

**THE PERILS OF THE LEARNING MODEL FOR
MODELING ENDOGENOUS TECHNOLOGICAL CHANGE**

By

William D. Nordhaus

January 2009

COWLES FOUNDATION DISCUSSION PAPER NO. 1685



**COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281**

<http://cowles.econ.yale.edu/>

The Perils of the Learning Model

For Modeling Endogenous Technological Change

William D. Nordhaus¹

Yale University

December 15, 2008

Abstract

Learning or experience curves are widely used to estimate cost functions in manufacturing modeling. They have recently been introduced in policy models of energy and global warming economics to make the process of technological change endogenous. It is not widely appreciated that this is a dangerous modeling strategy. The present note has three points. First, it shows that there is a fundamental statistical identification problem in trying to separate learning from exogenous technological change and that the estimated learning coefficient will generally be biased upwards. Second, we present two empirical tests that illustrate the potential bias in practice and show that learning parameters are not robust to alternative specifications. Finally, we show that an overestimate of the learning coefficient will provide incorrect estimates of the total marginal cost of output and will therefore bias optimization models to tilt toward technologies that are incorrectly specified as having high learning coefficients.

¹ This study expands on earlier presentations at the Energy Modeling Forum, the Santa Fe Institute, and the National Academies. I am grateful for comments from Christopher Magee, Nebosja Nakicenovic, and John Weyant. [Version is LearningPerils_v8.doc.]

I. The Fundamental Identification Problem

Models of learning and experience have a long history in studies of manufacturing productivity.² Because of their perceived successes in technological forecasting, they have recently been introduced in policy models of energy and global warming economics to make the process of technological change endogenous. This approach has serious dangers. The present note has three points. First, it shows that there is a fundamental statistical identification problem in trying to separate learning from exogenous technological change and that the estimated learning coefficient will generally be biased upwards. Second, we present two empirical tests that illustrate the potential bias in practice and show that learning parameters are not robust to alternative specifications. Finally, we show that an overestimate of the learning coefficient will generally underestimate the total marginal cost of output; because of this underestimate, optimization models tend to tilt toward technologies that are incorrectly specified as having high learning coefficients.

² The literature on learning curves is vast, going back more than a century, and no single reference can adequately capture the major issues. The original concept of an experience curve was documented with telegraph operators in W.L. Bryan and N. Harter, "Studies on the Telegraphic Language: The Acquisition of a Hierarchy of Habits," *Psychology Review*, 6:345-75, 1899. Two particularly influential articles were T.P. Wright, "Factors Affecting the Cost of Airplanes," *Journal of Aeronautical Sciences*, Vol. 3, No. 4, 122-128, 1936; and K. J. Arrow, "The Economic Implications of Learning-By-Doing," *Review of Economic Studies*, Vol. 29, 155-173, 1961. A recent comparison of alternative approaches is in Boyan Jovanovic and Yaw Nyarko, "A Bayesian Learning Model Fitted to a Variety of Empirical Learning Curves," *Brookings Papers on Economic Activity. Microeconomics*, (1995), pp. 247-305. A comprehensive survey of learning curves is contained in Louis E. Yelle, "The Learning Curve: Historical Review and Comprehensive Survey," *Decision Sciences*, 10, 302-328, 1979).

The basic idea is that productivity improves or costs decline as workers or firms gain experience with a production process. While there can be little doubt that productivity benefits from experience, the exact mechanism is poorly understood. In particular, it is unclear whether the learning is embodied in individual workers and firms, whether there are interindustry or international spillovers, and whether the improvements lead to durable technological changes, and even whether the learning effects can be distinguished from other technological changes.

In this section, we focus on the problem of identifying differences in productivity due to learning from exogenous changes. We begin by showing why it is impossible without further identifying assumptions to distinguish learning from exogenous technological change, and why the learning coefficient is generally biased upwards. To simplify for this exposition, we assume that all processes are exponential. Output (Q_t) is assumed to grow at constant growth rate g , so $Q_t = Q_0 e^{gt}$. Cumulative output at time t (Y_t) is therefore:

$$(1) \quad Y_t = \int_{v=-\infty}^t Q_0 e^{gv} dv = Q_0 e^{gt} / g$$

Taking the logarithmic derivative of (1) shows that the growth rate of Y_t is g .

The experience curve is assumed to have a true experience coefficient, b . In addition, there is an assumed constant rate of exogenous technological change at rate h . The cost function is therefore:

$$(2) \quad C_t = C_0 e^{-ht} Y_t^{-b}$$

“Exogenous technological change” in this context denotes all sources of cost declines other than the learning-curve-determined technological change. It would include *inter alia* spillovers from outside the industry, the returns to research and development, economies of scale and scope, as well as exogenous fundamental inventions.

Assume that prices are proportional to current instantaneous marginal cost, so the rate of decline in cost (c_t) equals the decline in price (p_t), which is given by:

$$(3) \quad p_t = c_t = h + bg_t$$

Because marginal cost is constant, price is under these assumptions exogenous to current demand. Demand is determined by a demand function with constant price elasticity (ε), elasticity of per capita demand with respect to aggregate output (λ), where the growth in aggregate per capita output is w_t , and constant population growth is n . These yield the growth in output (demand) as:

$$(4) \quad g_t = \varepsilon p_t + \lambda w_t + n$$

We can substitute $z_t = \lambda w_t + n =$ the autonomous (non-price-induced) growth rate of demand. Solving (3) and (4), we get the following reduced-form equations for the rate of cost (price) decline and the rate of output growth. Since the growth rates are constant, we suppress the time subscripts. Price decline is:

$$p = h + bg = h + b(\varepsilon p + z)$$

or

$$(5) \quad p = \frac{h + bz}{1 - b\varepsilon}$$

which is the equation for cost (price). The equation for output growth is:

$$g = \varepsilon(h + bg) + z$$

or

$$(6) \quad g = \frac{\varepsilon h + z}{1 - b\varepsilon}$$

From equations (5) and (6) we can calculate the slope of the *behavioral learning curve*, β . This is equal to p/g , or:

$$(7) \quad \beta = p / g = \frac{h + bz}{\varepsilon h + z} = b(1 + \phi)$$

where

$$(8) \quad \phi = \frac{h(1 - \varepsilon b)}{b(\varepsilon h + z)}$$

Note how difficult it would be to separate the true learning parameter (b) from the tangle of coefficients in (7).³ To obtain the true learning parameter, we

³ E. R. Berndt, *The Practice of Econometrics: Classic and Contemporary*, Addison-Wesley, New York, 1991 is a useful survey of the literature and econometric issues. Zvi Griliches, "R&D and Productivity: Econometric Results and Measurement Issues," in P. Stoneman (Ed.), *Handbook of Economics on Innovation and Technological Change*, Blackwell, Oxford, 1995 reviews the more general questions of the econometrics of technological change.

would need to have reliable estimates of the rate of exogenous technological change, the demand elasticity, and the rate of autonomous growth of demand.

Furthermore, the empirical experience parameter will be unbiased ($\beta = b$) only when exogenous technological change is zero ($h = 0$). It will be biased if $h \neq 0$. The size of the bias is determined by the sign of $(1 - \varepsilon b)$. For demand elasticities that are relatively low (less than 4), we would expect that $\varepsilon b > 1$, in which case the bias is upwards.

A numerical example will illustrate the result. A typical industry might have a price elasticity of $\varepsilon = 1$, an exogenous demand growth of $z = 0.04$ per year, and a rate of exogenous technological change of 0.01 per year. With a zero learning effect (true $b = 0$), we have from (7) that:

$$\beta = \frac{0.01 + 0 \times 0.04}{1 \times 0.01 + .04} = \frac{0.01}{0.05} = 0.2$$

In this case, therefore, the empirical learning coefficient is 0.2 even though the actual learning coefficient is zero.

Next, change the assumption to add a true experience curve coefficient of $b = 0.25$. Then the empirical learning coefficient becomes:

$$\beta = \frac{0.01 + 0.25 \times 0.04}{1 \times 0.01 + .04} = \frac{0.02}{0.05} = 0.4$$

In this case, the empirical experience coefficient is biased upwards from 0.25 to 0.4.

The general conclusion is that because of the interaction of demand, output growth, exogenous technological change, and learning, behavioral learning curves will generally have an upward bias in estimated learning coefficients. The only general case in which the coefficient is unbiased is when exogenous (non-learning) technological change is zero.

II. Some Empirical Tests

We can take actual data on output and productivity to show the difficulty in measuring learning from historical data. For a first example, we take U.S. data on multifactor productivity in the non-farm business sector for the period 1948-2007. While this is a highly aggregated data set, it has very high quality input and output data that are constructed by U.S. statistical agencies. The data are available at www.bls.gov.

I construct cumulative output for the period 1800 to 2007. These are based on government data for 1929 to 2007 and extrapolate backwards from 1929 using an assumed constant growth rate of 3.9 percent per year. I then estimate bivariate learning functions and exogenous technological change rates. The estimated rate of exogenous technological change for the 1948-2007 period is $0.0105 (\pm .00038)$. A learning equation without exogenous technological change has a learning coefficient of $0.278 (\pm 0.0093)$. If we combine the two variables, learning and time, the learning coefficient rises to the implausibly high level of $2.1491 (\pm 0.27)$. This simple example shows that learning models generally have coefficients that are not robust to specification changes.

As a second example, we estimate learning parameters for 34 major industry groups.⁴ For this purpose, we use data on industry output from the U.S. national income and product accounts. We select only those industries where the output and prices are “well measured.”⁵

The basic approach is to assume that output is produced by a constant returns to scale production function either with or without learning, as above. Average hourly earnings (AHE) are assumed to represent a reasonable proxy for the cost of production with unchanging technology.⁶ Prices are proportional to average instantaneous costs, and learning is assumed to be excluded from

⁴ These are Accommodation and food services; Administrative and waste management services; Air transportation; Apparel and leather and allied products; Broadcasting and telecommunications; Chemical products; Computer and electronic products; Computer systems design and related services; Electrical equipment, appliances, and components; Fabricated metal products; Farms; Food and beverage and tobacco products; Food services and drinking places; Forestry, fishing, and related activities; Furniture and related products; Information and data processing services; Machinery; Motion picture and sound recording industries; Motor vehicles, bodies and trailers, and parts; Nonmetallic mineral products; Oil and gas extraction; Paper products; Petroleum and coal products; Plastics and rubber products; Primary metals; Printing and related support activities; Publishing Industries (includes software); Rail transportation; Retail trade; Securities, commodity contracts, and investments; Textile mills and textile product mills; Transportation and warehousing; Truck Transportation; Utilities; Wholesale trade; and Wood products.

⁵ The notion of well-measured industries is discussed in William Nordhaus, “Baumol's Diseases: A Macroeconomic Perspective,” *Berkeley Journal of Macroeconomics*, Volume 8, Issue 1, 2008, Article 9, pp. 1-37.

⁶ The reference on the use of wages as a proxy for technological change is analyzed in William Nordhaus, “Baumol's Diseases: A Macroeconomic Perspective,” *Berkeley Journal of Macroeconomics*, Volume 8, Issue 1, 2008, Article 9, pp. 1-37.

pricing, or if included to be a multiplicative factor.⁷ Under these assumptions, the rate of change in the ratio of the AHE to the product price (call this the real price decline) will equal the rate of cost decline given by equation (2). We can then examine the relationship between the real price decline and factors such as learning as represented by cumulative output or time.

The data on output and inputs have been prepared by the BEA and provide both gross output and price indexes for the period 1947-2007 for major industry groups and for 1959-2007 for all other industry groups. For each group, we estimate cumulative output for the initial year as the first-year output index divided by the growth rate for the first seven years. The BEA data are available at <http://www.bea.gov/industry/index.htm#annual>.

For these data, two results are clear. First, the learning coefficients are wildly variable. Figure 1 shows a histogram of the empirical learning coefficients estimated from the linearized version of equation (2) with zero exogenous technological change. Clearly, the estimates are highly variable. For the 34 industries, assuming zero exogenous technological change, only 4 have estimated empirical learning coefficients in the plausible range between 0 and 0.5. The mean and median are well above the plausible range, which suggests a positive bias in the estimation. Additionally, the results are highly sensitive to an AR1 correction, which produce yet different learning coefficients.

⁷ Prices might reflect life-cycle costs if the firm believes that there are learning effects (see for example C. Lanier Benkard, "Learning and Forgetting: The Dynamics of Aircraft Production," *American Economic Review*, 2000, vol. 90, pp. 1034-1054). However, as long as the processes are exponential and the learning coefficient is constant, the effect of learning on the ratio of marginal cost to price would be invariant over time.

Figure 2 shows a scatter plot of the learning coefficients without and with exogenous technological change [i.e., with h equal to zero or with h as estimated in the linearized version of equation (2)]. The estimates with exogenous technological change are very different from those without exogenous technological change. The correlation between the two estimates is 0.009. Clearly, the estimates are highly sensitive to the specification.

These examples show that the estimates of the learning coefficients are not robust to specifications. Moreover, the estimates are often well outside the theoretically acceptable range.⁸

III. The perils of learning in optimization models

Learning has become a favorite tool for modeling technological change in many models of the energy sector and of global warming. It is convenient because learning-by-doing is one of the few “theories” of technological change that is easily included in models because of its simple specification.⁹ It is a

⁸ As a technical note, readers might wonder about whether it is appropriate to use such aggregated data to estimate learning equations. One advantage of the industry data is that we have very carefully prepared indexes of output and price, so the measures are close to the ideal. The appendix to this study shows the conditions under which elemental processes can be aggregated to determine an aggregate learning function. The results indicate that there are three sources of potential bias. The first bias would arise if the growth rate in individual industries deviated from exponential growth; the second bias relates to the correlation between the learning coefficient and output growth; and the third bias comes from potential biases in productivity measurement.

⁹ Examples of studies that use learning in energy and global warming models are T. Barker, H. Pan, J. Köhler, R. Warren, and S. Winne, “Decarbonizing the Global Economy

dangerous modeling technique, however, because the estimated learning rates are biased upwards and because these approaches therefore seriously underestimate the marginal cost of output. We showed the first point above and address the second in this section.

The danger in using learning to model exogenous technological change arises when the models select technologies on the basis of their cost characteristics. Learning models have total marginal costs that are lower than current marginal costs because an additional unit of output lowers all future costs as producers move down the learning curve. We can see this point by starting with a total cost function, which is defined as the present value of all current and future production (V_t), given by the following:¹⁰

$$(9) \quad V_t = \int_{v=t}^{\infty} Q_v C_v e^{-r(v-t)} dv = \int_{v=t}^{\infty} Q_v [C_0 e^{-hv} Y_v^{-b}] e^{-r(v-t)} dv$$

with Induced Technological Change: Scenarios to 2100 using E3MG," *The Energy Journal Special Issue, Endogenous Technological Change and the Economics of Atmospheric Stabilization*, 241-258, 2006; and M. Grubb, J. Köhler, and D. Anderson, "Induced Technical Change in Energy and Environmental Modeling: Analytic Approaches and Policy Implications," *Annual Review of Energy and Environment*, Vol. 27, 271-308, 2002.

¹⁰ The formula for marginal cost was developed in the context of industrial organization in Pankaj Ghemawat and A. Michael Spence, "Learning Curve Spillovers and Market Performance," *The Quarterly Journal of Economics*, Vol. 100, Supplement (1985), pp. 839-852 and Saman Majd and Robert S. Pindyck, "The Learning Curve and Optimal Production under Uncertainty," *The RAND Journal of Economics*, Vol. 20, No. 3 (Autumn, 1989), pp. 331-343. The marginal cost is empirically estimated and used in the context of a dynamic programming formulation in C. Lanier Benkard, "Learning and Forgetting: The Dynamics of Aircraft Production," *American Economic Review*, 2000, vol. 90, pp. 1034-1054.

Taking the derivative of (9) with respect to current output yields total marginal cost:

$$(10) \quad \partial V_t / \partial Q_t = C_t - b \int_{v=t}^{\infty} Q_v C_0 e^{-hv} Y_v^{-b-1} (\partial Y_v / \partial Q_t) e^{-r(v-t)} dv$$

Equation (10) has two terms. The first term is just the marginal cost (equal to the average cost) of output for the current (instantaneous) period. The second term is the learning impact. It shows the impact on cost through the effect of current output on future cumulative output. As long as $b > 0$, this learning impact is negative. The size of the negative term depends upon the learning parameter, the discount rate, and the exogenous growth term.

There is no obvious analytical expression that can be derived from equation (10). However, we can calculate the expression numerically for different parameters. Table 1 shows the calculation of the true total marginal cost as calculated by equation (10) for different parameters using a 200-year horizon. We have normalized the equation so that the instantaneous marginal cost is unity, and this is clearly the total marginal cost when there is no learning.

The table shows how the total marginal cost declines sharply as the assumed learning rate increases. The decline increases in the learning coefficient and the growth rate and decreases in the discount rate, although the growth rate does not affect the bias in a major way. For example, assume a conventional learning coefficient is 0.2 and a discount rate of 5 percent per year. If the true process is exogenous technological change, then this would bias downward the

total marginal cost by between 39 and 51 percent for the two assumed growth rates. With low discount rates, the bias is even larger.

This bias becomes particularly important in energy and global warming models which are designed to choose among different emerging technologies and where the technology is assumed to have an important learning component. For example, the model solve for future paths of solar and wind technologies based on current cost and different learning coefficients. Based on high learning rates, the model might suggest that technology A is a good bet for research and development. But this recommendation would be incorrect if the learning coefficient is based on a biased estimate of learning.

The point to emphasize here is that, in analyses that pick technologies on the basis of total discounted cost of production (as is entirely appropriate), then an upward bias in the learning rate will have a major impact on the apparent benefit of technologies with learning. The estimated costs can easily be underestimated by a factor of two. This danger is reinforced because, as shown in the first section, of the tendency to estimate learning rates in bivariate relationships, which will generally lead to strong upward biases in the learning coefficient.

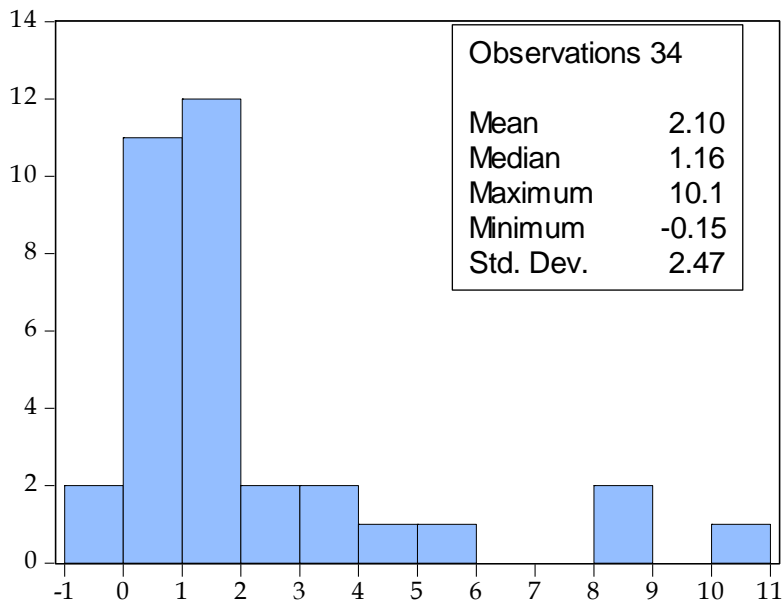


Figure 1. Estimated learning coefficients for 34 major industries, with zero exogenous technological change, 1947 – 2007 or subperiod

The histogram shows the estimated learning coefficients for 34 major industries. Most are well outside the usual assumed range of between 0 and ½. For a description of the methods, see the text.

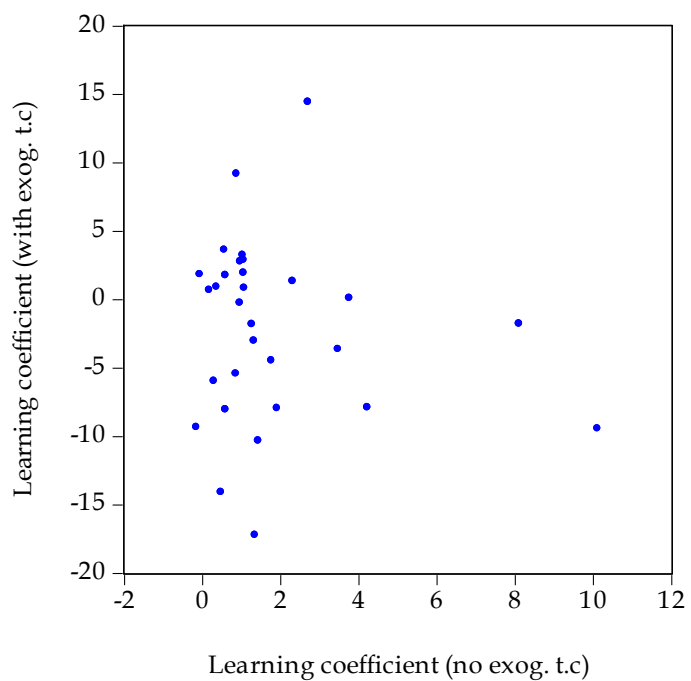


Figure 2. Comparison of estimated learning coefficient with and without exogenous technological change

Note that the graph has truncated the plot to the limits shown. 4 observations lie outside that range. The correlation coefficient between the two sets of estimated coefficients is 0.009.

Growth rate = 3 percent per year

Learning coefficient	Discount rate			
	0.01	0.03	0.05	0.1
0	1.000	1.000	1.000	1.000
0.1	0.554	0.638	0.700	0.698
0.2	0.317	0.410	0.490	0.487
0.3	0.186	0.265	0.343	0.340
0.5	0.069	0.113	0.168	0.166

Growth rate = 10 percent per year

Learning coefficient	Discount rate			
	0.01	0.03	0.05	0.1
0	1.000	1.000	1.000	1.000
0.1	0.420	0.597	0.780	0.774
0.2	0.221	0.377	0.608	0.599
0.3	0.129	0.246	0.475	0.463
0.5	0.050	0.111	0.289	0.277

Table 1. Total marginal cost with learning as function of growth, discount rate, and learning coefficient

This table shows how the assumed learning curve affects the total marginal cost as calculated by equation (10), conditional on discount rates and growth rates. To see the impact of different assumed learning rates, read down the table. For example, with a discount rate of 0.05 per year and a growth rate of 0.10 per year, the total marginal cost is 1.0 with no learning but 0.608 with a true learning coefficient of 0.2.

Appendix. Aggregation of Learning Curves

Empirical studies of learning or experience curves often aggregate from elemental processes. For example, an experience curve for photovoltaic cells might aggregate over plants, firms, and countries. The results from this study are for major industry groups. An important question analyzed in this appendix involves the conditions under which elemental learning processes can be aggregated.

The aggregation of process learning curves proceeds as follows. It will be convenient to work with productivity growth rather than cost decline because output and inputs measurements are more precisely defined. The growth in aggregate productivity is

$$(A.1) \quad \dot{A}_t / A_t = \dot{a}_t = \dot{Q}_t / Q_t - \dot{M}_t / M_t = \dot{q}_t - \dot{m}_t$$

A_t is productivity, Q_t is a Tornqvist measure of output, M_t is inputs, and the lower-case letters are rates of growth. We can rewrite (A.1) as the sum of the individual processes:

$$(A.2) \quad \dot{a}_t = \sum_{i=1}^n \dot{q}_{i,t} \theta_{i,t}^Z - \sum_{i=1}^n \dot{m}_{i,t} \theta_{i,t}^M$$

where $\dot{q}_{i,t}$ and $\dot{m}_{i,t}$ are the growth rates of $Q_{i,t}$ and $M_{i,t}$, and $\theta_{i,t}^Z$ and $\theta_{i,t}^M$ are the shares in nominal output and physical inputs for each industry $i = 1, \dots, n$.

Next, consider the rate of change of the learning function in equation (2):

$$(A.3) \quad \dot{a}_{i,t} = h_{i,t} - b_i \dot{y}_{i,t}$$

For this analysis, we simplify by assuming that each process has a constant growth rate.

This implies that output growth equals the growth in cumulative output, so $\dot{y}_{i,t} = \dot{q}_{i,t}$.

Calculate an aggregate learning curve by weighting the individual productivity growths in (A.3) by the shares of nominal output:

$$(A.4) \quad \sum_{i=1}^n \dot{a}_{i,t} \theta_{i,t}^Z = \sum_{i=1}^n h_{i,t} \theta_{i,t}^Z - \sum_{i=1}^n b_i \dot{q}_{i,t} \theta_{i,t}^Z$$

The left hand side of (A.4) is:

$$(A.5) \quad \begin{aligned} lhs(A.4) &= \sum_{i=1}^n \dot{a}_{i,t} \theta_{i,t}^Z = \sum_{i=1}^n [\dot{q}_{i,t} - \dot{m}_{i,t}] \theta_{i,t}^Z = \sum_{i=1}^n \dot{q}_{i,t} \theta_{i,t}^Z - \dot{m}_{i,t} \theta_{i,t}^M + \dot{m}_{i,t} \theta_{i,t}^M - \dot{m}_{i,t} \theta_{i,t}^Z \\ &= \dot{a}_t + \sum_{i=1}^n \dot{m}_{i,t} [\theta_{i,t}^M - \theta_{i,t}^Z] \end{aligned}$$

This is aggregate productivity plus a bias term that weights the growth rates of inputs by the difference between the shares of physical inputs and the shares of nominal outputs. Where outputs and inputs are correctly measured and under conditions of competition and constant returns, that bias term is zero.

After adding and subtracting a term, the right hand side of (A.4) is:

$$(A.6) \quad rhs(A.4) = \sum_{i=1}^n h_{i,t} \theta_{i,t}^Z - \sum_{i=1}^n \bar{b} \dot{q}_{i,t} \theta_{i,t}^Z + \sum_{i=1}^n \bar{b} \dot{q}_{i,t} \theta_{i,t}^Z - \sum_{i=1}^n b_i \dot{q}_{i,t} \theta_{i,t}^Z = \bar{h}_t - \bar{b} \dot{q}_t + \sum_{i=1}^n [\bar{b} - b_i] \dot{q}_{i,t} \theta_{i,t}^Z$$

The first term of (A.6) is the weighted average rate of exogenous technological change, using nominal output weights. The second term is the average learning rate times the aggregate growth rate of output, where the average learning rate, \bar{b} , is the weighted average learning rate using nominal output weights. The third term is a bias term equal to the deviation of the learning rate from the average times the weighted growth rate.

Putting all these together, we get the aggregate learning curve:

$$(A.7) \quad \dot{a}_t = \bar{h}_t - \bar{b} \dot{q}_t + \sum_{i=1}^n [\bar{b} - b_i] \dot{q}_{i,t} \theta_{i,t}^Z - \sum_{i=1}^n \dot{m}_{i,t} [\theta_{i,t}^M - \theta_{i,t}^Z]$$

This derivation shows that the learning curve can be aggregated from the individual process to the aggregate level subject to three potential errors. The first arises if the growth rate of output is not constant (which is not shown in the equation). The second is the bias term in equation (A.7) relating to the correlation between the learning coefficient, the rate of growth of output, and the output weights. As long as there is little correlation between the learning rate and the rate of growth of output, then the second bias term will be zero. The third bias comes from productivity measurement and

relates to the relation between input growth and the shares of inputs and nominal output.

As long as these three biases are small, the aggregate learning coefficient will be the weighted average of the individual processes where the weights are those of nominal output.