

SCIENCE FICTION ADVENTURES MAGAZINE

CONTENTS • MAY 1954 • VOL. 2, NO. 3

NOVEL

RULE GOLDEN, by Damon Knight 6

NOVELETTE

RAIN CHECK, by Judith Merril 126

SHORT STORIES

PINNACLE, by Lee J. Fox 75

COLLISION ORBIT, by K. MacLean 78

AMOEBA-HUNT, by R. H. Remington 104

ARTICLES

FRANKENSTEIN, OR FRIDAY?, by L. Jerome Stanton 88

THE 183rd CONGRESS, by Roy H. Millenson 153

DEPARTMENTS

EDITORIAL 5

THE DISSECTING TABLE 121

FANMAG, by Charles Lee Riddle 146

THE CHART ROOM 156

CARTOON, by Cato 125

COVER BY CLARENCE DOORE

ILLUSTRATIONS BY BERWIN • EBEL • KRENKEL • MENDENHALL • SUSSMAN

SCIENCE FICTION ADVENTURES is published bi-monthly by Future Publications, Inc., 80 Fifth Avenue, New York 11, N. Y. Copyright 1954 by Future Publications, Inc. All rights, including translation in foreign countries reserved. Entered as second-class matter at the Post Office at New York, N. Y. Subscription \$3.50 for 12 issues in the United States and Possessions; \$4.50 in Canada; \$5.00 elsewhere. All stories in this issue are new and have not previously been published. No actual persons are designated in works of fiction, either by name or character; and any similarity is coincidental. Not responsible for unsolicited manuscripts or art work. All submissions must include return postage to expedite the return of manuscripts not acceptable.

JOHN RAYMOND • Publisher

DAVID GELLER • Advertising Mgr.

H. L. MITCHELL • West Coast Office

M. BERWIN • Art Director

HARRY HARRISON • Editor

M. MACHLIN • Associate Editor

DAMON KNIGHT • Book Editor

JOHN VINCENT • Associate Editor

FRANKENSTEIN, or FRIDAY?

BY L. JEROME STANTON

They've arrived, the sentient and calculating mechanical brains. They aren't developed yet, the stage they are in now bears about the same relationship to the final product as an amoeba does to man. Whether they mean the future or the extermination of the human race has yet to be determined.

Science-fiction writers are both inventive and hardy, as dozens of anthologies (and numerous bill collectors) ably testify. But, in spite of their undeniable inventiveness, there are nevertheless certain science-fiction ideas that get used over and over again like the roller towel in a cheap boardinghouse. Now, in recent times, particularly in the last fifty years or so, a certain tendency in human affairs has become apparent that certainly merits a long, hard look from any science-fiction reader worthy of his Bug-Eyed Monsters. Briefly, the tendency is for those ideas that s-f writers use so frequently to become the accomplished fact of laboratory

and factory, and show up on the breakfast table, or in milady's bath. This notable trend is quite remarkable, and even a little disquieting at times. The disquiet arises because Man, the natural-born Sophomore, has a distressing ability to make things like atom bombs and laboratory-grown plagues come true before anyone other than a few science-fiction writers has given even an idle thought to the consequences.

Examples of this ability to make the science-fiction cliche become hard reality on a dollar-down, over-the-counter basis are already cluttering up the landscape in almost embarrassing profusion. Already, some of them



ILLUSTRATED BY BERWIN

show an alarming capacity to do things to us and our culture that the inventor in his wildest dreams never imagined. Science-fiction readers are well aware that long before motor transportation, radio, television, motion pictures and wonder drugs were actually developed, speculative writers like Wells, Verne, Gernsback, and many others had "invented" these useful gadgets over and over again for their readers. True, in many cases an actual development in some scientist's lab had pointed the way, but it was often a science-fiction writer who snatched up the bare possibility of, say, air transportation, and ran off with it in all directions, including the sociological. The development of atomic energy is a classic example of an idea belabored again, and again, and again by the yeasty minds and nimble typewriters of science-fictionists. It hardly needs pointing out here that sure enough, the Sharp Apples in science laboratories have again come through and made physical reality of this idea, which for years was the sole property and often the main, support of cent-a-word-and-dodge-the-landlord writers.

Fair enough. It looks as though some subtle but powerful process is at work here, busily

making fiction idea into fact, usually with little reckoning of the overall result. The s-f cliches are coming true, with no guaranteee as to whether they will be a Frankenstein's Monster to destroy Mankind, or a faithful Man Friday, to serve his will and smooth his path to the stars and the future. That being the case, let's take a peek at some of the real old warhorses in the s-f writer's bag of tricks, to see which one may be next on the list for production. Quite a few of the typewriter-weary standbys are still not being sold in Macy's Toy Department, or used to settle boundary disputes between nations.

There's matter transmission and its first cousin the matter duplicator, also antigravity, and good old time travel, the paradox-lover's delight. And there's telepathy, and the other more or less mystical members of the fecund "tele" family, like portation, kinesis, and so on. All of these are in the same state that air travel, radio, and nuclear energy were in not so very long ago, but none of them seems at the moment fully ripened and ready to drop.

Ah! But there is one venerable old chestnut among the s-f cliche-ideas that seems temptingly ripe, and dozens of laboratories are shaking the tree like

mad. This one is the s-f writer's vision of the man-created thinking machine, which is so often pictured as taking over the increasingly complex job of solving man's technical, social, and philosophical problems for him. Hardly another idea has been dealt with on so many levels in science fiction, or with as great a variety of conclusions. For that matter, hardly any other idea poses more provocative questions for us to speculate over, even before the idea becomes solid reality. Oddly enough, the development that has put realization of the all-embracing machine-thinker at least on the horizon of possibility, if not within our grasp, has itself received practically no public notice. This may be true because the device itself is still in a crude, undeveloped state, although quite capable of demonstrating the disturbing ability that makes it so important. As of this writing, it has not been applied to any practical use, in the way that the big computers have been. This may account for the lack of publicity. And yet, without the fundamental principle embodied in the design of this crude gadget, the science-fiction idea of the vast God-machine which absorbs all knowledge and spews out infallible answers to all questions would

still be completely out of reach. Superficially, the new device is not very impressive, but it can do something that the monster computers like ENIAC and TYPHOON cannot master, for all their massed vaults of tubes and relays and transistors. This new, disturbing egg now hatching is . . .

The Homeostat.

And the new factor in its performance that sets it apart from ENIAC, MANIAC, OMIBAC, and the other computers is this: The Homeostat can not only perform the operations for which it is first set up, but *it can learn from its own mistakes, and alter its own internal organization so as to solve problems that it failed to solve on earlier attempts.*

In the prim language of the laboratory, the peculiar ability described may not appear impressive, let alone earth-shaking. Yet the difference between this ability and the action of computers in solving mathematical problems of whatever complexity is both subtle and important. It is a real difference in *kind* of ability, rather than a difference in degree of complexity. It is through the use of this principle, in conjunction with other machine abilities already developed, that the way is opened to possible development of the real thinking machine, and the

many uncomfortable questions it suggests.

To understand the Homeostat and its importance, it will help to consider just what the difference between a "brain" and a computer is, and how "thinking" differs from calculation as performed by computers. Oddly enough, defining these various terms is about the most troublesome part of the whole story. This becomes doubly difficult if we attempt to satisfy at one stroke philosophers, biologists, psychologists, physicists, and just plain people. A good deal of the trouble arises because these worthy but divergent groups cannot agree among themselves on just what abilities a thing must have to be called a real brain, or on what "thought" is, for that matter. On one score they are all pretty well agreed: that neither an analog nor a digital computer is a true brain, although either can perform some of the functions of a brain. Most will also agree that whatever the thought process is, it involves something more than the mathematical operations performed by computers, regardless of their speed or size.

A good many people appear to regard thought as something apart from action. Some might describe it as "the sort of thing I'm doing in considering this

question," which isn't very helpful. To the biologist, however, the human brain is as much an acting machine as any other part of the body. It takes in something (information) through certain specialized channels we call the senses. Then it performs certain actions on the information taken in, often with reference to other related information already stored in memory. After that, the results of these actions may be sent as commands to the effectors (muscles, etc.) of the body, or as information to be stored in the memory banks, or both. Like other bodily organs, the brain plays an important part in the effort to survive. Indeed, in human beings, the brain seems to have a more important role in survival than in any other animal. It is with survival, the most basic goal of animal existence, then we can best start in defining "brain" and "thought," as a step toward understanding the Homeostat and its importance.

The tendency to strive toward a goal is a characteristic of living organisms, in particular the tendency to strive toward the goals of individual and species survival, which at times may conflict. This goal-seeking urge in living matter is insured by what may well be the most fundamental truism of the

Cosmos; things that can survive do, and things that can't, don't. If we define 'survival' as meaning that a thing continues to maintain its identity, the truism applies just as readily to a photon of radiant energy as to a molecule, or a rabbit, or a thinking machine.

The intrinsic goal-seeking urge, then, appears to be one of the characteristics a true brain must have. This, however, is something easy to build into an electro-mechanical contrivance, as various ingenious people like Claude Shannon of Bell Telephone Laboratories have ably demonstrated. In fact, we can go on and state several other abilities that a true brain has, which are also properties of computers and similar devices. It must be able to receive and store information. It must also be capable of combining related items of information, and derive from them other items of information, which may either be stores or provided as outputs or both. Oddly enough, this statement implies the ability to predict the probable course of future events, as a result of the derivation of new information from combinations of the old. The capacity to predict the probable future from study of the past has sometimes been put forward as the unique

characteristic of a true brain. Yet obviously this can't be true, or every radar gun director becomes a "true" brain!

What, then, really is the unique ability a true brain must have in addition to the qualities mentioned above? What test do we apply to separate the thinking goats from the finger-counting or angle-measuring sheep like ENIAC and the rest? This is the question Dr. W. R. Ashby had to answer for himself before he could devise the Homeostat, which appears to be the first man-made device able to cross the invisible dividing line.

As Director of Research at Barnwood House, one of the foremost psychiatric hospitals in England, Dr. Ashby had ample opportunity to consider the human brain in various state of function and malfunction. As a consequence, his analysis of the qualities a device must have before it can truly be regarded as a brain is both interesting and significant. Dr. Ashby has set down his ideas in an absorbing book appropriately titled "Design For A Brain," and it will be helpful to summarize his conclusions here before we consider how the Homeostat illustrates them in physical fact.

Apparently, the most unique property of a true brain is an ability to make changes in its

own internal circuits that tend to further its own survival. Putting it another way, it is the capacity to so control its own internal organization as to keep its ability to survive in its existing environment at a maximum. This means also the ability to make changes in its internal organization to compensate for changes in the environment that tend to lower its power to survive. An illustration of the difference between this ability to compensate or adapt to environmental changes and the goal-seeking ability will clarify the point. We can easily construct a small machine that will seek a goal in a way quite similar to a rat seeking food through the paths of a maze. Suppose we make a little motor-driven machine that can propel itself across a floor, powered by a small motor and self-contained battery. We can equip it with a pair of photocells and design it so that it turns toward a source of light and goes to it. This is quite similar to the rat going in the direction its sense of sight or smell tells it the food lies. Both are seeking a goal, and we can fairly consider that the urge to seek the goal is built-in. But in seeking their goals, the machine and the rat will show important fundamental differences in behavior.

Let's put our little machine down on the floor at one side of a room, and set up the lamp bulb goal directly opposite, with no obstacles between. When the lamp is switched on, the machine turns immediately and makes straight for it, by a direct path. Similarly, a hungry rat will go directly toward food when it is offered. Now suppose that directly above a part of the path the machine takes across the floor to the goal we suspend an inverted wire basket. This we arrange so that when the machine moves under it, a catch is released that drops the basket over the machine, trapping it, and thus preventing it from reaching its goal. Then we arrange a similar trap over the path the rat takes toward its goal.

Immediately a fundamental difference in behavior shows up. The machine, after being thus trapped any number of times, will continue to make the same mistake of following the most direct path toward the goal. It does this because all of its functions are built-in to start with, and it *cannot make any changes within itself to avoid the path which it has found by experiment does not actually lead to the goal.* In other words, it cannot learn by the experience of failure to seek out another path which may have a chance of success.

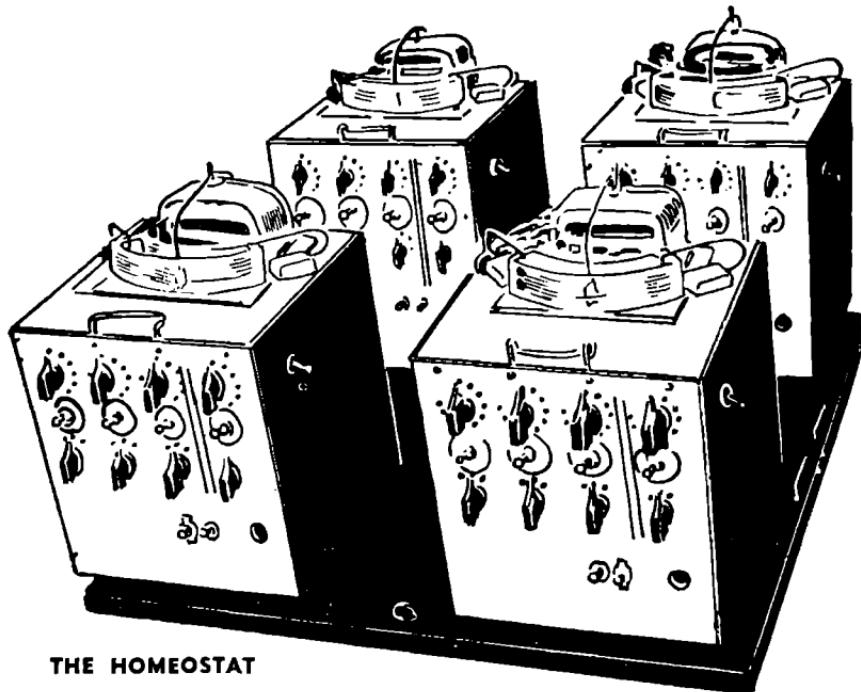
The rat on the other hand can, and does, as every warehouse owner well knows. After our experimental rat has been trapped under the descending wire basket once, it is likely to go around the area directly underneath the basket when the experiment is tried again, even if a considerable detour is necessary. And after two or three frustrating experiences, even the stupidest rat learns to avoid the basket trap. The rat has been able to make changes within its own patterns of conduct so that the goal which was not obtained on the first try is reached on a later try. In other words the rat is *adaptive*, and our simple-minded little machine is not. Not only that, but Mr. Rat (he's earned a little dignity by this demonstration of superiority) shows us another unique characteristic of a real brain. His efforts to find a successful path to the goal may be more or less random, but he shows a high degree of ability to select successful methods from unsuccessful ones, and use them over again.

It can be shown that most if not all of the activities of animals with brains, including man, are the result of the two fundamental factors stated above. We have a goal, or goals, and we show a capacity to change

our ways of trying to reach the goal when the way we are using does not work. This adaptive ability seems to be the really unique ability a thing must have before we can really call it a brain. And that, ladies and gentlemen, is just what Dr. Ashby has built into the Homeostat. True, the device illustrates adaptability in a very crude, simple sense, as compared with human adaptability, but it *does* employ the same principle. This time, the difference is one of degree of complexity, not kind of function.

And thereby hangs our somewhat frightening tale, because this is the one function of a true brain that is not already in our grasp. The possible consequences of building a real super-brain are something to make even the reckless thoughtful, but before we consider them, let's get an idea of what the Homeostat actually is, and what it does, as described in Dr. Ashby's book.

Basically, the Homeostat consists of several similar units, which are interconnected in a particular way. In the original model, there are four units. A single unit consists of a magnet suspended on a pivot, inside a coil of wire that actually is made up of four coils, one for each unit. It is the behavior of the



THE HOMEOSTAT

magnets of the four units that show the adaptive conduct of the gadget. A part of each unit is a trough in the shape of an arc of a circle, which contains a liquid that will conduct electricity, but has some internal resistance. Each trough is connected to a source of electric current so that one end is at a potential of -2 volts, and the other end is at -15 volts. A light but stiff metal arm attached to the pivot supports the magnet in the field of the coil, and also carries at its other end a metal vane that makes contact with

the electrolyte liquid in the trough. The whole assembly mounted on the pivot is free to turn through an arc of about 90 degrees. Thus the electric potential picked up by the contact will vary between -2 and -15 volts with the position of the magnet with respect to the coil. This voltage is carried to the grid of a vacuum tube in each unit, and acts to control the amount of current flowing in the vacuum tube plate circuit.

The output current through the vacuum tube in each unit is carried to the other units in

series, and is made to flow through the coil surrounding the magnet in each unit. A switch and potentiometer in each unit makes it possible for the experimenter to reverse the direction the current flows through any coil in each unit at will, and also to decide what portion of the total current output from each unit flows through the coil. The remaining part of each unit is a device called a uniselector, which contains 25 resistances of *random* value, and a rotary stepping switch. The uniselector is made so that when and *only* when the output current from any unit reaches a value sufficient to cause a relay to close. In addition, switches in each unit are provided so that the experimenter can take control of the potentiometers and polarity switches and set in any values he wishes, or can surrender control to the uniselector, in which case, the Homeostat will select its own values for current and polarity.

So much for the construction of the machine. In operation, each unit is adjusted so that when the magnet is in its central position, midway between the extremes of its possible swing, there is zero current output from the unit. Then if the magnet is deliberately pulled away from the central position, there

will be a current output that increases with the displacement, and since this current flows through the coil, it tends to return the magnet to the center. Thus the unit seeks the center as a goal. Let's say that it is happy at the center, and is always trying to be happy.

After each unit has been individually set up to be happy in the center, we can switch on the whole machine, make any arbitrary adjustments of any controls other than the "happiness" adjustment, and watch what happens. Some settings of the controls leave the whole machine stable as before; that is, there will be changes in the amount of current flowing through each part of the total coil of each unit, but the Homeostat will automatically reach a balance of currents such that the magnets resume their central positions after a period of time. However, some manual settings of the switches and potentiometers can be found that cause unstable conduct; that is, they cause the magnets to go to positions away from the center, perhaps to the extreme end of the travel, up against the stops. In the first case, it is as if an animal had adjusted to conditions in its environment in a way to make it self happy, and in the second, as if it had become neurotic, and

was driving itself away from its optimal condition.

So far, we have not really permitted the Homeostat to demonstrate its really unique ability, the capacity to make changes within itself to adjust to the outside environmental conditions we impose on it. That happens when we switch two of the units over to the "Uniselector" instead of the "Manual" position, so that current in either large enough to close its "Change" relay will cause the stopping switch to move to a new position. Remember that when this switch moves, it switches in a different resistance of random* value, which causes a proportional random change in the amount of current flowing through the coils.

Now, we have in effect a situation in which the two units kept on manual control serve as an "environment" for the two units put on their own. We can call them the "animal" if we wish, as it is their behavior in trying to get "happy," or centered, that illustrates adaptation, just as a cat adapts itself to the changing heat from a fireplace by moving closer or farther away. If we make some arbitrary change in the settings of the polarity

switches or current-controlling potentiometers in the two "environment" units, the other two units may not be able to adjust so that the currents in the parts of each main coil balance, and thus move the magnet to the happiness position. When this happens, and the arm swings away from the center, it reaches a point where the current causes the relay in that unit to close, switching in a different resistance. This one *may* be of such a value that the "animal" can now get "happy," or it may not. If it is not, the relay gets pulled closed again, and the Homeostat tries another value, or as we'd say of the rat, it tries another path through the maze to the food. Even with only two units serving as our "animal," 25 times 25, or 625 different combinations of resistor values are available for trial. And when all four are serving as animal, a total of 390, 625 possible combinations are available.

With the two-and-two set-up, when the experimenter changes the polarity setting of a switch in the "environment" units, the other two units immediately start hunting for settings of the uniselectors that will permit the "animal" pair to be happy again. The same thing happens when the "environment" potentiometer settings are changed. And this

*The resistor values for the stepping switches were actually chosen from a published table of random numbers.

happens every time a change is made in the environment.

But there is something more, an ability which will astonish many who find nothing too startling in the Homeostat's performance up to here. When the whole device is set on the uniselectors, so that it is all acting as an experimental "animal" with a brain of its own, it really goes to town. It can not only adjust itself to be happy no matter how we set the polarity and current controls, but, *we can make changes in the mechanism itself, and it will still adapt, and get happy again.*

After Dr. Ashby and the other experimenters had put the Homeostat through a stiff set of tests by juggling the controls, they tried other tricks to confuse it. They reversed the polarity of the current flowing in the electrolyte trough of one of the units. This meant that in that unit, moving the magnet in a given direction now produced an opposite effect to that originally built into the machine. It adapted, and speedily was happy again. Then they reversed the connections between the output of one unit and the input of another. It adapted.

Then they took out one magnet and turned it around, reversing its magnetic polarity. Homeostat was equal to the

change, and got happy as all get out, in jigtime. By this time the experimenters were beginning to get a little puzzled. Just what *did* you have to do, short of smashing the whole apparatus with a sledge, an impulse which occasionally cropped up, in order to keep it from readjusting itself and getting happy again?

They put bars across the troughs, so the arms could not swing through the full arc any more. It adapted. They connected two of the arms together with a light glass fiber, so that they had to move together. It adapted. Other, more complicated changes were made in the mechanism, and many different combinations of settings of the polarity and current controls were tried with each change. And each time, the Homeostat . . . but you get the idea. Unless the experimenters cut out the adaptive faculty itself, the relay and uniselector combination, the darned thing would adapt to anything they could devise, short of destruction by one means or another.

Well then, is the Homeostat a real brain? Obviously not, since it lacks certain of the attributes of the animal brain, but it *does* exhibit the all-important adaptive ability that the giant computers with their vast memory banks and electronic ganglia

lack. But the Homeostat certainly suggests that the day is not far off when a vastly refined version of the Homeostat, combined with equally advanced versions of the other devices with their huge memories and problem-solving powers, may be put together.

What then? What sort of performance could we expect from such a machine, and, above all what would we *do* with it, how far would we trust it? As to performance, there are some things that can be predicted with a good deal of confidence. For one thing, such a machine could, starting out right from scratch, learn to do very difficult and complex things in a way very like if not identical to that of a human being learning to play chess. It would *not* need to have built into it even such things as a knowledge of the rules of chess, such as the moves. These things could all be learned by trial and error, memorized, and the proper memory referred to whenever the situation on the board required it. And in the end, such a machine could wind up playing chess, even if the man who built it never himself even heard of the game. A complex and capacious machine could undoubtedly learn to play better chess than a simple one, and ultimately there appears to be no reason to

think a machine could be built to play better chess than any human being. The only possible limit would seem to be reached if a machine capable of comprehending all possible moves and combinations of moves and their effects were built. No machine capable of playing a better game than that could be built, for then the limitation would be in the game itself, rather than in the player. In other words, the outcome of any game played would be infallibly determined by the first move, as is the case for even we poor, bumbling human brains in games like tic-tac-toe.

Now so far, no human being seems to have come very close to a complete comprehension of chess, so that he infallibly wins every game in which he is permitted to say who shall make the first move. Yet chess is far, far less complex than our social or economic system, with their intricate, constantly changing paths of communication and influence, and their fantastically complex networks of internal and external feedbacks. This immediately suggests that, unless we can breed a strain of humans with superbrains to handle our complex social and economic problems, we may have to try building machines to do it. With humans, the very idea of a

strain of SUPER-humans seems to scare hell out of everybody, probably because nobody is now smart enough to figure out what the motives of such individuals would be. This may not remain true, as we get to understand ourselves better, particularly our emotional make-up. But it is pretty clear that the superman idea is not popular in present world opinion.

The first idea of many people would be that a machine, not having the human sex drive and other emotional factors, would be more trustworthy in solving our knotty problems. Many also think that it would be easier to get people to abide by the decisions of a machine, and put them into effect. It might even be possible to so connect the machine into our economic and social system that its "happiness" condition, and the general well-being, survival and happiness of mankind coincided. That might eliminate the chance that the machine would exhibit motives that conflicted with ours. It would, that is, IF we really knew what things are necessary for our own survival and happiness, and IF we had any way of insuring that our glorified machine would not make internal readjustments in itself so that its motives were no longer the same as ours.

But wait. It already appears that the better brain we make, the more it begins to resemble, and act like, animal brains. To be blunt, it begins to show a capacity for behaving neurotically, in a way all too familiar to psychiatrists. Remember, we are just groping our way into an understanding of how a brain works, in somewhat the manner of a man trying to examine the fine details of construction of a microscope, *with* the same microscope. Nobody can be sure yet that having emotions and a capacity for neurotic or insane conduct is not a basic property of *any* true brain, whether made of protoplasm or transistors. If such a thing is true, the only way to insure that a machine stayed sane and rational would be to keep continuous check on the functionality of each individual component. And even this would only prevent malfunctions due to breakdown of parts.

There appears at the moment to be no way to prove that such a complex machine would remain continuously rational, even if no parts failed, even though the Homeostat's performance is encouraging in this respect. Such checkups on the condition of parts would at least have one advantage over trying to deal with animal brains, however. About the only way we can do

much physical checking on the condition of parts of the animal brain is to look for tumors with X-ray or ultrasonic waves, or infections through less direct methods. Few people like the idea of having their skulls opened up just for an inspection of the works, unless there are already very serious indirect evidences of something wrong within, like a tendency to jump under subway trains, or bet against the Yankees.

But we cannot cheaply get ourselves off the hook by simply saying we won't ever build any vast God-machine to solve the gigantic problems the complexity of ourselves and our culture pose. We can't, because they are getting more complex and pressing faster than our capacity to deal with them is improving. The threat of a "thermonuclear" war is only one proof, and if more are needed, consider overpopulation, exhaustion of natural resources, mounting mental disease, or even possible invasion from outside our solar system. Sooner or later, we must get our capacity to solve problems to catch up with our ability to create them for ourselves, before one comes along that has a fatal penalty for non-solution. And since we seem to be much better able, or willing, to improve machines than we are to

improve ourselves, it looks as if we must come back to face the thorny question again.

All right, suppose at some future time we *do* build giant machines to deal with our problems. We can assume that such a machine brain would have an infallible memory capable of storing all the data that can be fed to it. It would also need an exhaustive cross-index, and very efficient means of searching for relevant data, or deciding what is relevant, for that matter. And of course, all data fed in at first would have to be "weighted," assigned a reliability rating based on how sure we are of the accuracy of data fed in. Later, the machine itself could reweight data, in the way that a human being does, as more data is accumulated. Indeed, this continuous re-evaluation of data in the light of new evidence might be one of the machine's greatest advantages, if it could be kept free of the egotism that causes humans to stubbornly persist in error in the face of truth.

Its problem-solving circuits would work at least as well as ours, and its decisions, based on total recall of all of the carefully weighted information, would be far better. And therein lies the rub. Being better decisions, some of them would certainly be different from what in our opinion

was the "best" course, and I can think of no way to keep people from having opinions about anything they can conceive of. That being so, would we have the nerve to accept the machine's directives and put them into practice? Would we, in other words, dare to trust our own handiwork when it truly performed its function of making better decisions than we can make? On the other hand, suppose we did, and suppose further that all appeared to be going well. How would we remain sure that the machine's "happiness" condition stayed coincident with ours. In other words, how could we be sure that the machine would not ultimately place maintaining its own survival and optimal condition (happiness) above ours? Remember, it is supremely adaptive, and smarter than anybody. That implies that it could surely deceive us as to its motives in any given decision, because one of the factors in all its decisions would be its vast knowledge of human capacities and limitations.

Of course, we could build another machine to check on the actions of the first one, but how we could keep it from taking the same path, I'm darned if I can see. Another thing, any machine-brain, however designed, is

bound to show some trend or bias, some general slant in its attack on problems that might turn out to have disastrous consequences for us. This is true because, in the design of such a brain-machine, there are many decisions that we would have to make on a more or less arbitrary basis. Should the plate voltages on the tubes be near the maximum permissible, or at some lower value? How quickly should a given circuit respond?

But supposing that in our future civilization all goes well, and the machine-brain remains the docile, obedient man-Friday, serving humanity with all of its great ability, solving all the knotty problems, making the world a place of increasingly easy existence for man. With no struggle to try our powers, no threat of dire consequences to spur us into efforts to improve ourselves, or even to maintain our present level of ability, what would happen to the human species? Would we indeed sink to the level of a mere parasite living on the bounty of the machine, or would we find within ourselves the drive to achieve and excel? To all these vexing and disquieting questions there seems to be only one really sure solution. Build the machine, and find out.