



Paper to be presented at the DRUID Academy 2013

on

DRUID Academy 2013
at Comwell Rebild Bakker, Rebild/Aalborg

Adder: a simplified model for simulating the search for innovations

Janne M. Korhonen

Aalto University
School of Business
janne.m.korhonen@aalto.fi

Julia Kasmire

Delft University of Technology

j.kasmire@tudelft.nl

Abstract

Computer simulations are increasingly used to study the development, adoption, and evolution of technologies. However, existing models suffer from various drawbacks that may not be easily corrected, among them lack of internal structure in technologies, static environments and practical difficulties of introducing search dynamics more representative of empirical findings from historiography of technology. In particular, the implicit assumption that distal search is required for radical breakthroughs from the state of the art is problematic. This paper discusses the theoretical background and rationale for an improved model, the Adder, and sketches out the model's main features.

Adder: a simplified model for simulating the search for innovations

Janne M. Korhonen and Julia Kasmire

School of Economics / Aalto Design Factory and TU Delft Faculty of Technology, Policy and Management

Betonimiehenkuja 5, 02150 Espoo, Finland and Jaffalaan 5, 2628BX, Delft, The Netherlands

E-mail: janne.m.korhonen@aalto.fi and j.kasmire@tudelft.nl

ABSTRACT

Computer simulations are increasingly used to study the development, adoption, and evolution of technologies. However, existing models suffer from various drawbacks that may not be easily corrected, among them lack of internal structure in technologies, static environments and practical difficulties of introducing search dynamics more representative of empirical findings from historiography of technology. In particular, the implicit assumption that distal search is required for radical breakthroughs from the state of the art is problematic. This paper discusses the theoretical background and rationale for an improved model, the Adder, and sketches out the model's main features.

INTRODUCTION

For the last two decades, researchers studying various aspects of change and evolution in organizations, strategy and technology have increasingly turned to computer simulations to better understand the dynamic and often complex interactions that are inherent in their fields of study. Much of this research has been undertaken using only few basic types of simulation models, resulting to the emergence of "dominant designs" that effectually serve as benchmarks against which novel contributions are evaluated.

To mention just a few historical examples, *in silico* modeling has been used to illuminate how tightly interlinked organizations have problems adapting to changing environments (Levinthal, 1997), explain why imitation of complex socio-technological systems is difficult (Rivkin, 2000) and why (and under which conditions) modularization may be a successful strategy (Ethiraj & Levinthal, 2004), help develop a theory of invention as a process of recombinant search over technology landscapes (Fleming & Sorenson, 2001), reproduce clustered patent citation patterns (Silverberg & Verspagen,

2005) and study the evolution of technologies (Frenken, 2006). A common theme in these models has been to see the process of novelty generation as a search for solutions in a specified "landscape" of potential solutions, with initial conditions (e.g. the starting point) often having a large impact on the outcome.

Although these model frameworks continue to be useful, they also suffer from certain limitations. First, NK and percolation frameworks have difficulties simulating endogenous situations, i.e. where attributes of solution space change over time (Ganco & Hoetker, 2009).

Second, they ignore the internal structure of technologies. Although many authors who have used simulations to study technological evolution (e.g. Frenken, 2006; Murmann & Frenken, 2006) do specifically argue that technologies have a hierarchical structure and consist of systems, subsystems and sub-subsystems (etc.), the current models do not easily allow for this. Thus, individual technologies are independent of each other, instead of existing in an interdependent ecosystem of combinatory possibility. Furthermore, the lack of internal structure makes the improvement of existing technologies through improvements in their components - a feature visible in any case study of technology - clearly impossible.

Third, introducing search behavior that more closely resembles human behavior is difficult in the existing models. While implementing advanced AI algorithms to search for solutions in the NK solution space or within percolation lattice is certainly possible, programming these algorithms is non-trivial for most researchers. Furthermore, complex AI algorithms pose their own challenges when applied to models that are intended to be simple and easily understandable "toy models" of innovation and technological change¹.

¹ It should be noted that most scholars of technology are in a broad agreement that "technologies" should be understood to mean not just material artefacts, but also nonmaterial methods, processes and devices that are means to fulfill a human purpose (e.g. Arthur 2009:28). This broader definition of technology thus includes fields such as management practices and strategies.

These limitations are particularly visible when the goal is to model a co-evolutionary system that evolves on several different levels simultaneously, a system whose complexity changes over time, or the influence of more or less rational actors. Obviously, these types of systems are of interest to many researchers from fields such as organization science, studies of science and technology, or management of innovations, to name just a few. Developing model frameworks that answer better to their requirements thus remains a challenge for simulation community.

THEORETICAL BACKGROUND

The notion that technologies "evolve" is not new. For example, Basalla's seminal work (1989) notes that the earliest attempts linking technology and evolution explicitly date from the 19th century. In addition to Basalla's work, the evolutionary processes have been explored by authors such as Frenken (2001, 2006) and Arthur (2009; 2007). All broadly agree that while technologies exhibit evolutionary features, strict 1:1 mapping of biological metaphors to technological evolution is not appropriate. Instead, technological evolution should be understood as an instance of "Universal Darwinism" (UD) as introduced by Dennett (1995). In principle, UD states anything that displays variation, selection and heredity will evolve through natural selection, whether or not it would be classified as "alive" in any traditional sense. This universal definition of evolution does not specify *how* variation, selection and heredity work; the sources of variation and the mechanisms of selection may result either from unconscious environmental pressures and random events, or from deliberate "tinkering" and rational selection. Similarly, heredity can operate through biological mechanisms such as DNA, or it can operate through information codified in drawings, patents, and operating manuals. In fact, all technologies are seen to descend in some way from technologies that preceded them (B. W. Arthur, 2009, p. 181).

However, Arthur (2009) emphasizes the need to consider *combinatorial evolution*, where novel technologies can arise by combination of existing technologies, in addition to incremental change. He finds this particularly useful when explaining developments such as jet engines or radar, that appear to be radical departures from the existing technologies (Arthur 2009:17). Arthur proposes a mechanism of *combinatorial evolution* where novel technologies can arise by combination of existing technologies. Combinatorial evolution (CE) does not abandon incremental variation, selection and retention: it acknowledges that these, too, have an important role to play. Rather, CE proposes an additional mechanism to explain how radical departures from existing technology might happen.

The idea of technologies being recombinations of existing technologies has been long accepted by scholars such as Gilfillan (1935), Schumpeter (1939), and Usher (1954), among others, who describe production or inventions as combinations of materials and forces (Schumpeter) or "new combinations of prior art" (Gilfillan). In Arthur's formulation, a CE mechanism for the evolution of technology would work as follows (Arthur 2009:21-24):

1. *Early technologies form using existing primitive technologies as components.*
2. *These new technologies in time become possible components, or building blocks, for the construction of further new technologies.*
3. *This implies that technologies have an internal structure, a hierarchy of subsystems and sub-subsystems.*
4. *The complex technologies form using simpler ones as components.*
5. *The overall collection of technologies bootstraps itself upward from the few to the many and from the simple to the complex.*

Furthermore, it is notable that few modelers have tried to test alternative search assumptions, preferring to assume relatively myopic search (taking into account only one decision at any given time) in the technology landscape. Although some efforts have been made (particularly with NK models) to introduce more "cognitive" search dynamics (e.g. Gavetti & Levinthal, 2000), the most common alternative to local, myopic search has so far been the so-called "distal" search. This is intended to break the deadlock caused in most models by local hill climbing type search, which leads to searching agents being stuck in local performance optima. Although the idea of distal search enjoys considerable support in theoretical literature (e.g. Katila & Shane, 2005), to us it seems that this idea is at best an oversimplification. The assumption underpinning distal search is that major advances require "thinking outside the box" and must be conceptually distant from the state of the art. This line of thinking is so prevalent within innovation literature that it leads to a somewhat circular reasoning: "radical" innovations (i.e. ones with radical effects) are almost by definition considered to be conceptually distant from the state of the art (e.g. Rosenkopf & Nerkar, 2001), and one of the most commonly suggested indicators used to measure the "radicalness" of innovations is usually its conceptual distance from the state of the art (e.g. Dahlin & Behrens, 2005).

However, history of technology has largely abandoned the idea of large conceptual leaps being necessary preconditions to successful innovation (Conway, 2008; Edgerton, 2010). Instead, most historians of technology nowadays maintain that major innovations tend to be the result of steady

accumulation of relatively minor advances. To mention just one piece of evidence in favor of this theory, the overwhelming prevalence of simultaneous or near-simultaneous independent discovery of key innovations (to take a recent example, Lemley (2012), in his non-exhaustive survey, finds 26 examples ranging from telephone and computer to steam engine and automobile) and the difficulty of pinpointing even a few examples of “true” independent discovery (Lemley, 2012) strongly suggest that true conceptual leaps in the sense of distal search happen only rarely, if at all. Instead, inventors typically take small, logical steps towards their goal; only when looked from the outside in and without appreciation of the details of the process, the results may seem to be a radical departure from the state of the art (e.g. commercialized inventions).

This type of search – where the agents have a certain knowledge of the goal they would wish to reach, but need to bridge the gap from here to there with small incremental steps – has been poorly represented in models used so far. Percolation and lattice-type models could be used to represent this type of search, but these models then largely ignore the hierarchical nature of technologies. While NK-type models represent adequately the internal structure and trade-offs of technologies, coding search algorithms that operate in this manner seems to be very difficult.

TOWARDS A BETTER MODEL: THE “ADDER” SUGGESTION

Given the drawbacks of the NK and percolation models, some scholars have introduced alternative models of technological evolution. Of particular interest is a model introduced by Arthur and Polak (2006) and latter used by Arthur (2007, 2009) to illustrate his formulation of the CE framework, as well as by scholars from other fields (e.g. White 2008). In their model, technological build-out begins from simple "primitive" technologies. These technologies are randomly combined to result to more complex technologies, which themselves are then potential components in future technologies. The system includes concrete needs instead of abstract fitness values, and technologies that better satisfy these needs or have fewer components - i.e. are cheaper - than their alternatives supersede older technologies.

The model is implemented using simple logic circuits (NAND gates in the most common version) as primitive technologies. Needs include simple logical functions, such as 2-bit adders. New technologies are evaluated against how closely they fulfill a need's truth table, and how few primitive components they use in doing so.

This model successfully captures the interlinked build-out of technologies from simpler components, and replicates stylized facts such as avalanches of "creative destruction" when a significant innovation suddenly makes obsolete many old technologies. However, the model is still relatively fixed, with no easy way to add many more needs than simple logical functions (e.g. NOT, IMPLY, *n*-way-or, 2-bit adder, 8-bit adder and so on) present in the original implementation. This limitation creates difficulties for studying settings where more or less every technology might fulfill *some* need.

In addition, introducing bounded rationality is – again – difficult. As a result, even Arthur and Polak (2006) are reduced to arguing that their model of random combinations *can* be representative of search through solution space, given certain assumptions. While this may be true, one would very much like to test how the assumption affects this otherwise excellent model.

A possible solution is the Adder model detailed in this paper. The Adder simplifies Arthur and Polak’s model by replacing logic circuits and Boolean arithmetic with real numbers and arithmetical expressions. Each experiment starts with elementary components (primitives) and arithmetical operators, usually with number 1 and operators plus (+) and minus (-).

During each time step of the experiment, developer agent(s) alter existing “technologies” by adding or removing components, hence the name of the model. The resultant technologies are then evaluated against goals and added to the repertoire of possible components. The objective of the system is to satisfy a certain set of needs or goals, expressed as real numbers. These needs can be thought to represent the needs that drive technology evolution, and as simplifications of logical operator needs used by Arthur and Polak (2006). The numbers can be either drawn randomly or according to some distribution pattern. Compared to a relatively fixed set of goals in the original model, this allows for certain flexibility in studying different technological landscapes (some landscapes might have feasible technologies more clustered in design space than others, for example).

As an example, let us assume that one of the goals is “10,” that this is the first step and therefore the available component is “1” and operators plus and minus, and that the selected method of alteration is random draw of 0 to 12 components and operators. A possible draw could be

$$1 + 1 - 1 + 1 + 1 - 1 + 1 + 1 \quad (1)$$

“producing” the value “4.” Although this technology did not fulfil the goal in itself, it is now added to the repertoire of possible components and is therefore available for use in the next step. Suppose that the next draw gets the components

$$-4 - 1 - 1 + 4 + 4 + 4 \quad (2)$$

that produce “10” and thus satisfy the first goal. The process continues until desired set of conditions is reached, for example, when all the goals are satisfied or the simulation has progressed for a predetermined time.

Evaluation by Cost and Fitness

The “goodness” or “fitness” of these technologies can be evaluated in a variety of ways, depending on the requirement of particular experiment.

One important evaluation criteria is the “cost” of the technology. In its simplest form, the cost is determined by counting the number of primitive elements required for the technology. To continue the above example, the primitive components (1’s) have a cost of 1. Component technologies, such as technology “4” above, have a cost equal to the cost of primitives within it. Thus, the above technology “4” would cost 8.

It is evident that such costing schemes neatly capture one of the important mechanisms in the evolution of technologies: that many technologies, when first developed, are very expensive, but become cheaper as R&D efforts are made towards improving the efficiency and manufacturing processes, for example. The “technologies” encapsulated in the model can be thought of as simplified idealisations of production recipes or assembly instructions, subject to improvements as more streamlined processes are found.

For example, the technology “4” in the model could be superseded by several generations of more efficient technologies, with the ultimate limit of efficiency being

$$1 + 1 + 1 + 1 \quad (3)$$

with a cost of 4. Determining the efficiency limits and the most efficient technology possible is always trivial (i.e. when Cost = Product).

Another possible evaluation criteria is the fitness-for-purpose, that is, how close the technology gets to the target. For example, the model may accept only those new technologies that are either i) closer to the target value than existing technologies or ii) cheaper than existing technologies.

Obsoleted Technologies

The Adder can model the obsolescence of technologies through basic mechanisms described above. If a new technology proves to be either fitter or cheaper than existing technologies, it takes their place in the repertoire of technologies that are used as a pool of possible components in the future. If the new technology is simply a cheaper version of already existing technology, the technologies currently using the old version are updated. As their costs are updated in turn, it is possible that a development of a new

component triggers an avalanche of replacements. The size-frequency distribution of these avalanches (see Fig. 1) shows hints of power law distribution, indicative of self-organized criticality (Bak & Wiesenfeld, 1988).

This obsolescence does not, however, necessarily obsolete other technologies that are already using the now-obsolete technology. This allows for “legacy” technologies, where otherwise obsolete components remain in use as parts of older systems. As an example, suppose that a target “10” is reached, while there exist technologies that use the tech 9 as a component. Although future technologies will not use tech 9 any longer, any technologies that have tech 9 as a component will retain it. It is even possible that tech 9 is incrementally developed towards a cheaper version, resulting to decreased costs for technologies that use it.

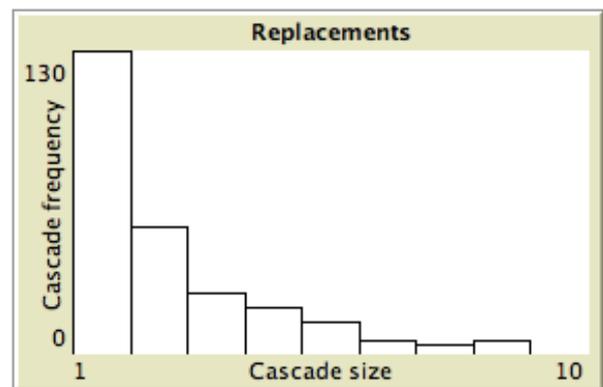


Figure 1: Sample Distribution of “Replacement Cascades”

Bounded Rationality

It is easy to see that implementing even fully rational agents is nearly trivial in this model. For example, it’s easy to have the agents determine the optimum components they’d need in order to reach a certain target. One can easily go further and implement bounded rationality (Simon, 1982) by e.g. introducing random uncertainty into the calculation.

What’s more, implementing search algorithms that better approximate the conceptions of historians of technology becomes possible. It is easy to conceive of several possible search dynamics that, for example, seek to satisfy a certain target but do so in an incremental manner.

Tuneable Difficulty of Search

The description above assumes that all technologies – all real numbers - are possible. This, however, is hardly the case in reality: one can imagine products such as chocolate coffee pots that may be feasible, but unviable.

A simple way of tuning the difficulty of the search in this landscape is the addition of “anti-targets” or “valleys.” These are simply numbers that are either not allowed or that incur some kind of a penalty. The density of these anti-targets can be easily adjusted, and thus different technological landscapes can be explored using the Adder model. It should be noted that these anti-targets correspond roughly to sites that are impossible to realize in the lattice percolation model described above.

An example of how the density of targets and anti-targets may be used to tune the difficulty of search for new technologies is shown below in Figure 2, where the y axis reports the highest technology reached at the moment of time.

Keystone technologies

Key variables of the model – the density and spread of targets and anti-targets – can be set to replicate an important finding of Arthur and Polak’s model, namely, that without early low-level needs, more advanced needs are difficult or impossible to satisfy. In other words, certain technologies may serve as “keystone” technologies, enabling further technological development. However, as the Adder is far more tuneable than the original, we have also found that these results depend on the setting of the variables in question. If (nearly) all technologies are feasible (no anti-targets), the lack of early needs does not stop the progress towards more complex technologies.

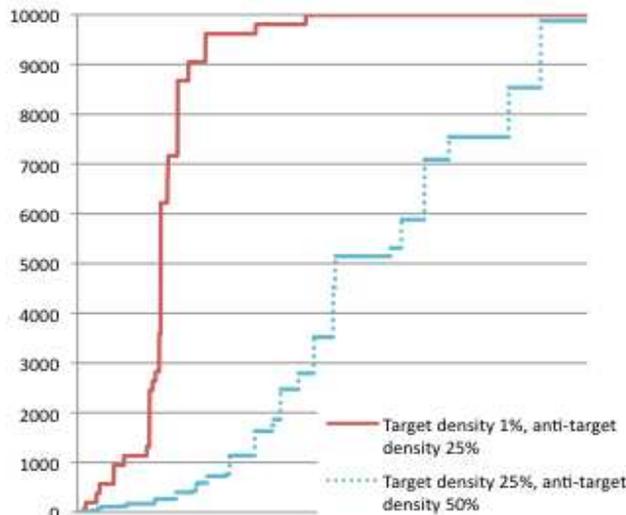


Figure 2: Differences in Performance With Two Different Target/Valley Ratios’

Incremental versus Radical Innovation

The Adder can be used to model either incremental improvement of existing technologies or radical departures from the existing state-of-the-art, or both at

the same time. Both incremental and radical improvement may take place by either randomly redrawing components and operators for existing technologies, or by more “rational” methods.

Multi-agent Industrial Ecosystems

The Adder can serve as a basis for a multi-agent simulation where agents have different search strategies. An industrial ecosystem can be modelled in this manner; for example, the Adder can model the division of labour between component producers and original design manufacturers. Features such as patents and knowledge sharing are possible to implement as well.

EXAMPLE CASE: RATIONAL AGENTS AND (UN)CERTAIN NEEDS

As an example demonstrating the Adder’s basic flexibility, we will briefly detail a somewhat modified, work-in-progress version of the basic Adder model. The goal of the following model is to study how the understanding of technological possibilities and user needs may affect technological development, and how fully rational agents would develop new technologies.

In a review of evolutionary theories of technological change, Nelson (2005) notes that two key variables – the strength of technological understanding and the knowledge of user needs - seem to control the rate and direction of technological advance. If both are very strong, technological advance can almost be planned.

To test this theory, we have coded a model where technologies are developed in a goal-seeking manner, instead of myopic or random searches of previous models. The developer has a perception (possibly incorrect) of user needs, i.e. target values, and a variable understanding of technologies it can use. In each turn of the simulation, a single developer agent attempts to satisfy a single perceived need by combining together available components. When developing a technology, the agent evaluates the expressions

$$T \pm U_t - P \pm U_p \quad (3)$$

and

$$T \pm U_t - C \pm U_c \quad (4)$$

where T is the target, P the product of the technology under development, C the contribution of a component under evaluation, and U_t , U_p and U_c associated uncertainties. If a component is found that

brings the new combination closer to target T , it is added to the combination (so that $P_{t+1} = P_t + C$), and the evaluation round starts again.

The goal of the agent is to get as close to the perceived target as possible using the technologies available at the time, using as few components as possible. The target T is selected from a space of adjacent possible targets. This non-monotonically increasing two-dimensional space is simply the space of those target values that neighbour already discovered target values, but have not been yet discovered (i.e. $T = P$). This represents in a stylized form the observation made by e.g. Arthur (2009) that technological “frontier” advances over time, as new technologies open new combinatory possibilities and create new needs.

The search for new technologies is made more difficult by introducing a variable density of “valleys.” These valleys represent combinations that are unfeasible for any reason. At no point in the development of technology can the combination’s real value, P , equal any valley value. However, the valleys can be “leapfrogged” by adding sufficiently “large” components. (For example, if “3” is a valley value but “4” is not, combinations 2+1 and 2+1+1 are not viable, but a combination 2+2 is.) Viable technologies are added to the repertoire of technologies and can be used as components in the following turns.

As the agent tests different components and repeatedly tries to satisfy a given target, it gains experience of both technology and the needs, enabling it to make more accurate assessments of what components are needed to satisfy a given need. The uncertainty associated with both technologies and needs diminishes according to a standard learning curve model,

$$U = U_i(x + 1)^{\log_2 b} \quad (5)$$

where U_i is the initial uncertainty, x the number of times a technology has been used as a component, and b the learning percentage.

Results

The simulation was run with varying parameter values for initial uncertainty and target and valley densities. All the simulations had learning percentage set to 80%, and one primitive component, “1,” and (+) operator were used. The plots of two representative simulation runs ($n = 5$) are displayed below:

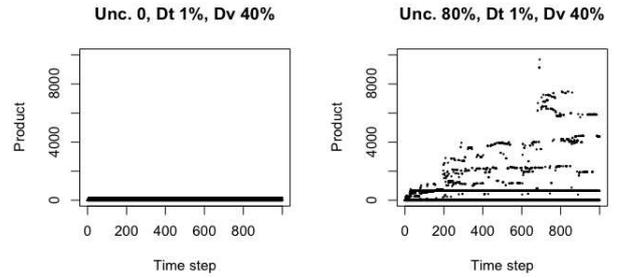


Figure 3: Representative Simulation Results, 5 Runs For Each Setting

On the left panel, the agent has perfect knowledge of user needs and technologies. On the right, it starts with an 80% initial uncertainty about target values and technology performance. In both simulations, the density of target values is 1% and the density of valleys or infeasible technologies 40% out of all the possible technologies. The results are robust to these parameters.

As the sample plots show, surprisingly, uncertainty seems to be a requirement for technological progress. When the agents have perfect knowledge of the needs and the performance of their technologies, they do not make mistakes that help them overcome technological barriers (valleys). Instead, they get stuck after a few non-primitive technologies, at most, are developed.

The preliminary results from this and other test runs indicate that the overall behavior of the model seems to be in line with stylized facts broadly used in other simulations of technological development, e.g. Silverberg and Verspagen (2005). For example, the technological development in the model often displays self-organized criticality and bursts of technological development: usually, the technological development is very slow over a number of turns, until a major valley is bypassed. After that, the speed of development may increase exponentially, until another large valley blocks development – for a while.

Although this example details a simple model, it is easy to see how a more complex model could be developed. For example, the model could be easily extended to n dimensions, and agents with different development strategies and differing knowledge can be added to compete with the single agent of above simulations. As the Adder’s basic premise – numbers and +/- operators – is so simple, the modeler is free to quickly test out variations.

DISCUSSION

In this paper, I have presented an alternative to the two simulation frameworks most commonly used to study technological evolution. The Adder model is designed to be a streamlined version of an alternative

model. Being easy to understand, implement, and modify, I believe it has potential to help researchers interested in the evolution of technologies to study settings hitherto unreachable with existing model frameworks.

Of course, the Adder is not meant to be everything for everyone. NK models are still very well suited for studying complex interactions in a relatively static environment. Similarly, percolation models are perfectly adequate for studying the spread of technologies or concepts over time. A benefit of these simpler models is that they are largely parameterized by a single variable – the K in NK models and the density of available sites in the percolation model. Such simplicity makes for parsimonious models, a desirable feature in most cases. However, the parsimony can also lead to too abstract models, and I believe that unmodified NK and percolation models are not very good fit for studying the evolution of “ideas” over time and in a dynamic environment, i.e. where the ideas themselves have an effect on the selection environment.

The basic framework of using real numbers and arithmetical operators to study technological evolution is extremely flexible and very easily extendable. Thus, it lends itself well to further research. In addition to some of the possibilities outlined above, the model can be used in e.g. studying the effects of constraints on technological development, the effect of different product development strategies, and perhaps even the workings of a broader, interlinked economy of technology-developing organizations. Simpler modifications, such as random mutations, are even easier to implement.

REFERENCES

- Arthur, B. W. (2009). *The Nature of Technology: What it is and how it evolves*. New York: Free Press.
- Arthur, W. B. (2007). The structure of invention. *RESEARCH POLICY*, 36(2), 274–287. doi:10.1016/j.respol.2006.11.005
- Arthur, W. B., & Polak, W. (2006). The evolution of technology within a simple computer model. *Complexity*, 11, 23–31.
- Bak, P., & Wiesenfeld, K. (1988). Self-Organized Criticality: An Explanation for 1/f Noise. *Physical Review A*, 38, 364.
- Basalla, G. (1989). *The Evolution of Technology*. Cambridge, UK: Cambridge University Press.
- Conway, E. M. (2008). *High-Speed Dreams: NASA and the Technopolitics of Supersonic Transportation, 1945–1999*. Baltimore: Johns Hopkins University Press.
- Dahlin, K. B., & Behrens, D. M. (2005). When is an invention really radical? Defining and measuring technological radicalness. *Research Policy*, 34(5), 717–737.
- Dennett, D. C. (1995). *Darwin's Dangerous Idea: Evolution and the meanings of life*. New York: Simon & Schuster.
- Edgerton, D. (2010). Innovation, Technology, or History: What Is the Historiography of Technology About? *Technology and Culture*, 51(July), 680–698.
- Ethiraj, S. K., & Levinthal, D. A. (2004). Modularity and innovation in complex systems. *Management Science*, 50(2), 159–173.
- Fleming, L., & Sorenson, O. (2001). Technology as a complex adaptive system: evidence from patent data. *Research Policy*, 30(7), 1019–1039. doi:10.1016/S0048-7333(00)00135-9
- Frenken, K. (2001). *Understanding Product Innovation using Complex Systems Theory*. University of Amsterdam.
- Frenken, K. (2006). *Innovation, Evolution and Complexity Theory*. Cheltenham and Northampton: Edward Elgar.
- Ganco, M., & Hoetker, G. (2009). NK Modeling Methodology in the Strategy Literature: Bounded Search on a Rugged Landscape. In D. Ketchen & D. Bergh (Eds.), *Research Methodology in Strategy and Management Volume 5* (pp. 237–268). Emerald. Retrieved from <http://www.emeraldinsight.com/journals.htm?articleid=1795453&show=abstract>
- Gavetti, V. D., & Levinthal, D. A. (2000). Looking forward and looking backward: Cognitive and experiential search. *Administrative Science Quarterly*, 45(1), 113–137.
- Gilfillan, S. C. (1935). *The Sociology of Invention*. Chicago: Follett Publishing.
- Katila, R., & Shane, S. (2005). When does lack of resources make new firms innovative? *Academy of Management Journal*, 48(5), 814–829.
- Lemley, M. A. (2012). The Myth of the Sole Inventor. *Michigan Law Review*, 110(5), 709–760. doi:10.2139/ssrn.1856610
- Levinthal, D. A. (1997). Adaptation on Rugged Landscapes. *Management Science*, 43(7), 934–950.
- Murmann, J., & Frenken, K. (2006). Toward a systematic framework for research on dominant designs, technological innovations, and industrial change. *Research Policy*, 35(7), 925–952. doi:10.1016/j.respol.2006.04.011
- Nelson, R. R. (2005). *Evolutionary Theories of Cultural Change: An Empirical Perspective*. New York.
- Rivkin, J. W. (2000). Imitation of complex strategies. *Management Science*, 46, 824–844.
- Rosenkopf, L., & Nerkar, A. (2001). Beyond local search: boundary-spanning, exploration, and

impact in the optical disk industry. *Strategic Management Journal*, 22, 287–306.

Schumpeter, J. A. (1939). *Business Cycles*. New York: McGraw-Hill.

Silverberg, G., & Verspagen, B. (2005). A percolation model of innovation in complex technology spaces. *Journal of Economic Dynamics and Control*, 29(1-2), 225–244.
doi:10.1016/j.jedc.2003.05.005

Simon, H. A. (1982). Reply: Surrogates for Uncertain Decision Problems. *Models of Bounded Rationality, Volume 1* (pp. 235–244). Cambridge, Mass.: MIT Press.

Usher, A. P. (1954). *A History of Mechanical Inventions*. Cambridge, Mass.: Harvard University Press.