LEARNING CURVES

Theory, Models, and Applications

Edited by
Mohamad Y. Jaber

CRC Press
Taylor & Francis Group
Boca Raton London New York

CRC Press is an imprint of the
Taylor & Francis Group, an informa business

CRC Press
Taylor & Francis Group
6000 Broken Sound Parkway NW, Suite 300
Boca Raton, FL 33487-2742

Printed in the United States of America on acid-free paper

International Standard Book Number: 978-1-4398-0738-5 (Hardback)

Visit the Taylor & Francis Web site at
http://www.taylorandfrancis.com

and the CRC Press Web site at
http://www.crcpress.com

23 Developments in Interpreting Learning Curves and Applications to Energy Technology Policy

Bob van der Zwaan and Clas-Otto Wene
In Memoriam: Leo Schrattenholzer

CONTENTS

INTRODUCTION

In 1999, a group of scientists and analysts from academia, industry, and government agencies who were participating in a workshop arranged by the International Energy Agency (IEA) observed that experience and learning curves "are underexploited for public policy analysis" (IEA/OECD 2000, Appendix B, 112). They recommended that experience and learning curves "are used to analyze the cost and benefits of programs to promote environment friendly technologies" and "are explicitly considered in exploring scenarios to reduce CO_2 emissions and calculating the cost of reaching emissions targets" (IEA/OECD, 2000, Appendix B, 114). The IEA Committee on Energy Research and Technology (CERT) supported the findings of the workshop and initiated an international collaboration.

McDonald and Schrattenholzer (2001) provided the first overview of experience curves for energy technology. Particular attention in their seminal publication was given to the distribution of learning rates for 48 cases of different energy technologies. While the statistical significance of a number of the included learning rates was not particularly high, several robust overall conclusions could be made on the basis of their distribution. One of the main findings was that their distribution had both commonalities and differences with respect to those published by Dutton and Thomas (1984) for (not exclusively energy-based) technologies manufactured by individual companies. Both of these studies found a median learning rate of close to 20%, but McDonald and Schrattenholzer (2001) found a higher frequency of smaller learning rates than Dutton and Thomas (1984) did.

The last 10 years have seen a steadily increasing amount of studies being carried out on technology learning as measured by experience curves, initially on renewable technologies, but now also including fossil, nuclear, hydrogen and end-use technologies (Junginger 2010; Schoots et al. 2008; 2010). Recent high-level policy documents embrace the insights from experience and learning curves into the crucial role of market deployment in order to make low-carbon energy technologies more cost-efficient (IEA/OECD 2006, 2008, 2010; Stern 2006; EESC 2009).

The work in the late 1990s on experience and learning curves at the IEA Secretariat established their political implications and global reach (IEA/OECD, 2000), but this work could rely on earlier pioneering efforts in order to apply to learning curves a political message in the energy field. Maycock and Wakefield (1975) and Williams and Terzian (1993) analyzed the reduction in cost of photovoltaic (PV) cells, and Neij (1999) did the same for wind power. Tsuchiya (1989) showed how niche markets can act as stepping stones for the riding down of the experience curve and inspired the launching of Japan's Residential PV System Dissemination and PV-Roof Program (IEA/OECD 2000, 64–74). Early efforts to introduce learning curves into energy-economic models failed because of the strong non-linearities of the curves and the increasing returns to scale that the introduction of these curves generated in these models (Manne 1994). Technology learning was then introduced in the technology modules of the U.S. National Energy Modeling System, but it proved the model could only be solved iteratively (Kydes 1999). Messner (1997), and Mattsson and Wene (1997) first managed to introduce learning curves in global bottom-up energy system models and solve their models in single-solution runs, while van der Zwaan et al. (2002) were the first to achieve this in a top-down version of energy-economic-environment (EEE) models. Following these efforts, technology learning is now a standard feature in most bottom-up energy systems and top-down EEE-integrated assessment models (Kahouli-Brahmi 2008).

The reason for the political relevance of experience and learning curves is that technology learning is a key factor behind technological change (Sagar and van der Zwaan 2006). The curves are therefore important tools in answering the question of how to retain economic growth while using technology to transform the global energy system so that it can sustain a low-carbon economy. The insights from experience and learning curves suggest that such a transformation can be made, but initial investments in learning will be in the order of hundreds of billions of U.S. dollars (IEA/OECD 2000, 2008; van der Zwaan and Rabl 2004). This raises the questions of how well we understand the technology learning process, and how well we can trust forecasts based on this understanding.

Although high-level policy reports acknowledge the existence of learning effects, they also point to the large uncertainties in the estimates of future learning phenomena. Because of the long time spans involved, these uncertainties translate into large uncertainties about the resources or learning investments that are needed to bring

the new technologies to breakeven with incumbent high-carbon technologies. The reports stop short of recommending experience and learning curves as operational policy tools. In their report to the 2006 G8 meeting of heads of state, the IEA finds:

> "Technology learning is the key phenomenon that will determine the future cost of renewable power generation technologies. Unfortunately, the present state of the art does not allow reliable extrapolations." (IEA/OECD 2006, 231)

The Stern review observes that:

> "There is a question of causation since cost reductions may lead to greater deployment; so attempts to force the reverse may lead to disappointing learning rates. The data shows technologies starting from different points and achieving very different learning rates." (Stern 2006, 362, 411–412)

A key criticism is that the curves appear to express purely empirical relations between cost, price or technical performance, and cumulative production or use. Theoretical grounding is needed to explain observed learning rates, limit uncertainties in extrapolations, and legitimize government deployment programs. Another important point of criticism is that the set of technologies with associated learning curves is biased: many technologies exist for which no learning has been observed, for which learning has inadvertently stopped, or for which, in fact, diffusion itself has halted (or has even returned) as a result of unforeseen (and unforeseeable) events (Sagar and van der Zwaan 2006). Furthermore, it is argued that reported learning rates often confuse at least two (and sometimes more) economic phenomena that are fundamentally different in nature. Learning-by-doing versus economies-of-scale is a proper point in case, since the former is realized by cumulative capacity or activity, while the latter is dominated by the size of the plant or sector under consideration at a given moment in time (for a recent example, related to hydrogen technology, in which these phenomena are disentangled, see Schoots et al. 2008, 2010). Several mechanisms have been proposed to explain technology learning and the observed relationships (Abell and Hammond 1979; Arthur 1988; Argote and Epple 1990; Adler and Clark 1991; Nemet 2006). Unfortunately, they generally fail to reconstruct the shape of the learning curves or explain the observed learning rates.

In this chapter we provide some interpretations of experience and learning curves starting from three different theoretical platforms. These interpretations are aimed at explaining learning rates for different energy technologies. The ultimate purpose is to find the role that experience and learning curves can legitimately play in designing efficient government deployment programs and in analyzing the implications of different energy scenarios. The "Component Learning" section summarizes recent work by the authors that focuses on the disaggregation of technologies in their respective components and argues that traditional learning for overall technology should perhaps be replaced by a phenomenology that recognizes learning for individual components. The "Learning and Time" section presents an approach that departs more strongly from the conventional learning curve methodology, by suggesting that exponential growth and progress may be the deeper underlying processes behind observed learning-by-doing. Contrary to this view, the cybernetic approach presented in the "Cybernetic Approach" section sees learning curves as expressing a fundamental property of organizations in competitive markets and applies the findings from second order cybernetics to calculate the learning rates for operationally closed systems. All three interpretations find empirical support. The "Conclusions" section summarizes the pros and cons of the three approaches from an energy technology policy perspective and provides some concluding remarks.

COMPONENT LEARNING

Ferioli et al. (2009) investigate the use of learning curves for the description of observed cost reductions for a variety of energy technologies, using as a starting point the representation of energy processes and technologies as the sum of different components. While it is widely recognized that in many cases learning-by-doing may improve the overall costs or efficiency of a technology, they argue that, so far, insufficient attention has been devoted to studying the effects of single-component improvements, which, taken together, may explain an aggregated form of learning. Indeed, for an entire technology, the phenomenon of learning-by-doing may well result from the learning of one or a few individual components only. They analyze under what conditions it is possible to combine learning curves for single components to derive one comprehensive learning curve for the total product. The possibility that, for certain technologies, some components (e.g., the primary natural resources that serve as essential input for their fabrication or use) do not exhibit cost improvements might account for the observed time-dependence of learning rates as reported in several studies. The learning rate might also change considerably depending on the (sub)set of data considered, a crucial issue to be aware of when one uses the learning curve methodology. These observations may have important repercussions for the extent to which learning curves can be extrapolated in the future.

Learning-by-doing has often been shown to slow down in the long term. A way to describe such slowing down is to consider a product, process, or technology as an aggregate of several components. The cost C of every industrial product can be expressed as the sum of the costs of its components. If one assumes that the cost of each component decreases over time according to a power law relation as a result of learning, it is possible to write the overall cost relation of a generic product as:

$$C(x_t) = \sum_{i=1}^{n} C_{0i}\left(\frac{x_{ti}}{x_{0i}}\right)^{-b_i} = C_{01}\left(\frac{x_{t1}}{x_{01}}\right)^{-b_1} + C_{02}\left(\frac{x_{t2}}{x_{02}}\right)^{-b_2} + \ldots + C_{0n}\left(\frac{x_{tn}}{x_{0n}}\right)^{-b_n}, \quad (23.1)$$

in which the index i represents a given cost component. Each component is, in principle, characterized by a different learning parameter b_i, and a different initial cumulative production x_{0i}. Aggregate learning may or may not be broken down into component learning according to Equation 23.1. At any rate, the value of cumulative production for each component is at least as important as the individual learning parameter. The reason for this is that, between components, x_{0i} may have widely diverging values, and along with b_i, x_{0i} determines how much scope exists for future learning. For example, the production of wind turbines has a negligible effect on the historic cumulative production of steel or aluminum, so that not much production cost reduction for these construction materials (needed for components like the support mast and turbine housing) can be expected by the deployment of wind turbines. On the other hand, continued cost improvements can be expected for the fabrication of (lightweight) rotor blades that so far have reached a much more limited cumulative production. It is therefore necessary to discuss, both in general and for each technology independently, under what conditions the conventional learning curve equation can be broken down into the component learning expression of Equation 23.1. Conversely, one may question when the validity of an equation of the form of Equation 23.1, can be demonstrated, and can it be approximated by the conventional expression for learning curves; that is, with only one term.

For ease of exposition, Ferioli et al. (2009) analyze the properties of a simplified model in which the cost for a product or technology is determined by only two components – one characterized by learning, and one for which the cost is constant in time (i.e., no cost reduction can be observed). If α is the share of the total cost that initially can be attributed to the learning component, then $1-\alpha$ is the beginning cost share of the non-learning component. The overall cost as a function of the cumulative production of the learning component can, in this simplified case, be expressed as:

$$C(x_t) = \alpha C_0\left(\frac{x_t}{x_0}\right)^{-b} + (1-\alpha)C_0, \quad (23.2)$$

where C_0 is the total cost at production level x_0. Equation 23.2 can be considered a special case of the more elaborate model presented by Equation 23.1 and can be useful to highlight some properties of the latter. A theoretical justification for this model is the observation (see, e.g., Schoots et al. 2008; Schoots et al. 2010; van der Zwaan and Rabl 2004) that some parts of a technology, such as the raw materials and labor, may not experience cost reductions or may become more expensive over time. The precise value of α can, in principle, be determined for each technology, which we find may be a valuable avenue for further work. Equation 23.2 assumes that the learning component is the innovative part of the new total product, so that the cumulated production of the learning component and the overall technology are the same. It is also supposed that the component does not improve from simultaneously being part of another technology, so that the capacity of the composite and its learning component evolve synchronously.

One can test whether data that is usually fitted with a traditional learning curve which does not distinguish between different components can be fitted on the basis of Equation 23.2. Figure 23.1 shows that data for the price of gas turbines from MacGregor et al. (1991) can be fitted in two different ways. The data points can be fitted with a learning curve over a range spanning three orders of magnitude of cumulative deployment. The learning rate (LR) is found to be 13% with $R^2 = 0.95$ (Figure 23.1, left plot). It can be observed, however, that the data present an inflection point. In the literature (see notably Seebregts et al. 1999), it has been proposed to fit such data with a piecewise-linear learning curve. Figure 23.1 (right plot) shows a two-piece learning curve applied to the present case: one obtains LR = 19% and LR = 10%, and $R^2 = 0.97$ and $R_2 = 0.94$, for the two learning-curve pieces,

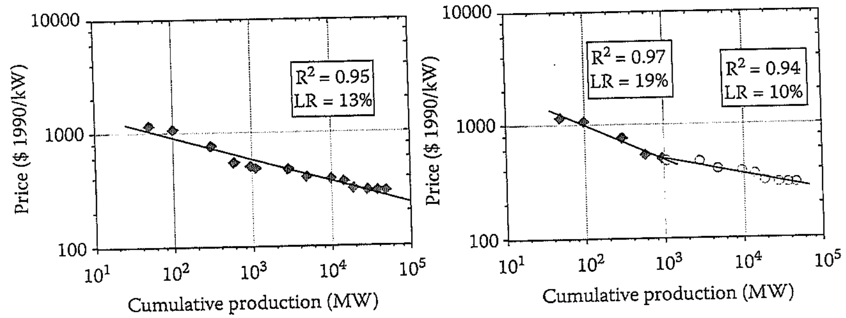

FIGURE 23.1 A set of gas turbine prices fitted in two different ways: linearly and piecewise linearly. (Data from MacGregor et al., *The market outlook for integrated gasification combined cycle technology*. General Electric Company, 1991.)

respectively. Given the empirical nature of learning curves, the fact that the learning rate changes over time leads to methodological issues: a constant learning rate is one of the fundamental assumptions of the learning curve methodology.

The same set of data can also be fitted with an expression of the form of Equation 23.2, as shown in Figure 23.2. For example, the system can be described as composed of a learning component (with LR = 24%) that makes up 80% of the total cost, and a non-learning part (hence with constant cost) that accounts for the remaining 20%. One could think of the costs associated with the steel as being necessary to fabricate the turbine. This fit is clearly better than both of those shown in Figure 23.1, since it represents overall an $R^2 = 0.97$, while applying to the same data set. In any case, if one finds learning curves acceptable when they possess an accuracy of $R^2 > 0.90$, one cannot discard any of these regressions—certainly not the two-component one.

While one may find all three fits to the gas turbine price data of Figures 23.1 and 23.2 acceptable, clear differences occur when one uses these learning curves for price

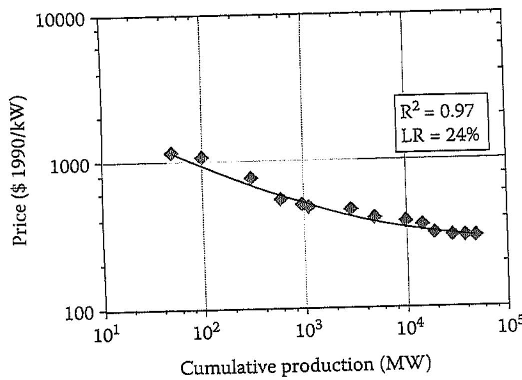

FIGURE 23.2 The fit proposed for the gas turbine price data from Figure 23.1, based on Equation 23.2 with LR = 24% (for the learning component) and $\alpha = 0.8$.

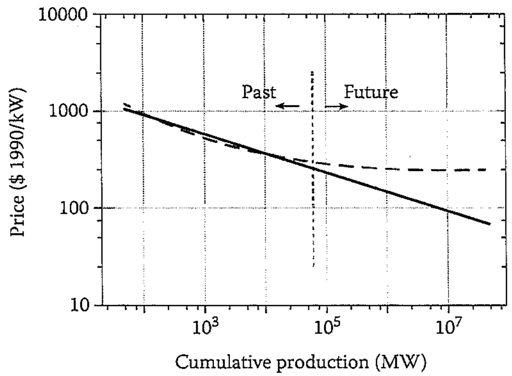

FIGURE 23.3 Two fits for gas turbine prices extrapolated over three orders of magnitude into the future: A linear regression based on the traditional learning curve (——) and one based on Equation 23.2 (– – –).

or cost estimates in the future. Extrapolating cost data over several orders of magnitude of cumulative production can lead to significant errors in the estimates of both the breakeven capacity and learning investment (needed to reach competitivity with the incumbent technology) when one uses the wrong learning model. We point this out in Figure 23.3, in which both the fit of the left plot of Figure 23.1 and that of Figure 23.2 are depicted and extrapolated over three orders of magnitude in the future. Indeed, we see that the two lines diverge rapidly for higher values of cumulative production, with obvious consequences in terms of, for example, the total learning investment needed to reach a given cost level in the future. As the cost of the innovative component is reduced, the non-learning component gains in relative weight in its contribution to the overall costs. Hence the composite learning process is slowed down.

The above argument suggests that technology cost reductions may not continue indefinitely and that well-behaved learning curves do not necessarily exist for every product. In addition, even for diffusing and maturing technologies that display clear learning effects, market and resource constraints can eventually reduce the scope for further improvements in their fabrication or use. It appears likely that some technologies, such as wind turbines and photovoltaic cells, are significantly more amenable than others to industry-wide learning. For such energy technologies, Ferioli et al. (2009) assess the reliability of using learning curves to forecast cost reductions. They argue that due attention must be paid to cost components that do not decrease over time, or may even increase.

LEARNING AND TIME

In another paper, Ferioli and van der Zwaan (2009) analyze the dynamics for the growth and cost reduction of innovative products in the energy sector. They provide a series of examples showing that simple exponential relations can be used to describe growth and cost reduction as functions of time for many types of technologies. These two simple models, for technological growth and progress respectively, when taken together, are shown to be a de facto equivalent to the well-known learning curve. They propose a stylistic computational component-based model that accounts for both these exponential relationships. The main novelty of this model is that it introduces time in the learning curve methodology. While there may be additional explanatory variables, they argue that accounting for time improves the understanding and use of learning curves.

In their article, they first point out that much attention has been paid over the course of the past two decades to how technology development feedback and inducement may be included in theories of economic growth (for an overview see Aghion and Howitt 1997). This "endogenous growth theory" may benefit greatly from the abundance of detailed technology assessment studies in the innovation literature (e.g., Grübler et al. 1999). As these technology assessments confirm, arguably the simplest model to mathematically represent (the early phase of) technological diffusion assumes growth with a fixed percentage, β, each year. Hence, the total output increases by a factor $\alpha = 1 + \beta$ each year, and the output in year $n + 1$, $y_n + 1$, is simply $y_n + 1 = \alpha y_n$ (with y_n being the output in year n). Time can be uniformly normalized by dividing by the unit of one year in order to ensure the dimensional consistency of the equations. Without a loss of generality, an output can be assumed of one in year 1. The annual output as a function of time can thus be written as:

$$y(t) = \alpha^t. \tag{23.3}$$

The cumulative production, $x(t)$, is found by integration of Equation 23.3, which (unsurprisingly) is again an exponential in t (with an extra multiplicative factor and a constant introduced by integration). Therefore the growth model for cumulative production can be approximated by a simplified exponential relation of the form:

$$x(t) = ab^t, \tag{23.4}$$

in which a and b are fitting parameters. The parameter b of Equation 23.4 is thus not exactly identical to the annual growth parameter α in Equation 23.3. For many products, companies and technologies, growth in cumulated production can be approximated by this reduced exponential model. Ferioli and van der Zwaan (2009) give several examples for non-energy technologies, based on data from Goddard (1981). Likewise, they show examples for energy technologies, which include PV modules, wind turbines, sugarcane-based ethanol production in Brazil, and large stationary fuel cells. The historic cumulative production for these technologies and products is fitted with Equation 23.4 based on data from Harmon (2000), Neij (2004), Goldemberg (1996), and Adamson (2006), respectively. The fits prove good to excellent.

Having developed a simple model for technological expansion, they then focus on the dynamics of progress. They assume that cost reduction occurs through increases in productivity, such that the unit cost is reduced by a fixed percentage every year. Hence, the cost in year $n + 1$ equals $C_{n+1} = \gamma C_n$, with $\gamma = 1 - \delta$ and δ the productivity gain. Cost as a function of time can thus be expressed by an exponentially decaying relation of the form:

$$C(t) = df^t, \tag{23.5}$$

in which d and f are fitting parameters ($f < 1$). For several technologies, it is demonstrated that costs as a function of time can be approximated by Equation 23.5, based on data from Goddard (1981), Harmon (2000) and Goldemberg (1996), respectively.

Ferioli and van der Zwaan (2009) show that the elimination of time from Equations 23.4 and 23.5 leads to a power law relation between the cost and cumulative output:

$$C(x) = d\left(\frac{x}{a}\right)^{\frac{\ln f}{\ln b}}. \tag{23.6}$$

Equation 23.6 is a de facto equivalent to the learning curve. Parameters a and d are equivalent to the usual normalization point (x_0, $C(x_0)$), and since $f < 1$ (that is, costs are decreasing over time) and $b > 1$ (output is increasing over time), the exponent in Equation 23.6 is negative, so that Equation 23.6 represents a learning curve (with LR $= 1 - 2^{\ln f/\ln b}$).

They next construct a concise model based on the exponential relations for growth and technical progress (cost reduction) as expressed by Equations 23.4 and 23.5. For the cost part of their numerical model, they consider the overall cost of a product as the sum of the costs of its components (as in Ferioli et al. 2009). For example, the cost of individual parts plus the cost of assembling and marketing may make up the total cost of a given consumer product. As with output growth, they suppose that cost reduction takes place through a stochastic process. Cost reduction occurs in a number of finite steps that reduce the cost of each component by a fraction. The magnitude and number of the cost reduction steps are not fixed, but vary randomly for each step respectively. The purchase of a new automated machine, for example, may in one step reduce the cost of the manufacturing part of a final product by x%. The cost reduction for each component at each step materializes randomly in time, that is, not as a function of the cumulative production. The cost part of the model is graphically depicted in Figure 23.4.

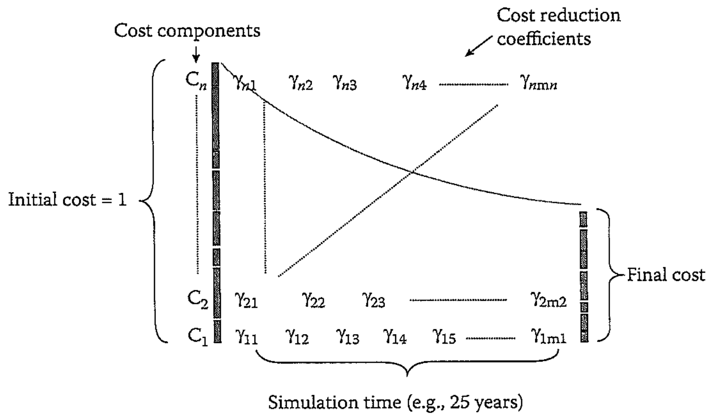

FIGURE 23.4 Illustration of the cost reduction model in Ferioli and van der Zwaan (2009). The overall technology consists of n components, the cost of each decreasing step-wise over time on the basis of a multiplication of positive coefficients $\gamma_{ij} < 1$.

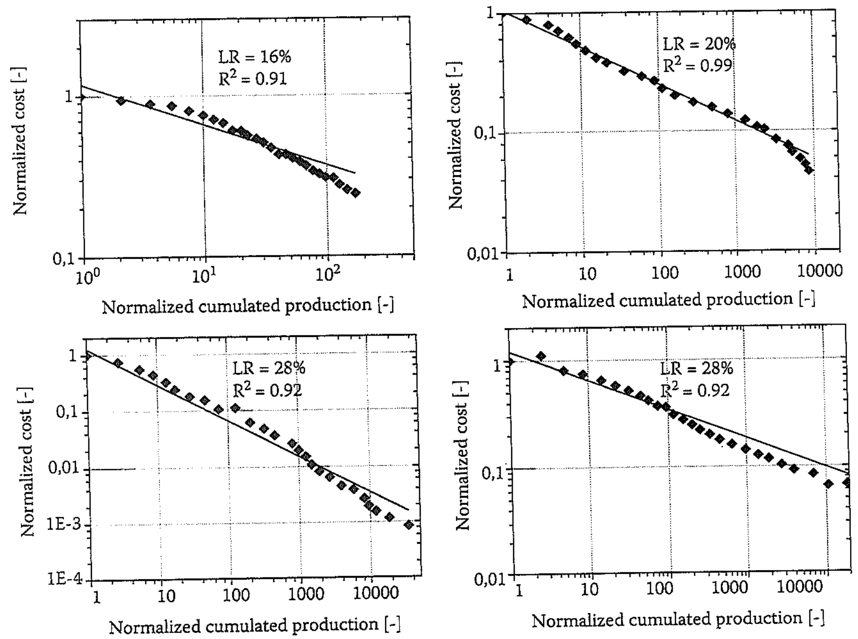

FIGURE 23.5 Cost reduction as a function of cumulative production, as calculated with the simple stochastic model in Ferioli and van der Zwaan (2009) (♦) and best fits with the power law relation of the traditional learning curve (—).

With the two generators of data for technological growth and progress in this simple numerical model, and once the few parameters involved have been appropriately chosen, the calculated trends for the output growth and the cost reduction prove to simulate observed evolutions of these variables for many products and technologies. It is demonstrated that, taken together, these two sub-models are equivalent to the learning curves as reported in the literature, which is confirmed by the results reported in Figure 23.5.

As an additional veracity check to validate their model, Ferioli and van der Zwaan (2009) compare their learning curve simulations directly to cost-versus-output data as available in the literature. Figure 23.6 shows two arbitrary examples of reported learning curve data, for the cost and cumulative output of PV modules (Harmon, 2000) and Brazilian ethanol (Goldemberg 1996) respectively. The numbers for PV, collected over more than two decades, represent an average annual growth rate during this period of as high as approximately 40% per year (Harmon 2000). During this interval the cost of PV modules decreased by more than an order of magnitude. For Brazilian ethanol, the average growth rate was significantly lower, at approximately 10% per year, during a little less than two decades. The ethanol production cost reduced by about a factor of three during these years (Goldemberg 1996). Ferioli and van der Zwaan (2009) contrast these historic data with results from their simulations. The main assumptions in their numerical model concern the average annual

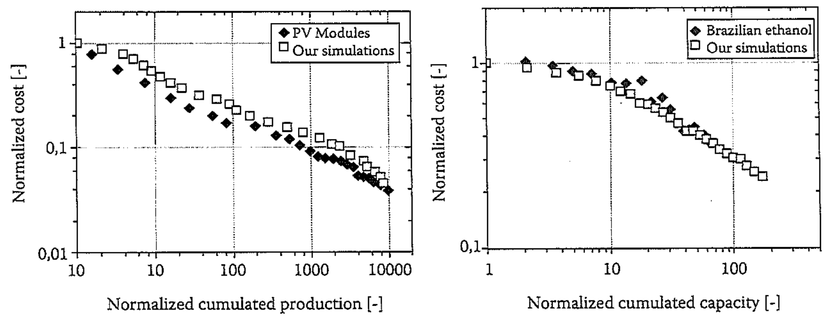

FIGURE 23.6 Comparison of learning curve data for PV (Harmon 2000) and Brazilian ethanol (Goldemberg 1996) (♦) with data generated by the model simulations by Ferioli and van der Zwaan (2009) (□).

growth rate over the whole period under consideration, the standard deviation of this yearly growth, and the final total cost reduction, which are all taken so as to stay close to the real-life values of these quantities. It proves that their model can properly reproduce the learning curves published for these two cases. It can be shown that other learning curves can be simulated in a similar way through using this model.

CYBERNETIC APPROACH

The learning curves for a technology show the performance of a human interactive system producing and/or operating the technology. Is this learning system best understood as an open system reacting to demands and opportunities in the environment, or as a closed system acting autonomously based on its internal structure? Most explanations of learning curves and technology learning focus on the role of environmental interactions in explaining the phenomenon. The cybernetic approach (Wene 2007, 2008a, 2008b) takes the opposite view by assuming that all performance improvements are the result of closed chains of internal operations. Features of the environment, and events and processes in the environment, appear as perturbations, and the system may decide to modify its operations, but these modifications are always conditioned by the internal structure of the system. In this approach, learning rates are the result of the eigenbehavior (Varela 1979; von Förster 2003) of the system. The underlying theory provides a spectrum of *eigenbehavior* values for the unperturbed system in a perfect competitive environment, while the basic eigenbehavior (or eigenvalue) corresponds to a learning rate of 20%. Perturbations from the environment result in the observed broad distribution of learning rate values (see, e.g., McDonald and Schrattenholzer 2001) around the unperturbed eigenbehavior value(s) (Wene 2010).

The cybernetic approach applies theoretical results for biological and social systems (von Förster 1980, 2003; Varela 1979, 1984; Luhmann 2002) to technology learning systems, among others in the energy field. This section exposes the basic features of this approach and provides the results for learning rate distributions. The autonomy of the system is expressed as *operational closure*, meaning that the system forms and controls all its operations. While the system is open to information, and in this case to material and energy flows, the network of internal operations closes on itself. The closure theorem of cybernetics states that in every operationally closed system, eigenbehaviors arise. The task of the analyst is to find the operational loops that represent learning and define the operators whose fixed points provide the values for the eigenbehavior. Guides for this task are the form of the learning curve and the well-known OADI-SMM (observe, assess, design and implement—shared mental model) model for organizational learning (Kim 1993; Espejo et al. 1996). Occam's razor, which is generally a good principle in theory design since it urges for a minimum of elements and assumptions, possesses particular relevance when specifying the loops and operators involved in OADI-SMM.

The OADI mnemonic emphasizes learning as a result of self-reflection; that is, making new designs after assessing the observations of one's own action, which in this case is the implementation of previous design efforts. This is consistent with the assumption of operational closure and is captured in the three feedback loops in Figure 23.7. The hypothesis is that these loops, taken together, provide the elements needed to explain technology learning as expressed by learning curves. The internal and external feedback loops express implementation and observation. The external loop closes over the market and the internal loop over producing. Together, with the self-reflecting, third loop (SRL), they provide the double closure proposed by von Förster (2003) as the process required for an organism or an organization to modify its behavior in order to manage environmental perturbations without losing operational closure. The internal and external loops reflect the double closure over production and sales as analyzed by Baecker (1996). The self-reflection loop represents assessment and design and operates on the internal state, Z. This internal state sets the transfer function of producing. In other words, through its actions on Z, the SRL determines the performance (i.e., the relation between the input and output) of the learning system. The computing block integrates the three loops. The SRL thus drives the learning of the system, but is for its operations totally dependent on the status of the two other loops.

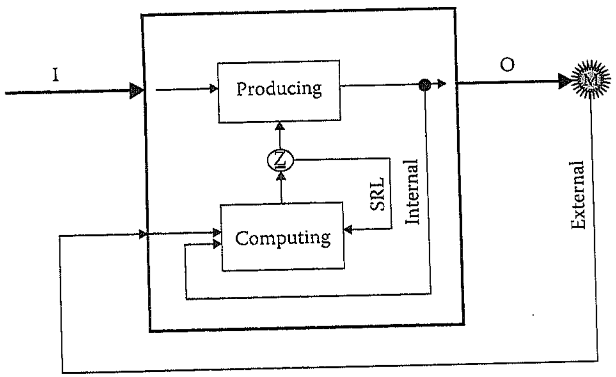

FIGURE 23.7 The elements of the technology learning system.

In order to find the mathematical expression for the operators describing the action of the three loops on the state function, we observe that the learning curve can be written as:

$$P(t) = \text{output/input} = C_0\ X(t)^E, \tag{23.7}$$

in which $P(t)$ is the performance of the system at time t. If all inputs are measured in monetary units, then the performance is the output generated from one monetary unit of input (which is equivalent to the inverse of the cost per unit output). In this equation, C_0 is a constant and $X(t)$ is the sum of all output until time t, while E is referred to as the "experience parameter." The progress ratio (PR) and learning rate (LR) are, as usual, given by:

$$\text{PR} = 1 - \text{LR} = 2^{-E}. \tag{23.8}$$

Taking the logarithm of Equation 23.7, and differentiating plus rearranging provides:

$$X/\text{d}X \times \text{d}P/P \times 1/E = 1 \tag{23.9}$$

Let us consider the unperturbed case. The external loop carries the urge from the perfect competitive market to improve performance by as much as possible. Without perturbations, this loop can be seen as a constant factor. The internal loop tracks output from production, which is represented by the right-hand-side of Equation 23.7. The SRL sets performance via Z, which is represented by the left-hand-side of Equation 23.7. Occam's razor suggests using two operators that work on two independent parts of the state function, which we will call C^+ and C_{SRL}.

One more aspect of the learning curve helps to constrain the choice of operators and initial state function. The learning curve requires that the identity of the output is retained, which can be interpreted as implying that the technical properties of the output pertaining to its original use remain the same. The importance of self-reflection and identity suggests searching for solutions based on complex algebra. For the unperturbed case, the identity requirement can be interpreted as constraining all solutions to the unit circle in the complex Argand plane. C^+ becomes a stereographic operator projecting $X(t)$ onto the unit circle, so that $X(t) \to \infty$ is represented by $-i$ in the complex plane. In the aggregated universe of the learning curve, the C^+ operator corresponds to corporate memory; that is, the SMM in the OADI-SMM.

Equation 23.9 suggests that the C_{SRL} operator is a function of the cumulative output (X) and additional output dX. Wene (2007, 2010) presents arguments for the mathematical form of this operator consistent with Equation 23.9 and with research on individual and organisational learning. Ultimately, however, as in all theoretical research, the unique form of C_{SRL} must be justified by its ability to explain and forecast empirical results.

This theory now yields values for the experience parameter, which correspond to the eigenbehaviors of the technology learning system in the unperturbed case:

$$E(n) = 1/\left[(2n+1)\cdot\pi\right] \qquad \text{with } n = 0,1,2,3,..., \tag{23.10}$$

while the corresponding learning rates are:

$$\mathrm{LR}(n) = 1 - 2^{-E(n)} \qquad \text{with } n = 0,1,2,3,..., \tag{23.11}$$

and the first four modes of learning are:

$$\mathrm{LR}(0, 1, 2, 3) = 20\%, 7\%, 4\% \text{ and } 3\%. \tag{23.12}$$

Changes in eigenbehavior due to perturbations in the external loop can be analyzed without introducing new operators. The double closure between internal and external loops, which modifies the operations in the self-reflecting loop, recursively acts on the system state with an operator matrix with the two basic operators as elements. The result is a distribution of learning rates around the eigenvalues for the unperturbed case. The form and width of this distribution depend on the nature of the perturbations. The latter may appear as negative in the form of environmental restrictions, or as positive in the form of learning spill-overs from other enterprises or industries.

Figure 23.8 from Wene (2010) compares theoretical and empirical distributions of experience parameters. The empirical distributions are from the compilations by, respectively, Dutton and Thomas (1984), based on cost measurements in individual enterprises (Figure 23.8a), and Weiss et al. (2010), for energy supply technologies based on market prices (Figure 23.8b). Negative and positive perturbations are assumed to be additive and Poisson distributed. The assumptions on perturbations are simple but provide good fits and some initial observations.

The Dutton and Thomas (1984) distribution of individual enterprise technologies shows no significant influence from higher order learning modes. The distribution of learning rates for energy supply technologies (Weiss et al. 2010), on the other hand, cannot be reproduced without higher learning modes. The theoretical curve depicted in Figure 23.8b only includes the distribution around the first higher learning rate mode, LR(1), but the fit could be improved by also including higher modes. One can speculate about the cause for the difference between these two distributions. One reason may be that the learning systems for some energy technologies are not properly closed. The network of operations may have side-loops outside of the system; that is, the chosen system boundaries are too narrow to provide full operational closure. A case in point may be wind turbines, for which measurements indicate learning rates of 4–8% that may be representative of higher learning modes (Neij 1999; Durstewitz and Hoppe-Kilpper 1999). Indeed, it proves that increasing the system boundary to encompass the complete wind power plant (Neij et al. 2004) or including the production of electricity from wind power (IEA/OECD 2000; Neij et al. 2004) provides learning rates representative of the zero learning mode. The results of Junginger et al. (2005) for wind farms point to the need to widen the boundaries from national to global learning systems. An analysis of the organisation of the entire wind industry, rather than of just the technology itself, may thus be required to understand the low learning rates observed for wind turbines. This example at least points towards the need for the systematic organisational studies of entire industrial sectors in order to understand learning systems that have enough autonomy to provide operational closure.

Cost analysis requires the assessment of the network of operations within a firm. This aids the investigator in avoiding boundaries which do not provide operational

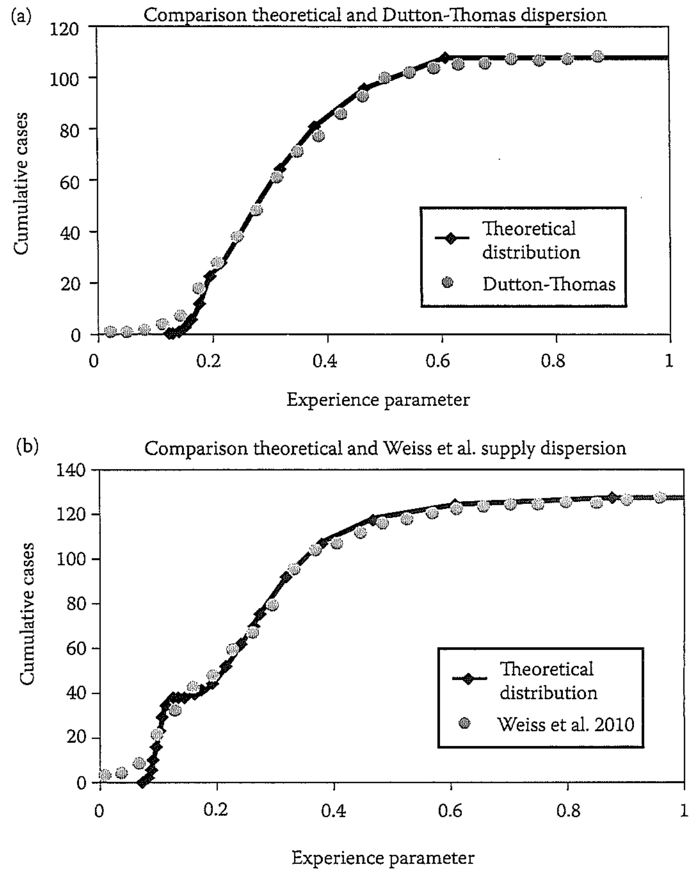

FIGURE 23.8 Comparison of theoretical and empirical distributions. (From Wene, C.O. 2010. Adaptation in technology learning systems. *The proceedings of the 11th IEAA European conference*. 25–28 August, 2010, Vilnius, Lithuania.)

closure. This may be one explanation for the absence of higher learning modes in the Dutton and Thomas (1984) distribution. Operations can, of course, be disturbed by external features, events, and processes. Why then should such disturbances influence the industry but not the firm? This suggests another explanation for the lack of higher learning modes in Figure 23.8a—namely the choice of firms and technologies. Many of the energy technologies in the distribution by Weiss et al. (2010) are regulated—especially fossil and nuclear technologies, which are exposed to insistent but changing external perturbations, related, for instance, to emissions and safety standards. Such perturbations may be too strong to manage through double closure, thus forcing the system into a higher learning mode. Wene (2008b) discusses government spending on public research and development (R&D) as another potentially disturbing factor in the system's environment for producers and users of energy technologies. The purpose of the spending is typically to produce knowledge that increases learning, but systemically this external production of knowledge represents a perturbation that disturbs the internal network of operations. Prospects for radical innovation through public R&D may yield important rethinking or second-loop learning, which may lead to swift cost reductions and improvements in efficiency. External R&D that represents a continuous series of minor improvements, however, may disrupt the system's own learning operations and leave it in a higher learning mode with a lower learning rate.

CONCLUSIONS

This chapter has presented three approaches, which at increasing theoretical depth search to explain the learning curve phenomenon. It is not possible at this point to strike a total balance between the approaches, but following the declared purpose of the chapter we here summarize some of their merits, and present potential deficiencies from an energy technology policy point of view.

Component learning provides a quality-controlled methodology to construct a learning curve for a new technology relying on components for which knowledge exists on learning parameters. A typical example is the fuel cell, but the approach applies in principle to any technology. The usefulness of this method for technology policy is obvious, since it provides refined estimates of learning rates for new technology, and can assist in estimating joint learning for two technologies with similar components. While it provides limited insight into the learning process itself, the component learning approach may increase the reliability of forecasts, particularly in the energy field. In its present form, however, it ignores the possibility of interference between components, which may lead to lower or higher rates of learning.

Factoring the learning curve into market expansion and cost dynamics provides new ways to analyze and extract meaning from market and cost time series. Learning curves for different technologies can be simulated through the use of a small set of fitted parameters. This method widens the theoretical base for learning curves to economic growth and innovation theories, but a further exploration and backing of this base is both desirable and necessary. This may further increase the reliability of technology forecasts based on this approach. In order to provide proper forecasts, however, cost parameters would need to be explained in terms of the measurable properties of new technology, or of the learning system behind new technology.

The cybernetic approach finds learning rates around 20% to be a natural consequence, or an eigenvalue, of the operations of a learning system and independent of any parameter value. This method argues for continuous and reliable learning over the lifetime of a technology. Measured dispersions around the eigenvalues can be reproduced with a few fitted parameters. Like the "learning and time" method, this approach widens the theoretical base for learning curves—in this case to second-order cybernetics and to models for organizational learning—but also here the theoretical foundation should be further investigated. Also, the approach has yet to explain how parameters fitted to the measured distributions of learning curves relate to actual perturbations to the learning system, and how learning parameters for a specific technology can be calculated from the properties of the learning system and its environment.

From the viewpoint of energy technology policy, the three approaches appear to be complementary rather than competing. They strengthen learning curves as policy tools but require further development and grounding in empirical results. Over the last decade, measuring, understanding, and applying learning curves for energy technologies has developed into a large field. Our ambition with this chapter has not been to provide a complete picture of the theoretical developments in this domain, nor have we been in a position to consistently list all the energy technologies for which learning curves have been determined. At the time, McDonald and Schrattenholzer (2001) provided such a listing, although since their seminal publication this methodology, as applied to the energy sector, has progressed substantially. However, we hope that this chapter provides the reader with an insight into at least some of the recent developments, and notably those with which the authors are most familiar. We have undoubtedly left out subjects that were close to the heart of one of our great inspirers in this field, but we are confident—having learned from him before his unexpected early departure—that he would have enjoyed the exploration of these alternative approaches.

ACKNOWLEDGMENT

This chapter was written in memory of Leo Schrattenholzer, to whom both authors are indebted. The authors acknowledge the support of research grants from the Netherlands Organization for Scientific Research (NWO, under the ACTS Sustainable Hydrogen Program, projects 053.61.304, 053.61.305 and 053.61.024) and from the Swedish Energy Agency (project 32179-1).

REFERENCES

Abell, D.F., and Hammond, J.S., 1979. Cost dynamics: Scale and experience effects. In: *Strategic planning: Problems and analytical approaches*. Prentice Hall: Englewood Cliffs.

Adamson, K.A., 2006. *Fuel cell today, market survey: Large stationary applications*. http://www.fuelcelltoday.org/online/survey?survey=2006-11%2F2006-Large-Stationary (accessed May 22, 2008).

Adler, P.S., and Clark, K.B., 1991. Behind the learning curve: A sketch of the learning process. *Management Science* 37(3): 267–281.

Aghion, P., and Howitt, P., 1997. *Endogenous growth theory*. Cambridge: MIT press.

Argote, L., and Epple, D., 1990. Learning curves in manufacturing. *Science* 247(4945): 920–924.

Arthur, B., 1988. Competing technologies: An overview. In *Technical change and economic theory*. eds. G. Dosi, C. Freeman R. Nelson, G. Silverberg and L. Soete. London: Pinter.

Durstewitz, M., and Hoppe-Kilpper, M., 1999. Using information of Germany's '250 MW wind'-programme for the construction of wind power experience curves. In *Proceedings IEA workshop on experience curves for policy making – The case of energy technologies*, eds. C.O. Wene, A., and Voss, T. Fried. 10–11 May 1999, Stuttgart, pp. 129–134, Forschungsbericht 67, Institut für Energiewirtschaft und Rationelle Energieanwendung, Universität Stuttgart.

Dutton, J.M., and Thomas, A., 1984. Treating progress function as a managerial opportunity. *Academy of Management Review* 9(2): 235–247

EESC, 2009. Towards an eco-efficient economy – Transforming the economic crisis into an opportunity to pave the way for a new energy era. Exploratory Opinion of the European Economic and Social Committee, TEN/398-CESE 1700/2009, Brussels, 5 November 2009.

Espejo, R., Schuhmann, W., Schwaninger, M., and Bilello, U., 1996. *Organisational transformation and learning – A cybernetic approach to management*. Chichester: Wiley.

Ferioli, F., Schoots, K., and van der Zwaan, B.C.C., 2009. Use and limitations of learning curves for energy technology policy: A component-learning hypothesis. *Energy Policy* 37(7): 2525–2535.

Ferioli, F., and van der Zwaan, B.C.C., 2009. Learning in times of change: A dynamic explanation for technological progress. *Environmental Science and Technology* 43(11): 4002–4008.

Goldemberg, J., 1996. The evolution of ethanol costs in Brazil. *Energy Policy* 24(12): 1127–1128.

Grübler, A., Nakicenovic, N., and Victor, D., 1999. Dynamics of energy technologies and global change. *Energy Policy* 27(5): 247–280.

IEA/OECD, 2000. *Experience curves for energy technology policy*. International Energy Agency/Organisation for Economic Co-operation and Development, Paris.

IEA/OECD, 2006. *Energy technology perspectives 2006—In support of the G8 plan of action, scenarios and strategies to 2050*. International Energy Agency/Organisation for Economic Co-operation and Development, Paris.

IEA/OECD, 2008. *Energy technology perspectives 2008—In support of the G8 plan of action, scenarios and strategies to 2050*. International Energy Agency/Organisation for Economic Co-operation and Development, Paris.

IEA/OECD, 2010. *Energy technology perspectives 2010—In support of the G8 plan of action, scenarios and strategies to 2050*. International Energy Agency/Organisation for Economic Co-operation and Development, Paris.

Junginger, M. (ed.), 2010. *Technological learning in the energy sector: Lessons for policy, industry and science*. Cheltenham, U.K.: Edward Elgar. (forthcoming)

Junginger, M., Faaij, A., and Turkenburg, W.C., 2005. Global experience curves for wind farms. *Energy Policy* 33(2): 133–150.

Kahouli-Brahmi, S., 2008. Technology learning in energy-environment-economy modeling: A survey. *Energy Policy* 36(1): 138–162.

Kim, D.H., 1993. The link between individual and organizational learning. *Sloan Management Review* 35(1): 37–50.

Kydes, A.S., 1999. Modeling technology learning in the national energy modeling system. In *Proceedings IEA workshop on experience curves for policy making – The case of energy technologies*, eds. C.O. Wene, A. Voss, and T. Fried. pp. 181–202, Forschungsbericht 67, Institut für Energiewirtschaft und Rationelle Energieanwendung, Universität Stuttgart, Stuttgart, Germany.

Luhmann, N., 2002. *Theories of distinction*. Stanford: Stanford University Press.

MacGregor, P.R., Maslack, C.E., and Stoll, H.G., 1991. *The market outlook for integrated gasification combined cycle technology*. General Electric Company.

Manne, A.S.,1994. Private communication.

Mattsson, N., and Wene, C.O., 1997. Assessing new energy technologies using an energy system model with endogenized experience curves. *International Journal of Energy Research* 21(4): 385–393.

Maycock, P.D. and Wakefield, G.F., 1975, *Business analysis of solar photovoltaic energy conversion*. Photovoltaic Specialists Conference, 11th, Scottsdale, Ariz., May 6–8, 1975, Conference Record. (A76-14727 04–44) New York, Institute of Electrical and Electronics Engineers, Inc., 1975, p. 252–255.

McDonald, A., and Schrattenholzer, L., 2001. Learning rates for energy technologies. *Energy Policy* 29(4): 255–261.

Messner, S., 1997. Endogenized technological learning in an energy systems model. *Journal of Evolutionary Economics* 7(3): 291–313.

Neij, L., 1999. Cost dynamics of wind power. *Energy* 24(5): 375–389.

Neij, L., Andersen, P.D., Durstewitz, M., Helby, P., Hoppe-Kilpper, M., and Morthorst, P.E., 2004. *Experience curves: A tool for energy policy assessment*. IMES/EESS Report No. 40, Department of Technology and Society, Lund University, Sweden, see www.iiiee.lu.se

Nemet, G.F., 2006. Beyond the learning curve: Factors influencing cost reductions in photovoltaics. *Energy Policy* 34(17): 3218–3232.

Sagar, A.D., and van der Zwaan, B.C.C., 2006. Technology innovation in the energy sector: R&D, deployment, and learning-by-doing. *Energy Policy* 34(17): 2601–2608.

Schoots, K., Ferioli, F., Kramer, G.J., and van der Zwaan, B.C.C., 2008. Learning curves for hydrogen production technology: An assessment of observed cost reductions. *International Journal of Hydrogen Energy* 33(11): 2630–2645.

Schoots, K., Kramer, G.J., and van der Zwaan, B.C.C., 2010. Technology learning for fuel cells: An assessment of past and potential cost reductions. *Energy Policy* 38(6): 2887–2897.

Seebregts, A.J., Kram, T., Schaeffer, G.J., Stoffer, A., Kypreos, S., Barreto, L., Messner, S., and Schrattenholzer, L. 1999. *Endogenous technological change in energy system models*. ECN-C--99-025, ECN, Petten, The Netherlands.

Stern, N., 2006. The economics of climate change. *The Stern Review*. Cambridge University Press: Cambridge

Tsuchiya, H.,1989. Photovoltaic cost based on the learning curve. *Proceedings of the international solar energy society clean & safe energy forever symposium*, Kobe City, Sep. 4–8, 1989, p. 402.

van der Zwaan, B.C.C., Gerlagh, R., Klaassen, G., and Schrattenholzer, L., 2002. Endogenous technological change in climate change modeling. *Energy Economics* 24(1): 1–19.

van der Zwaan, B.C.C., and Rabl, A., 2004. The learning potential of photovoltaics: Implications for energy policy. *Energy Policy* 32(13): 1545–1554.

Varela, F., 1979. *Principles of biological autonomy*. Elsevier-North Holland, New York.

Varela, F., 1984. Two principles for self-organization. In *Self-organization and management of social systems*. eds. H. Ulrich and J.B. Probst, p. 25–32. Berlin: Springer.

von Förster, H., 1980. *Observing systems*. The systems inquiry series. Intersystems: Seaside.

von Förster, H., 2003, *Understanding*. Berlin: Springer.

Weiss, M., Junginger, M., Patel, M.K., and Blok, K., 2010. A review of experience curve analyses for energy demand technologies. *Technological forecasting and social change* 77(3): 411–428.

Wene, C.O., 2007. Technology learning system as non-trivial machines. *Kybernetes* 36(3/4): 348–363.

Wene, C.O., 2008a. A cybernetic perspective on technology learning. In *Innovations for a low carbon economy: Economic, institutional and management approaches*, eds. T.J. Foxon, J. Köhler and C. Oughton. Cheltenham: Edward Elgar.

Wene, C.O., 2008b. Energy technology learning through deployment in competitive markets. *The Engineering Economist* 53(4): 340–364.

Wene, C.O., 2010. Adaptation in technology learning systems. *The proceedings of the 11th IEAA European conference*. 25–28 August, 2010, Vilnius, Lithuania. Available at http://www.iaee2010.org/?q=node/84

Williams, R.H., and Terzian, G., 1993. *A benefit/cost analysis of accelerated development of photovoltaic technology*. PU/CEES Report No. 281, Princeton University, N.J.