Murray Gell-Mann
Santa Fe Institute, 1660 Old Pecos Trail, Suite A, Santa Fe, NM 87501 and Los Alamos National Laboratory, Los Alamos, NM 87545

Complex Adaptive Systems

INTRODUCTION

The various groups at the Santa Fe Institute studying complex adaptive systems (CAS) have somewhat different points of view and have adopted different vocabularies. Some of us speak of "artificial life" or "artificial social life" or "artificial worlds," while others, of whom I am one, prefer to consider natural CAS and computer-based systems together. The latter include methods for adaptive computation as well as models and simulations of natural CAS.

Even the term CAS has different meanings for different researchers. As one distinguished professor remarked, "a scientist would rather use someone else's toothbrush than another scientist's terminology." For example, my nomenclature differs from that of John Holland, from whom I have learned so much. He calls something a CAS only if it is a collectivity of interacting adaptive agents, each of which I would refer to as a CAS. Likewise, John uses the term "internal model" to mean what I call a schema.

There are additional possible sources of misunderstanding as well, stemming from the relation between computer-based and natural systems. At one of our Science Board Symposia, a speaker asked, "Are we using computation as an aid in understanding biology (e.g., evolution, thinking, etc.) or are we using biology as

a metaphor for work on computation?" That is an important question. At some institutions where computation and neural systems are studied, there is real confusion on this issue. For example, success in designing a computing system based on "neural nets" is sometimes taken as evidence that such nets furnish a serious model of the human brain, with the units or nodes corresponding to individual neurons.

I favor a comprehensive point of view according to which the operation of CAS encompasses such diverse processes as the prebiotic chemical reactions that produced life on Earth, biological evolution itself, the functioning of individual organisms and ecological communities, the operation of biological subsystems such as mammalian immune systems or human brains, aspects of human cultural evolution, and adaptive functioning of computer hardware and software. Such a point of view leads to attempts to understand the general principles that underlie all such systems as well as the crucial differences among them. The principles would be expected to apply to the CAS that must exist on other planets scattered through the universe. Most of those systems will of course remain inaccessible to us, but we may receive signals some day from a few of them.

As to successful adaptive computational methods and devices, we have examples such as neural net systems, based on a perceived similarity, even though it may be rather remote, to the functioning of the human brain, and genetic algorithms, based on a resemblance to evolutionary processes. Surely these sets of methods belong, together with many others, mostly as yet undiscovered, to a huge class of computational CAS, with common features that will be well worth identifying and understanding. Some of the new computational methods may exhibit similarities to the operation of natural CAS that we know, such as the immune system, but others may be quite unlike any natural process familiar to us.

A CAS gathers information about its surroundings and about itself and its own behavior, at a certain level of coarse graining. The time series that represents this information can sometimes be approximated by a steady one, although in general it is changing with time, frequently in ways that depend on the system's behavior, and the surroundings are often coevolving. The following[1,2] are general characteristics of a CAS:

1. Its experience can be thought of as a set of data, usually input $\rightarrow$ output data, with the inputs often including system behavior and the outputs often including effects on the system.
2. The system identifies perceived regularities of certain kinds in the experience, even though sometimes regularities of those kinds are overlooked or random features misidentified as regularities. The remaining information is treated as random, and much of it often is.
3. Experience is not merely recorded in a lookup table; instead, the perceived regularities are compressed into a schema. Mutation processes of various sorts give rise to rival schemata. Each schema provides, in its own way, some combination of description, prediction, and (where behavior is concerned) prescriptions for action. Those may be provided even in cases that have not been encountered

before, and then not only by interpolation and extrapolation, but often by much more sophisticated extensions of experience.

4. The results obtained by a schema in the real world then feed back to affect its standing with respect to the other schemata with which it is in competition.

Now the feedback process need not be a clear-cut one in which success is well defined and leads to survival of the schema while failure, equally well defined, results in its disappearance. Fitness may be an emergent or even an ill-defined feature of the process; the effect on the competition among schemata may be only a tendency; and a demoted schema may be kept for use in a subordinate capacity or retained in memory while not utilized (it might, after all, produce useful variants). The important thing is the nature of the selection pressures exerted in the feedback loop, whether or not they are expressible in terms of a fitness function. (Similarly, physical forces may or may not be derivable from a well-defined potential.)

An excellent example of a CAS is the human scientific enterprise, in which the schemata are theories, giving predictions for cases that have not been observed before. There is a tendency for theories that give successful predictions (and exhibit coherence with the body of successful theory) to assume a dominant position, although that is by no means a simple, mechanical procedure. Older, less successful theories may be retained as approximations for use in restricted sets of circumstances. Even wrong theories are not necessarily wholly forgotten, since they may inspire some useful theoretical work in the future.

In its application to the real world, a schema is in a sense reexpanded, reequipped with some of the arbitrariness of experience, some of the random material of the kind that was stripped away from the data when regularities were identified and compressed. For instance, a theory must be combined with boundary conditions in order to give a prediction. The additional data adjoined to the schema may simply be part of the continuing stream of incoming data, which contain, in general, the random along with the regular.

In most CAS the level of the schemata and the level at which results are obtained in the real world are entirely distinct. In the realm of biological organisms, that is the distinction between genotype and phenotype, where the phenotype depends not only on the genotype but on all the accidents of development that intervene between the DNA and the adult organism. However, in some cases, such as Tom Ray's world of digital organisms, the genotype and phenotype are not physically different, but distinguished only by function. His sequences of machine instructions play both roles. As Tom Ray remarks, certain theories of the origin of life on Earth assert that RNA once behaved that way, both as bearer of information and as agent of chemical activity, before the appearance of organisms exhibiting separate genotype and phenotype.

Some new computer simulations of evolution try to include distinct genotypic and phenotypic levels. One that is under development at UCLA even simulates sexual reproduction, with haploid and diploid generations, and tries to test William Hamilton's idea that the principal utility of the male lies in helping to outrace

enemies, especially parasites, by providing the offspring with genetic diversity that would be lacking in parthenogenesis.

Complex adaptation is to be contrasted with simple or direct adaptation, as in a thermostat, which just keeps mumbling to itself, "It's too cold, it's too cold, it's too hot, it's just right, it's too cold," and so forth. In the 1940s, the chemist Cyril (later Sir Cyril) Hinshelwood put forward a direct adaptation theory of the development of bacterial resistance to drugs. Genetic variation and selection were rejected in favor of a straight negative feedback process in chemical reactions in the cell. The drug interfered at first with the chemistry of the cell, but then the deleterious effects were mitigated as a result of reaction dynamics, and the mitigation was transmitted mechanically by the bacteria to their progeny in the course of cell division. There was no compression of regularities, no competition of schemata.

Hinshelwood's theory lost out, of course, but it has not been totally forgotten, and it now serves my purpose as an example of direct adaptation rather than the operation of a CAS. Direct control mechanisms are common in nature and in human industry, and they formed the subject matter of cybernetics half a century ago.

The cybernetic era was followed by the era of the expert system, employing a fixed "internal model" designed using the advice of experts in a field, for instance a decision tree for medical diagnosis. The expert system did not learn from the results of its work, however. It remained fixed until it was redesigned. (Only if the human redesigners are included can the expert system be regarded as a CAS, of the kind that involves "directed evolution" or "artificial selection," with humans in the loop.) The new era of CAS in robotics and other such fields is the age of constructed systems that actually learn, by formulating schemata subject to variation and to selection according to results in the real world.

It is useful to distinguish various levels of adaptation. In particular, we can take the example of human societies, where a schema is a set of customs, traditions, myths, laws, institutions, and so forth, what Hazel Henderson calls "cultural DNA." (The biologist Richard Dawkins has invented the word "meme" for a unit of that DNA analogous to a gene.)

The schemata include prescriptions for collective behavior. A culture operating on the basis of a given schema reacts to altered circumstances such as climatic change, invasion, and so forth, in ways prescribed by that schema. If the climate turns warmer and drier, the response of a group of villages may be to move to higher elevations. In the event of attack by outsiders, the inhabitants of all the villages may retire to a fortified site, stocked with food and water, and sustain a siege. What happens at this level is something like direct adaptation.

On the next level, the society may change its schema when the prevailing one does not seem to have given satisfactory results. Instead of migration to the high-lands, the villagers may try new crops or new methods of irrigation or both. Instead of retreating to a fort, they may respond to invasion with a counterattack aimed at the enemy's heartland.

Finally, there is the level of Darwinian survival of the fittest (as in population biology). In some cases, not only does a schema fail, but the whole society is wiped out. (The individual members need not all die, but the society ceases to exist as a functioning unit.) At this level the successful schemata are the ones that permit the societies using them to survive.

Not only are these three levels of adaptation distinct, but the time scales associated with them may be very different. Nevertheless, discussions of adaptation in the social science literature sometimes fail to discriminate among the levels, with unfortunate results for clarity.

The disappearance of societies is somewhat analogous to the death of organisms or to the forgetting of ideas. Such phenomena are, of course, universal and not unrelated to the second law of thermodynamics. Still, over a given period of time, the importance of mortality can vary from one domain to another.

In cases where death is very important at the phenotypic level, a crucial measure of success for a schema is phenotypic survival, and reproduction assumes great significance. Moreover, population can then supply a rough quantitative measure of fitness. In biology, one often follows the population of a cluster of genotypes such as a species or subspecies, and the clustering phenomenon is itself of very great interest. One can also follow subpopulations characterized by particular alleles of certain genes.

By contrast, there are situations where death is comparatively unimportant, whether at the genotypic or the phenotypic level. One schema can dominate another without the losing one disappearing; reproduction is not of overwhelming interest; and population is not of critical importance as a measure of fitness. Consider individual human thinking, for example. If we try to grasp an issue more clearly than before, we may succeed in getting an idea that dispels a great deal of previous confusion and displaces, to a considerable extent, earlier ideas. (That is not so easy, by the way, because existing ideas entrench themselves and we have a tendency to interpret new information as confirmatory, so that we dig ourselves deeper and deeper into what may be a quite unsuitable hole.) Over time scales such that forgetting is not a crucial factor, replication and population are not particularly relevant concepts to the success of an idea in the thinking of an individual person. What matters most is that at the real world level one idea has received more positive feedback than another and thus assumed a comparatively dominant position. Over a very long time scale, of course, every system eventually has to get rid of clutter in some way, so that erasure, forgetting, or some other kind of grim reaper has to come into the picture.

Looking at CAS overall, we see that fitness is a rather elusive concept when it is endogenous. If an exogenous criterion is supplied, as in a machine that is designed and programmed to win at chess, then of course the feedback loop involves a well-defined fitness. But when fitness is emergent, it is not so easy to define without a somewhat circular argument in which whatever wins is fit by definition, and whatever is fit is likely to win.

As everyone recognizes, fitness is even less well defined when it is acknowledged that the surroundings of the system are themselves undergoing change and often coevolving. In the latter case, fitness "landscapes," even to the extent that they could be defined for fixed surroundings, now give way to a picture of shifting and interdependent landscapes for the different adaptive components of the total system.

The greatest difficulty in discussing features of a system that are "adaptive" (or that render it "fit") is the distinction between what is adaptive and what has resulted from a process of adaptation. The latter may often be maladaptive. Let us discuss some common reasons for that.

The simplest reason is, of course, that a CAS engages, under the influence of selection pressures in the real world, in a search process over the abstract space of schemata that is necessarily imperfect. Even if fitness is well defined, a system that merely searches for local maxima by "hill climbing on a landscape" would most often get stuck on a molehill. To have a chance to find mountains nearby, the search process must include other features, such as noise (but not too much noise) or else pauses in climbing to allow for free exploration. Naturally, schemata that are more or less maladaptive are often selected.

Apparently maladaptive schemata often occur for another reason, namely that the system is not defined broadly enough to encompass all the important selection pressures that are operating on the schemata concerned. For example, in the scientific enterprise, it would be a mistake to ignore the pressures other than purely scientific ones that affect the viability of a schema, especially in the short run. Scientists often exhibit human frailty, and issues of jealousy, greed, and the misuse of power may play a role in the fate of theories; even observational data are occasionally falsified. Of course it is equally foolish to exaggerate the importance of these extra-scientific selection pressures and to ignore the powerful correcting effect that comparison with nature keeps supplying.

The prevalence of prescientific theories, such as those associated with sympathetic magic, provide even more striking examples of the breadth of selection pressures. Suppose the members of a tribe believe in the efficacy of bringing rain by pouring out on the ground water obtained in a special place in the mountains. Clearly it is not carefully controlled comparison with results that sustains faith in the procedure, but selection pressures of very different kinds. For instance, the authority of powerful individuals or groups may be enhanced by the prevalence of belief in the ceremony, which may, in addition, be part of a whole set of customs that cement the bonds holding the society together.

More generally, it is significant that any CAS is a pattern-recognition device that seeks to find regularities in experience and compress them into schemata. Often it will find fake regularities where there is in fact only randomness. A great deal of superstitious belief can probably be attributed simply to that effect, which might be labeled the "selfish schema." (I have already mentioned how new data are often interpreted so as to strengthen an existing belief.)

Of course, a CAS will often err in the other sense and overlook regularities. Both types of error are presumably universal. In the realm of human beliefs, overlooking obvious regularities can usually be identified with denial. It is striking that in human beings both superstition and denial are typically associated with the alleviation of fear: in the former case fear of the random and uncontrollable and in the latter case fear of regularities that are all too evident, like the certainty of death.

Another example of the breadth of selection pressures comes up in studying the evolution of human languages. Here one should first of all distinguish several different CAS, at different levels and on different time scales. One is the evolution, over hundreds of thousands or millions of years, of the biological capacity to use languages of the modern type. Another is the evolution of those languages themselves, over thousands or tens of thousands of years. Yet another is the learning of a native language by a child. Consider the second of these three systems, concentrating for example on the evolution of grammar and phonology. One encounters, of course, the usual mixture of fundamental rules (in this case the "innate" constraints on grammar and phonology determined by biological evolution), frozen accidents or founder effects (in this case arbitrary choices made in ancestral languages that may have been transmitted to their descendants), and what is adaptive (in this case features that make for more effective communication). However, the selection pressures in linguistic evolution are not wholly linguistic. A great deal depends on whether a people speaking one language is more advanced culturally or stronger militarily than a people speaking another language. Such matters may easily have a greater effect on the fates of the two languages than which one is more convenient for communication.

Another common reason why maladaptive features arise from a process of adaptation is that time scales are mismatched. When circumstances change much more rapidly than the response time of the CAS, traits occur that may have been adaptive in the past but are so no longer. For instance, global climate change on a scale of a few decades will not permit the same kind of ecological adaptation that would be possible in the case of much slower change.

The human tendency to form groups that don't get along with one another, based on what are sometimes rather minute differences that an outsider would barely perceive, may be to a considerable extent an inherited tendency, even though it is fortunately subject to modification through culture. If a hereditary component is really involved, it may have been adaptive under the conditions that prevailed many tens of thousands of years ago. For example, it could have served to limit the size of the population in a given area to a number that the area could support. Nowadays, in a world of destructive weapons, the tendency seems quite maladaptive.

The phenomenon of imprinting provides an extreme case of the mismatch of time scales. A greylag goose that glimpsed Konrad Lorenz instead of its mother when it was first hatched was condemned to treat Lorenz as its mother ever after. The process of imprinting, which works fine in the more common case when the

gosling sees its real mother, compromised forever the chances of a normal goose life for any gosling that saw Lorenz instead.

A milder phenomenon is that of windows of maturation. Béla Julesz emphasizes that certain abnormalities in vision have to be corrected early in childhood if they are to be corrected at all. In the case of learning deficits, it is important for public policy to know the extent to which they must be remedied during the first two years or so of life and the extent to which plasticity of the central nervous system permits them to be dealt with later by such programs as Head Start. (Of course the chances of success of Head Start are in any case compromised if the duration and intensity of the program are insufficient, as they often are.)

We must pay attention to time scales for other reasons as well. Fundamental rules on one scale of space and time may reveal themselves to be the results of frozen accidents on a larger scale. Thus the rules of terrestrial biology (such as the occurrence of DNA based on the nucleotides abbreviated A, C, G, and T) may turn out to represent just one possibility out of very many. On a cosmic scale of space and time, the earthly rules would then have the character of a frozen accident or founder effect. That is already widely believed to be the case for the occurrence of certain right-handed molecules in important biological contexts where the corresponding left-handed molecules do not occur. (Attempts to derive that asymmetry from the left-handedness of the weak interaction for matter, as opposed to antimatter, do not seem to have succeeded.)

Some of the most interesting questions about CAS have to do with their relations to one another. We know that such systems have a tendency to spawn others. Thus biological evolution gave rise to thinking, including human thought, and to mammalian immune systems; human thought gave rise to computer-based CAS; and so on. In addition, CAS are often subsystems of others, as an immune system forms part of an organism. Often, a CAS is a collectivity of adaptive agents, each a CAS in its own right, constructing schemata describing one another's behavior. One of the most important branches of the emerging science of CAS concerns the inclusion of one such system in another and the functioning of collectivities such as ecological communities or markets.

One class of composite CAS of particular interest consists of natural or computer-based systems with human beings "in the loop," as in the breeding of animals or plants (what Darwin called artificial selection as opposed to natural selection) or as in a computer system that creates pictures by presenting a human being with successive choices of alterations in an initial pattern.

Pure computer-based CAS can be used for adaptive computation, for modeling or simulating in a crude fashion some natural CAS, and for study as examples of CAS. In all three capacities, they illustrate that astonishingly great apparent complexity can emerge from simple rules, alone or accompanied by a stochastic process. It is always a fascinating and useful exercise to try to prune the rules, making them even simpler, while retaining the apparent complexity of the consequences. Such investigations will gradually lead to a mathematical science of rules

and consequences, with theorems initially conjectured on the basis of examples and then proved.

Applications to natural or behavioral sciences require, at a minimum, not just those abstract propositions about rules and consequences but also additional information specifying situations simulating in some convincing way ones that arise in the science in question.

Still more information must be supplied if the computer model is to have any relevance to policy. Conditions prevailing on the planet Earth, including human institutions as well as features of the biosphere, have to be at least vaguely recognizable in the model. Even then, it is critical to use the results mainly as "prostheses for the imagination" in forecasting or in discussing policy options. Trying to fit policy matters into the Procrustean bed of some mathematical discipline can have most unfortunate consequences.

It is a major challenge to the Santa Fe Institute to try to construct bridges connecting these different levels of abstraction, while maintaining the distinctions among them.

When we ask general questions about the properties of CAS, as opposed to questions about specific subject matter such as computer science, immunology, economics, or policy matters, a useful way to proceed, in my opinion, is to refer to the parts of the CAS cycle,

I. coarse graining,

II. identification of perceived regularities,

III. compression into a schema,

IV. variation of schemata,

V. application of schemata to the real world,

VI. consequences in the real world exerting selection pressures that affect the competition among schemata,

as well as to four other sets of issues:

VII. comparisons of time and space scales,

VIII. inclusion of CAS in other CAS,

IX. the special case of humans in the loop (directed evolution, artificial selection), and

X. the special case of composite CAS consisting of many CAS (adaptive agents) constructing schemata describing one another's behavior.

Here, in outline form, is an illustrative list, arranged according to the categories named, of a few features of CAS, most of them already being studied by members of the Santa Fe Institute family, that seem to need further investigation:

I. Coarse Graining

1. Tradeoffs between coarseness for manageability of information and fineness for adequate picture of the environment.

II. Sorting Out of Regularities from Randomness

1. Comparison with distinctions in computer science between intrinsic program and input data.
2. Possibility of regarding the elimination of the random component as a kind of further coarse graining.
3. Origin of the regularities in the fundamental laws of nature and in shared causation by past accidents; branching historical trees and mutual information; branching historical trees and thermodynamic depth.
4. Even in an infinite data stream, it is impossible to recognize all regularities.
5. For an indefinitely long data stream, algorithms for distinguishing regularities belonging to a class.
6. Tendency of a CAS to err in both directions, mistaking regularity for randomness and vice versa.

III. Compression of Perceived Regularities into a Schema

1. If a CAS is studying another system, a set of rules describing that system is a schema; length of such a schema as effective complexity of the observed system.
2. Importance of potential complexity, the effective complexity that may be achieved by evolution of the observed system over a given period of time, weighted according to the probabilities of the different future histories; time best measured in units reflecting intervals between changes in the observed system (inverse of mutation rate).
3. Tradeoffs between maximum feasible compression and lesser degree that can permit savings in computing time and in time and difficulty of execution; connection with tradeoffs in communication theory—detailed information in data base versus detailed information in each message and language efficiency versus redundancy for error correction.
4. Oversimplification of schema sometimes adaptive for CAS at phenotypic (real world) level.
5. Hierarchy and chunking in the recognition of regularities.

IV. Variation of Schemata

1. In biological evolution, as in many other cases, variation always proceeds step by step from what already is available, even when major changes in organization occur; vestigial features and utilization of existing structures for new functions are characteristic; are there CAS in which schemata can change by huge jumps all at once?
2. Variable sensitivity of phenotypic manifestation to different changes in a schema; possibility in biological case of long sequences of schematic changes with little phenotypic change, followed by major phenotypic "punctuations;" generality of this phenomenon of "drift."
3. Clustering of schemata, as in subspecies and species in biology or quasispecies in theories of the origin of life or word order patterns in linguistics—generality of clustering.
4. Possibility, in certain kinds of CAS, of largely sequential rather than simultaneous variants.

V. Use of the Schema (Reexpansion and Application to Real World)

1. Methods of incorporation of (largely random) new data.
2. Description, prediction, prescribed behavior—relations among these functions.
3. Sensitivity of these operations to variations in new data.

VI. Selection Pressures in the Real World Feeding Back to Affect Competition of Schemata

1. Concept of CAS still valid for systems in which "death" can be approximately neglected and reproduction and population may be correspondingly unimportant.
2. Exogenous fitness well-defined, as in a machine to play checkers; when endogenous, a elusive concept: attempts to define it in various fields, along with seeking maxima on "landscapes."
3. Noise, pauses for exploration, or other mechanisms required for the system to avoid getting stuck at minor relative maxima; survey of mechanisms employed by different systems.
4. Procedures to use when selection pressures are not derivable from a fitness function, as in neural nets with (realistic) unsymmetrical coefficients.
5. Possible approaches to the case of coevolution, in which the fitness concept becomes even more difficult to use.
6. Situations in which maladaptive schemata occur because of mismatch of time scales.
7. Situations in which maladaptive schemata occur because the system is defined too narrowly.

8. Situations in which maladaptive schemata occur by chance in a CAS operating straightforwardly.

VII,VIII. Time Scales; CAS Included in Others or Spawned by Others

1. Problems involved in describing interactions among CAS related by inclusion or generation and operating simultaneously on different levels and on different time scales.

IX. CAS with Humans in the Loop

1. Information about the properties of sets of explicit and implicit human preferences revealed by such systems.

X. CAS Composed of Many Coadapting CAS

1. Importance of region between order and disorder for depth, effective complexity, etc.
2. Possible phase transition in that region.
3. Possibility of very great effective complexity in the transition region.
4. Possibility of efficient adaptation in the transition region.
5. Possibility of relation to self-organized criticality.
6. Possible derivations of scaling (power law) behavior in the transition region.
7. With all scales of time present, possibility of universal computation for the system in the transition region.

ACKNOWLEDGMENTS

REFERENCES

1. Gell-Mann, Murray. "Complexity and Complex Adaptive Systems." In *The Evolution of Human Languages*, edited by M. Gell-Mann and J. A. Hawkins. Santa Fe Institute Studies in the Sciences of Complexity, Proc. Vol. X. Reading, MA: Addison-Wesley, 1992
2. Gell-Mann, Murray. *The Quark and the Jaguar.* New York: W. H. Freeman, 1994 (in press).

DISCUSSION

COWAN: If the language is dynamic and adaptive, to what extent does the grammar permit you to describe such changes?

GELL-MANN: I don't know exactly how to answer that, but I can answer a related question that may be of interest.

There are a lot of things called grammatical universals, rules that are true of grammars of all known human languages. They are usually of the "if-then" form, and they often refer to things like word order. Others are very simple and are not of the "if-then" form. For example, all known human languages have pronouns, and every known human language has a genitive construction of some kind. Some rules are like the following. If a language has a special form for *three* objects, then you may be sure it also has a special form for *two* objects. There's no language with a singular, a form for three objects, and a plural, without a form for two objects as well.

These grammatical universals are supplemented by grammatical near-universals, or statistical universals, which start out "For all known human languages except one or two, the following rule holds. . . ." Now, some of the people who study universals are very pure, rigid, dogmatic Chomskyans, who concentrate on innate, preprogrammed, biologically evolved rules and insist that grammatical universals must reflect only those rules. Other linguists study the utility of some of these rules in communication, and they say many of them could simply keep arising in the course of linguistic evolution. Other universals may be the results of frozen accidents. If modern human languages go back only something of the order of 10^5 years, then they may have features that are simply inherited from a unique acestral tongue.

All of these three mechanisms may be present, as well as mixtures of them. But the most important property of linguistic evolution is that language isn't a closed system. A society has lots of characteristics associated with it besides language,

and the death of a language may be caused by events that have nothing to do with language as all. Linguistic evolution is a complicated business, and people who try to simplify it, by looking at only one aspect, are doing the subject an injustice.

WALDROP: In terms of taking experience and compressing that into a schema as opposed to a look-up table: look-up tables are fast, whereas expanding a schema can be slow, depending on the processing power of the machine you're doing it on.

GELL-MANN: Correct. I refer you to my abstract where I wrote the following (this is under "compression of perceived regularities"): "2. Tradeoffs between maximum feasible compression and lesser degree of compression which can permit savings in computing time and in time and difficulty of execution. Connection with tradeoffs in communication theory, detailed information in the data base, versus detailed information in each message. Language efficiency versus redundancy for error correction." Here I've consulted with Charlie Bennett, and John Holland, who really know about these matters, and they tell me that such tradeoffs are of very great importance. So I completely agree with you.

KAUFFMAN: Concerning the fundamental distinction you're making between direct adaptation and complex adaptive systems. Let me take what might be a form of the strong artificial life claim, that anything you can do with neurons and computer chips I can do with molecules. Consider *E. coli*, which has something called the lac operon in it. The function of the lac operon is that when lactose comes into *E. coli*, it binds to a receptor molecule. Normally the receptor molecule sits on the gene—called the operator—and stops transcription of the lactose gene. But when lactose is in the cell, the lactose binds to the repressor molecule and pulls it off the operator thereby allowing transcription so that *E. coli* now metabolizes lactose. The point here is when you're trying to draw the distinction between something called "direct adaptation" of *E. coli* becoming resistant to a toxin, which it might do for example by evolving a new protein, or evolving new regulatory circuitry to switch on in the presence of that toxin, which commonly happens, then if *E. coli* can do it, what's the distinction between *E. coli* doing direct adaptation to toxins, and the distinction you want.

GELL-MANN: If it happens genetically, it is not direct adaptation.

KAUFFMAN: So why isn't that a complex adaptive system?

GELL-MANN: It is. I said genetic adaptation *was*.

KAUFFMAN: I thought you said that direct adaptation is not complex.

GELL-MANN: I said Hinshelwood's kind of direct adaptation was not. If you are discussing genetic adaptation which I do not call direct, then you are dealing with complex adaptation.

KAUFFMAN: Oh. Well, what's excluded?

GELL-MANN: What is excluded is a case where you have simple cybernetic feedback without compression, without what, in biology, usually goes through the genes. In biology, compressed information is usually genetic.

COWAN: Hinshelwood did it by a chemical reaction in one generation.

KAUFFMAN: But suppose that that worked, suppose that you pulled out a modification in a protein molecule so that it is now a new protein. A Lamarckian-modified molecule would not have compressed information.

GELL-MANN: That's all right. You can imagine a biological process that involves compression of information and that *does* go on in the same generation, without genes, and that would still be a complex adaptive system. It would be a new, hitherto-undiscovered one, or maybe one that's been discovered already but is obscure.

KAUFFMAN: So your distinction is between whether or not the process that does the recognizing and reacting somehow compresses its description. . .

GELL-MANN: I distinguish between compressing a lot of experience into a small message and a look-up table. Most of the cybernetic devices that we discussed 50 years ago did not have the faculty of compression.

ANDERSON: Couldn't the degree of compression be a criterion, rather than a measure of complexity? In the case of scientific theories it very often is, if one takes the Bayesian attitude that a simple theory is intrinsically better than a complex theory.

GELL-MANN: Yes, I agree. I think you could use it as a quantitative measure, also. But I don't know exactly how. I don't know whether to take ratio of numbers of bits, or what. But there probably is some quantitative measure that we use.

PINES: A comment: As you remarked, the scientific enterprise is a good example of a complex adaptive system, and so are we all. And in a sense, I think it's interesting to view this workshop as a complex adaptive system in which, in a sense, the selection pressures have to do with which, if any, of those guiding principles have general applicability. Then I can imagine several ways of trying to test this. One might be that we arrive at a consensus that some of these work, some don't. But I'm not sure that's enough. I wonder if we shouldn't—as we look at each of these principles—ask the question, "Can I compare it to an experiment, or series of experiments? Can I compare it to a series of observations?" Or finally, "Can I carry out a set of computer simulations to test the principles involved?" I make this comment in the hope that as we go along in the meeting, we'll address these issues, and each of the principles, from this point of view: How does it work? Does

it really work when applied to practice? And finally, a further question: Do any of the principles possess any predictive power in dealing with a particular simple system?

GELL-MANN: I don't think it's such a simple matter as just saying, "These are theories, and they should be tested by observation." I think it's more subtle. We structured the workshop so that during the first few days, some notions of general ideas of complex adaptive systems, and also nonadaptive complexity, would be presented. Then there would be a few days of discussion by people who are experts in fields such as immunology, economics, adaptive computation, and the origin of life—various fields where the rubber hits the road. It's, so to speak, our phenotypic arena. Although we don't ourselves do experiments, nevertheless, these are people who are in contact with experiment and whose theories *are* intended to predict correctly the results of observations. Then we come back and reexamine the discussions that we've had at the beginning and see how we want to modify our original ideas in the light of criticism by the other general theorists, but especially in the light of what we've heard from the people who are doing the work in the specific fields. Now that's not quite the same as what you said.

These proposals are not exactly theories. These are suggestions for how to organize the work, and the test is not whether they're, so to speak, "true" or not (although I may have lied here and there); the test is whether these are useful in organizing our thinking about all the things we're going to hear about in the general session, and especially in the individual sessions on the individual subjects. That's what we'll come back and discuss during the last day or so: how we want to modify our general ways of talking, in the light of their utility not so much in the lab as in the discussion. Those are my views of the selection pressures on the general ideas.

PINES: Now my question: why do you think so much more attention has been paid to trying to arrive at quantitative measures of complexity, and so little attention, relatively speaking, to quantitative measures of adaption?

GELL-MANN: I think that we have not had a lot of attention to either, in the sense of complexity that I've tried to illustrate here. What's happened is we've had a lot of attention to quantitative measures of complexity as defined by mathematicians, for mathematicians, in mathematical contexts. And that's no wonder. But the kind of complexity we're talking about here is still a bit ill-defined, it's quite subjective, as we've discussed. And likewise adaptation, evolution, and so on are tricky to work with. These are less clear-cut mathematical problems, and there has been less attention to them. However, in certain fields, like mathematical population biology, we find attention paid to clear-cut discussion of the issues.

FELDMAN: One of the answers to that is that each of the disciplines comes at the definition of adaptation in a different way, and in biology you don't often measure adaptation. They measure, as you said, fitness differences—relative fitnesses. In the experimental computer science that I've seen at Santa Fe Institute

the criterion for adaptation is given to you before you begin, namely, you want to see the process which leads to the best version of this program, or the best instruction. You know what the criterion for adaptation is. Biologists, on the whole, don't do that.

GELL-MANN: But I think there isn't yet much of a general theory, apart from biology, apart from computer science, apart from other particular disciplines. I don't think there's been much work on pattern recognition compression, and especially variation and selection, as general phenomena rather than in connection with biology, thinking, computers, or whatever. What I'm proposing here is that it *might* be useful to think of them in that way.

JEN: I was confused by the distinction you were making in the context of speciation and evolution of individuality between physical and biological systems. It led me to think about what other sort of fundamental distinctions there are, what other contexts in which you think there are actually fundamental distinctions between physical and nonphysical systems (though I'm not even sure what a physical system is). Given the fact that you think that there are such distinctions, and given the fact that we are as a group, by and large, mostly people coming from the physical sciences looking at the nonphysical sciences, how should this affect our choice of what problems to work on, given the fact that we have so much enthusiasm now for looking at everything. You believe, for instance, that a process like evolution of human language is in fact tractable, or susceptible to these methods of analysis which is not at all *a priori* obvious, because of the fact that it is so different from what we know from physical science. What is the scope of things that we can look at, and what is the scope of things that we're really presumptuous to be thinking about at this point?

GELL-MANN: I spoke rapidly about differences, and of course I was eliding differences which in my writing I've tried to make more explicit.

There is no nonphysical system in the universe, according to the beliefs of everybody in this room. I don't think that anyone here would believe that there are fundamental "vital forces" which are outside of physics and chemistry, and govern life. I don't imagine there's anyone in this room who believes that there are fundamental "mental" processes that govern thinking and that are outside of biology and, therefore, outside of physics and chemistry. When I made the distinction between physical evolution and the kind we're talking about here I meant only that there are properties of complex adaptive systems—which I tried to describe—which are largely absent as far as we know in whole classes of physical evolution: evolution of galaxies, of stars, of planets, and so on. We have no evidence of compression, schemata, variation of schemata, selection; *maybe* they exhibited these phenomena, but if so we don't know about it.

I've used the issue of turbulence, for example, as a case. We know that turbulence in a complicated pipe with changing shape has little eddies in it, and the

little eddies spawn smaller eddies, and the smaller eddies spawn smaller eddies, and certain eddies find their way through the pipe successfully and live to reproduce with lots of little tiny eddies. Other eddies don't make it through the pipe. Now, are we talking there about selection and evolution in the biological sense? Well, probably *not*, because we don't have any evidence that these eddies are doing the work of perceiving regularities, compressing them, and then constructing variant compressed schemata that undergo selection. Rather the selection of eddies seems to be taking place on the surface, with look-up tables. Of course, we don't *know* that, for sure.

So you are right to be skeptical, if that is the attitude you're expressing. Doyne Farmer has phrased it very well in saying that he would like some day to understand how, from the equations or the rule-based mathematics for a system, one could tell whether it was making a compressed model of its environment and of its own behavior. In other words, could you tell from a mathematical description whether the system is going beyond the physics and chemistry that everything shares to complex adaptive behavior. We don't know how to do that, and it would be an interesting challenge.

LANGTON: You made the statement that the difference between learning and culture and evolution is that evolution doesn't make big leaps.

GELL-MANN: No, no. *Biological* evolution, I said, tends not to make big jumps but works with what it has and proceeds by modifications of what is there—organs that are there, for example—for new functions. You can see that in societies, also. The British, for example, are very good at this. They have the Privy Council, which used to supply advice at the highest level to the ruler. The ruler doesn't rule anymore, but simply reigns; they've still got the Privy Council, however. So what do they do with it? Well, they make it an advisory committee on science, for example. Human thinking may have the possibility, occasionally, of operating in a different mode where you make a big jump. However, some investigators think that isn't so. . .

ANDERSON: I don't think it's so. I don't believe that the brain makes such big leaps. It always uses something it already has.

LANGTON: To flip the coin, I also think it's also true that evolution some-times (rarely) does make very big leaps. For instance, the Cambrian explosion, the evolution of multicellularity, or the origin of eukaryotic cells. . .

GELL-MANN: Those look like very rapid processes viewed from the present, because they took place billions of years ago, but they took awhile.

LANGTON: I'm not sure I see such a big difference. . .

GELL-MANN: In my opinion there is a significant difference. I think that an engineer designing something—a hypothetical engineer, maybe not a real graduate of an engineering school—could make a more rapid jump than a gene. But I could be wrong. We should try to find some quantitative measure of jumping and see whether it's true.

FONTANA: The basic difference that I believe to have caught in the first two talks. . .your talk was centered around the genotype-phenotype distinction. This is actually the essential distinction that was lacking, it seemed to me, in the considerations that Phil was making.

GELL-MANN: *Until* today I always said it was essential, and today I partially took it back, in the sense that there are cases where the two are *physically* the same.

FONTANA: I'd like to convince you to draw this distinction more radically, actually. Because that is exactly the distinction that physics is lacking (I mean a traditional way of thinking in physics). This distinction actually has its counterpart in mathematics, and in particular in a theory of mathematics called recursive function theory. There the distinction goes under the heading "Object and Function." A function, for a mathematician—prior to recursion theory—or for a set theoretician, was essentially a phenotype. So all functions were phenotypes. There were infinite arrays of facts, all of which were mostly accidental. Now if you take such an infinite array of facts, of input-output pairs, and you can express it in terms of a rule by capturing a pattern in this infinite series of facts, then you have constructed a computable function. So recursion theory strikes me as being the most basic, and the first mathematical theory about compression at all, because it tells you what you can express with finite means. It tells you which phenotypes have a genotype, so to speak, in an abstract sense, in a very general sense.

In this respect I would say that it is not true that, for example, if you'd like to include Tom Ray's work in your definition of complex adaptive systems, I think you can still do it because I believe that in Tom's work there *is* a clear-cut distinction between a genotype and phenotype because the *program* is a genotype, but a program is also a function. It's a series of input-output pairs that can be captured by a pattern which is the program, but nevertheless the program is just a specification of actions so it has a phenotype, obviously.

GELL-MANN: But in a sense it's the same thing.

FONTANA: No, it's not the same thing. If I write down the function "x^2," that's not the same thing as the function that maps any natural number x^2.

GELL-MANN: What you're saying is that the distinction has been compressed very thin, but it's not absolutely zero.

FONTANA: I think the distinction becomes very important particularly when we live in a finite world where a function, or an object that expresses a function, never sees the entire domain over which it is defined, but only a tiny part. So you can have a function that is the identity only on a very specific domain but not in general. It depends on what other objects are there. This leads me to a second brief comment. . . .

ANDERSON: First this comment; you mentioned my talk. I think that this represents a misunderstanding of my assigned role in giving this talk and I do not like being used as a straw man in this way. Essentially you're saying it's the difference between two historical eras, which I would certainly agree with. I was describing all the historical approaches to complex adaptive systems and carefully avoiding any discussion of CAS's themselves

FONTANA: Let me add a brief remark. The interesting fact about recursion theory, if we would like to take this as a model of an abstract genotype-phenotype distinction, is that it is constructive. That's why I like it so much. It tells you how you build new things—in your terms, schemas—out of available ones in a totally nonrandom fashion, without mutation or recombination and so forth. Clearly, mutation and recombination are events that are most important for an adaptive system, and a complex adaptive system seems to me also to have a constructive part where you can't buy the actual numbers individually—you get them only as a package and you get the entire structure implied by them. So, if you have elements in complex adaptive systems that are objects and functions at the same time, that are genotypes and phenotypes, then by virtue of their being functions that can act on these phenotypes, you get constructive effects that are not taken into account by purely focusing on random mutation and recombination.

GELL-MANN: What you're saying, if I may just put it in very simple terms so I can follow you, is that while the distinction has practically disappeared in Tom Ray's work, and also in the RNA theory of the origin of life—because we're dealing with the same agents, really, that are the chemical agents, and are the bearers of the information—nevertheless, it's worth still making the distinction between these two *roles* and in that way we can preserve the genotype-phenotype distinction in a useful way. Because in all the cases where it does exist, it is a very important distinction, preserving it in its degenerate form is a very good idea.

FONTANA: Exactly. I would say it occurs already with molecules, where you would say there is no such distinction. I would still make a distinction between the structure of a protein and the set of chemical reactions it can undergo, and these are just two examples. You can think the one as being a representation of the others, but they're two different functions, one is extensional and one is intensional.

BROWN: Early on in your talk—you went over this very rapidly—I think I heard you say that complex systems can exist, and actually can evolve and

develop if the environment, essentially, is a constant time series and doesn't change. I wonder if that is really correct...

GELL-MANN: By time series with constant properties I mean that it does have certain reactions to what's done to it and reacts back on the system, but it does it in a constant way, so to speak. So we're not including secular variation like changing of the earth's atmosphere, and we're not including coevolution.

BROWN: I want to challenge you on that and ask whether these major extrinsic changes may not be very important in the development of these kinds of systems. If we get to the question that Chris Langton asked about the leaps, for example, in biology and in ideas—I wonder if many times those leaps aren't triggered and in some fundamental way caused by a change in the environment. One way to get off the hill, to use the metaphor, is to change the landscape so the peak of the hill no longer is the top of the mountain. That involves potentially a major environmental change.

GELL-MANN: I think you're right in many cases, and in fact I have a whole lecture—which I will not inflict on you—on getting creative ideas, which depends precisely on that point: that one can search for artificial stimuli to get one out of traps, out of idea traps, to other regions where the better ideas may lie, through a changed environment. But what I wanted to say was something slightly different. I wanted to say that you can still maintain the definition of a complex adaptive system and study some of its properties by idealizing, by ignoring the change in the time series represented by the environment, and by ignoring coevolution. You will not get many of the most important properties of real ones, but you will still have something that would fall into the rubric of complex adaptive systems and would adapt to its environment. And that's worth studying, I believe, as a simple case that is still part of the general subject. But you're right; in most practical cases the changes are very important.

ANDERSON: The fitness-nonfitness distinction is already available at the rather primitive level of spin glasses. You can deal with them as neural nets, and deal with them—as John Hopfield did—with a fitness landscape. And in fact all of the mathematics works perfectly well with asymmetric coefficients.

GELL-MANN: He was able to translate what was done into the language of some kind of physics and that was very nice. But it was also a step backwards because the people working on psychological models were already using the correct, nonsymmetric coefficients.

ANDERSON: But one can now deal with this mathematically and formally without assuming a fitness (Lyapunov) function.

BUSS: I have a comment on these definitional issues and then a question for the speakers. The suggestion was made by Murray that one of the strengths of science is that there's some mild selection pressure on being correct. I would contend that another strength of science is that the definition of what you are working on is not a problem internal to the science itself. I think that if you look at fields where definition of your activity is in fact internal to the discipline—for example, philosophy—there is not quite the same illusion of progress as you have in other disciplines.

GELL-MANN: We've seen a lot of that in recent years in the penetration of what are fundamentally political disputes into the philosophical bases of various subjects. In archeology, for example, it's been a terrible problem, now receding. You're right. That kind of debate in a scientific field can be devastating to scientific progress.

That was why I mentioned three different ways in which a society can "adapt," and how futile disputes about which of those is *really* adaptation, which of those is *really* evolution, can cause the subject to grind to a halt. Is that the sort of thing you mean?

BUSS: I think it speaks to how much discussion there need be about what complexity is.

Now my question. I was struck by the difference between George's introduction, which focused on levels of complexity, the hierarchy of complexity—how do you get chemistry out of physics, how do you get biology out of chemistry, how do you get social systems out of biology—from the two largely methodological talks that followed. It wasn't obvious to me that the material covered in those two presentations in any sense went to the heart of the emergence of new organizational levels that George was focusing us on.

GELL-MANN: I think that's an excellent question. I don't have much that's intelligent to say about it, but that's not because I don't think it's a very important subject.

There are, presumably, thresholds of complexity allowing certain kinds of systems to function. And there is presumably some threshold of complexity for a complex adaptive system which is why we're allowed to call it a complex adaptive system. Or, if not a threshold, at least a mathematical relationship, so that if we define the degree of compression, the degree of variation or selection, and so on, we can relate the complexity of the system to those degrees. In other words, either a sharp threshold or a gentle one. Likewise, there's presumably some sort of threshold for self-awareness, which would be very exciting to understand. And again, that may be sharp or it may be gentle. At a place like this, one is presumably allowed to mention consciousness—although there are a number of campuses where saying the "C word" would lead to ostracism.

ANDERSON: I would agree with what Murray says very much; it's not that we don't agree, but we know less about that. Really, you build hierarchies—of course you build hierarchies.

GELL-MANN: Now on the other subject of what is fundamental, and how fundamental science gives rise to sciences on other levels; that I have written about a lot (but not here), and it's an important subject too—how particle physics and cosmology give rise to chemistry, giving rise to biology, and so on and so on, with the addition of new information at each stage.

PINES: I want to go back to one of Murray's points, which has been discussed a bit, that the difference between, as he suggested, a biological system— which goes along pretty much as it has been, within a certain set of regularities, small changes, and so forth—as compared to a thought, human thought, where one can suddenly get great leaps, or so one thinks.

GELL-MANN: I said I thought that there was a quantitative difference between the two, but people have disputed it.

PINES: I think there is, and I would suggest that it's the difference between learning and innovation. [In] learning, in some sense, you're adapting to be able to recognize an existing set of patterns, which are sort of spelled out for you by your environment, whereas with innovation you may suddenly view things differently. The complex adaptive system may be capable of saying, "Hey, there may be a quite different set of patterns out there other than those which the environment is providing to me." And I wonder if, in some poetic sense at least, this is also connected with the notion of emergent behavior. Namely, that you're going along with a system which seems to be, you know, bouncing up and down in some regular, or irregular, kind of way, and suddenly you find—subject to a given set of stimuli— that it quite changes its behavior pattern.

GELL-MANN: Let me argue on the other side, for a moment, since you've taken my side. Chris Langton isn't here, today; I'll argue on his side.
        I've thought some, and read a lot, and written a little bit, about innovation in human thinking—in particular in connection with science, but also art and other subjects—and talked with a lot of practitioners of all sorts of subjects (in the arts, engineering, and the sciences). The methods seem to be very similar, and one thing that stands out is that usually the leap (if there is one), begins in a negative way, by getting rid of some unnecessary prohibition that was adopted a long time before along with a useful idea. Getting rid of that prohibition allows more freedom than was thought to exist, and it can lead quite rapidly to progress in solving or formulating a problem. Perhaps the wrong *Verbot* holds back to normal form of progress by small tips, so that when the brain is removed, a longer tip can be taken.

ANDERSON: Let's say that our answers to Leo kind of implied that evolution can make big jumps, because we were saying that there are quite different levels on the hierarchy and that they had to happen sometime. We have to have consciousness for the first time, in some sense.

GELL-MANN: Well maybe. That depends on how gentle the threshold is. We don't know how sharp these thresholds are.

PINES: Jim Brown makes the point that it may have a lot to do, in evolution at least, with a rapid change in the environment. And of course you can ask the same question about human innovation.

LLOYD: Saltation (the theory of evolution by leaps and bounds), rather than gradualism, is a viable theory of evolution, at any rate.

GELL-MANN: There may be some planets on which it usually goes that way. On this planet it doesn't seem to be true.

FELDMAN: This issue was addressed extensively by Saul Wright in the '20s and '30s, when he developed what he called the "shifting balance" theory. And it was designed to handle this sort of thing. You can have an evolutionary system—a biological evolutionary system—going along, changing—in the Darwinian sense, gradually—but with a number of different attracting points. Environmental change has the property of cutting down the population's size in a nonrandom way, so that each segment is now put into the domain of attraction of a different attractor. That's the shifting balance theory. But the long run of that is that you get an array of ultimate phenotypes, caused by these genotypes, and it might have been a stimulus from an environmental change, but it wasn't a large biological change at the level of the genotype that was necessary for this.

GELL-MANN: I think that much discussions of this topic from now on should refer to Tom Ray's program, and also to other programs that people are developing, in some of which there is a physical distinction between phenotype and genotype. For example at UCLA, they're working on such computer models of evolution, trying to test by computer modeling the idea of Hamilton and others, the idea that the value of sex—the role of the male—may have something to do with resistance to short-generation parasites. But to take Tom Ray's only, he has some very nice examples of cases where the situation is stable for a very very long time—apparently—but changes may be taking place that don't matter much for survival, and then all of a sudden, after a long period, huge changes in populations take place, as in the model of punctuated equilibrium.

FELDMAN: The key issue is, is there a biological organism that does that?

GELL-MANN: Yes that is the key issue. We have other cases where organisms are very similar to what they used to be billions of years ago, like extremophiles in

hot deep acidic waters, with sulfur. Now we don't know whether those are *genetically* like the very ancient extremophiles, or whether they've undergone a substantial amount of genetic change while remaining phenotypically very similar. It could be either way, and either way would be theoretically satisfactory. We just don't know. But in any case it's a fact that there is a case where the environment has changed very little, and where the phenotypic response to that environment has stayed more or less constant. It's like one of John Holland's cases where the problem is so simple that it's *solved* by the computer. I mean, playing tic-tac-toe, for example. It's all very well to design a checkers automaton, or a chess automaton, but a tic-tac-toe automaton converges quite rapidly to something that plays tic-tac-toe perfectly. It may change genotypically after that, but it's not going to change phenotypically because it already plays tic-tac-toe perfectly.

BUSS: I'd like to disassociate the comments that I've made with respect to hierarchy of complexity from this issue of whether biological evolution can make leaps. There are certain classes of leaps that are simply undeniable, and likely do involve something more than our conventional apparatus. For example, when we go from a prokaryotic cell to a eukaryotic cell, when we're combining two self-replicating entities to make a third class of self-replicating entities—that is an organizational shift that has happened biologically. It is clearly on a different grade than whether in fact you make a long bristle fast. There's a long history of the relative magnitude of changes within a given organizational grade, and the rate of their appearance, and there are a wide number of scenarios that can predict the rate of your choice. But I would like to say that that's a different set of problems than the set of problems of how you get to new organizational classes in evolution, and those are necessarily rapid.

KAUFFMAN: I'm struck, Phil, by the fact that you're drawing a distinction between throwing away irrelevant degrees of freedom, a system losing entropy and contracting down to a region of its phase volume—and Murray's statement about throwing away clutter, and compression. If you have a reversible dynamical system, it's reversible, and you can't in that sense throw away degrees of freedom. It seems to me that there's a connection between the point that you've made and Murray's statements about compression. Then there's the question of what you mean by the relevant degrees of freedom in keeping them. The fundamental question is: relevant for whom? Attempting to describe that is inevitably going to get us into a notion of agency. For whom is it good and for what purposes?

For example, if you've got bugs—Darwinian evolution with self-reproducing things—then I know how to answer the question "for whom?" Just, is it toxin, or food, for this bug? But more generally, one needs some sort of notion of relevance for whom, for what: what is becoming? And I think that it's fundamental, but I don't think it's well-posed.

GELL-MANN: I don't think that deriving the selection forces from a potential—or even treating them necessarily as forces rather than a statistical distribution of forces—is always possible. In specific cases it may be, but as a general rule I think it's not, and therefore I wouldn't use a potential function as a fitness. I would just study the selection pressures as such.

KAUFFMAN: You're restating that it's a dynamical system.

GELL-MANN: Yes. The second thing I would say is that in connection with entropy, that it's by now well known, as a result of the work of Landauer, and Bennett, and Zurek, and Lloyd (and even me, to some extent), if you're going to consider entropy from outside the closed system, that's one thing. Then it just increases in the usual way. But if you're going to consider it from inside, with an observer, then every time that observer learns something, on the average over which alternative occurs and is learned, the effective entropy of the system *decreases*, and keeps on decreasing, as more and more stuff is learned, provided you define entropy in the old way. And it's useful, therefore, to have a newer definition of entropy in which you add in the algorithmic complexity of the record of what's been learned. In that case, on the average, it continues to increase just as before. And that helps to clear up one or two points of confusion to which you alluded.

In order to keep the second law of thermodynamics going, from the point of view of an observer inside the system, you have to modify the definition of entropy in this way.

ANDERSON: These are points you can argue indefinitely, and my answer is different from Murray's and much more practical: namely, there are no closed systems in the universe, and therefore what would happen in a closed system is irrelevant. We're always radiating to somewhere and you're always adding in food, and this process of turning one into the other is what's happening. The second different answer I would make is that the actual entropy in the information that you're using is so negligible compared to what you're using to function—what you must use to function, to function at some reasonable rate—that you don't need to keep count of that.

GELL-MANN: It is certainly negligible, compared with the usual entropy that we talk about in chemical processes. However, it's not negligible if you're concentrating on that issue.

ANDERSON: Yes. One should not concentrate on that issue, is my answer. But we have not really answered Stu's question which is a fascinating and provocative one—is there an analogy between Liouville's theorem in phase space and something like it in information space?

WALDROP: If I understand what John Holland has been telling us, the real distinction between learning as in the mind, and evolution as in a biological

system, is the difference between implicit and explicit models. An implicit model is something like a rule, "If this is the situation, then do that." It's encoded a model of the world into a set of rules that are useful for the organism to do. The prototypic example is a bacterium swimming upstream in a glucose gradient. It is in effect executing a rule, "If the gradient is such-and-such, swim upstream." The bacterium has no brain, so it can't really know much about its world, but it does function as if it did. And this is a useful thing to do.

The explicit model is something much more like consciousness, where we do have an explicit model of, say, the physics of building a building so that we can reason about it, model it, and come to conclusions about it.

GELL-MANN: You're just talking again about the distinction between self-awareness and the lack of self-awareness, right?. At this point, we don't know all that much about what self-awareness is. It's true that human beings are supposed to have an unusual degree of self-awareness, and therefore human learning would be characterized, in many cases, by the properties of self-awareness. And we don't believe that's the case with biological evolution. Therefore that will be a difference. But I don't understand how that has to do with the question of whether there are significant jumps, or not. The fact that you mentioned, which is certainly true—that in human learning it's often related to self-awareness—doesn't change the dispute. People were asking whether in biological evolution there may not *also* be big leaps, and Leo Buss clarified very nicely the fact that there had been *certain* big leaps, of an organizational character. But then one can still ask: "Omitting those, does biological evolution proceed by big leaps, or by little ones?"

WALDROP: I'm saying that that might not be the important difference between learning and biological evolution, that both of them may in fact go by big leaps and small steps.

EPSTEIN: I've been having a crazy thought for a while. I don't know whether I'm the only one or not. But I've been thinking we've been talking about the scientific enterprise as an example of a complex adaptive system, with linguistic evolution as another example. And I've been thinking that maybe artistic evolution, and musical evolution, pose interesting problems for the construction of a good definition of a complex adaptive system. And what I'm thinking of is this: that in a certain strain of musical development, particularly western European classical music history since the sixteenth century forward, there's a very distinct line of development based on the emergence of the tonal system, and the gradual relaxation of that system with the Romantics, with Wagner, with extreme chromaticism, and finally the atonal movement. There seems to be directionality, the competition of schemata (in the form of different composers, and the search for different organizing principles); definite selection of some sort going on, in the dominant schools of musical thought, but no noncircular definition of fitness that I can come up with. And I'm thinking, "Here's a nice example of a complex adaptive system, that seems

to evolve, where fitness is a strange notion." And I'm wondering what we do with things like that, and whether the problem of sort of the endogenous...that fitness itself is an emergent thing that the systems have to somehow come to discover. And I wonder whether this isn't an interesting model...

GELL-MANN: The selection pressures there have to do in part with peer evaluation, in part with audience evaluation, in part with historical evaluation in part with subjective evaluation by people of their own work, and so on. That this would all be summarized in a potential—which is what fitness ideas would say— seems somewhat unlikely. I don't believe that we should look everywhere for fitness. It works here and there, and I don't expect it to work everywhere. But the selection pressures certainly exist everywhere.

Something I find interesting about the kind of development you describe is the following: Archeologists these days are terribly reluctant to engage in what they imagine are "value judgments" about the past by talking about periods in terms of florescience and decline in the arts.

Nevertheless, it's my belief that one can identify certain properties of many periods and many kinds of art—phenomena that occur over and over again. Namely, people formulate, in a given art at a given time, certain requirements that hadn't been very important before. Then there's a period of challenge when people are trying to meet those requirements in all sorts of interesting ways. That's an archaic period. Then the requirements are not. The artists are able to do whatever it is, and the art *flourishes*. Then comes a period when they begin to go off into variant complications. And these are often called "archaic art," "high art," and "rococo art." I think it could be a scientific fact that there exists this kind of sequence.

EPSTEIN: I think in the case of music history that this happens; that in the case of Bach, for example, the art of the fugue is basically an attempt to construct large musical structures based on very rigorous limitations on thematic material. There's this tiny snippet, this little string of material, you're going to be permitted to use, and you construct this enormous thing. You run the string back- wards at great, long expanse; you dilate it; you shrink it; you turn it upside down; you run it backwards. But the whole thing is very self-similar in that, anywhere you look, it's some variant of this thing. And then that impulse is in fact relaxed. And then it returns in twentieth century music, with serial composition that again says "Thou shalt use only this tone row, this string of material." So I agree.

BROWN: Isn't that just frequency-dependent selection played out over a long time course? Once you have a dominant thing, the rare allele gets favored and then, when that becomes dominant, the other gets favored...

EPSTEIN: Yes, except that it's part and parcel of a larger process, which is the gradual destruction of tonal music over four centuries, that has definite di- rection to it. I mean, you can document how this took place.

ANDERSON: I've already mentioned, in discussing George's comments, that there are many examples of evolution where the institution has its own fitness function that is irrelevant to the ostensible value and function. Modern music is clearly a selfish meme, because the populace liked Bach, the populace does not like John Cage or Schoenberg or Milton Babbitt. The populace invented a totally new atonal music, a very complicated music (often using the ideas of academic music) and abandoned some of the old rules, and this music is rapidly taking over from this other music—formal music—which is very much controlled by a selfish meme that has only institutional value.