Energy Policy 37 (2009) 2525–2535

Contents lists available at ScienceDirect

Energy Policy

journal homepage: www.elsevier.com/locate/enpol

Use and limitations of learning curves for energy technology policy: A component-learning hypothesis

F. Ferioli [a,*], K. Schoots [a], B.C.C. van der Zwaan [a,b]

[a] *Energy Research Centre of The Netherlands (ECN), Policy Studies Department, Petten/Amsterdam, The Netherlands*
[b] *The Earth Institute, Lenfest Center for Sustainable Energy, Columbia University, New York, USA*

ARTICLE INFO

Article history:
Received 18 February 2008
Accepted 21 October 2008
Available online 5 April 2009

Keywords:
Energy technology
Learning by doing
Experience curve

ABSTRACT

In this paper, we investigate the use of learning curves for the description of observed cost reductions for a variety of energy technologies. Starting point of our analysis is the representation of energy processes and technologies as the sum of different components. While we recognize that in many cases "learning-by-doing" may improve the overall costs or efficiency of a technology, we argue that so far insufficient attention has been devoted to study the effects of single component improvements that together may explain an aggregated form of learning. Indeed, for an entire technology the phenomenon of learning-by-doing may well result from learning of one or a few individual components only. We 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) do not exhibit cost improvements might account for the apparent time dependence of learning rates reported in several studies (the learning rate might also change considerably over time depending on the data set considered, a crucial issue to be aware of when one uses the learning curve methodology). Such an explanation may have important consequences for the extent to which learning curves can be extrapolated into the future. This argumentation suggests that cost reductions may not continue indefinitely and that well-behaved learning curves do not necessarily exist for every product or technology. In addition, even for diffusing and maturing technologies that display clear learning effects, market and resource constraints can eventually significantly 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 technologies we assess the reliability of using learning curves at large to forecast energy technology cost reductions.

1. Introduction

Given the current reliance of our economy on fossil fuels, substantial efforts will be needed to decarbonize it on a global scale. Some peculiar aspects of the energy sector contribute to the complexity of the challenge. First, given the size of it, proven CO_2 abatement measures should be deployed on scales unseen for most other environmental problems. Second, persistent difficulties in internalizing the environmental costs, or externalities, of energy use heavily distort the market in favor of the incumbent fossil-based technologies. Therefore, market forces alone cannot be expected to deliver the required fundamental change. It is the purpose of carefully crafted public policy to optimize the transition to a sustainable energy system. In addition to the continued development of new energy technologies, the deployment of existing clean ones is essential to this transition. To assess as accurately as possible the economic implications of the necessary profound technological transformation, quantitative tools such as learning curves have been developed. Key to strategically planning the deployment and estimating the potential capacities of alternative energy technologies are attempts to forecast their future costs, and the learning curve methodology is one of the instruments available to achieve this task.

As new energy technologies translate into new commercial products, the focus of the industry-concerned shifts from R&D to deployment. At this stage significant cost reductions may be brought about by accumulating experience merely as a result of deployment activity. The lessons learned on the field usually yield a variety of process improvements, among which cost reductions. Learning curves have extensively been used to describe this phenomenon of "learning-by-doing" for the deployment of energy technologies (Wene, 2000; McDonald and Schrattenholzer, 2001).

* Corresponding author.
E-mail address: ferioli@ecn.nl (F. Ferioli).

0301-4215/$ - see front matter © 2008 Published by Elsevier Ltd.
doi:10.1016/j.enpol.2008.10.043

In contrast to direct cost-estimate techniques, learning curves have the potential to describe cost reductions (or more generally progress) for a product over a range spanning a volume growth of orders of magnitude. While learning curves are regularly used for strategic planning at the firm level (Dutton and Thomas, 1984), in the context of the design of energy policy they are often used for, e.g., estimating the future potential of emerging technologies (Neij, 1997; van der Zwaan and Rabl, 2004) and providing input for comprehensive energy system modeling. Even if learning curves have proved useful for a number of purposes, they need to be handled carefully in order to derive reliable and robust lessons for energy policy making. In this paper, examples from wind power, photovoltaic (PV) cells and hydrogen production are used to illustrate some important methodological issues and caveats related to the use of learning curves. In particular, a detailed discussion is provided for the way error margins ought to be accounted for when applying this technique.

In Section 2, we briefly introduce the formulation of learning curves. In Section 3 a model is presented that describes a product, process or technology as the sum of several components, each of which learns at a different pace. The model derived is used to discuss the reliability of learning curves for the long-term forecasting of energy technology costs. Section 4 proposes an analysis of the possibility to derive industrywide learning curves for energy technologies. In this section, the accuracy of learning curves for energy technologies is assessed with wind power as an example. Finally, we report our main conclusions in Section 5.

2. Method and framework

The concept of learning-by-doing expresses that accumulating the deployment or use of a technology increases the corresponding experience, which typically results in the optimization of the process involved. In particular, technology improvements are often economic in nature and thus result in cost reductions, so that changes in cost or price are usually used as a proxy for learning-by-doing. Already in the 1930s, it was observed that costs may decrease by approximately a fixed percentage with each doubling of cumulated production (Wright, 1936). This quantitative relation can be written as

$$C(x_t) = C(x_0)\left(\frac{x_t}{x_0}\right)^{-b}, \tag{1}$$

in which x_t is the cumulated production (or capacity), b a positive learning parameter, and $C(x_t)$ the cost (or, as it is used in many cases, the price) of a product, process or technology at x_t. The variables $C(x_0)$ and x_0 are, respectively, the cost and cumulated production at an arbitrary starting point.[1] Learning curves are derived by fitting Eq. (1) to cost and production data observed in the past. The starting point then ideally corresponds to the first unit of production. In practice, however, it often proves more appropriate to choose a later (but still early) stage of deployment for $t = 0$, and for the purpose of estimating future cost reductions on the basis of learning curves, it can be convenient to use the present cumulative production as starting point. The learning rate (LR) is defined as the relative cost reduction (in %) after each doubling of cumulative production, that is

$$LR = 1 - 2^{-b}. \tag{2}$$

[1] One can easily see that this starting point is arbitrary in principle by realizing that, given two levels of cumulative production x_0 and x_1, $C(x_i) = C(x_0)(x_t/x_0)^{-b} = C(x_0)(x_1 x_t/x_1 x_0)^{-b} = C(x_0)(x_1/x_0)^{-b}(x_t/x_1)^{-b} = C(x_1)(x_t/x_1)^{-b}$.

Learning curves have been developed for many products, processes and technologies in several industrial fields, which thus constitute empirical evidence for the phenomenon of learning-by-doing and the existence of LRs. Studies have been undertaken that propose a more theoretical clarification of learning-by-doing, some of which are more established and accepted than others (see, e.g., Arrow, 1962; Wene, 2007). These analyses, however, still remain far from a broadly agreed explanation of the apparently robust cost–production relation. This article attempts to contribute to opening the black box of learning curves.

In a comprehensive survey Dutton and Thomas (1984) analyzed the results of 108 studies that report LRs in 22 industrial sectors, among which are the electronics, machine tools, paper-making, steel and automotive industries. We have normalized the distribution of the LR values from their data set to obtain the relative probability of each LR and have fitted it with a normal distribution as shown in Fig. 1. The observed LRs are approximately normally distributed with a mean $\mu \approx 19\%$ and a standard deviation $\sigma \approx 8\%$. Our Gaussian fit describes the variance in the data with a coefficient of determination that is statistically significant ($R^2 = 0.76$), so that there is a 95% probability of finding values for LR between 3% and 34% (truncated to the closest integer). Note that the Dutton and Thomas data refer to learning at the firm level and hence do not include LR values that cover entire industrial sectors or fields at large. Data from other studies (most notably McDonald and Schrattenholzer, 2001) report a comparable central value and a similar distribution for the observed LRs (increasing the confidence in the statistical relevance of the data set).

On a double-logarithmic scale the exponential relation of Eq. (1) is represented by a straight line with slope $-b$. This is shown in Fig. 2, in which the normalized costs (C/C_0) are plotted as function of the normalized cumulated capacity (x/x_0) for the average LR of Fig. 1 and the corresponding 95% confidence level (CL) values (for brevity, here we use $C_0 = C(x_0)$). It can readily be observed that the cumulative production required to reach a given relative cost reduction depends strongly on the value of the LR. For any technology entering the market, we define the breakeven capacity, x_b, as the cumulated production or deployment necessary to reach a given cost target, C_b, e.g., to become competitive with an incumbent technology that delivers the same or similar service. Fig. 2 demonstrates that if $LR = 19\%$ and the technology under consideration needs to reach a cost one-tenth of the current

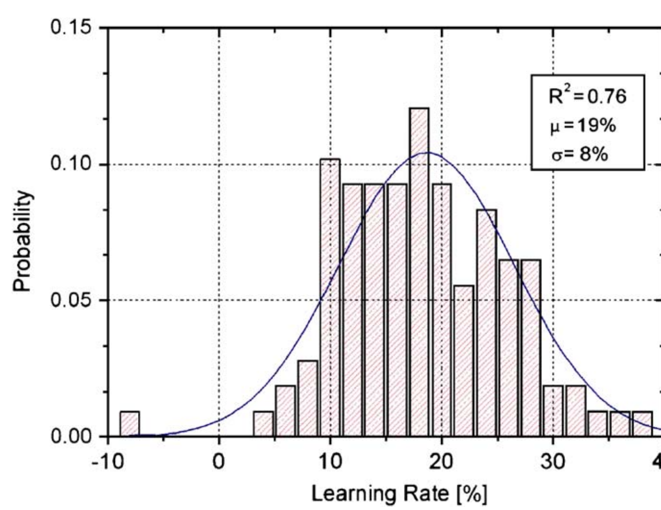

Fig. 1. Distribution of observed learning rates (bars) and fit with a normal distribution (solid curve) based on the mean (μ) and standard deviation (σ) of an observed set of learning rates. Data from Dutton and Thomas (1984).

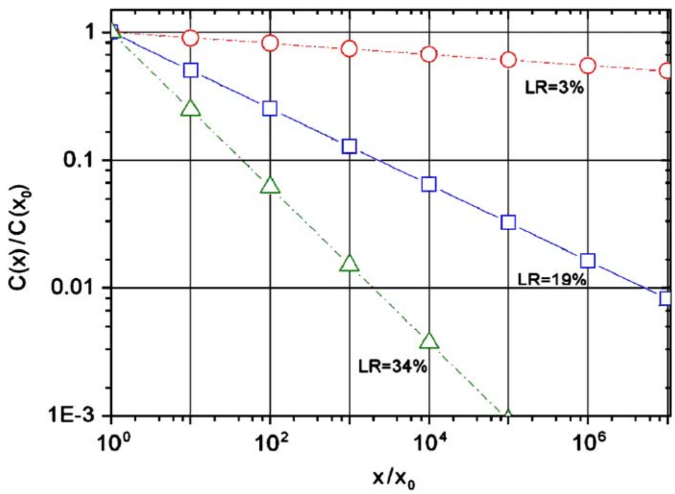

Fig. 2. Normalized costs (C/C_0) as function of the normalized cumulated capacity (x/x_0) for three different learning rates (19%, 3% and 34%) corresponding, respectively, to the average (□) and 95% CL values (○ and ▲).

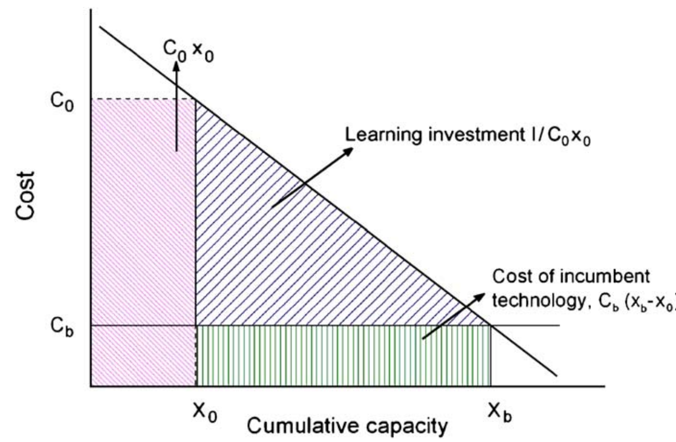

Fig. 3. Graphical illustration of the learning investment, I, required to reach breakeven cost level C_b.

level in order to become competitive, the cumulated capacity ought to be expanded by three orders of magnitude.

If the technology costs are expressed per unit of cumulative production, the total cost of deploying a given capacity can be found by integrating Eq. (1). The learning investment, I, is defined as the additional cost required for reaching the competitive breakeven capacity, that is, the total deployment costs minus the costs that the same capacity of conventional technology would have incurred. Written explicitly, the breakeven capacity and learning investment as function of the normalized breakeven cost, C_b/C_0, are, respectively

$$x_b = x_0\left(\frac{C_b}{C_0}\right)^{-1/b}, \tag{3}$$

and

$$I = C_0 x_0\left\{\frac{1}{1-b}\left[b\left(\frac{C_b}{C_0}\right)^{(b-1)/b} - 1\right] + \frac{C_b}{C_0}\right\}. \tag{4}$$

Fig. 3 illustrates the geometrical meaning of the quantities in Eqs. (3) and (4): I, x_0, C_0, x_b and C_b. In particular, the monetary value of the learning investment depends on the capacity already deployed (or, more precisely, on C_0x_0). In other words, translating the learning curve horizontally to the left or to the right will result in an investment that is, respectively, smaller or bigger.[2]

Table 1 shows the normalized breakeven capacity, x_b/x_0, and the normalized learning investment, I/C_0x_0, necessary to achieve a given relative cost reduction for three different values of the LR. This table confirms our observation from Fig. 2 that, with an average LR of 19%, it is necessary to deploy approximately three orders of magnitude times the current installed capacity ($x_b \approx 10^3 x_0$) in order to reduce the current cost by one order of magnitude ($C_b/C_0 = 0.1$). It also shows, for example, that to reach this dramatically reduced cost level one needs to invest the equivalent of approximately 100 times the current cumulative installed capacity at current costs ($I \approx 10^2\ C_0x_0$). In the case of energy technologies the unit costs can be considerably high compared to e.g., consumer electronics (certainly if entire power plants are considered, but even for modules of PVs). Hence, the required learning investment can become prohibitively large, especially if a significant capacity has already been deployed. Therefore, the current cumulative production or installed capacity proves to be one of the parameters that are fundamental for estimating the maturity of a technology or product. Note that Eq. (4) shows that the learning investment depends linearly on C_0x_0. A systematic error in the determination of the cumulative production, e.g., as a result of the omission of early production values for which data might be lacking, thus typically produces an error in the calculation of the learning investment that is limited in comparison to the one caused by an uncertainty in the value of the LR (given the power-law dependence on the latter).

Table 1
Normalized breakeven capacity (x_b/x_0) and normalized learning investment (I/C_0x_0) necessary to reach a given normalized breakeven cost (C_b/C_0) for different values of the learning rate.

Breakeven cost	Breakeven capacity			Learning investment		
	$LR = 3\%$	$LR = 19\%$	$LR = 34\%$	$LR = 3\%$	$LR = 19\%$	$LR = 34\%$
C_b/C_0	x_b/x_0	x_b/x_0	x_b/x_0	I/C_0x_0	I/C_0x_0	I/C_0x_0
0.5	>1E+6	10	3	1.6E+5	1	0.4
0.2	>1E+6	199	15	>1E+6	16	5
0.1	>1E+6	1943	47	>1E+6	84	9
0.05	>1E+6	1.9E+4	148	>1E+6	1362	16
0.02	>1E+6	3.8E+5	682	>1E+6	3374	18
0.01	>1E+6	>1E+6	2168	>1E+6	1.7E+4	30

Essentially, for all products or technologies a maximum production or capacity limit exists due to either market or resource constraints. If the market for a given product saturates, new capacity is only needed for the replacement of aged products, which significantly reduces the scope remaining for increasing the cumulated capacity and thus limits the opportunities for learning-by-doing. In particular energy technologies are in addition often bounded by constraints related to the availability of natural resources. In fact, for energy technologies resource constraints are usually more common than those related to the size of the market, given the large role energy plays in our world economy. When market constraints are reached, learning phenomena usually come to a halt, whereas when the constraints reached are related to the limited presence of some natural resource, costs of the technology often tend to rise.[3] Wind energy, for example, may be

[2] To realize this, just consider that for a given current cost C_0 it is cheaper to double the cumulative capacity from 1 to 2 MW than from 1 to 2 GW. On a log–log plot the geometrical area under the curves is the same in the two cases.

[3] Note that 'natural resource' can have a meaning as diverse as fossil fuels, heavy metals, wind, sun, or waste disposal options like the atmosphere or the geological underground.

limited by the availability of sufficient windy sites. If wind turbines are placed in a sub-optimal location, the cost of wind electricity consequently rises. If wind turbines are placed offshore because of a lack of space on land, electricity costs may increase due to an augmentation of installation and operation costs. Biomass is ultimately limited by the availability of land as well as by issues like competition with food crops. Also the use of fossil fuels clearly possesses a limiting capacity. During the last decades we have witnessed impressive improvements in exploration technology, e.g., through the exploitation of 3D-seismic detection techniques, and learning curves have been proposed for the costs associated with oil extraction and pipeline installation activities (see notably McDonald and Schrattenholzer, 2001). Nevertheless, corresponding reductions in oil prices have not been observed, or at least, if available, they have been balanced or shadowed by resource-related cost increases. It is an accepted notion that, as cheap oil reserves are being depleted, oil in new resources will be more expensive to extract, which is likely to offset the effects of learning-by-doing. For several alternative energy technologies it appears possible to estimate what in each respective case the limiting factor or capacity may be (see, e.g., IEA, 2006). While below we come back to the issue of resource and market constraints in the context of specific technologies, a detailed analysis of such limiting capacities for energy technologies at large is beyond the purpose of this paper.

3. From innovation to products

A way to describe the possible long-term slowing down of learning-by-doing (including effects of potential resource constraints) is to consider a product, process or technology as an aggregate of several components. Naturally, the cost 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

$$C(x_t) = \sum_{i=1}^{n} C_{0i}\left(\frac{x_{ti}}{x_{0i}}\right)^{-b_i} = C_{01}\left(\frac{x_{t1}}{x_{01}}\right)^{-b_1} + C_{02}\left(\frac{x_{t2}}{x_{02}}\right)^{-b_2} + \cdots + C_{0n}\left(\frac{x_{tn}}{x_{0n}}\right)^{-b_n} \quad (5)$$

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}. For whether aggregate learning can be broken down into component learning according to Eq. (5), the value of the cumulative production of each component is at least as important as the individual learning parameter. The reason is that between components x_{0i} may have widely diverging values, and along with b_i also 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 cost reduction for these construction materials (needed for notably components like the support mast and turbine housing) can be expected by the deployment of windmills. On the other hand, continued improvements can be expected for the fabrication of (light-weight) 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 Eq. (1) can be broken down into the component learning expression of Eq. (5). Vice-versa, one may question when an equation of the form of Eq. (5), if its validity can be demonstrated, can be approximated by the expression of Eq. (1), i.e. with only one term.

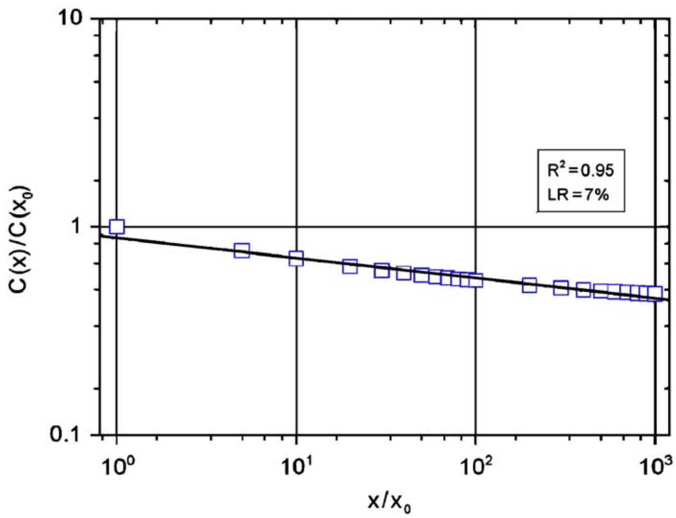

Fig. 4. Data points (□) calculated with Eq. (6) for parameter values $b = 0.3$ (i.e., $LR \approx 19\%$ for the learning component) and $\alpha = 0.6$, as well as a linear fit (—) through these points.

For ease of exposition we 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–α is in the beginning the cost share of the second component. The overall cost as a function of the cumulative production of the learning component can, in this simplified case, be expressed as

$$C(x_t) = \alpha C_0\left(\frac{x_t}{x_0}\right)^{-b} + (1-\alpha)C_0, \quad (6)$$

in which C_0 is again the total cost at production level x_0. Eq. (6) can be considered a special case of the more elaborate model presented by Eq. (5), useful to highlight some properties of these functional relations. Theoretical justification for this model is the observation (e.g., Schoots et al., 2008; van der Zwaan and Rabl, 2004) that some parts of a technology, such as raw materials and labor, may not experience fast cost reduction or even become more expensive in time.[4] In principle, the value of α can be accurately calculated for each technology, but this task is left to further work. Furthermore, we assume 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. We also suppose 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. Fig. 4 shows a set of data points calculated through Eq. (6) and plotted on a double-logarithmic scale, based on assumptions for $b = 0.3$ (that is, $LR = 19\%$ for the learning component) and $\alpha = 0.6$ (that is, 40% of the initial cost can be attributed to a component that does not involve any cost reductions).

Fig. 4 also depicts a linear fit through these data points. One can see that over three orders of magnitude, Eq. (6) can be accurately fitted with a straight line, that is, a learning curve of the form of Eq. (1). Indeed, the calculated regression accuracy ($R^2 = 0.95$) is comparable to that of the observed learning curves

[4] Similar arguments are to some extent addressed also in Carlson (1973).

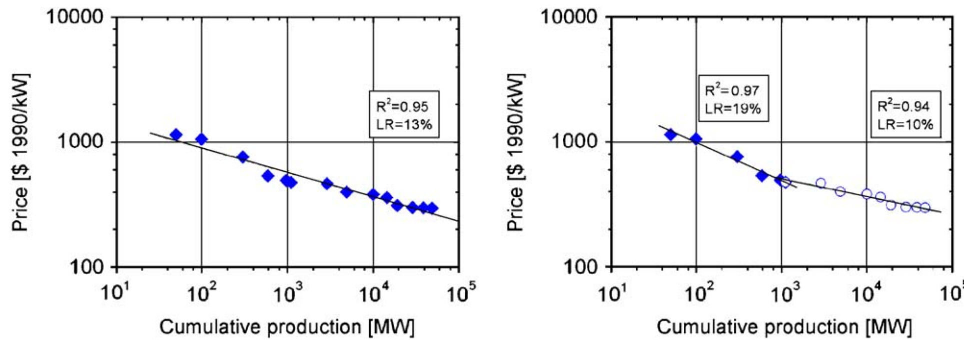

Fig. 5. As set of gas turbine prices fitted in two different ways, linearly and piecewise linearly. Data from MacGregor et al. (1991).

reported in the literature. The exponent of this linear fit, however, corresponds to an $LR = 7\%$, which is considerably smaller than the exponent used in Eq. (6) corresponding to an $LR = 19\%$ for the learning component only. In other words, the combination of a learning and non-learning component can be approximated by a single technology that also yields learning, but the corresponding *LR* is lower. Similarly, it is possible to make linear fits of data obtained through Eq. (6) on the basis of a wide range of values for parameters *b* and α. For α in between 0.1 and 1, and for many different *LR*s (for the learning component), it is possible to interpolate data generated through Eq. (6) with a straight line on a double-logarithmic scale with good accuracy. The R^2 for the fit becomes less than 0.9 only for low vales of α (typically <0.1), i.e. when the non-learning component dominates the total cost of the technology, and for high values of the *LR* of the learning component. One can turn this argument around, by stating that if a product exhibits a high *LR* it is generally more difficult to describe it as the sum of a learning and non-learning component, especially if there are empirical indications that the component(s) that learn to contribute modestly to the total technology costs. Indeed, one can only obtain a good linear fit of data along a line with curvature when this curvature is limited, that is, when the non-learning component contributes modestly to overall costs or when the *LR* of the learning component is low. Otherwise, the data along the curve (Eq. (6)) and the linear fit (Eq. (1)) will significantly diverge if further extrapolated.

Hence, overall learning can be seen as the result of learning of one (or several) component(s) while the other component(s) do(es) not learn in certain circumstances only. We can test these considerations by inspecting some of the learning curves published in the literature. Fig. 5 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 with an $LR = 13\%$ and $R^2 = 0.95$ (Fig. 5, left plot). It can be observed, however, that the data present an evident inflection point. Therefore, in the literature (see notably Seebregts et al., 1999), fitting to such data with a piecewise-linear learning curve is proposed (Fig. 5, right plot). In this case, one obtains $LR = 19\%$ and 10%, and $R^2 = 0.97$ and 0.94, for the two learning curve pieces, respectively. We find this solution rather unsatisfactory, however, since, given the empirical nature of learning curves, the fact that the *LR* changes over time leads to unavoidable methodological issues: a constant *LR* 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 Eq. (6) as shown in Fig. 6. For example, the system can be described as composed of a learning component (with a relatively high $LR = 24\%$) that makes up 80% of the total cost, and a non-learning part (hence with constant cost) that accounts for

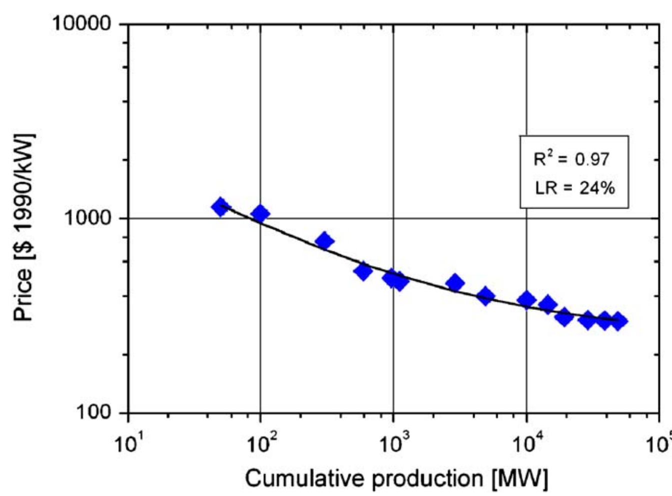

Fig. 6. The fit we propose of the gas turbine price data from Fig. 5, based on Eq. (6) with $LR = 24\%$ (for the learning component) and $\alpha = 0.8$.

the remaining 20% (one can think of the costs associated with the steel necessary to fabricate the mast and turbine). This fit is arguably better than both the ones shown in Fig. 5, since it represents a higher $R^2 = 0.97$, while applying to exactly the same data set. In any case, if for example one deems any fit on the basis of these available data with an accuracy $R^2 > 0.90$ acceptable,[5] then one cannot convincingly discard any of these potential regressions from the "phase space of possibilities", and certainly not the two-component fit proposed by us. Note that the value we chose for α derives merely from the fact that it represents the best fit of Eq. (6) to the data points and is not based on an analysis of the specific cost components of gas turbines.

While one may find all three fits to the available gas turbine price data of Figs. 5 and 6 acceptable, clear differences occur when one uses the corresponding different learning curves to the future. Extrapolating carelessly cost data over several orders of magnitude of cumulative production can lead to significant errors in both the breakeven capacity and the learning investment when one uses the wrong learning model. We point this out by Fig. 7, in which both the fit of the left plot of Fig. 5 and the one of Fig. 6 are depicted, and further extrapolated over three more orders of magnitude. Indeed, we see that the two lines diverge rapidly for

[5] The average R^2 for the learning curves reported by McDonald and Schrattenholzer (2001) for energy technologies is approximately 0.81. Therefore, a fit with $R^2 = 0.9$ is well-above average for a learning curve and can hardly be discarded on the basis of statistical considerations alone. Nonetheless, the choice of $R^2 = 0.9$ remains arbitrary and should not be interpreted as a threshold value to assess the validity of learning curves.

higher values of the cumulative production, with obvious repercussions in terms of such notions as the total capacity or learning investment needed to reach a given level of deployment in the future. As the cost of the innovative component is reduced, the non-learning component gains more in relative weight in terms of its contribution to the overall cost, and hence slows down the composite learning process.

For the capacity of energy technologies like wind turbines and power plants, due to the intrinsic scale of engineering involved, the first available cost data are normally minimally in units of MW (with PV modules being one of the very few exceptions). The fitting exercises shown in Figs. 4–6 suggest that in order to obtain a reliable interpolation one should be in a position to evaluate cost data over an extensive set typically spanning several orders of magnitude of cumulative production. The gas turbines example illustrates that data covering two orders of magnitude might not be enough to reveal important trends such as the possibility of having non-learning components. Extending the data set used to derive the learning curve presents several difficulties. Even assuming that a greater number of data points will allow to derive a reliable learning curve, it will likely be more difficult to compare costs collected over a longer interval of time (e.g., due to inflation or fluctuations of the exchange rate between currencies). Being aware of the complexity of the problem, we assume that it is possible to collect and compare cost-cumulative production data over a wider range of, for example, three orders of magnitude. Indeed, the three orders of magnitude depicted on the horizontal axes of Figs. 4–6 allow the determination of learning curves with an R^2 close to 1. Three orders of magnitude down the learning curve from the unit MW one arrives at capacities that more conveniently can be expressed in GW, which is the unit commonly used for current large-scale power plants. Today a couple of TW power generation capacity is installed worldwide. It is estimated that this value may expand several times, up to an order of magnitude, over the 21st century. Suppose one wants to use the learning curve methodology to estimate what the expected future cost reduction could be with respect to today, under certain values for the observed *LR*, for the main technologies that currently contribute to this global power production capacity. Then one readily concludes, on the basis of the limits of the total needs worldwide and given that the dominant power technologies are currently installed in terms of hundreds of GW (certainly for nuclear, hydro and fossil-based power generation, but today almost also for wind power), that for these technologies only a rather restricted subset of cells in Table 1 apply, with the corresponding implications for the expected cost reductions. Hence, apart from the limits that exist in terms of the economies-of-scale of single power plant due to upper boundaries associated with the relevant engineering problem, there also exists a clear limit to the totality of the sector or industry. This limit represents the boundaries of the learning system and should thus be explicitly accounted for when employing learning curves for the design of future energy technology policy.

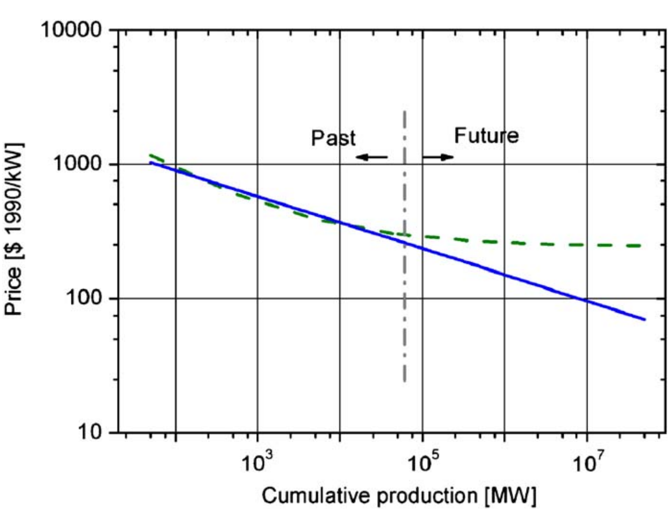

Fig. 7. Two fits for gas turbine prices extrapolated over three more orders of magnitude: a linear one based on Eq. (1) (—) and one based on Eq. (6) (- - -).

Cost data exhibiting a high *LR* cannot be fit accurately with an expression of the form of Eq. (6). If, as we claim, the learning curves reported in the literature can often be described as composed of multiple components, higher values for the *LR* should be relatively less abundant than lower ones. It appears that the distribution shown in Fig. 1 is slightly right-skewed (i.e. positively skewed), which indeed in principle means (barring statistical fluctuations) that more studies have determined relatively low *LR*s. An over-abundance of low *LR*s could point towards learning of composite systems in which parts of the overall technology learns with a higher *LR* than other systems (that could learn at zero rate). Like in the gas turbine example, a system with a non-learning component can often be described with a single learning curve with an overall *LR* that is lower than that of the individual learning component. The fact that in Fig. 1, we observe a higher frequency of *LR*s in the 10–19% range than in the 19–28% range suggest that our supposition may be correct. Given a lack in *LR* data, however, this conclusion so far cannot be rigorously demonstrated. A common measure of the asymmetry of a statistical distribution is the so-called skewness.[6] For a random normally distributed data set the skewness tends to zero as the number of its element tends to infinity. A finite set may show some level of skewness due to random statistical fluctuations. It proves that the data in Fig. 1 yield a positive skewness of about 0.23. We believe this value lies in the range of fluctuations of a random normally distributed set of comparable size. If in the future the set of *LR*s is expanded, however, we may be able to demonstrate that the skewness we then find is statistically meaningful, which would support our proposal that overall system learning can often be decomposed into multiple component learning.

High *LR*s have of course been observed for several products. Such elevated levels of learning may especially apply to radical innovations possessing new features that did not exist before. These characteristics make them much more valuable than, e.g., the raw materials they are made of. The first airplanes built constitute an appropriate example in case. Other more recent examples are several high-tech products that today have become so important in modern technology-based society, such as semiconductors. These are so much based on fundamentally new concepts, related to multiple activities from manufacturing to engineering, that opportunities for learning multiply, while the relative weight of non-learning components is minimal. For many of these radical innovations there is empirical evidence that cost reductions can be sustained for several decades and over a wide range of cumulative production spanning several orders of magnitude, needed to develop a reliable learning curve.

Fig. 8 shows the well-known cost curve for PV modules (from Harmon, 2000), which at their conception constituted a fundamentally new method to generate electricity from solar radiation. The learning curve shows multiple inflections, but cost reductions are sustained over as much as four orders of magnitude of cumulative production. Such inflections were

[6] For n data with values x, mean μ and standard deviation σ, the skewness is defined as $\sqrt{n}\sum_{i=1}^{n}(x_i-\mu)^3/\sigma^3$.

explained with the "shake-out" phenomenon (see, e.g., Wene, 2000; Schaeffer et al., 2004): since cost data are normally not available and price is used to construct learning curves, changes in the production costs-to-price ratio (due e.g., to the effect of competition in the market) might affect the observed *LR*. The corresponding overall *LR* is relatively high, given the innovative nature of PV technology when it was introduced for household purposes in the 1970s. One can explain the inflection points also by realizing that competition stimulates innovation, and that the materialization of possible improvements is naturally focused on the most costly components. Substituting or improving a critical expensive component, e.g., step by step through incremental innovations, allows at each modification a reinvigoration of the learning process and thus overall for continued learning. This description remains consistent with the "shake-out" phenomenon often used to describe the inflections in a curve such as that of Fig. 8. Indeed, the PV data can also be fit piecewise with expressions of the form of Eq. (1) or (6). Especially, at these stages of technology development interactions with continued R&D prove important, since research may provide the information needed on how to replace or improve a critical component. The high cost of crystalline silicon, for example, has stimulated research on amorphous and thin film modules. The effect of such R&D, and especially the contribution of incremental innovations, cannot usually be readily disentangled from other learning effects. Interestingly, it has been pointed out that economies-of-scale have so far probably been the greatest factor for cost reductions of PV modules (Nemet, 2006). Many analysts claim that such economies-of-scale should not be included in the phenomenon of learning-by-doing per se, while of course importantly contributing to the realizable cost reductions. It should be noted, however, that the difference between production cost and price (empirical data are normally available only for the latter) may complicate the analysis of "shake-out" phenomena.

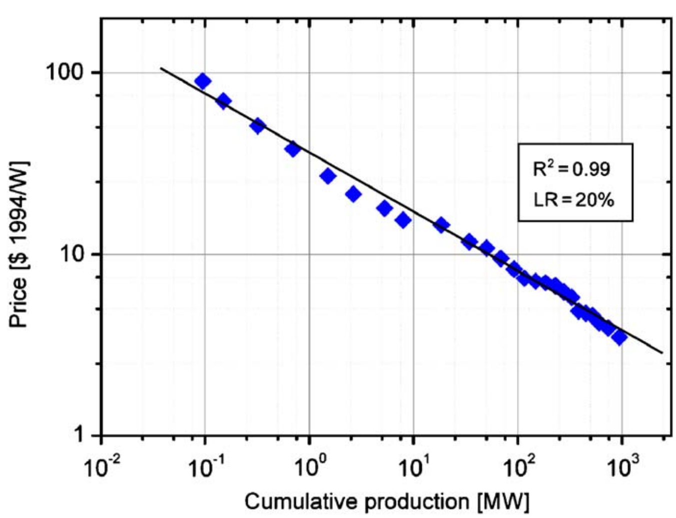

Fig. 8. Price data for PV modules after Harmon (2000).

Given the expected importance of economies-of-scale for energy technologies at large, the expectation about future demand and, ultimately, public opinion also play a key role in bringing about cost reductions. Dutton and Thomas (1984) do not distinguish between scale effects and learning-by-doing per se, so that it is likely that scale effects are included in at least some (and probably many) of the studies reported in Fig. 1. Of course, it is not granted that the costs for PV modules will continue to improve indefinitely with a 20% *LR* (see, for generic arguments in this context, Sagar and van der Zwaan, 2006). For PV there are clear signs that component learning is at work. If learning cost equations apply for single components, then necessarily an overall expression similar to Eq. (5) must be employed. Thus, over a wider scale the cost of PV modules can exhibit the behavior described by the dotted line in Fig. 7. This supports the notion that learning-by-doing may fade out as production or time proceeds, as suggested by Sagar and van der Zwaan (2006).

As also described in Grubler et al. (1999), the full learning cycle for a new product can now tentatively be described with the graph in Fig. 9. In the first stage, the innovative components, for which most opportunities for improvements exist, dominate the overall cost of the technology. The learning process develops fast and the cumulated capacity remains relatively small initially, which simplifies in principle the collection of reliable data. In this stage, it is easiest to develop a well-behaved learning curve. In the maturity stage the cost share of non-learning components, such as raw materials, becomes significant, so that the learning curve diverges from a straight line on a double-logarithmic plot, as was shown in the examples of Figs. 4 and 6. In some cases (but not necessarily in all), the non-learning components can be substituted or improved, which implies that, after a slowdown, the cost reduction process may continue. If, at this stage, the new product becomes competitive on the market, or establishes some sort of niche market, the incentives for continued innovation might be reduced and the learning curve bends towards the horizontal. When a product starts to reach the limiting capacity in terms of natural resource availability or market constraints, a similar effect may take place and the further deployment may be halted at current prices. At this stage limited cost reductions for single components may still persist, but do not necessarily translate into observable improvements for the final product. Fig. 9 shows qualitatively the overall learning process, divided into three main stages, both on a double-log scale (left) and a double-linear scale (right). The latter could also be interpreted as showing

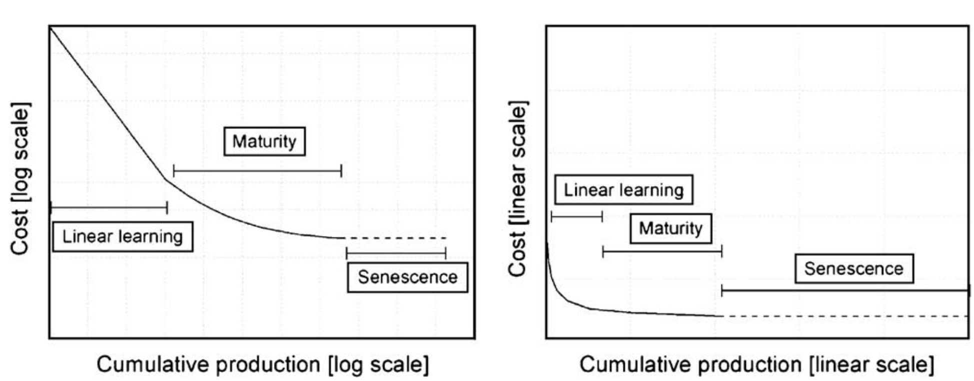

Fig. 9. Qualitative description of the learning cycle as a function of cumulated production (left, logarithmic scale; right, linear scale).

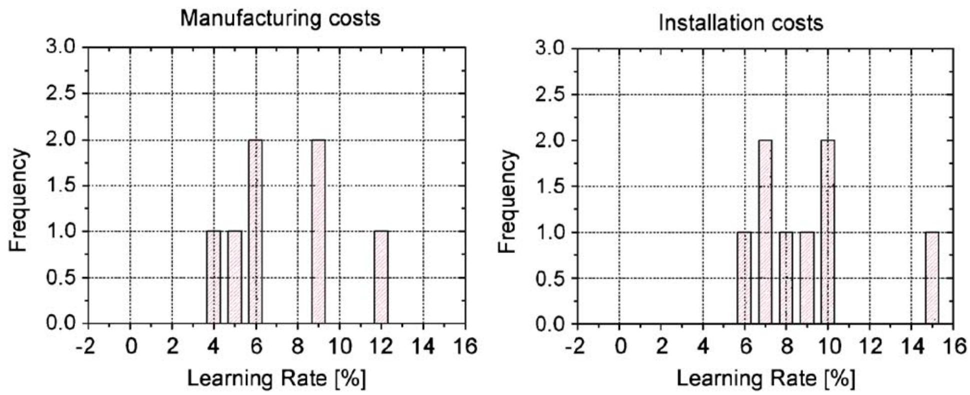

Fig. 10. Frequency of observed learning rates for wind turbines, respectively, for manufacturing costs (left, sample of 7) and turnkey costs (right, sample of 8) for different manufacturers, models and countries. *Source*: (Neij et al., 2003).

cost reductions as a function of time, rather than the cumulative production.

We observe that the later stages of product evolution, maturity and notably senescence, are normally not reported in the learning curve literature. For energy policy making, however, it is important to point out their existence, since they may have considerable implications for the design of specific policy instruments and their effects. We also note that bottom-up and top-down energy–economy–environment modelers may well want to consider accounting for these latter stages. In fact, van der Zwaan et al. (2002) in their DEMETER model do so, through their assumption that in the long run a minimum price for energy options exists below which learning-by-doing cannot fall. Another relevant issue in this field is that, due to market fluctuations and the long-time span often required to increase significantly the cumulated production, it might be difficult to observe a clear trend for the corresponding part of the learning curve, even while learning may be at work. Indeed, it proves that in the case of hydrogen production total cost data are too scattered and the available capacity scale too limited (even while the time frame inspected is large) to determine any meaningful *LR* (Schoots et al., 2008).

4. Estimating future costs

Many of the learning curves reported in the literature were developed for a specific firm or production plant only. For example, the distribution of observed *LR*s shown in Fig. 1 was compiled from studies that essentially all focused on a specific product within a given company (Dutton and Thomas, 1984). For energy policy purposes, however, we are often interested in evaluating the potential of a new technology for an entire country or region, or even the world at large. In view of our assessment whether learning curves can be used as a reliable tool for strategic planning in the energy sector, we therefore here discuss if the learning curves as observed for single firms can be aggregated into an industry-wide power-law for learning-by-doing that may be applicable on a national or international level. Wind energy technology offers an appropriate example in case.

Learning curves have been proposed for the costs of several activities related to wind energy, among which are notably the production costs of wind turbines from different manufacturers and the installation costs of distinct types of turbines in different countries. A recent study reports several such kinds of learning curves for wind power (Neij et al., 2003). Fig. 10 summarizes the values of a selection of *LR*s from this analysis, for, respectively, the production costs and total installation costs of wind turbines (Neij et al., 2003). The installation costs include the combined contributions from manufacturing, assembling and putting in place, together referred to as the turn-key costs. Both in the case of manufacturing and installation costs the R^2 associated with the learning curve fits (>0.9) is high compared to the typical value observed for learning curves (see e.g. McDonald and Schrattenholzer, 2001), which implies a good statistical correlation between costs and cumulated production. As we see in Fig. 10, the average *LR* is slightly higher for the total installation costs than for the manufacturing costs only: 9% against 7%.

This study also reports a set of learning curves developed for total wind electricity generation costs (including turn-key plus operation and maintenance costs) in two different countries, Denmark and Germany (Neij et al., 2003). Of all learning curves reported, the total power production costs exhibit the fastest progress through learning-by-doing, with an average *LR* of approximately 13%.[7] The observation that power production learns faster than wind turbine installation and that the latter learns (somewhat) faster than the wind turbine manufacturing suggests that the more activities one considers together, the more learning this combined activity reveals. Indeed, the combined processes of manufacturing, assembling, placing and using (operating and maintaining) wind turbines in this example prove to involve a steeper learning curve than when one omits one or more of these single activities or "components". For wind energy the cost components manufacturing, assembling, placement and use of wind turbines (the latter also referred to as variable costs) are each part of an overall learning process that together multiply the possibility for optimizing the final product, electricity in this case. Even when the overall design of a new wind turbine involves the introduction of a more costly component, the use of the total system may imply stronger learning than before this component was included, and learning for the total technology may be higher than for each of the individual constituents. For example, a more advanced blade can be more expensive to fabricate, but produce cheaper electricity because of improved aerodynamics. This is a striking observation, because in the previous section we argued that non-learning components in principle lead to the opposite result, that is, a lower *LR* for the compound system. Our conclusion is that which of these two mechanisms is at work depends very much on the technology or system under consideration. Depending on the way cost components interact, the overall *LR* can be higher than that of the components, as in the above example of wind power, or lower, as in the cases discussed

[7] Note that installation maintenance and electricity costs are normally measured in different units (e.g., €/MW and €/kWh). However, the learning curve in Eq. (1) does not depend on the units used.

in Section 3. Note that, more generally, our examples demonstrate that the overall *LR* cannot necessarily be derived from the *LR*s of the individual components, activities, economic sub-sectors or individual firms—in any case not immediately or in a straightforward manner.

Another appropriate example that at present receives great attention from scientists and policy-makers alike, relates to carbon dioxide capture and storage (CCS). Several cost components can be identified for CCS technology, among which notably separation of CO_2 from the exhaust stream, transport of CO_2, its injection in the geological reservoir and monitoring of the storage site. The relative shares of individual cost components depend on several parameters. For example, CO_2 separation from the outlet stream of large point sources usually constitutes the greatest share of the total costs of CCS implementation, if the transport costs can be held within a range of typically hundreds of kilometers. The overall *LR* for total CCS application will probably depend on the interaction between its individual constituents in a non-trivial way. As in the case of wind power, the learning-by-doing for CCS technology could be significantly higher than for its single components, and it may well be that the costs of some of them may need to increase in order to minimize the costs or maximize the efficiency of the whole system. For example, higher transport costs may be incurred in order to reach a secure storage site, that is, one with negligible risk for leakage of CO_2 from the underground repository, but farther from the emissions source than close-by less-safe options. This might nevertheless induce an overall cost decrease, if hereby the monitoring costs that guarantee safe storage are minimized. On the other hand, as discussed in Section 3, it may also be well possible that a non-improving component reduces the cost reduction possibilities for the overall CCS process. For instance, the capture part of CCS equipment may so closely resemble similar separation technologies such as flue gas desulphurization (FGD) that little improvement can be expected to take place for CCS deployment, since FGD has already reached a high level of maturity. The large total cumulative capacities installed for related technologies like FGD probably need to be taken into account when estimating the learning potential for a technique like CO_2 capture. At this stage of development of CCS, it might not be possible to estimate the overall *LR* for this technology, given that today many of these sorts of effects are likely to interfere.

The results for PV (Fig. 8) and wind power (Fig. 10) proffer evidence for the notion that at least some products are characterized by industry-wide learning. Hence, for such technologies it should be possible to develop industrial learning curves, i.e. linear power-law fits of industry-wide costs plotted against national, regional or global cumulated production. Several studies have attempted to estimate the reliability of learning curves for forecasting possible future cost reductions of a range of different technologies. Often these analyses focus on specific products or factories only. Alchian (1963), for example, derives learning curves for the amount of labor required for the production of several types of airframes (i.e. essentially airplanes without their turbines) at different facilities. He evaluates learning curves by using fairly consistently available data from the first production phase of each model investigated. This start-off period lasts for about one year, during which typically a few hundred airframes are produced of each model. The *LR*s derived on the basis of this first phase are used to forecast future labor requirements for each of the different airframes beyond that phase. Thus, the learning curves are extrapolated and the results compared with the actual labor (cost) data recorded for the second production phase, during which typically thousands of airframes are produced. The total size of the data set available differs for each model, but, on average, the learning curves can be drawn over 2.7 doublings of cumulative production during that second period.[8] Alchian then points out that extrapolating learning curves involves considerable uncertainty: the estimated amount of labor required is either significantly in excess or falls short of the true level of labor employed, and the average absolute value of the error is as high as approximately 22% (with variations depending on the airframe model under consideration). Of course, Alchian's analysis covers learning curves derived for specific products (airframes) in particular manufacturing facilities (of the respective airplane construction companies). His findings, however, bear great relevance for attempts to estimate possible uncertainties in future expected cost reductions on the basis of industry-wide learning curves, notably for new energy technologies candidate to succeed incumbent products.

The *LR*s for wind power shown in Fig. 10 are derived from fairly recent data reported in Neij et al. (2003), but an earlier reference, with the same first author, provides detailed cost data for wind power in the year 1995 (Neij, 1997). This prior data set proves more suitable to test the reliability of extrapolating learning curves. The cost of wind power in 1995 was estimated at 0.066 \$(1995)/kWh, which is possibly the central value of a range of cost figures (although such a range is not reported in Neij, 1997). The total installed capacity in that year was reported to be 5 GW worldwide, and the installation costs amounted to 1333 \$(1995)/kW (Neij, 1997). This same reference derives *LR*s as observed in 1995 for wind power installation and electricity costs of 4% and 9%, respectively. These *LR*s are significantly smaller than those reported by the same author in the more recent paper 9% and 13%, respectively (Neij et al., 2003). The European Wind Energy Association estimates that in 2004, the globally installed capacity of wind turbines had grown to some 47 GW (Morthorst, 2004), that is, 3.2 doublings of the cumulated installed capacity with respect to the 1995 level. By 2004 the installation costs were in the range of 900–1100 €(2004)/kW depending, among other factors, on the country where the turbines were deployed (Morthorst, 2004). The cost of wind electricity in medium wind areas in the same year were estimated in the 0.05–0.06 €(2004)/kWh range. We now use the 1995 *LR*s to extrapolate the costs observed in that year, and correct for inflation and currency exchange rates. This implies a forecast for 2004 of 1160 €(2004)/kW for installation costs and 0.049 €(2004)/kWh for electricity costs. Table 2 summarizes these figures.

As we see from the numbers listed in Table 2, for both the installation and electricity costs of wind energy, the estimates for 2004 based on the cost data known and *LR*s determined in 1995 lie outside the ranges that actually materialized in 2004. Recently the installation costs of wind turbines were observed to progress with a *LR* of 9% (Neij et al., 2003). If this *LR* had been used instead of the 4% as calculated in 1995, then the cost level of 1160 €(2004)/kW would not have been reached at cumulative installed capacity of 47 GW but of about 14 GW instead. Also, based on an installation cost target of 1160 €(2004)/kW, the learning investment required to reach this level as forecasted in 1995 with a *LR* of 4% was about 3 billion €(2004). But given the observed *LR* of 9% the actual learning investment has only been about 600 million €(2004). In particular, the industry-wide learning curves derived for wind energy in these two references appear to be at least as reliable as the firm-level learning curves investigated by Alchian. Yet this example demonstrates that one has to be wary of uncertainties. Overall, with the data available in 1995, the docking of wind energy to the market could reliably be estimated over approximately 2–3 doublings of the cumulative installed capacity.

[8] For the model with the largest available data set the learning curve can be extrapolated over 4.2 doublings of cumulated production.

Table 2
Comparison between learning curve estimates and actual installation and electricity costs for wind power (1995–2004).

	1995	2004	Learning curve estimate
Cumulative capacity	5 GW	47 GW	–
Installation costs	1333 $(1995)/kW	900–1100 €(2004)/kW	1160 €(2004)/kW
Electricity cost	0.066 $(1995)/kWh	0.05–0.06 €(2004)/kWh	0.049 €(2004)/kWh

Interestingly, in this case the estimate for the investment costs of wind power proved to be higher than the actual costs in 2004, while the predicted cost of wind electricity was lower than the actually measured cost range for that same year (or in any case hovering around the lower side of that cost range). A comparison between the forecasted and actual data suggests that in this case the overall induced policy error is dominated as much by difficulties encountered in converting cost data and estimating cost ranges (depending here on, e.g., wind availability) as by the uncertainties with which the *LR* is known. The differences between learning curve estimates and the actual data would change significantly if, for example, a different base year was chosen to compare cost quotes, or if a different value for the average wind speed was assumed. In other words, over the cost and time range we considered here, the errors resulting from observed uncertainties in values for the *LR* are comparable to (and may even be shadowed by) those from annual variations in, for instance, the exchange rate between the US dollar and the Euro.

This example for wind energy shows that extrapolating cost reductions over long-time frames or capacity expansions, while providing valuable insight, requires caution. For PV technology these issues become even more apparent, since an increase of several orders of magnitude (rather than several doublings) in installed capacity is needed in order for this electricity option to reach a competitive cost target. Over the necessary expansion range associated with PV, uncertainties in estimates of the *LR* carry a larger weight than in the case with wind power. Even if cost data can be reliably fitted with a straight line on a double-logarithmic scale, this does not necessarily mean that, given the statistical errors involved, the slope of the fit is constant for every subset of the data. If for the case of PV depicted in Fig. 8 the cost data for only the first two orders of magnitude of installed cumulated capacity are used, the estimated *LR* is 22.5% rather than 20.2%. Assuming a cost target of 0.05 €(2004)/kWh, the breakeven capacity can be calculated to amount to 90 GW for $LR = 22.5\%$ and 190 GW for $LR = 20.2\%$, corresponding to a learning investment of 42 billion € and 71 billion €, respectively (for comparable calculations of the learning investment for PV required to reach competitive breakeven, see notably Nemet, 2006 and van der Zwaan and Rabl, 2004). In this case for PV, due to the large extrapolation involved, it is evidently more important to have a precise estimate for the value of the *LR*, and uncertainties have a larger impact on cost forecasts and thus energy policy making, than in the case of wind power.

5. Conclusions

In this paper, we have reviewed some of the possible caveats of the use of learning curves for energy policy purposes. Learning curves may provide insight in future cost trends for energy technologies and are, once the limitations of the methodology are taken into account, an attractive tool for both scientific analysts and public policy-makers (see also Neij, 2008, for a discussion of the limitations of learning curves for energy policy making). Learning curves provide a phenomenological description of the relation between past costs and cumulated production, and thus allow for the estimation of future cost reductions by simple extrapolation. Likewise, they can be employed to calculate the investments needed to bring a technology down to a competitive level, which may be welcome or even necessary to diffuse it on a large scale in the market. Many examples of learning curves have been reported in the literature. To good approximation, it is found that the observed *LR*s are normally distributed. We confirm that, based on a large set of investigated energy technologies, on average the *LR* amounts to approximately 19%, with 3% and 34% as lower and upper levels spanning a 95% confidence interval. It should be pointed out, however, that the sample of technologies and products for which learning curves are available might be biased (as discussed for example in Sagar and van der Zwaan, 2006) and more studies should be included before rigorous conclusions can be drawn about the distribution of *LR*s.

One of our important observations is that in order to derive reliable learning curves, one ought to apply the analysis to an extensive set of cost-cumulative production data (e.g., greater then two orders of magnitude of cumulative output). For example, the learning curve for PV cells depicted in Fig. 8 could involve a rather different *LR* if data had been fitted over two orders of magnitude of cumulated capacity only (i.e. a subset of the total data set shown). In other words, determining a learning curve on the basis of two orders of magnitude of cumulative capacity data, or less, implies a sizeable uncertainty in the corresponding *LR*, with a comparable error in the energy policy gauged on this *LR*.

Our primary finding is that, even when the learning curve is evaluated over a wide range (i.e., three orders of magnitude of cumulated production) quite different fits of the same set of data are imaginable and at least equally justifiable. We point out that products can often be described as the sum of a learning component and one for which no cost reductions occur. This claim can be based on the inspection of the technology under consideration and an appreciation of its specific features, including for instance the resources and materials needed to build or operate it. In this article we proffer a second argument supporting this thesis, which we demonstrate through a study of publicly available cost data for gas turbines. This example shows that representing a product as the sum of two components, each learning at a different rate (one of which, for gas turbines, with zero value), yields a better fit than considering the technology as one indivisible entity. The learning of individual components separately may also explain why learning curves are often observed to bend towards the horizontal i.e. slow down when a technology matures.

This result yields important implications for the evaluation of the prospects of energy technologies on the basis of learning curves. Indeed, our component learning hypothesis produces cost estimates significantly different from those with the standard learning-by-doing framework, especially when the learning curve is extrapolated far into the future. Evaluating the relative weight of different cost components might ameliorate the estimation of a technology's cost reduction potential. For example, if non-learning components such as the required fossil fuel feedstock constitute a great share of the overall cost, then the prospect for learning is probably limited. The possibility of skewness observed for the probability density function of *LR*s—although statistically not (yet) significant—may support our component learning proposal.

For an analysis of the future costs and potential of a new energy technology, an understanding of the scope for future deployment is important and should thus be carefully evaluated. We demonstrate that for quite some energy technologies, like state-of-the-art wind turbines, the potential for growth is typically limited to only a couple of doublings with respect to the currently installed capacity. This is due to the fact that their cumulative deployment is already approaching the 100 GW level and the overall electricity industry has a finite size. This means that the potential capacity increase is limited to few cells of Table 1, that is, only under a high *LR* can one still expect fairly sizeable cost reductions but never by an order of magnitude. For other energy technologies similar qualitative arguments regarding their growth potential can be made, which implies that it might more generally not be possible to observe cost reductions once they approach a technology-specific upper bound for deployment. Of course, for the learning curve methodology not only the value of the *LR* is primordial but also the cumulated capacity already deployed. The present cumulated capacity can be considered a proxy for the maturity of a technology. If a large cumulative production has already been realized, the investments required to achieve further substantial cost reductions can become prohibitive. Furthermore, the time required to increase significantly the cumulated production may become too long to observe cost reductions distinguishable above all sorts of 'background' effects like market fluctuations of resource inputs.

Acknowledgements

This research was funded by the Netherlands Organization for Scientific Research (NWO) under the ACTS Sustainable Hydrogen program (No. 053.61.304). The authors would particularly like to acknowledge Gert Jan Kramer for his valuable suggestions on the analysis underlying this paper. They are also grateful to many colleagues of the ACTS program and at ECN for their useful comments during the presentation of the findings reported in this article. The authors are responsible for all remaining errors.

References

Alchian, A., 1963. Reliability of progress curves in airframe production. Econometrica 31 (4), 679–693.
Arrow, K.J., 1962. The economics implications of learning by doing. Rev. Econ. Stud. 29 (3), 155–173.
Carlson, J.G., 1973. Cubic learning curves: precision tool for labor estimating. Manuf. Eng. Manage. 71 (5), 22–25.
Dutton, J.M., Thomas, A., 1984. Treating progress function as a managerial opportunity. Acad. Manag. Rev. 9, 235–247.
Grubler, A., Nakicenovic, N., Victor, D., 1999. Dynamics of energy technologies and global change. Energy Policy 27, 247–280.
Harmon, C., 2000. Experience curves of photovoltaic technology, interim report, IIASA, Laxenburg, Austria.
International Energy Agency (IEA), 2006. Renewable Energy, RD&D Priorities. International Energy Agency (IEA), Paris, France.
MacGregor, P.R., et al., 1991. The market outlook for integrated gasification combined cycle technology. General Electric Company.
McDonald, A., Schrattenholzer, L., 2001. Learning rates for energy technologies. Energy Policy 29, 255–261.
Morthorst, P.E., 2004. Wind energy the facts. in: Cost&Prices, vol. 2. The European Wind Energy Association, 〈http://www.ewea.org/index.php?id=91〉 (last accessed 3.10.2007).
Neij, L., 1997. Use of experience curves to analyze the prospects for diffusion and adoption of renewable energy technology. Energy Policy 23 (13), 1099–1107.
Neij, L., 2008. Cost development of future technologies for power generation—a study based on experience curves and complementary bottom-up assessments. Energy Policy 36 (6), 2200–2211.
Neij, L., et al., 2003. Experience Curves: a tool for energy policy assessment, extool project final report, environmental and energy systems studies, Lund, Sweden.
Nemet, G.F., 2006. Beyond the learning curve: factors influencing cost reductions in photovoltaics. Energy Policy 34 (17), 3218–3232.
Sagar, A.D., van der Zwaan, B.C.C., 2006. Technology innovation in the energy sector: R&D, deployment, and learning-by-doing. Energy Policy 34, 2601–2608.
Schaeffer, G.J., et al., 2004. Learning from the Sun, photex project final report, Energy research Centre of The Netherlands (ECN), Petten, The Netherlands.
Schoots, K., Ferioli, F., Kramer, G.J., van der Zwaan, B.C.C., 2008. Learning curves for hydrogen production technology: an assessment of observed cost reduction. Int. J. Hydrogen Energy 33 (11), 2630–2645.
Seebregts, A.J., et al., 1999. Endogenous Technological Change in Energy System Models, Energy research Centre of the Netherlands, (ECN), Petten, The Netherlands.
Wene, C.O., 2000. Experience Curves for Energy Technology Policy. International Energy Agency (IEA), Paris, France.
Wene, C.O., 2007. Technology learning system as non-trivial machines. Kybernetes 36 (3/4), 348–363.
Wright, T.P., 1936. Factors affecting the cost of airplanes. J. Aeronaut. Sci. 3 (2), 122–128.
van der Zwaan, B.C.C., Rabl, A., 2004. The learning potential of photovoltaics: implications for energy policy. Energy Policy 32, 1545–1554.
van der Zwaan, B.C.C., Gerlagh, R., Klaassen, G., Schrattenholzer, L., 2002. Endogenous technological change in climate change modelling. Energy Econ. 24 (1).