ISIS Opening Address: Complexity versus Simplicity

Murray Gell-Mann

Citation: AIP Conference Proceedings **755**, 1 (2005); doi: 10.1063/1.1900386
View online: http://dx.doi.org/10.1063/1.1900386
View Table of Contents: http://scitation.aip.org/content/aip/proceeding/aipcp/755?ver=pdfcov
Published by the AIP Publishing

ISIS Opening Address: Complexity versus Simplicity

Prof. Murray Gell-Mann

Santa Fe Institute, Santa Fe, New Mexico 87501-8943 USA

Abstract. There are a lot of different concepts to cover all the meanings that we attach intuitively to the word complexity, and to its opposite simplicity. There is one kind of complexity that corresponds best to what is meant by the word complexity in ordinary conversation, and in most scientific dialog. It's what I call *effective complexity*. Roughly, *effective complexity* refers to the length of a very precise description of the regularities of an entity. Not the features that are treated as random or incidental, but the features that are treated as regularities.

Keywords: effective complexity, randomness, phenomenon, simplicity, order, chaos
PACS: 80.0

EFFECTIVE COMPLEXITY

Complexity does not mean randomness. For example if we take a sheet of paper with a lot of dots arranged randomly that is not complexity, in fact it's rather simple. A complex novel, for example, would have many different characters, many different scenes, many different sub-plots, each of them taking quite a long description. Here is Snoopy's take on that. Sally says: “That too many characters in the book, too much going on… I can't keep track of them all.” Snoopy agrees, Snoopy says: “I like a book when it's only one character and nothing happens to him.” So that would be a simple novel, not a complex novel. The U.S. Tax code is complex. Every rule in the book is a regularity because it's a rule, and there are a great many of them. The U.S. Tax code fills up a large heavy book.

We also see complexity used in the world of advertising. There is an advertisement for a cognac, a brandy. Both the brandy and the young lady shown in the picture are labeled appropriately complex. This is a very good illustration of what people in the advertising business mean by complex, it's slightly sexist unfortunately but here's the young lady rated according to a number of characteristics, instigator, muse, daughter, klutz, accomplish, gardener, slow kisser, journalist, optimist, pessimist, flirt, insomniac and in each one she's rated numerically so there are a great many different traits and a rather lengthy description of the measure of each one. Clearly she is complex and presumably the brandy is complex as well, but you see that they mean the same thing that we mean by effective complexity.

You may notice that I'm defying convention here by not wearing a neck tie. But I brought along a whole packet. Suppose we look at the patterns of these various ties. Here is one that is relatively fashionable two or three years ago among people who wear ties. You see it is very simple, especially for those of you in the back. You see a rather coarse gray version of the tie, a dark blue stripe, a thin white stripe, a red stripe and the same pattern repeating over and over and over again. So it could be described very simply and in a sort description. For those of you in the front, perhaps you can see inside each red stripe there is some thin blue lines and one red seam. And inside the blue there is also some lines with a red thick. That makes the description slightly longer but only very slightly. It's still a simple neck tie. We can look at one that's somewhat more complex. This one has a chain and some other decorations besides the stripes and a little bit more color, but is relatively simple, too. This one it's rather simple, especially for those of you in the back, and needs a very short description. Now here's one that's genuinely complex, this was designed by Jerry Garcia, when he was still a visual artist. See now, we take a long time to describe adequately the regularities of this tie. So this is definitely a complex one. Now you notice we've concentrated on the pattern of the ties and asked whether they're simple or complex but we've been ignoring the soup stains, the wine stains, the milk stains and so on, on the ties. We just looked at the patterns. But how do we know what is important? Is it the pattern, or the stain? Suppose you're a dry cleaner, then you're probably interested mainly in the stain and not very much in all the patterns. So what is the regularity it depends some how on the kind of the judge. A judge of what is important and what isn't important, not directly a judge of regularity and randomness but a judge of least what is to be treated as important and what is to be treated as unimportant. That judge doesn't have to be human, doesn't even have to be alive. But it has to be something to make it distinction between what is treated as important and what is not. We see that understanding regularity and randomness is the key to understanding the distinction between the simple from the complex. Everything we see around us displays this delicate interplay of regularity and randomness. The universe does, and as we get to smaller and smaller and smaller things in our experience, the same is true: delicate interplay of regularity and randomness, delicate interplay of law and chance.

We see that the distinction between the regular and the random is often context-dependent and can even be subjective. When we listen to music on the radio and there's a lot of static, a lot of noise on the radio we describe the static as random and the music as regular. But in the 1930s, at the Bell Telephone laboratories in New Jersey, Dr. Janski and his associate Dr. Bailey were instructed to try understand something about where static came from. And what they discovered way back then in 1930s was that there were important sources of static in the sky located in particular constellations in the night sky and corresponding places in the daytime sky. What were they? Well, they were signals from distant stars in a particular distant galaxy, and that is what gave rise to radio astronomy. So static is not just noise, it contains important regularities, actually it is the basis for the whole science of radio astronomy. And we saw in a similar way that neck ties can exhibit regularities not only in the pattern but also in the stain. Some people might be more interested in the stains then in the pattern. Now in science we search for regularities. Natural phenomenon tend to obey laws and as Sir Issac Newton said “It is the business of natural philosophy to find them out” to find the little regularities in nature, the non-random phenomenon in nature, science is a search. In the 17th Century science was called usually natural philosophy. As in that sentence we just quoted from Newton. It was natural philosophy distinct from what we might call armchair philosophy. Natural philosophy involves comparing of nature, getting ones ideas by looking at nature, by making theories, testing the theories by observation of nature and so on. Constantly going back to the natural world to check on ones ideas whether they're right or they're relevant. This is the opposite of armchair philosophy.

There has been some significant controversies about regularities: Are there regularities of a certain kind or are there not? Take one example: for many years there were Wall Street, and else where in the financial world so called chartists who made charts of prices in financial markets prices of stocks, for example. Graphs of stock prices versus time or prices and commodities versus time or whatever. And then claim they could from the fluctuations in these graphs deduce something about future fluctuations and thereby make money for themselves or their clients. Of course they could not deduce the future regularities. All they could do would be to supply probabilities for certain kinds of fluctuations in the future. But the point is they could extract non-random results and apply them, or so they claim. That there were lots of famous neo classical economists saying that the fluctuations around fundamentals were nothing but a random walk. In fact one of them wrote a book called "A Random Walk Down Wall Street". After a random walk you can not extract any useful information: it's just random. By looking at the historical data, past behavior of markets, one can see that there is a non-trivial correlation between price fluctuations of one time and price fluctuations of another time. So from a historical record you could show that the random walk down Wall Street people were simply mistaken.

What gives rise to complexity, where does complexity come from in nature? Well, the fundamental law of physics which govern the behavior of all matter in the universe and of the universe itself. They seem to be simple. You can write them down very concisely as far as we know. Of course we do not exactly have these laws yet, but we are getting closer, and all the indications are that we will get the laws that govern the behavior of elementary particles, the basic building block of all matter, and the boundary condition at the beginning of the expansion of the universe. When we will get these two fundamental laws they will be simple. They will be describable in a brief message. The first thing is the unifying quantum theory of all the elementary particles everywhere in the universe. And the other one is the initial condition of the universe beginning, near the beginning of its expansion around 13 billion years ago. As we said both of these laws seem to be coming out simple, so where does effective complexity come from? It doesn't come, we believe from the fundamental laws. Well let's ask this question, if we know the exact fundamental laws of physics, the theory of the elementary particles and the initial condition of the universe, can we then in principal (not of course in practical) predict the behavior of everything in the universe? A hundred years ago many people would have said yes. If you really know the laws and you know the initial condition, you can predict everything in principal. But it is not true, because we know that the universe is governed by quantum-mechanical theory, and in quantum-mechanical theory all you get is a set of probabilities for various alternative histories of the universe. You do not get a prediction of a particular history of the universe but a set of probabilities for many different alternatives. So the fundamental laws are probabilistic and not fully deterministic. And the history of the universe is co-determined by the fundamental laws which we believe to be simple and an unimaginable long sequence of chance events which we call accidents which are governed by probabilities. There are a number of different outcomes for each accident or chance event and each one has a probability. But you do not know in advance which one you are going to get.

Let us take a simple laboratory example. Take a radioactive nucleus that emits alpha-particles that is helium nuclei. In advance of the emission, before the radioactive nucleus disintegrates, you have no idea in what direction the alpha-particle will come out. After it comes out of course you would know, but before hand there is no way to tell. In this case the probabilities are all equal. That is not always true, but in this example the probabilities are all equal, every direction is equally probable, all directions are equally probable. Only afterwards can you specify in what direction the alpha-particle went. That is just one simple example of what accidents are chance invested. Theorists may be coming close to a description of the fundamental unified theory of all the elementary particles. That is the research on a particular model of superstring theory and search for a possible generalization which is named m-theory. This kind of research may be coming very close to the unified theory of the elementary particles and their interactions, in other words, a complete theory of the behavior of matter in terms of its fundamental constituents. Already this body of theory based on so-called superstrings has scored a remarkable triumph. Some of this was done actually in my group from Cal Tech, when I was still there, although I didn't do it, but I brought the people there who did. And what they found was the superstring theory, which may be a part of the true underlined theory of elementary particles. In the superstring theory one can derive Einstein's famous general-relativistic theory of gravitation, and further more, one can derive it within quantum-mechanics and without the absurd infinite corrections that have plagued all previous attempts to reconcile general relativity with quantum-mechanics. Here they are fully reconciled and there are no infinite corrections. Everything looks perfectly ok. Well that triumph seems to be the indication that the people who are looking for the correct unified theory starting from superstring theory maybe on the right track.

The other fundamental law is the initial commission of the universe near the beginning of its expansion around 13 billion years ago. In that, it is a simple initial condition, and that is ultimately responsible for the so called arrow of time. That distinguishes the past from the future. If you see a film of a lot of little bits of egg, yolk, white, shell and so on. Starting out scattered around and gradually assembling to form an egg. You would conclude that was a film being shown backwards. Nobody has ever seen a situation where a lot of little bits of egg, yolk, white, and shell assemble to form an egg. But many times you dropped an egg and seen it come apart into little bits of yolk, white, and shell, and you can easily tell which is the movie going forward and which is the movie going backwards. That's the so called arrow of time or at least one of the arrows of time. It's related of course to the second law of thermodynamics, a little fancy expression which means that the average disorder has the tendency to increase in a closed system. So the egg breaking is an example of increased disorder. The egg reassembling would be an example of disorder decreasing in a closed system, which we don't see. So the second law of thermodynamics depends crucially on this orderly state at the beginning of the expansion of the universe 13 billion years ago. Now you can say you observed the phenomenon of the egg in the laboratory, or in your house. What does it have to do with the universe? Well, you can trace all these phenomena in different parts of the universe, back to a fundamental reason why time goes forward, and that has to do with the universe. The universe goes forward in time, therefore the various parts go forward in time, therefore people for example live forward in time, you can not remember what happened tomorrow, and so on. Therefore when you drop an egg, you drop it forward in time and not backwards. So the arrow of time for the universe is ultimately responsible for the little local arrows of time that we see in parts of the universe.

Now, as we take into account all these probabilistic situations, and all these accidents, all these random events we see the alternative possible histories of the universe to be represented as a branch in a tree. You have one accident with various possible results, different probabilities, and another accident with different probabilities, another accident with different probabilities. This way you get a branching tree for all the alternative histories of the universe with a probability, with probabilities in all the branches and with a probability for each history in the tree. Jorge Luis Borges had a story called "El jardin de senderos que se bifurcan" ("The Garden of Forking Paths") it's about a man who left behind a mysterious garden with some message involved, people couldn't figure out at first what it was but then they realized it was a map of the alternative histories of the universe represented as forking paths in the garden. Now some people refer to the still sought unified theory of the elementary particles and their interactions even when it's supplemented by the initial condition of the universe and the other fundamental law. They refer to it as the theory of everything. You may have heard that expression used often by very distinguished scientist, but it's a stupid name, because most things have to do with accidents. There is very little that depends only on the fundamental laws. Elementary particle physics and cosmology are the two sciences that depend on the fundamental laws, but all the other sciences depend on accidents as well. Geology for example depends on the existence of the solar system and the planets, the various histories of the planets all depending on numerous accidents. Biology depends on even more accidents, all the accidents that have taken place in the course of biological evolution in the last 3.9 billion years or whatever it is since life started on Earth. Even chemistry, some of which is derivable directly from elementary particle physics, depends to some extent on accidents because you have chemistry only when conditions of temperature and pressure and so on are such to permit atoms to exist, atoms and molecules then you get chemistry. In the center of the sun for example there is essentially no chemistry, it's too hot. You do have nuclear physics in the center of the sun, but not much chemistry. So, all the other sciences, except elementary particle physics and cosmology, depend on accidents. Think of all the accidents that produced the various people in this room. Some nasty little quantum fluctuation that gave rise to our galaxy long ago billions of years ago, and then gave rise to the evolution of all the stars in the galaxy, including this very ordinary star that we call the Sun. And along with the sun, when the matter condensed, a lot of planets were formed, too, and all that depended on a large number of accidents. One of those was the third planet from the sun, the Earth. And it underwent all sorts of changes, many of which depended on accidents. Then life came about in a particular way on Earth almost 4 billion years ago. But that involved accidents too then, in the course of biological evolution there were enormous numbers of accidents that took place, all though of course things like natural selection created regularities among those accidents, but they were still accidents. Think of all the other things that had to happen to produce us, two people meeting, sperm meeting egg and so on in a particular manner determining our various genomes, but it's not only the genomes that determine the person. Identical twins have the same genomes, but they're different people, different experiences in the womb, different experiences in early childhood, and different experiences in adulthood, and so on. So, there were enormous numbers of accidents that produced all of us the way we are today.

Some of the chance events, accidents, or branching in the tree produce more future regularities than others, in finite regions of space and time, and we call those *frozen accidents*. They must be the main source of effective complexity because the fundamental laws are thought to be simple. So it is the frozen accident that creates most of the regularities we see over and above the fundamental laws. Take some examples of life on Earth, many right handed molecules claim working roles while the corresponding left handed molecules do not. For example, left handed amino acids and right handed sugars are very important in biology where its' mirror image molecules play very little role. Now it has been possible to explain why the sugars would all have the same handedness and why the amino acids had the opposite handedness. People have succeeded in giving theoretical explanations of that, but why it does it a particular way that it does it, left handed amino acids and right handed sugars, nobody has explained that, it seems to be an accident. People have tried very hard to contribute it to the fact that for matter is opposed to antimatter, that we interact with left handed but they never succeeded in doing that. So, it looks as if this is an accident, an accident of the earliest form of life on Earth, and one that has been propagated through the whole process of biological evolution and all living things, the characteristics of all living things today. We can also look at accidents in human history, and ask this what consequences they had, where they important frozen accidents.

Now of days, so many distinguished historians are much more tolerant than they used to be, asking what if so and so had happened instead of so and so. It's not something you can easily test of course, but that kind of speculation has proved to be interesting and historians are more and more tempted by it. They call it Counter Factual History or Contingent History and many of them like to talk about an incident that took place in 1889. While Buffalo Bill's Wild West Show was touring Europe and made a stop in Berlin, the star of the show was the famous female marksman, Annie Oakley. Annie would ask for a male volunteer from the audience to smoke a cigar and have her stand there with the cigar in his mouth while Annie shot the ash off the end of the cigar. Normally there were no male volunteers, then her husband who was himself a distinguished marksman, would step up, he would smoke the cigar, leave some ash at the end, and have his wife Annie shoot the ash of the end of the cigar. But on this one occasion in 1889 in Berlin, there was a male volunteer from the audience, the Kaiser, William the second, who had just ascended the throne a year before on the untimely death of his father. And there he was all dressed up in a uniform, this very elegant uniform, he took out an expensive Havana cigar and clipped off the end and took off his band, lit it and waited for some ash to accumulate at the end of it and then stood at attention on the stage waiting for Annie to shoot the ash of the end of the cigar. Well Annie was worried, she had been drinking heavily the night before, and her husband was one thing but the Kaiser was another. But it all went off all right, she shot the cigar and not the Kaiser, we know that because otherwise we would have read about it in history and other places. But what if she had actually killed the Kaiser? William the Second was a difficult character as most you know. He canceled the reinsurance treaty with Russia, he engaged in naval competition with Britain, which Bismarck advised against, fired Bismarck, who was trained to construct the stable order in Europe, and his work led in many ways to the First World War. Would the First World War have been quite different if he had been killed in 1889? Would it perhaps never have happened… we don't know, but people can speculate about it. Anyway it may well be an example of a very important frozen accident. And historians argue about whether the results of major accidents like deaths of prominent figures are eventually healed by grand historical forces or whether they continue into the future to have very important effects for a very long time.

In any case since the fundamental laws are thought to be simple most effective complexity anywhere in the universe can be traced to frozen accidents. And as time goes on in many domains of experience, we see that entities of greater and greater effective complexity come into being. Why is that? Well, any given entity can become less complex. People die, civilizations die, and they become less complex. But the point is that the envelope of effective complexity keeps getting pushed out as more and more complex things come into existence, even though many things decline in complexity. The boundaries of complexity keep being expanded. The envelope of effective complexity is pushed out when the accumulation of the results of frozen accidents outstrips the forgetting or erasure of the results of frozen accidents. And as time goes on you get this accumulation. It doesn't contradict the famous second law of thermodynamics that average disorder in a closed system tends to increase because mechanisms of self-organization can cause local order to increase while order is decreased elsewhere, like for example your refrigerator makes ice cubes which are very regular and very ordinary but if you go around to the back of a refrigerator there's a huge amount of heat coming out which is very disorderly and which makes up for the order that's being created in the freezing compartment, more than makes up for it. That's how local order can create spiral arms of galaxies, different forms of snowflakes and so on. As many theories believe, protons disintegrate with a half-life of 10^{33}-10^{34} years. Then, after the 10^{36}-10^{37} years there'll be almost no matter left of the kind that we know, made of atoms and molecule, and so on, instead there'll just be a soup of electrons, positrons, photons, neutrinos, and antineutrinos. Very few regularities as they are now conceived. At this point the envelope of effective complexity might shrink, but this is not something to worry about right away, 10^{36} years is a long time.

Complexity versus simplicity

Above, we talked about effective complexity, the length of a very concise description of the regularities, and about where it comes from, since the fundamental laws of nature seem to be simple. And we found it comes mostly from accumulating accidents, accumulating frozen accidents. What is a lengthy concise description? What do we mean by regularities? Is there some mathematical way of distinguishing regularities from random futures? We pointed out that everything around us exhibits mixture of regularities and randomness including the whole universe. And the job of the scientist is to locate the regularities and understand how they interact with randomness. Now since we're in the age of computers we'll have to talk about the bit string, the language that computers use. So a bit string is a string of zeroes and ones, that's what information is these days. And if we're talking about some entity and we want to talk about its effective complexity then we represent it by a bit string. And in order to do that we have to specify various characteristics of the entity we were talking. One thing is the level of detail, which we are calling coarse-grained. In physics it's usually called coarse graining at which it's being described. I think the phrase coarse graining probably comes from photography, you have a very grainy photograph then you're seeing only certain features out of many, if you have a much finer grain photo you're seeing much more features and so on and so forth. Besides that we need the language in which the entity is being described, and then most important--and often neglected--, the knowledge and understanding of the world that is assumed.

Imagine for instance being an explorer running across a new hitherto un-contacted Indian tribe in the Amazonian jungle. It is un-contacted, but it speaks a language that's the same as that of a neighboring group that has been contacted. And you have actually learned that language, so you can talk to these people in their own language even though they have not encountered people from the outside directly before. Now your job is to explain to these Indians in their own language what a tax manual’s mutual fund is. You would have to give a lot of background and that is the idea here, the knowledge and understanding of the world is assumed very much influences the characters of the description. Finally, there is the system of coding from the language from which the entity is being described to bit string. Given all that, you can represent the entity by bit string and it is one of many different possible bit strings and one of many possible bit strings of the same length. Now this quantity called Algorithmic Information Content of a bit string or the like of the entity that the bit string is describing. It's the length of the shortest program that will call a given universal computer U to print out the bit sting and then halt. Universal computer of course is either capable of doing any calculation, and has infinite memory, or more plausibly, it has the capacity to create memory when ever needed in order to solve problems. So it can solve any problem, although it may take a very long time. It's something of an idealization in universal computers. When I say solve any problem, it means solve any solvable problem. And of course often it takes so long that it's not practical. So here it is again, they talk about the entity, e, and the bit string, s_e, and we call the A.I.C. the Algorithmic Information Content $K_u(s_e)$ or $K_u(e)$, where K means the Algorithmic Information Content and the u is the particular universal computer. As this gets larger and larger and larger, it becomes more and more independent of the particular universal computer that's in use, so it kind of covariance problems. Now our job is describing effective complexity. Remember, effective complexity was defined roughly as the length of a very concise description of a regularity, not the features treated as random or incidental. So our job is to split K, the A.I.C. of the entity, or the bit string that describes it, into two pieces. The A.I.C. of the regularities and the A.I.C. of the rest of the features treated as random or incidental. Then first part would be the effective complexity, the second we can call the random information. Now we are using A.I.C. here as a way to represent the idea of concise description. What do we mean by the length of a concise description? Well, we can say the A.I.C. Otherwise we might deal with a lot of redundancy. For example, in my book I quote the story of a school teacher, who assigns to her primary school class writing a three hundred word essay on something that recently happened in your household. Now let us assume that there is a student who did what I would have done in those circumstances years ago. Which is to fool around outside all weekend and finally scribble something on Monday morning to please the teacher. So what this ridiculous student wrote was yesterday the neighbors had a fire in their kitchen and I leaned out of the window and yelled: "Fire, fire, fire, fire…" Now that is an example of something that can be compressed. If the teacher had not insisted on a three hundred word essay, he could have said I leaned out of the window and yelled fire two hundred and eighty-three times, the equivalent, but much shorter. So in the same spirit, we're using Algorithmic Information Content to represent the length of a concise description of what we're talking about.

Let us give a couple of examples at both ends of the spectrum so to speak. Lets take a bit string that is perfectly regular, like 1111111111 or 00000000 it has very little A.I.C. because it's so regular. The regularities have very little A.I.C., all they have to say is it's all ones or it's all zeroes and the length. The length of the bit string is about the only thing that is important or otherwise the regularity takes a trivial time to describe, all ones or all zeroes. So it's effective complexity if very low, the effective complexity is the A.I.C. of the regularity. At the other end of the spectrum, we take a bit string with no regularities, an incompressible string. It has maximum A.I.C. for its life because it's incompressible. But the A.I.C. of the regularities is again very low because it doesn't have any regularities, be about just the length and nothing else. So in both cases, at both ends of the spectrum so to speak the A.I.C. of the regularities, the effective complexity is very low. So you can draw the diagram in Figure 1.

The A.I.C. is very high in the middle, that's a string with no regularities at all. Close to the origin O there is a string with very low A.I.C. just all ones or all zeroes. At both ends the A.I.C. of the regularities, the effective complexity has to be very small, very close to zero at both ends. In the middle it can get bigger. So we see that effective complexity can be there in a considerable amount only in the middle, only in this intermediate range between perfect order and perfect disorder. That region is the region where you can have a lot of effective complexity. I say that here, can be high only in the intermediate region between perfect order and perfect disorder. It is not a particular place; it is a sort of anywhere in here.

Now, how do we represent the regularities in the first place, what do we mean mathematically by regularities as opposed to features that are treated as random or incidental? Remember that there is always a judge, not a necessarily human, not necessarily alive, but some kind of judge that makes a distinction between the important and the unimportant. Only then do we define regularities and randomness.

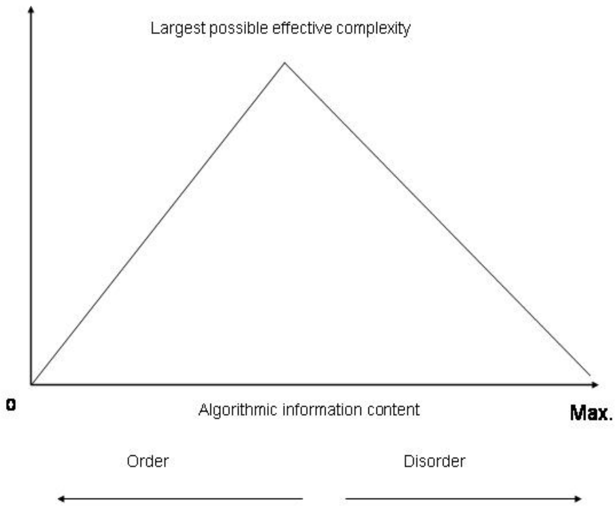

Figure 1. Algorithmic Information Content (A.I.C) and effective complexity

Now we can borrow from statistical mechanics, this way of treating regularities. Regularities of an entity e are represented by embedding the real entity in a set of things, an ensemble Ẽ the rest of which are just imagined. Our entity is real, we talked about it, describing it, but we bury it in a mass of things, all the rest of which have just been made up for the purpose. And we assign probabilities to all the members of the set, making it what's called an ensemble. An ensemble is just a set with probabilities assigned for all members of the set. Here this is written out, to describe the regularities attributed to an entity, embedded it conceptually in a set of entities with probabilities for the member that makes it an ensemble. It's very important though that the entity itself should be a typical member of the ensemble, otherwise the ensemble is hardly describing the entity we're talking about. It has to be a typical member, that is one does not have an abnormally low probability for that set. The probability distribution that will reflect the regularities, that's the way we do it in statistical mechanics, physics, and it seems to be something that quiet generalizing.

This idea of burying the thing we want to talk about, embedding it in a set of things, the rest of which are imaginary and imagined. It's very familiar in the arts, for example in fiction or drama and in many kinds of poetry. What we're doing is constructing an alternative world that we set along side our own. And if there are many works of fiction, many dramas, many poems then it's a whole set of different worlds. You could have imaginary families, but sometimes the novelist or the draftsman makes them so real, that we debate what they would do in various conditions and so we feel we know those people, they represent additional families beyond the ones we actually encounter in the real world. The characters represent additional people that we don't find in the real world, but we find in these books or in the dramas that we see them perform. So it's not an unfamiliar idea even outside of science to understand more the regularities of something by embedding it in a set of things, the rest of which are made up. So how do we do that? Well in statistical mechanics, we always do that. For example, suppose you have a sample of gas, sample of oxygen gas for example. It is a sizable sample, microscopic sample; it may have 10^{24} molecules of oxygen. 10^{24} is a lot: a trillion trillion, trillion trillion molecules in this sample of gas, to specify exactly the state, even classical. Classical you don't have the uncertainty principal to worry about but you have to specify the position in three dimensions of every single molecule. So that there are new molecules in the sample, that's six new coordinates that we have to specify. Up to some accuracy they were good depending on the coarse graining. Where could we possibly get all that information? How could we acquire that information, even with the crude approximation, for the values of the positions and momentum? It is an unbelievable amount of information. If we ever acquired it , where could we store it? If we ever stored it, how could we read it? And if we ever read it, how do we make use of it? So, clearly we don't try to do that. What we do in fact is, in Statistical Mechanics is embed this sample of gas in a huge set of samples, the rest of which are imaginary. And we specify some probability distribution for this set and that's our description then of gas. In this way we can define temperature, we can define pressure, we can define a whole bunch of other variables if we want them and describe the system by specifying these properties only. What we can do then in general, that is a very general procedure. So here I have written it again in Fig. 2. Regularities of an entity can be described embedding e in an ensemble E, as a typical member, and the various elements of this ensemble exhibit the variations in Hypothetical Entities sharing the regularities attributably.

For example, the various states of the gas at certain temperatures. Then we can look at the A.I.C. of the ensemble of the probabilities distributions $\{p_r\}$ which depends for example on how many parameters we are specifying. If it is just temperature and pressure this would be rather low. If we are specifying a lot of things it will be high. So it is a length of a very concise description of the ensemble, in other words, the members and their coarse-grained probabilities. Then we define Y to be the A.I.C. of the ensemble, that's our way of giving a very concise description of the length. So it is a candidate for the effective complexity. Why do I call it a candidate? Because we have not yet specified the ensemble. When we do, Y will be the effective complexity. Now we think the other quantity we are interested in is the ignorance or information or intrepid or whatever we want to call it. It has various names. If you have various outcomes of things and they all have the same probability then that probability is one over the number of possible outcomes, $1/\{p_r\}$. The number of bits of information in it is $\log_2(1/p_r)$, the number of possibilities. So if you have three flips of the fair coin, for example you have eight possibilities that are the same probability and the number of bits is three. If you have five flips of the fair coin then the number of bits is five and so on and so forth. It's the law of the base 2 of number of possibilities and therefore the law of the base 2 of 1/probability, the common probability. Now if the probabilities are different, then we have the average. I, the ignorance, is the average over the probabilities of the law: $I=\Sigma\, p_r \log_2(1/p_r)$. It is ignorance, but can also be given as information. Ignorance and information are exactly the same thing. The ignorance of a letter is exactly the same thing as the information you derive from when you open a letter. It just depends whether you are talking about the time before you open the letter or the time after you open the letter. Before you open it, it's ignorance. After you open it, it's knowledge, information. So ignorance or information: fortunately they both begin with the same letter I. Now of days people are experimenting very much with alternative formula patterns, but I am not going to go into that here today. Mostly this is the right formula, it is certainly the one used in campus. So we have then two quantities we're concerned with, the A.I.C. in the ensemble, and that is the effective complexity once we choose the right ensemble, and the other is the ignorance or information, I, which will be the random information once we have chosen the correct ensemble (Fig. 2).

Now suppose that what we are doing is looking at a set of data and trying to make up a theory to account for those data, that is something many of us do, those of us that who at least part time theoreticians do that all the time. Now to explain a given set of data, we can go to two extremes, though we try to avoid those extremes, one is to make the theory have a huge number of parameters and a huge number of bells and whistles so that it really specifies those data and not any other data. Well that is very good in zeroing in on the data, but it is not very good as a theory because it has a huge number of adjustable parameters… used to say that way you could fit in an elephant. So we do not want that to be, we do not want the Y to be to large. If the ensemble is very complicated with lots of bells and whistles and lots of parameters, Y would be much too large. We also do not want I to be to large, because I is the ignorance, and we do not want a huge amount of ignorance. In other words we do not want the theory to predict the data that we see, but numbers of other sets of data at the same time. We do not want the theory to allow almost any set of data, which would not be very good either. But we can trade of between I and Y because what we can do if the theory has a lot of ignorance, we can add enough bells and whistles so as to get rid of the ignorance but then we have a huge amount of Y, which we do not want.

The other way around also, if we lots and lots of bells and whistles and extra parameters, we can get rid of a lot of them but often at the cost of increasing the ignorance, in other words not zeroing it very well on the particular data that we have in hand. Since there are all these possible trade offs, what we really need is to minimize the sum of the two. Try to make K=Y+I as small as possible. Later we can worry about the trade offs. Occasionally the theorist is in a very fortunate situation, a win win situation, where he's reducing both I and Y together. For example, James Clark Maxwell, when he formulated Maxwell's equation of the electromagnetic field was doing a great job both ways, he was simplifying the equation and he was also zeroing in better on the data, the electromagnetic fields. So occasionally we can do that, but also we have to make use of trade offs. In any case what we want to do is minimize the sum of the two, the ignorance plus the Algorithmic Information Content, plus the effective complexity. Now if the ensemble consisted of just the entity itself, we aren't burying it in the midst of a whole bunch of other things, we're just looking at the entity itself, one element, it's a one element ensemble, with probability one for that element and zero for anything else. Then Y would just be K, where K is the Algorithmic Information Content of the entity because that's all there is in the ensemble. So the Algorithmic Information Content of the ensemble is the same as that of the entity, it's just K, and I is zero in that case, because there is just one entity and that's all we have. So for the probability 1, p is just zero. Now it turns out that K, which is equal to Y+I in this very special case, is actually the minimum possible value of this of Y+I. We couldn't have a smaller value. But there are many different ensembles that have obtained that minimum least roughage. We have to decide which one to use.

So here we are, we have information or ignorance which is also proportionate to entropy, we have the Algorithmic Information Content of the ensemble which is the candidate for effective complexity and we have the, this is the candidate for the random information if we choose the right ensemble. This is the candidate for effective complexity if we choose the right ensemble. We talked about the sum of the two, which we can call total information. We want that to be its minimum and we already said what the minimum is, it's just K. So here's what we do, to choose the ensemble, we're finally getting down to choosing it, so we're actually describing our entity. We keep the total information equal to its minimum value, which is K the A.I.C. of the entity itself, and we maximize I and minimize Y, while keeping fixed particular other quantities beside this, the ones that are judged to be important. And what those are, are the average values, the ensemble averages of various quantities. For example, in statistical mechanics we usually keep fixed the ensemble average of the energy, and if we do that we define the absolute temperature. It's just one parameter in the Maxwell-Boltzmann distribution, E to the minus the energy, that's a constant and in that case the only quantity judged to be important besides I+Y is just the energy, the average energy. Fix the average energy, otherwise we maximize I and minimize Y and we get the maximum of both sides. We can, if we choose to treat a lot of things as important and keep them all fixed while maximizing the ignorance, the measure of ignorance, which is maximizing the entropy. Then we are proceeding just the way we do in statistical mechanics, we are maximizing the ignorance and keeping certain important quantities fixed. And we have the average energy held fixed, and we get this reduced distribution, the Maxwell distribution. So here is what we have in a diagram, the final diagram, describing what we are doing. I am plotting Y versus I, where Y is the candidate for the effective complexity, I is the candidate for the random information (the measure of ignorance). We want to hold Y+I=K, that is this straight line here called minus one fixing Y plus I equals K. As we go down to very small I and very large Y, this line ends in the allowable space of ensemble, one resolves here, this line one resolves here toward the maximum. So this is as far as we can go, we want to minimize or maximize I while staying on this line; we have to come down to this point. Provide we don't keep anything else fixed, but if we keep other things fixed like the average energy, the average this and the average that and so on, we can gradually move up this line further, further and further attributing more and more and more properties to our ensemble and knowing more and more and more therefore about our energy. The ignorance is being reduced, effective complexity is being increased, because we have a more and more complex description of the entity but it involves less and less ignorance, we go up here so. Now the problem with being down here at the bottom not specify nothing besides Y+I equals K, then maximizing Y and minimizing I going to this point. The problem is: this point lies very low. We have shown, many

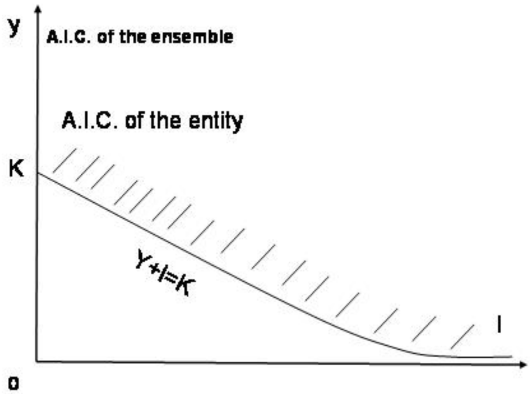

Figure 2. Ingorance, Information, and Effective complexity.

people have shown that this point lies very, very low, which means the effect complexity if very low here. Everything is simple, everything would come out simple, and it’s not really satisfactory, what's the use of a theory of simplicity and complexity where everything is simple. So we do not want that, in fact we keep a number of quantities fixed and move up this line to some reasonable place like this. Now the Great Russian mathematician Kolmogorov, one of the people who invented A.I.C., fell into this trap. He called it the *minimum sufficient statistics*, but one of his students, David, who is now a famous professor of computer science at Boston University, told him that it is wrong: “not good everything simple”, but Kolmogorov would not listen. So he kept talking about this minimum sufficient statistic here, where as in fact we need to keep a lot of important quantities fixed, while maximizing I and minimizing Y and keeping this on this straight line. We have to keep a number of things fixed and then we move up here. We have a more complex description than with one of less ignorance. So here we have the mathematics of describing effective complexity.

We just maximize the measure of ignorance while holding this at its minimum value and Y always comes out tiny, that is that point with the curve departs from the straight line, everything comes out simple. In one case out of the universe, it is ok. But in every other situation it is no good. Here we are, say this once more, just for emphasis. More regularities recognized, the entity is allowed more individuality instead of being treated merely a statistic, move up the straight line just a bit, your characterizing the number of properties of the entity not just say the temperature but the number of those. Now another way look this whole matter which may appeal to many of you, another way to look at Y and I as that in terms of the program, you remember A.I.C. was the length of the shortest program, it causes the universal computer to print out the bit string under consideration and that's all. But we can think of a program in two parts, a basic program and then another program that supplies the data for the basic program. We can divide it into two parts, a more fundamental program and then data that you fed into it. You can think of Y as the length of the basic program and I is the length of the program that feeds in data to the system. So we have a general program that can accommodate various data sets and I will describe that message to be condensed. For most people that's a more congenial way of talking about this. Now the judge can come up with all sorts of different ideas about what is important and what is not. Suppose you're a anthropologist for example, then the kind of complexity you're interested in is social complexity. How many roles are there in society, how many different professions are there in this society and over. How complicated are those professions, do they have lawyers in that society, do they have doctors in that society, do they have witch-doctors in that society. How complicated is the stuff the lawyers deal with, or the witch-doctors deal with. That's social complexity. And of course a lot depends on the data that are available. Lets take Mars for example, the planet Mars, before the voyages by NASA and then after them. In the old days there was an astronomer looking through a telescope, peering through a telescope hoping for a night of good seeing, when the boiling of the atmosphere wasn't so bad and trying to get a clear picture of Mars and understand what kind of features there were on the surface of Mars. The astronomer Perceval Loel, who worked in Arizona, but obviously his ancestors were from Boston, who claimed he could on nights of especially good seeing spot lines on Mars, straight lines. Now back in 1877 they have seen lines on Mars also, and he called them *canali*, channels in Italian. But Perceval Loel interpreted them as canal dubbed by intelligent creatures. How could you possibly get a straight line on Mars, if there were not intelligent organisms some how engineering these straight lines? They must be canals and there must be intelligent Martians. NASA photos even the early ones, which were not very fine grained, showed clearly that this was all imaginative; this was just the human mind seeking patterns, which it does as you know. Perceval Loel was just looking for a pattern, and he could see these straight lines, which the earliest NASA pictures show was not there. They were just pieced together with bits of other things. Now ever since the time of Linden Johnson, the Vice President of The United States, who has been interested in the space program and has been chairman for the space council and when Dan Quail became vice president, he assumed that role also and became very enthusiastic about manned voyages to Mars. And he was invited to open a NASA meeting on that subject, give a little talk at the beginning of the meeting. So somebody in NASA prepared a nice two page speech for Dan Quail. And the first page was about the old Mars, Perceval Loel, the canal, the intelligent Martians and so on. The second page was about all the things that NASA had discovered to be actually true of the surface of Mars. The old stuff turned out to be wrong but unfortunately Dan Quail lost the second page. So his entire speech was about the old Mars, the canals, the intelligent Martians and so on and so forth. You see how much difference greater course graining and greater fine graining can be. Now the latest wrinkle of course is astrobiology which is very exciting and has to do with what life or something like it would look like on other planets and evidence could possible be collected to bear about that. It seems extremely likely to me that there must be life on other planets, because there are so many planets in the universe. The number of stars is gigantic as you know and planets are not unusual. When I was a undergraduate, the usual theory then was that planets had to be created in a rare triple collision star, but very shortly afterwards scientist went back to the old time investigatable hypotheses of the late eighteenth century, which was a condensation of dust, the big condensation given to the star and the little condensation to the planet, very simple, that's what people believe today. And it seems to be very common, now that people are actually able to see large planets, discovering a lot of which of course theory predicted along time ago and when they are able to see smaller planets they will find a lot of others too. Life does not seem to be anything very special, once you have the right conditions as we did on the earth some four billion years ago, life presumably springs up, or something very life like, we don't know how to exactly define it and it's up to the astrobiologist to define it as that. So it is difficult but not crazy to think about what form, how different life could be. How much depends on accidents, which we discussed so much this morning and how much depends on the fundamental laws of physics. How much of biochemistry is rooted in physics and how much is the result of a historical accident. People at the Santa Fe Institute are very anxious, very eagerly pursuing that question and it's crucial for astrobiology. The most interesting question is about intelligence life, that is, if you define us to be intelligent, which I am not sure I would. And it seems to me there too, that there is likely intelligent life because we know evolution works toward more and more complex things, usually and solve more and more difficult problems, provided the relevant creatures don't kill themselves. So I think very, very, likely in an enormous number of planets in the universe, there are others that contain what we would call intelligent life. But Pogo said the last word in this subject, Pogo is the cartoon character called Kelly who lived in the Okey Ponokey Swamp on the border of Florida and Georgia. Along with a lot of other swamp creatures, alligators, bugs and so on. And one day Pogo was talking to one of the other swamp creatures and said, “Out there on some planet orbiting some other star, there maybe entities that are more intelligent than we are, we humans or on the other hand, maybe we humans are the most intelligent entities in the universe. Either way it's a mighty sobering thought.