Research Policy 45 (2016) 647–665

Contents lists available at ScienceDirect

Research Policy

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

How predictable is technological progress?

J. Doyne Farmer[a,b,c], François Lafond[a,d,e,*]

[a] Institute for New Economic Thinking at the Oxford Martin School, University of Oxford, Oxford OX2 6ED, UK
[b] Mathematical Institute, University of Oxford, Oxford OX1 3LP, UK
[c] Santa-Fe Institute, Santa Fe, NM 87501, USA
[d] London Institute for Mathematical Sciences, London W1K 2XF, UK
[e] United Nations University – MERIT, 6211TC Maastricht, The Netherlands

ARTICLE INFO

Article history:
Received 10 March 2015
Received in revised form
30 September 2015
Accepted 2 November 2015
Available online 6 January 2016

JEL classification:
C53
O30
Q47

Keywords:
Forecasting
Technological progress
Moore's law
Solar energy

ABSTRACT

Recently it has become clear that many technologies follow a generalized version of Moore's law, i.e. costs tend to drop exponentially, at different rates that depend on the technology. Here we formulate Moore's law as a correlated geometric random walk with drift, and apply it to historical data on 53 technologies. We derive a closed form expression approximating the distribution of forecast errors as a function of time. Based on hind-casting experiments we show that this works well, making it possible to collapse the forecast errors for many different technologies at different time horizons onto the same universal distribution. This is valuable because it allows us to make forecasts for any given technology with a clear understanding of the quality of the forecasts. As a practical demonstration we make distributional forecasts at different time horizons for solar photovoltaic modules, and show how our method can be used to estimate the probability that a given technology will outperform another technology at a given point in the future.

1. Introduction

Technological progress is widely acknowledged as the main driver of economic growth, and thus any method for improved technological forecasting is potentially very useful. Given that technological progress depends on innovation, which is generally thought of as something new and unanticipated, forecasting it might seem to be an oxymoron. In fact there are several postulated laws for technological improvement, such as Moore's law and Wright's law, that have been used to make predictions about technology cost and performance. But how well do these methods work?

Predictions are useful because they allow us to plan, but to form good plans it is necessary to know probabilities of possible outcomes. Point forecasts are of limited value unless they are very accurate, and when uncertainties are large they can even be dangerous if they are taken too seriously. At the very least one needs error bars, or better yet, a distributional forecast, estimating the likelihood of different future outcomes. Although there are now a few papers testing technological forecasts[1] there is as yet no method that gives distributional forecasts based on an empirically validated stochastic process. In this paper we remedy this situation by deriving the distributional errors for a simple forecasting method and testing our predictions on empirical data on technology costs. To motivate the problem that we address, consider three technologies related to electricity generation: coal mining, nuclear power and photovoltaic modules. Fig. 1 compares their long-term historical prices. Over the last 150 years the inflation-adjusted price of coal has fluctuated by a factor of three or so, but shows no long term trend, and indeed from the historical time series one cannot reject the null hypothesis of a random walk with

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

[1] See e.g. Alchian (1963), Alberth (2008). Nagy et al. (2013) test the relative accuracy of different methods of forecasting statistically but do not produce and test a distributional estimate of forecast reliability for any particular method. McCrory, cited in Jantsch (1967), assumes a Gaussian distribution and uses this to calculate the probability that a targeted level of progress be met at a given horizon. Here we assume and test a Gaussian distribution for the natural log.

http://dx.doi.org/10.1016/j.respol.2015.11.001

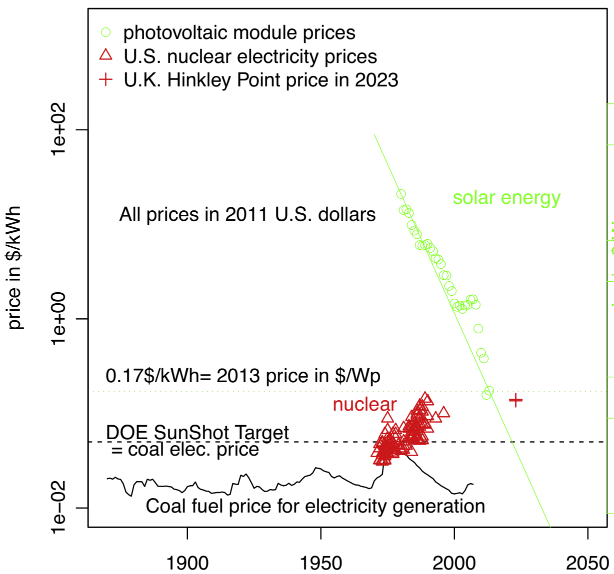

Fig. 1. A comparison of long-term price trends for coal, nuclear power and solar photovoltaic modules. Prices for coal and nuclear power are costs in the US in dollars per kilowatt hour (scale on the left) whereas solar modules are in dollars per watt-peak, i.e. the cost for the capacity to generate a watt of electricity in full sunlight (scale on the right). For coal we use units of the cost of the coal that would need to be burned in a modern US plant if it were necessary to buy the coal at its inflation-adjusted price at different points in the past. Nuclear prices are Busbar costs for US nuclear plants in the year in which they became operational (from Cooper (2009)). The alignment of the left and right vertical axes is purely suggestive; based on recent estimates of levelized costs, we took $0.177/kW h = $0.82/Wp in 2013 (2013$). The number $0.177/kW h is a global value produced as a projection for 2013 by the International Energy Agency (Table 4 in International Energy Agency (2014)). We note that it is compatible with estimated values (Table 1 in Baker et al. (2013), Fig. 4 in International Energy Agency (2014)). The red cross is the agreed price for the planned UK Nuclear power plant at Hinkley Point which is scheduled to come online in 2023 (£ 0.0925 ≈ $0.14). The dashed line corresponds to an earlier target of $0.05/kW h set by the U.S. Department of Energy.

no drift[2] (McNerney et al., 2011). Nuclear power and solar photovoltaic electricity, in contrast, are both new technologies that emerged at roughly the same time. The first commercial nuclear power plant opened in 1956 and the first practical use of solar photovoltaics was as a power supply for the Vanguard I satellite in 1958. The cost of electricity generated by nuclear power is highly variable, but has generally increased by a factor of two or three during the period shown here. In contrast, a watt of solar photovoltaic capacity cost $256 in 1956 (Perlin, 1999) (about $1910 in 2013 dollars) vs. $0.82 in 2013, dropping in price by a factor of about 2330. Since 1980 photovoltaic modules have decreased in cost at an average rate of about 10% per year.

In giving this example we are not trying to make a head-to-head comparison of the full system costs for generating electricity. Instead we are comparing three different technologies, coal mining, nuclear power and photovoltaic manufacture. Generating electricity with coal requires plant construction (whose historical cost has dropped considerably since the first plants came online at the beginning of the 20th century). Generating electricity via solar photovoltaics has balance of system costs that have not dropped as fast as that of modules in recent years. Our point here is that different technologies can decrease in cost at very different rates.

Predicting the rate of technological improvement is obviously very useful for planning and investment. But how consistent are such trends? In response to a forecast that the trends above will continue, a skeptic would rightfully respond, "How do we know that the historical trend will continue? Isn't it possible that things will reverse, and over the next 20 years coal will drop in price dramatically and solar will go back up?".

Our paper provides a quantitative answer to this question. We put ourselves in the past, pretend we don't know the future, and use a simple method to forecast the costs of 53 different technologies. Actually going through the exercise of making out-of-sample forecasts rather than simply doing in-sample regressions has the essential advantage that it fully mimics the process of making forecasts and allows us to say precisely how well forecasts would have performed. Out-of-sample testing such as we do here is particularly important when models are mis-specified, which one expects for a complicated phenomenon such as technological improvement.

We show how one can combine the experience from forecasting many technologies to make reliable distributional forecasts for a given technology. For solar PV modules, for example, we can say, "Based on experience with many other technologies, the probability is roughly 5% that in 2030 the price of solar PV modules will be greater than or equal to their current (2013) price". We can assign a probability to different price levels at different points in the future, as is done later in Fig. 10 (where we show that very likely the price will drop significantly). We can also compare different technologies to assess the likelihood of different future scenarios for their relative prices, as is done in Fig. 11.

Technological costs occasionally experience structural breaks where trends change. Indeed there are several clear examples in our historical data, and although we have not explicitly modeled this, their effect on forecast errors is included in the empirical analysis we have done here. The point is that, while such structural breaks happen, they are not so large and so common as to over-ride our ability to forecast. Every technology has its own story, its own specific set of causes and effects, that explain why costs went up or down in any given year. Nonetheless, as we demonstrate here, the long term trends tend to be consistent, and can be captured via historical time series methods with no direct information about the underlying technology-specific stories.

In this paper we use a very simple approach to forecasting that was originally motivated by Moore's Law. As everyone knows, Intel's ex-CEO, Gordon Moore, famously predicted that the number of transistors on integrated circuits would double every two years, i.e. at an annual rate of about 40%. Making transistors smaller also brings along a variety of other benefits, such as increased speed, decreased power consumption, and less expensive manufacture costs per unit of computation. As a result it quickly became clear that Moore's law applies more broadly, for example, implying a doubling of computational speed every 18 months.

Moore's law stimulated others to look at related data more carefully, and they discovered that exponential improvement is a reasonable approximation for other types of computer hardware as well, such as hard drives. Since the performance of hard drives depends on physical factors that are unrelated to transistor density this is an independent fact, though of course the fact that mass storage is essential for computation causes a tight coupling between the two technologies. Lienhard, Koh and Magee, and others[3] examined data for other products, including many that have nothing to do

[2] To drive home the point that fossil fuels show no long term trend of dropping in cost, after adjusting for inflation coal now costs about what it did in 1890, and a similar statement applies to oil and gas.

[3] Examples include Lienhard (2006), Koh and Magee (2006, 2008), Bailey et al. (2012), Benson and Magee (2014a,b), Nagy et al. (2013). Studies of improvement in computers over long spans of time indicate super-exponential improvement (Nordhaus, 2007; Nagy et al., 2011), suggesting that Moore's law may only be an approximation reasonably valid over spans of time of 50 years or less. See also e.g. Funk (2013) for an explanation of Moore's law based on geometric scaling, and Funk and Magee (2014) for empirical evidence regarding fast improvement prior to large production increase.