

Learning Curves and Technology Assessment

Alan McDonald
International Atomic Energy Agency, Vienna, Austria

Leo Schrattenholzer
International Institute for Applied Systems Analysis, Laxenburg, Austria

RR-03-002
March 2003

Reprinted from *International Journal of Technology Management*, **23**(7/8):
718–745 (2002).

Research Reports, which record research conducted at IIASA, are independently reviewed before publication. Views or opinions expressed herein do not necessarily represent those of the Institute, its National Member Organizations, or other organizations supporting the work.

Reprinted with permission from *International Journal of Management Technology*, **23**(7/8):718–745 (2002).

Copyright © 2002 Inderscience Enterprises Ltd.

The *International Journal of Management Technology* is published in both print and electronic format. For more information visit their web site www.inderscience.com

All rights reserved. No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopy, recording, or any information storage or retrieval system, without permission in writing from the copyright holder.

Learning curves and technology assessment

Alan McDonald and Leo Schrattenholzer

Environmentally Compatible Energy Strategies (ECS) Project,
International Institute for Applied Systems Analysis (IIASA),
Schlossplatz 1, A-2361 Laxenburg, Austria
Fax: +43-2236-807-488 E-mail: leo@iiasa.ac.at

Abstract: This paper uses the formal concept of learning curves to analyse regular behaviour of performance improvements in various energy technologies. The concept allows the estimation of a single indicator of technological progress, the learning rate, which expresses the constant percentage improvement (usually in terms of cost reductions) in a technology for each doubling of the technology's cumulative installed capacity. We present 42 energy-related learning rates, either calculated directly from available data or assembled from the literature. We elaborate briefly on eight of these to illustrate issues addressed by technology assessments to convert these raw historical learning rates into prospective learning rate distributions for use in long-term energy models. The paper includes a sensitivity analysis of policy-relevant variables with respect to learning rates, a discussion of possible extensions and limitations of the approach and an outlook on future work in the field.

Keywords: Technological learning; technology assessment; energy; energy modelling.

Reference to this paper should be made as follows: McDonald, A. and Schrattenholzer, L. (2002) 'Learning curves and technology assessment', *Int. J. Technology Management*, Vol. 23, Nos. 7/8, pp.718-745.

Biographical notes: Alan McDonald is the Nuclear Information Officer in the Planning and Economic Studies Section of the International Atomic Energy Agency. From 1997 to 2000 he was the Program Officer of IIASA's Environmentally Compatible Energy Strategies Project. From 1982 to 1997 he was Executive Directive of the US Committee for IIASA at the American Academy of Arts and Sciences. In 1974 he received his undergraduate degree in General Engineering, and a Master of Science degree in Aeronautical and Astronautical Engineering from Stanford University. In 1979 he received his Master of Public Policy degree from Harvard University's John F. Kennedy School of Government. He has worked on refuelling system design and reactor safety policy for General Electric's Fast Breeder Reactor Department and on energy facility siting procedures for the California Energy Resources Conservation and Development Commission.

Leo Schrattenholzer holds degrees in Mathematics (MSc) and Energy Economics (PhD) from the Technical University of Vienna, Austria. He has been affiliated with IIASA since 1973, currently as Project Leader of the Environmentally Compatible Energy Strategies (ECS) Project. He has worked as a consultant to the World Bank and UNDP as well as to governmental institutions on national strategies to reduce greenhouse gas emissions. Since 1981, he has served as one of two co-directors of an international network of energy analysts (the International Energy Workshop). He has contributed to the IPCC's Second Assessment Report as a lead author. The focus of his present work in the ECS project is in the field of energy technology assessment,

including the analysis of the role of research and development in enhancing technological progress.

1 Introduction: energy technology assessment

The 1998 IIASA-WEC study *Global Energy Perspectives* [1] concluded that technological change will be critical for future energy systems. The Council on Foreign Relations' Study Group on Global Warming Technology Policy for the USA states as its starting premise that "cutting emissions will require massive technological change toward clean, carbon-free fuels". *Factor Four*, von Weizsäcker *et al.*'s [2] new report to the Club of Rome describes itself as "herald[ing] nothing less than a new direction for technological progress". The 2000 *IPCC Special Report on Emissions Scenarios* (SRES) [3] concluded that technological change is effectively the most important driving force of the energy system, more important than demographic change and economic development.

The emphasis given to technology in these and other studies reflects the increasing research on the dynamics of technological change and on incorporating the results in long-term large-scale energy models. The ultimate objective behind using such models is to assess alternative policies for efficiently steering technological change in desirable directions.

One concept we consider particularly useful for incorporating technological change in energy models is that of learning curves. A learning curve describes technological progress (measured generally in terms of decreasing costs for a specific technology) as a function of accumulating experience with that technology. This is a significant departure from most earlier energy models (both top-down and bottom-up), in which the standard input assumptions include prescribed cost reductions for various technologies as functions of time. Early work on incorporating learning curves in energy models was carried out particularly by IIASA [4], at Chalmers University [5] and in the US Department of Energy [6] and technological learning curves have since been included in several E3 (energy-economy-environmental) models [7]. As a result, such models are capable of identifying longer-term optimal strategies that would not be optimal if viewed myopically, considering relative technology costs only at a given point in time. The longer-term view, taking learning into account, can thus uncover possibilities for profitably investing in technologies that are expensive today, but will become significantly cheaper as experience with their production and use accumulates.

Ultimately to draw useful policy conclusions, we need reasonably reliable estimates of the learning curves that these new expanded models should use in extrapolating into the future, plus model features that can deal with remaining uncertainties. Learning curves, after all, hardly represent a physical law. Rather they describe a persistent empirical phenomenon with still significant uncertainties surrounding both the estimation of specific learning rates and their extrapolation in long-term energy scenarios. Plausibility must therefore take the place of predictability and the quality of the fit with which a learning curve is estimated for a given set of data is the principal measure of the reliability of the tool. Both because the quality of fit is sometimes low and because of the uncertainty inherent in projecting empirical patterns 100 years into the future (the planning horizon of many long-term energy models), energy models that can take as

inputs distributions of learning rates rather than single values have an advantage. Against this advantage one must weigh the numerical complexity of model calculations and the associated computing requirements, as well as the difficulty of generally comprehending and fully assimilating voluminous complex model outputs and their subtle multiple sensitivities. But even if we opt for simpler, more comprehensible models, the variability of observed learning rates remains important. As shown in the next section, the sensitivity of policy-related variables to such variability is highly non-linear.

Section 2 of this paper introduces the concept of learning curves in a formal mathematical way and presents a first step in sensitivity analysis. Section 3 presents a total of 42 learning rates for a variety of energy technologies ranging from primary energy supply and conversion to end-use and energy services. We have estimated 26 of these (or re-estimated to assure that the results are consistent) directly from data sets we have assembled. The remaining 16 are estimates reported in the literature, for which we do not have the original data. We elaborate briefly on eight of these to illustrate issues addressed by technology assessments to convert these raw historical learning rates into prospective learning rate distributions for use in long-term energy models. Section 4 provides a methodological discussion of the learning curve concept. Alternative causes of technological learning, some of them closely related to experience accumulation, are also discussed in Section 4. Finally Section 5 provides a summary and offers an outlook on further work.

2 The concept of technological learning

2.1 Definitions

The concept of technological learning was first introduced over 60 years ago [8]. As used in this paper, it assumes that a technology's performance improves as experience with the technology accumulates. The concept can be used with a variety of different indicators of technological performance and experience, but we will start with specific capital costs as the performance indicator and total cumulative installed capacity as the experience indicator. In this case, technological learning is defined by the following power function.

$$\text{Cost} = A * \text{Ccap}^b \quad (2.1)$$

where:

- Cost ... Specific capital costs (e.g., \$/kW),
- A ... Specific capital costs at a total (initial) cumulative capacity of 1,
- Ccap ... Total cumulative installed capacity (e.g, gigawatts) and
- b ... Learning elasticity (a constant).

From this definition it follows that a doubling of total cumulative capacity reduces specific costs by a factor of 2^b . In the usual case where b is negative, 2^b (labelled the *progress ratio*, pr) is between zero and one. The complement of the progress ratio ($1-pr$) is called the *learning rate* (lr) [9]. A learning elasticity of -0.32 , for example, yields a progress rate of 0.80 and a learning rate of 20% . This means that the specific capital cost of newly installed capacity decreases by 20% for each doubling of total installed capacity. On a double-logarithmic scale, the decrease in costs appears as a straight line.

Most of the results and discussion in this paper use specific capital costs (or prices) and cumulative capacity as performance and experience indicators respectively. However, we include among the technologies singled out for additional discussion in Section 3.3, examples in which the performance indicator is efficiency or the cost of production (i.e., cost *per unit output*). In the case where performance is measured in terms of production costs (wind power), we use total cumulative *production* as the most appropriate matching experience indicator.

Note that time does not figure explicitly in the concept of learning curves. Nonetheless, empirical learning curves are usually derived from time series of costs and capacities and thus carry invisible time labels with potentially interesting information. These become important if we want to address the possibility of knowledge depreciation (see Section 4.1), i.e., the possibility that experience acquired a long time ago is not as valuable as experience acquired yesterday. Moreover, in the hypothetical case of exact exponential growth, the logarithmic transformations of the cost functions over time are linear, a property that is useful in cases where information on total installed capacity is difficult to obtain or unreliable. If, in such a case, exponential capacity growth can be plausibly assumed, the validity of the 'learning model' can be judged by checking whether the logarithm of cost is a linear function of time. However, to calculate the learning *rate* for that learning curve we would have to know the rate of exponential capacity growth.

As in any one-formula model, the application of equation 2.1 to the data sets in Table 1 is a substantial simplification. But as was the case with other simple tools for statistical measurement (such as a data set's mean and standard deviation), if the usefulness of calculated learning rates becomes widely accepted, the importance of the concept's 'truth value' becomes secondary. Nonetheless, in this paper we discuss the shortcomings of equation 2.1 to help identify the practical limits of its validity. Exploring such shortcomings leads to suggested modifications (see Section 4) and to insights about translating empirical learning rates into appropriate inputs to energy models that look ahead as much as 100 years (Section 3.3). Note also that our use of the term 'learning curve,' and a single learning rate, to describe cost or performance in the aggregate differs from the way others sometimes use the terms. We will not review here all the different ways in which the terms have been used, but to avoid confusion we should note that our usage differs, in particular, from that of the Boston Consulting Group [10]. It uses the term 'learning curve' to refer to performance improvements related to one specific input, such as labour or raw materials. The aggregate pattern it refers to as an 'experience curve.'

2.2 Sensitivity of policy-relevant variables

To illustrate the sensitivity of policy-relevant variables to variations in estimated learning rates, we define two variables for a new technology, the *break-even capacity* and the *technology maturing costs*, as follows. If we assume for the moment that the new technology's *competitive cost target* is unchanging over time and if we know the technology's learning rate and current cost, we can calculate how much more capacity will have to be installed for the technology to become competitive. The capacity at which this happens is called the *break-even capacity*. The extra expenditures required to bring costs down to the competitive cost target are called the *technology maturing costs*.

Ignoring for the moment any discounting effects, the technology maturing costs equal the integral between the technology's cost curve and the horizontal line representing the competitive cost target.

Table 1 Learning rates of energy technologies

<i>Technology</i>	<i>Country/ Region</i>	<i>Time Period</i>	<i>Estimated Learning Rate</i>	R^2 ^a	<i>Performance Measure</i>	<i>Experience Measure</i>	<i>Reference/Data Source</i>
Oil extraction	North Sea	—	≈25%	—	labour (man-hrs)	units (platforms)	[11]
Crude oil at the well	USA	1869-1971	5%	—	sale price (\$/bbl)	production	[12]
Gas pipelines, onshore	USA	1984-1997	3.7%	0.09	investment price (\$/mile-inch ²)	capacity (mile-inch ²)	[13]
Gas pipelines, offshore	USA	1984-1997	24%	0.76	investment price (\$/mile-inch ²)	capacity (mile-inch ²)	[13]
Gas turbines	World ^b	1958-1963	22%	—	investment cost (\$/kW)	capacity (MW)	[14]
Gas turbines	World ^b	1963-1980	9.9%	—	investment cost (\$/kW)	capacity (MW)	[14]
Gas turbines	World ^b	1958-1980	13%	0.94	investment cost (\$/kW)	capacity (MW)	[1,14]
Coal for electric utilities	USA	1948-1969	25%	—	sale price to utility (\$/ton)	production (tons)	[12]
Electric power production	USA	1926-1970	25%	—	sale price (\$/kWh)	production (kWhr)	[12]
Nuclear power plants	OECD	1975-1993	5.8%	0.95	investment cost (\$/kW)	capacity (MW)	[15]
Hydropower plants	OECD	1975-1993	1.4%	0.89	investment cost (\$/kW)	capacity (MW)	[15]
Coal power plants	OECD	1975-1993	7.6%	0.90	investment cost (\$/kW)	capacity (MW)	[15]
Electricity from supercritical coal power	USA	n.a.	3%	—	production cost (\$/kWh)	production (TWh)	[16,17]
Coal power plants	USA	1960-1980	1.0%-6.4% ^b	—	investment cost (\$/kW)	capacity (units)	[17]
Lignite power plants	OECD	1975-1992	8.6%	0.96	investment cost (\$/kW)	capacity (MW)	[15]
GTCC power plants	OECD	1984-1994	34%	0.78	investment cost (\$/kW)	capacity (MW)	[15]
GTCC power plants	World	1981-1991	-11% ^c	0.41	investment price (\$/kW)	capacity (MW)	[18]
GTCC power plants	World	1991-1997	26% ^c	0.90	investment price (\$/kW)	capacity (MW)	[18]
GTCC	EU	n.a.	4%	—	production cost (\$/kWh)	production (TWh)	[16,18]
Wind power plants	OECD	1981-1995	17%	0.94	investment cost (\$/kW)	capacity (MW)	[15]
Wind power	Germany	1990-1998	8%	—	investment price (\$/kW)	capacity (MW)	[16]
Wind power	Denmark	1982-1997	4% ^d	—	investment price (\$/kW)	capacity (MW)	[16]
Wind power plants	Germany	1990-1998	8%	0.95	investment price (\$/kW)	capacity (MW)	[19]
Wind turbines	Denmark	1982-1997	8%	n.a.	investment price (\$/kW)	capacity (MW)	[20]
Wind power (electricity)	California	1980-1994	18%	0.85	production cost (\$/kWh)	production (TWh)	[21,22]
Wind power	USA	1985-1994	32%	—	production cost (\$/kWh)	production (TWh)	[16]
Wind power	EU	1980-1995	18%	—	production cost (\$/kWh)	production (TWh)	[16]
Electricity from biomass	EU	1980-1995	15%	—	production cost (\$/kWh)	production (TWh)	[16]

Table 1 Learning rates of energy technologies (continued)

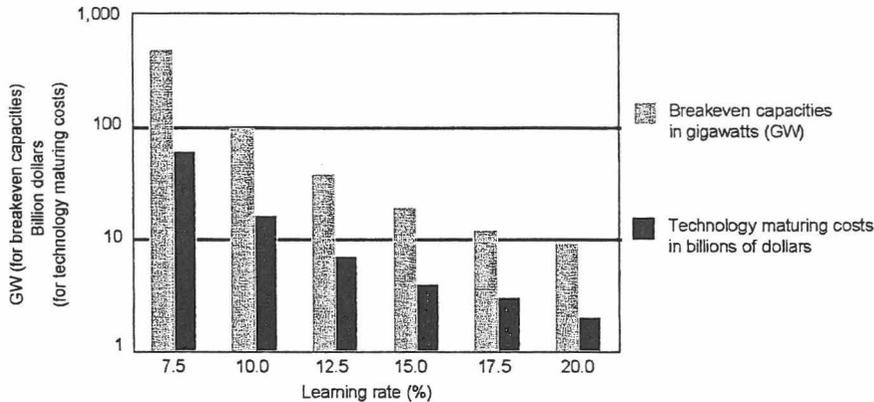
Technology	Country/ Region	Time Period	Estimated Learning Rate	R ² ^a	Performance Measure	Experience Measure	Reference/Data Source
Solar PV modules ^e	World	1968-1998	20%	0.99	investment price (\$/W _{peak})	capacity (MW)	[23]
Solar PV panels	USA	1959-1974	22%	0.94	sale price (\$/W _{peak})	capacity (MW)	[24]
Solar PV systems	EU	1985-1995	35%	–	production cost (ECU/kWh)	production (TWh)	[16]
Solar PV modules	World	1976-1992	18%	–	sale price (\$/W _{peak})	sales (MW)	[16]
Solar PV modules	EU	1976-1996	21% ^f	–	sale price (\$/W _{peak})	sales (MW)	[16]
DC converters	World	1976-1994	37%	0.35	conversion losses (%)	capacity (units)	[25]
Ethanol	Brazil	1979-1995	20%	0.89	sale price (\$/boe)	production (m ³)	[26]
Ethanol	Brazil	1978-1995	22% ^g	–	sales price (\$/boe)	production (m ³)	[16]
Retail petrol processing	USA	1919-1969	20%	–	production cost (\$/bbl)	production (bbl)	[12]
Model-T Ford	USA	1909-1918	14%	0.96	sale price (\$ per car)	production (cars)	[27,28]
Compact fluorescent lamps, integral-electronic type	USA	1992-1998	16%	0.66	sale price(\$ per lumen)	production (units)	[29]
Air conditioners	Japan	1972-1997	10%	0.82	sale price (Yen per unit)	sales (units)	[30]
4-function pocket calculators	USA	Early 1970s	30%	n.a.	sale price (\$ per unit)	production (units)	[24]
SONY laser diodes	–	1982-1994	23%	0.95	production cost (Yen per unit)	production (units)	[27]

Source: McDonald and Schratzenholzer [31]

- a Two cautions are in order concerning values for R². For each line in the Table, R² expresses the quality of the fit between the data and the estimated learning curve. However, R² values in different lines should not be compared because sample sizes are different. Second, R² measures the correlation for a straight-line fit to the logarithms of the dependent and independent variables. As linear regression minimises the sum of error squares, this means that relative rather than absolute errors are minimised.
- b The geographical scope of the data is not reported explicitly. The context suggests it is the whole world.
- c Note that these learning rates are based on prices, and one explanation of the negative 1981-1991 'learning rate' could be oligopolistic pricing behaviour.
- d Based on Neij [20]. The learning rate of 4% considers only wind turbines equivalent to 55 kW or larger. The 8% learning rate reported two entries below includes all Danish wind turbines.
- e Based on preliminary data.
- f 21% is the learning rate for the 'stability' stage shown in Figure 7. For the 'development' and 'price umbrella' stages the learning rate is 16%. For the 'shakeout' stage it is 47%.
- g 22% is the learning rate for to the 'stability' stage shown in Figure 7. For the 'development' and 'price umbrella' stages the learning rate is 10%. For the 'shakeout' stage it is 53%.
- h Joskow and Rose estimate a range of learning rates for different utilities, architect-engineering firms, and technology categories, after accounting for inflation, plant size, the inclusion of scrubbers or cooling towers, whether certain structures are indoors or out, and whether a unit is the first on a site.

With these definitions, we can demonstrate the sensitivity of break-even capacities and technology maturing costs in Figure 1. The Figure shows a hypothetical situation in which the competitive cost target is \$1 per watt (\$1000/kW) and initial costs are twice as high at \$2/W. Break-even capacities and technology maturing costs are calculated for learning rates between 7.5% and 20%.

Figure 1 Sensitivity of break-even capacities and technology maturing costs



Note that the vertical axis in Figure 1 is logarithmic. Because the horizontal axis is linear, the most important observation from the figure is that with decreasing learning rates, technology maturing costs and break-even capacities grow faster than exponentially. In terms of absolute values, going from a learning rate of 20% to one of 10% increases technology maturing costs from \$2 billion to \$16 billion and break-even capacities from 9 to 96 gigawatts (GW). Although these cost figures are illustrative, they are nonetheless in a realistic range and the calculated results are thus indicative of realistic orders of magnitude.

This is particularly important for novel technologies, for which learning curves are not statistically estimated but constructed directly from a postulated learning rate and initial values for costs and the total installed capacity at time zero. Estimating the initial capacity can be difficult, making the sensitivity of estimated technology maturing costs and break-even capacities with respect to this parameter especially important. Because both technology maturing costs and break-even capacities depend linearly on the initial capacity, their sensitivity to variations in the initial capacity is linear. If the initial capacity is actually twice the value assumed, for example, that doubles the technology maturing costs and break-even capacities.

While the simple definitions of break-even capacities and technology maturing costs introduced here indicate the sensitivity of policy variables to variations and uncertainty in estimated learning rates, we recognise that ultimately more detailed calculations will be needed. In particular, competitive cost targets drop over time as other technologies also improve. Thus more detailed energy models are needed to answer completely policy questions as to whether near-term subsidies to speed currently expensive but promising technologies down their learning curves will be fully compensated by eventual cost reductions (including external costs if those are known).

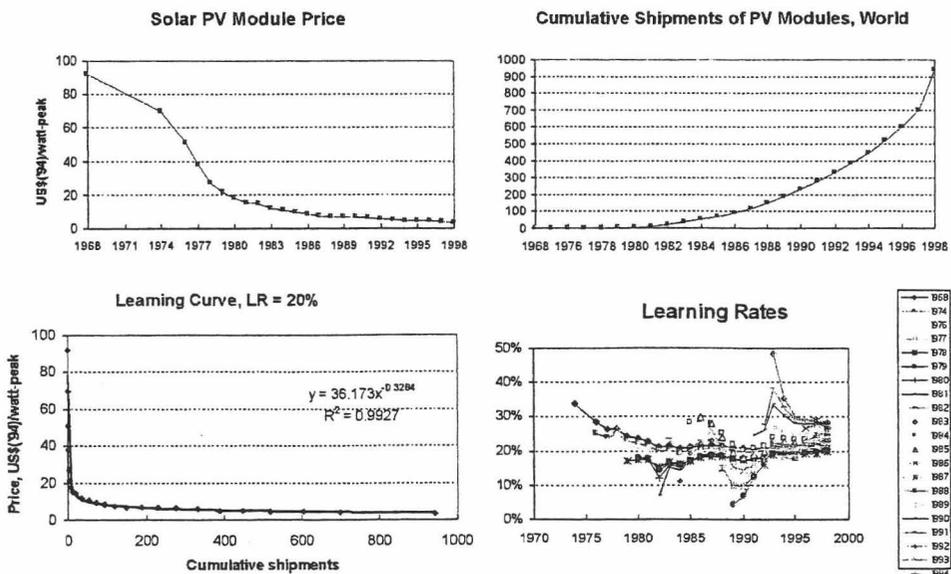
3 Energy technology learning curves

This section describes the formats we have used for assembling data, estimating learning rates and presenting these together with estimates published by others. Table 1 summarises the results, both calculated and collected. These are compared in Section 3.2 with the range of learning rates in manufacturing activities that Dutton and Thomas published in 1984 [32]. Section 3.3 provides additional observations on selected technologies from Table 1.

3.1 A format for presenting learning curves

As noted above, empirical learning curves are usually derived from time series of costs and capacities and thus carry invisible time labels. To keep track of these time labels, among other things, we used the format shown in Figure 2 for assembling the database underlying many of the learning rates presented below in Table 1. The top panels of Figure 2 show the time series for the performance indicator (specific prices of solar PV modules, in this case, on the left) and the experience indicator (cumulative shipments, on the right). The lower left panel shows the fitted learning curve, with its estimated learning rate and correlation coefficient (R^2) [33]. The lower right panel represents the goodness-of-fit in further detail, showing all the learning rates that can be calculated from any two points in the data set. For example, the curve labelled '1968' (the top label in the legend box) describes the learning rates between 1968 and the year described by the value on the horizontal axis. In this example, the curves in the lower right panel seem to show more variability than revealed by the lower left panel's R^2 value of 0.99.

Figure 2 Solar PV modules: prices (upper left), cumulative shipments (upper right), learning curve (lower left) and learning rates (lower right)



Source: Harmon [23]

Both time series at the top of Figure 2 are relatively smooth. In cases where the curves are less smooth, the timing of irregular ups and downs may help interpret results. Particularly when performance is measured by specific *prices* (rather than costs), fluctuations around the long-term trend may reflect market shocks or non-equilibrium pricing strategies (Sections 3.3.3, 3.3.6 and 3.3.7 address specific examples). Looking at the underlying time series can help identify such instances. These may reduce values of R^2 , but such fluctuations are largely irrelevant for long-term energy scenarios and their effects on R^2 should probably be ignored if estimated learning rates are to be used in such scenarios.

Keeping careful track of the underlying data, as in Figure 2, can also help avoid potential analytic pitfalls, particularly in terms of mismatches between time series for the performance and experience variables. Two examples emphasised in a recent International Energy Agency report [16] are the following. The first involves combining total installed *capacity* of a technology as the experience variable with *production* costs as the performance indicator. Consider the case of wind turbines. Accumulated installed capacity reflects the experience of companies producing wind turbines. As their experience increases, the costs of producing turbines should drop, resulting in a corresponding drop in the specific investment costs for utilities buying turbines. But the utilities' total costs of producing electricity include the additional impact of improving capacity factors – improvements related to utility experience best measured in terms of cumulative electricity *production*, rather than capacity. The second example involves misleading mismatches between *world* market prices as a performance indicator and cumulative *national or regional* capacities as the experience indicator [34].

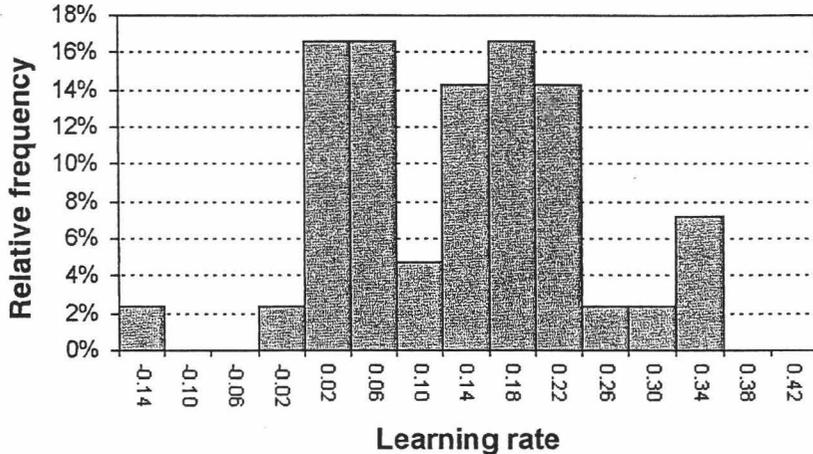
Together, the formats of Figure 2 and Table 1 (which identifies the geographic scope of each dataset) provide a succinct summary of much of the information needed to go beyond estimating historical learning curves to more comprehensive technology assessments. One objective of such assessments is a set of learning rates appropriate for long-term future energy scenarios. In addition to the summary of important ancillary information provided by Figure 2 and Table 1, there is a further motive for presenting these formats here. By so doing we hope to encourage other authors to use similar formats in assembling and presenting their own data in future publications.

3.2 Overview

Table 1 presents 42 learning rates that we have either calculated for different energy technologies (26) or collected from the literature (16) [31]. The distribution of all 42 learning rates is shown in Figure 3. The overall pattern is quite irregular. Even if we disregard the outlying negative 'learning' rate, the distribution has 'holes' that cannot be plausibly explained and three apparent modes. One possible reason is that our sample size is simply still too small and that as more learning rate estimates are added to the distribution, it will become more regular. Another reason could be that in those cases where we could not estimate learning rates ourselves, different methods or conventions were used. One observation supporting this possibility is that the distribution of just those learning rates for which we did the original estimation, or re-estimated published learning rates, is more regular. In any case, the observed irregularities were an important motivation for us to 'dig deeper', i.e., to look at each technology in more detail – some examples are given in Section 3.3 – to understand additional factors (like market shocks

and predatory pricing strategies) that might cause bias or variability in the underlying data sets.

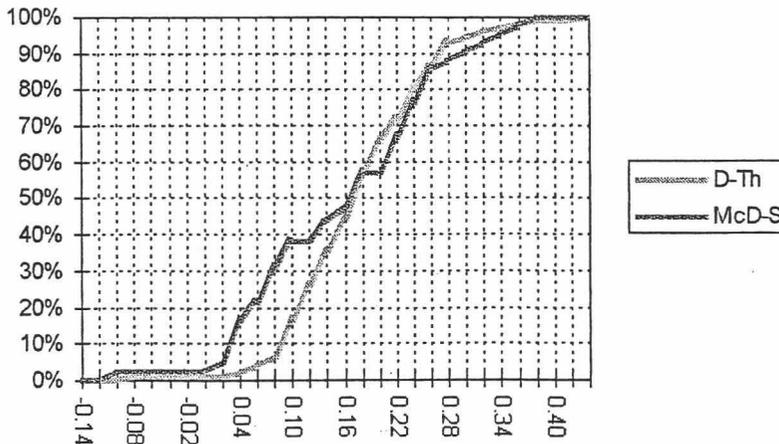
Figure 3 Learning rates of 42 energy technologies



Source: McDonald and Schratzenholzer [31]

A smoother representation of the distribution of learning rates in Table 1 is provided by the cumulative curve, as shown in Figure 4. For comparison, Figure 4 also shows the distribution of manufacturing learning rates from Dutton and Thomas. Figure 4's presentation leads to a smoother image and shows that, by and large, the two distributions are similar, in particular above the median – near a 16% learning rate for both curves. The biggest divergence between the curves is between the 10th and 40th percentiles. The 20th percentile for energy technologies is at a learning rate of 6%. For manufacturing it is at 13%.

Figure 4 Cumulative distributions of two surveys of learning rates



Source: Dutton and Thomas [32] and McDonald and Schratzenholzer [31]

Given the irregular shape of the distribution in Figure 3, however, it is important to look at the data behind Table 1's learning rates for possible misleading biases. Identifying such biases is the first step in moving beyond the exercise of estimating historical learning rates towards a technology assessment using past learning rates as just one input among several for generating prospective learning rate distributions for use in long-term energy scenarios. An examination of all Table 1's entries is beyond the scope of this paper. But the next section highlights possible biases in eight selected entries, chosen to cover the technology range from primary energy extraction to final energy use.

3.3 *Towards technology assessment*

3.3.1 *North Sea oil extraction*

Technological progress in oil extraction is one of the remarkable success stories of the past 20 years. It has enabled producers in the North Sea to stay in business through, in particular, the drastic decline in international crude oil prices in 1985 and 1986. Unfortunately for our efforts in estimating a learning rate, oil extraction technology costs are not very transparent and extraction cost reductions are therefore difficult to compile. Our estimate is based on a small amount of performance-related numerical information presented by Blackwood [11], who reports reductions in the average number of man-hours needed to construct one ton of three platform jackets for the Eastern Trough Area Project. According to Blackwood, 'significant improvement' was achieved in comparison with an earlier project (Andrew) requiring 66 man-hours per ton and a knowledgeable source 'would not confirm' a proposed estimate of 50 man-hours per ton for the Eastern Trough Area Project. If we assume that the 50 man-hours per ton conjecture is correct and that the construction was the second of its kind (i.e., that it constitutes the first doubling), we get a progress rate of $50/66 = 0.75$ and a learning rate of 25%. There are two principal uncertainties in this calculation. First is the uncertainty about the 50 man-hours per ton as the right number to attach to the phrase 'significant improvement.' Second is the assumption that the improvement from 66 to 50 man-hours corresponds to the first capacity doubling for this technology. If Andrew and the Eastern Trough Area Project represent points beyond the first capacity doubling, then the improvement from 66 to 50 man-hours would correspond to a learning rate higher than 25%. Both sources of uncertainty would have to be taken into account when developing a learning-rate distribution for this technology that might be input into long-term energy scenarios.

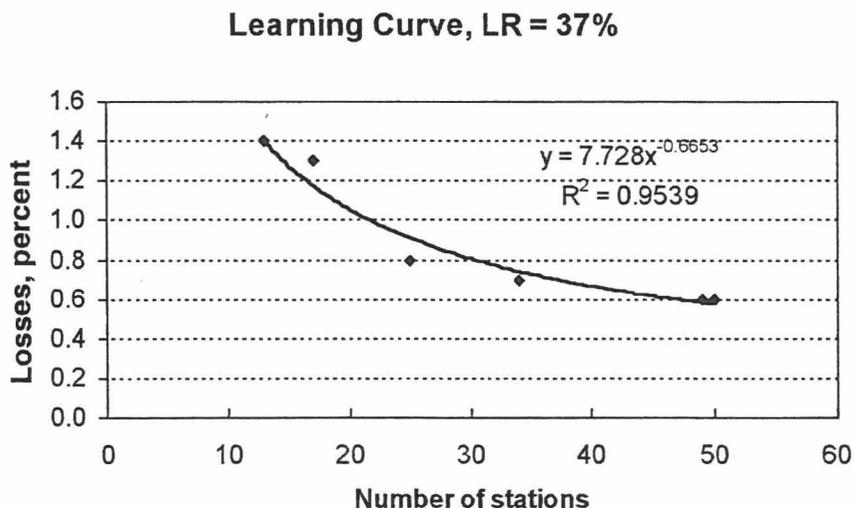
Two other caveats concern the fact that the improvement measure in this case is not specific cost. In all such cases with non-cost performance indicators, we must remember that, at any point in time, higher efficiency versions of a technology (from light bulbs to power plants) are generally available at a higher cost. Thus if preferences (or government regulations) shift toward higher efficiency versions of a technology, for example and this shift is not taken into account, resulting estimates of learning rates may be misleadingly high. The second caveat is that improvements in parameters like efficiency and productivity may be significantly motivated by *increasing* costs in various factors of production. The learning rate calculated during such periods of market pressure may well be higher than the learning rate most appropriate for long-term energy scenarios. We have no reason to believe that market pressure or preference shifts toward higher efficiency technology played distorting roles in the case of North Sea oil, but our use of a

non-cost performance indicator in this case makes it the appropriate place to raise a warning flag.

3.3.2 AC/DC converters

Moving from primary energy extraction to energy transmission, we have included in Table 1 one learning rate based on a study by Rabitsch of electricity grids and transformers in Eurasia [25]. Rabitsch also uses a non-cost performance indicator, specifically losses in converter stations connecting AC and DC lines. Losses are currently in the range of 0.5 to 0.7% of the rated power, down from 1.4% in 1976. Based on these values and data from several ABB projects, Rabitsch estimates a learning rate of 37%, with a comparatively high correlation coefficient of 0.95 (Figure 5).

Figure 5 Converter losses for a single station versus the cumulative installed capacity of all DC systems in the world



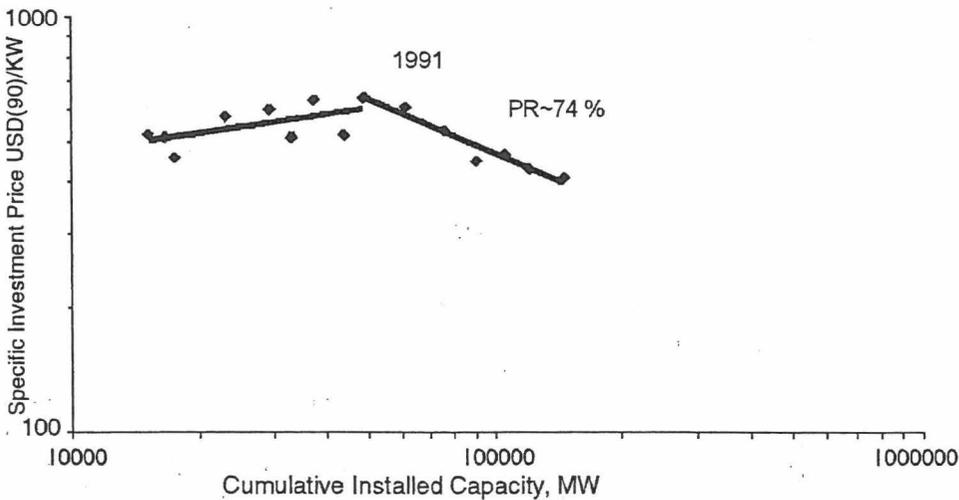
Source: Rabitsch [25]

As was the case for platform jackets in the North Sea, these non-cost performance improvements may have been partly caused by changing market circumstances. In particular the optimal design of a converter is usually driven by expectations (or intentions) about the appropriate interest rate, economic lifetime, utilisation factor and costs of electricity, among other factors. Changes in these factors will influence converter designs and the resulting losses. However, in the absence of a more complete description from Rabitsch or others of changes in these factors during the period covered by his data, we essentially assume that their effects are either consistent or largely random with respect to experience accumulation. In either case, for long-term energy modelling it would be appropriate to anchor a learning rate distribution on the value of 37% estimated by Rabitsch.

3.3.3 Gas turbines

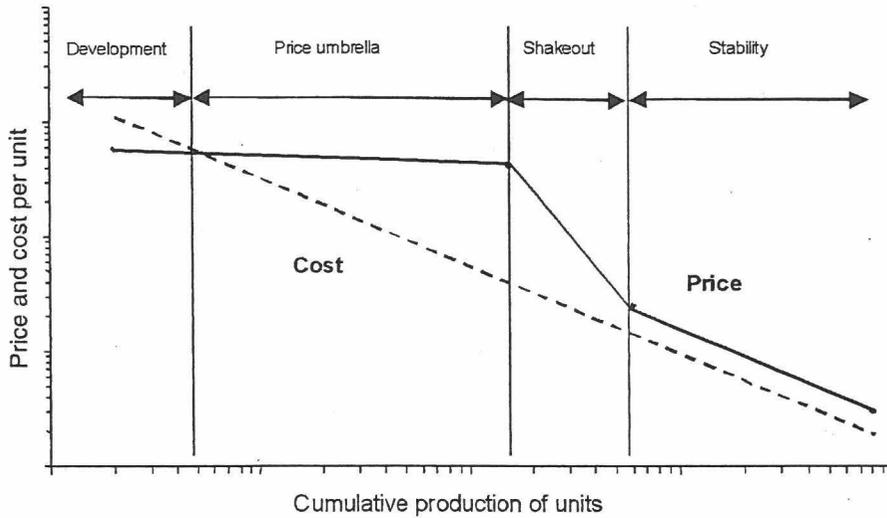
Figure 6 presents data assembled by Claeson [18] for gas turbine combined-cycle power plants (GTCC). She fits a two-stage learning curve over the period from 1981 to 1997 using specific investment prices as the performance indicator. The first part of the curve (1981-1991) shows an *increasing* trend in GTCC investment prices, corresponding to a 'learning' rate of -11% and a bad fit ($R^2 = 0.37$). The second part, after 1990, implies a learning rate of 26% with $R^2 = 0.90$. (Estimating a single learning curve for the full data set results in a correlation that is close to inconclusive ($R^2 = 0.18$) and a 'learning rate' of 5.4%.) One explanation for the irregular data pattern could be the high variability of the original data. Due to different conditions prevailing in different cases (with respect to the size, location, changing pollution abatement requirements, the exact definition of costs, etc.), strict comparability of data from different years cannot be assumed. Claeson does average all prices for each year to get one average data point per year and this smoothes out part of such variability. But it is unlikely to compensate for it all. A second partial explanation arises from the fact that technological performance in this case is measured in terms of prices rather than costs. The use of prices instead of costs is discussed in more detail in the next section on wind energy. As pointed out by Claeson, for GTCCs, this raises the possibility that the negative 1981-1991 'learning' rate could be a result of oligopolistic pricing behaviour during that period. The subsequent price reduction could then be the 'shakeout' described in Figure 7.

Figure 6 Two-stage learning curve of GTCCs, 1981-1997



Source: Claeson [18]

Figure 7 Price-cost relations for a new product



Source: IEA [16] and BCG [10]

3.3.4 Wind energy

Wind turbines for electricity production are an increasingly popular technology as a result of their substantial external environmental benefits. Although they are still, in many circumstances, more expensive than their competition, electricity production costs have dropped by more than an order of magnitude within less than 20 years and many expect that there is still a lot of room for further improvement.

Given the current popularity of wind power, Table 1 includes seven different wind data sets and their calculated learning rates. These fall into two groups: one for electricity production (with performance and experience indicators based on production) and one for wind turbines (with performance and experience indicators based on capacity). The learning rates based on production are significantly higher than those based on capacity. One explanation is the significant technological progress that has occurred in turbine design to lower the threshold wind speed for power production. Thus one unit of capacity generates an increasing number of kilowatt-hours.

Another possible cause of differences among wind power learning rates is illustrated by comparing the 32% learning rate for 1985-1994 US wind electricity production, reported by IEA [16], to the much lower rate of 18% that we estimated for California for 1981-1994. If we re-estimate our learning curve for California for the same period as the IEA estimation (1985-1994), we get a learning rate of 34%, a figure much more in line with IEA. Because both the data series in question are in terms of prices rather than costs, one explanation for this result may be the hypothesis suggested originally by the Boston Consulting Group and shown in Figure 7. The Figure presents a possible relationship between a straight-line learning curve for costs (the dashed line) and price reductions (the kinked solid line) that are driven by assumed changes in market structure as well as by declining costs. The market structure is assumed to move through four characteristic

stages. The first two stages ('development' and 'price umbrella') reflect an oligopolistic market that manages to keep the learning rate for price reductions below the learning rate for costs. In the 'shakeout' stage the learning rate for prices is higher than that for costs and in the 'stability' stage, learning rates for prices and costs are identical. The 1981-1994 data set for California on which our calculations are based suggests a transition similar to the first kink in the solid line in Figure 7. To the extent that such a transition explains the higher learning rate estimates of 34 and 35% from post-1985 data, such estimates would be too high to use as approximations for the underlying constant learning rate for costs (i.e., corresponding to the dashed line in Figure 7).

3.3.5 Solar photovoltaic modules

As was the case with wind energy, a number of learning curves have already been published for solar photovoltaic technology. One recurring problem with past estimates is that the technology has not always been defined precisely. Given that for decision and policy makers, it is probably electricity costs that will matter most in their economic assessment of solar photovoltaic technology, analyses should focus on the development of total PV system costs rather than just modules. Still, most reported learning curves have analysed only solar cells or modules.

Harmon [23] has surveyed the existing literature on solar PV and has estimated the learning curve shown in the lower left panel of Figure 2. It is based on a compilation of the data behind currently published learning curves, mainly Ayres *et al.* [35]. The Figure presents cost data for modules, rather than total solar systems as recommended above, for two reasons. One was simply that a lot more data are available for modules. The second reason – which also partly explains why module data are more available – is that systems come in several variants, which serve different purposes and have different costs. Harmon includes a detailed discussion of the many possible non-module, or balance-of-system (BOS), cost factors and their possible future developments. Among other things, these include structural components, AC converters, other power conditioning equipment and energy storage.

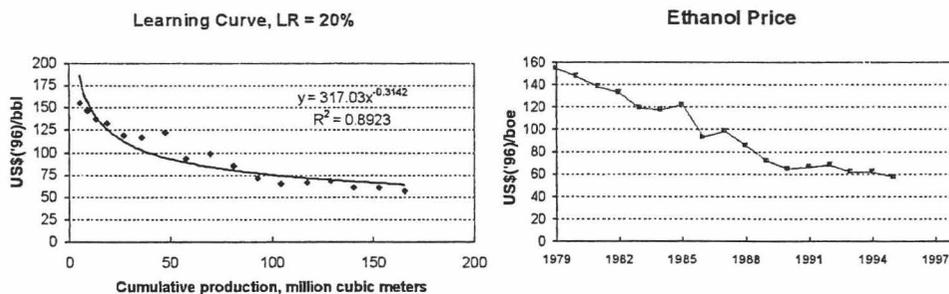
The correlation coefficient ($R^2 = 0.99$) in Figure 2 is rather high. Moreover, the estimated learning rate of 20% is quite consistent with other published PV learning rates based on shorter time intervals. For long-term energy modelling, therefore, we would be justified in using a distribution for the learning rate for PV modules (around a mode of 20%) that is tighter than those for the other technologies discussed in this section. Harmon also found little difference among the different types of PV modules (e.g., mono and poly-crystalline silicon, or thin-film, plus their sub-classes) that have dominated the industry at different times. Similarly, the geographical coverage of the data suggests significant geographical invariance, i. e., that PV learning rates are quite similar in Japan and the USA.

3.3.6 Ethanol in Brazil

Goldemberg [26] presented an analysis of technological learning in ethanol production for Brazil. From his raw data on ethanol prices paid to producers (which serve as an approximation to costs) and cumulative ethanol production, we have plotted the learning curve shown in the left part of Figure 8. The results are a learning rate of 20% and a correlation coefficient of 0.89. The fit appears to be quite good, except perhaps for prices

between about 80 and 120\$/boe. Figure 8 shows the underlying time series for ethanol prices. From this we can see that the divergence of the data between 80 and 120\$/boe from the fitted learning curve corresponds to high ethanol prices in particularly 1984 and 1985, followed by an abrupt drop in 1986. This suggests that the decline in ethanol prices between 1985 and 1986 was partly driven by the concurrent sharp drops in the prices of international oil and, subsequently, petrol, ethanol's main competitor. Similarly, the rise in the price of ethanol after 1990 could be connected with general uncertainties in the international oil market following the 1990-91 Gulf crisis. Both of these causes would be unrelated to technological change in ethanol production and would therefore be best ignored when developing an ethanol learning rate distribution for long-term energy modelling.

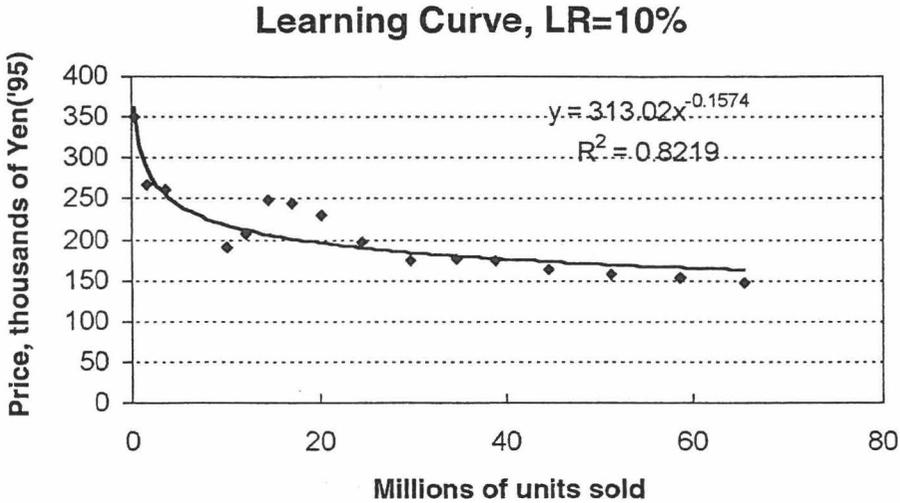
Figure 8 Learning curve estimated for ethanol production in Brazil (left) and prices paid to producers (right), 1979-1995



Source: Goldemberg [26]

3.3.7 Air conditioners in Japan

Akisawa [30] has studied the technological and market development of room air conditioners in Japan since 1961. Figure 9 shows a learning curve estimated from price data for a 'heat-pump' design (i.e., it can work both as an air conditioner and a heat pump) that currently dominates the Japanese air conditioning market. The estimated learning rate for the entire data series is 10%, but the correlation coefficient (R^2) of 0.82 is not very high. Most of the problem appears to be due to irregular price behaviour in the early periods [36], corresponding to prices between 250,000 and 200,000 Yen. From 1990 onwards (when prices drop – and stay – below 200,000 Japanese Yen per unit), the fit becomes much better. Considering just the period from 1990-1997, we get a learning rate of 17% and a correlation coefficient of 0.94.

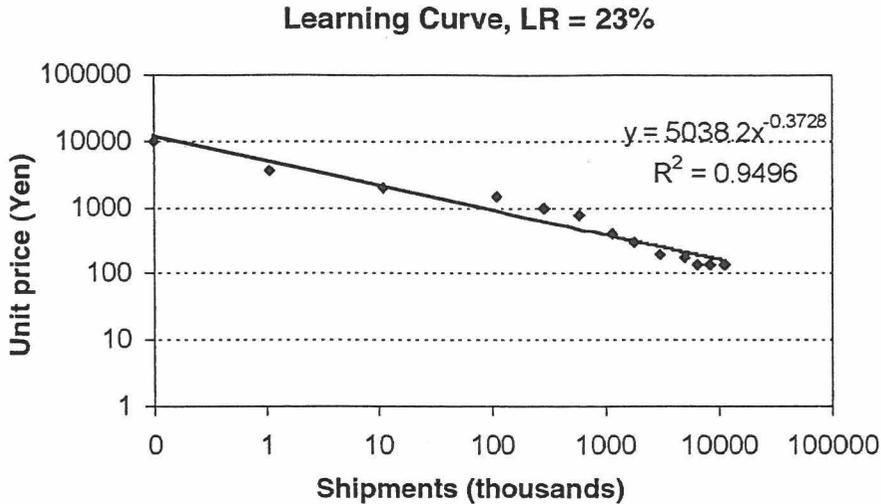
Figure 9 Learning curve (10% learning rate) of air conditioners (heat pump type), Japan, 1972-1997

Source: Akisawa [30]

Again we are dealing not with costs, but with prices and should therefore consider plausible factors other than costs that might have an impact on prices. Because the early erratic price behaviour in Figure 9 corresponds to the period when the now dominant design faced the most competition, some of the ups and downs could reflect pricing strategies related to a changing market structure analogous, but not identical, to Figure 7. In this case the post-1990 learning rate of 17% would seem the most appropriate for long-term energy modelling.

3.3.8 SONY laser diodes

The final item from Table 1 on which we want to make additional observations is laser diodes. These constitute not so much an energy end-use technology as we normally define the term, as they do an important component in an increasing number of true energy end-use technologies, such as compact disc players. We have chosen laser diodes for two reasons. First, the data (see Figure 10) stretch over almost six orders of magnitude of cumulative production and are suggestive of the hypothesis that a technology's learning rate decreases as it moves down its learning curve. Lipman and Sperling [28], whose data are presented in Figure 10, argue that the data seem to follow a learning rate of first 25% and then 20% and then eventually reach a floor. On the other hand, while the unchanging prices over the last three data points do support the notion of 'floor costs,' the data also show a wave-like pattern that would suggest an eventual return to more rapid learning. This raises the second reason for highlighting laser diodes – given current production of six million units per year and the present cumulative production level of 10.8 million, more doublings in the near future will make it possible to test particularly the 'floor' hypothesis against real data.

Figure 10 Learning curve (23% learning rate) of Sony laser diode manufacturing costs 1982-1994

Source: Lipman and Sperling [27]

4 Variations on the basic concept and discussion

As noted in several of the examples above, learning is one of several possible factors contributing to observed cost decreases. Other possibilities include:

- 1 R&D expenditures
- 2 unit size (economies of scale)
- 3 unit lifetimes (technologies that turn over frequently have more opportunities to incorporate learning)
- 4 market size (technologies where annual production is of the order of millions of units have more opportunity for learning than technologies produced in the tens or hundreds)
- 5 the stage of the technology's lifecycle (perhaps technologies still in the R&D phase learn faster than commercialised technologies)
- 6 time (ageing equipment may increase production costs and thus bias performance indicators; experience may depreciate over time and general technological progress in subcomponents like computers may reduce overall costs)
- 7 progress in closely related technology clusters (progress in stationary fuel cells may have spillover effects on mobile fuel cells)
- 8 market structure and prices (more competitive markets and prices increase incentives for cost reductions) and

- 9 government regulations (environmental mandates, for example, may slow or reverse cost reductions).

Even more detail is considered in many econometric analyses of particularly US power plant costs. These include factors such as regional labour costs or whether sites at different latitudes require different amounts of construction to be done indoors [17,37–41].

The additional possible causes of cost reductions that are most important to examine depend on one's objectives. If you are a decision maker in a US utility pondering when and where to build your next power plant, you probably care about cost variations between northern and southern building sites, or between union and non-union states. But these are of less interest to our principal audience – developers and users of long-term, global or national energy models. Moreover, for long-term energy models it is not clear how much effort should be put into trying to distinguish among factors that may be highly correlated. For example, given the data that are available, model inputs in which learning and scale economies are lumped into a single estimated learning rate may be simpler, as reliable and therefore more useful than efforts to extract the two separate effects from the empirical data and then treat them separately in long-term energy models.

More generally, the descriptive power of estimated learning curves is still important and valuable even if their explanatory power is imperfect. In much the same way that the standard error of a given data sample can be informative whether or not the sample is distributed normally, learning rates can be informative indicators of time series relating technological performance to experience. A similar example is the extraordinarily rich set of logistic curves collected by Marchetti [42], Nakićenović [43], Grübler [44] and others. Their regularity is impressive even in cases where the original model (the differential equation defining the logistic function, describing species growth in a limited environment) may not be a satisfactory representation of the market substitutions reflected in the data. From this perspective it is less necessary to analyse in detail the separate effects of the items above than it is to find descriptive phenomenological correlations. For example, Neij [20] classified technologies into three categories – plants, module technologies and continuous processes – for which she found average learning rates of 10, 20 and 22% respectively.

Given our ultimate objective of compiling learning rates useful for long-term energy modelling, the above list is particularly useful in two ways. For some items, it helps us to pre-process the raw data to extract a learning rate appropriate for long-term scenarios. An example is the model in Figure 7, explaining differences between prices and costs as a function of market structure. For long-term energy modelling we are interested in the learning rate corresponding to the 'stability' portion of the curve and it is indeed such learning rates that are reported in Table 1 for two items, solar PV modules in the EU between 1976 and 1996 and Brazilian ethanol between 1978 and 1995 as calculated by the IEA (see footnotes to Table 1).

Secondly, some of the items appear to have long-term effects and should be considered for direct inclusion in long-term modelling of learning. Examples are the first, fifth and sixth items, which concern possible learning rate changes as time passes and a technology matures from its R&D stage to widespread adoption. Several authors address this possibility using kinked learning curves that are only piece-wise linear on a double-logarithmic scale (see, for example, Nakićenović *et al.* [1]). Although this is motivated

partly by efforts to fit specific data sets better, it has been justified by qualitative arguments that learning rates should be higher for technologies still in their R&D phase, then decrease as technologies become commercialised and then decrease again when they can be labelled 'mature.' (Note that the argument that learning rates should be higher in the R&D phase suggests that R&D expenditures could be a mechanism for steepening learning curves and thus complement policy mechanisms discussed in Section 4.2 for moving faster down a given learning curve of unchanged slope.)

Section 4.1 presents an alternative to kinked learning curves that models decreasing learning rates as the result of continuous knowledge depreciation, rather than sharp transitions among life-cycle stages. Turning to potential policy considerations, Section 4.2 then addresses one anticipated potential result of fully endogenising technological learning in energy models – technological lock-in.

4.1 Knowledge depreciation

Knowledge and the relevance of knowledge depreciate with time. A well-known example of knowledge depreciation, or 'forgetting by not doing,' is the one case of negative learning reported by Dutton and Thomas [32]. It describes Lockheed's experience producing the L-1011 TriStar aircraft. As production increased from 1972 to 1975, costs reductions were generally consistent with a standard learning curve [37]. But production cuts began in late 1975. In 1976 only six planes were produced, compared to 41 in 1974. Although Lockheed geared back up to 24 planes in 1979 and 25 in 1979, real unit costs rose after the 1975 cuts and appear to have remained above the selling price until Lockheed decided to phase out production altogether in 1981. The overall result was a negative 'learning' rate, much of which is often attributed to the loss of knowledge associated with laying off many experienced workers starting in 1975, only to have to hire new recruits later to ramp up production.

Argote models knowledge depreciation by including knowledge (K) as a factor of production in a Cobb-Douglas production function, with knowledge defined as

$$K_t = \lambda K_{t-1} + q_t, \quad 0 \leq \lambda \leq 1 \quad (4.1)$$

where q_t is production at time t and λ is a knowledge depreciation parameter to be estimated. If $\lambda=1$, there is no knowledge depreciation and knowledge equals exactly cumulative production. At the other extreme, $\lambda = 0$ implies no knowledge accumulation.

Argote estimates values of λ for data sets describing Liberty Ship construction in the USA during World War II, a North American truck plant and fast-food franchises. The estimates for λ are 0.75 for the ships (for monthly time steps), 0.989 for the trucks (for weekly time steps) and 0.83 for the fast-food franchises (for *weekly* time steps). Corresponding learning rates for the ships and trucks are 26% and 17% respectively. (No learning rate is given for the fast-food franchises.) The estimated values of λ mean that over the course of a year, experience depreciates to only 3% of its initial value in the shipyards and to 56% of its initial value in the truck plant. For the fast-food franchises, knowledge loses over half its value each *month*.

Given the differing estimates of knowledge depreciation across studies, Argote (whose focus is more on organisations than technologies) hypothesises that more technologically sophisticated organisations have less knowledge depreciation. She notes

that for high-technology organisations a greater portion of their knowledge is embedded in hardware and computer software where it may be more resistant to depreciation.

If we adjust equation 2.1 to include depreciated experience in a manner analogous to Argote's production function formulation, we get the following:

$$C_t = A * (Q_t)^b$$

where

$$Q_t = \lambda Q_{t-1} + q_t \quad (4.2)$$

How the adjustment in equation 4.2 bends the straight line that equation 2.1 leads to when plotted on double-log scales, depends on how production changes with time. If production grows exponentially, equation 4.2 still defines a straight line on a double-log plot, but if annual production stops growing and becomes a constant, equation 4.2 starts to flatten and relatively quickly becomes horizontal. This reflects the fact that for uniform annual production each capacity doubling takes longer than the last. That extra time means that knowledge depreciation has a bigger effect on later doublings. The percentage cost reduction for each doubling would consistently decrease and the learning curve would bend toward the horizontal. Such a learning curve would be effectively a continuous sister of the 'kinked' learning curves mentioned above. Kinked learning curves have the disadvantage relative to equation 4.2 that they require largely arbitrary definitions of the boundaries between different life-cycle stages.

With respect to the data reported in Table 1, knowledge depreciation may account for part of the variation in the EC-TEEM [7] learning rates across different power production technologies. The learning rates for established technologies (nuclear, hydro, coal and lignite) are well below that for today's preferred choice for new power capacity, GTCC. If capacity doublings are occurring more slowly for the established technologies than for GTCC (which they are), knowledge depreciation should have a greater opportunity to depress their calculated learning rates relative to that of GTCC.

4.2 Implications for policy assessment

The ultimate objective of this research is useful advice to policy makers. To illustrate how we envisage that goal, this section begins by outlining one possible policy assessment application that could be undertaken once all the data and model extensions are in place.

The premise is that people with limited planning horizons – i.e., energy businesses looking only at the next 10 or 20 years, or consumers looking at only the next one or two years – will tend to underinvest (from the long-term global perspective) in new technologies that are currently expensive. The fact that, thanks to learning effects, these technologies have the potential to become important inexpensive clean contributors to the energy system does not influence their purchasing decisions. Advocates of solar power, wind power and fuel cells would all consider their technologies to fall at least partly into this class, where consumers and companies, left to their own devices, are likely to underinvest relative to the long-term social interest. Where the market fails to serve perceived social interests, we all naturally turn to governments to compensate. This is the logic behind government subsidies, in all their myriad forms, for new technologies such as wind, solar and fuel cells. It is also partly the logic behind government procurement, e.g., new-technology buses for public transportation systems. And it is part of the logic behind government technology mandates – e.g., green certificates to show that by a given date,

say, 10% of the electricity we consume comes from renewable sources. In the first instance, subsidies will lower the consumer's price and encourage use. Expanded use means quicker progress down the learning curve. Government purchases directly increase use and thus speed progress down the learning curve. And mandates that force consumers to buy more of a new technology than economic considerations would warrant also increase use and accelerate progress down the learning curve.

To compare such policies and different levels for subsidies, purchases and mandates, one would start by running a model that includes full learning effects and perfect foresight. For such a case, the imposition of carbon constraints would be likely to result in scenarios that include earlier investments and/or R&D expenditures in currently expensive low-carbon technologies than in a situation with exogenous technological progress. This is because the model – thanks to perfect foresight – ‘knows’ that these will pay off in the long run. To incorporate the limited time horizon of purchasing and investment decisions, the model would then be run with limited foresight. Depending, among other things, on the assumed degree of risk aversion and discount rate, this would lead to a different portfolio of ‘optimal’ investments in technologies and a different time pattern and portfolio of R&D expenditures. The final step would be a series of model runs, again with limited foresight, but this time incorporating various alternative technology acceleration policies. The resulting costs and environmental impacts would be used to compare the relative efficiency and effectiveness of different policies in steering an imperfect myopic world towards the benchmark provided by the scenario with perfect foresight.

One potential difficulty faced by this analytic approach is the possibility of technological lock-in. To illustrate technological lock-in consider the example of the QWERTY keyboard – the standard keyboard design in which the first six letters on the top row spell QWERTY. This is clearly not the only possible arrangement of keys and Paul David and others have argued that there exist alternatives that were found ergonomically superior [45–47]. But the QWERTY keyboard is as pervasive as ever. At least two arguments are given for its success. First, when the QWERTY keyboard was introduced, the mechanical hammers that actually struck the typewriter ribbon and paper were less likely to jam together than was the case with other designs. Second, early chance events gave the QWERTY keyboard an initial advantage and thus an essential headstart toward becoming transformed into a standard. Typing courses began to teach typing on QWERTY typewriters, businesses purchased QWERTY typewriters and expected new recruits to be proficient on them, suppliers supplied them and manufacturers manufactured them. Now, particularly in computer applications, hammer-lock has disappeared as an issue. But we are locked in to the QWERTY design and the alternative designs advertised as ergonomically superior have been locked out of the market.

This is an example of lock-in largely through so-called network effects – i.e., the value of QWERTY keyboards and proficiency for each of us increases as more and more people use them. Technological learning can also contribute to a form of lock-in. As experience with early technological innovations improves their costs and performance, it widens their advantage over later competitors and makes it increasingly difficult for late-starting alternatives to catch up. This means that technologies entering into the energy market early have a significant increasing advantage – other things being equal. Future developments are partially locked in to the direction set by early innovations. (This is not

necessarily bad. Environmentalists might be very happy to see technological learning help lock in a chance head start enjoyed by some renewable technologies – e.g., for synthetic fuels – over fossil alternatives.)

For policy makers this result raises the stakes surrounding early investments. Get them right, the logic seems to argue and thanks to learning the ‘system’ will spread the right technologies to all corners of the world. Get them wrong and you may never be able to undo the harm. While this is an overstatement, it emphasises the point that learning can magnify the impacts of early cost-reducing investments. Nonetheless several caveats are in order. A good place to start is a critique of exactly the example we used to introduce the notion of lock-in, the QWERTY keyboard.

In 1990 Liebowitz and Margolis published a paper entitled ‘The fable of the keys’ [48] and have more recently produced a full-length book criticising mistaken analyses of path dependence and lock-in [49]. They argue, in contrast to David, that the QWERTY keyboard is *not* an example of the market inefficiently locking in an inferior design due to a few largely chance events a long time ago. They fault David’s analysis on three points. First, the evidence is very weak that QWERTY is inferior to other keyboards. If QWERTY is not inferior, its popularity as a standard introduces no economic inefficiency. Second, the initial competition among alternative keyboard layouts from which QWERTY emerged the victor was long and extensive. Therefore QWERTY provides no economic analogy to the proverbial flap of a remote butterfly’s wings that leads to a hurricane thousands of miles away. Third, it is only inefficient to stick with an inferior standard when the gains from switching to the new standard exceed the transition costs (e.g., retraining, retooling, or replacing people, procedures and technology). Liebowitz and Margolis argue, with a number of examples, that where a new standard’s proponents believe it will generate gains in excess of transition costs, they have a market incentive to cover all or part of the transition costs in order to profit from the remaining gains. Mechanisms for doing so include discounts for early adopters, guarantees of satisfaction, rental offers, rebates for those turning in equipment using the old standard, subsidised or free training with the new aspiring standard (as took place during the early competition among typewriter keyboards) and subsidised conversion of old technologies to the new standard (as was also done with typewriters). Liebowitz and Margolis direct particular scorn at David’s citation of studies alleging that the costs of retraining QWERTY typists on Dvorak keyboards would be paid back in ten days. Noting that this amounts to returns on the order of 2,200%, Liebowitz and Margolis consider the failure of any entrepreneur, in more than 60 years, to exploit such a huge profit opportunity as persuasive evidence that it does not exist.

Liebowitz and Margolis are concerned with network effects, i.e., situations where the value one consumer places on a product grows with the number of other people using the product. The more people who own telephones or fax machines, the more valuable a telephone or fax machine is to me. The more video rental stores and relatives who use VHS-standard equipment, the more valuable VHS-standard equipment is to me. David’s argument was that the QWERTY keyboard is an example of network effects locking in an inferior technology only because it got a bit of a head start on its rivals many years ago. Liebowitz and Margolis strongly and persuasively disagree for the reasons summarised above.

Path dependency and lock-in tendencies arising from learning have little to do directly with network effects. But the general observation above is still valid for learning effects. That is, while they create a tendency for the early technological leader to lock in

its initial advantage over time, such lock-in is hardly inevitable or permanent. Partly this is because the relative benefits and costs of a technology can change with time and its scale of use. For example, local pollution from internal combustion engines and even global pollution (from their CO₂ emissions) are now more serious concerns than when the automobile was first introduced. Government regulations have already internalised some of these pollution costs and may internalise more, thereby changing the internal combustion engine's relative costs and benefits among prospective competitors. Another reason learning does not make lock-in inevitable is that experience is not the only determinant of cost and performance. Competitors may prove to be better managed, more inventive, more responsive to consumers, or better able to undertake and exploit scientific research. And where there is a profit to be made by persuading consumers to switch products, entrepreneurs will use the sorts of transition cost-sharing mechanisms listed above to encourage customers to switch. However, other factors being equal, learning effects mean that the hurdle that competitors need to overcome tends to grow with time.

In weighing tendencies toward lock-in and their different degrees of reversibility, it is important to note that Liebowitz and Margolis (in their 1999 book) argue that software applications are 'instantly scalable,' i.e., production can expand essentially instantaneously to meet rapid demand growth. The ease of reversing lock-in is higher for such technologies. Even where energy technologies display network effects, most energy technologies do not have this instant scalability. And while some end-use technologies, like light bulbs and small motors, may cost less than software and be as easy to change, other energy technologies are much more expensive and long-lived. (Path dependence due to the durability of expensive long-lived acquisitions Liebowitz and Margolis label first-degree path dependence, with which they have no quarrel.) First-degree path dependence does not mean that competitors cannot overcome the learning effects and network effects that currently favour, for example, the internal combustion engine. But it does mean that the hurdle that needs to be overcome is relatively higher for long-lived energy technologies than for instantly scalable software applications.

We have spent a fair amount of time on lock-in because some expect that sufficiently sophisticated modelling of technological learning will lead to policy results with significant lock-in effects. That is, a small change in initial assumptions may give one technology a small head start that, thanks to learning effects, eventually locks the system in to that technology. Another small change in assumptions may give a different technology the initial head start and lead to a totally different end result, again because of learning and lock-in. If uncertain initial conditions are modelled stochastically, learning and lock-in could transform smooth probability distributions for uncertain input variables into highly non-smooth, unpredictable output distributions. Similarly, outcomes could be extremely sensitive to small policy interventions early in the planning horizon and quite insensitive to much larger interventions later.

We are so far unaware of research where smooth input distributions have produced non-smooth results due to lock-in driven by technological learning. But that may just be because the models are not yet sufficiently sophisticated. Our purpose here has been to anticipate the possibility of such results and caution that, should they materialise, they could well be overly sensitive to early uncertainties, overly reinforcing of early policy actions and overly unforgiving of early policy 'mistakes.' It will thus be important to include within our models various unlocking mechanisms and creative options for

sharing transition costs in ways that make reversing undesirable lock-in profitable for consumers and suppliers alike. Yes, early decisions are important, but their importance should not be overestimated. Next year's choices will not totally determine the future energy system. There will be good reasons and opportunities to continue to intervene creatively.

5 Conclusions and outlook

In this paper we have presented an overview of learning rates of energy-related technologies, mainly from the perspective of energy modelling in support of energy policy making. The importance of learning curves for policy making has been recognised by several research groups which, following an initiative by the International Energy Agency, have recently formed a 'virtual network' with the acronym EXCETP (*Experience Curves for Energy Technology Policymaking*). One visible output of this group is a system of WWW pages [50].

We believe that the incorporation of technological learning in energy modelling for policy making represents significant progress over previous analyses in which technological progress was an exogenous factor whose beneficial influence just unfolded automatically with time. Although future learning rates remain uncertain and unpredictable, simply recognising the concept should motivate policies to invest in new R&D as even cautious assumptions about future learning can make support for certain technologies worthwhile. Because such investments may only pay off in the long run, however, there is a particularly important role for governments with their responsibility for the larger long-term interests of society and (limited) experience in successfully managing trade-offs between short-term costs and longer-term benefits.

In terms of future research, the learning-curve formulation presented in equation 2.1 may be regarded as overly deterministic. With cumulative capacity as the only explanatory variable, the only policy options for accelerating cost or performance improvements involve speeding capacity growth – through subsidies or mandates to spur demand or through direct government procurement programs. Additional policy options are, however, suggested by Section 4's list of additional possible causes of cost reductions. One of the most obvious items on the list relevant to energy policy making is R&D. This has been the motivation for a new set of activities, partly supported by the European Community, which aims to quantify the impact of R&D support on technological progress in a fashion similar to what has been done for cumulative experience. This effort, labelled the SAPIENT project, began in March 2000 and is scheduled for two years. SAPIENT involves, as did its predecessor, the TEEM project [7], the joint research of the major European groups working on technological learning in long-term energy scenarios. SAPIENT goes beyond TEEM in also including Japanese and US groups.

Acknowledgments

The authors would like to thank an anonymous reviewer for very helpful comments and suggestions. The research reported here was generously supported in part by Japan's Central Research Institute of the Electric Power Industry (CRIEPI).

References and Notes

- 1 Nakićenović, N., Grübler, A. and McDonald, A. (Eds.) (1998) *Global Energy Perspectives*, Cambridge University Press, Cambridge, UK, New York, NY, USA, ISBN 0-521-64569-7, (http://www.iiasa.ac.at/cgi-bin/ecs/book_dyn/bookcnt.py).
- 2 von Weizsäcker, E., Lovins, A.B. and Lovins, L.H. (1997) *Factor Four: Doubling Wealth, Halving Resource Use*, Earthscan Publications Ltd., London.
- 3 Working Group III of the Intergovernmental Panel on Climate Change (IPCC) (2000) 'Summary for policymakers – special report on emission scenarios', based on a draft prepared by Nakićenović, N., Davidson, O., Davis, G., Grübler, A., Kram, T., La Rovere, E.L., Metz, B., Morita, T., Pepper, W., Pitcher, H., Sankovski, A., Shukla, P., Swart, R., Watson, R. and Dadi, Z. (http://www.ipcc.ch/pub/SPM_SRES.pdf).
- 4 Messner, S. (1997) 'Endogenized technological learning in an energy systems model', *Research Report RR 97-15*, IIASA, Laxenburg, Austria.
- 5 Mattsson, N. (1997) 'Internalizing technological development in energy systems models', *Thesis for the Degree of Licentiate of Engineering, ISRN CTH-EST-R97/3SE*, Energy Systems Technology Division, Chalmers University of Technology, Göteborg, Sweden.
- 6 NEMS (2000) *The National Energy Modelling System: An Overview 2000*, Energy Information Administration, Office of Integrated Analysis and Forecasting, US Department of Energy, Washington, DC 20585, also: www.eia.doe.gov/oiaf/aeo/overview/index.html
- 7 Capros, P. (Ed.) (2000) 'Energy technology dynamics and advanced energy system modelling', Report of the TEEM Project (TEEM: a programme partially funded by the European Commission, Directorate General Research), *International Journal of Global Energy Issues*, Vol. 14, Nos. 1–4.
- 8 Wright, T.P. (1936) *Journal of Aeronautical Science*, Vol. 3, p.122.
- 9 For some data sets, estimating a learning curve leads to values of b equal to or greater than zero. Thus costs stagnate or increase with cumulative experience. In these cases, the terms 'progress ratio' and 'learning rate' are still used, although they are no longer as intuitively descriptive.
- 10 Boston Consulting Group (BCG) (1968) *Perspectives on Experience*, Boston Consulting Group Inc.
- 11 Blackwood, D. (1997) 'Eastern trough area project' (ETAP), *Engineer*, June, p.4.
- 12 Fisher, J.C. (1974) *Energy Crises in Perspective*, Wiley, New York.
- 13 Zhao, J. (2000) *Diffusion, Costs and Learning in the Development of International Gas Transmission Lines*, IR-00-054, IIASA, Laxenburg, Austria.
- 14 MacGregor, P.R., Maslak, C.E. and Stoll, H.G. (1991) *The Market Outlook for Integrated Gasification Combined Cycle Technology*, General Electric Company, New York.
- 15 Kouvaritakis, N., Soria, A. and Isoard, S. (2000) 'Modelling energy technology dynamics: methodology for adaptive expectations models with learning by doing and learning by searching', *International Journal of Global Energy Issues*, Vol. 14, Nos. 1-4, pp.104–115.
- 16 IEA (2000) *Experience Curves for Energy Technology Policy*, International Energy Agency, Paris, France (ISBN 92-64-17650-0).
- 17 Joskow, P.L. and Rose, N.L. (1985) 'The effects of technological change, experience and environmental regulation on the construction cost of coal-burning generating units', *Rand J. Economics*, Vol. 16, No.1.
- 18 Claeson, U. (1999) 'Using the experience curve to analyze the cost development of the combined cycle gas turbine', in: *Proceedings of the IEA International Workshop, Experience Curves for Policy Making: The Case of Energy Technologies*, Stuttgart, Germany, 10-11 May.
- 19 Durstewitz, M. (1999) 'Using information of Germanys '250 MW Wind' – programme for the construction of wind power experience curves', in: *Proceedings of the IEA International*

- Workshop, Experience Curves for Policy Making: The Case of Energy Technologies*, Stuttgart, Germany, 10-11 May.
- 20 Neij, L. (1999) 'Cost dynamics of wind power', *Energy – The International Journal*, Vol. 24, pp.375–389.
 - 21 CEC (California Energy Commission) (1997) *Wind Project Performance*, 1995 Summary Staff Report.
 - 22 Loiter, J.M. and Norberg-Bohm, V. (1999) *Energy Policy*, Vol. 27, p.85.
 - 23 Harmon, C. (2000) *Experience Curves of Photovoltaic Technology*, IR 00-014, IIASA, Laxenburg, Austria.
 - 24 Maycock, P.D. and Wakefield, G. F. (1975) *Business Analysis of Solar Photovoltaic Energy Conversion*, Texas Instruments Incorporated, Dallas, TX.
 - 25 Rabitsch, H. (2000) 'Electricity transmission networks in Eurasia', unpublished data.
 - 26 Goldemberg, J. (1996) 'The evolution of ethanol costs in Brazil', *Energy Policy*, Vol. 24, p.1127.
 - 27 Lipman, T.E. and Sperling, D. (1999) 'Forecasting the costs of automotive PEM fuel cell systems using bounded manufacturing progress functions', in: *Proceedings of the IEA International Workshop, Experience Curves for Policy Making: the Case of Energy Technologies*, Stuttgart, Germany, 10-11 May.
 - 28 Abernathy, W.J. and Wayne, K. (1974) *Harvard Business Review*, Vol. 52, p.109.
 - 29 Iwafune, Y. (2000) *IR-00-009*, IIASA, Laxenburg, Austria.
 - 30 Akisawa, A. (2000) 'Technological development and market substitution of heat pumps – a case of room air-conditioners in Japan,' unpublished data.
 - 31 McDonald, A. and Schrattenholzer, L. (2001) 'Learning rates for energy technologies', *Energy Policy*, Vol. 29, pp.255–261.
 - 32 Dutton, J.M. and Thomas, A. (1984) 'Treating progress functions as a managerial opportunity', *Academy of Management Review*, Vol. 9, p.235.
 - 33 Estimated learning curves are often presented with cumulative experience plotted on a logarithmic scale, or with both parameters plotted logarithmically. For the presentation of our own analysis, we prefer to use absolute scales unless there is a specific point to be made that is best illustrated by using logarithms.
 - 34 If one were to make the assumption that learning rates are 'intrinsic' to a technology and thus identical in different geographical areas, this difference of geographical scope could still lead to unbiased learning rates – if the two time series are synchronized.
 - 35 Ayres, R., Hourcade, J.C. and Hélioui, K. (1998) 'Expected diffusion of photovoltaic systems in southern Europe, Integrating technology diffusion micro models for assessing sustainable development policy options: Innovation and economic conditions for a climate double-dividend of policy in Europe', prepared by the Commission of the European Communities Directorate-General for Science, Research and Development, *Project #: ENV4-CT96-0292*, October.
 - 36 The first four data points in Figure 9, which cover the period from 1972 to 1985, do not correspond to annual intervals. From then onwards, the data correspond to annual time series.
 - 37 Argote, L. (1999) *Organizational Learning: Creating, Retaining and Transferring Knowledge*, Kluwer, Dordrecht, Netherlands.
 - 38 Easterling, R.G. (1982) 'Statistical analysis of US power plant capacity factors through 1979,' *Energy*, Vol. 7, No. 3, pp.253–258.
 - 39 Joskow, P.L. and Rozanski, G.A. (1979) 'The effects of learning by doing on nuclear plant operating reliability', *The Review of Economics and Statistics*, May, Vol. LXI, No. 2, pp.161–168.
 - 40 Krautmann, A.C. and Solow, J.L. (1988) 'Economies of scale in nuclear power generation', *Southern Economic Journal*, Vol. 55, pp.70–85.

- 41 Zimmerman, M.B. (1982) 'Learning effects and the commercialization of new energy technologies: the case of nuclear power,' *The Bell Journal of Economics*, Autumn, Vol. 13, No. 2.
- 42 Marchetti, C. (1986) 'Stable rules in social behavior', *IBM Conference*, Academia Brasileira de Ciências, Rio de Janeiro, 11 November.
- 43 Nakićenović, N. (1984) 'Growth to limits: long waves and the dynamics of technology', *International Institute for Applied Systems Analysis*, Laxenburg, Austria.
- 44 Grübler, A. (1990) *The Rise and Fall of Infrastructures. Dynamics of Evolution and Technological Change in Transport*, Physica-Verlag, Heidelberg, Germany.
- 45 David, P.A. (1985) 'Clio and the economics of QWERTY,' *American Economic Review*, Vol. 75, pp.332-337.
- 46 David, P.A. (1986) 'Understanding the economics of QWERTY: the necessity of history,' in W.N. Parker (Ed.), *Economic History and the Modern Economist*, Basil Blackwell, New York.
- 47 David, P.A. (1992) 'Heroes, herds and hysteresis in technological history: 'the battle of the systems' reconsidered', *Industrial and Corporate Change*, Vol. 1, pp.129-180.
- 48 Liebowitz, S.J. and Margolis, S.E. (1990) 'The fable of the keys,' *Journal of Law and Economics*, April, Vol. 33, No. 1, pp.1-25.
- 49 Liebowitz, S.J. and Margolis, S.E. (1999) *Winners, Losers & Microsoft*, The Independent Institute, Oakland, California.
- 50 The URL of the EXCETP home page is: <http://www.iiasa.ac.at/Research/ECS/docs/EXCETP/exet-home.html>