are viable technologies (state 2), i.e., discovered and connected to the baseline. The red line represents the best-practice frontier (BPF) around which search is taking place in a band of radius 6. A typical pattern is shown of ‘overhanging cliffs’ of yellow sites on the left and in the middle.

3. Simulation results

In the following, we will compare the behavior of the model with firms either evenly distributed over columns and fixed or self-organizing and locally moving to more attractive sites in the manner described above. Figure 3 shows the average rate of innovation generated per period (defined as the area swept out by the BPF per period) in runs of 15,000 periods for the two regimes as a function of the mean $\langle q \rangle$ of the lognormal distribution for various values of the search radius m (ten runs per parameter value are generated with different random seeds, with $\sigma=2$ and $\pi=1$). A typical pattern is evident in all panels. While the innovation rate falls off weakly with $\langle q \rangle$ for fixed firms, it declines much more rapidly for moving ones, but from a higher initial level. This results in a pattern of innovation rate switching, with self-organizing firms being more ‘innovative’ for low values of $\langle q \rangle$ but much inferior for higher values. Evidently, when innovations are generally more difficult to make, the herding behavior of the self-organizing firms is counterproductive. When the landscape is generally ‘thinner’, in contrast, the herding behavior rather successfully identifies and exploits the avenues of easy progress in the landscape. Not unexpectedly the rate of innovation of both regimes increases with the search radius m , a result already known from the original percolation model.

A rather similar result holds for fixed $\langle q \rangle$ but variable σ , as shown in Figure 4 for $\langle q \rangle=2$, $m=3$, and five runs per data point. Here the self-organizing regime proves to be superior for larger values of σ and inferior for smaller ones. The crossover point is just visible on the right hand side. In this model σ plays a role a bit like the percolation probability in the original model, with higher values of σ indicating a rougher landscape. The crossover point shifts downward as we decrease $\langle q \rangle$, as can be seen in Figure 5 for $\langle q \rangle=0.4$ (only one run shown per data point).

A similar crossover occurs if we vary the payoff rate π to successful innovations (Figure 6, with only one run per data point). Larger payoffs evidently favor the self-organized regime, so that if firms receive greater individual rewards for innovation, it pays both individually and collectively for them to choose their location in technology space, even if this carries the danger of duplication of effort.

Figure 7 displays how the firms cluster over time in the self-organized regime by plotting the clustering index d as a function of time, where d is defined as

$$d = \sum_{i=1}^{N_c} n_i^2 - n$$

with n_i the number of firms in the i th technology column and n the total number of firms (which by assumption is equal to the total number of columns N_c). From an initial value of zero for uniformly distributed firms it climbs to a high ‘steady state’ value within about 4000 periods but thereafter displays fluctuations with more ‘long memory’ than just a simple stochastic process would suggest. Figure 8 shows the final values of d after 15,000 periods for different values of $\langle q \rangle$ and search radius m . The average values over ten runs are rather stable for small values of m , but increase nearly linearly with $\langle q \rangle$ for higher values of m . A trough around $\langle q \rangle \approx 0.6-0.9$ also appears to be present.

In Figure 9 we compare the innovation size distributions resulting from a run with fixed firms and the equivalent run with moving ones ($\langle q \rangle = 0.2$, $\sigma = 4$, $m = 3$, with the first 5,000 periods deleted to eliminate the effects of transients). These are Pareto plots showing the number of observations greater than or equal to a certain size, on a double-log scale. Pareto-distributed observations will fall on a straight line in such a plot. For fixed firms we observe a definite curvature (indicating that, while the distribution is highly skewed, it is not fat-tailed and more resembles a lognormal distribution). In contrast, in the moving firm regime we observe a striking region of linearity over at least two decades of observations. The slope of this curve is almost exactly -1, as indicated by doing an ordinary least squares fit to the observed curve. The tail index α can more properly be calculated by making use of its maximum likelihood estimator due to Hill (1975). This is defined using the largest k values of the rank order statistics of the observations as follows:

$$1/\alpha(k, n) = H(k, n) = \frac{1}{k} \sum_{i=1}^k (\ln X_{[i]} - \ln X_{[k+1]}).$$

Figure 10 shows the Hill plots of the tail index for the fixed and the moving firm cases, respectively. The monotonic decline of the Hill plot for the former indicates that it is not fat tailed. The near-perfect plateau at a value of α of almost exactly one between the 100th and nearly the 10,000th largest observations is striking. However, there are significant and systematic deviations from linearity for both the smallest and very largest innovations.

What is most remarkable about this scaling behavior in the self-organized regime is that it appears to be insensitive to the values of the principal parameters. The panels of Figure 11 and 12 display a selection of Pareto plots for various constellations of parameter values. The general form remains unchanged: an intermediate region of very precise scaling behavior over two decades with tail index one. Because this regularity emerges without the need to tune the system to a critical value of an exogenous parameter, as would be the case in a pure percolation model, it appears to be an instance of self-organized criticality. This is perhaps not completely surprising, since in some respects our model resembles the well-known model of interface growth due to Sneppen (1992).² Why the tail index always assumes a value of one remains something of a mystery. However, it falls into the same ballpark as the empirical estimates of the tail index of monetary measures of the returns to innovation found in Silverberg and Verspagen (2004), where values near or just below one were observed. Of course our

² It differs from the Sneppen model in that interface growth (the advance of the BPF) takes place at sites selected by a local, probabilistic ‘extremal’ rule (firms move only locally and with a probability less than one to the most active previous sites) rather than employing straightforward extremal dynamics.

measure of innovation size does not map directly to any of the monetary measures employed in the empirical literature. Nevertheless, this congruence is intriguing.

Another example of scaling behavior emerges when we plot the average innovation rate against the variance of the BPF realized after 15,000 periods (the latter is a measure of the extent to which different technological categories evolve at different rates). There is a generally positive relationship between these two values for various values of $\langle q \rangle$ and randomizations (holding other parameters fixed), as one would expect, but no scaling (Figure 13). When we vary σ , however, leaving $\langle q \rangle$ fixed, this changes radically (Figure 14). Even in the case of fixed firms, where the range of variation is quite limited, scaling is evident. This is even more the case for moving firms, where the range of variation is quite broad. For comparison, if the individual column heights were evolving according to a random walk with drift (assumed identical and independently distributed across columns), then the variance would be a linear function of the mean rate of advance.

4. Conclusions and future research

In this paper we have introduced endogenous R&D search by economically motivated firms in a percolation model of innovation dynamics. A previous model without endogenous R&D search has already proven to be useful in explaining some of the stylized facts about innovation, specifically with regard to the temporal clustering of innovation and the skewed nature of innovation size distributions. The introduction of the ‘Toyota’ landscape in place of the binary percolated technology landscape does not seem to change the basic properties of the model in terms of its ability to generate these stylized facts. However, the ability of firms to move and redeploy their R&D efforts as a result of previous advances on the landscape introduces a quality of self-organization into the model. The most striking change in comparison with fixed firms is a much more highly skewed distribution of innovations and scaling over a considerable range, with a characteristic tail index of one. While the moving firms ostensibly display a higher degree of ‘rationality’ than the fixed firms, collectively this freedom is not always advantageous. In particular, for higher values of the average difficulty of the landscape ($\langle q \rangle$) or smoother landscapes (lower σ), or very low rewards to previous success (low payoff parameter π), it is collectively more sensible to retain the distributed search of fixed firms than try to exploit the ability of moving firms to independently focus on hot ‘lodes’ of technological richness (whether this is individually rational for firms is another question). When the world has the opposite character, however, for example when technologies vary greatly in the ease with which they can be developed, but overall are not too difficult to uncover, and previous success is sufficiently rewarded, self-organization yields significant payoffs in terms of more rapid rates of technical change. In this respect our model offers food for thought regarding the ‘rationality’ of such phenomena as the Internet bubble and gold rushes.

The rate at which firms actually change their locations and move toward more attractive sites can be tuned using an additional parameter β in the multinomial transition probabilities between the current column i and a neighboring column j :

$$p_{ij} = e^{\beta(h_j - h_i)} / \left[\sum_{k=-m}^m e^{\beta(h_k - h_i)} \right].$$

By varying β we can make firms more or less ‘rational’ or responsive to disparities in the state of neighboring columns in technology space. The parameter β would play a role similar to that in the Brock and Hommes (1998) model of routes to chaos in financial market with heterogeneous agents. What effect this might have in the context of our model is a subject for future research.

References

- Bak, P., 1996, *How Nature Works: The Science of Self-Organized Criticality*, New York: Springer-Verlag.
- Brock, W.A. and Hommes, C.H., 1998, ‘Heterogeneous beliefs and routes to chaos in a simple asset pricing model’, *Journal of Economic Dynamics and Control*, **22**: 1235-1274.
- Dosi, G., 1982, "Technological Paradigms and Technological Trajectories", *Research Policy*, **11**: 147-162.
- Foray, D. and Grübler, A. 1990, ‘Morphological analysis, diffusion and lock-out of technologies: ferrous casting in France and the FRG’, *Research Policy*, **19**: 535–550.
- Frenken, K and Leydesdorff, L., 2000. ‘Scaling trajectories in civil aircraft 1913–1997’, *Research Policy*, **29**: 331–348.
- Harhoff, D., Narin, F., Scherer, F. M. and Vopel, K., 1999, "Citation frequency and the value of patented inventions", *Review of Economics and Statistics*, **81**: 511-515.
- Hausman, J., B. H. Hall and Z. Griliches 1984. “Econometric Models for Count Data with an Application to the Patents-R&D Relationship.” *Econometrica*, **52**: 909-938.
- Hill, B. M., 1975, ‘A Simple General Approach to Inference about the Tails of a Distribution’, *The Annals of Statistics*, **3**: 1163-1174.
- Nelson, R.R. and Winter, S.G., 1977, “In Search of a Useful Theory of Innovation”, *Research Policy*, **6**: 36-76.
- Sahal, D., 1981, *Patterns of Technological Innovation*, New York: Addison-Wesley.
- Saviotti, P.P. 1996. *Technological Evolution, Variety and the Economy*, Cheltenham and Brookfield: Edward Elgar.
- Scherer, F. M., 1998, ‘The Size Distribution of Profits from Innovation’, *Annales d'Economie et de Statistique*, **49/50**: 495-516.
- Scherer, F.M. and Harhoff, D., 2000, “Technology Policy for a World of Skew-Distribution Outcomes”, *Research Policy*, **29**: 559-566.
- Scherer, F.M., Harhoff, D., and Kukies, J., 2000, “Uncertainty and the Size Distribution of Rewards from Innovation”, *Journal of Evolutionary Economics*, **10**: 175-200.
- Schumpeter, J. A. 1939. *Business Cycles: A theoretical, historical and statistical analysis of the capitalist process*. New York, McGraw-Hill (page numbers quoted in the text refer to the abridged version reprinted in 1989 by Porcupine Press, Philadelphia).

- Silverberg, G., 2002, "The Discrete Charm of the Bourgeoisie: Quantum and Continuous Perspectives on Innovation and Growth", *Research Policy*, **31**: 1275-1289.
- Silverberg, G. and Lehnert, D., 1996, 'Evolutionary Chaos: Growth Fluctuations in a Schumpeterian Model of Creative Destruction', in Barnett, W. A., Kirman, A. and Salmon, M., (eds), *Nonlinear Dynamics in Economics*, Cambridge: Cambridge University Press.
- Silverberg, G. and B. Verspagen 2003a, "Brewing the future: Stylized facts about innovation and their confrontation with a percolation model," Paper prepared for the EMAEE Conference, Augsburg, April 10-12, 2003, ECIS Working Paper 80.
- Silverberg, G. and B. Verspagen, 2003b, 'Breaking the Waves: a Poisson Regression Approach to Schumpeterian Clustering of Basic Innovations', *Cambridge Journal of Economics*, **27**: 671-693.
- Silverberg, G. and B. Verspagen, 2004, "The size distribution of innovations revisited: an application of extreme value statistics to citation and returns measures of patent significance", Maastricht: MERIT Research Memorandum 2004-021, forthcoming in *Journal of Econometrics*.
- Silverberg, G. and B. Verspagen, 2005, "A Percolation Model of Innovation in Complex Technology Spaces," *Journal of Economic Dynamics and Control*, **29**: 225-244.
- Sneppen, K., 1992, 'Self-Organized Pinning and Interface Growth in a Random Medium', *Physical Review Letters*, **69**: 3539-3542.
- Trajtenberg, M., 1990, "A Penny for your Quotes: Patent Citations and the Value of Innovations", *Rand Journal of Economics*, **21**, 172-187.
- Winter, S.G., 1984, 'Schumpeterian Competition in Alternative Technological Regimes', *Journal of Economic Behavior and Organization* **5**:287-320.




Figure 1 Clusters of innovations occur when disconnected islands of inventions are joined to the BPF by cornerstone innovations.

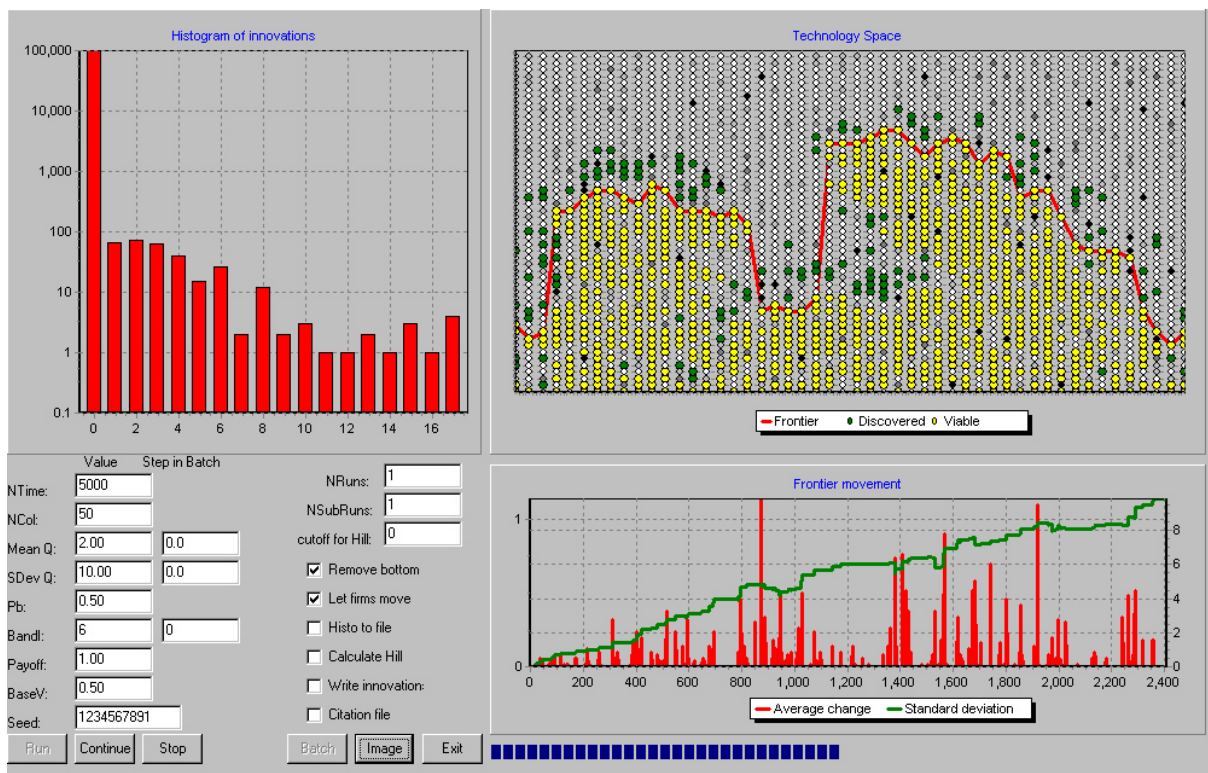


Figure 2 Screenshot of the computer implementation of the model.

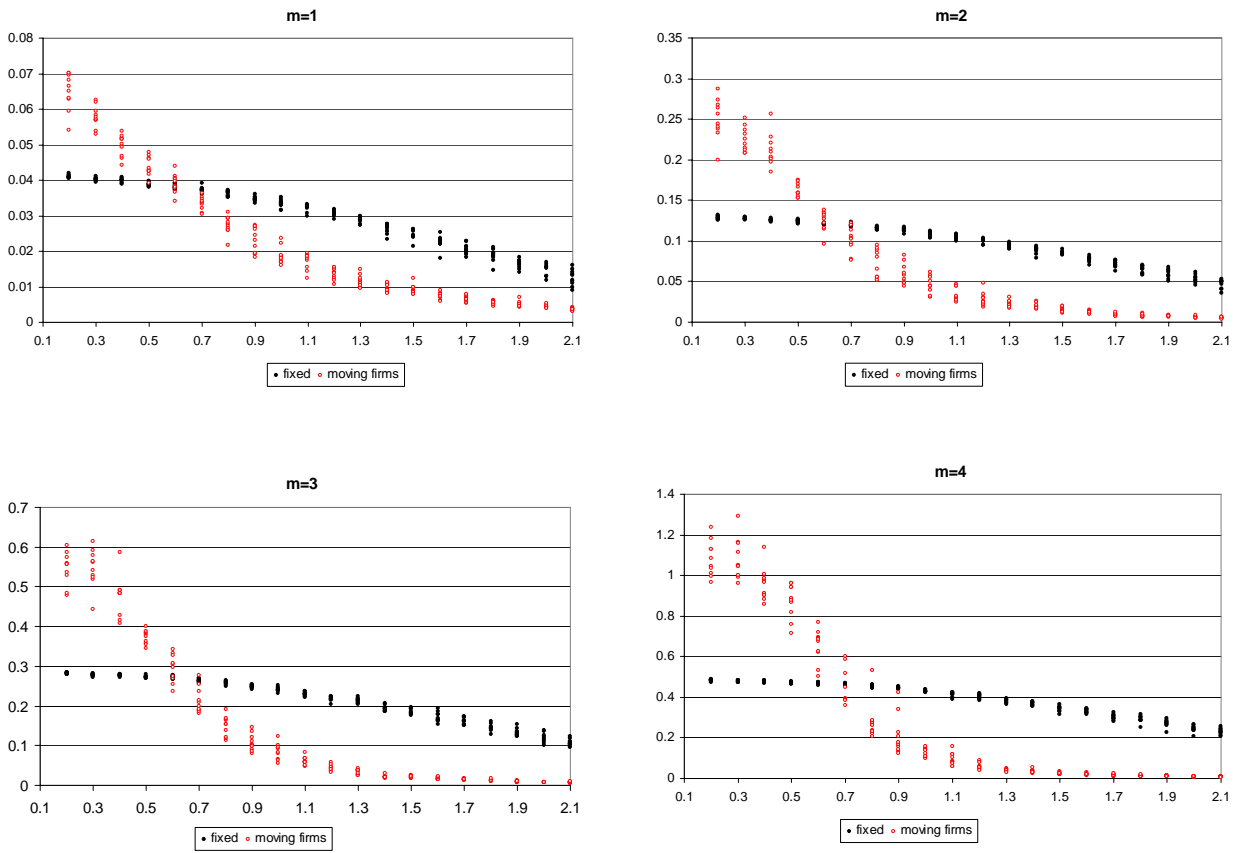


Figure 3 The innovation rate as a function of the mean of the generating lognormal distribution for four values of the search radius m and fixed and moving firm regimes.

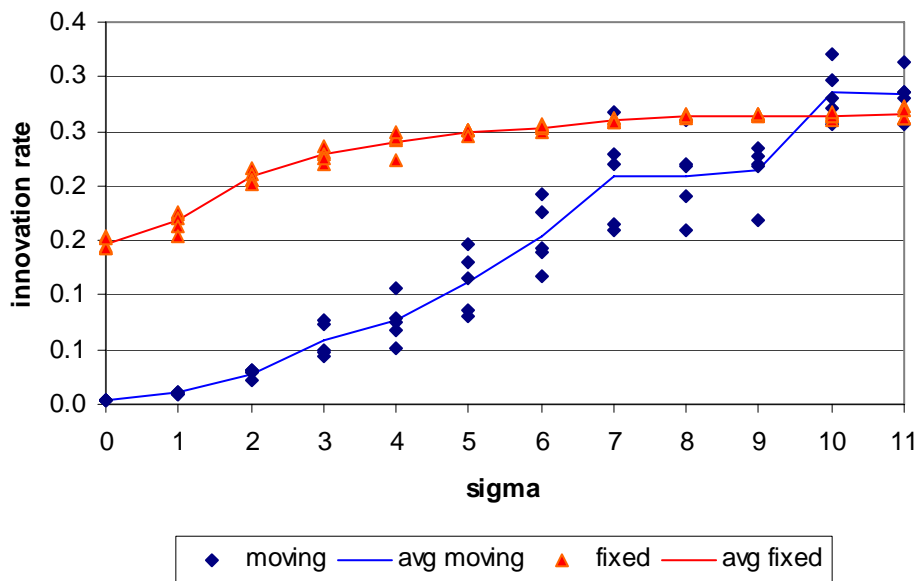


Figure 4 The innovation rate as a function of the standard deviation of the generating lognormal distribution, five data points per value, for fixed and moving firm regimes. $\langle q \rangle = 2$, $m=3$.




Figure 5 The innovation rate as a function of the standard deviation, for $\langle q \rangle = 0.4$, $m=3$, one data point per value.

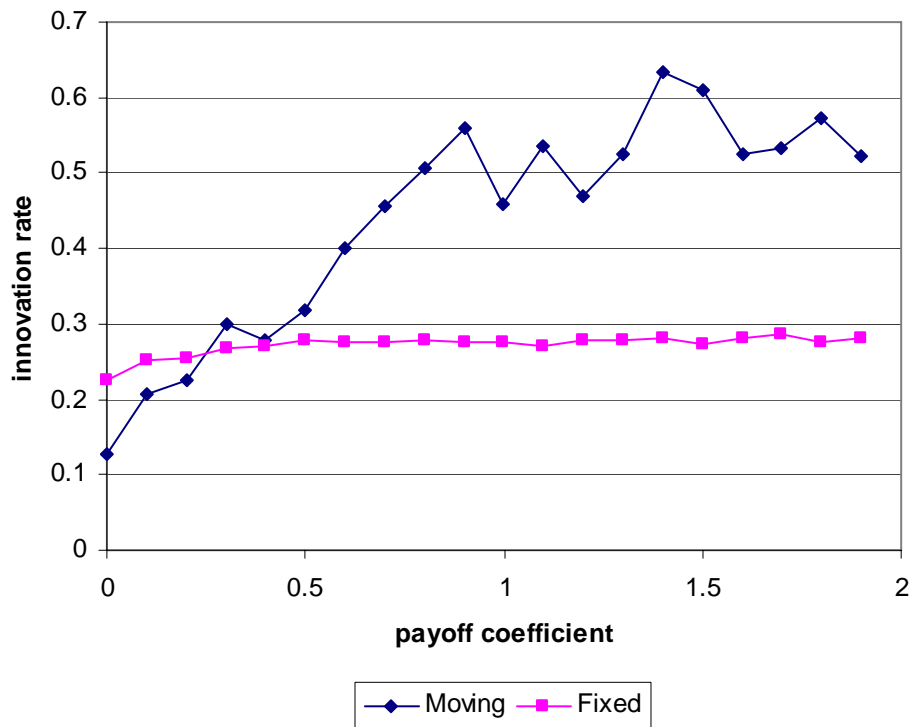


Figure 6 The innovation rate as a function of the payoff coefficient. $\langle q \rangle = 0.4$, $\sigma=2$, $m=3$, one data point per value.




Figure 7 Time series of the clustering index d for a single run with $\langle q \rangle = 1$, $\sigma = 5$, $m = 3$, $\pi = 1$.




Figure 8 The clustering index after 15,000 periods for various values of $\langle q \rangle$ and m . $\sigma = 2$, $\pi = 1$, ten data points per value with lines representing average value.




Figure 9 Pareto plots of innovation size distributions for fixed (left panel) and moving (right panel) firm regimes. $\langle q \rangle = 0.2$, $\sigma = 4$, $m = 3$, $\pi = 1$.




Figure 10 Hill plots for fixed (left) and moving (right) firm regimes corresponding to the Pareto plots of Figure 9.






Figure 11 Pareto plots of runs with $\langle q \rangle = 0.4$, $m = 3$, $\pi = 1$, and $\sigma =$ (a) 2, (b) 3, (c) 5.





Figure 12 Pareto plots for $\langle q \rangle = 0.4$, $m = 3$, $\sigma = 2$, and (a) $\pi = 0$, (b) $\pi = 1.9$.




Figure 13 Scatter plot of the variance of BPF vs. innovation rate for $\sigma=2$, $\pi=1$, several values of m , and various values of $\langle q \rangle$ between 0.2 and 2.1, pooling multiple runs.

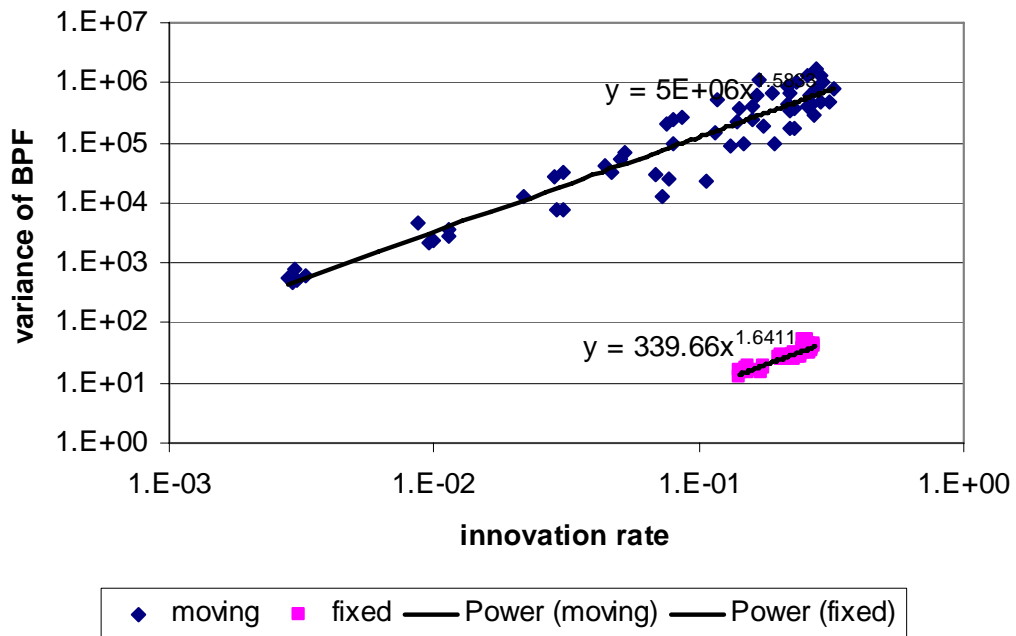


Figure 14 Scatter plot of variance of BPF vs. innovation rate for $\langle q \rangle=1$, $m=3$, $\pi=1$, and various values of σ between 0 and 11, pooling multiple runs.