Carcinogenic Effects of Chronic Exposure to Very Low Levels of Toxic Substances

by Richard Peto*

Dogma

The sort of situation that I wish to consider is contamination of the air around vinyl chloride plants with low levels of vinyl chloride monomer. People living nearby would be exposed to it continuously for however many decades they lived there. If some cancers are caused by ambient levels of (say) 0.01 ppm, I want to discuss the likely effects on the numbers of cancers and the age distribution of these cancers that would result from a different ambient concentration—0.001 ppm, or 0.1 ppm, for example. I believe that what should in general be expected in such circumstances is that the age distribution of the cancers that are caused will be the same whatever the dose rate may be; lower dose rates will simply cause fewer cancers than higher dose rates would do, but the age distribution will be almost identical. Moreover, the actual number of extra cancers will, when the dose rate is sufficiently low, simply be proportional to the dose rate: if the dose rate is low, halving it will halve the numbers of cancers, quartering it will just quarter the number of cancers, and so on; no "threshold" should be anticipated.

I would, of course, expect to find the same sort of results in other circumstances where humans are exposed to chronic low doses of other toxic agents, and I would also expect to find this in animal experiments involving chronic exposure of animals to low dosage rates.

*Department of the Regius, Professor of Medicine, Radcliffe Infirmary, Oxford OX2 6HE, England.

Rationalization of Dogmatic Beliefs

Let us consider, for a given level of chronic exposure to a certain agent at an average dose rate of d units per day, a graph (Fig. 1) of $P(a, d)$, the probability that a person exposed for life will get the cancer of interest at age a years, if he does not die before that age. [We have to include the proviso about not dying previously; $P(a, d)$ as defined is determined by cancer induction processes, but the absolute chance of cancer at age a, which equals $P(a, d)$ multiplied by the probability of still being alive at age a, would depend on all sorts of other things—road traffic accidents, homicide, heart disease, and so on—which are irrelevant to the biology of cancer induction.]
To help cause cancer, an agent must help alter a target cell in such a way that this alteration, together with other alterations which the cell suffers (previously or subsequently), transforms the cell's phenotype from normal to neoplastic. Whatever the nature of this alteration, there must be other ways in which it (or an alteration equivalent to it) could have been effected; mammalian biochemistry is so complex that almost any biochemical state could be reached by several different routes, and moreover other carcinogens which act similarly to the test carcinogen are likely to be present in trace amounts.

Let the total effective dose of agents causing this particular cellular change be called D, where $D = d_0 + d$; d_0 is due to background carcinogens, spontaneous cellular accidents and so on, while d, as before, is the dose rate of the toxic substance we are concerned with. For example, if the carcinogen of interest acts by being a mutagen, d_0 might be a mea-

P(a,d)
for a given
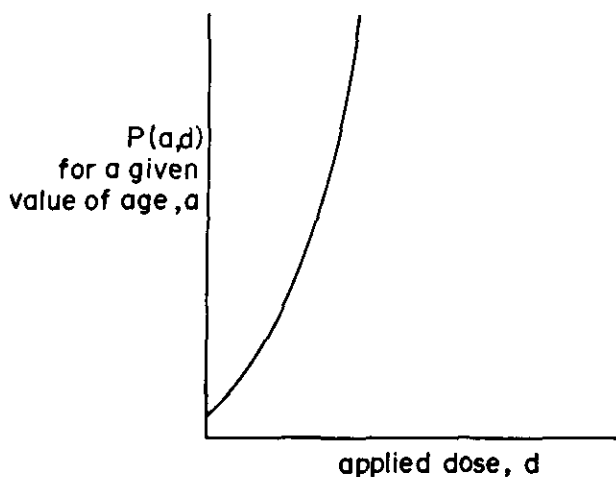value of age, a

applied dose, d

FIGURE 1. Hypothetical dose/response relationship, giving cancer incidence at a particular age against daily exposure rate to a particular carcinogen. The exact shape of the graph as the dose-rate increases does not matter much; it could just as well be a straight line or a line curving in a completely different way, and the later arguments about low doses would still hold. The only important points to note are that there is still some risk even at zero dose, and that the line does not become horizontal as zero dose is approached.

P(a,d)
for a given
value of age, a
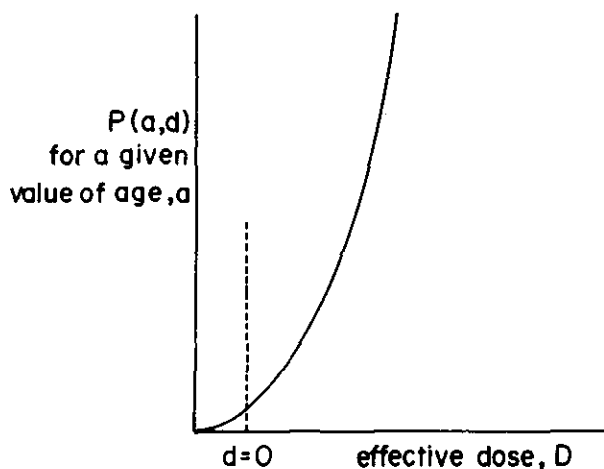
d=0   effective dose, D

FIGURE 2. The same hypothetical dose/response relationship as in Figure 1, but now with the age-specific risk plotted against the total effective dose D, which includes not only the applied dose d of the test carcinogen but also a quantity representing background processes, either other carcinogens or spontaneous cellular accidents, which can produce biochemical effects which are functionally equivalent to those by which the test agent causes cancer. (That portion of this graph below $d = 0$ cannot, of course, be observed simply by varying the applied dose of the test carcinogen, unless the background can also be reduced.)

sure of the frequency of equivalent mutations which occur either spontaneously or because of other agents contaminating the environment.

Now consider a graph of P (a, d) against D (Fig. 2). The point $D = d_o$ (corresponding to $d = 0$) on this graph is off both the axes, and so as the graph of P against D crosses the line $D = d_o$ it will in general have positive slope. Positive slope at $d = 0$ is sufficient (see Appendix) to establish my dogma; (1) that at low dose rates, the age distribution of extra tumors is independent of the dose rate and (2) that, at low dose rates, the expected number of extra tumors is simply proportional to the dose rate.

Qualifications

Some substances exert their carcinogenic effect by producing a pathological change (e.g., formation of a fibrous capsule, formation of bladder calculi, destruction of sperm ducts, suppression of ovulation, or something), such that this macroscopic change then predisposes to cancer. For some such processes "threshold dose rates," below which the change will definitely not occur, might exist. However, it should be remembered, for a hypothetical agent which is carcinogenic because it suppresses ovulation, that although, for example, 1 μg of daily estrogen would not suppress ovulation in most females, 1 μg of daily estrogen to everyone in the population might make the women who are already almost completely anovulatory completely so, and might make those with irregular ovulation ovulate a little less frequently. Thus, because of the wild, outbred heterogeneity of humanity, even agents with threshold-type action in most individuals will sometimes comply with these dogma.

There may thus be a few exceptions, but the arguments are very general and difficult to circumvent, so regulatory agencies should, in any particular case, expect these dogma to apply unless there is specific evidence to the contrary. (It is difficult to imagine what evidence could contradict these dogma for the wide class of substances which are carcinogenic by virtue of their mutagenicity.) These dogma should, therefore, be a foundation of regulatory action.

Counter Arguments

Two points are often made by people who do not agree with these arguments. One is intuitive; instead of discussing whether thresholds exist, this question is bypassed and instead the question of where the thresholds exist is discussed. The observable dose range (of, say, 10^{20} or 10^{25} molecules),

where some risk exists, is contrasted with a dose range so low that it is, in common parlance, "safe"—10^0 or 10^2 molecules, for example. (Actually, the risk from 10^2 molecules might in fact be about 10^{-18} times the risk from 10^{20} molecules—that is certainly "safe," although it is not zero.) The argument then proceeds to ask where, between 10^0 and 10^{20} molecules, safety ceases. To expound this argument is to expose it, especially if we add 10^{20} molecules of other (background) carcinogens to each number cited; the argument depends wholly on verbal confusion between a "safe" risk and a zero risk.

The second counterargument is more interesting. Both Blum and Druckrey have observed that when various (high) dose levels of a carcinogenic agent are given repeatedly to laboratory animals, the time T_{50} until half the animals have cancer is approximately proportional to $d^{1/3}$. This has led to the hope, which I believe to be misguided, that at low enough dose levels almost all tumors would occur long after age 100 and so would be irrelevant to the human condition. (In addition, Blum observed that the variance of the logarithms of the times at which tumors arose was independent of dose; only the mean of the logarithms of the tumor times changed as the dose rate varied.)

The point is that the doses given in these experiments were so high that, had we been able to excise each tumor as it arose, each animal would in expectation have developed several tumors during its natural lifespan. If an animal would have developed several tumors, each with a similar age distribution, then what Blum or Druckrey would observe would only be the first of these several tumors. As with decreasing dose, fewer tumors per animal arise, so the expected time of the first of the tumors on this animal will increase. However, this increase will not continue indefinitely; if the expected number of tumors per animal is only 0.1, for example, then so few animals would get two or more tumors that the effects of the first tumor obscuring later tumors will be negligible, and at all total incidence levels below 10% there will be no further such effect of dose level on the age distribution of those tumors which are observed. It is noteworthy that if Armitage and Doll's multistage model of cancer is accepted, leading to an incidence rate of cancer among cancer-free survivors which is proportional to $(d + d_0)^n$ (age)^k, then Blum and Druckrey's quantitative results—median proportional to $(d + d_0)^{(n/1+k)}$, and constant variance of the logarithm of tumor induction time—would be predicted, as would my dogma. (This doesn't prove my dogma, but it does prove the Blum and Druckrey's experimental results are consistent with them.) These ar-

guments, with references, are developed more fully in a recent review on multistage models (1).

I shall now outline the extent to which experimental results can be used to characterize the dose levels at which linearity begins to obtain, and the risks that then exist.

Uses of Experimental Results to Infer Dose–Response Relationships for Low Doses

Because of statistical uncertainties in the differences between the numbers of tumors found in a control group of animals and a treated group of animals, cancer risks lower than about 10% cannot be accurately characterized, even by large experiments. The experimentally observable range of attributable risk is therefore the range 10% and upwards, and the dosage necessary to give an attributable risk of 1% cannot be determined by direct experiment—likewise, nor can the effects of still lower doses be characterized by direct experiment.

Because of this, much work has gone into the fitting of mathematical models: at this conference, Brown and Crump and Guess have presented their work on this problem, and Hartley would have done so had he been able to attend. These models all attempt to use experimental data (in the risk range 10% and upwards) to estimate the dose–response relationship in the unobservable range (1% and downwards). It seems to me that these authors have all reached qualitatively similar conclusions, and that although the detail of their results in any particular case will depend on the details of the models they happen to fit (none of which models, we must remember, will precisely fit reality), the qualitative conclusions they reach would have been reached whatever class of models they had fitted, and that these qualitative conclusions are therefore valid as a basis for legislation.

The statistical method is to consider a class of "possible" graphs of risk against dose, and then to suppose that the true graph of risk against applied dose plus background dose is one of these possibles. Selection of a particular graph and a particular background dose leads to a specific prediction of the risk at each dose level. Among the range of selections (of graph and background) which fit the observed data as well, or nearly as well, as possible, what will the range of predicted risks at very low doses be? Rather curiously, the range of possible risks at low doses does not depend very strongly on the class of graphs we choose to consider as "possible" graphs of risk against dose. Three main types

of experimental result may occur, and the consequences of these statistical methods for each type of data are outlined below.

Particular Case No. 1

The dose–response relationship in the observable range is found to be roughly a straight line (as, for example, with cigarette smoke and lung cancer in man). Here the graph of attributable risk against dose rate will go straight through zero, with a slope approximately equal to the slope seen in the observable range and confidence limits approximately equal to the confidence limits on a straight line through zero and the observed data points. The upper and lower confidence limits on the risk at low doses will thus usually differ from each other by less than a factor of 2, and the risk at low doses is therefore known sufficiently accurately to assist in rational legislation.

Particular Case No. 2

We find no significant effect at any of the dose levels: the substance is, as far as we know, not a carcinogen. In this case the lower confidence limit is obviously zero, and (to within a factor of about 2) the upper confidence limit is a straight line through zero with slope such that at the highest experimental dose the risk roughly equals one or two S.E.'s of the difference in risk between the highest dose group and the control group. (This will be true even if the observed risk in the high-dose group happens to be lower than the observed risk in the controls: the maximum likelihood estimate of attributable risk due to the high dose is then zero, and the quoted rule follows from staying within an appropriate distance of the maximum.) This gives us a triangular confidence region which will usually be no use to legislators, thus expressing mathematically the common-sense idea that animal experiments cannot demonstrate noncarcinogenicity, even for animals.

Particular Case No. 3

We find such striking upward curvature of the dose–response relationship in the observable range as to suggest to an optimist either that a threshold exists, or that the risk at (say) 0.1 or 0.01 of the lowest dose which has a statistically significant effect may be very much less than 0.1 or 0.01 times the effect at that dose level.

Unfortunately, in this case the models all concur that the relationship could become approximately linear just below the experimentally observable dose range.

(Intuitively, this may be thought of as being due to uncertainty in the effective "background" dose; usually, several particular background doses can be selected such that, with a suitable selection of one of the possible graphs, the data are adequately fitted. Taking the largest of these possible background doses and its associated graph, the possible linearity of the graph of "extra risk" versus 'actual dose' follows.) Because of this, the upper confidence limit for attributable risk is roughly a straight line joining the attributable risk at the lowest statistically significant dose to zero risk at zero dose. In general, the lower confidence limit may be some orders of magnitude below this, and near-total uncertainty will result. However, for a particular reasonable class of models, Crump (2, 3) proved if the excess risk of a particular type of cancer at the lowest dose which produces a statistically significant excess of such cancers is of the same order of magnitude as the risk in the controls, then the lower confidence limit does not differ widely from the (approximately linear) upper confidence limit, and again we have adequate accuracy to assist in rational legislation.

An example of Crump's theorem might be if the net risk of cancer was 10^{-5} times the square of the total effective dose of PAH (pg/day), and the total effective background dose was 10 pg/day of PAH. An applied dose of PAH of 4 pg/day would increase the total dose to 14 pg/day, approximately doubling the background. In the range 0–4 pg/day of applied dose, we would have the extra risk shown in Table 1, and this is sufficiently similar to simple proportionality for legislative purposes.

Table 1.

Applied dose of PAH, pg/day	Extra risk per 10^5
4	96
3	69
2	44
1	21
x, where $x < 1$	$20x < $ risk $ < 21x$
0	0

Appendix

If at each particular age the graph of $P(a, d)$ against d has positive slope at $d = 0$, then at low dose rates it follows, first, that the expected number of extra cancers is simply proportional to the dose rate and, second, that the age distribution of these extra cancers is independent of the dose rate.

Proof: Let the probability of attaining age a if $d = 0$ be written $x(a)$. In a population of N individuals exposed to the chronic dose level d, the probability of attaining age a is $x(a) \exp\{-\int P(a, d)\, da\}$, which equals $x(a)[1 - \int P(a,d)\, da + o(d)]$, the last term in which denotes a quantity which becomes relatively negligible for small enough d. (In other words, as d becomes small, the age distribution comes to be determined by diseases other than cancer induced by the substance of interest.) The expected number of cancers induced at age a is therefore

$$N\,P(a, d)\,[x(a)]\,[1 - \int P(a, d)\, da + o(d)]$$

Writing $k_a = N\,x(a)\,(P/dd)_{d} = {}_o$, $k_a > o$ and this expected number reduces to $k_a\,d + o(d^2)$, which is adequately approximated by $k_a d$ for small enough d. The actual number of extra cancers at each age is thus proportional to d, and the required results follow; for example, doubling d will simply double the number of extra cancers at each age, thereby doubling the total number of extra cancers but leaving the age distribution of the induced cancers unaltered, etc.

REFERENCES

1. Peto, R. Epidemiology, multisstage models and short-term mutagenicity tests. In: Origins of Human Cancer, Cold Spring Harbor Publications, New York, 1977, p. 14033.
2. Crump, K. S., et al. Fundamental carcinogenic processes and their implications for low-dose rate assessment. Cancer Res. 36: 2973 (1976).
3. Guess, H. A., Crump, K. S., and Peto, R. Uncertainty estimates for low dose rate extrapolations of animal carcinogenicity data. Cancer Res., in press.