Innovation in the Computer Industry

T. MODIS and A. DEBECKER

ABSTRACT

A learning-curve approach has been employed to study inventiveness in the computer industry. The appearance of new models as well as the appearance of new computer manufacturers have been following typical logistic S-curves over the last 26 years of computer history. Innovation seems to come in S-waves, and even though many of today's major manufacturers are close to exhausting their maximum innovation potential in their present S-wave, the overall computer market is rather "young." In contrast, the personal computer market is already beyond its maturity phase. Remarkably invariant over the 26 years considered, and through the complete range of computer sizes, remains the fact that for every five new computer models appearing on the market, there is one new computer company emerging also.

I. Introduction

A learning process, like any other growing process, typically follows an S-curve logistic pattern with time. That is, it starts slowly, then progressively increases in rapidity to develop into an exponential growth. Later on, the rate slows down to asymptotically approach a maximum value. The constraining maximum carries the hallmark of capacity (potential), while the evolution of the process is subject to the dynamics of competition. Learning a game is limited by the complexity of the game. Improving further on tick-tack-toe becomes soon impossible, while learning to better play chess can go for a long time before it slows down. And when this happens, it may be so because of the learner's limitations—exhausting the niche of his intellectual capacity—rather than the number of different configurations permitted by the game. In either case, however, it will happen logistically because it will be a niche being consumed to completion.

We propose here to look at innovation in the computer industry as a learning process. As in the games mentioned earlier, there must be a finite number of reasonably distinguishable computer models that can be put forth from a certain set of technological breakthroughs. Alternatively, on the analogy of the chess game, the number of new computer models realized may not be limited by the set of breakthroughs. We may allow for an almost continuous stream of small successive breakthroughs (one can further visualize these as a chain of smaller S-curves), in which case the number of models put

THEODORE MODIS is an experimental high energy physicist who has carried out research experiments for over 15 years at CERN, in Geneva, Switzerland; at Brookhaven National Laboratory, Upton, New York; and at Nevis Laboratories, Irvington-on-Hudson, New York. He has also been on the teaching faculties of Columbia and Geneva Universities. ALAIN DEBECKER is a mathematician and has been teaching in various technical schools in France and Switzerland since 1982.

Address reprint requests to Mr. Theodore Modis, Digital Equipment Corporation, Mgt. Science Dept., 65, chemin de'Etang, 1219 Chatelaine, Geneva, Switzerland.

out may just be limited by the ability of the organization to fill up its own niche, e.g., exhausting the ingenuity of its engineers who happen to have been all trained in a particular school of thought, etc. In either case, at any given time some niche will be in the process of being consumed.

The usefulness of the approach is in forecasting the remaining of a "well established" S-curve. Once the filling up of a niche has followed a logistic path for the historical range available, forecasting becomes possible for the remaining of the curve. Under-performing is "unnatural"—no niche was ever left partially filled in nature for natural reasons. Over-performing can only be envisaged with a subsequent S-curve whose characteristics we do not know and which will start at the "end" of this one. Thus, there is a window for forecasts.

If one considers that in the computer industry, the historical periods treated here cover from 10 to 25 years, the forecasting horizon is a rather long-term one, reaching the year 2000. We gained confidence in our endeavor to describe innovation in computers as a logistically growing process from a twofold success. The first success is a good logistic description of the historical data for the industry's most representative manufacturers. The second is a successful forecast when considering history to stop 5 years ago. Moreover, in the process, we stumbled across an idiosyncrasy of the computer market that can, in itself, be used as a forecasting tool—a ratio between new models and new manufacturers that is invariant with time over the whole range of computer sizes.

The historical data we used were drawn from the International Data Corporation (IDC) *1985 Processor Data Book* covering the period from the beginning of 1958 up to the end 1984. All computer models other than personal computers, including different versions, whenever reported separately by IDC, are considered. We treated personal computers separately in Section IV because we felt that they were probably belonging to a niche distinctly different from other computers. Appearance date for a given model is taken as the first installation date reported. Models reported with zero installations have been left out.

II. The Model

The idea to study markets with growth curves is not new [1–3]. The diffusion of innovation itself has been described by growth curves in several publications [4,5]. In fact, there have been attempts to further develop the use of growth curves. To just give two examples, the dynamic potential model by Mahajan and Peterson [6] allows for a nonconstant ceiling, and Meade's adaptive radiation [7] parametrises market development as a function of past observations rather than time. In spite of the fact that these approaches may have improved the applicability of S-curves in certain cases, which justifies the undertaking of the cited works, we decided to use the simplest possible logistic version. Physicists have painstakingly found out that the simpler a law, the more fundamental it is and the broader the spectrum of application it covers. But most important, being simple has the additional advantage of high probability of usage. The correlation, if any, is yet to be established between successful entrepreneurship and the use of elaborate formulations of sophisticated models and forecasting procedures produced abundantly in academic economic institutions.

The logistic form we are employing is the same as that of Fisher and Pry [4] and Marchetti [5], which is based on the simple fact that the fractional rate at which a niche is being filled (or emptied) is proportional to the remaining section of the original niche. In mathematical terms, if P is the population at time t:

$$\frac{1}{P}\frac{dP}{dt} = (\text{Max} - P) \qquad (1)$$

or

$$\ln\left(\frac{P}{\text{Max} - P}\right) = \alpha(t - t_0) \qquad (2)$$

where Max, the original niche size (capacity), α, and t_0 are constants.

The solution of the differential equation (1) is an S-curve, while the logarithmic representation of the ratio "Actual"/"Remaining" is a straight line.

Equation (2) is employed in a three-parameter fit via an χ^2 minimization technique to determine Max, α, and t_0. If we suspect that data are missing in the early part of the S-curve, a fourth parameter C is introduced, and equation (2) becomes

$$\ln\left(\frac{P - C}{\text{Max} - (P - C)}\right) = \alpha(t - t_0) \qquad (3)$$

The χ^2 minimization process has now to provide determinations for the four parameters Max, α, t_0, and C.

Deviations from the logistic path are expected at the low-end, i.e., when P is only a few percent of Max. We identify this region as the "infant mortality" region where competition dynamics are not yet naturally established.

A weakness of the model is that fits, based on data limited to the range where P has not reached 50% of Max yet, carry a large uncertainty in the determination of Max, since over this range an S-curve differs minimally from a simple exponential. Furthermore, poor quality data, which are often the case in industrial reporting, will introduce additional uncertainties.

It is for that reason that we undertook an extensive simulation study to be published elsewhere, on the uncertainties involved in S-curve logistic fitting, which enables quantitative determination of the expected errors on Max as a function of data errors and range covered. A rule-of-thumb general result is that, given at least half of the S-curve and a precision of better than 10% on each historical point, the uncertainty on Max will be smaller than 20% with 90% confidence level [8].

Finally, the niche definition is of crucial importance. If the members considered do not belong to the same niche, their collective population will not follow a logistic pattern with time. Sometimes subdividing data sets according to computer families, e.g., VAXes, Cybers, etc., can provide better fits to S-curves than an overall fit. At the same time, a niche within a niche, causing an S-curve to lead into another one with bigger Max and different α and t_0—or a larger niche being made up of several smaller niches—are altogether situations that can be envisaged and that we encountered in the case studies that follow (companies D and A respectively).

III. Case Studies

Nine major computer manufacturers are analyzed in detail. They are all well described by the model. They altogether represent about 80% of the computer industry. Results are reported here in terms of Figures 1 and 2 and Table 1 on five of those cases, which bring out interesting aspects of the logistic pattern. For one case, that of company B, the
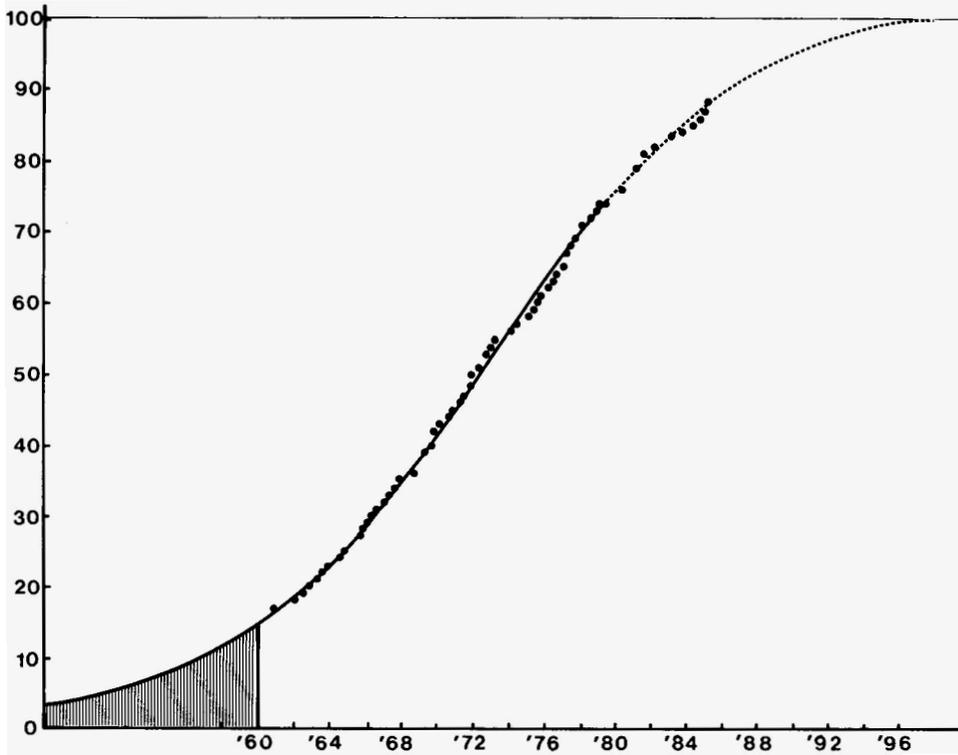
Fig. 1(a). Data points and logistic fit for the company described in "Case Studies" (company E). The fit is on data points up to the end of 1979. A 4-parameter fit revealed 14 models missing from the beginning of the historical period.
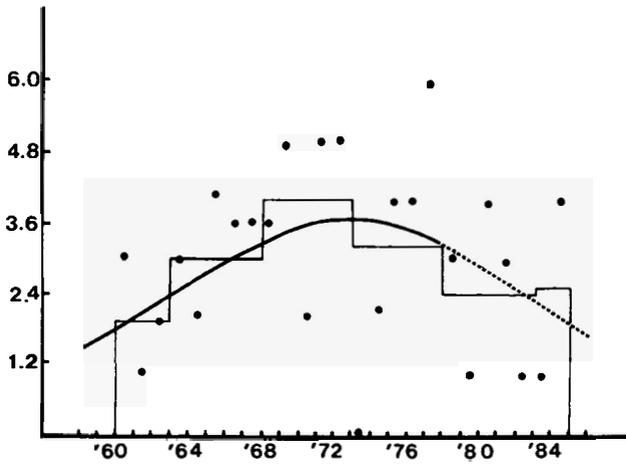

Fig. 1(b). Life-cycle plot showing yearly rates of new models (dark circles) and 5-year averages (histogram). The smooth line is the derivative of the S-curve of Figure 1(a).

TABLE 1
Fit Parameters for the Five Companies

Company	Max. No. of Models	1985 Position (% of Max.)	1985 Yearly Growth Rate	Appearance 1st Model	Max. Yearly Rate at Date (% of Max.)
A	76 ± 8	39 (51%)	3.4	12/1959	3.4 (4.4%) 1984
B1	27 ± 1.5	25 (93%)	0.6	10/1964	1.5 (5.4%) 1975
B2	13 ± 1.5	9 (70%)	1.4	5/1978	1.5(11.6%) 1984
C	87 ± 18	25 (29%)	3.6	11/1966	4.3 (4.9%) 1989
D	231 ± 45	62 (27%)	3.8	8/1960	4.8 (2.1%) 1996
E	101 ± 2	88 (87%)	1.5	10/1960	3.6 (3.5%) 1973

For company E, 14 units have been added, which were missing from the early part (see Figure 1).

data have been split into two families, thus achieving better logistic descriptions. The uncertainties quoted are based on look-up tables in our simulation study mentioned above, if one assumes a conservative 5% error per historical data point and asks for a confidence level of 90%. Here below, one case is presented in detail. It concerns a computer manufacturer with substantial market share and long tradition in the industry. The logistic fit we obtained was significantly improved when the fourth parameter C was allowed to vary. Its determination implied a missing section of about 14 early models. The data points fit well into the logistic (Figure 1[a]), giving a final maximum of 103 models with a present rate of 1.5 new models per year. The peak of innovation was back in 1973 when 3.6 new models per year were being produced. Presently this manufacturer is at 85% of this wave's maximum potential.

In order to test the forecasting power of the logistic description, we tried to progressively reduce the historical period. The S-curve superimposed to the data is determined by using historical data only up to the end of 1979. The parameters determined in this interval are within 1% of the parameters determined by the full set of data. The implications, however, are much more striking when we look at the life-cycle plot, shown in Figure 1(b). Here we are plotting the number of new models per year, and the points (dots) fluctuate to such an extent that no exponential smoothing or other time-series approach could possibly forecast the downward trend, much less the exact turnover point. Averaging over 5 years at a time, the thin histogram line already shows a better agreement with the smooth curve. The precision of the determination of the turning point, however, rests solely with the logistic fit on the cumulative data.

In this light, the forecasts made are expected to agree well with the long-term cumulative number of new models. When one forecasts the number of new models per year, fluctuations around the predicted rate, as large as the ones shown on Figure 1(b), should not be surprising.

For each case study, a logistic graph is given in the logarithmic representation whereby the suitability of the model can be qualitatively judged by the quality of the straight line character. For the sake of easy visual interpretation, the vertical scale is calibrated to directly display the actual population, $P(t)$. As mentioned earlier, deviations from the straight line in the very early stages, a few percent of Max, are to be expected. Straight lines have been superimposed to guide the eye. All graphs depict the innovation potential realized, i.e., cumulative number of different models, as a function of time. The units are percentages of maximum innovation potential and calendar date. Table 1 lists the perceived innovation potential, i.e., the maximum number of different computer models expected in this S-wave, the present position with respect to this maximum, and
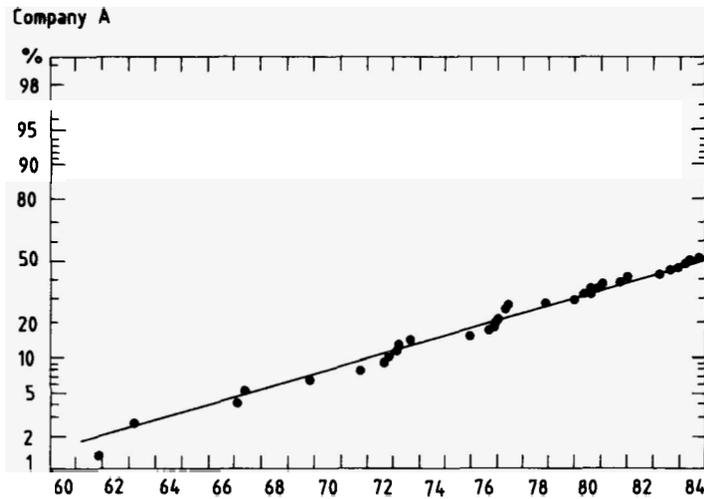
Fig. 2(a)–to 2(d). Logistic patterns for companies A, B, C, and D, respectively, in the logarithmic representation. What is plotted is the population of new models expressed in percentage of maximum. Company B has been split in the two families of products that it produces (see text). Straight lines have been superimposed to guide the eye.

the maximum and present rates of growth. The date of the first model in the S-wave is also given. This is the installation date of the first model as reported by IDC. The company described in detail is included in Table 1 as company E.

For company A, in Figure 2(a), two clusters of innovative activities around 1972–1973 and 1976–1977 could be analyzed in terms of shorter logistic curves, the case of niche-within-a-niche mentioned earlier, and possibly correlated to technological or other breakthroughs. For company D as shown in Figure 2(d), the straight-line logistic description of the data points is excellent for the historical period starting in January 1966 and continuing to date. Earlier data is amenable to different logistic curves, which may be interpreted as cases of small early successive niches.
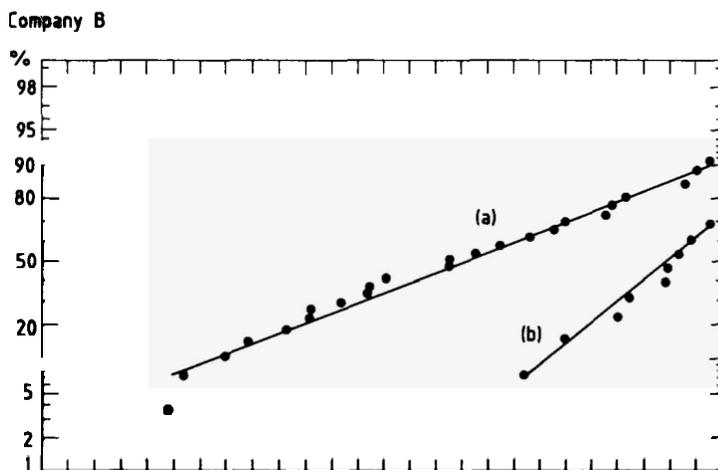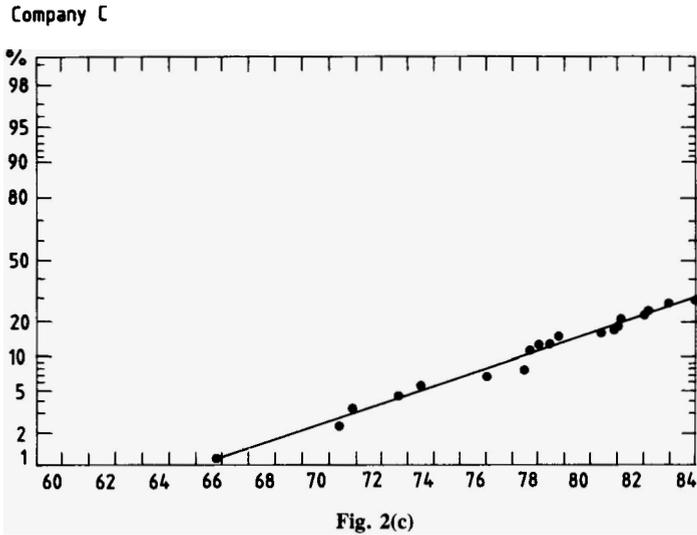

Fig. 2(b)

Company C

Fig. 2(c)

IV. The Overall Computer Market

The overall computer market was looked at by considering all new models of all manufacturers together. Figure 3 incorporates many more manufacturers than the nine studied in detail—947 models enter in the graph of Figure 3. In Figure 4 the reasoning applied to models is extended to manufacturers, i.e., a logistic growth hypothesis. In the historical period covered, i.e., since 1958, there have been 185 new manufacturers making an entry. Both graphs show a remarkable straight-line logistic character; there is no need to superimpose straight lines to guide the eye. Table 2 summarizes the parameters for which the best logistic paths are achieved.

It is of interest to point out that the computer market looked at from this viewpoint is a rather young market occupying—by the end of 1984—the lower third of its maximum potential both in terms of new models and new manufacturers. The maximum rate of

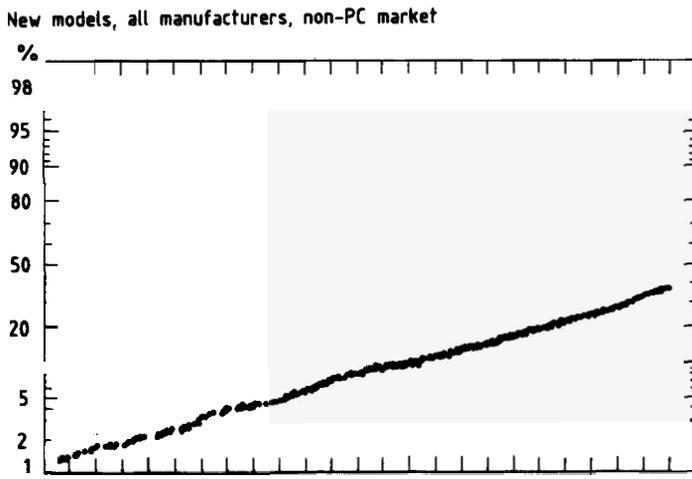New models, all manufacturers, non-PC market
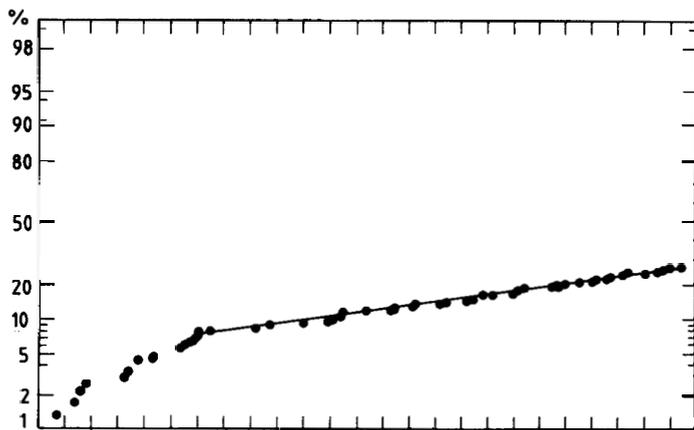
Fig. 3. Populations of new models expressed in percentage of maximum in the non-PC market.
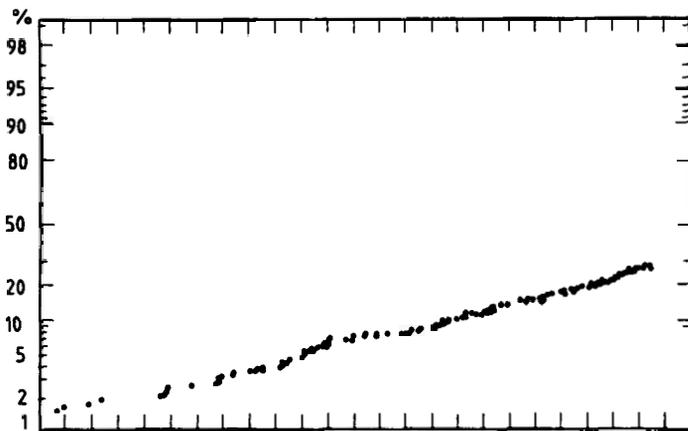

TABLE 2
The Whole Computer Market (PCs Excluded)

	Maximum Potential	1985 Position (% of Max.)	1985 Yearly Growth Rate	Max. Yearly Rate at Date (% of Max.)
New Models	2800 ± 350	947 (33%)	102	112 (4.0%) 1988
New Manufac.	696 ± 145	185 (27%)	19	26 (3.7%) 1980

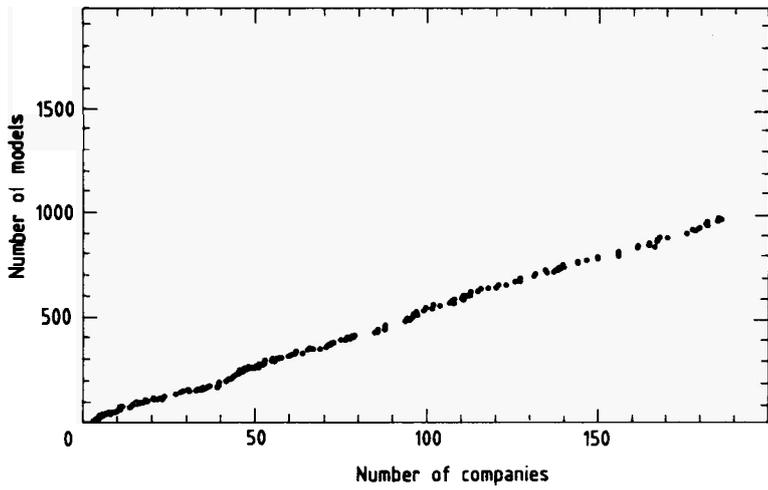First model (and first manufacturer) July, 1958.

Fig. 5. New models versus new manufacturers in the non-PC market. Linear scales.

productivity is around 1989, when we will see about 112 new models and 26 new companies appearing in a year.

The remarkably logistic way in which new models and new manufacturers evolved in time suggested that a strong correlation may exist between the two. Figure 5 depicts these two variables plotted one against the other in terms of raw data, with linear scales,
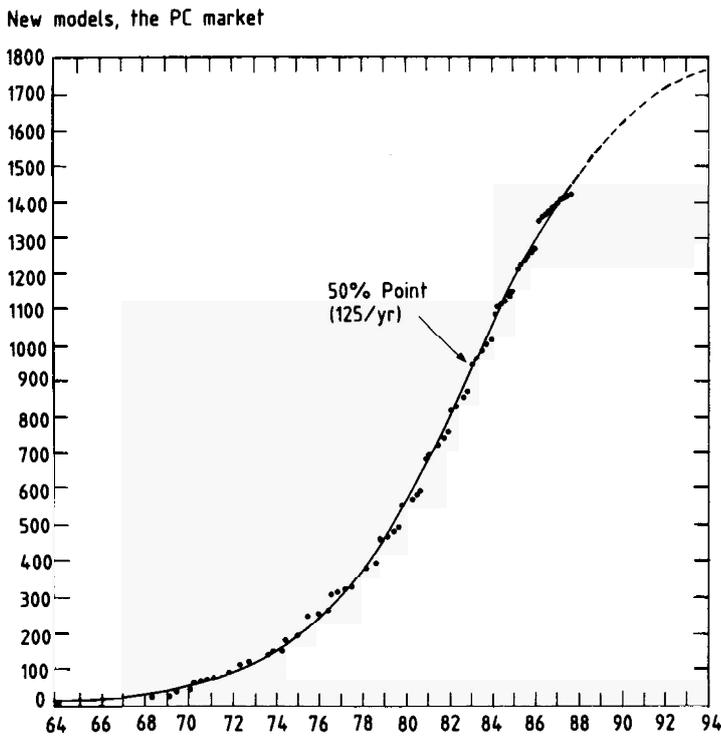

Fig. 6. Logistic growth (linear scales) for the appearance of new models in the PC market. The solid line is a 3-parameter fit.
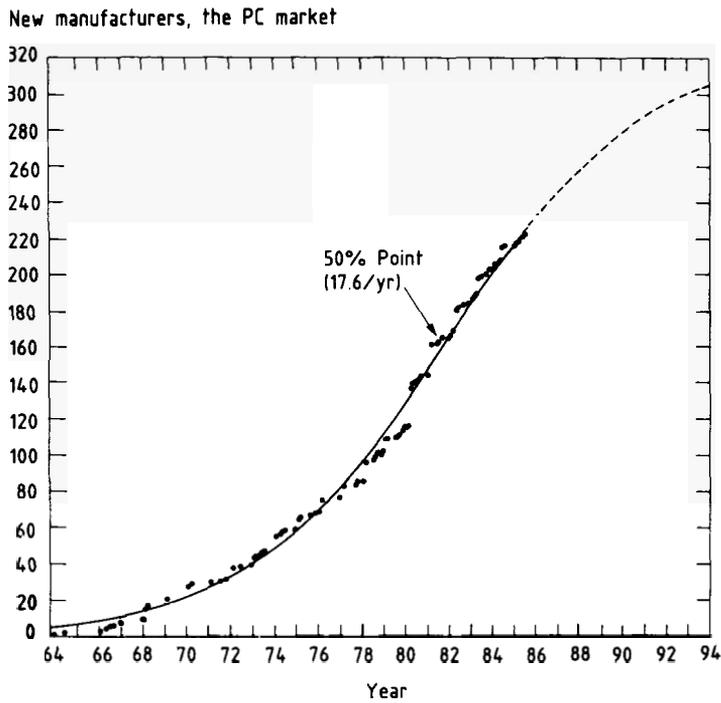
New manufacturers, the PC market

Fig. 7. Logistic growth (linear scales) for the appearance of new PC manufacturers. The line is a 3-parameter fit.

no fits, and no other manipulations whatsoever. The 947 points entering in this graph follow a strictly obeyed linear relationship over 25 years. The emerging invariant, the slope, is 5.1, which implies that on the average, there will be a new computer manufacturer emerging for every five new models that appear on the market.

This strict correlation makes us believe that there should be no difference in the realization rates between these two variables. Indeed, within the uncertainties quoted in Table 2, the ratio 5.1 seems to be well respected far into the future.

V. The PC Market

Personal computers (PC) were grouped all together and were treated as one species growing logistically. As before, PC manufacturers were treated similarly. The results are presented in linear scales since both niches are penetrated sufficiently to visually show the familiar bends of the logistic curve. Figures 6 and 7 show the fits for PCs and PC

TABLE 3
The PC Market

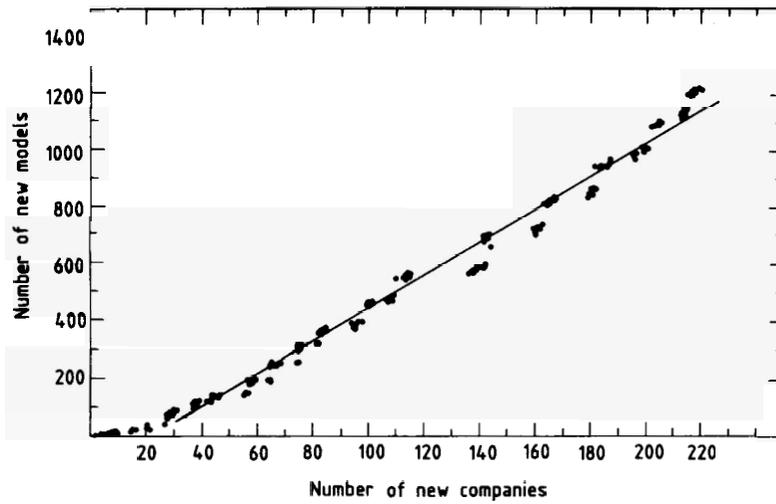|	Maximum Potential	1987 Position (% of Max.)	1987 Yearly Growth Rate	2000 Yearly Growth Rate
PCs	1864 ± 90	1377 (74%)	82	5.6
PC Manufac.	321 ± 15	234 (75%)	14	1.3

Fig. 8. New models versus new manufacturers in the PC market. Linear scales.

producers respectively. Table 3 gives the numerical results. Both curves seem to be about 75% exhausted by 1987. The proliferation of models and companies is astounding. At their peak, around 1982–1983, there were 125 distinguishable new PCs per year introduced by almost 18 new companies per year. In 1987 these rates have dropped to 82 new models for 14 companies. By the year 2000, less than six new models and less than two producers per year are expected. Plotting models versus manufacturers, again we can see (Figure 8) a straight line emerge with a slope of 5.7.

VI. Conclusions

Innovation in the computer industry has been treated as a logistic growth process, and results are presented in detail for five major computer manufacturers and their models as well as for personal computers.

The case studies presented reflect three different categories of manufacturers. We could call them "young," "middle-aged," and "old" as far as innovation is concerned, although this may be little related to their actual age of existence on the market. The "old" ones are like company E, as they have already achieved almost 90% of their potential. Among the "young" ones, we find cases such as company C, well below its mid-point, and consequently still growing at an increasing rate. In the "middle-aged" category, cases like that of company A may be found, which create new models at a maximum rate at present.

The most spectacular, diametrically opposed cases, however, are those of company D and B, both very successful on the market place. D has by far the highest potential (231 different models), with the largest maximum rate of growth 4.8 models per year to be achieved by 1996. This rate of growth represents the smallest realization rate (2.1%) of max. per year. Moreover, in spite of its longest history, D definitely belongs to the "young" category, with 27% present penetration, and it will probably remain among the "young" for a long time because of its low realization rate.

Company B, on the other hand, is opposite to D in every respect—smallest potential, 27 and 13 models in its two lines of products, with the smallest maximum rate of growth

at 1.5 models per year. One family is already exhausted, and the other, being realized at a rate of 11.6% of max. per year, will also be exhausted very quickly; four more models are expected before the line is finished. Company B seems to be most in need to start a new S-wave of innovation.

All of the above observations refer to these companies' present wave. New waves based on new technologies or other conceptual breakthroughs are most certainly in the making, and we are unable to forecast these developments now. These developments are certain to come since we see with confidence in Section IV that the overall computer market is far from exhausting its innovation potential both in terms of new models and in terms of new manufacturers. The PC market, on the other hand, seems to be already into its declining phase. Only one or two companies will manage to enter the market with a handful of new models during the year 2000. These are sombering numbers compared to the over 100 models and over 15 new producers per year in the early 1980s.

In Section IV we see that we must expect one new manufacturer for every fifth new model appearing on the market, a result that was corroborated for the PC market in Section V. This is more than amusing. It indicates that the expected number of different models per manufacturer should not be far from five! How, then, are we to understand prolific giants with hundreds of different models to their credit?

To answer this question, we should look closer at what happens around these gigantic citadels of computer know-how. Successful models, concepts, and/or engineers them- selves migrate and branch off to give rise to small companies that are often based around a single successful idea or a single person. In other words, the mother company, through its technological brewing, produces not only new models but also new companies. There will always be small companies coming and going, and never reaching five models. But the norm is five, and the more a large manufacturer surpasses this number, the more will there be a need for the appearance of small companies to counterbalance. The harmonious way, and perhaps the one of least resistance, is to simultaneously produce new models and new companies at the ratio of 5 to 1. In fact, a good rule for ambitious entrepreneurs who wish to set up new companies could be to aim either at quick gains to be achieved by their fifth model; or, if they intend to go beyond this point, to plan and invest in the necessary infrastructure, ensuring sufficient technological fertility to give birth not only to models but to daughter companies as well.

References

1. Fildes, R., Forecasting and Planning. R. Fildes and D. Wood, eds., Gower Press, Hampshire, 1987.
2. Balachandra, R., Perceived Usefulness of Technological Forecasting Techniques, *Technological Forecasting and Social Change,* 16, 155–156, (1980).
3. Meade, N., The Use of Growth Curves in Forecasting Market Development—A Review and Appraisal, *Journal of Forecasting,* 3, 75–88.
4. Fisher, J. C., and Pry R. H., A Simple Substitution Model of Technological Change, *Technological Forecasting and Social Change,* 3, 75–88 (1971).
5. Marchetti, C., On the Role of Science in the Postindustrial Society, *Technological Forecasting and Social Change,* 24, 197–206.
6. Mahajan, V., Peterson, R. A., Innovation Diffusion in a Dynamic Potential Adopter Population, *Management Science,* 24(15), 1589–1597, (1978).
7. Meade, N., Forecasting Using Growth Curves—An Adaptive Approach, *Journal of Operational Research Society,* 36(12), 1103–1115, (1985).
8. Debecker, A., and Modis, T., Determination of Uncertainties in S-Curve Logistic Fits. Presented at the 6th International Symposium on Forecasting, Paris, June 1986. In press.