ARTICLE IN PRESS

Technological Forecasting & Social Change xxx (xxxx) xxx–xxx

Contents lists available at ScienceDirect

Technological Forecasting & Social Change

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

How well do experience curves predict technological progress? A method for making distributional forecasts

François Lafond[a,b,c,*], Aimee Gotway Bailey[d], Jan David Bakker[e], Dylan Rebois[b], Rubina Zadourian[f,a,g], Patrick McSharry[b,h,i,j], J. Doyne Farmer[a,f,k,l]

[a] Institute for New Economic Thinking, Oxford Martin School, UK
[b] Smith School for Enterprise and the Environment, University of Oxford, UK
[c] London Institute for Mathematical Sciences, UK
[d] U.S. Department of Energy, USA
[e] Department of Economics, University of Oxford, UK
[f] Mathematical Institute, University of Oxford, UK
[g] Max Planck Institute for the Physics of Complex Systems, Germany
[h] Carnegie Mellon University Africa, Rwanda
[i] African Center of Excellence in Data Science, University of Rwanda, Rwanda
[j] Oxford Man Institute of Quantitative Finance, University of Oxford, UK
[k] Computer Science Department, University of Oxford, UK
[l] Santa Fe Institute, USA

ARTICLE INFO

JEL classification:
C53
O30
Q47

Keywords:
Forecasting
Technological progress
Experience curves

ABSTRACT

Experience curves are widely used to predict the cost benefits of increasing the deployment of a technology. But how good are such forecasts? Can one predict their accuracy a priori? In this paper we answer these questions by developing a method to make distributional forecasts for experience curves. We test our method using a dataset with proxies for cost and experience for 51 products and technologies and show that it works reasonably well. The framework that we develop helps clarify why the experience curve method often gives similar results to simply assuming that costs decrease exponentially. To illustrate our method we make a distributional forecast for prices of solar photovoltaic modules.

1. Introduction

Since Wright's (1936) study of airplanes, it has been observed that for many products and technologies the unit cost of production tends to decrease by a constant factor every time cumulative production doubles (Thompson, 2012). This relationship, also called the experience or learning curve, has been studied in many domains.[1] It is often argued that it can be useful for forecasting and planning the deployment of a particular technology (Ayres, 1969; Sahal, 1979; Martino, 1993). However in practice experience curves are typically used to make point forecasts, neglecting prediction uncertainty. Our central result in this paper is a method for making *distributional forecasts*, that explicitly take prediction uncertainty into account. We use historical data to test this and demonstrate that the method works reasonably well.

Forecasts with experience curves are usually made by regressing historical costs on cumulative production. In this paper we recast the experience curve as a time series model expressed in first-differences: the change in costs is determined by the change in experience. We derive a formula for how the accuracy of prediction varies as a function of the time horizon for the forecast, the number of data points the forecast is based on, and the volatility of the time series. We are thus able to make distributional rather than point forecasts. Our approach builds on earlier work by Farmer and Lafond (2016) that showed how to do this for univariate forecasting based on a generalization of Moore's law (the autocorrelated geometric random walk with drift). Here we apply our new method based on experience curves to solar photovoltaic modules (PV) and compare to the univariate model.

Other than Farmer and Lafond (2016), the two closest papers to our contribution here are Alberth (2008) and Nagy et al. (2013). Both papers tested the forecast accuracy of the experience curve model, and performed

* Corresponding author.
E-mail addresses: francois.lafond@inet.ox.ac.uk (F. Lafond), doyne.farmer@inet.ox.ac.uk (J.D. Farmer).

[1] See Anzanello and Fogliatto (2011), Dutton and Thomas (1984), Yelle (1979) and for energy technologies Candelise et al. (2013), Isoard and Soria (2001), Junginger et al. (2010), Kahouli-Brahmi (2009), Neij (1997), Nemet (2006).

https://doi.org/10.1016/j.techfore.2017.11.001
Received 17 March 2017; Received in revised form 28 October 2017; Accepted 1 November 2017
0040-1625/

Please cite this article as: Lafond, F., Technological Forecasting & Social Change (2017), https://doi.org/10.1016/j.techfore.2017.11.001

comparisons with the time trend model. Alberth (2008) performed forecast evaluation by keeping some of the available data for comparing forecasts with actual realized values.[2] Here, we build on the methodology developed by Nagy et al. (2013) and Farmer and Lafond (2016), which consists of performing systematic hindcasting. That is, we use an estimation window of a constant (small) size and perform as many forecasts as possible. As in Alberth (2008) and Nagy et al. (2013), we use several datasets and we pool forecast errors to construct a distribution of forecast errors. We think that out-of-sample forecasts are indeed good tests for models that aim at predicting technological progress. However, when a forecast error is observed, it is generally not clear whether it is "large" or "small", from a statistical point of view. And it is not clear that it makes sense to aggregate forecast errors from technologies that are more or less volatile and have high or low learning rates.

A distinctive feature of our work is that we actually calculate the expected forecast errors. As in Farmer and Lafond (2016), we derive an approximate formula for the theoretical variance of the forecast errors, so that forecast errors from different technologies can be normalized, and thus aggregated in a theoretically grounded way. As a result, we can check whether our empirical forecast errors are in line with the model. We show how in our model forecast errors depend on future random shocks, but also parameter uncertainty, as is only seldomly acknowledged in the literature (for exceptions, see Vigil and Sarper, 1994 and Van Sark, 2008).

Alberth (2008) and Nagy et al. (2013) compared the forecasts from the experience curve, which we call Wright's law, with those from a simple univariate time series model of exponential progress, which we call Moore's law. While Alberth (2008) found that the experience curve model was vastly superior to an exogenous time trend, our results (and method and dataset) are closer to the findings of Nagy et al. (2013): univariate and experience curves models tend to perform similarly, due to the fact that for many products cumulative experience grows exponentially. When this is the case, we cannot expect experience curves to perform much better than an exogenous time trend unless cumulative experience is very volatile, as we explain in detail here.

We should emphasize that this comparison is difficult because the forecasts are conditioned on different variables: Moore's law is conditioned on time, while Wright's law is conditioned on experience. Which of these is more useful depends on the context. As we demonstrate, Moore's law is more convenient and just as good for business as usual forecasts for a given time in the future. However, providing there is a causal relationship from experience to cost, Wright's law makes it possible to forecast for policy purposes (Way et al., 2017).

Finally, we depart from Alberth (2008), Nagy et al. (2013) and most of the literature by using a different statistical model. As we explain in the next section, we have chosen to estimate a model in which the variables are first-differenced, instead of kept in levels as is usually done. From a theoretical point of view, we believe that it is reasonable to think that the stationary relationship is between the increase of experience and technological progress, instead of between a level of experience and a level of technology. In addition, we will also introduce a moving average noise, as in Farmer and Lafond (2016). This is meant to capture some of the complex autocorrelation patterns present in the data in a parsimonious way, and increase theoretical forecast errors so that they match the empirical forecast errors.

Our focus is on understanding the forecast errors from a simple experience curve model[3]. The experience curve, like any model, is only an approximation. Its simplicity is both a virtue and a detriment. The virtue is that the model is so simple that its parameters can usually be estimated well enough to have predictive value based on the short data sets that are typically available[4]. The detriment is that such a simple model neglects many effects that are likely to be important. A large literature starting with Arrow (1962) has convincingly argued that learning-by-doing occurs during the production (or investment) process, leading to decreasing unit costs. But innovation is a complex process relying on a variety of interacting factors such as economies of scale, input prices, R&D and patents, knowledge depreciation effects, or other effects captured by exogenous time trends.[5] For instance, Candelise et al. (2013) argue that there is a lot of variation around the experience curve trend in solar PV, due to a number of unmodelled mechanisms linked to industrial dynamics and international trade, and Sinclair et al. (2000) argued that the relationship between costs and experience is due to experience driving expectations of future production and thus incentives to invest in R&D. Besides, some have argued that simple exponential time trends are more reliable than experience curves. For instance Funk and Magee (2015) noted that significant technological improvements can take place even though production experience did not really accumulate, and Magee et al. (2016) found that in domains where experience (measured as annual patent output) did not grow exponentially, costs still had an exponentially decreasing pattern, breaking down the experience curve. Finally, another important aspect that we do not address is reverse causality (Kahouli-Brahmi, 2009; Nordhaus, 2014; Witajewski-Baltvilks et al., 2015): if demand is elastic, a decrease in price should lead to an increase in production. Here we have intentionally focused on the simplest case in order to develop the method.

2. Empirical framework

2.1. The basic model

Experience curves postulate that unit costs decrease by a constant factor for every doubling of cumulative production[6]. This implies a linear relationship between the log of the cost, which we denote y, and the log of cumulative production which we denote x:

$$y_t = y_0 + \omega x_t. \tag{1}$$

This relationship has also often been called "the learning curve" or the experience curve. We will often call it "Wright's law" in reference to Wright's original study, and to express our agnostic view regarding the causal mechanism. Generally, experience curves are estimated as

$$y_t = y_0 + \omega x_t + \iota_t, \tag{2}$$

where ι_t is i.i.d. noise. However, it has sometimes been noticed that residuals may be autocorrelated. For instance Womer and Patterson (1983) noticed that autocorrelation "seems to be an important problem" and Lieberman (1984) "corrected" for autocorrelation using the Cochrane-Orcutt procedure.[7] Bailey et al. (2012) proposed to estimate Eq. (1) in first difference

$$y_t - y_{t-1} = \omega(x_t - x_{t-1}) + \eta_t, \tag{3}$$

where η_t are i.i.d errors $\eta_t \sim \mathcal{N}(0, \sigma_\eta^2)$. In Eq. (3), noise accumulates so that in the long run the variables in level can deviate significantly from the deterministic relationship. To see this, note that (assuming $x_0 = \log(1) = 0$) Eq. (3) can be rewritten as

[2] Alberth (2008) produced forecasts for a number (1,2, … 6) of doublings of cumulative production. Here instead we use time series methods so it is more natural to compute everything in terms of calendar forecast horizon.

[3] We limit ourselves to showing that the forecast errors are compatible with our model being correct, and we do not try to show that they could be compatible with the experience curve model being spurious.

[4] For short data sets such as most of those used here, fitting more than one parameter often results in degradation in out-of-sample performance (Nagy et al., 2013).

[5] For examples of papers discussing these effects within the experience curves framework, see Argote et al. (1990), Berndt (1991), Isoard and Soria (2001), Papineau (2006), Söderholm and Sundqvist (2007), Jamasb (2007), Kahouli-Brahmi (2009), Bettencourt et al. (2013), Benson and Magee (2015) and Nordhaus (2014).

[6] For other parametric models relating experience to costs see Goldberg and Touw (2003) and Anzanello and Fogliatto (2011).

[7] See also McDonald (1987), Hall and Howell (1985), and Goldberg and Touw (2003) for further discussion of the effect of autocorrelation on different estimation techniques.