Fitting a Gompertz Curve

Philip Hans Franses

The Journal of the Operational Research Society, Vol. 45, No. 1 (Jan., 1994), 109-113.

Stable URL:
http://links.jstor.org/sici?sici=0160-5682%28199401%2945%3A1%3C109%3AFAGC%3E2.0.CO%3B2-I

The Journal of the Operational Research Society is currently published by Operational Research Society.

J. Opl Res. Soc. Vol. 45, No. 1, pp. 109–113
Printed in Great Britain.

0160-5682/94 $9.00+0.00

Fitting a Gompertz Curve

PHILIP HANS FRANSES
Erasmus University Rotterdam, The Netherlands

In this paper, a simple Gompertz curve-fitting procedure is proposed. Its advantages include the facts that the stability of the saturation level over the sample period can be checked, and that no knowledge of its value is necessary for forecasting. An application to forecasting the stock of cars in the Netherlands illustrates its merits.

Key words: time series, forecasting

INTRODUCTION

Forecasting the stock of vehicles or the prospective sales of a new product are examples of practical occasions in which a univariate time series can be usefully, although roughly, described by a Gompertz trend curve. Loosely speaking, this curve has an S-shape which, in contrast to the often-applied logistic curve, is non-symmetrical. More precisely, the Gompertz curve assumes that the period of increasing growth of sales or stock is shorter than the period in which this growth is decreasing and in which the process is adjusting to its saturation level. See Harrison and Pearce[1], Mar Molinero[2] and Meade[3] for surveys of the distinct types of trend curves.

The saturation level is one of the three unknown parameters in the model, and its value is usually assumed *a priori* or estimated iteratively. Clearly, this value plays a central role in the forecasting of future values of a time series process. Further, given that the variability of time series like stocks and sales are typically related to the level of the series, it is common practice to weight the observations in the Gompertz curve estimation procedures. Unfortunately, these weights have to be estimated from the same data, see Harrison and Pearce[1], and can therefore also be influential for forecasting. In the present paper I propose a simple Gompertz curve fitting procedure which deals with these issues. It is applied to forecasting the stock of cars in the Netherlands.

FITTING A GOMPERTZ CURVE

A Gompertz curve is depicted in Figure 1. Two characteristics of this curve are clearly observable. The first is the point of inflexion, i.e. the point in time at which the rate of growth changes from increasing to decreasing. This occurs before half the saturation level is reached. Second, the rate of growth is always larger than (although it decreases to) zero. These characteristics establish the usefulness of fitting a Gompertz growth curve for time series processes like the stock of cars or the sales of a new product. The mathematical representation of the process x_t depicted in Figure 1 is given by

$$x_t = \alpha \exp(-\beta \exp(-\gamma t)), \tag{1}$$

where α, β and γ are unknown positive-valued parameters, the first of which is the value of the saturation level. The t is a linear deterministic time trend defined by $t = 0, 1, 2, \ldots$

To fit model (1) to an empirical series, current estimation procedures often consider three observations of the process x_t, see Meade[4] and Oliver[5] for alternative approaches. Substituting these and their corresponding values for t in (1) gives three equations with three unknown parameters. The three observations are usually linear combinations of the first, middle and last observations on x_t. One may also want to minimize the sum of squared errors in a grid-search in which the value of α is fixed and the remaining parameters are estimated. Since empirical observations on

Correspondence: P. H. Franses, Econometric Institute, Erasmus University Rotterdam, PO Box 1738, NL-3000 DR Rotterdam, The Netherlands

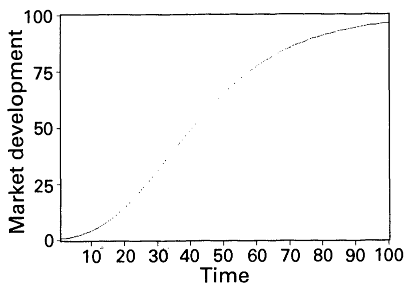

FIG. 1. *A Gompertz curve.*

a sales or stock process are typically non-stationary, i.e. their mean and variance are not constant over time, it is common practice to weight the observations according to their level. There are also examples in which, instead of weighting, the model in (1) is enlarged with a first-order autoregressive error process, see for example Mar Molinero[2]. In that study[2] it was found that either the corresponding parameter is close to 1, or that the autoregressive process contains a unit root. In the latter case, standard statistical inference is not valid, and it is difficult to establish the significance of the parameter. See Granger and Newbold[6] for a survey of the definitions and tests for non-stationarity. In summary, there are several choices to be made for the empirical fitting of a Gompertz curve, and it seems that these can have a large impact on forecasting performance.

There is, however, a simple strategy to circumvent the above problems. This is based on transforming the process in (1) via taking first differences and logarithms, see Harvey[7] for a related approach in the case of a logistic trend curve. Denoting log for the natural logarithm, consider a transformed version of (1),

$$\log x_t = \log\alpha - \beta\exp(-\gamma t). \tag{2}$$

Taking first differences of $\log x_t$, or $\Delta\log x_t = \log x_t - \log x_{t-1}$, gives

$$\begin{aligned}\Delta\log x_t &= -\beta\exp(-\gamma t) + \beta\exp(-\gamma t + \gamma) \\ &= \exp(-\gamma t)(\beta\exp\gamma - \beta),\end{aligned} \tag{3}$$

where now the value of α is removed from the equation. A linearization of (3), and adding an error term, yields the equation

$$\log(\Delta\log x_t) = -\gamma t + \log(\beta\exp\gamma - \beta) + \varepsilon_t \tag{4}$$

of which the parameters and their standard errors can easily be estimated using standard non-linear least squares (NLS) techniques, given an appropriate choice of starting values. A sequence of saturation levels can be estimated via

$$\hat{\alpha}_t = \exp(\log x_t + \hat{\beta}\exp(-\hat{\gamma}t)) \tag{5}$$

which is also a simple calculation.

There are several interesting aspects to equations (4) and (5). First, the problem of the non-stationarity of empirical x_t series, and its effect on parameter estimates has been overcome since $\log(\Delta\log x_t)$ is a trend stationary variable, see also Figure 12 in Harrison and Pearce[1]. Hence, inference on the parameters β and γ, such as hypothesis tests on their values, can be carried out using standard procedures. Also, distributional assumptions on ε_t in (4), for example normality, are easily made, cf. Bewley and Fiebig[8]. Second, all parameters can easily be estimated. Other forms of (1), such as

$$\log x_t - \log\alpha = \exp(-\gamma)(\log x_{t-1} - \log\alpha) \tag{6}$$

or

$$\Delta\log x_t = \exp(-\gamma)\Delta\log x_{t-1} \tag{7}$$

suffer from empirical regressors which are likely to be non-stationary as in (6), or from the fact that only one of the three parameters is estimated as in (7). Thirdly, although the level α is assumed to be constant over the sample, there is an opportunity to check its stability by considering the range of estimated α_t values obtained from (5). Structural breaks in the saturation level because of, for example, technology shocks can now also be detected. To obtain one estimate for the value α, $\hat{\alpha}$, one could take, an average of the, possibly smoothed, $\hat{\alpha}_t$ series. Similarly, one can calculate the standard error of $\hat{\alpha}$. An alternative procedure for estimating α is to calculate a forecast for x_t via (4) when allowing t to be very large. That value should then come close to the saturation level. Finally, although of interest for the understanding of the process under consideration, the values of α and β are not required for forecasting future values of x_t. The forecasts can simply be obtained from a recursion formula implied by (4). Of course, this also applies to (7).

There are several possibilities to calculate forecast intervals for x_t, or to calculate the forecast errors. Given that the model is usually estimated for only a small number of observations, and also considering the complicated non-linear function of x_t in (4), the most simple method may be to use a bootstrap-technique. Suppose that (4) is estimated with T_1 observations, that the estimated variance $\hat{\varepsilon}_t$ is $\hat{\sigma}^2$, and that forecasts are required for T_2 periods ahead. The bootstrap method boils down to, first, generating $T = T_1 + T_2$ observations for ε_t^* from a normal distribution with mean 0 and variance $\hat{\sigma}^2$, then generating observations on y_t^* via

$$y_t^* = -\hat{\gamma}t + \log(\hat{\beta}\exp\hat{\gamma} - \hat{\beta}) + \varepsilon_t^*, \tag{8}$$

re-estimating (4) for the first T_1 observations only, and finally forecasting x_t^* for T_2 data points. When this is repeated B times, one can calculate the mean and variance of the B forecasts for each horizon.

AN APPLICATION

To illustrate the proposed Gompertz curve fitting method, I have chosen to consider the stock of cars in the Netherlands, 1965–1989. The raw and smoothed observations are given in Table 1, and the graph of the smoothed series is depicted in Figure 2. It can be seen that a Gompertz curve may be a useful description of these data. Furthermore, given that long range forecasts of the stock of cars are quite important for a country as small as the Netherlands, one can imagine that an estimate of the saturation level may also have policy implications.

TABLE 1. *The (smoothed) stock of cars in the Netherlands, 1965–1989 (× 1000)*

Year	Stock (raw series)	Stock (smoothed)	Year	Stock (raw series)	Stock (smoothed)
1965	1273	1284	1978	4056	4066
1966	1502	1492	1979	4312	4281
1967	1696	1700	1980	4515	4455
1968	1952	1921	1981	4594	4600
1969	2212	2184	1982	4630	4659
1970	2465	2444	1983	4728	4729
1971	2702	2716	1984	4818	4828
1972	2903	2916	1985	4901	4921
1973	3080	3123	1986	4950	5035
1974	3214	3252	1987	5118	5095
1975	3399	3406	1988	5251	5213
1976	3629	3618	1989	5371	5322
1977	3851	3865			

The stock data are smoothed by regressing the stock in year t on a constant, on the stock in year $t - 1$ and on the new car sales in year t, and by taking the fitted stock series as the smoothed series.

Denoting s_t as the smoothed stock of cars, the estimation results of the model in (4) are (with standard errors in parentheses): $\hat{\beta} = 1.500$ (0.099) and $\hat{\gamma} = 0.104$ (0.009). This model is estimated with 24 observations, and tests for residual autocorrelation and normality indicate no misspecification. The value of the R^2 for (4) equals 0.857. The 25 estimates of the saturation level

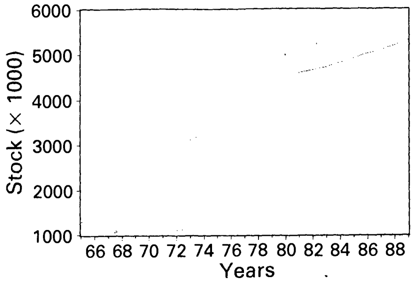

FIG. 2. *The smoothed stock of cars in The Netherlands, 1965–1989 (×1000).*

TABLE 2. *Smoothed estimated saturation level*

Year	Stock (× 1000)	Year	Stock (× 1000)
1968	5850	1979	6039
1969	5855	1980	6069
1970	5965	1981	6079
1971	6039	1982	6054
1972	6089	1983	5959
1973	6011	1984	5901
1974	5956	1985	5925
1975	5848	1986	5941
1976	5794	1987	5969
1977	5889	1988	5946
1978	6014	1989	5978

mean = 5962 standard deviation = 83

$\hat{\alpha}_{max} = 6089$ $\hat{\alpha}_{min} = 5794$

The estimated saturation level is smoothed by regressing it on a constant and on the levels when they are lagged by one and three periods. The fitted series is taken to be the smoothed saturation level series.

α are smoothed using a third-order subset autoregression, and they are displayed in Table 2 and Figure 3. It appears that these estimates are quite constant over the sample, and that the mean of the smoothed $\hat{\alpha}_t$ is 5962. This implies a saturation level of about 6 million cars in the Netherlands. The parameter estimates for α, β and γ can be used to calculate the fitted values for x_t, $\hat{x}_t$. Comparing these with the smoothed x_t observations shows that there are 12 negative and 13 positive forecasting errors, and that the corresponding R^2 is equal to 0.9998.

To calculate the forecasts of the level and also the corresponding error variances for the stock of cars for the period 1990–2010, I have chosen to use the bootstrap-technique. The results are displayed in Table 3 and Figure 3. It can be observed that the saturation level is almost reached in 2010.

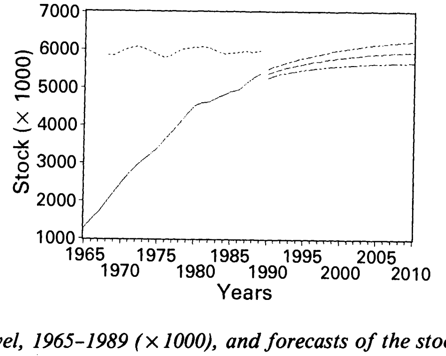

FIG. 3. *The estimated saturation level, 1965–1989 (×1000), and forecasts of the stock of cars 1990–2010 (×1000) with a two standard error forecasting interval.*

TABLE 3. *Forecasting the stock of cars, 1990-2010 (× 1000), with standard deviation of the forecast errors, both obtained using a bootstrap technique (200 replications)*

Year	Forecast of stock	Standard deviation of forecast error
1990	5383	61
1991	5443	63
1992	5499	67
1993	5549	71
1994	5595	75
1995	5637	80
1996	5675	85
1997	5709	91
1998	5741	96
1999	5769	101
2000	5795	106
2001	5818	110
2002	5840	115
2003	5859	119
2004	5877	123
2005	5892	127
2006	5907	131
2007	5920	135
2008	5932	138
2009	5942	141
2010	5952	144

CONCLUDING REMARKS

In this paper I propose a simple Gompertz curve fitting procedure which does not face the problems of current fitting methods such as the effects of the non-stationary behaviour of an empirical time series and the assumptions on the value of the saturation level. The advantages are that estimates of this level can yield insights in the stability of the saturation level over the sample period, and that knowledge of its value is not necessary for forecasting. Furthermore, the other two parameters can be estimated via standard methods as non-linear least squares. A simple bootstrap-technique can be used to calculate forecast intervals. An application to forecasting the stock of cars in the Netherlands illustrates its merits.

A possible drawback of the procedure is that it assumes that the growth of the process is always positive. For practical series this may, however, not be the case for a small number of data points, even when a Gompertz curve may be useful. One way to circumvent this problem is to delete the observations that invalidate this assumption, and to estimate the model for the remaining observations. Unreported experiences with the data sets given in Harrison and Pearce[1] indicate the usefulness of this approach.

Acknowledgements—The author wishes to thank Ronald Bewley, Nigel Meade, Marius Ooms, participants at seminars in Rotterdam and Perth, and two anonymous referees for helpful comments.

REFERENCES

1. P. J. HARRISON and S. F. PEARCE (1972) The use of trend curves as an aid to market forecasting. *Ind. Mktg Mgmt* **2**, 149-170.
2. C. MAR MOLINERO (1980) Tractors in Spain: a logistic analysis. *J. Opl Res. Soc.* **31**, 141-152.
3. N. MEADE (1984) The use of growth curves in forecasting market development—A review and appraisal, *J. Forecasting* **3**, 429-451.
4. N. MEADE (1985) Forecasting using growth curves—An adaptive approach. *J. Opl Res. Soc.* **36**, 1103-1115.
5. R. M. OLIVER (1987) A Bayesian model to predict saturation and logistic growth. *J. Opl Res. Soc.* **38**, 49-56.
6. C. W. J. GRANGER and P. NEWBOLD (1986) *Forecasting Economic Time Series*. Academic Press, San Diego.
7. A. C. HARVEY (1984) Time series forecasting based on the logistic curve. *J. Opl Res. Soc.* **35**, 641-646.
8. R. BEWLEY and D. C. FIEBIG (1988) A flexible logistic growth model with applications in telecommunications. *Int. J. Forecasting* **4**, 177-192.