of more than one event occurring simultaneously, but in Mr Lewis's case we have sometimes multiple failures. I suspect this is a very small effect, but in theory it does complicate the application of the Poisson theory.

When the engineer comes to a really tough problem, such as air traffic control, where we cannot tolerate failures every hundred hours, the engineer adopts other methods: he uses either redundancy or a stand-by system (Wilcox and Mann, 1962), so that the actual mean time between failures is no longer so critical. What is then critical is the proportion of time for which the computer is out of action for repair, as against the proportion of time for which it is working well. So when the engineer looks at a set of test data such as Mr Lewis has presented, he first asks whether we have a goodly number of large intervals and not too many short intervals, because, if we have a number of short intervals perhaps due to failure to locate a defective component, then we must lump them together as computer failure time. For this reason I have extracted the actual time intervals and plotted a distribution of time intervals, as distinct from a survival curve. I then fitted an exponential from the tail end, that is to say, by matching it to the residual number of points after the last interval which had been classified. For Machine 1, Fig. M/C 1 shows that the points are mostly clustered around the exponential distribution of times, except that for very short times they rise quite sharply above the line. Figs. M/C 2 and M/C 3 show the same for Machines 2 and 3. Thus the effect of the branching process appears empirically to be what one would have hoped: it increases the number of short intervals, but does not rob us of long intervals.

I ought not to challenge a statistician's view of what is significant, but I cannot help being a little worried about discarding two items in the data, the two very long periods. The probabilities are so low for long intervals that one hesitates to say that anything is impossible.

I would also have liked, in order to complete the picture, some data on repair times, because if we take Machine 3 and arbitrarily take 100 units as the limit of a useless period, which is suggested by Fig. M/C 3, then this Machine appears to be unserviceable for about 0·16 per cent of the time, or 140 seconds out of a 24-hour day. Since it may perhaps take 10 to 30 minutes to rectify a fault, is not the actual repair time going to be significant after all?

I am very reassured to find that we are not being robbed of the long periods in our distribution, for it has been suggested (at the Symposium on *The Use of Redundancy in System Design*, held by the Society of Instrument Technology in London on February 14th, 1964) that the human-error effect would impose a cut-off at the end of the curve. Now we can hopefully go ahead expecting to get better and better computers which can eventually be applied even to air traffic control. I would like to second the vote of thanks to Mr Lewis for his paper.

The vote of thanks was put to the meeting and carried unanimously.

Professor A. RÉNYI: I would like to comment on two points, both concerning the possibility of other models for computer failure, or failure in other complicated mechanisms or machines consisting of a large number of parts, and on possible refinements of the model discussed by Mr Lewis. My first remark is that on pp. 398–399 there are given four main assumptions underlying the simple renewal model. The second of these is that failures in the separate components are independent, and Mr Lewis later remarks that while this assumption is certainly not exactly true in practice, one would still expect the output to be a Poisson process. Now in this connection I would like to mention that I can very well imagine such a type of correlation between failures in different components, that this correlation itself causes serious departure from a Poisson process. It would be interesting to work out a model on these lines.

My second remark concerns the third point on p. 400. Here Mr Lewis says that he has ignored repair times which alternate with the periods of correct operation of the computer, because these are usually short compared to the mean times between failures, and although the length of these periods appears to be unrelated to the length of the working periods. Now I think that it would be possible to give a more refined model in which a certain type of correlation between these two things is taken into account, because by a very simple reasoning it seems that such a correlation must exist. I cannot tell how strong it is, but it must exist, and the simple argument for this is the following. The repair time really consists of two parts. The first part is, so to say, the search period, until locating the failure, and the second is the repairing. Now what I think is that the time for finding the failure is correlated with the next interval of correct operation in the following way. If the length of time required to locate the failure was long, this means that a large part of the computer has been checked, and then the probability that there remains another mistake is smaller than in the case when the failure was detected in a short time. In this case the probability that there is another undetected failure is larger. This can be seen by a very simple argument, by using Bayes's formula, and considering the converse problem. Take any apparatus consisting of a large number of parts, m, say, with k of them bad parts (there is failure in k of the parts) and then try to find bad parts. If k is large then the time to find one of them is shorter. This time decreases when k increases. Then conversely it follows that the posterior probability of there being still another failure is larger if the time for finding one failure was short. For this reason it would be perhaps interesting to consider a model in which the repair times are not ignored but included in the model, and some dependence on the size of this cluster on the result in the search time could be taken into account.

I am interested in this question because recently I did some work on the theory of random search, and I was especially interested in searching for failures in a complicated system such as a computer. I cannot go into the details of explaining the results one could get in this way, while what I did up to now is only a first step, and the models are very crude and very rough simplifications of the real situation. The main character of the results, however, is that by a random search the time needed to locate the failure is usually of the same order of magnitude as is the best systematic method; and if the number of components is very large then the systematic method of search is for practical reasons usually very difficult to carry out. These results may be of some interest, but I have to emphasize that my models are very simplified models and cannot at the moment be applied to real situations such as failures in computers; I think, however, that it would be interesting to work further on this line.

These are simply remarks on possibilities of further research; I found the paper itself very interesting and stimulating.

Professor WALTER L. SMITH: I would like, first, to compliment Mr Lewis for giving us this evening a most interesting paper in which a model of some importance is given a very exhaustive examination. It is hard to believe that any worth-while stones have been left unturned in the area of the present discussion. However, there occur to me one or two observations which may be worth making.

In Section 4.7 Mr Lewis discusses the asymptotic normality of the counting process $N(t)$; I feel it needs to be explained that $N(t)$ is, in fact, a *cumulative process* (as defined in my 1955 paper in *Proc. Roy. Soc.*, A) and that this asymptotic normality, and various other asymptotic theorems, automatically flow from the general theorems I have given for such processes. To see that $N(t)$ is as claimed, let us define an L-point ("L" for Lewis) as a time instant at which an original failure occurs and such that every subsidiary process which started before this instant has also ended before this instant. It would not be difficult to show that, under reasonably general conditions, the intervals between successive L-points are independent and identically distributed random variables with, say, finite means and variances. Thus the L-points are *regeneration points* of the branching Poisson process and the successive increments in $N(t)$ between successive L-points are a further sequence of independent and identically distributed random variables. This establishes the main requirement for $N(t)$ to be a cumulative process, and, in view of the monotone