- Although the evolution of technologies is modelled as a random walk process, the model displays a rich interplay of competing and co-operating trades.
- Successful trades must be closely connected to other trades (at least some α_v must be high); isolated trades will quickly decay. They should not depend on their own products (α_u should be low), and they need not be able to exploit natural resources (N_i may — or even should — be low). Chernenko's industrial subpopulation takes over in the end, but it must be able to exploit other trades effectively (κ_i must be high).

Successful populations are heterogeneous. In all the experiments which were carried out with this model (of which only two could be reported here in detail), extinction is imminent when the number of extant subpopulations goes down. In this case, new subpopulations have not enough partners to be exploited, leaving them no chance to grow, and even the extant subpopulations lack new partners which they could exploit. So one of the most important results of the model is that diversity is important. Darwin's principle, "that the greatest amount of life can be supported by great diversification of structure" (Darwin 1987: 113), thus seems to apply for this model of an evolutionary process.

8 Simulating paradigm shifts using a lock-in model

Koen Frenken and Okke Verbart[*]

A number of scholars have drawn attention to the similarities between the dynamics of scientific and technological development (Constant 1980; De Solla Price 1984; Clark 1987). Most strikingly, both scientists and engineers self-organise themselves into networks which are characterised by shared cognitive and social codes that make up a *paradigm*. Following Kuhn's notion of paradigm in science, economists have started to study technologies in terms of technological paradigms (Dosi 1982).

The dynamics of paradigms can be studied using simulation models of networks. Arthur's *lock-in* model provides a simulation of competing paradigms which are selected on the basis of network externalities that cause increasing returns to adoption (Arthur 1983, 1988, 1989). Once the relative size of the network of one paradigm *vis-à-vis* another reaches a critical value, all new adopters will opt for this paradigm, causing a lock-in.

In the following, we propose two extensions of Arthur's model. Firstly, we argue that returns to adoption of a paradigm are not solely dependent upon the size of its network, but also on its intrinsic quality. We extend the model by specifying the evolution of a paradigm's problem-solving capacity as a life-cycle. Secondly, we argue that agents are members of generations of adopters. In the model, we define generations as cohorts of agents that do not take into account the choices of previous generations.

Combining the two extensions within one simulation model, we show that the probability of paradigm shift is crucially dependent on the timing of new generations of adopters. In short, revolutions tend to take place if the occurrence of a new generation coincides with the crisis stage of a paradigm's life-cycle. The model illustrates the historical finding that paradigms are robust to competition during their growth phase but sensitive to competition during their crisis phase.

Suggestions for future modelling efforts are made in the final section with special reference to recent models of complex co-evolving systems. Such systems

[*] Suggestions from Loet Leydesdorff, Peter van den Besselaar, Vincent Mangematin

consist of parallel operating subsystems that are loosely coupled (Simon 1969). Their dynamics show self-organising behaviour over several time horizons (e.g., Kauffman 1993). Such models can be used to study complex systems in which paradigm shifts in one subsystem affect the probabilities of paradigm shifts in other subsystems. Their collective dynamics illustrate long waves in scientific and technological development.

8.1 Lock-in

Arthur's lock-in model has given rise to a proliferation of simulation models which analyse the conditions of lock-in processes using a variety of modelling techniques (David 1993; David and Foray 1994; Bruckner et al. 1994; Dalle 1998; Windrum and Birchenhall 1998). In this section, we go back to the original model and recapitulate on its main features. It will be argued that the generality of Arthur's model is higher than commonly recognised. Apart from its application to competing technical standards, the model can be viewed as a simulation of competing paradigms. In the succeeding sections, we will develop an extended version of Arthur's model to simulate both lock-in and *lock-out*. In this way, the probability of paradigm shifts can be analysed for different parameter regimes.

Arthur's model

Arthur's model consists of two types of agent (R and S) which are randomly drawn from a mixed pool. Each type of agent chooses a technology i according to the highest returns to adoption. Table 8.1 shows the model for the case of two competing technologies (i = A,B). Returns to adoption are determined by the agent-specific value of natural preference (a_r, b_r and a_s, b_s) and the number of agents n_i that adopted technology i previously.

Lock-in occurs when parameters r and s are positive. In that case, returns to the adoption of technology i increase as n_i increases. Random drift may push the system towards one of the two basins of attraction. For example, if R-agents — choosing according to their highest natural preference a_r — are drawn far more often than S-agents, then S-agents may choose against their natural inclination b_s. Normally S-agents would choose technology B, but due to the fact that a considerable number of R-agents have chosen technology A, they opt for A too.

	returns technology A	returns technology B
R-agent	$a_r + rn_a$	$b_r + rn_b$
S-agent	$a_s + sn_a$	$b_s + sn_b$

$$a_r > b_r;\ b_s > a_s;\ a_r = b_s;\ a_s = b_r;\ r = s;\ r, s > 0$$

Table 8.1. Arthur's lock-in model

One can derive from table 8.1 the conclusion that lock-in of technology A occurs when the value of (a_s+sn_a) exceeds (b_s+sn_b) and lock-in of technology B when (a_r+rn_a) exceeds (b_r+rn_b). We can rewrite the formulas as follows:

lock-in technology A if $n_a - n_b > (b_s - a_s) / s$

lock-in technology B if $n_a - n_b > (b_r - a_r) / r$

Symmetry of parameter values assures that each technology has the same probability of being selected (here, 0.5). Because of the stochastic nature of the drawings of agents, lock-in occurs irrespective of the specific parameter values. The parameter values only affect the rapidity of lock-in. The higher the values for r and s and the lower the values for a_s, b_s, a_r, and b_r, the more rapidly lock-in occurs and *vice versa*. Figure 8.1 shows a simulation run in which technology B was selected. The x-axis represents the number of adopters and the y-axis measures the distribution of market shares of technologies. The curvature of market shares shows that monopoly of B is obtained asymptotically.

Two features of Arthur's model are crucial. Firstly, the process that leads to a lock-in is a *path-dependent* process: whether technology A or B is selected depends on the specific sequence of random drawings of R-agents and S-agents. Secondly, once the market share of one technology has obtained a critical size causing a lock-in, the system goes through an *irreversible* transition: thereafter, both types of agents choose the same technology with unit probability.

Appreciation

Several historical studies support the empirical relevance of the lock-in model. For example, competing standards in video technology gave rise to a situation of two temporary networks of users of either Betamax or VHS that eventually locked into one VHS network (Rosenbloom and Cusumano 1987). Other examples of competing standards and lock-in phenomena in aircraft, automobile and traffic systems, in electricity systems, gas turbines and nuclear power stations, as well as in typewriting, software and telecommunications, are well documented (David 1975, 1985, 1992; Utterback and Abernathy 1975; Constant 1980; Sahal 1981, 1985; Cowan 1990; Mangematin and Callon 1995; Cowan and Gunby 1996; Davies 1996; Islas 1997).

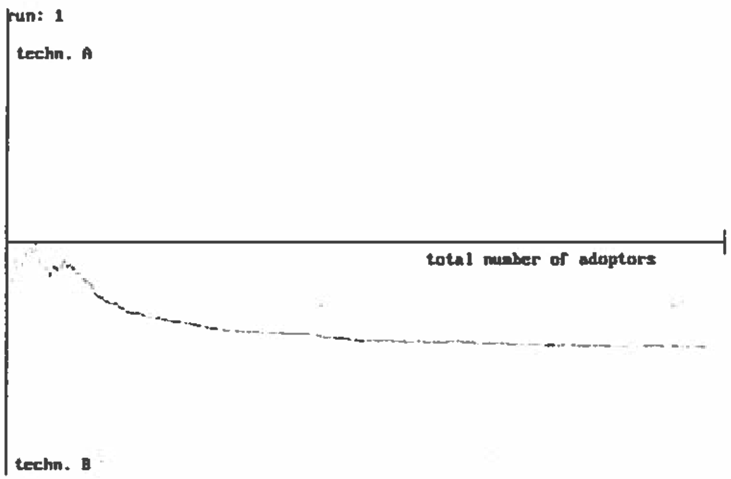

Fig. 8.1. Lock-in

Arthur's lock-in model has been important because it provided a formal explanation for the emergence and persistence of technical standards. The model shows that lock-in processes are due to increasing returns to adoption caused by *network externalities*: the more people already use a particular standard, the higher the benefits for new users to enter this network. Agents organise themselves by choosing among standards that are incompatible (Katz and Shapiro 1985). Importantly, networks of agents are not distinguished on the basis of the intrinsic properties of the agents, e.g., their natural preferences, but on the basis of their adoption of different mutually exclusive standards.

We can appreciate the generality of the model once standards are conceptualised as one possible axis along which technologies can differ. In general, lock-in may occur in all selection processes where the formation of networks causes increasing returns to adoption. For example, the geographical proximity of users may explain innovation clusters (Arthur 1988). In that case, types of firms are distinguished in terms of their 'natural' geographical preferences, and the size of networks among firms is distinguished on the basis of the number of firms that are already located in a region. Then, lock-in occurs if a critical number of firms are located in one region for stochastic reasons causing substantial spillovers within that region. Thereafter, all firms choose to locate in the same region. The crux of the argument holds that agents have to make a fixed and irreversible investment in order to join a network. The investment allows them to profit from externalities which are network-specific but restrains them from gaining from externalities present in other networks.

Given its generality, the lock-in model is also of value as a simulation of competing paradigms defined as communication networks that differ in terms of their code of communication (Leydesdorff 1995; Cowan and Foray 1997). The more people use a particular code, the higher the benefits for new agents of joining this particular communication network since the size of knowledge spillovers is positively related to the size of the network. Like technical standards and geographical proximity, paradigms facilitate exchange between users of a similar code but restrains them from communicating with agents using a code specific to a different paradigm.[1]

8.2 Simulating paradigms

After Kuhn's (1962) seminal study, a large number of histories of scientific theories and technological systems have been written using the concept of a paradigm (for an overview, see Clark 1987). In this section we formalise two important notions about paradigms and we incorporate these into Arthur's original model. These are that the problem-solving capacity of a paradigm follows a life-cycle pattern, and that agents are members of generations of adopters. The two extensions will prove to be sufficient to model *lock-out* and paradigms shifts.

Life-cycles

Arthur's model can be criticised because it does not take into account the cognitive changes which take place within a paradigm (cf. Ahrweiler and Wolkenhauer, this volume). In his model, returns to adoption are only related to the size of the adoption network. However, agents are expected to base their choice also on the effectiveness of a paradigm to solve problems. Thus, apart from the size of the network, the state of problem-solving capacity constitutes a second component of returns to adoption.

A paradigm's *problem-solving capacity* is dependent upon the number of previous adopters that have contributed to its growth. However, this dependency relation is not likely to be linear (if it were linear, Arthur's original model would do). Instead, the problem solving capacity of paradigms follows a *life-cycle* pattern (Utterback and Abernathy 1975, Sahal 1981, 1985; Sterman 1985; cf. Kuhn 1962). Initially, as the number of adopters increases, problem-solving capacity increases exponentially, as a result of the division of labour on the producer's side, and new fields of application on the user's side. Then, a phase of development begins in which problem-solving capacity grows fastest. This phase corresponds

[1] The idea of increasing returns in scientific development goes back to Merton's (1968)

to the heyday of *normal problem-solving* activity. Thereafter, as a paradigm matures, the division of labour on the producer's side encounters a limit because coordination costs become too high. In addition, fields of application on the user's side become exhausted. Anomalies signal open questions that are considered to be beyond solution within the common framework, and the growth of problem-solving capacity comes to a standstill.

Formally, the evolution of problem-solving capacity can be modelled as an S-curve representing the life-cycle of a paradigm. Empirical studies of scientific and technological development tend to confirm the life-cycle pattern of paradigm development using a variety of performance indicators (Sahal 1981, 1985; Marchetti 1991; Klepper 1996).[1] The S-shaped curve representing problem-solving capacity has itself been the object of simulation attempts (Sterman 1985). Here, we treat the S-shaped life-cycle as a *stylised fact* and introduce it by simply specifying the mathematical function.

The specification of an S-shaped curve which represents the evolution of the problem-solving capacity termed the *learning effect* T of paradigm i is given by the sigmoid function:

$$T_i(n_i) = 1/(1+e^{-(\beta/\alpha)z}), \text{where } z = n_i - \alpha/2 \quad (\alpha > 0, \beta \geq 0)$$

In the formula, parameter α refers to the length of a paradigm's life-cycle in terms of the number of adopters and parameter ß determines the skewness of the shape of the S-curve relative to the length of the S-curve α.

If no adopters have chosen technology i ($n_i = 0$), the learning effect is close to zero. If half the length of the learning curve has been reached ($n_i = \alpha/2$), T is equal to ½. When the total learning curve has been reached ($n_i = \alpha$), T approaches 1. The S-shaped curvature of the life-cycle represents progress accelerating during the first half of the cycle and slowing down during the second half of the cycle. Figure 8.2 shows the graph for $\alpha = 2000$ and $\beta = 20$.

[1] The specification of problem-solving capacity as an S-curve must not to be confused

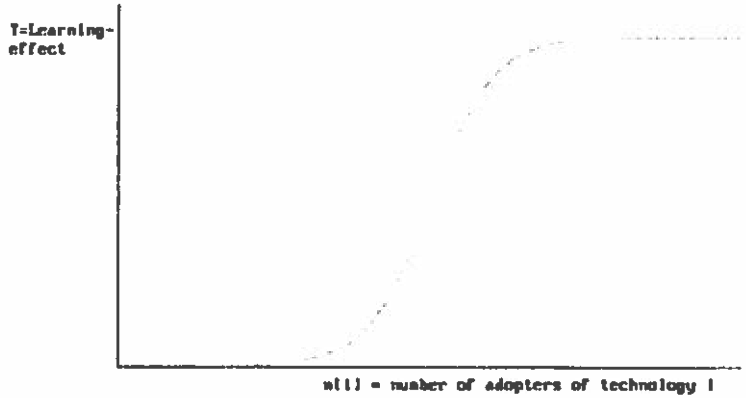

Fig. 8.2. The evolution of problem-solving capacity T as a S-shaped curve

Generational change

The second extension concerns the notion of generational change. In Arthur's original model adopters take *all* previous decisions into account. The assumption ensures that lock-in will take place eventually with probability 1.0. Historical studies stress, however, that group dynamics are not only present in networks of adoption, but also within generations of adoption, i.e., on a temporal scale (Kuhn 1962; Constant 1980). Thus, the historical time-horizon that agents take into account is subject to change. During some periods new adopters are influenced by a large group of previous adopters (the old generation), and during other periods they base their decision only on a small group of recent adopters (new generations).

Introducing the assumption of generational change implies that adopters base their decisions on a reference group of adopters which belong to the same generation. This is modelled by specifying cohorts of adopters that only take into account the previous decisions made by members of the same generation. In the model, generations come and go according to a fixed frequency and therefore have a fixed size (in a more elaborated setting, the size of generations can also be made subject to stochastic variation).

Model specification

Table 8.2 shows the extended model that combines the notions of a S-shaped life-order to secure an equal probability of lock-in for both paradigms. The size of each generation is labelled K and the number of adopters of generation G that have adopted technology A and B are labelled n_{aG} and n_{bG}. If an adopter is the first agent of a new generation, the network effect is absent because the members in reference groups n_{aG} and n_{bG} equal zero. On the other hand, if an adopter is the last agent of a generation, the network effect is at its maximum.

	returns technology A	returns technology B
R-agent	$a_r + rn_{aG}T_a$	$b_r + rn_{bG}T_b$
S-agent	$a_s + sn_{aG}T_a$	$b_s + sn_{bG}T_b$

$$a_r > b_r;\ b_s > a_s;\ a_r = b_s;\ a_s = b_r;\ r = s;\ r, s > 0;\ K > 0$$

Table 8.2. Extended lock-in model

In the extended model in Table 8.2, each adoption increases both the returns to adoption via the network dynamic (linearly) and via the learning dynamic (following an S-curve). If the life-cycle dynamics are linear ($\beta = 0$) and if generational change is absent (i.e., $G = 1$), the extended model perfectly matches Arthur's original model. In that case, T is a constant and takes on the value ½. Thus, Arthur's model is a special case of the extended lock-in model presented here.

We use the model to analyse the conditions under which generational change leads to a paradigm shift (i.e., lock-in into the alternative one) or to paradigm stability (lock-back into the existing one). In the former case, the paradigm shows sensitivity to generational change, whereas in the latter case the paradigm shows robustness to generational change.

Results

The model has been run for different values of α and β in order to analyse the conditions that are likely to bring about paradigm shifts. We restrict the analysis to the case of two generations ($G = 1,2$). To allow for fair competition, we have to assume that once a lock-in occurs into a technological paradigm, an alternative paradigm is available which has the same level of problem-solving capacity T, but which has not yet been adopted. For example, if technology A locks-in, the learning effect of technology B is reset to the same level as A's, and n_b is reset to zero (no adopters yet). In other words, once a technological paradigm locks-in, an alternative paradigm is 'discovered' that equals the problem-solving capacity of the prevailing paradigm. Therefore, the second generation will always temporarily lock-out the paradigm that locked in during the first generation. This causes a transition from lock-in of the prevailing paradigm to competition between the old paradigm and the alternative paradigm. Thereafter, the alternative paradigm may

lock-in (paradigm shift) or the old paradigm may lock-in again (lock-back). Each time a new paradigm becomes available, it starts at the bottom of its S-curve.

In this setting, the stochastic sequence of drawings of types of agents determines whether the alternative paradigm will be adopted, or whether the old paradigm survives. Figure 8.3 shows a simulation run in which a paradigm shift took place: technology B locked-in during the first generation and technology A locked-in during the second generation (the term network effect refers to parameters r and s).

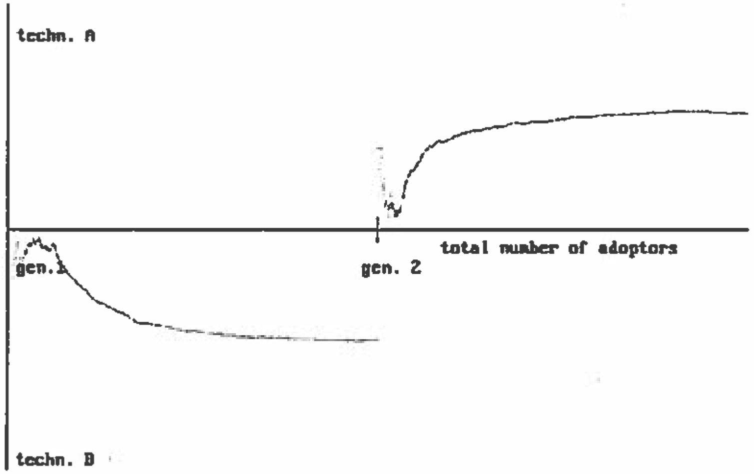

Fig. 8.3. Paradigm shift due to generational change

We simulated the model systematically for different parameter settings to analyse the likelihood of paradigm shifts under different regimes. Table 8.3 shows the results for different settings of the length of the life-cycle α and the skewness of the life-cycle β. The probability values refer the number of paradigm shifts divided by the total number of runs.

The results show that the probability of lock-in into the alternative paradigm versus the probability of lock-back into the old paradigm depends critically on both parameters.

	$\beta = 10$	$\beta = 15$	$\beta = 20$
$\alpha = 1000$	0.838	0.750	0.706
$\alpha = 1500$	0.268	0.384	0.504
$\alpha = 2000$	0.038	0.012	0.004

K=2000; r = s = 0.025; number of runs per parameter setting = 500

Table 8.3. Probability of paradigm shifts for different regimes

For all the runs, we kept the size of generations at $K = 2000$. This implies that if $\alpha = 1000$ for each paradigm, the paradigm that locked in during the first generation has gone through the whole of its life-cycle. Even if it locked in after the last adoption during the first generation (i.e., after 2000 adopters), the selected paradigm will have completed its cycle (for $\alpha = ½ K$). Then, at the time of arrival of the new generation, the prevailing paradigm finds itself in its crisis phase ($dT/dn_i = 0$).

In contrast, if $\alpha = 2000$, the paradigm that locked in during the first generation is still in its growth phase: its problem-solving capacity grows at a fairly high rate (dT/dn_i is close to its maximum). This case corresponds to the phase of normal problem-solving. Setting α to 1500 simulates the in-between case: the paradigm is close to crisis, but its problem-solving capacity is still growing slightly (dT/dn_i is close to zero).

If we compare the probabilities of paradigm shifts for different regimes of α, the results show that if $\alpha = 1000$ (crisis phase), paradigms are very sensitive to competition. In most cases, the system shifts from one paradigm to the other. If α is 2000 (growth phase), paradigms are more often insensitive to competition as is reflected by the low probability values. This represents the robustness of phases of normal-problem solving. As expected, the in-between case ($\alpha = 1500$) shows in-between probability values.

The comparison of simulations for different regimes of ß provides us with a more complicated picture. If the S-curve is relatively linear (ß = 10), paradigm shifts occur more often in the two extreme cases of α, whereas paradigm shifts occur less often in the in-between case. And, if the S-curve is relatively skewed (ß = 20) paradigms shifts are less likely in the extreme cases of α, but more likely in the in-between case. In other words, if the old paradigm is faced with competition during its crisis phase or growth phase, as is the case for the two extreme values of α, the skewness of the life-cycle may protect it from being taken over by the new paradigm: a more skewed life-cycle implies that the problem-solving capacity of the *new* paradigm grows only very slowly at the beginning of its adoption. The latter feature is advantageous for the older paradigm. However, if the older paradigm is located in between the growth phase and crisis phase, a more linear life-cycle may protect it from being taken over, because the growth of the problem solving capacity of the *old* paradigm (dT/dn_i) is still substantial.

Summarising, the model shows that paradigms in their crisis phase are very

whereas paradigms at the stage of normal problem-solving are quite robust in the face of competition. In addition, the skewness of a paradigm's life-cycle influences its sensitivity to competition differently for different phases of its problem-solving capacity.

8.3 Towards a co-evolution model

The extended lock-in model discussed above has been developed to specify and analyse the basic structure of a single system in which two paradigms compete for adoption. The model can serve as an illustration of the evolution of paradigm competition within a single industry or scientific discipline. However, both science and the economy are composed of several, parallel subsystems in which paradigms come and go. Paradigms shifts in one system may have a considerable impact on the competitive advantage of paradigms in other systems because of the existence of epistatic links. For example, whether the aviation industry is locked into a piston propeller paradigm or a turbojet paradigm has important consequences for the competitive strength of particular paradigms in related industries such as the engine industry, the helicopter industry and the aerospace industry. The same reasoning can be applied to competing paradigms in economics that favour their counterparts in sociology and political sciences. Thus, paradigm shifts in one industry/discipline may spillover to another ones, possibly creating an *avalanche* of shifts throughout the whole system (cf. Bak *et al.* 1988). Such a system constitutes a *complex* system of co-evolving subsystems of which the functions are epistatically related (Simon 1969).

Following Kauffman's (1993) *NK*-model of complex systems, we envisage a model consisting of N subsystems (i.e., industries/disciplines) which are related by K epistatic links. The subsystems operate independently from each other in the sense that for each subsystem agents are drawn from independent pools. The epistatic relations between subsystems can then be modelled by varying the network effect relative to the collective problem-solving capacity of the K related paradigms in other subsystems. The groups of paradigms in different subsystems start to co-evolve as a collective.

In such a model, generational change may fail to bring about any paradigm shift during phases of early collective development in which paradigms are shielded by high growth in the problem-solving capacity of related paradigms. However, during later phases of collective development, generational change may cause a snowball of paradigm shifts changing the entire system within a short period of time. Such a model may provide a basis for the simulation of the rise and fall of techno-economic paradigms which are composed of epistatically related paradigms creating long waves in the economy (Freeman and Perez 1988).

Springer
Berlin
Heidelberg
New York
Barcelona
Budapest
Hong Kong
London
Milan
Paris
Singapore

Petra Ahrweiler · Nigel Gilbert
Editors

Computer Simulations in Science and Technology Studies

With 51 Figures
and 10 Tables

Table of Contents

Part I

Introducing simulation in social studies of science and technology

Part II

Simulating the logic of discovery

Part III

Evolutionary models of science and technology

Part IV

Simulating conditions and dependencies of scientific work