

Statistical Basis for Predicting Technological Progress

Béla Nagy^{1,*}, J. Doyne Farmer¹, Quan M. Bui^{1,2} Jessika E Trancik^{1,3,*}

1 Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501, USA

2 St. John's College, 1160 Camino Cruz Blanca, Santa Fe, NM, 87505, USA

3 Engineering Systems Division, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

* E-mail: Corresponding bn@santafe.edu, trancik@mit.edu

Abstract

Forecasting technological progress is of great interest to engineers, policy makers, and private investors. Several models have been proposed for predicting technological improvement, but how well do these models perform? An early hypothesis made by Theodore Wright in 1936 is that cost decreases as a power law of cumulative production. An alternative hypothesis is Moore's law, which can be generalized to say that technologies improve exponentially with time. Other alternatives were proposed by Goddard, Sinclair et al., and Nordhaus. These hypotheses have not previously been rigorously tested. Using a new database on the cost and production of 62 different technologies, which is the most expansive of its kind, we test the ability of six different postulated laws to predict future costs. Our approach involves hindcasting and developing a statistical model to rank the performance of the postulated laws. Wright's law produces the best forecasts, but Moore's law is not far behind. We discover a previously unobserved regularity that production tends to increase exponentially. A combination of an exponential decrease in cost and an exponential increase in production would make Moore's law and Wright's law indistinguishable, as originally pointed out by Sahal. We show for the first time that these regularities are observed in data to such a degree that the performance of these two laws is nearly tied. Our results show that technological progress is forecastable, with the square root of the logarithmic error growing linearly with the forecasting horizon at a typical rate of 2.5% per year. These results have implications for theories of technological change, and assessments of candidate technologies and policies for climate change mitigation.

Introduction

Innovation is by definition new and unexpected, and might therefore seem inherently unpredictable. But if there is a degree of predictability in technological innovation, understanding it could have profound implications. Such knowledge could result in better theories of economic growth, and enable more effective strategies for engineering design, public policy design, and private investment. In the area of climate change mitigation, the estimated cost of achieving a given greenhouse gas concentration stabilization target is highly sensitive to assumptions about future technological progress [1].

There are many hypotheses about technological progress, but are they any good? Which, if any, hypothesis provides good forecasts? In this paper, we present the first statistically rigorous comparison of competing proposals.

When we think about progress in technologies, the first product that comes to mind for many is a computer, or more generally, information technologies. The following quote by Bill Gates captures a commonly held view: "Exponential improvement – that is rare – we've all been spoiled and deeply confused by the IT model" [2]. But as we demonstrate here, information technologies are not special in terms of the functional form that describes their improvement over time. Information technologies show rapid rates of improvement, but many technologies show exponential improvement. In fact, all the technologies we study here behave roughly similarly: Information technologies closely follow patterns of improvement originally postulated by Wright for airplanes [3–8], and technologies such as beer production

or offshore gas pipelines follow Moore’s law [9, 10], but with a slower rate of improvement [8, 11–15].

It is not possible to quantify the performance of a technology with a single number [16]. A computer, for example, is characterized by speed, storage capacity, size and cost, as well as other intangible characteristics such as aesthetics. One automobile may be faster while another is less expensive. For this study we focus on one common measure of performance: the inflation-adjusted cost of one “unit”. This metric is suitable in that it can be used to describe many different technologies. However, the nature of a unit may change over time. For example, a transistor in a modern integrated circuit today may have quite different performance characteristics than its discrete counterpart in the past. Furthermore, the degree to which cost is emphasized over other performance measures may change with time [17]. We nonetheless use the changes in the unit cost as our measure of progress, in order to compare competing models using a sizable dataset. The crudeness of this approach only increases the difficulty of forecasting and makes it particularly surprising that we nonetheless observe common trends.

Analysis

We test six different hypotheses that have appeared in the literature [3, 9, 18–20], corresponding to the following six functional forms:

$$\begin{aligned}
 \text{Moore} \quad \log y_t &= at + b + n(t) & (1) \\
 \text{Wright} \quad \log y_t &= a \log x_t + b + n(t) \\
 \text{lagged Wright} \quad \log y_t &= a \log(x_t - q_t) + b + n(t) \\
 \text{Goddard} \quad \log y_t &= a \log q_t + b + n(t) \\
 \text{SKC} \quad \log y_t &= a \log q_t + c \log(x_t - q_t) + b + n(t) \\
 \text{Nordhaus} \quad \log y_t &= at + c \log x_t + b + n(t)
 \end{aligned}$$

The dependent variable y_t is the unit cost of the technology measured in inflation adjusted dollars. The independent variables are the time t (measured in years), the annual production q_t , and the cumulative production $x_t = \sum_{i=1}^t q_i$. The noise term $n(t)$, the constants a , b and c and the predictor variables differ for each hypothesis.

Moore’s law here refers to the generalized statement that the cost y of a given technology decreases exponentially with time, i.e.

$$y_t = B \exp(-mt), \quad (2)$$

where $m > 0$ and $B > 0$ are constants [9, 12]. (We assume throughout that $t > 0$, and we have renamed $a = -m$ and $b = \log B$ in Eq. (1)). Moore’s law postulates that technological progress is inexorable, i.e. it depends on time rather than controllable factors such as research and development.

Wright’s law, in contrast, postulates that cost decreases at a rate that depends on cumulative production, i.e.

$$y_t = Bx_t^{-w}, \quad (3)$$

where $w > 0$ and $B > 0$ are constants, and we have renamed $a = -w$ and $b = \log B$ in Eq. (1). Wright’s law is often interpreted to imply “learning by doing” [5, 21]. The basic idea is that cumulative production is a proxy for level of effort, so that the more we make the more we learn, and knowledge accumulates without loss.

Another hypothesis is due to Goddard [18], who argues that progress is driven purely by economies of scale, and postulates that:

$$y_t = Bq_t^{-s}, \quad (4)$$

where $s > 0$ and $B > 0$ are constants, and we have renamed $a = -s$ and $b = \log B$ in Eq. (1).

We also consider the three multi-variable hypotheses in Eq. (1): Nordhaus [20] combines Wright’s law and Moore’s law¹, and Sinclair, Klepper, and Cohen (SKC) [19] combine Wright’s law and Goddard’s law. For completeness we also test Wright’s law lagged by one year. Note that these methods forecast different things: Moore’s law forecasts the cost at a given time, Wright’s law at a given cumulative production, and Goddard’s law at a given annual production.

We test these hypotheses on historical data consisting of 62 different technologies that can be broadly grouped into four categories: Chemical, Hardware, Energy, and Other². The data are sampled at annual intervals with timespans ranging from 10 to 39 years. The choice of these particular technologies was driven by availability – we included all of the data that we could find to assemble the largest database of its kind. For a detailed description see the Supporting Information.

To compare the performance of each hypothesis we use hindcasting, which is a form of cross-validation. We pretend to be at time i and make a forecast $\hat{y}_j^{(f,d,i)}$ for time j using hypothesis (functional form) f and data set d , where $j > i$. The parameters for each functional form are fit using ordinary least squares based on all data prior to time i , and forecasts are made based on the resulting regression³. We score the quality of forecasts based on the logarithmic forecasting error

$$e_{fdij} = \log y_j^{(d)} - \log \hat{y}_j^{(f,d,i)} \quad (5)$$

The quality of forecasts is examined for all datasets and all hypotheses (and visualized as a three-dimensional error mountain, as shown in the Supporting Information). For Wright’s law an illustration of the growth of forecasting errors as a function of the forecasting horizon is given in Fig. ??.

Developing a statistical model to compare the competing hypotheses is complicated by the fact that errors observed at longer horizons tend to be larger than those at shorter horizons, and errors are correlated across time and across functional forms. After comparing many different possibilities (as discussed in detail in the Supporting Information), we settled on the following approach. Based on a search of the family of power transformations, which is known for its ability to accommodate a range of variance structures, we take as a response the square root transformation of the logarithmic error. This response was chosen to maximize likelihood when modeled as a linear function of the hindcasting horizon = target – origin = $j - i$, using a linear mixed effects model.

Specifically, we use the following functional form to model the response.

$$r_{fdij} \equiv |e_{fdij}|^{0.5} = \alpha_f + a_d + (\beta_f + b_d)(j - i) + \epsilon_{fdij}, \quad (6)$$

where r_{fdij} is the expected root error. The parameters α_f and β_f depend on the functional form and are called *fixed effects* because they are the same for all datasets. α_f is the intercept and β_f is the slope parameter.

The parameters a_d and b_d depend on the dataset, and are called *random effects* because they are not fitted independently, but are instead treated as dataset-specific random fluctuations from the pooled data. The quantities a_d and b_d are additive adjustments to the average intercept and slope parameters α_f and β_f , respectively, to take into account the peculiarities of each dataset d .

In order to avoid adding 62 a_d parameters plus 62 b_d parameters, we treated the $\begin{pmatrix} a_d \\ b_d \end{pmatrix}$ pair as a two-dimensional random vector having a bivariate normal distribution with mean $\begin{pmatrix} 0 \\ 0 \end{pmatrix}$ and variance-

¹Note that the conclusions presented do not work against Nordhaus’ point about the difficulty in separating learning from exogenous sources of change [20].

²All data can be found in the online Performance Curve Database at pcdb.santafe.edu.

³An alternative is to adjust the intercepts to match the last point, which produces better short term forecasts. For example, for Moore’s law this corresponds to using a log random walk of the form $\log y_{t+1} = \log y_t - \mu + n(t)$, where $n(t)$ is an IID noise term. We have not done this here to be consistent with the way these hypotheses have been historically presented. The method we have used here also results in more stable errors. Our purpose here is not to propose an optimal forecasting method, but rather to compare existing hypotheses.

covariance matrix $\begin{pmatrix} \psi_a^2 & \psi_{ab} \\ \psi_{ab} & \psi_b^2 \end{pmatrix}$. This approach dramatically reduces the number of parameters. We parameterize the dataset-specific adjustments as random deviations from the average $\begin{pmatrix} \alpha_f \\ \beta_f \end{pmatrix}$ at a cost of only 3 additional parameters instead of $2 \times 62 = 124$. This parsimonious approach makes maximum likelihood estimation possible by keeping the number of parameters in check.

Finally, we add an ε_{fdij} random field term to take into account the deviations from the trend. This is assumed to be a Gaussian stochastic process independent of the $\begin{pmatrix} a_d \\ b_d \end{pmatrix}$ random vector, having mean 0, and given a_d and b_d , having variance equal to a positive σ^2 times the fitted values:

$$\text{Var}(\varepsilon_{fdij} | a_d, b_d) = \sigma^2 \text{E}(r_{fdij} | a_d, b_d) \quad (7)$$

We also define an exponential correlation structure within each error mountain (corresponding to each combination of dataset and hypothesis, see the Supporting Information), as a function of the differences of the two time coordinates with a positive range parameter ρ and another small positive nugget parameter η quantifying the extent of these correlations:

$$\text{Corr}(\varepsilon_{fdij}, \varepsilon_{fd'ij'}) = \delta_{ff'} \delta_{dd'} (1 - \eta) \exp\{-(|i - i'| + |j - j'|)/\rho\}, \quad (8)$$

where the two Kronecker δ functions ensure that each error mountain is treated as a separate entity.

Equations (7) and (8) were chosen to deal with the observed heteroscedasticity (increasing variance with increasing logarithmic forecasting error) and the serial correlations along the time coordinates i (hindcasting origin) and j (hindcasting target). Based on the likelihood, an exponential correlation function provided the best fit. Note that instead of a Euclidean distance (root sum of the squares of differences), the Manhattan measure was used (the sum of the absolute differences), because it provided a better fit in terms of the likelihood.

Using this statistical model, we compared five different hypotheses. (We removed the Nordhaus model from the sample because of poor forecasting performance. This model gave good in-sample fits but generated large and inconsistent errors when predicting out-of-sample, a signature of over-fitting.) Rather than the $62 \times 5 \times 2 = 620$ parameters needed to fit each of the 62 datasets separately for each of the five functional forms, there are only 16 free parameters: $5 \times 2 = 10$ parameters α_f and β_f , three parameters for the covariance matrix of the bivariate random vector (a_d, b_d) , and three parameters for the variance and autocorrelation of the residuals ε_{fdij} .

Results and Discussion

We fit the error model to the 37,745 different r_{fdij} data points using the method of maximum likelihood. In Fig. 2 we plot the expected root error $r_{fij} = \alpha_f + \beta_f(j - i)$ for the five hypotheses as a function of the hindcasting horizon. While there are differences in the performance of these five hypotheses, they are not dramatic. The intercept is tightly clustered in a range $0.16 < \alpha_f < 0.19$ and the slope $0.024 < \beta_f < 0.028$. Thus all the hypotheses show a large initial error, followed by a growth in the root error of roughly 2.5% per year⁴.

The error model allows us to compare each hypothesis pairwise to determine whether it is possible to reject one in favor of another at statistically significant levels. The comparisons are based on the intercept and slope of the error model of Eq. (6). The parameter estimates are listed in Tables S1 and S3 and the corresponding p -values in Tables S2 and S4. For example, at the 5% level, the intercept of Goddard is significantly higher than any of the others and the slope of SKC is significantly greater than that of Wright, lagged Wright and Goddard. With respect to slope, Moore is at the boundary of being rejected

⁴This is a central tendency for the pooled data.

in favor of Wright. Fig. 2 makes the basic pattern clear: Goddard does a poorer job of forecasting at short times, whereas SKC and to a lesser extent Moore do a poorer job at long times.

We thus have the surprising result that most of the methods are quite similar in their performance. Although the difference is not large, the fact that we can eliminate Goddard for short term forecasts indicates that there is information in the cumulative production not contained in the annual production, and suggests that there is a learning effect in addition to economies of scale. But the fact that Goddard is not that much worse indicates that much of the predictability comes from annual production, suggesting that economies of scale are important. (In our database technologies rarely decrease significantly in annual production; examples of this would provide a better test of Goddard’s theory.) We believe the SKC model performs worse at long times because it has an extra parameter, making it prone to overfitting.

Although Moore performs slightly worse than Wright, given the clear difference in their economic interpretation, it is surprising that their performance is so similar. A simple explanation for Wright’s law in terms of Moore’s law was originally put forward by Sahal [22]. He noted that if cumulative production grows exponentially⁵, i.e.

$$x_t = A \exp(gt), \tag{9}$$

then eliminating t between Eqs. (2) and (9) results in Wright’s law, Eq. (3), with $w = m/g$. Indeed, when we look at production vs. time we find that in almost every case the cumulative production increases roughly exponentially with time. This is illustrated in Fig. 3, where we show three representative examples for production and cost plotted as a function of time. Fig. 3 also shows histograms of R^2 values for fitting g and m for the 62 datasets. The agreement with exponential behavior ranges from very good to rather poor, but of course these are short time series and some of them are very noisy.

We test this in Fig. 4 by plotting the measured value of w_d against the derived value $\hat{w}_d = m/g$ for each data set d . The values cluster tightly along the identity line, indicating that Sahal’s conjecture is correct.

The differences in the data sets can be visualized by plotting a_d and b_d as shown in Fig. 5. All but one data set is inside the 95% confidence ellipsoid, indicating that the estimated distribution of (a_d, b_d) is consistent with the bivariate normal assumption. The intercepts vary in a range roughly $-0.10 < a_d < 0.17$ and the slopes $-0.018 < b_d < 0.015$. Thus the variation in the corresponding logarithmic forecasting error for the different datasets is comparable to the average error for all datasets (Fig. 5) and about an order of magnitude larger than the difference between the hypothesized laws (Fig. 2).

To illustrate the practical usefulness of our approach we make a forecast of the cost of electricity for residential scale photovoltaic solar systems (PV). Fig. 6 shows the best forecast (solid line) as well as the expected error (dashed lines)⁶. The expected cost in 2020 is 6 cents/kWh, with a range (3, 12), and in 2030 it is 2 cents, with a range (0.4, 11)⁷. The current cost of the cheapest alternative, coal-fired electricity, is roughly 5 cents/kWh⁸. In contrast to PV, coal-fired electricity is not expected to decrease in cost, and will likely increase if there are future penalties for CO₂ emissions [23].

The key postulate that we have made here is that the processes generating the costs of technologies through time are generic except for technology-specific differences in parameters. This hypothesis is powerful in allowing us to view any given technology as being drawn from an ensemble. This means that we can pool data from different technologies to make better forecasts, and most importantly, make

⁵ Note that if production grows exponentially, cumulative production also grows exponentially with the same exponent.

⁶ Note that these are not confidence limits, but rather projected absolute log deviations from the best forecast, calculated from Eq. (6) using α_{Moore} , β_{Moore} , $a_{\text{Photovoltaics2}}$, and $b_{\text{Photovoltaics2}}$. The sharp drop in the one year forecast relative to the last observed data point comes from the fact that forecasts are based on the average trend line, and because this data series is particularly long. PV costs have risen recently due to increased material costs and other effects, but many industry experts expect this to be a short-lived aberration from the long-term cost trend. See footnote 4 and Section 5 of the Supporting Information.

⁷ This does not include the additional cost of energy storage technologies. Note also that this is for residential scale PV. Industrial scale PV is typically about two-thirds the cost of electricity from residential scale systems.

⁸ This is the wholesale cost at the plant (busbar), which may be most directly comparable to industrial scale PV.

error estimates. This is particularly useful in studying technology trends, where available data is limited. Of course we must add the usual caveats about making forecasts – as Neils Bohr reputedly said, prediction is very difficult, especially of the future.

Our analysis reveals that decreasing costs and increasing production appear to be closely related, and that the hypotheses of Wright and Moore are more similar than they might appear. We should stress, though, that they are not the same. For example, consider a scenario in which the exponential rate of growth of PV production suddenly increased, which would decrease the current production doubling time of roughly 3 years. In this case, Wright predicts that the rate at which costs fall would increase, whereas Moore predicts that it would be unaffected. Distinguishing between the two hypotheses requires a sufficient number of examples where production does not increase exponentially, which our current database does not contain. The historical data shows a strong tendency, across different types of technologies, toward constant exponential growth rates. Recent work has, however, demonstrated super-exponential improvement for information technologies [24] over long time spans, suggesting that Moore’s law is only a reasonable approximation over short spans of time. This evidence from information technology, and the results presented here, suggest that Moore may perform significantly worse than Wright over longer time horizons.

Acknowledgments

This work was funded under National Science Foundation grant NSF-0738187. All the conclusions and opinions are those of the authors and do not necessarily reflect those of the NSF. Partial funding for this work was provided by The Boeing Company (Contact: Scott H. Mathews, TF). We thank all contributors to the performance curve database (pcdb.santafe.edu).

References

1. Gillingham K, Newell RG, Pizer WA (2008) Modeling endogenous technological change for climate policy analysis. *Energy Economics* 30: 2734 - 2753.
2. Fried I (2010) Gates: We’ve been spoiled by Moore’s Law. *CNET News* .
3. Wright TP (1936) Factors affecting the costs of airplanes. *Journal of Aeronautical Sciences* 10: 302-328.
4. Argote L, Epple D (1990) Learning curves in manufacturing. *Science* 247: 920-924.
5. Dutton JM, Thomas A (1984) Treating progress functions as a managerial opportunity. *Academy of Management Review* 9: 235-247.
6. IEA (2000) Experience Curves for Energy Technology Policy. *OECD/IEA*.
7. McDonald A, Schrattenholzer L (2001) Learning rates for energy technologies. *Energy Policy* 29: 255-261.
8. Alberth S (2008) Forecasting technology costs via the experience curve - myth or magic? *Technological Forecasting and Social Change* 75: 952-983.
9. Moore GE (1965) Cramming more components onto integrated circuits. *Electronics Magazine* 38.
10. Moore GE (1975) Progress in digital integrated electronics. *IEEE International Electron Devices Meeting* : 11-13.

11. Kurzweil R (2005) *The Singularity Is Near: When Humans Transcend Biology*. Viking Penguin. <http://singularity.com/>.
12. Koh H, Magee CL (2006) A functional approach for studying technological progress: Application to information technology. *Technological Forecasting and Social Change* 73: 1061-1083.
13. Nordhaus WD (2007) Two centuries of productivity growth in computing. *The Journal of Economic History* 67: 128-159.
14. Koh H, Magee CL (2008) A functional approach for studying technological progress: Extension to energy technology. *Technological Forecasting and Social Change* 75: 735-758.
15. Amaya MA, Magee CL (2008) The progress in wireless data transport and its role in the evolving internet. Working paper, Massachusetts Institute of Technology.
16. Dosi G, Nelson RR (2010) Technical change and industrial dynamics as evolutionary processes. In: Hall B, Rosenberg N, editors, *Handbook of the Economics of Innovation*, Elsevier/North Holland.
17. Abernathy WJ, Wayne K (1974) Limits of the learning curve. *Harvard Business Review* : 109-119.
18. Goddard C (1982) Debunking the learning curve. *IEEE Transactions on Components, Hybrids, and Manufacturing Technology* 5: 328-335.
19. Sinclair G, Klepper S, Cohen W (2000) What's experience got to do with it? sources of cost reduction in a large specialty chemical producer. *Management Science* 46: 28-45.
20. Nordhaus WD (2009) The perils of the learning model for modeling endogenous technological change. Technical report, National Bureau of Economic Research.
21. Thompson P (2010) Learning by doing. In: Hall B, Rosenberg N, editors, *Handbook of the Economics of Innovation*, Elsevier/North Holland.
22. Sahal D (1979) A theory of progress functions. *AIIE Transactions* 11: 23-29.
23. McNerney J, Farmer JD, Trancik JE (2011) Historical costs of coal-fired electricity and implications for the future. *Energy Policy* 39: 3042-3054.
24. Nagy B, Farmer JD, Trancik JE, Gonzales JP (2011) Superexponential long-term trends in information technology. *Technological Forecasting and Social Change* 78: 1356 - 1364.

Figure Legends




Figure 1. An illustration of the growth of errors with time using the Wright model. The mean value of the logarithmic hindcasting error for each dataset is plotted against the hindcasting horizon $j - i$, in years. An error of $10^{0.5} \approx 3$, for example, indicates that the predicted value is three times as big as the actual value. The longest data-sets are: PrimaryAluminum (green), PrimaryMagnesium (dark blue), DRAM (grey), and Transistor (red).




Figure 2. An illustration of the growth of errors of each hypothesized law vs. time. The plot shows the predicted root absolute log error $r_{f_{ij}}$ vs. forecasting horizon $(j - i)$ using each of the functional forms (see Eq. (6)). The performance of the five hypotheses shown is fairly similar, though Goddard is worse at short horizons and SKC and Moore are worse at long horizons.

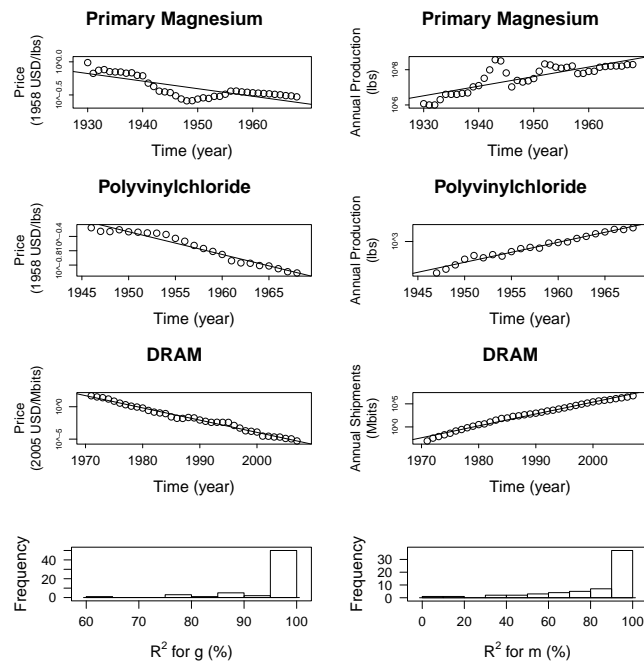


Figure 3. Three examples showing the logarithm of price as a function of time in the left column and the logarithm of production as a function of time in the right column, based on industry-wide data. We have chosen these examples to be representative: The top row contains an example with one of the worst fits, the second row an example with an intermediate goodness of fit, and the third row one of the best examples. The fourth row of the figure shows histograms of R^2 values for fitting g and m for the 62 datasets.

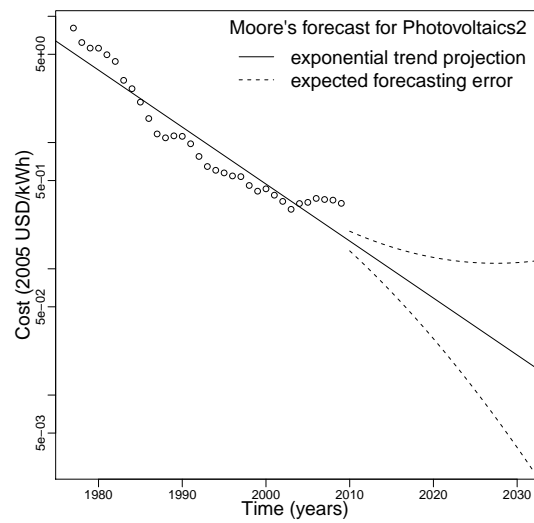


Figure 6. A projection of future PV electricity costs from the Photovoltaics2 historical data set (1977 – 2009) using Moore's exponential functional form. The solid line is the expected forecast and the dashed line is the expected error.