Preferred citation: M. Eigen. Experiments on biogenesis. In **The role of fundamental research in paper-making**, *Trans. of the VIIth Fund. Res. Symp. Cambridge, 1981*, (Fundamental Research Committee, ed.), pp 19–33, FRC, Manchester, 2018. DOI: 10.15376/frc.1981.1.19.

EXPERIMENTS ON BIOGENESIS

Manfred Eigen
Max-Planck-Institut fur biophysikalische Chemie,
Göttingen, FRG

Abstract

A living system however primitive, must include mechanisms for conservation as well as for variation of genetic information and for its adaptation to optimal phenotypic performance under given environmental conditions.

In order to maintain such a state of optimal performance the system has to metabolise steadily energy-rich material thereby remaining in a (meta-) stable `far from equilibrium´ position. If such material prerequisites are fulfilled `life´ should emerge inevitably as a `regularity among material events´ and such a regularity, if it exists, should be testable under laboratory conditions.

Virus replicases have been shown to provide efficient mechanisms for assembling de novo macromolecular RNA-templates from energy-rich monomeric precursors and for adapting them quickly to optimal self-reproductive performance. With the help of such systems the theory of molecular self-organisation (e.g. the quasi-species model) has been tested quantitatively. The experiments, in particular, have addressed the question "first template - first enzyme?" including adaptation of target and source function of the genotypes. A machine is described in which experiments on self-reproduction and translation can be carried out in a continuous and systematic way. Questions regarding the origin of life can thereby be scrutinised by experimental falsification.

Thank you, Dr. Corte, for your kind invitation to speak at this symposium.

There are two important relationships between my subject and yours, besides the several already mentioned by Dr. Corte in his introduction. I don't refer to the obvious one of the material we use for our papers, but to the facts that without life there would be no paper, and without human life there would be no interest in it.

So my subject relates to yours at the most fundamental level possible.

From these remarks you might be forgiven for expecting that this talk, entitled `Experiments in Biogenesis' is somehow related to the old idea of trying to produce a homunculus in a test tube.

Those of you familiar with German literature, especially that of Goethe, will recall that he delegated the task of creating the homunculus to Faust's former famulus, Wagner, who was by then a professor of Chemistry. I have often wondered why Goethe did this, bearing in mind his own wonderful work on plant and insect morphology, and deep understanding of the nature of life. (Goethe is almost the only German quoted by Darwin in `Origin of Species'.)

I think it must have been because Goethe was at the time angry with the professors at the University of Jena, and so assigned them, by proxy, a task he knew to be doomed to failure. Just as the reaction is nearing completion and the glow of life is seen in the test tube, Wagner says (Faust part II, laboratory scene) "...What nature must slowly organise, that have we taken, and will now crystallise."

But the homunculus scatters into pieces on the first occasion, and Goethe's sympathy lies with Mephisto who says "...Much have I seen in my years of travel, including crystallised humanity."

So I am not going to speak about trying to produce a homunculus, though I shall show you that we are now in a position to try to perform such an experiment.

In performing a laboratory experiment you expect a result, one moreover which you try to ensure will be reproducible. However, you can expect exact reproducibility only if there is a physical principle behind the process you are investigating, if, in other words, there is some regularity governing the material events. Without such regularity, your experiment will be unscientific and the result unreproducible.

The laws of physics predict the evolution of systems. Starting with certain initial conditions, they answer the question `What then?´ But they can only do this if the regularity governing the process can be established.

The history of life´s evolution can only be reconstructed if we can discover the regularity principle governing it and the appropriate starting conditions. When we know these we can try to build a realistic model, whose results we can test in comparison with genuine historical witnesses. We must have witnesses; theory alone is not sufficient.

So before attempting to perform experiments on life and its origins, we need to discover what physical principle lies behind it, and whether we can measure it.

Now we do know that all forms of life, from the simplest virus or single-cell bacterium, through plants, animals, to humans, share the one property of being enormously complex. And so we ask ourselves if this very complexity is in fact a physical principle, and, if so, how should it be measured.

Modern information theory provides us with a way. If we can describe the exact state of a system and know also the number of other possible states, composed of the same elements, that it might be in, we can say what the probability of finding that state is, which is another way of describing its complexity.

I have said that all life exhibits this phenomenon of complexity. Indeed, complexity is important at a much simpler level than that of a complete living creature. Even the smallest molecular sub-units of life, the proteins, nucleic acids and enzymes, the pre-stages through which the historical process must have passed, show enormous levels of complexity.

As an example of this, consider a typical enzyme, a kind of trypsin, one of the smallest sub-units associated with life.

Even the smallest of these comprise some hundred amino-acid residues, of which histidine is a typical example. Some of these enzymes may contain two hundred such residues, of which only four or five actually form part of the active centre. All the rest of the molecule is used to construct, by folding and distortion, the best possible geometrical configuration of the four or five active residues, whose relative positions are thus defined to within fractions of an Ångström unit, and are in the best possible positions to produce the co-operative catalysis which is the function of the enzyme. Every single one of those residues is necessary to produce this optimal catalyst, so given that there are twenty possible amino-acid residues, we can calculate the combined probability of finding exactly the right molecule at each site. With a hundred possible sites, and twenty possible residues able to occupy each, the combined probability of synthesising this molecule by chance is 20^{100}, or approximately 10^{130}. This gives an indication of the complexity level of this molecule.

It isn't easy to appreciate the meaning of a number as large as this, until we look at one or two of the known facts about our universe. Even close-packing the entire known universe, a sphere of radius some 10 billion light years, with protein molecules would provide us with only 10^{105} of them.

Since a simple bacterial cell makes use of a few thousand such enzymes, while the human body needs about ten thousand, I think you will agree that there is absolutely no chance that all this complexity could have come about randomly. Even making allowance for time doesn't help much, since the age of the universe, 15 billion years, is only about 10^{17} seconds.

If therefore there is a physical explanation for the origin of life, then we must try to discover the regularity behind the goal-directed optimisation which must have been responsible for this enormous complexity. Once we have this, then we can begin re-creating the process in the laboratory.

It has become clear over the past decades that the logic governing the physical principle was introduced by Charles Darwin, though not to molecules because he didn't know about them. He used his principle of Natural Selection to explain the evolution of species, which was a popular subject at the time. (Goethe too believed in evolution, fifty years before Darwin.) He proposed Natural Selection in order to establish a principle behind evolution, and the way he put it was "The more complex evolves from the less complex, according to the conditions of Natural Selection".

Following Darwin, population geneticists like Haldane, Fisher and White, working in the first half of this century, made it clear that Natural Selection, with the emphasis on the natural, is a consequence of self-reproduction. They showed, using the theories of population genetics, that such reproduction is a sufficient condition for Natural Selection. We have found more recently, applying their statistics to molecules and making use of theory developed over the past twenty years, that it is also a necessary condition.

The logical statement of our present position is summarised in the following three sentences:

1. Systems of the order of complexity under consideration must have evolved, since there is no other way of having made them.

2. This evolution is based on the principle of Natural Selection.

3. Natural Selection as a regularity principle governing the patterns among events appears only in self-reproductive systems.

These statements have all been proved theoretically, but I don't want to discuss theory now and shall go on to describe some experimental results supporting them.

Evolution by Natural Selection can occur only in self-reproductive systems, so the experiments I am going to describe are all concerned with such systems.

The simplest living organisms are the viruses, occupying the border between life and non-life. Amongst these the simplest are the single-stranded RNA viruses, whose genome consists of one, single-stranded, RNA molecule. The information contained in such a molecule is usually enough to encode up to four proteins: one of these will be a co-protein for protection, another will be the penetration protein that the virus needs to penetrate its chosen host cell. We are not concerned with these. The most important protein encoded is that of the enzyme responsible for replicating the virus genome.

Thus within the virus' own RNA sequence exists the coding for the enzyme needed to catalyse the reaction that copies it.

Indeed, the virus we have studied, Q_{beta} Bacteriophage, is more economical than that, since it doesn't store the code even for the entire enzyme. Its RNA contains the code for just one of the four protein sub-units of which the enzyme is composed: the other three are made, conveniently, by the host cell.

The association of these four proteins thus forms an enzyme, but not the one that catalyses the reproduction of the entire virus genome directly: instead it catalyses the production of palindromic sections which must then associate to form the complete chain.

Thus, the virus genome is eventually translated completely by the machinery of any cell it infects, by making firstly the replicases, and finally, the whole genome, causing the demise of the cell.

We have studied this sytem very thoroughly, and know that the sub-unit which is encoded directly by the virus genome was once itself a factor of the host cell. Because we have found the precursor, we can even go so far as to say that the virus itself was once part of the host cell; somehow it separated off and became its worst enemy. Indeed the precursor too can reproduce inside the host cell, but with different specificity and lower activity.

Sol Spiegelmann was the first person to demonstrate that this is really the infectious mechanism of the virus, and that this is the principal of the virus' life cycle. In his experiments he isolated viral genome RNA and some of the protein sub-unit encoded by it, and associated these with some of the ribosome protein from the host cell and some energy-rich nucleoside tri-phosphate. He was then able to demonstrate how enzymes were produced that were capable of linking up the nucleoside tri-phosphate substrate molecules to produce large numbers of viral genome replicas. He was then able to recover complete particles of the original virus from the plaque left after infecting a coli culture with these replicas.

This principle, of constructing an intermediate replica from the enzyme, the RNA template, and energy-rich substrate, and later re-introducing the intermediate replica into the life cycle in order to recover complete particles of the original virus, is very powerful.

Spiegelmann always made use of a satellite compound of the virus genome, a molecule consisting of only 220 nucleotides rather than the 4500 of the original, and even though this is just about the simplest system known you can imagine how complicated it is.

The replication reaction sequence proceeds serially, using the template-enzyme complex as a base on which to link nucleoside tri-phosphate molecules. The phosphate-diester bonds between adjacent tri-phosphate molecules are formed before the next one is added.

When the new molecule has reached the same length as the original template, linking stops, and because the linking is vectorial, the enzyme complex must disassociate from the replica before it can take part in a new construction.

In a typical experiment we introduce radio-active phosphorus, P^{32}, into the tri-phosphate molecules and monitor the progress of the reaction by observing the radio-activity of the RNA.

Starting with a milli- or micro-litre of 10^{-8} molar enzyme solution, which, though very dilute, still contains some 10^{14} molecules per milli-litre, we add an equal volume of template

RNA, plus appropriate substrate material, and observe the progress of the reaction. Invariably the rate of production of new material is linear, showing that the template-enzyme affinity is so great that all available enzymes must be in use from the beginning. Were this not the case, then because the reaction product is also one of the starting materials (template RNA) we would expect to see exponentially increasing quantities of new material. The linearity of our production curves shows us that no matter how much template is present, all the enzymes are in use, so that all must have been in use from the start.

We have tried diluting the template to low enough concentrations to be able to detect an initial exponential section of the production curve, but the affinity is so great that whenever measurement begins the rate is found to be linear.

So to detect the initial exponential section of the production curve, which must exist, we have recourse to serial dilution methods.

Beginning with a solution containing about a million template molecules, we so dilute it that ultimately there is a whole sequence of solutions containing zero, one, two, or more molecules, according to a Poisson distribution. We can tell how many molecules were originally in each solution by what happens when enzyme is added: those solutions originally containing none, remain with none: those where there was a single original show a particular production rate: those where there were two originals show a faster production rate: and so on.

Figure 1 illustrates the different production curves obtained with different initial numbers of template RNA molecules.

By careful analysis of the results of these experiments we have been able to show that equal factors of dilution produce equal factors of production rate of new template RNA material. Such behaviour can only be represented by a logarithmic relationship, so there must be an exponential part of the template amplification curve at extremely low concentrations, though the enzyme-template affinity is so great that we never see it directly. (One is indeed shown on the curves in figure 1.)

In addition we measured all the pertinent reaction rate parameters and concentrations in order to build up a complete picture of what is happening in this very complex reaction.

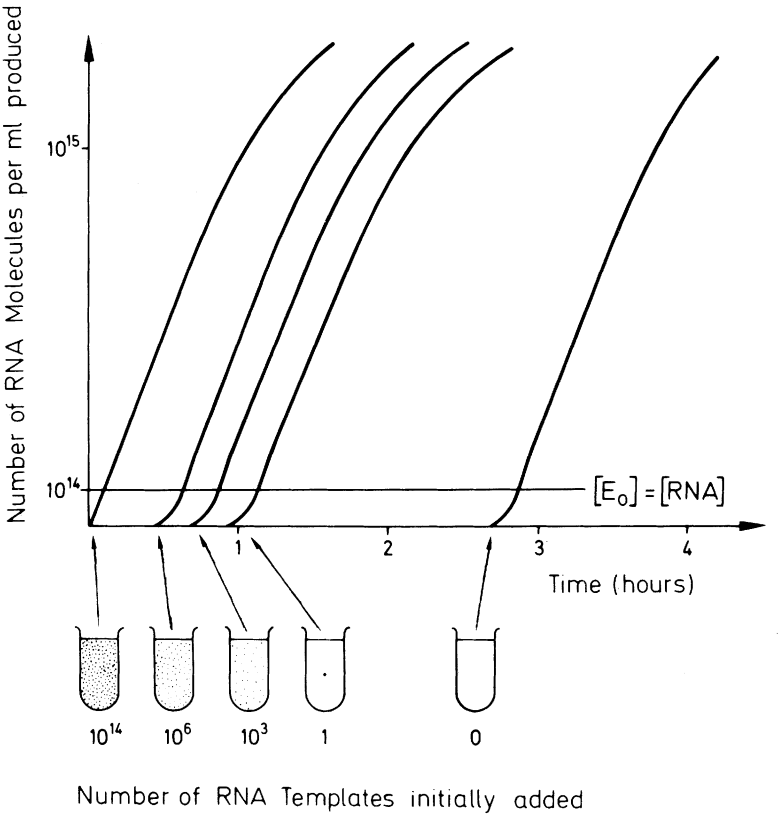

Fig 1

To interpret all this information we had to make use of computer simulation, because with 220 nucleotides in the RNA chain. (You will remember that even the satellite molecule used by Speigelmann was that long) there are over 1000 rate equations. (The actual number is given by the expression $4n+4$, where n is twice the number of nucleotides.)

From this simulation we found that growth of a particular mutant RNA continues exponentially until the enzyme is saturated with that particular mutant, not just generally saturated (as could be the case in a multi-mutant system).

We also find that positive and negative strands contribute differently to the growth and that it is the geometric mean of the two rate parameters that governs the overall rate. Production of strands of both signs occurs simultaneously, so that systems whose positive and negative strands show some palindromic symmetry have a great selective advantage, because then the reaction sequence is the same for both strands, and enzyme optimisation vastly simpler.

But in speaking of the origins of life, it isn't perhaps terribly appropriate to describe the results of some experiments in enzyme chemistry, since there were no enzymes when life began.

So we have been anxious to discover whether these complex enzymes can be replaced by other, simpler catalysts that might have been present in the early stages of evolution.

Leslie Orgel, of the Stock Institute at La Jolla, California, found that simple zinc ions, which are co-factors of all replicases, are excellent catalysts for the 3'-5' bond in nucleotides, especially in Guanine-Cytidine systems, where they can catalyse molecules of up to 200 units long. So it looks very likely that the first enzymes were zinc ions, perhaps chelated with poly-peptides to improve their activity.

Thus we know there can definitely be nucleic acid polymerisation without enzymes, so the other question to be asked, or maybe it is another side of the same question, is whether there can be nucleic acid polymerisation without templates. You will remember we always used template RNA material in our experiments, even when there was only one molecule of it present.

But during one of these experiments, Manfred Sumper, a colleague of mine, made an accidental discovery of enormous importance.

He had finished a serial dilution experiment late one evening, and, because of the lateness, had left the washing up to the following day. On returning in the morning he was astonished to find that in those solutions where there had been no original template material, there were now large populations of identical RNA polymers. Of course this occurrence was treated with the utmost scepticism by all our co-workers, who suggested that we go and prepare our reagents more carefully in future.

But we were not convinced that impurities had been responsible for this odd result, because we were by then very familiar with the production rates for even single molecules, and those samples had definitely not amplified at single molecule rates.

So when we told Fritz Lippmann at the Rockefeller Institute of this result he cabled us back congratulations on having broken the dogma of molecular biology. He meant by this that biologists had always believed that the direction of information transfer must be from the complex nucleic acid to the simpler protein, because how could a simple protein molecule hope to encode a complex polymer, 220 units long?

Well, it did. And because we were studying evolution we guessed that we had witnessed a Darwinian optimisation.

We didn't expect to be able to reproduce this result, since a polymer 220 units long has 4^{220} possible sequences. But then if it had really been a Darwinian optimisation that we had witnessed, perhaps we would be able to repeat the result.

So another colleague of mine, Christof Biebricher, performed the following experiment.

He incubated ultra-carefully purified (to satisfy the sceptics) enzyme and nucleoside tri-phosphate solutions for long enough that any single impurity RNA molecule present would have been multiplied a million-fold. The time needed to achieve this is well-known to us now, and it is too short to produce any RNA de novo.

He then quenched the reaction, performed a serial dilution, and incubated each resulting solution long enough for de novo production, if it exists, to occur. He found at the end of this stage of the experiment that each sample had produced its own population of RNA molecules, each of which, according to fingerprint and sequence analysis, was different. This fact of their all being different really proved how impurities could not have caused the generation, for the populations would in that case all have been the same as the impurity. So we were sure that we were witnessing de novo production of a series of RNA mutants.

Having ascertained all this, he then recombined his samples and found that he was left very soon with only a single RNA type, the one with the highest rate of growth, the one used to prepare the original enzyme.

It is quite clear from this experiment that the enzyme has a weak ability to link together nucleotide polymers from the triphosphates without needing an RNA template. We believe that it lines up the substrate molecules at certain recognition sites and uses the resulting sequences as templates since analysis of the products obtained from Biebricher's experiment showed that certain blocks occurred very frequently. The presence of these blocks limits the number of possible protein sequences in the product, the complexity, to about 10^{12}. Since there were about 10^{14} enzyme molecules in each reaction solution then the chances of trying all the possibilities and adapting them to find the best polymer were brought down to a reasonable level.

And now I shall describe a striking difference between these de novo experiments and the earlier ones on template instructed replication. Where the latter are deterministic, with well defined rate constants, the former are erratic, showing considerable variability in their rates. It seems that the long induction times observed in de novo generation stem from the difficulty in making the first molecule. Once this is achieved, the experiment proceeds normally.

The different reaction kinetics are summarised in figure 1. We realised that the results of this experiment contained an illustration of the statement that natural selection is a

consequence of self (ie template instructed) reproduction. During the de novo incubation period different templates are made in all the different solutions. But only one of these, that with the most rapid growth rate, survives when the solutions are recombined: selection has occurred.

We also managed to show that this selection was subject to certain constraints, though I shall not go into details here.

An experiment performed by Karl Weitmann at the University of Zurich showed how the target of this selection turns out to be a quasi-species, mutant distribution.

He performed a serial dilution on a Y-type solution containing a number of mutants, each with different protein sequences.

He was able to establish, analytically, the modal protein sequence, that is to say the sequence built up of the most frequently occurring protein at each site. This modal sequence didn't actually occur amongst his mutant distribution.

So, having performed his dilution down to single molecule level and added his enzyme and substrate, he observed each mutant reproduced rapidly, resulting of course in a collection of large populations of each mutant, all different.

Then he left them all alone long enough for de novo production to occur, and they all became the same, the one composed of the modal sequence, which is also the one with the highest rate of reproduction.

From the results of experiments like this it is possible to calculate the individual mutant distribution, and, more importantly, each one's information content. This we have done for both RNA and DNA replicases and found the values actually observed in nature. So not only are we sure that the principles of natural selection given at the beginning of the paper are correct, but also that natural products are made according to these principles, since they can be used as quantitative predictors.

We now understand the limitations on each selection step, and know that evolution proceeds rapidly for short templates.

But there remains still the problem of how the best product for the fastest evolution is arrived at.

The answer is at first surprising.

Imagine a mechanism that, by self-reproduction and selection, can climb uphill in a valleyed landscape. You might expect that, starting at any point, it would reach the top of the nearest hill, but that because it would need to descend again, it would never reach the highest summit.

This turns out wrong, because our nucleotide system is not in an imaginable two or three dimension landscape.

A sequence of N nucleotides evolves in an N-dimensional space and by making a mutation at any of the positions has N points from which to proceed. A binary sequence therefore has 2^N possible alternative arrangements, while an N-dimensional topography has 2^N possible saddle points.

Thus there is an outcome to cater for each and every possible decision, and there will always be a route up to the highest peak, at least for molecules with the dimensions of proteins. Whether this would be true for larger chains is not yet clear, but at least for proteins, the world appears to be the best of all possible worlds.

Finally, because I imagine you to be interested in commercial applications, I shall speak a little about some that I think will soon be possible.

For example we could quite easily modify our experiments such that they can self-adapt to produce any desired protein sequence, such as might be required to produce a gene in selective factor hybridisation (genetic engineering). That is fairly simple.

A really useful application would be to produce a molecule that after translation gave a useful protein, a catalyst, a drug, or whatever.

This would be much harder because ordinary translation is not sufficiently specific, it will operate not just on the gene you want but also on all others.

A machine sufficiently specific is still science fiction, though we are trying to build one.

Based on the principle of the cell-sorter, it would be a reactor of the first type in which the reagent is so dilute that when it is atomised each droplet contains on average only a single mutant RNA strand.

It should be possible to produce a million droplets in an hour or so, and focus them electrically onto a tape, say about 10m long.

Then each droplet would be incubated for amplification, so that each mutant strand would be multiplied a million- or billion-fold, to produce a useful population.

Then each mutant population would be incubated for translation and, using some automatic test device, we would find out which population has the most useful characteristics. This would form the input to a second stage of the process, and we have calculated that within 40 generations we should be able to produce the best protein sequence for our needs.

So, there is a brief insight into a new field of research that may one day even be useful as well.

Thank you.