Prometheus
Vol. 28, No. 1, March 2010, 5–13

RESEARCH PAPER

An inductive model of technological progress

John Enos* and Alison Etheridge

Magdalen College, Oxford, OX1 4AU, UK

The purposes of this paper are threefold: to suggest a conceptual scheme which encompasses innovations and improvements in industry; to express this scheme in the mathematical form of Poisson jump processes; and, finally, to illustrate it with a sequence of cost data drawn from three-quarters of a century of company operations. In the stochastic process, three parameters summarise the events: (i) random intervals between innovations; (ii) different random intervals between improvements; and (iii) the rate at which the benefits of improvements decay over time. The fit between the model's performance, on the one hand, and the cost data, on the other, is good: the theoretical deductions and the actual cost series bear a close resemblance. Most models of technological progress are based upon the mere passage of time, whereas this model stems from the events.

Definition of terms

A technology itself is a coherent and interrelated set of ideas, and, possibly, if one wants a broad definition, the social, political and economic institutions within which the operation is performed. In the case of industrial innovations, the set of ideas comprising the technology is based upon fundamental physical and chemical principles; a different set of ideas, based upon different physical and chemical principles, would constitute a different technology.

We divide technical change, the motive force, into two elements: by the first – innovation – we mean the invention and initial development of a new technology, terminating when it (the technology) is incorporated into a set of artefacts whose operation yields useful products. Assembling a proper collection of artefacts and integrating them so as to reflect, in the form of a physical analogue, the principles inherent in the technology may well yield an innovation; but unless the assembly yields either an entirely novel product, or a familiar product more cheaply, progress is halted at the stage of the innovation. Historically, there have been examples of technologies realised as innovations, but not passing the test of economic profitability. It is our wish to note the possibility of technological success followed by economic failure that leads us to make the distinction.

The second of our terms is improvement, meaning reduction in the amount of resources needed to carry out its application. Although, as we shall see, an improvement may have a major effect economically, in the sense of reducing substantially the cost of applying the technology, its technological effect is not so grand as to alter the basic set of ideas or the conjunction of artefacts within which these ideas are made manifest.

*Corresponding author. Email: john.enos@oxfordenergy.org

ISSN 0810-9028 print/ISSN 1470-1030 online

DOI: 10.1080/08109020903554294
http://www.informaworld.com

Innovations and improvements in petroleum refining

The statistical evidence that provides the basis for our inductive model of growth is derived from the four innovations and sets of improvements that, since their origin, have formed the heart of the major petroleum refining processes; these processes are collectively called 'cracking'. (In the cracking process, the large hydrocarbon molecules comprising heavy petroleum oils, useful in their crude form as industrial fuels only, are split or 'cracked' into smaller molecules, useful chiefly as motor gasoline or diesel fuel.) The first cracking innovation occurred in 1911/2, and improvements to the fourth cracking innovation have continued up to the last year of our inquiry, 2000. Our statistics cover the period from the initial date until 1987; historically they represent an uninterrupted sequence of 75 years.

The data relevant to this paper appear in Table 1, recording the intervals of time elapsing between innovations (four in total, occurring in 1911/12, 1921, 1936 and 1942), and between the improvements succeeding each innovation (20 all told, occurring in the intervals between innovations). Table 2 records the successive reductions in costs of manufacturing motor gasoline by cracking as consequences of the application of the innovations and improvements.

There are three conclusions that we wish to draw from the data in Tables 1 and 2. The first is that the sizes of the time intervals between innovations and improvements, in Table 1, and of the cost reductions following the innovations (i.e., the data in all rows but the first), in Table 2, appear to be random. Successive intervals between improvements are read down the columns in Table 1, and successive reductions in cost down the columns of Table 2.

The second conclusion that we draw from the data is that the appearance of an innovation does not immediately lower the cost of manufacture. This conclusion is evident in the first row of numbers in Table 2, the first of which is irrelevant (since the technology – the cracking of gas oil at elevated temperature and pressure for gasoline – had never been attempted before), and the next three of which are zero (indicating that the profitability of operating the plant that incorporated the innovation was the same as that of operating the most efficient contemporary plant incorporating the previous technology).

The third conclusion is that, within a given technology, successive improvements are, on average, of diminishing importance in terms of reducing costs. This conclusion is supported by the averages of the individual cost data reported in the final column of Table 2, and is consistent with the arguments of Enos (1958) and Rosenberg (1982)

Table 1. Time intervals in years between improvements, by regime, 1911/12 to 1987

Regime	I	II	III	IV	
Dates	1911/2–1921	1921–1936	1936–1942	1942–1987	Average
End previous regime to innovation	1	1	0.8	1.1	1
Innovation to first improvement	1	1	3	12	4
First improvement to second	1	3	0.5	5	2.4
Second improvement to third	0.25	5	2	4	2.8
Third improvement to fourth	2	5	–	13	6.7
Fourth improvement to fifth	–	–	–	10	10

Sources: Enos (1962), particularly Table 1, p. 42; and Enos (2002), Tables 8.1 and 8.6a.
Improvements for Regimes II and III are calculated from engineering reports and refinery records.

Table 2. Cost reductions between regimes and between improvements, 1911/2 to 1987

Regime	I	II	III	IV	
Dates	1911/2–1921	1921–1936	1936–1942	1942–1987	Average
End previous regime to innovation	13(a)	0	0	0	0
Innovation to first improvement	17	31	17	22	22
First improvement to second	8	4	29	22	16
Second improvement to third	8	2	19	25	13.5
Third improvement to fourth	2	20	10	13.5	11
Fourth improvement to fifth	12	11	–	6	9.7
Fifth improvement to sixth	–	–	–	0.83(b)	0.8

Notes: (a) The raw material was not used, and the main product was not produced, before the innovation; (b) average of several simultaneous improvements, estimated from contemporaneous records.
Sources: As for Table 1

that, with the passage of time and effort, the further potential for improvement of a technology diminishes. At the time of its introduction, any technology has a potential for improvement, the most significant part of which is quickly exploited, usually by the innovators themselves (Tyre, 1992). As time passes, less and less significant parts of the potential remain to be exploited, yielding smaller and smaller increments of gain. The individual observations of reductions in cost may seem to be random, but the general trend, as shown by the averages in the last column of Table 2, indicate a steady decline over the life of a technology.

The model

The inductive model detailed in Appendix 1 is the simplest that reflects these three observations. The times of innovations and the times of improvements will each be modelled as a Poisson process. An innovation in our model will not be accompanied immediately by an increase in productivity. This reflects the equality of costs of operating the most efficient plant under the previous technology, and the first plant using the new technology. The logarithm of the increase in productivity at the time of an improvement will be modelled as a random variable (independently for each improvement) that we deflate exponentially with the time since the most recent innovation.

Through time, the trajectory exhibited by the model follows a staggering path. If the measure of advance is the increase in the productivity with which resources are applied (in our case, the inverse of the reductions in cost), and if one starts on the path at the time of the first innovation, the pace will be fastest (that is, the jump will be biggest) soon afterwards, when the first improvement occurs. Thereafter, there will be a random series of improvements, yielding random increases in productivity. In the calculations in Appendix 1 we calculate the expected increase in the logarithm of the productivity over any given time interval. The results are equivalent to those obtained by approximating this by a smooth progression of improvements (by substituting the expected value of the random variable that dictates the size of the improvement). Interrupting the progression of improvements will be the next innovation, whose plant exhibits a productivity exactly equal to the productivity attained by the improved version of the plant incorporating the technology associated with the previous innovation – in this case, the initial innovation that began the journey along the path. In Figure 1 the vertical axis represents some measure of productivity with which resources are used; the horizontal axis represents some measure of the course of events (usually the passage of time). The three phenomena captured by the model can be seen in any graph of productivity: first, the irregular appearance of new technologies (the timing of the leaps from a previous to a subsequent regime) and of improvements (the timing of the upward steps in productivity); secondly, the random height of successive steps (improvements) within a given regime; and, thirdly, the equal productivities of the plant providing the last observation of the previous regime and of the plant providing the first observation of its successor (the horizontality of leaps; 'long jumps', not 'high jumps'). This is illustrated by Figure 1, where we have plotted increases in productivity in the petroleum cracking industry for the first two periods 1912 to 1936.

Once suitable simplifications have been made, the distribution of the stochastic model of innovations and improvements can be determined. The improvements occur at random intervals, approximated by a Poisson process with the same average rate of occurrence; similarly innovations occur at the Poisson rate determined by their average rate of occurrence. The increases in productivity at successive improvements are modelled by a sequence of independent random variables whose diminishing size is modelled by a negative exponential function of whatever measure of the course of events (say, the passage of time) has been selected. Once values for these parameters have been chosen, the overall sequence of events is specified.

Summarising the experience in petroleum refining

In mathematics, the literature that attracts us is that of stochastic processes, particularly Markov processes. Within Markov processes, a natural class to focus upon is

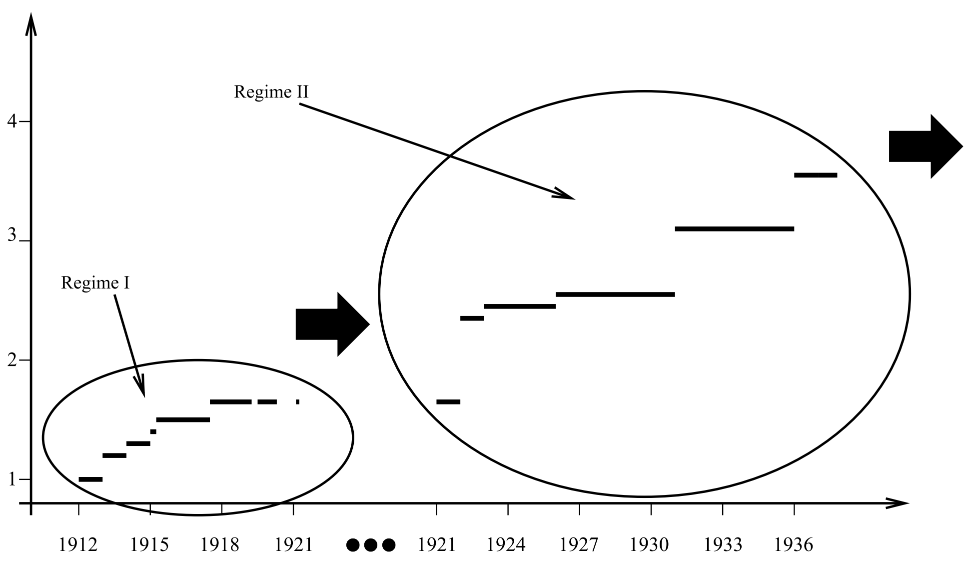

Figure 1. Evolution of productivity through innovations and improvements in the petroleum cracking industry for the period 1912 to 1936

those which exhibit 'jumps' from one regime to another. Such processes, usually called 'piecewise deterministic processes' (Davis, 1984) or 'jump Markov processes' arise naturally in many applications (such as dam theory, queueing theory, etc). They have been studied extensively in the area of stochastic control, in which a process is controlled so as to achieve a stated objective (see Vermes (1985) and Davis (1986) for a comprehensive introduction). A systematic mathematical study of the processes can be found in Jacod and Skorohod (1996). The statistical data upon which we draw for this paper indicate that the underlying industrial activities can be reasonably approximated by such a jump Markov process. Here we focus on elementary aspects of the process so that the necessary mathematical techniques are more than adequately covered by elementary texts, such as Grimmett and Stirzaker (2001).

Given our mathematical formulation, in order to summarise the experience of the petroleum refining industry in carrying out its innovations and improvements, we need merely to specify the date at which the statistical process commenced and the average values of four parameters: the average rate of innovation, the average rate of improvement, the average size of improvement if it were to occur immediately after an innovation, and the rate at which the expected value of subsequent improvements diminishes in size. In the case of cracking in petroleum refining, the date of the first innovation was 1911/1912; it is this date on which the model's trajectory takes off. The values of two of the four parameters, in averages per year, are estimated from the data in Tables 1 and 2: there were four innovations over the entire period of 75 years, yielding an average of one innovation for every 19 years; or, as a rate, 0.05 innovations per year. Over the same interval, there were 20 improvements, or, more accurately, when the multiple sets of improvements in the middle of the fourth regime are admitted, 19 single improvements and a single cluster of improvements. Counting the cluster as one improvement, the average rate of improvement was 0.26 per year.

The third and fourth parameters can also be calculated from the data in Tables 1 and 2. We model each improvement as an independent realisation of a positive random variable X multiplied by $\exp(-\rho\tau)$, where τ is the time since the most recent innovation and ρ is a positive constant. We make no assumptions about the distribution of X, as the only parameter that we need for our calculations is its mean. The parameter ρ is estimated by a linear regression of the logarithm of the increase in productivity against the time since the most recent innovation at each improvement. The value obtained this way is 4%. We then use this value of ρ to 'correct' the increases in productivity at each improvement and take the average to produce an estimate for the expected value of X. This yields a value of 22%, representing the expected size of the first improvement occurring immediately after the innovation. This corresponds to an expected value of 21% if the improvement occurs one year after the innovation.

The final equation of the mathematical model in Appendix 1 (Equation (1)) specifies the average rate of increase in productivity predicted by the stochastic process generating innovations and improvements. As one would expect, this predicted rate increases if the rate of improvements or the rate of innovations increases. It decreases if we increase the parameter ρ that measures the decline in the size of improvements with time since the most recent innovation. In Table 3, we compare estimates from this equation with true values at four dates: 10 years, 25 years, 50 years and 75 years after the beginning of the process.

Table 3. Comparison of predicted and actual rates of increase of productivity

Length of time interval	Predicted rate of increase	Actual rate of increase
10 years	0.05	0.05
25 years	0.04	0.05
50 years	0.04	0.06
75 years	0.04	0.04

Notes: This table is reported to only one significant figure and so disguises a systematic downwards bias in our estimates. In fact, in these data, the final regime appears to be of a somewhat different nature to previous regimes. This is not, of course, reflected in our model, which treats all regimes as having the same (random) description.

Conclusions

The correspondence between the theoretical and actual data on innovations and improvements in petroleum refining is surprisingly good (although the figures in Table 3 are only reported to one significant figure, a more careful analysis would reveal a systematic underestimate of the true rate of increase in productivity). What is clear is that our model does capture the essential features of the data.

It would be interesting to see if the development and application of other industrial technologies followed the same trajectory as that of cracking in the petroleum refining industry, and, if so, if their parameter values – their summary statistics – were of similar orders of magnitude. However, even if economists are not to be graced with comparative data from other industries, they might settle for less expansive results derived from studies of smaller entities than an entire industry. Considering only petroleum refining, it might be possible to chart the technological history of another geographical area (besides the United States, where the 'best practice' provided us with our observations) or of a single company (rather than the universe of innovators and improvers). But these are, perhaps, matters for future investigations into the nature of technological change.

Acknowledgement

The first author wishes to thank Paul David for his comments.

References

Davis, M. (1984) 'Piecewise-deterministic Markov processes: a general class of non-diffusion stochastic models', *Journal of the Royal Statistical Society, Series B,* 46, pp.353–88.

Davis, M. (1986) 'Control of piecewise-deterministic processes via discrete time dynamic programming' in *Stochastic Differential Systems,* Lecture Notes in Control and Information Science, Springer, Berlin.

Enos, J. (1958) 'A measure of the rate of technological progress in the petroleum refining industry', *Journal of Industrial Economics,* June, pp.187–94.

Enos, J. (1962) *Petroleum Progress and Profits: A History of Process Innovation,* MIT Press, Cambridge, MA.

Enos, J. (2002) *Technical Progress and Profits: Process Improvements in Petroleum Refining,* Oxford University Press, Oxford.

Grimmett, G. and Stirzaker, D. (2001) *Probability and Random Processes,* 3rd edition, Oxford University Press, Oxford.

Jacod, J. and Skorohod, A. (1996) 'Jumping Markov processes', *Annales de l'Institut Henri Poincaré, Probabilités et Statistiques,* 32,1, pp.11–67.

Rosenberg, N. (1982) *Exploring the Black Box: Technology, Economics and History,* Cambridge University Press, Cambridge.

Tyre, M. (1992) 'Managing the introduction of new process technology: an international comparison' in Kochan, T. and Useem, M. (eds) *Transforming Organisations,* Oxford University Press, Oxford.

Vermes, D. (1985) 'Optimal control of piecewise deterministic Markov processes', *Stochastics,* 14, pp.165–208.

Appendix 1. The mathematical model

We model the increase in productivity as a stochastic process. This process exhibits jumps whenever an improvement takes place. The size of the jump depends on the time since the last innovation. We model the improvements as taking place according to a Poisson process of rate μ. Innovations take place at a rate λ. Suppose that an improvement takes place at a time t. Let us write $P(t–)$ for the productivity immediately before the improvement, then $P(t) = e^{\tau(\tau(t))}$ $P(t–)$, where $\tau(t)$ is the time since the most recent innovation and the random variable r is to be specified. We see from Table 2 that there is no obvious pattern to the size of improvements except that, on average, they decay as the time since the most recent innovation increases. We therefore model them as $r(s) = Xe^{-\rho s}$, where X is a positive random variable and ρ is a constant.

Let us denote the times of improvements by t_i. We shall calculate the expected value of $R(T) \triangleq \sum_{t_i<T} r(\tau(t_i))$. This then provides the estimate $P(T) \approx P(0)e^{R(T)}$ for large times T. It is worth making some remarks about our choice of model. First, it might seem more natural to model directly percentage increases at each improvement. However, the corresponding model, in which percentage increases are modelled in the same way as $r(s)$ above, does not have a closed-form expression for the expected productivity at time T, or indeed for any obvious function of it. The disadvantage of this approach is that, although over long time periods the error in the estimate for the average rate of increase in productivity may be a small percentage of the true value, the corresponding estimate for the actual increase in productivity can be large (because of the exponentiation). By estimating $R(T)$ instead of $e^{R(T)}$ we are introducing a systematic downward bias in our estimate for productivity. We have followed this route because it provides a closed-form solution.

Let us turn then to the calculation of the expected value of $R(T)$, denoted $\mathbb{E}[R(T)]$. We proceed in two stages:

Step 1: a single regime of length t

The first step is to consider just a single regime of length t. Let us write $R_0(t)$ for the sum $\triangleq \sum_{t_i<t} r(\tau(t_i))$, representing the logarithm of the increase in productivity over the regime and N_t for the number of improvements during the regime. In our model N_t has a Poisson distribution with parameter μt. If we condition on $N_t = n$, then the times of the improvements can be realised as n independent uniformly distributed random variables on $[0,t]$. Now if U is uniformly distributed on $[0,t]$, then

$$\mathbb{E}\,[r(\tau(U))] = \mathbb{E}\left[Xe^{-\rho U}\right] = \mathbb{E}\,[X]\frac{1}{\rho t}(1-e^{-\rho t}),$$

as, in our model, X is assumed to be independent of the time of the improvement. Since $\mathbb{E}\,[N_t] = \mu t$, we obtain

$$\mathbb{E}\,[R_0(t)] = \frac{\mu\mathbb{E}[X]}{\rho}(1 - e^{-\rho t}).$$

Step 2: multiple regimes

We now suppose that innovations occur according to a Poisson process with a rate λ. Suppose that M_T innovations occur during the time interval $[0,T]$. Notice that this results in $M_T + 1$ regimes and we denote the corresponding values of R_0 by $R_0^{(1)}, R_0^{(2)}, \ldots, R_0^{(M_T+1)}$. Then, conditioning on M_T, we obtain

$$\begin{aligned}
\mathbb{E}\,[R(T)] &= \sum_{n=0}^{\infty}\mathbb{E}\left[\sum_{i=1}^{n+1} R_0^{(i)} \middle| M_T = n\right]\mathbb{P}[M_T = n] \\
&= \sum_{n=0}^{\infty}(n+1)\,\mathbb{E}\left[R_0^{(1)} \middle| M_T = n\right]\mathbb{P}\,[M_T = n] \\
&= e^{-\lambda T}\,\frac{\mu\mathbb{E}[X]}{\rho}\left(1 - e^{-\rho T}\right) \\
&\quad + \sum_{n=1}^{\infty}(n+1)\int_0^T \frac{n}{T^n}(T-s)^{n-1}\,\frac{\mu\mathbb{E}[X]}{\rho}(1 - e^{-\rho s})\,ds\,\frac{(\lambda T)^n e^{-\lambda T}}{n!},
\end{aligned}$$

where we have used the fact that, conditional on $M_T = n$, the first innovation occurs at a time distributed as the minimum of n independent uniformly distributed random variables on $[0,T]$.

It is a tedious, but easy, exercise to evaluate this expectation. Note first that, for a positive parameter α,

$$\sum_{n=1}^{\infty}(n+1)\frac{\alpha^n e^{-\alpha}}{n!} = \alpha + 1 - e^{-\alpha}.$$

The first term in the summation in the last line above is given by this expression with $\alpha = \lambda T$. To evaluate the second term we make the substitution $s \to T - s$ in the integral, integrate by parts and then interchange the order of summation and integration to obtain

$$\begin{aligned}
&\sum_{n=1}^{\infty}(n+1)\int_0^T \frac{n(T-s)^{n-1}}{T^n}(e^{-\rho s})\,ds\,\frac{(\lambda T)^n e^{-\lambda T}}{n!} \\
&= e^{-(\lambda+\rho)T}\sum_{n=1}^{\infty}(n+1)\frac{(\lambda T)^n}{n!}\int_0^T \frac{n s^{n-1}}{T^n}e^{\rho s}\,ds \\
&= e^{-(\lambda+\rho)T}\sum_{n=1}^{\infty}(n+1)\frac{(\lambda)^n}{n!}\left(T^n e^{\rho T} - \int_0^T \rho s^n e^{\rho s}\,ds\right) \\
&= \sum_{n=1}^{\infty}(n+1)\frac{(\lambda T)^n e^{-\lambda T}}{n!} - \rho e^{-(\lambda+\rho)T}\int_0^T \sum_{n=1}^{\infty}\frac{(n+1)(\lambda s)^n e^{-\lambda s}}{n!}e^{(\lambda+\rho)s}\,ds \\
&= \lambda T + 1 - e^{-\lambda T} - \rho e^{-(\lambda+\rho)T}\int_0^T (\lambda s + 1 - e^{-\lambda s})e^{(\lambda+\rho)s}\,ds.
\end{aligned}$$

Another integration by parts yields the value of this final integral and then combining the above and rearranging yields

$$\mathbb{E}\,[R(T)] = \frac{\mathbb{E}[X]\mu}{\lambda+\rho}\left(\lambda T + \frac{\rho}{\lambda+\rho}\left(1-e^{-(\lambda+\rho)T}\right)\right).$$

In order to obtain a prediction for the rate of increase in productivity over a time interval of length T years beginning at the start of our data, we simply divide this by T to obtain

$$\mathbb{E}\left[\frac{R(T)}{T}\right] = \frac{\mathbb{E}[X]\mu}{T(\lambda+\rho)}\left(\lambda T + \frac{\rho}{\lambda+\rho}\left(1-e^{-(\lambda+\rho)T}\right)\right). \tag{1}$$

It is this prediction that appears in Table 3. Notice that as the time T increases to infinity this predicted rate converges to a constant, $\mathbb{E}[X]\lambda\mu/(\lambda+\rho)$.