

THE CYBERNETIC BRAIN

6

STAFFORD BEER

FROM THE CYBERNETIC FACTORY
TO TANTRIC YOGA

Our topic changes character here. Grey Walter and Ross Ashby (and Gregory Bateson) were first-generation cyberneticians, born in the 1900s and active until around 1970. Stafford Beer and Gordon Pask were central figures in the second generation of British cybernetics, twenty years younger and active in cybernetics until their deaths in 2002 and 1996, respectively. What the two generations had in common was the defining interest in the adaptive brain. Where they diverged was in the question of how the brain fitted into their cybernetics. To a degree, Beer and Pask carried forward the attempt to build synthetic brains that they inherited from Walter and Ashby, in their work on biological and chemical computers discussed in this chapter and the next. Even there, however, the emphasis in Beer and Pask's work was not on understanding the brain per se, but in putting these "maverick machines," as Pask called them (Pask and Curran 1982, chap. 8), to work in the world. More generally, psychiatry was not a central concern for either Beer or Pask. Instead, they found inspiration in ideas about the adaptive brain in their extensions of cybernetics into new fields: Beer in his work in management and politics and even in his spiritual life; Pask in his work on training and teaching machines, and in the arts, entertainment, theater, and architecture. This is what interests me so much about the cybernetics of both men: the many projects they engaged in help us extend our range of examples of ontology in action. What also interests me is that, like Bateson and Laing, and unlike Ashby in his understanding of clinical psychiatry, Beer and Pask took the symmetric fork

in the road. The referent of their cybernetics was always reciprocally adapting systems.

I should add that Beer and Pask were extraordinary individuals. Beer displayed fabulous energy and creativity. Reading a diary that he kept during his first visit to the United States, from 23 April to 12 June 1960, leaves one limp (Beer 1994 [1960]); his career in management was accompanied by awesome literary productivity (in terms of quality as well as quantity), mainly on cybernetic management and politics, though he was also a published poet (Beer 1977); he painted pictures, and some of his works were displayed in Liverpool Cathedral and elsewhere (Beer 1993b); he also taught tantric yoga, loved women, slept only briefly, and drank continually (white wine mixed with water in his later years).

After an outline biography, I turn to Beer's work in management and politics, focusing in turn on his early work on biological computers, his viable system model of organizations, and the team synteegrity approach to decision making. Then we can examine the spiritual aspects of Beer's cybernetics and the cybernetic aspects of his spirituality. The chapter ends with an examination of the relation between Beer's cybernetics and Brian Eno's music.

— — — — —

Stafford Beer was born in Croydon, near London, on 25 September 1926, nearly five years the elder of two brothers (his younger brother Ian went on to be headmaster of Harrow Public School and on his retirement wrote a book called *But, Headmaster!* [2001]).¹ Like Ashby, Walter, and Pask, Stafford had a first name that he never used—Anthony—though he buried it more deeply than the others. His brother's third name was also Stafford, and when Ian was sixteen, Stafford "asked me to sign a document to promise that I would never use Stafford as part of my name. I could use it as I. D. S. Beer, or, indeed, using the four names together but he wanted the 'copyright' of Stafford Beer and so it was forever more."² Early in World War II, their mother, Doris, took Stafford and Ian to Wales to escape the German bombing, and at school there Stafford met Cynthia Hannaway, whom he married after the war. In 1942 the family returned to England, and Stafford completed his education at Whitgift School, where "he was a difficult pupil as he was found to be unsuitable for certain Sixth Form courses or he demanded to leave them for another. He could not stand the specialization and talked all the time of holistic teaching and so on. He wanted to study philosophy but that was not taught at school. He was precocious to a degree. A letter written by him was published in the

Spectator or the Economist, no-one could understand it.” He went on to study philosophy and psychology at University College London—which had then been evacuated to Aberystwyth, back in Wales—for one year, 1943–44.³ At University College he swam for the college team and was English Universities backstroke champion as well as getting a first in his first-year examinations. In 1944 he joined the British Army as a gunner in the Royal Artillery. In 1945 he went to India as a company commander in the Ninth Gurkha Rifles and later became staff captain intelligence in the Punjab. In 1947 he returned to Britain, remaining with the Army as army psychologist with the rank of captain.

Back in England, Beer married Cynthia, and they had six children together, though the first was stillborn. Following a divorce, Beer married Sallie Steadman, a widow and mother of a daughter, Kate, and they had two more children, for a total of eight, but this marriage, too, ended in divorce, in 1996. From 1974 onward Beer lived alone in Wales for much of the year (see below). In 1981 he met and fell in love with another cybernetician, Allenna Leonard (then a mature graduate student and later president of the American Society for Cybernetics), and she was Beer’s partner for the remainder of his life.

Leaving the army, Beer hoped to do a PhD in psychology at University College, but when told that he would have to recommence his studies as a first-year undergraduate he turned his back on the academic life, and in 1949 he began work for Samuel Fox in Sheffield, a subsidiary company of United Steel, where he created and ran its Operational Research Group (probably the first such group to exist in Britain outside the armed forces). From 1956 until 1961 he was head of the Operational Research and Cybernetics Group of United Steel, with more than seventy scientific staff based in the appropriately named (by Beer) Cybor House in Sheffield. In 1961 he founded Britain’s first operational research consulting firm, SIGMA (Science in General Management). In 1966 he moved on to become development director of the International Publishing Corporation (IPC), then the largest publishing company in the world, where his work largely concerned future initiatives around computing and information systems. In 1970, Beer left IPC “following a boardroom disagreement about development policy.” From 1970 until his death in Toronto on 23 August 2002 he operated as an independent consultant in a variety of arenas, some of which are discussed below.

Besides his career in management and consultancy, Beer was a prolific writer of scholarly and popular works, including more than two hundred publications and ten books on cybernetics, which he referred to as “ten pints



Figure 6.1. Beer as businessman. Source: Beer 1994a, facing p. 1. (This and other Beer images in this chapter, where otherwise unattributed, are courtesy of Cwarel Isaf Institute and Malik Management Zentrum St. Gallen [www.management.kybernetik.com, www.malik-mzsg.ch].)

of Beer” (Beer 2000). After 1970, he occupied many institutional roles and gained many honors. At different times he was president of the Operational Research Society, the Society for General Systems Research, and the World Organization of Systems and Cybernetics. He had several footholds in the academic world, though none of them full-time. His most enduring academic base was at the Business School of Manchester University, where he was visiting professor of cybernetics from 1969 to 1993. He was research professor of managerial cybernetics at University College Swansea from 1990 to 1997, visiting professor of management science at the University of Durham from 1990 to 1995, visiting professor of cybernetics at the University of Sunderland and life professor of organizational transformation at Liverpool John Moores University, both from 1997 until his death. And so on, including visiting professor-

ships at many other universities in Britain, Canada, Sweden (Stockholm), and the United States dating from 1970 onward. He was awarded major prizes for his work in operations research and cybernetics by the Operations Research Society of America, the American Society for Cybernetics, the Austrian Society for Cybernetics, and the World Organization of Systems and Cybernetics. A festschrift in Beer's honor was published in 2004 (Espejo 2004), and two volumes of his key papers have also appeared (Beer 1994a; Whittaker 2009).

Figure 6.1 is a photograph of Beer in the early 1960s when he was director of SIGMA—the smartly trimmed hair and beard, the three-piece suit, the cigar: the very model of a successful English businessman. In the early 1970s, however, Beer changed both his lifestyle and appearance. Partly, no doubt, this was in disgust at events in Chile with which he had been deeply involved, culminating in the Pinochet coup in 1973 (as discussed below). But also, as he told me, approaching his fiftieth birthday, he was moved to take stock of his life—"I had had two wives, I had eight children, a big house and a Rolls-Royce"—and the upshot of this stock taking was that in 1974 Beer renounced material possessions and went to live in a small stone cottage in a remote part of Wales.⁴ He retained the cottage for the rest of his life, but after the mideighties he divided his time between there and a small house he shared with Allenna Leonard in Toronto. This break in Beer's life was registered by a change in his appearance (fig. 6.2) and also in his writing style. Until this change, Beer's writing took a fairly conventional form. His first book in its wake was *Platform for Change: A Message from Stafford Beer*, printed on paper of four different colors, signaling different modes of argument and presentation. The introduction, printed on yellow paper, begins thus (Beer 1975, 1):

HELLO

I would like to talk to you
 if you have the time
 in a new sort of way
 about a new sort of world.

It ends (6):

I am fed up with hiding myself
 an actual human being
 behind the conventional anonymity
 of scholarly authorship.




Figure 6.2. Beer after the move to Wales. Source: Beer 1994a, 315. Photo: Hans-Ludwig Blohm. © Hans-Ludwig Blohm, Canada.)

From Operations Research to Cybernetics

Beer's route into cybernetics began with his work in operations research (OR) which in turn grew out of his work in the British Army in India. We do not need to delve deeply into the history of OR, but some brief remarks are relevant. As its name suggests, OR developed in World War II as a scientific approach to military operations. "Scientific" is to be contrasted here with traditional approaches to tactical and strategic planning based on the accumulated expertise of military commanders, and wartime OR can be broadly characterized in terms of a quantifying spirit aimed at modelling military activities with an eye to optimizing performance. One could try to calculate, for example, the optimal U-boat search pattern to be flown by a specified number of aircraft of given speed and range. OR was first developed in Britain

in conjunction with new radar technologies but was also taken to a high art in the United States.⁵

Beer was not himself involved in the wartime development of OR. On his own account, he rather wandered into it while he was in the army, first by attempting to use symbolic logic, which he had studied at University College, to organize large numbers of men into functioning systems.⁶ He first heard of OR as a field on his return to England and plunged himself into it as he moved into civilian life. Two early papers, published in 1953 and 1954, for example, outline novel statistical indices for measuring the productivity of manufacturing processes which he developed and implemented at the Samuel Fox steel company. These papers have a very practical bent, including ideas on how the sampling of productivity should be done and how the information could be systematically and routinely collected, assembled, and presented. The aim of the indices in question was the ability to forecast how long it would take to perform any given operation, a topic of interest both to the managers and customers of the mill (Beer 1953, 1954).

Beer's career in OR was very successful, as is evident from the biographical sketch above, and OR continued to play an important part throughout his subsequent work, both as an employee and as a consultant. But at an early stage he began to look beyond it. The second of the OR papers just mentioned is largely devoted to the development and use of performance measures for individual production operations in the factory, but it concludes with a section entitled "The Future Outlook" (also the title of Grey Walter's novel in its English publication two years later) looking forward to the development of "models . . . which would embrace the whole complex manufacturing structure of, say, an integrated steelworks." Beer notes that such models would themselves be very complex to construct and use and mentions some relevant mathematical techniques already deployed by OR practitioners, including game theory and linear programming, before continuing, "Advances in the increasingly discussed subject of cybernetics, allied with the complex models mentioned, might result in a fully mechanized form of control based on the technique described here" (1954, 57).

What did cybernetics mean, in assertions like that, for Beer, and how did it differ from OR? This takes us straight back to questions of ontology and a concept that I have been drawing on all along, that of an exceedingly complex system. Here we need only return briefly to its origin. In his first book, *Cybernetics and Management* (1959), Beer distinguished between three classes of systems (while insisting that they in fact shaded into one another): "simple," "complex," and "exceedingly complex" (fig. 6.3). He gave six examples of the

SYSTEMS	<i>Simple</i>	<i>Complex</i>	<i>Exceedingly complex</i>
Deterministic	Window catch	Electronic digital computer	EMPTY
	Billiards	Planetary system	
	Machine-shop lay-out	Automation	
Probabilistic	Penny tossing	Stockholding	The economy
	Jellyfish movements	Conditioned reflexes	The brain
	Statistical quality control	Industrial profitability	THE COMPANY

Figure 6.3. Beer's classification of systems. Source: S. Beer, *Cybernetics and Management* (London: English Universities Press, 1959), 18.

first two types (subdividing them further into “deterministic” and “probabilistic” systems). Under “simple” came the window catch, billiards, machine shop layout, penny tossing, jellyfish movements, and statistical quality control; under “complex” we find electronic digital computers, planetary systems, automation, stockholding, conditioned reflexes and industrial profitability. What those examples have in common, according to Beer, is that they are in principle knowable and predictable, and thus susceptible to the methods of the traditional sciences. Physics tells us about billiard balls; statistics about penny tossing; OR about stockholding and industrial profitability—this last, of course, being especially relevant to Beer. OR was, then, a classical science of production, a science appropriate to those aspects of the world that are knowable and predictable, in the same space as modern physics. However, under “exceedingly complex” systems (which, according to Beer, can have only probabilistic forms) we find just three examples: the economy, the brain, and the company. And Beer's claim was that these are “very different” (Beer 1959, 17):

The country's economy, for example, is so complex and so probabilistic that it does not seem reasonable to imagine that it will ever be fully described. The second, living, example—the human brain—is also described in this way. Moreover, it is notoriously inaccessible to examination. . . . Inferential investigations about its mode of working, from studies such as psychiatry and electroencephalography, are slowly progressing.

Probably the best example of an *industrial* system of this kind is the Company itself. This always seems to me very much like a cross between the first

two examples. The Company is certainly not alive, but it has to *behave* very much like a living organism. It is essential to the Company that it develops techniques for survival in a changing environment: it must adapt itself to its economic, commercial, social and political surroundings, and it must learn from experience.

Beer's exceedingly complex systems, were, then, as discussed already, in a different ontological space from the referents of OR (or physics). They were not fully knowable or adequately predictable, and they were "the province of cybernetics" (18). Beer's enduring goal was precisely to think about management cybernetically—to inquire into how one would run a company, or by extension any social organization, in the recognition that it had to function in and adapt to an endlessly surprising, fluctuating and changing environment.⁷

Toward the Cybernetic Factory

MY GOD, I'M A CYBERNETICIAN!

STAFFORD BEER, ON FIRST READING WIENER'S *CYBERNETICS* (BEER 1994C)

Beer first read Norbert Wiener's *Cybernetics* in 1950 and plunged into the field, establishing an individual presence in it and close personal connections as he went. By 1960, "I had known McCulloch for some years, and he would stay at my house on his Sheffield visits. . . . The British pioneers in cybernetics were all good friends—notably Ross Ashby, Frank George, Gordon Pask, Donald MacKay and Grey Walter" (Beer 1994 [1960], 229). "Norbert Wiener, as founder of cybernetics, was of course my great hero," but Beer did not meet him until his first trip to the United States when, on 25 May 1960, Wiener "almost vaulted over his desk to embrace me," greeting Beer with the words "I have become increasingly conscious that the growing reputation of my work [Wiener's] in Europe derives in large measure from your lectures and writings, and from the fact that you have built Cybor House. For this I should like to thank you" (Beer 1994 [1960], 281, 283).

In what follows, we will be largely concerned with connections between Beer's cybernetics and Ashby and Pask's. Beer and Pask actively collaborated in the work on biological and chemical computers discussed below and in the next chapter, and one can trace many parallels in the development of their work. But the defining features of Beer's cybernetics were Ashby's homeostat

as key model for thinking about adaptive systems and Ashby's law of requisite variety, as a tool for thinking realistically about possibilities for adaptive control. Much of what follows can be understood as a very creative extension of Ashby's cybernetics into and beyond the world of organizations and management. During the 1950s, Beer experimented with a whole range of cybernetic approaches to management (e.g., Beer 1956), but two ideas quickly came to dominate his thinking. First, one should think of the factory (or any complex organization) in analogy with a biological organism. Second, and more specifically, to be adaptive within an unknowable environment, the factory as organism should be equipped with an adaptive brain.

Beer laid out an early and striking version of this vision in a paper he presented to a symposium on self-organization held at the University of Illinois's Allerton Park on the 8 and 9 June 1960 (Beer 1962a). He opened the discussion with the notion of the "automatic factory," then attracting great interest, especially in the United States. This was a vision of industrial automation taken, one might think, to the limit. In the automatic factory, not only would individual machines and productive operations be controlled by other machines without human interference, but materials would be automatically routed from one operation to the next. In the "lights out" factory, as it was sometimes called, the entire production process would thus be conducted by machines, and human labor made redundant—literally as well as metaphorically.⁸

Beer was not at this stage a critic of the automatic factory, except that he did not feel it was automatic enough. He compared it to a "spinal dog"—that is, a dog whose nervous system had been surgically disconnected from the higher levels of its brain. The automatic factory (1962a, 164) "has a certain internal cohesion, and reflex faculties at the least. [But] When automation has finished its work, the analogy may be pursued in the pathology of the organism. For machines with over-sensitive feedback begin to 'hunt'—or develop ataxia; and the whole organism may be so specialized towards a particular environment that it ceases to be adaptive: a radical change in the market will lead to its extinction." Beer's argument was that to make it adaptive and to avoid extinction in market fluctuations, the automatic factory would need a brain.

At present, such an automatic factory must rely on the few men left at the top to supply the functions of a cerebrum. And . . . the whole organism is a strange one—for its brain is connected to the rest of its central nervous system at discrete intervals of time by the most tenuous of connections. The survival-value of such a creature does not appear to be high. . . .

This will not do. The spinal dog is short of a built-in cerebrum; and the automatic factory is short of a built-in brain. The research discussed in this paper is directed towards the creation of a brain artefact capable of running the company under the evolutionary criterion of survival. If this could be achieved, management would be freed for tasks of eugenics; for hastening or retarding the natural processes of growth and change, and for determining the deliberate creation or extinction of whole species. (Beer 1962a, 165)

The reference to eugenics is provocative to say the least, but the idea is an interesting one. The cybernetic factory, as Beer imagined it, would be *viable*—a key term for Beer: it would react to changing circumstances; it would grow and evolve like an organism or species, all without any human intervention at all. The role of humans in production would thus become that of metamanagement—managers would survey the field of viable production units and decide on which to promote or retard according to metacriteria residing at a level higher than production itself. Figure 6.4 is Beer’s schematic vision of what the cybernetic factory should look like, and much of his essay is devoted




Figure 6.4. Schematic of the cybernetic factory. Source: Beer 1994a, 192.

to a formal, set-theoretic definition of its contents and their relations. This is not the place to go into the details of the formalism; for present purposes, the important components of the diagram are arranged around the circumference: the T- and V-machines at the left and right, bridged by the U-machine and “states of the world” at the bottom. The symbols within the circumference represent processes internal to the U-machine.

Beer envisaged the T-machine as something like Pitts and McCulloch’s scanning device (Pitts and McCulloch 1947, discussed in chap. 3) updated in the light of more recent neurophysiological research. The “senses” of the T-machine would be numerical inputs representing the state of the factory’s environment (supplies and orders, finance) and its internal state (stocks, performance measures, etc.). The function of the T-machine was “scansion, grouping and pattern recognition” (Beer 1962a, 173). It would, that is, turn atomistic raw data into a meaningful output, in much the same way as the human brain picks out “universals” from our sensory data. The V-machine was conceived essentially as a T-machine running backward. Its inputs would be framed in the language of T-machine outputs; its outputs would be instructions to the motor organs of the plant—directing production operations and flows, ordering stock, or whatever.

Between the T- and V-machines lay, yes, the U-machine. The U-machine was to be “some form of Ashbean ultrastable machine” (Beer 1962a, 189)—a homeostat, the brain artifact of the firm. The job of the U-machine was continually to reconfigure itself in search of a stable and mutually satisfactory relationship between the firm and its environment. The U-machine was thus the organ that would enable the factory to cope with an always fluctuating and changing, never definitively knowable environment. It was the organ that could take the automatic factory to a level of consciousness beyond that of a spinal dog. Figure 6.5 summed up Beer’s abstract presentation, accompanied by the words “The temptation to make the outline look like a coronal section of the living brain was irresistible and I apologize to cerebra everywhere for such insolence” (197).⁹

The second major section of Beer’s essay was a progress report on how far he had gone toward realizing a cybernetic factory at the Templeborough Rolling Mills, a division of United Steel engaged in the manufacture of steel rods.¹⁰ This can help us think more concretely about the cybernetic factory, and here we need to refer to figure 6.6. The top level of the diagram represents various material systems relating to the flow of steel within the plant and their interconnections: the “Supplying system” feeds the “Input stocking system” which feeds the “Producing system,” and so on. The next level down,




Figure 6.5. The cybernetic factory as brain. Painting by Stafford Beer. The *T*, *U*, and *V* machines are labeled on the smaller painting in the bottom left. Source: Beer 1994a, 198, fig. 3.

“Sensations,” is the most important. Nineteen “sensations” are shown in the diagram, running from “a. tons bought” to “s. tons requested.” Each of these sensations should be understood as taking the form of numerical data relating to aspects of the plant or its environment—the current state of production, the profit and loss account, the balance sheet, as shown in lower levels of the figures. The “sensation” aspect of this diagram relates to the T-machine of Beer’s formal discussion, and his claim was to have sufficiently *simulated* a T-machine to make it clear that an automatic one could be built. The grouping of data into nineteen categories, for example, entailed “a large number of decisions . . . which, ultimately, the brain artefact itself is intended to take by its multiple multiplexing techniques. The research team in the field has, however, taken these decisions on an informed basis, by operational research methods” (Beer 1962a, 202).

The “sensations,” then, were to be considered inputs to the T-machine, and further numerical transformations were supposed to correspond to the




Figure 6.6. The steel mill as cybernetic factory. Source: Beer 1994a, 200–201, fig. 4.

functioning of “the T-Machine proper” (Beer 1962a, 203). These transformations, derived in practice from OR studies, first recombined the nineteen sensations into twelve “functions”—six referring primarily to the company and six to its environment. The functions all depended on ratios of expected behavior to actual behavior of precisely the form of the indices developed in Beer’s earlier OR work, discussed above. “This last point,” Beer wrote (204–5),

is important, since it incorporates in this exemplification the essential “black box” treatment of unknowns and imponderables common to all cybernetic machines. For a model of performance in any field may be inadequate: predictions and judgements based upon it will be effectual only insofar as the model is adequate. But in exceedingly complex and probabilistic systems no analytic model can possibly be adequate. The answer to this paradox, which I have used successfully for 10 years, is to load the raw predictions of any analytic model with a continuous feedback measuring its own efficiency as a predictor. In this way, everything that went unrecognized in the analytic work, everything that proved too subtle to handle, even the errors incurred in making calculations, is “black boxed” into an unanalyseable weighting which is error-correcting.

Here, then, we have an example of one way in which Beer's cybernetics tried to handle the unknown—a predictor that reviewed its own performance in the name of predicting better.¹¹

The values of the twelve parameters were measured daily in the steel mill and “were plotted on boards in an Operations Room for the benefit of management, as a by-product of this research” (Beer 1962a, 205). A plot of a year's readings is shown in figure 6.7, which Beer referred to as an encephalogram (205). He was reaching here for a suggestive connection between his work in management and brain science à la Grey Walter, referring to emergent periodicities in the data and noting that the “encephalographer finds this structural component of information (the brain rhythm) of more importance than either its amplitude or voltage” (182). This tempting idea seems to have proved a red herring, alas; I am not aware of any subsequent development of it, by Beer or anyone else. Several other, readily automatable statistical and mathematical transformations of these data then followed, and the work of the T-machine, as simulated at Templeborough, was said to be complete. Given that “the T-Machine was said to be set-theoretically equivalent to a V-Machine,” the problem of constructing the latter could be said to have been shown to be soluble, too (208). But figure 6.4 also shows the intervention of the U-machine, the homeostatic brain, into the life of the cybernetic factory: what about that?




Figure 6.7. An EEG of the firm. Source: Beer 1994a, 206, fig. 5.

The outputs of the simulated T-machine in successive time steps were recorded at Templeborough as a “generalized gestalt memory” indicated in the lower left and right of figure 6.6, the left portion relating to inner states of the factory, the right to its environment. These memories could be thought of defining “two phase spaces in which the company and the environment can respectively operate.” And the U-machine was intended to search for a set of “preferred states” within this space via a “mutually vetoing system by which the homeostatic loop in the diagram continues to operate until both company and environmental points in phase-space (representing vectors of functions) lie in the appropriate preferred states set” (Beer 1962a, 208).¹² This notion of mutual or reciprocal vetoing was very important in Beer’s work (and Pask’s), so I want to digress briefly here to explain it.

The idea of mutual vetoing came directly from Ashby’s cybernetics, and here Beer, like Bateson and Pask, took the symmetric fork in the road. Imagine an interconnected setup of just two of Ashby’s homeostats, both of which are free to reconfigure themselves. Suppose homeostat 1 finds itself in an unstable situation in which its essential variable goes out of whack. In that case, its relay trips, and its uniselector moves to a new setting, changing the resistance of its circuit. Here one can say that homeostat 2—with its own internal parameters that define the transformation between its input from and output to homeostat 1—has *vetoed* the prior configuration of homeostat 1, kicking it into a new condition. And likewise, of course, when homeostat 2 finds itself out of equilibrium and changes to a new state, we can say that homeostat 1 has vetoed the first configuration of homeostat 2. Eventually, however, this reconfiguration will come to an end, when both homeostats achieve equilibrium at once, in a condition in which the essential variables of both remain within limits in their mutual interactions. And this equilibrium, we can then say, is the upshot of a reciprocal vetoing: it is the condition that obtains when the vetoing stops and each machine finds a state of dynamic equilibrium relative to the other’s parameters.

This is enough, I think, to unravel the above quotation from Beer. One can think of the U-machine and the firm’s environment as two reciprocally vetoing homeostats, and the U-machine itself attempts to find a relation between its inputs from the T-machine and its outputs to the V-machine that will keep some essential variable standing for the “health” of the company within limits. Beer never reached the stage of defining exactly what that essential variable should be at this stage in his work. For the sake of concreteness, we could imagine it as a measure of profitability, though Beer proposed interestingly different measures in subsequent projects that we can review below.

It was clear enough, then, what the U-machine should do, though in 1960 Beer still had no clear vision of how it should be made, and at Templeborough “management itself,” meaning the actual human managers of the plant, “plays the role of the U-Machine” (Beer 1962a, 208). The state of the art was thus that by that date a cybernetic factory had been simulated, though not actually built. Beer was confident that he could construct automated versions of the T-machine, as the factory’s sensory organ, and the V-machine, as its motor-organ equivalent. Neither of these had actually been constructed, but their working parts had been simulated by OR studies and data collection and transformation procedures. The U-machine, which figured out the desirable place for the factory to sit in the factory-environment phase space, continued to be purely human, simulated by the managers who would review the “gestalt memory” generated by the T-machine and figure out how to translate that into action via the inputs to the V-machine. The U-machine, then, was the key (209):

As far as the construction of cybernetic machinery is concerned, it is clear that the first component to transcend the status of mere exemplification must be the U-Machine. For exemplifications of T- and V-input are already available, and can be fed to a U-Machine in parallel with their equivalent reporting to management. . . . Having succeeded in operating the cybernetic U-Machine, the research will turn to constructing cybernetic T- and V-Machines. . . . After this, management would be free for the first time in history to manage, not the company in the language of the organism, but the T-U-V(R) control assembly in a metalanguage.

But what was the U-machine to be? Beer ended his talk at Allerton Park with the words “Before long a decision will be taken as to which fabric to use in the first attempt to build a U-Machine in actual hardware (or colloid, or protein)” (212). Colloid or protein?

Biological Computing

Beer’s thinking about the U-machine was informed by some strikingly imaginative work that he and Pask engaged in in the 1950s and early 1960s, both separately and together—work that continued Ashby’s goal of a synthetic brain but with an original twist. Ashby had built an adaptive electromagnetic device, the homeostat, which he argued illuminated the go of the adaptive brain. Following his lead, Beer and Pask realized that the world is, in effect, already full of such brains. Any adaptive biological system is precisely an

adaptive brain in this sense. This does not get one any further in understanding how the human brain, say, works, but it is an observation one might be able to exploit in practice. Instead of trying to build a superhomeostat to function as the U-machine—and Beer must have known in the mid-1950s that Ashby's DAMS project was not getting far—one could simply try to enroll some naturally occurring adaptive system as the U-machine. And during the second half of the 1950s, Beer had accordingly embarked on “an almost unbounded survey of naturally occurring systems in search of materials for the construction of cybernetic machines” (Beer 1959, 162). The idea was to find some lively system that could be induced to engage in a process of reciprocal vetoing with another lively system such as a factory, so that each would eventually settle down in some agreeable sector of its environment (now including each other).

In 1962 Beer published a brief and, alas, terminal report on the state of the art, which makes fairly mind-boggling reading (Beer 1962b), and we can glance at some of the systems he discussed there to get a flavor of this work. The list begins with quasi-organic electrochemical systems that Beer called “fungoids,” which he had worked on both alone and in collaboration with Pask. This was perhaps the aspect of the project that went furthest, but one has to assume Pask took the lead here, since he published several papers in this area in the late 1950s and early 1960s, so I postpone discussion of these systems to the next chapter. Then follows Beer's successful attempt to use positive and negative feedback to train young children (presumably his own) to solve simultaneous equations without teaching them the relevant mathematics—to turn the children into a performative (rather than cognitive) mathematical machine. Beer then moves on to discuss various thought experiments involving animals (1962b, 28–29):

Some effort was made to devise a “mouse” language which would enable mice to play this game—with cheese as a reward function. . . . In this way I was led to consider various kinds of animal, and various kinds of language (by which I mean intercommunicating boxes, ladders, see-saws, cages connected by pulleys and so forth). Rats and pigeons have both been studied for their learning abilities. . . . The *Machina Speculatrix* of Grey Walter might also be considered (with apologies to the organic molecule). . . . However no actual machines were built. . . . By the same token, bees, ants, termites, have all been systematically considered as components of self-organizing systems, and various “brainstorming” machines have been designed by both Pask and myself. But again none has been made.




Figure 6.8. The *Euglena* homeostat. Square, *Euglena* culture, with tropism displayed as shown; solid diamond, stimulus; circle, sensory receptor; hatched triangle, inhibiting influence, and, open triangle, stimulating influence, of *a*'s sensation on *b*'s stimulus. Source: Beer 1994a, 30, fig. 2.

Beer had, however, devoted most of his own efforts to systems composed from simpler organisms: colonies of *Daphnia*, a freshwater crustacean (Pask had considered *Aedes aegypti*, the larva of the yellow fever mosquito), of *Euglena* protozoa, and an entire pond ecosystem. The key question with all three systems was how to interest these biological entities in us, how to couple them to our concerns, how to make a U-machine that would respond to and care about the state of the cybernetic factory. And this coupling was where Beer's attempts foundered (1962b, 29):

Many experiments were made with [*Daphnia*]. Iron filings were included with dead leaves in the tank of *Daphnia*, which ingested sufficient of the former to respond to a magnetic field. Attempts were made to feed inputs to the colony of *Daphnia* by transducing environmental variables into electromagnets, while the outputs were the consequential changes in the electrical characteristics of the phase space produced by the adaptive behaviour of the colony. . . . However, there were many experimental problems. The most serious of these was the collapse of any incipient organization—apparently due to the steadily increasing suspension of tiny permanent magnets in the water.

Euglena are sensitive to light (and other disturbances) in interesting ways, and Beer sought to achieve optical couplings to a tank full of them “using a point source of light as the stimulus, and a photocell [to measure the absorption of light by the colony] as the sensory receptor” (fig. 6.8).

However, the culturing difficulties proved enormous. *Euglena* showed a distressing tendency to lie doggo, and attempts to isolate a more motile strain failed. So pure cultures were difficult to handle. Moreover, they are not, perhaps, ecologically stable systems. Dr. Gilbert, who had been trying to improve the *Euglena* cultures, suggested a potent thought. Why not use an entire ecological system, such as a pond? . . . Accordingly, over the past year, I have been conducting experiments with a large tank or pond. The contents of the tank were randomly sampled from ponds in Derbyshire and Surrey. Currently there are a few of the usual creatures visible to the naked eye (*Hydra*, *Cyclops*, *Daphnia*, and a leech); microscopically there is the expected multitude of micro-organisms. [The coupling is via light sources and photocells, as in the *Euglena* experiments.] . . . The state of this research at the moment is that I tinker with this tank from time to time in the middle of the night. (Beer 1962b, 31–32)

Clearly, however, Beer failed to enroll the pond ecosystem, too, as a U-machine. The cybernetic factory never got beyond the simulation stage; we do not live in a world where production is run by *Daphnia* and leeches, and Beer's 1962 status report proved to be a requiem for this work. I now want to comment on it ontologically and sociologically, before moving on to later phases in Beer's career in management.

Ontology and Design

The sheer oddity of trying to use a pond to manage a factory dramatizes the point that ontology makes a difference. If one imagines the world as populated by a multiplicity of interacting exceedingly complex systems, as modelled by Ashby's homeostats, then one just might come up with this idea. It follows on from what has gone before, though even then some sort of creative leap is required. In contrast, it is hard to see how one would ever come to think this way from a modern technoscientific perspective. One would think instead of trying to program a computer to do the job of management, but that is a very different approach, in ways that are worth pondering.

We could start with issues of representation and performance. In the discussion that followed Beer's presentation at the 1960 Allerton conference, Beer made an interesting contrast between digital and biological computing in just these terms. When the subject of the former came up, he remarked that "this analogy with computers I do not like for two reasons." One had to do with the dynamics of memory and whether memory should be understood like the

storage of “a parcel in a cloakroom” or as a “path of facilitation through phase space.” The other went like this (1962a, 220–21):

The big electronic machines . . . are preoccupied with digital access. Now why is this? It is always possible, given an output channel which you can fit on somewhere, to say what is happening just there, and to get an enormous printout. Now we [Beer and Pask] are not concerned with digital access, but with outcomes. Why do we pay so much money to make it [digital output] available? In the sort of machines that Gordon and I have been concerned with, you cannot get at the intermediate answer. If you take out [one?] of Gordon’s dishes of colloid, you may be effectively inverting a matrix of the order 20,000. The cost of the computer is perhaps 10 cents. The only trouble is you do not know what the answer is. Now this sounds absurdly naïve, but it is not, you know, because you do not *want* the answer. What you do want is to *use* this answer. So why ever digitise it?

We are back to the notion of representation as a detour away from performance. Digital computing, in this sense, is an enormous detour away from its object—the functioning of a factory for example—into and through a world of symbols. In the previous chapter we discussed the discovery at Kingsley Hall and Archway that this detour could be drastically shortened or even done away with in therapeutic practice. But Beer *started* from this realization: in a world of exceedingly complex systems, for which any representation can only be provisional, performance is what we need to care about. The important thing is that the firm adapts to its ever-changing environment, not that we find the right representation of either entity. As ontological theater, then, Beer and Pask’s biological computers stage this performative ontology vividly for us, dispensing entirely with representation, both exemplifying an ontology of sheer performance and indicating how one might go on in computing if one took it seriously. I could note here that this concern for performance and a suspicion of representation per se is a theme that ran through all of Beer’s work.¹³

There is second and related sense of a detour that also deserves attention here. As Beer put it (1962a, 209, 215), “As a constructor of machines man has become accustomed to regard his materials as inert lumps of matter which have to be fashioned and assembled to make a useful system. He does not normally think first of materials as having an intrinsically high variety which has to be constrained. . . . [But] we do not want a lot of bits and pieces which we have got to put together. Because once we settle for [that], we have got to

have a blueprint. We have got to design the damn thing; and that is just what we do not want to do.” The echoes of Ashby on DAMS and the blueprint attitude are clear. We are back to the contrasting conceptions of design that go with the modern ontology of knowable systems and the cybernetic ontology of unknowable ones. Within the frame of modern science and engineering, design entails figuring out what needs to be done to achieve some result and then arranging “inert lumps of matter” to achieve those specifications. Digital computers depend on this sort of design, specifying material configurations right down to the molecular level of chemical elements on silicon chips. Beer’s idea instead was, as we have seen, to find lively (not inert) chunks of matter and to try to enroll their agency directly into his projects. This gets us back to the discussion of the hylozoist quality of biofeedback music (chap. 3) and the idea that it’s all there already in nature (as in the extraction of music from the material brain). We could say that the modern stance on design has no faith in matter and relies upon human representations and agency to achieve its effects. The cybernetic ontology, as Beer staged it, entailed a faith in the agency of matter: whatever ends we aim at, some chunk of nature probably already exists that can help us along the way. We don’t need these long detours through modern design. We can explore Beer’s hylozoism further later in the chapter in a broader discussion of his spirituality.

There is, of course, yet a third sense of detour that comes to mind here. The mastery of matter, from the molecular level upward, required to build a digital computer has been painstakingly acquired over centuries of technoscientific effort. Beer’s argument was, in effect, that perhaps we didn’t need to make the trek. Just to be able to suggest that is another striking manifestation of the difference that ontology makes.

Now, Heidegger. It makes sense to see modern computer engineering as operating in the mode of enframing. It is not that semiconductor engineers, for example, have actually achieved some magical mastery over matter. For all their representational knowledge, they remain, like the rest of us, in *medias res*, obliged to struggle with the performance of obstinate stuff (Lécuyer and Brock 2006). Nevertheless, a successful chip is one that fits in with our preconceived plans: the chip either manipulates binary variables in a regular fashion, or it does not—in which case it is junk. Bending matter to our will like that is just what Heidegger meant by enframing. And then we can begin, at least, to see that the cybernetic ontology in this instance has more in common with a stance of revealing. Beer wanted to find out what the world—assemblages of mice, *Daphnia*, his local pond—could offer us. Against this, one might argue that Beer had some definite end in view: a replacement for the

human manager of the factory. But the important point to note is that the pond was not envisaged as an identical substitute for the human. We will see in the next chapter that Pask, who thought this through in print further than Beer, was clear that biological computers would have their own management style, not identical to any human manager—and that we would, indeed, have to find out what that style was, and whether we could adapt to and live with it. This is the sense in which this form of cybernetic design in the thick of things is a stance of revealing rather than enframing.

One last thought in this connection. Somewhere along the line when one tries to get grips with Beer on biological computing, an apparent paradox surfaces. Beer's goal, all along, was to improve management. The cybernetic factory was supposed to be an improvement on existing factories with their human managers. And yet the cybernetic brain of the factory was supposed to be a colony of insects, some dead leaves for them to feed on, the odd leech. Did Beer really think that his local pond was cleverer than he was? In a way, the answer has to be that he did, but we should be clear what way that was. Recall that Beer thought that the economic environment of the factory was itself an exceedingly complex system, ultimately unknowable and always becoming something new. He therefore felt that this environment would always be setting managers problems that our usual modes of cognition are simply unable to solve. This connects straight back to the above remarks on Beer's scepticism toward representational knowledge. On the other hand, according to Beer, biological systems *can* solve these problems that are beyond our cognitive capacity. They *can* adapt to unforeseeable fluctuations and changes. The pond survives. Our bodies maintain our temperatures close to constant whatever we eat, whatever we do, in all sorts of physical environments. It seems more than likely that if we were given conscious control over all the parameters that bear on our internal milieu, our cognitive abilities would not prove equal to the task of maintaining our essential variables within bounds and we would quickly die. This, then, is the sense in which Beer thought that ecosystems are smarter than we are—not in their representational cognitive abilities, which one might think are nonexistent, but in their performative ability to solve problems that exceed our cognitive ones. In biological computers, the hope was that “solutions to problems simply grow” (1962a, 211).

The Social Basis of Beer's Cybernetics

At United Steel, Beer was the director of a large operations research group, members of which he involved in the simulation of the cybernetic factory at

the Templeborough Rolling Mills. This was a serious OR exercise, supported by his company. The key ingredient, however, in moving from the simulation to the cybernetic reality, was the U-machine, and, as Beer remarked in opening his 1962 status report on biological computing, “everything that follows is very much a spare time activity for me, although I am doing my best to keep the work alive—for I have a conviction that it will ultimately pay off. Ideally, an endowed project is required to finance my company’s Cybernetic Research Unit in this fundamental work” (1962b, 25). I quoted Beer above on tinkering with tanks in the middle of the night, evidently at home, and Beer’s daughter Vanilla has, in fact, fond childhood memories of weekend walks with her father to collect water from local ponds (conversation with the author, 22 June 2002). We are back once more on the terrain of amateurism, ten years after Walter had worked at home on his tortoises and Ashby on his homeostat.

Again, then, a distinctive cybernetic initiative sprang up and flourished for some years in a private space, outside any established social institution. And, as usual, one can see why that was. Beer’s work looked wrong. Tinkering with tanks full of pond water looked neither like OR nor like any plausible extension of OR. It was the kind of thing an academic biologist might do, but biologists are not concerned with managing factories. The other side of the protean quality of cybernetics meant that, in this instance, too, it had no obvious home, and the ontological mismatch found its parallel in the social world. I do not know whether Beer ever proposed to the higher management of United Steel or to the sponsors of his consulting company, SIGMA, that they should support his research on biological computing, but it is not surprising that he should be thinking wistfully of an endowed project in 1962, or that such was not forthcoming. We should, indeed, note that Beer failed to construct a working U-machine, or even a convincing prototype. This is, no doubt, part of the explanation for the collapse of Beer’s (and Pask’s) research in this area after 1962. But it is only part of the explanation. The electronic computer would not have got very far, either, if its development had been left solely to a handful of hobbyists.

Of course, Beer did not carry on his cybernetics in total isolation. As mentioned above, having read Wiener’s *Cybernetics* in 1950, he sought out and got to know many of the leading cyberneticians in the United States as well as Britain. In the process, he quickly became a highly respected member of the cybernetics community which existed transversely to the conventional institutions to which its members also belonged. It was Beer who first brought Ashby and Pask together, by inviting both of them to a lecture he gave in the

city hall in Sheffield in 1956, and his recollection of the meeting sheds some light on the characters of both (S. Beer 2001, 553): “Gordon was speaking in his familiar style—evocative, mercurial, allusory. He would wave his arms about and try to capture some fleeting insight or to give expression to some half-formed thought. I was used to this—as I was to Ross’s rather punctilious manner. So Ashby would constantly interrupt Gordon’s stream of consciousness to say, ‘Excuse me, what exactly do you mean by that?’ or ‘Would you define that term?’ Both were somewhat frustrated, and the evening was close to disaster.” Beyond his personal involvement in the cybernetics community, Beer appreciated the importance of establishing a reliable social basis for the transmission and elaboration of cybernetics more than the other British cyberneticians. Ross Ashby also presented his work at the 1960 conference at which Beer presented “Towards the Cybernetic Factory,” and while there Beer conspired with Heinz von Foerster to offer Ashby the position that took him to the University of Illinois (Beer 1994 [1960], 299–301). In the second half of the 1960s, when Beer was development director of the International Publishing Corporation, he conceived the idea of establishing a National Institute of Cybernetics at the new Brunel University in Uxbridge, London, aiming to create academic positions for both Gordon Pask and Frank George. Beer persuaded the chairman of IPC, Cecil King, to fund part of the endowment for the institute and a fund-raising dinner for the great and good of the British establishment was planned (with Lord Mountbatten, the queen’s uncle, and Angus Ogilvy, the husband of Princess Alexandra, among the guests). Unfortunately, before the dinner could take place there was a palace coup at IPC—“in which, ironically, I [Beer] was involved”—which resulted in the replacement of King by Hugh Cudlipp as chairman.

I had never managed to explain even the rudiments of cybernetics to him [Cudlipp]. Moreover, it is probably fair to say that he was not one of my greatest fans. . . . At any rate the dinner broke up in some disorder, without a single donation forthcoming. Dr Topping [the vice-chancellor at Brunel] went ahead with the plan insofar as he was able, based on the solitary commitment that Cecil King had made which the new Chairman was too late to withdraw. Gordon was greatly disappointed, and he could not bring his own operation (as he had intended) [System Research, discussed in the next chapter] into the ambit of the diminished Institute which soon became a simple department at Brunel. The funding was just not there. However, both he and Frank George used their Chairs on the diminished scale. (S. Beer 2001, 557)

Though Beer had not fully achieved his ambition, the establishment of the Department of Cybernetics at Brunel was the zenith of the institutional career of cybernetics in Britain, and we shall see in the next chapter that Pask made good use of his position there in training a third generation of cyberneticians. Characteristically, the trajectory of cybernetics in Britain was further refracted at Brunel, with Pask's PhD students focusing on such topics as teaching machines and architecture. The Brunel department closed down in the early 1980s, and, given the lack of other institutional initiatives, these students were once more left to improvise a basis for their careers.¹⁴

In the 1960s, then, Beer helped find academic positions for three of Britain's leading cyberneticians and played a major role in establishing an academic department of cybernetics. Conversely, as remarked already, in 1974 Beer effectively deinstitutionalized himself in moving to a cottage in Wales. Partly, as I said, this was an aspect of an overall shift in lifestyle; partly it was a response to events in Chile. Partly, too, I think, it was a reflection of his failure in the later 1960s to persuade Britain's Labour government of the importance of cybernetics. He wrote of his "disappointment in the performance of Harold Wilson's 'white heat of technology' government. This was operating at a barely perceptible glow, and the ministers with whom I had been trying to design a whole new strategy for national computing failed to exert any real clout. There were five ministers involved—the Postmaster General himself (John Stonehouse) 'did a runner' and was discovered much later in Australia" (S. Beer 2001, 556). Beer was an exceptionally well connected spokesman for cybernetics in the 1960s, but the fruits of his efforts were relatively few. As he once put it to me, speaking of the sixties, "the Establishment beat us" (phone conversation, 3 June 1999).¹⁵

The Afterlife of Biological Computing

Neither Beer nor Pask ever repudiated his biological computer work; both continued to mention it favorably after the 1960s. In his 1982 popular book, *Micro Man*, Pask discusses a variety of "maverick machines," including his electrochemical systems, which he describes as "dendritic." He mentions that improved versions of them have been built by R. M. Stewart in California and comments that "there is now a demand for such devices, which are appropriate to non-logical forms of computation, but dendrites . . . are physically too cumbersome for such demand to be met practically. It now seems that

biological media may perform in similar fashion but on a more manageable scale” (Pask and Curran 1982, 135). A few pages later he actually reproduces a picture of a pond, with the caption “A real-life modular processor?” Likewise, Beer in the text he wrote for a popular book on the history of computing, *Pebbles to Computers*: “Some thirty years ago, some scientists began to think that biological computers might be constructed to outpace even electronic achievement. At that time it was not clear that transistors themselves would become reliable! Attempts were made to implicate living cells—microorganisms—in computations. In England in the ’fifties, one such computer solved an equation in four hours that a bright school girl or boy could solve in (maximum) four minutes. Its time had not yet come!” (Blohm, Beer, and Suzuki 1986, 13).

Biological computing enjoyed a happier fate in science fiction, making its way into the popular imagination. With Beer’s experiments on mice with cheese as a “reward function” we are surely in the presence of the mouse-computer that turns up in both Douglas Adams’s *The Hitchhiker’s Guide to the Galaxy* (1979) and Terry Pratchett’s *Discworld* series of fantasy novels.¹⁶ The most convincing representations of biological computing that I have come across include the obviously organic control systems of alien space ships that featured in various episodes of *Doctor Who* and, more recently, in Greg Bear’s novel *Slant* (1997), which includes a biological computer called Roddy (recombinant optimized DNA device) that is an entire ecosystem of bees, wasps, ants, peas, and bacteria (and which succeeds in subverting the world’s most sophisticated conventional AI, Jill).

And back in the material world biological computing has, in fact, recently been experiencing a resurgence. Figure 6.9 shows a cockroach-controlled robot, recently built by Garnet Hertz in the Arts, Computing, Engineering Masters Program at the University of California, Irvine. A giant Madagascan cockroach stands on the white trackball at the top of the assembly, attached by Velcro on its back to the arm which loops above the other components. Motions of the cockroach’s legs rotate the trackball, which in turn controls the motions of the cart (much as a trackball can be used to control the motion of the cursor on a computer screen). Infrared sensors detect when the cart is approaching an obstacle and trigger the appropriate light from an array that surrounds the roach. Since roaches tend to avoid light, this causes the roach to head off in another direction. The entire assemblage thus explores its environment without hitting anything or getting stuck—ideally, at least. The cybernetic filiations of this robot are obvious. From one angle, it is a version of Grey

Walter's tortoise, five decades on. From the other, a lively biological agent replaces the precisely designed electronic circuitry of the tortoise's brain, exemplifying nicely the sense of "biological computing."¹⁷ Figure 6.10 shows another biorobot, this one built by Eduardo Kac as part of his installation *The Eighth Day*. This time, the robot is controlled by a slime mold. These machines have no functional purpose. They are artworks, staging for the viewer a cybernetic ontology of entrained lively nonhuman agency. We can return to the topic of cybernetic art at the end of this chapter. For now, we might note that back in the 1950s and early 1960s Beer and Pask were aiming at something much more ambitious than Hertz and Kac, to latch onto the adaptive properties of biological systems, rather than their basic tropic tendencies.¹⁸



Figure 6.9. Cockroach-controlled robot. (Photograph by Garnet Hertz. Used by permission.)



Figure 6.10. Eduardo Kac, *The Eighth Day*, 2001. Transgenic artwork with biological robot (biobot), GFP plants, GFP amoebae, GFP fish, GFP mice, audio, video, and Internet (dimensions variable). The photograph shows the biobot in the studio, with its internal amoebae already in place, before it was introduced into the transgenic ecology that constitutes *The Eighth Day*. Source: www.ekac.org/8thday.html. Used courtesy of Eduardo Kac.

The Viable System Model

When Beer's dreams of biological computing came to an end in the early 1960s, this implied not an abandonment of his vision of the cybernetic factory but a transformation of it. Beginning in 1972, a trilogy of books developed his account of what he called the viable system model—the VSM for short: *Brain of the Firm* (1972; 2nd ed., 1981), *The Heart of the Enterprise* (1979), and

Diagnosing the System for Organizations (1985). The VSM was at the forefront of Beer's thinking and consulting work from the 1960s to the 1990s and attracted a considerable following. A two-day workshop on the VSM held at the Manchester Business School in January 1986 led to the production of an edited volume describing further interpretations and applications of the VSM by a range of academics, consultants, and people in industry and the military (Espejo and Harnden 1989), and variants of the VSM are still practiced and taught today.

The VSM transformed Beer's earlier vision of the cybernetic factory along two axes. First, the simulation of the cybernetic factory discussed above, where human management filled in for the not-yet-built U-machine, became in effect the thing itself. Beer continued to look forward to as much computerization of information gathering, transmission, and transformation as possible (as in the T- and V-machines). But the ambition to dispense with the human entirely was abandoned. Instead, human managers were to be positioned within purposefully designed information flows at just those points that would have been occupied by adaptive ponds or whatever (e.g., the position they in fact occupied in the earlier simulations).

Second, Beer extended and elaborated his conception of information flows considerably. In *Brain of the Firm*, the first of the VSM trilogy, he argued thus: The aim of the firm had, as usual, to be to survive in an environment that was not just fluctuating but also changing—as new technologies appeared in the field of production and consumption for example. How was this to be accomplished? What would a viable firm look like? The place to look for inspiration, according to Beer, was again nature, but now nature as the source of inspiration in the design of viable organizations, rather than nature as the immediate source of adaptive materials. Beer's idea was to read biological organisms as exemplary of the structure of viable systems in general, and to transplant the key features of their organization to the structure of the firm. In particular, he chose the human nervous system as his model. In the VSM, then, Beer's strategy was to transplant the organic into the social, but not as literally as before. The firm would no longer contain trained mice or *Daphnia* at its heart; instead, information flows and processing would be laid out as a diagram of human bodily flows and transformations.

The spirit of the VSM is strikingly expressed in the juxtaposition of two figures from *Brain of the Firm*. Figure 6.11A is a schematic of the body; figure 6.11B is a schematic of the firm. *Brain* goes into considerable detail in rehearsing the then-current understanding of human neurophysiology—the pathways both




Figure 6.11. Control systems: *A*, in the human body; *B*, in the firm. Source: S. Beer, *Brain of the Firm*, 2nd ed. (New York: Wiley, 1981), 131, figs. 23, 22. Permission: John Wiley & Sons.

nervous and biochemical along which information flows, the operations performed upon it at different points—and how it might be transcribed into the organization of the firm. I am not going to attempt an extensive review. But some key features need to be singled out. The VSM divided the “nervous system” of the firm into five subsystems, numbered 1–5 in figure 6.11B. Although the VSM was supposed to be applicable to any organization (or any viable system whatsoever), for illustrative purposes figure 6.11B is a diagram of a company having four subsidiaries, labeled 1A–1D, and as indicated in figure 6.11A, one can think of these in analogy to systems in the body, controlling the limbs, the heart, the kidneys, and so on. A notion of autonomy arises here, because such systems in the body largely control themselves without reference to the higher levels of the brain. The heart just speeds up or slows down without our ever having to think about it. It adapts to the conditions it finds itself in by reflex action largely mediated somewhere down the spinal column. Beer’s contention was that subsidiaries of the firm should be like that. They should act in the world and on one another (supplying materials to one another, say) as indicated by the circles with wavy lines and arrows moving off to the left from them, and their performance would be monitored at appropriate points on the “spinal column”—the square boxes labeled 1A and so on. This monitoring would consist of a comparison of their performance in relation to a plan already given by the higher management of the firm, and deviations could

be compensated for by appropriate adjustments to their behavior. The model here would be a simple servo-controlled negative feedback mechanism.

But even at this level of the body, autonomy is not complete. Figure 6.11B shows direct connections between the control systems, 1A, 1B, and so on. The idea here is that if something unusual is happening in subsidiary 1A, say, which supplies parts to 1B, then 1B should know about it so that it can take steps to allow for that. There must, that is, be some information channel linking these subsidiaries, as there is between the heart and the lungs. And further, Beer observed, in the human body there are usually several different channels linking levels of the nervous system. Figure 6.11A thus distinguishes two further channels—the sympathetic and the parasympathetic systems—and figure 6.11B shows their equivalents—lines of information flow upward, from the controllers on the spinal cord (the squares) and from the operational sites (the circles). The equivalent of the sympathetic system is system 2 of the VSM. Beer understood this as attempting to damp conflicts that could arise at the system 1 level—the various subsidiaries trying to hoard some material in short supply to each other's detriment, for example. This damping, which Beer knew enough not to expect to be necessarily successful, would be accomplished by reference to system 3. Corresponding to the pons and the medulla at the base of the brain, system 3 would be basically an operations research group, running models of the performance of the entire ensemble of subsidiaries, and thus capable, in principle, of resolving conflicts between subsidiaries in the light of a vision available to none of the subsidiaries alone.¹⁹

At this stage, no information has traveled upward beyond system 3 into higher layers of management. The parasympathetic system, however, was envisaged to act somewhat differently. This traveled straight up to system 3 and was intended to transmit an “algedonic” “cry of pain.” Less metaphorically, production data would be monitored in terms of a set of dimensionless ratios of potential to actual performance of the kind that Beer had introduced in his 1953, paper discussed earlier. If one of those ratios departed from a predecided range, this information would be automatically passed onward to system 3, which, in the light of its models, would act as a filter, deciding whether to pass it on to levels 4 and possibly 5.²⁰

I am inclined to think that system 4 was Beer's favorite bit of the VSM. The equivalent of the diencephalon and ganglia of the human brain, this had access to all the information on the performance of the firm that was not filtered out by system 3; it was also the level that looked directly outward on the state of the world. If the level 1 systems had access to information directly relating to their own operations, such as rising or falling stockpiles or order books,



Figure 6.12. World War II operations room, near London, during the Battle of Britain. Source: Beer 1968a, 23.

level 4 had much wider access, to national economic policies and changes therein, say, to the price of money, the results of market research, and what have you. System 4 was, then, the T-U-V system of Beer's earlier model, with the humans left in.

Beer envisaged system 4 as a very definite place. It was, in fact, modelled on a World War II operations room, of the kind shown in figure 6.12 (taken from Beer's 1968 book *Management Science*), as developed further by NASA at "Mission Control in the Space Centre at Houston, Texas, where the real-time command of space operations is conducted" (Beer 1981, 193–94), and updated with all of the decision aids Beer could think of (194–97). All of the information on the state of the firm and of the world was to be presented visually rather than numerically—graphically, as we would now say. Dynamic computer models would enable projections into the future of decisions made by management. Simply by turning knobs (197), managers could explore the effects of, say, investing more money in new plant or of trends in consumption. Feedbacks that had passed the level 3 filters would also arrive at system 4 from the lower levels, "signalled appropriately—that is, if necessary, with flashing red lights and the ringing of bells" (194), alerting management to emerging production problems, perhaps to be passed on again to level 5. In terms of social organization, "I propose a control centre for the corporation which is in continuous activity. This will be the physical embodiment of any System 4. All senior formal meetings would be held there; and the rest of the

time, all senior executives would treat it as a kind of club room. PAPER WOULD BE BANNED FROM THIS PLACE. It is what the Greeks called a *phrontisterion*—a thinking shop” (194).

System 4, then, differed from system 3 in the range of its vision, a vision which now encompassed the future as well as the present, and Beer imagined system 4 as a primary locus for decision making on change. If the levels below strove to implement given production plans for the firm, level 4 was the level at which such plans were drawn up and modified.

Finally we arrive at system 5, the equivalent of the human cortex. This was the level where policies were deliberated upon and the most consequential decisions were made (Beer 1981, 201). This was where the human directors of the firm had been imagined to continue to exist in the original blueprint for the cybernetic factory. The firm’s viability and continued existence, and even growth and evolution, were maintained by systems 1–4. The job of system 5 was, therefore, to think big at a metalevel superior to questions of mere viability.²¹

This outline of the VSM is now almost complete, but two points need to be added. First, the various levels of the viable system were intended to be coupled *adaptively* to one another. The 3 and 4 systems, for example, would engage in the process of reciprocal vetoing discussed earlier. Level 4 might propose some change in the overall operating plan for the firm; this would be run through the OR models at level 3 and might be rejected there—perhaps it would place excessive strain on one of the subsidiaries. Level 3 could then propose some modified plan back to level 4, which could run it through its models. Perhaps the plan would be vetoed again, once more transformed, and returned to level 3. And so on, back and forth, until some operating plan agreeable to both systems 3 and 4 was discovered.

Second, we should note a *recursive* aspect of the VSM. Beer argued that firms were themselves parts of bigger systems—national economies, say. The entire 1–5 structure of the firm would thus appear as a single system 1 on a diagram of the national economy. This in turn should be a viable system with its own levels 2–5 overseeing the ensemble of firms. Proceeding down the scale instead of up it, each subsidiary of the firm should also be a viable system in its own right, meaning that the level 1 systems of figure 6.11 should actually have their own levels 1–5 within them. Figure 6.13 shows what became Beer’s standard diagram of the VSM, depicting two levels of such recursion. The two level 1 subsidiaries in square boxes at the lower end of the spinal column (running up the right-hand side) are shown as having their own 1–5 structure projecting downward at an angle of 45 degrees (and each has two subsidiary




Figure 6.13. The viable system model showing recursive embeddings. Source: S. Beer, *Diagnosing the System for Organizations* (New York: Wiley, 1985), 136, fig. 37.

“operations” within its large circle). Likewise, the 3-4-5 brain at the top of the spinal column is itself enclosed in a square box, indicating that it is part of a level 1 system of some bigger system. Beer felt that such recursivity was a necessary property of viable systems—they had to be nested inside one another “like so many Russian dolls or Chinese boxes” in a chain of embeddings which “descends to cells and molecules and ascends to the planet and its universe” (Beer 1989a, 22, 25).

The VSM was thus a vision of the firm in the image of man. Especially, functions of management and control were envisaged on the lines of the human brain and nervous system. The brain and nervous system were simulated by a combination of information technologies and real human beings appropriately arranged. One could say that the VSM is one of the most elaborated images of the cyborg in postwar history, though the word “cyborg” is tautologous here, standing as it does for “cybernetic organism.” Any viable system was exactly that, according to Beer. We should also note that the VSM was the “circuit diagram” (Beer 1981, 123) of a time machine, an adaptive system accommodating itself to the exigencies of the unknown in real time, ranging from mundane disturbances at the level of production to world-historical changes.

The VSM as Ontology and Epistemology

The basic ontological vision that the VSM conjures up is the same as that of the cybernetic factory before it: the world as an ungraspable and unmasterable space of becoming; the organization as open-endedly and performatively adaptable. The VSM, however, also suggests some refinements to that picture. First, my portrayal of the cybernetic factory was centered on the brain of the firm as a unitary entity, the U-machine, in dialogic conversation with the firm’s environment. Beer’s conception of the VSM, in contrast, was one in which the overall behavior of the firm was the upshot of an interplay of many active but quasi-autonomous elements, the VSM’s systems 1–5, themselves also interacting with different aspects of the firm’s environment. The recursive aspect of the model adds an indefinite sequence layers of elements to this picture. The VSM thus moves us toward a vision of *ontological multiplicity*, a multiplicity which is, furthermore, *irreducible*: the system 3 of a given organization is not reducible to the organization’s system 4, say, or to the system 3 of another organization.²²

Second, we can return to the question of goals. Walter’s and Ashby’s devices had fixed goals that organized their adaptation: the homeostat reconfigured itself so as to keep its essential variables within preset limits. Beer’s concep-

tion of the VSM, in contrast, specified no goals whatsoever, except adaptation itself. And we could think of Heidegger: adaptation in the VSM was a process of revealing rather than enframing. The process that Beer called reciprocal vetoing between levels of the system, for example, was by no means as negative as the phrase suggests. A veto from one level to another was at the same time an invitation for a novel counterproposal, a way of finding out what the other had to offer.

The VSM was explicitly about information flows and transformations, so we can return now to a consideration of cybernetic epistemology as well as ontology. In Beer's vision, viable systems do contain knowledge—representations of their own inner workings and of their environment—principally enshrined in the OR models at level 3 of the VSM and the projective models at system 4. What should we make of this? First, we could recall that in his work on truly biological controllers, Beer had sought to avoid this detour through representation. His biological computers did not contain any representational or symbolic elements; they were intended simply to do their adaptive thing. The VSM, then, one might say, was a concession to representation as a response to the failure of biological computation. And it is appropriate to recall that, as I remarked before, Beer did not much trust representational models. He did not think, for example, that one could arrive at a uniquely correct model of the firm and its environment that could function unproblematically at level 4 of the VSM. This is a direct corollary of the idea that both the firm and its environment are exceedingly complex.

Beer did not, however, take this to imply that the construction of representational models was a useless endeavor. His idea, instead, was that the models in question should be continually examined and updated in relation to performance—“*continuously adapted*” (Beer 1981, 185) or even always “aborting” (Beer 1969 and 1994b, 151). The company should act in the light of the future projections of the model at level 4, but then actual developments in time should be compared with expectations from the model's simulations. These would not, in all probability, match, and the model should be adjusted accordingly.²³ The VSM thus stages for us an image of a performative epistemology—a more elaborated version of what we have seen in the preceding chapters. The “knowledge components” of the VSM were not an end in themselves; they were geared directly into performance as part of the mechanism of adaptation, and they were revisable in performance, just like the other components of the VSM; they were not the controlling center of the action.

Here I need to enter a caveat. What might adaptation of these models in practice mean? I just described adaptation in the VSM as open ended, but

Beer imagined and was prepared to implement something less than this in his models. He understood them as sets of mathematical equations linking long lists of variables such as demand, revenue, technological and economic change, dividends, share prices, and the money market. And the basic form of these sets of equations was not, in itself, revisable, at least as part of Beer's description of the regular functioning of a viable system. What could be revised in practice were the parameters figuring in these equations which specified the intensity of the couplings between variables. Beer's models were thus adaptive, but only to a degree, within a fixed overall form.²⁴

One further point. The symbolic models of the VSM were envisaged as conventional simulations programmed on digital computers. In this respect, there was no distinctively cybernetic aspect to the VSM. But it is still instructive to review Beer's thoughts on the computerization of industry. It is important to note that Beer was himself an enthusiast for computers. As early as 1956 at United Steel he had had installed one of the first computers in the world to be dedicated to management science, a Ferranti Pegasus (Harnden and Leonard 1994, 4). He was nevertheless a consistent critic of the way in which computers were being introduced more generally into industry and business. His argument was that "the first and great mistake" was that "people set out to automate the procedures and therefore the organisations they already knew. These themselves were frozen out of history and fixed by professionalism." Computers were, in other words, being used to automate existing clerical tasks while leaving the overall structure of the traditional organization untouched: "Companies have exchanged new lamps for old, and set them in the window as marks of progress. . . . We are using a powerful control instrument competent to reorganise the firm, its departments and functions, and encapsulating it in a received system geared to the quill pen." Instead, Beer argued, we should ask, "What should my enterprise be like, now that computers exist?" (Beer 1967, 214–17).

Beer was especially critical of the use of computers in business to automate and augment record keeping, and this gets us back to the ontological question. If the world is beyond our capacity to know it, and if, even worse, it continually changes, knowing the past is of limited utility. Our information processing should therefore be forward looking, as in the system 4 model of the VSM. "It is worth making a tremendous effort to burst through the barrier marked 'now,' and to make managers concern themselves with what *can* be managed—namely the future, however near—rather than peruse a record of what can be managed no longer—namely the past, however recent. We may learn from that past record of course, but we cannot influence it in retrospect. . . .

Look straight ahead down the motorway when you are driving flat out. Most enterprises are directed with the driver's eyes fixed on the rear-view mirror" (1981, 127, 199). Beer's idea in the VSM was thus that most of the information that one can collect on an organization is useless and can be discarded. This was what the filtering operations at the various levels did, keeping only anomalous signals for transmission to higher levels.

Seen from this angle, the object of the VSM was to reorganize the firm around the computer—to effect a transformation that was social as well as technological, to rearrange the human components as part of an adaptive technosocial system of information flows and transformations. Here, too, then, social relations and ontology hung together. And this contrast between the VSM and the traditional structure of the organization is another nice example of how ontology can make a difference in practice. Further aspects of this are apparent below.

The VSM in Practice

The VSM was a normative vision of the organization. Organizations had to look like the VSM if they were to survive and grow in time. The obvious implication of that would seem to be that they needed to be remade from the ground up to exemplify the VSM. Beer had one serious chance at that, which is reviewed in the next section. But Beer could hardly claim that all existing organizations were nonviable—some of them had been around for a long time, the Catholic Church, for example. He therefore made a more nuanced argument. Just like organisms, organizations could be more or less viable—some struggling to survive, others actually dying, others springing happily into the future: “The amoeba succeeded, the dinosaur failed, the coelacanth muddles along” (Beer 1981, 239). And the problem was that organizations had no way to discuss this temporal viability; they lacked any language or conceptual apparatus for it.

What organizations had instead was organization charts of hierarchical power relationships running downward from the board of directors through vertical chains of command devoted to production, accounting, marketing, and so on. Beer's claim was that such charts did not, and could not, represent how firms actually worked. They functioned, at most, as devices for apportioning blame when things went wrong.²⁵ Already, then, whether anyone recognized it or not, the VSM was a better description of how the firm really worked, and Beer's pitch was that the formal VSM could therefore function as a diagnostic tool (1981, 155). One could examine the firm, or any other

organization, and see just which bits of it corresponded to the five levels of the VSM, and one could examine the ways in which they were connected together. Certain aspects of the firm might thus be identified as especially deficient as compared to the VSM diagram and made the targets of therapeutic intervention. Beer claimed that an experienced VSM practitioner could often walk into a factory and identify the major problems within a day or two and that, once pointed out, management would recognize the veracity of the judgment—such problems having already been subconsciously recognized and papered over (Beer 1989a, 27). Of course, addressing the problems thus identified might take much longer—conceivably a period of years. “My guess would be that organizations cannot face up to more than a quarter of the reshaping that their long-term viability demands. This is of course the reason why so many enterprises are in a state of *continuous . . . reorganisation*” (Beer 1981, 239).

One simple example of what might be at stake here, of continuing practical and scholarly interest, concerns automation and the interlinkages between the systems 1 of figure 6.13. Beer noted that such linkages between different subsidiaries of a single company found no formal representation on the typical organization chart. But (Beer 1981, 107) “I have collected scores of examples of this. Sometimes, very often perhaps, the foremen in the related departments make it their business to keep in intimate touch. Maybe they walk across the road and drink tea together; maybe they telephone: ‘You’d better know, Charlie, that . . .’ In a few extreme case, it was not possible to discover how the messages were transmitted—but transmitted they certainly were.” Beer was quite happy with such informal channels of communication; his only concern was that Ashby’s law should be respected—that there should be enough variety at each end to cope with that at the other, and that there be enough bandwidth between them to mobilize and organize those varieties appropriately. Instead, Beer argued, the introduction of computers as information-processing devices often acted to sever such channels completely. Because the channels did not appear on the organization chart, they did not become automated; at the same time, their human conduits—the foremen, in this example—might be forbidden to step outside their own domains, or their positions eliminated entirely. “In the limiting case where the departmental outstation is fully automated, there is no possible way in which the social link can be maintained. Computers do not just happen to develop the trick of shouting to each other across the void, as human beings always do” (108). A technological transformation which appeared progressive on the surface might thus be regressive as seen from the perspective of the VSM.²⁶

Beer also claimed that many organizations were entirely lacking a system 2 (1981, 175), and in the absence of the “sympathetic” damping generated by the 1-2-3 system would thus always be prone to pathological competition and “oscillations” between their subsidiaries. More generally, Beer worried about the higher levels of the brain of the firm. Pieces of the organization which he felt should lie directly on the “command” axis were often found to be misplaced. This was true especially of parts of the organization that had grown up since World War II, including management accounting, production control (Beer’s first job in the steel industry), and operations research (his first love in management). These had no place on prewar organization charts and thus found themselves a position almost at random (Beer 1981, 82–83). OR groups, for example, might be found buried in subsidiaries and thus serving the overall organization asymmetrically—to the benefit of some subsidiary rather than the whole firm. The moral of the VSM was that there should be an OR group on the command axis itself, at level 3. Beer also argued that “in most firms System 4 is a fiasco” (153–54). Elements of system 4—the monitoring and planning organ at the base of the conscious brain—were usually to be found in any large organization, but they tended to be dispersed across the organization instead of grouped coherently together on the command axis. Certainly very few clubby operations rooms were to be found in industry in this period.

We need to remember that from 1970 onward Beer made his living primarily as an independent management consultant, and his writings on the VSM were integral to that. In 1989, he produced a list of consultancies he had been engaged in (Beer 1989a, 35):

Small industrial businesses in both production and retailing, such as an engineering concern and a bakery, come to mind; large industrial organizations such as the steel industry, textile manufacturers, ship-builders, the makers of consumer durables, paper manufacturers are also represented. Then there are the businesses that deal in information: publishing in general, insurance, banking. Transportation has figured: railways, ports and harbours, shipping lines. Education, and health (in several countries), the operations of cities, belong to studies of services. Finally comes government at all levels—from the city, to the province, to the state and the nation itself—and the international agencies: the VSM has been applied to several.

Obviously . . . these were not all major undertakings, nor is “success” claimed for massive change. On the other hand, none of these applications was an academic exercise. In every case we are talking about remunerated consultancy, and that is not a light matter. The activities did not necessarily last for very long

either, since speedy diagnosis is a major contribution of the whole approach. On the other hand, some of them have lasted for years. Undoubtedly the major use of this work to date was in Chile from 1971–1973.

Chile is next. Here I can just emphasize what is obvious from this list: Beer operated not only at the level of commercial companies; many other kinds of social organizations were likewise open to his interventions. We should also remember what was noted earlier—that by the 1980s the VSM had gained a significant following among management consultants and their academic counterparts, leading to the publication of at least one multiauthor book on the VSM (Espejo and Harnden 1989). The interested reader can look there for case studies written up by Beer and his followers, including Beer’s sixty-page account of his association over nine years with a mutual life assurance company (Beer 1989a), as well as for various methodological and substantive reflections on and extensions of the VSM. The VSM was never one of those great fads that seem to have periodically overtaken the world of management since the Second World War. Given its subtlety and complexity, to which I have done scant justice here, this does not seem surprising. But it has been at the heart of a significant movement.

Chile: Project Cybersyn

In Chile in the autumn of 1970 Salvador Allende became the world’s first democratically elected socialist president. The new government started nationalizing the banks and major companies operating within Chile, operating through an existing organization known as CORFO (Corporación de Fomento de la Producción). On 13 July 1971, the technical general manager of CORFO, one Fernando Flores, wrote to Beer (Beer 1981, 247): “This letter spoke of ‘the complete reorganization of the public sector of the economy,’ for which it appeared its author [Flores] would be primarily responsible. He had read my books, and had even worked with a SIGMA team ten years before. He went on to say that he was now ‘in a position from which it is possible to implement, on a national scale—at which cybernetic thinking becomes a necessity—scientific views on management and organization.’ He hoped that I would be interested. I was.” Beer’s commitment to the project became “total” (245), and he subsequently published a long account of the project’s evolution and termination, in five chapters added to the second edition of *Brain of the Firm* (Beer 1981, 241–399). Beer’s chapters are, as usual, very dense, and I can only attempt an overview of his account as a way of sketching in the

main features of what was undoubtedly the world's most striking cybernetic project.²⁷

Taking up Flores's invitation, Beer flew into the capital of Chile, Santiago, on 4 November 1971, remaining for eight days and returning to London on 13 November. In Santiago he met Flores and his collaborators, and together they made plans to implement the VSM at the level of the national economy. Beer had just completed the manuscript of *Brain of the Firm*; the Chileans studied it while he was there, and it became the basis for their vision of Chile's future. On 12 November Beer met President Allende himself and explained the VSM to him. When Beer drew the box for system 5 of the VSM diagram, he was thinking of it as representing the president, but Allende "threw himself back in his chair: 'at last,' he said, 'el pueblo'" (Beer 1981, 258)—the people. Beer was so impressed by this that he told the story often. Allende was apparently similarly impressed with Beer and the VSM: "The President says: Go ahead—fast" (257).

What did the plan sketched out on Beer's first visit look like—Project Cybersyn, for "cybernetic synergy," as it became known? Beer felt that speed was of the essence—"within a year . . . the foreign reserves would run out" (251)—so he aimed to begin by installing a cut-down version of the VSM by, astonishingly, 1 March 1972. This was less than four months after his first visit, and he promised to return on 13 March 1972. The initial plan aimed to achieve real-time (meaning daily) communications between system 1 productive activities at the level of individual factories, and a system 4 control room to be constructed in Santiago.

OR teams were charged "to construct a quantitative flow chart of activities within each factory that would highlight all important activities" (253). OR models would then be used in consultation with management—typically workers' committees, foreign managers having fled the country—to construct indices of performance analogous to those Beer had devised in the steel industry and reported upon in the 1953 OR paper discussed above (163).²⁸ "In practice, it turned out that some ten or a dozen indices were adequate to monitor the performance of every plant" (253). Among these was to be an index to measure morale as a ratio depending inversely on absenteeism (253).

The question of what to do with all the data thus generated, how to handle it, then arose. Ideally, every plant should have its own computer to "process whatever information turned out to be vital for that factory's management" (252)—this, thinking of each plant as a viable system in its own right. "But such computers did not exist in Chile, nor could the country afford to buy them. . . . Therefore it was necessary to use the computer power available in

Santiago: it consisted of an IBM 360/50 machine and a Burroughs 3500 machine” (252). The remaining technical problem was to connect plants all over the country up to Santiago. This was to be accomplished by requisitioning telex machines, augmented by microwave and radio links whenever necessary. “The plan allowed just four months for this to be accomplished (and it was)” (252). This national information system was known as Cybernet; the data it brought to Santiago were processed there “and examined for any kind of important signal. . . . If there were any sort of warning implied by the data, then an alerting signal would be sent back to the managers of the plant concerned” (253). Beer himself took two tasks back to England with him (256): “I had to originate a computer program capable of studying tens of thousands of indices a day, and of evaluating them for the importance of any crucial information which their movements implied. . . . I had done this kind of system building many times before. . . . Secondly, I should need to investigate prospects for a simulation system in the operations room that could accept the input of real-time data. This would be a completely novel development in operational research technique.” The basic blueprint and timetable for Cybersyn were thus set. Beer’s own account covers subsequent developments in some detail; we can review some of the main features.

As indicated above, the Cybernet national information system was indeed established by the deadline of March 1972. The first computer program mentioned in the above quotation took longer than hoped to construct, partly because of the incorporation of very new OR techniques in forecasting. A temporary version was indeed implemented in March 1972, but the permanent version only became operational in November that year. By that time “something like seventy percent of the socio-industrial economy was operating within this system, involving about four hundred enterprises” (Beer 1981, 262, 264).²⁹

These “Cyberstride” programs sat at the system 3 level, contributing to the homeostasis of the 1-2-3 economic assemblage while at the same time filtering data upward into the 3-4-5 system. A key element of the latter was a computer model of the Chilean economy and its national and global environment. This was to be the centerpiece of system 4 planning, intended to enable future projections according to different inputs and assumptions. This program was also Beer’s responsibility. Lacking time to design such a model afresh, Beer announced in a January 1972 report that he had decided “to make use of the immediately available DYNAMO compiler extensively developed by J. W. Forrester of MIT. I have directed three projects in the past using this compiler, and have found it a powerful and flexible tool” (266). Forrester’s



Figure 6.14. Operations room of Project Cybersyn. Source: Beer 1974a, 330, fig. 12.1.

work had grown by devious routes out of his World War II work at the Servo-mechanisms Laboratory at MIT and was just about to become famous, or notorious, with the publication of the Club of Rome's *Limits to Growth* report, which, on the basis of DYNAMO simulations, predicted an imminent collapse of the global economy and ecosystems.³⁰ Work in London and Chile under Chilean direction had developed a tentative version of the Checo (for Chilean economy) program by June 1972, and by September a better model was running. "I wanted to inject information *in real time* into the Checo program *via Cyberstride*. Thus any model of the economy, whether macro or micro, would find its base, and make its basic predictions, in terms of aggregations of low-level data—as has often been done. But Checo would be updated every day by the output from Systems 1-2-3, and would promptly rerun a ten-year simulation; and this has never been done. This was one of my fundamental solutions to the creation of an effective Three-Four homeostat; it remains so, but it remains a dream unfulfilled" (268). This continual updating was the way in which Checo simulations were foreseen as evolving in time, responsively to real-time input, thus exemplifying the performative epistemology of the VSM discussed in general terms in the previous section.

The system 4 operations room loomed ever larger as potentially the visible symbol, the icon, of Project Cybersyn (fig. 6.14). Detailed design was turned over to Gui Bonsiepe in Chile, from which emerged a plan for an octagonal room ten meters wide that would serve as an "information environment." Information on any aspect of the functioning of the economy at the desired level of recursion would be displayed visually on panels on the walls, including flashing warning signals that registered the algedonic "cries of pain" from lower levels, mentioned above, and an animated Checo simulation of the

Chilean economy that could be used to predict the effects over the next decade of decisions taken today. These days, computer graphics could handle what was envisaged with ease, but in the early 1970s in Chile the displays included hand-posted notes (of algedonic warnings), banks of projectors, and slides prepared in advance of meetings (showing quantified flow charts of production). The Checo display “certainly worked visually; but the computer drive behind it was experimental and fragmentary” (Beer 1974a, 329–32). The target date for completion of the control room was set as 9 October 1972; in fact, it was in “experimental working order” by 10 January 1973 (Beer 1981, 270).

Project Cybersyn evolved very quickly, but so did other developments (Beer 1981, 307):

As time wore on throughout 1972, Chile developed into a siege economy. How ironic it was that so many eyes were focussed with goodwill on the Chilean experiment in all parts of the world, while governments and other agencies, supposedly representing those liberal-minded observers, resisted its maturation with implacable hostility. The nation’s life support system was in a stranglehold, from financial credit to vital supplies; its metabolism was frustrated, from the withholding of spare parts to software and expertise; literally and metaphorically, the well-to-do were eating rather than investing their seed-corn—with encouragement from outside. Even more ironic, looking back, is the fact that every advance Allende made, every success in the eyes of the mass of the people (which brought with it more electoral support) made it less likely that the Chilean experiment would be allowed to continue—because it became more threatening to Western ideology.

Before Allende came to power, copper had been Chile’s major source of foreign exchange, and “we were to see the spectacle of the ‘phantom ship’ full of copper that traipsed around European ports looking for permission to unload” (307). Economic collapse was imminent, and Beer’s thought was to “search for novel and evolutionary activity whereby the Chilean economy might very rapidly enhance its foreign earnings” (308). His answer was indigenous crafts, wine, and fish, and in 1972 and 1973 he sought to mobilize his contacts in Europe to expand those markets—without success. There was nothing especially cybernetic about those efforts, but they do indicate Beer’s commitment to Allende’s Chile.

In 1973 the situation in Chile continued to worsen. In September 1973, the Cybersyn team received its last instruction from the president, which

was to move the control room into the presidential palace, La Moneda. “By the 11 September 1973, the plans were nearly ready. Instead La Moneda itself was reduced to a smoking ruin” (Beer 1974a, 332). Salvador Allende was dead, too, in the ruin: the Pinochet coup—Chile’s 9/11—brought a definitive end to the Chilean experiment with socialism, and with it went Cybersyn. Beer was in London at the time but had prepared for the end by devising three different codes in which to communicate with his collaborators and friends in Chile, who were, through their association with the Allende government, in very serious trouble. Beer did what he could to help them. On 8 November 1973, he wrote to von Foerster at the University of Illinois: “My dear Heinz, I think you know that I am doing everything possible to rescue my scientific colleagues (at the level of Team Heads) from Chile. It is going well—10 families. There is another problem. My main collaborator is held in a concentration camp, and is coming up for trial. There is a real risk that he will be shot, or sent down for life.”³¹ The collaborator in question was Fernando Flores, who had risen to become Chile’s minister of finance before the coup. Beer enclosed the draft of his personal statement to be read at Flores’s trial and urged von Foerster to send his own. In the event, Flores was imprisoned for three years, until Amnesty International helped to negotiate his release, when he moved to the United States, completed a PhD in Heideggerian philosophy, and became a highly successful management consultant.³²

The Politics of the VSM

THE PROBLEM IS FOR CYBERNETICS TO DISCOVER, AND TO MAKE ABUNDANTLY CLEAR TO THE WORLD, WHAT METASYSTEMS TRULY ARE, AND WHY THEY SHOULD NOT BE EQUATED WITH THE SUPRA-AUTHORITIES TO WHICH OUR ORGANIZATIONAL PARADIGMS DIRECT THEM. IT IS AN APPALLING [SIC] DIFFICULT JOB, BECAUSE IT IS SO VERY EASY TO CONDEMN THE WHOLE IDEA AS TOTALITARIAN. HENCE MY USE OF THE TERM: THE LIBERTY MACHINE. WE WANT ONE THAT ACTUALLY WORKS.

STAFFORD BEER, “THE LIBERTY MACHINE” (1975 [1970], 318)

Beer’s daughter Vanilla recalls that “Stafford and I generally ran Jesus and Marx together in an attempt to produce metanoyic possibilities,” so I turn now to Beer’s politics and its relation to his cybernetics; later sections will focus on his spiritual beliefs and practices.³³

As a schoolboy, Beer shared a bedroom with his brother, Ian, who recalled that Stafford “painted the whole wall . . . with extraordinary apparitions. In the centre of the wall was the original ‘Towering Inferno’—a huge skyscraper with flames all around the bottom licking their way up the tower.” Vanilla Beer adds that the picture was called *The Collapse of Capitalism*. In the late forties, Stafford fell out with his father, who pressured him into admitting that he had voted for the Labour Party in the recent election (Ian Beer, letter to Stafford’s family, 25 August 2002). Later in life, Beer sometimes described himself as “an old-fashioned Leftist” (Medina 2006) or even as “somewhat to the left of Marx,” though it would be a mistake to think of him within the conventional frame of British Marxism: “Stafford was fond of telling the story about Marx that had him saying ‘Thank God I’m not a Marxist.’ He didn’t usually describe himself in this context but Stafford had a great deal of admiration for Marx, especially his early writings on alienation. He wasn’t much of a fan of *Das Capital* mostly on the grounds of dull and repetitive.”³⁴

Little of this found its way into Beer’s early writings. Until 1970, his books, essays, and talks were largely couched in a technical idiom and addressed to a management readership. But in 1969 (Beer 1975, 3)

I had come to the end of the road in my latest job . . . and re-appraised the situation. What was the use of seeking another such job all safe and sound pensions all that from which haven to speak and write as I had done for years about the desperate need for drastic change and how to do it in a sick world? Not even ethical. How to begin? It was almost 1970. A decade opened its doors for business. There were speeches to be made already committed throughout that first year and I must see them through. What’s more these platforms gave me the opportunity if I could only seize it to collect my thoughts for a new life and to propound ARGUMENTS OF CHANGE.

This series of talks, with assorted explanatory material, was published in 1975 as *Platform for Change: A Message from Stafford Beer*. In 1973, just before the Pinochet coup, Beer continued to develop his thinking in public, this time in the Canadian Massey Lectures on CBC radio, which were published the next year as *Designing Freedom* (Beer 1974b). The focus of these works, and many to follow, was on liberty, freedom, and democracy. Marx is not mentioned in them, nor any of the classic Marxist concerns such as class struggle. Instead, Beer attempted a distinctly cybernetic analysis, which is what interests me most. Here we can explore another dimension of ontology in action: cybernetics as politics.

The distinctly cybernetic aspect of Beer's politics connected immediately to the ontology of unknowability. Other people, at any scale of social aggregation, are exceedingly complex systems that are neither ultimately graspable nor controllable through knowledge. And along with that observation goes, as I noted in chapter 2, a notion of respect for the other—as someone with whom we have to get along but whom we can never possibly know fully or control. And this was Beer's normative political principle: we should seek as little as practically possible to circumscribe the other's variety, and vice versa—this was the condition of freedom at which Beer thought politics should aim. This, in turn, translated into an explicit view of social relations. If the ontology of knowability sits easily with an image of hierarchical command and control, in which orders are transmitted unchanged from top to bottom, then Beer's notion of freedom entailed a symmetric notion of adaptive coupling between individuals or groups. In a process of reciprocal vetoing—also describable as mutual accommodation—the parties explore each other's variety and seek to find states of being acceptable to all. The ontological and practical resonances here among Beer and Bateson and Laing are obvious, though Beer was operating in the space of organizations rather than psychiatry.

Beer recognized, of course, that any form of social organization entailed some reduction in the freedom of its members, but he argued that one should seek to minimize that reduction. In reference to viable systems, his thought was that freedom was a condition of maximal "horizontal" variety at each of the quasi-autonomous levels, coupled with the minimum of "vertical" variety reduction between levels consistent with maintaining the integrity of the system itself. Hence the notion of "designing freedom": as Beer explained it, the VSM was a diagram of social relations and information flows and transformations that could serve to guarantee the most freedom possible within organized forms of life. As we need to discuss, that view did not go uncontested, but let me emphasize now two features of Beer's vision.

First, there are many absorbing books of political theory which go through immensely subtle arguments to arrive at the conclusion that we need more freedom, fuller democracy, or whatever—conclusions which many of us would accept without ever reading those books. Beer was not in that business. He took it for granted that freedom and democracy are good things. The characteristic of his work was that he was prepared to think through in some detail just how one might arrange people and information systems to make the world freer and more democratic than it is now. Beer's specific solutions to this problem might not have been beyond criticism, but at least he was prepared to think at that level and make suggestions. This is an unusual

enterprise, and I find it one of the most interesting and suggestive aspects of Beer's cybernetics. Second, we should note that, as already remarked, Beer's talks and writings did not foreground the usual substantive political variables of left-wing politics: class, gender, race. They foregrounded, instead, a generic or abstract topology in which the exercise of politics, substantively conceived, would be promoted in a way conducive to future adaptations. We should perhaps, then, think of Beer as engaging in a particular form of subpolitics rather than of politics as traditionally understood.

That said, Cybersyn was the only cybernetic project discussed in this book to be subjected to the political critique I mentioned in the opening chapters. I therefore want to examine the critique at some length, which will also help us get Beer's subpolitics into clearer focus and serve to introduce some more features of Cybersyn.

The Political Critique of Cybernetics

The early phases of Project Cybersyn were conducted without publicity, but public announcements were planned for early 1973. Beer's contribution to this was "Fanfare for Effective Freedom," delivered as the Richard Goodman Memorial Lecture at Brighton Polytechnic on 14 February 1973 (Beer1975b [1973]). The previous month, however, reports of Cybersyn had appeared in the British underground press and then in national newspapers and magazines (Beer 1981, 335), and the media response had proved hostile. The day after Beer's "Fanfare" speech, Joseph Hanlon wrote in the *New Scientist* that Beer "believes people must be managed from the top down—that real community control is too permissive. . . . The result is a tool that vastly increases the power at the top," and concluded with the remark that "many people . . . will think Beer the supertechocrat of them all" (Hanlon 1973a, 347; and see also Hanlon 1973b). Hanlon's article thus sketched out the critique of cybernetics discussed in chapter 2: cybernetics as the worst sort of science, devoted to making hierarchical control more effective.

Beer replied in a letter to the editor, describing Hanlon's report as a "hysterical verbal onslaught" and resenting "the implied charge of liar" (Beer 1973a). One H. R. J. Grosch (1973) from the U.S. National Bureau of Standards then joined in the exchange, explicitly calling Beer a liar: "It is absolutely not possible for Stafford Beer, Minister Flores or the Chilean government or industrial computer users to have since implemented what is described." Grosch further remarked that this was a good thing, since Cybersyn "well merits the horror expressed by Dr Joseph Hanlon. . . . I call the whole concept

beastly. It is a good thing for humanity, and for Chile in particular, that it is as yet only a bad dream.” Beer’s reply (1973b) stated that the Cybersyn project had indeed achieved what was claimed for it, that “perhaps it is intolerable to sit in Washington DC and to realise that someone else got there first—in a Marxist country, on a shoestring,” and that “as to the ‘horror’ of putting computers to work in the service of the people, I would sooner do it than calculate over-kill, spy on a citizen’s credit-worthiness, or teach children some brand of rectitude.”

The political critique of Cybersyn and the VSM was further elaborated and dogged Beer over the years, and I want now to review its overall form, rather than the details, and how one might respond to it. The critique is fairly straightforward, so I shall present it largely in my own words.³⁵

In 1974, Beer said of Cybersyn that it “aimed to acquire the benefits of cybernetic synergy for the whole of industry, while devolving power to the workers at the same time” (Beer 1974a, 322), and there is no doubt of his good intentions. His critics felt that he was deluding himself, however, and Hanlon’s description of Beer as a “supertechocrat” presaged what was to follow. I find it useful to split the critique into four parts.

1. The VSM undoubtedly was a technocratic approach to organization, inasmuch as it was an invention of technical experts which accorded technical experts key positions—on the brain stem of the organization at levels 3 and 4. No one had asked the Chilean workers what sort of a subpolitical arrangement they would like. Nor, I believe, did Beer ever envisage the basic form of the VSM changing and adapting once it had been implemented in Chile. There is not a lot one can say in reply to this, except to note that, on the one hand, the fixity of the overall form of the VSM can be seen as a noncybernetic aspect of Beer’s cybernetic management. As ontology in action, the critics seized here on a nonexemplary feature of Beer’s work. But we might note, too, that expert solutions are not necessarily bad. Beer’s argument always was that cyberneticians were the experts in the difficult and unfamiliar area of adaptation, and that they had a responsibility to put their expertise to use (see, e.g., Beer 1975 [1970], 320–21). To say the least, Cybersyn was a new and imaginative arrangement of socioinformatic relations of production, which might, in principle—if the Pinochet coup had not happened—have proved to have increased the freedom of all concerned. Beyond this, though, the critics found more specific causes for concern within the structure of the VSM itself.

2. Another thread of the critique had to do with the algedonic signals that passed upward unfiltered to higher levels of the VSM. Beer spoke of these as “cries for help” or “cries of pain.” They were intended to indicate that

problems had arisen at the system 1 level which could not be addressed there, and which therefore needed assistance from higher levels in their resolution. Beer assumed that the upper levels of the system would adopt a benevolent stance relative to the lower ones and would seek to provide genuine assistance on the receipt of an algedonic signal. Critics pointed out instead that such signals could also constitute a surveillance system that would sooner or later (not necessarily under Allende) be used against the lower levels. A profit-maximizing higher management might readily translate too many algedonic warnings into a rationale not for assistance with problems but for plant closures. Again, it is hard to spring to Beer's defense. He might have replied that to think this way is to denature and degrade the biological model behind the VSM. Brains do not jettison arms and legs every time we get pins and needles, but the obvious reply would be that this simply brings into question Beer's biological model for social organizations. For Beer, this was a normative aspect of the model, but no one could guarantee that higher management would accede to this.

A more detailed version of this same critique acknowledged that there must be some vertical communication within organizations but questioned the automaticity of "cries for help." In the VSM, this was simply a matter of statistical filtration of data. If production indices remained anomalous after an agreed period of time, the algedonic signal automatically passed on to the next level. Werner Ulrich (1981, 51–52) pointed out that in a less automated system there would be a place for management learning—managers come to recognize patterns in the signals arriving at their level and thus to discriminate between which needed to be passed on and which did not—thus protecting the lower levels to some extent from vindictiveness above. I do not know whether Beer ever addressed this point, but, again, the VSM was not exemplary of the cybernetic ontology in action to just the degree to which this automaticity was a fixed part of the VSM.

3. Following the lines set down by Hanlon, the VSM's critics asserted that the VSM prescribed a "top-down" mode of organizational control: management or government gave orders that the workers were then expected simply to implement. Cybersyn "has some kind of built-in executive power. . . . Its strongly hierarchical organisation and its concept of 'autonomy' one-sidedly serve the top decision maker, the government" (Ulrich 1981, 52, 54). As before, there is something to this critique, but it is worth taking it slowly. Though the critics seem to have read Cybersyn as implementing a classic "command and control" form of organization, with a unilinear flow of orders descending

from on high, in this they were wrong. Beer did not think of viable systems in that way. This was precisely the significance of the adaptive couplings that pervaded the VSM, especially the couplings between the various levels. As discussed earlier, these were modelled on the reciprocal vetoing in Ashby's multihomeostat setups and implied that the parties at different levels had to cast around for mutually agreeable initiatives and plans, precisely *not* the traditional command-and-control mode. These adaptive couplings were the most definitively cybernetic component of the VSM, and it is significant that the critics failed to get to grips with them or even to recognize their distinctive character. Beer often complained that outsiders erred in a similar way concerning all sorts of cybernetic machines and contrivances, utterly failing to grasp their adaptive aspects, and this seems to have been the case here. If ontology makes a difference, then that difference eluded the VSM's critics. But more needs to be said.

Cybersyn was, on one occasion, operated in both a surveillance and a command-and-control mode. This was the time of the *gremio* strike in October 1972, a "CIA-instigated trucker's strike" in Chile (Ulrich 1981, 54n; Beer 2004 [2001], 860) which threatened to halt flows of goods around the country.³⁶ The Cybernet information system was then switched temporarily to monitoring shortages around the country and figuring out how to use the transportation available to overcome them. Beer was very pleased that this approach worked and that the strike was defeated (Beer 1981, 312–15), but there was no homeostatic give-and-take involved in this episode in negotiating plans between different levels, and it serves to show just how readily the organic quality of the VSM could be conjured away, and, indeed, this possibility seems to have appealed to Allende's enemies.³⁷ "At the end of July [1973] . . . several strange messages reached me. . . . They were coming from the political opposition. It seemed that this [Cybersyn] was the best project undertaken under Allende's aegis, and that his (self-assumed) successor would continue it in his own way. This would not, of course, involve any 'nonsense' about worker participation. . . . I found these overtures obnoxious; but our strategies were well prepared" (Beer 1981, 345). The strategies, I believe, were intended to render Cybersyn useless in the event of a coup, but three comments are called for. First, in its genuinely cybernetic aspect—the adaptive couplings between levels—the VSM did serve to undo hierarchies of command and control. Second, these adaptive couplings could easily be "switched off" and replaced by asymmetric ones. It is fair to say, then, that the VSM was hardly a potent bulwark against the institutional arrangements that Beer wanted to obviate. This,

too, was much on his critics' minds. But third, as Beer might have advised, we should be concerned here with the future more than the past. Even if key components of the VSM were readily erasable, the VSM remains interesting as a model for a democratic subpolitics.

4. We can return to the question of goals. In chapters 3 and 4 we looked largely at systems with fixed goals. Ashby's homeostats adapted open-endedly, but so as to keep their essential variables within given limits. According to Beer, the quasi-organic viable system likewise had goals that patterned its adaptation. But, unlike Ashby, Beer was not attempting to construct models of the adaptive brain, and he therefore did not have to take a sharp position on what the goals of a viable system are. I said earlier that one could think of the profitability of an enterprise as the sort of thing at issue, but actually Beer had something different and more interesting in mind, which we can get at via the critique of the VSM. At the heart of Werner Ulrich's (1981, 35) long critique, for example, is a contrast between "purposive" and "purposeful" systems, which relates to a more familiar distinction between means and ends: a "purposive" system is a means to some extrinsically specified end, while a "purposeful" one can deliberate on its own ends. Ulrich criticized the VSM as purposive, and at one level this is correct. Beer was keen not to try to build any substantive goals beyond adaptability into the VSM; this is an aspect of what was entailed in my earlier description of the VSM as a form of subpolitics.

Ulrich, however, went on from this observation to claim that because the VSM had no substantive goals, then whatever goals a system came to manifest would have to be supplied in a top-down fashion, from systems 4 and 5 of the model—we are back to technocracy from a different angle. But here there are some complications worth discussing. One reply would be that Beer was working for a democratically elected government responsive to "the will of the people," but that is an observation about the specific context of Cybersyn rather than an intrinsic feature of the VSM in general. Another reply would go along the lines indicated above: that the adaptive couplings between the VSM's levels are reciprocally adaptive, not one-way. But here, still, some asymmetry remained in the VSM. Beer does not seem to have envisaged the formulation of new plans and goals from below; the higher levels of management and government do seem to have held the advantage here in his thinking (though this assertion will be qualified below when we come to his work on "syntegration," which indeed focused on inclusive processes of goal formation). Nevertheless, Project Cybersyn, as it evolved, did at least try to close the

loop between government initiatives and their popular reception in various ways, and I want to examine just one of these.

On Goals

In March 1972 . . . we addressed the basic issue of the organization of the state that is not economic but societal. . . . I wrote a second paper about a project to examine:

“the systems dynamics
of the interaction
between government and people
in the light of newly available technology
such as TV
and discoveries in the realm
of psycho-cybernetics”

(Beer 1981, 278)

There were, of course, many channels by which the Chilean government could communicate with the Chilean population at large and vice versa. But the reference to TV immediately suggests an asymmetry. Governments could transmit information over the television in great detail and length—a high-variety channel, in the language of information theory. The people, in contrast, could not reply via the TV at all—an exceedingly low-variety channel. Of course, the people could communicate via other channels, such as forming political parties and voting in elections, but Beer felt that it was necessary to do something to increase the information flow from people to government if a homeostatic equilibrium was to be achieved. He also, as usual, felt that the channel from people to government should be a real-time one, so that the latter could react to how the former felt today rather than last week or last month or last year.³⁸ The solution Beer proposed, novel and endearing, is shown in figure 6.15. The aim here was to supplement the economic algedonic feedback of the VSM with social feedback. TV viewers, for example, would be provided with very simple “algedonic meters” of the form shown in the lower left of figure 6.15. These would be simple semicircular devices in which a partition could be rotated clockwise (toward “happy”) or counterclockwise (“unhappy”) in response to whatever was happening before them—a televised political speech, say. Some simple wiring arrangements would aggregate





Figure 6.15. Feedback from the people. Source: S. Beer, *Brain of the Firm*, 2nd ed. (New York: Wiley, 1981), 281, fig. 45.

these algedonic signals (the precise arrangements being left open in the initial proposal) and transmit them for display in real time on the TV screen. In this way, the politicians would get instantaneous feedback on their proposals or arguments. And—this is the clever bit—the viewers could also see how the politicians would react to the feedback, and so on in a cascade of feedbacks between the TV studio and its audience (Beer 1981, 285). In effect, some

channel, however crude, would thus be opened for mass debate—or, better, a dance of agency—with the government. Again, policy making could thus emerge in real-time interaction.

Like many of the cybernetic devices we have been exploring, these algedonic meters of Beer's were at once serious and amusing, and even startling in spanning the gap between the two. Their origins, I would guess, lay in the clapometers and swingometers of the BBC's popular music TV shows and election reporting.³⁹ An interesting feature is that they were truly algedonic in being able to register pleasure as well as pain, unlike the algedonic signals in the basic VSM, which were regarded as warnings that something was wrong. Though Beer initially conceived their use in mass communication, they could obviously be deployed in much more limited contexts—in small meetings, say, where some planning group reported to its constituents, or at factory gates as feedback from the workers to management.

Beer's son Simon, an electrical engineer, built a prototype system “of ten algedonic meters, linked by a single wire in a loop through a large summation meter” (Beer 1981, 284), and took it out to Chile, where experiments were done on its use with a group of fifteen friends. These friends, however, “rapidly learned how to rig the system. They joined in plots to ‘throw’ the lecturer by alternating positive and negative responses, for instance” (286). The algedonic meter was, in this instance, too much fun. And one can easily imagine less amusing forms of rigging—the political party instructing its supporters to slam the indicator to the left whatever an opponent said—or even arguments about whether “unhappy” should be at the left or the right. This form of feedback was thus never introduced in Chile, leaving Beer to reflect that its design was a tricky problem and that more cybernetic research was needed. Nevertheless, it is interesting to stay with them just a little longer.

Beer contrasted his algedometers favorably with another and more familiar form of quasi-real-time feedback from the people to government: questionnaires and opinion polls (Beer 1974a, 334–38). From Beer's perspective, the great thing about the algedonic meters was that they were inarticulate, wordless. They measured “happiness,” but the nature of happiness and its causes were left undefined. They simply indicated a positive or negative response on some undefined scale. Beer's enthusiasm for this mode of communication had to do with his intense interest in performance and his associated suspicion of representational knowledge. The trouble with opinion polls, Beer argued, is that the domain of inquiry is circumscribed by the questions asked (themselves framed by politicians, journalists, academics, and so on) and lacks variety. Articulated questions might therefore be able to determine how people

feel about specific government policies, but they can never find out whether people's real concerns lie entirely elsewhere. Polls can never contribute, then, to the emergence of real novelty in real-time politics, only to a fine-tuning of the status quo. In contrast, the algedonic meters constituted an open invitation to genuine experiment. If a politician or journalist were to float some wild idea and the integrated meter reading went from lethargically neutral to wildly positive, there would be reason to think that some genuine but hitherto unthought-of social desire had been tapped.

And here we can return to Ulrich's critique of the VSM as purposive rather than purposeful. Though Beer did not try to build into the VSM any substantive goals, he did try to think through the ways in which the system could articulate its own goals, in practice, in a nonhierarchical fashion. We can think of the algedonic meters as expanding the VSM as a subpolitical diagram of social relations and information flows in such a way as to enable any organization to become purposeful, rather than purposive, on its own terms. Ulrich is wrong here about the VSM, at least in principle, though, as above, practical concerns are not hard to find: it would have been less difficult for General Pinochet and his friends to eliminate algedonic meters than, say, rifles in the hands of the workers.

One last thought about the algedonic meters. What did they measure? At the individual level, an unanalyzed variable called "happiness." But for the aggregated, social, level Beer coined a new term—*eudemony*, social well-being (Beer 1974a, 336). Again he had no positive characterization of eudemony, but it is important that he emphasized that it is not any of the usual macrovariables considered by politicians and economists. Eudemony is not, or not necessarily, to be equated with GNP per capita, say, or life expectancy (Beer 1974a, 333). Eudemony is something to be *explored* in the adaptive performance of a viable social system, and, obviously, Beer's algedonic meters were an integral part of that. This thought is perhaps the most radical aspect of Beer's subpolitics: the idea that social systems might continually find out what their collective ends are, rather than, indeed, having those ends prescribed from above (the wonders of the free market, for example). And this remark gets us back to the general question of cybernetics and goals. Beer's cybernetics, unlike that of Walter and Ashby, did not enshrine any idea of fixed goals around which adaptation was structured. Goals, instead, could *become* in Beer's (and Pask's) cybernetics. As ontological theater, then, the VSM staged a vision of open-ended becoming that went an important step beyond that of the first-generation cyberneticians. Beer had not, of course, solved the prob-

lem of building a machine that could mimic the human facility of formulating goals; his systems could be adaptive at the level of goal formation precisely because they contained human beings within themselves.

— — — — —

Where does this leave us? After reviewing the critiques of the VSM and Project Cybersyn, I continue to think that we *can* see the VSM as enshrining a very interesting approach to what I have called subpolitics. The VSM offers a considered topology of social locations and relations, information flows and transformations that, to a considerable degree, promises a dispersal of autonomy throughout social organizations. The key elements of the VSM, from this perspective, are the adaptive, homeostat-like couplings between the various levels of the VSM, and the algedonic signals that travel back up the system. Like Beer's earlier experimentation with biological computing, his work on the VSM seems original and singular to me. It is hard to think of any equivalents in more conventional approaches to political theory and practice. And for this reason I am inclined to point to the VSM as another item on my list of striking examples of the cybernetic ontology in action, in politics and management. Here again we can see that the cybernetic ontology of unknowability made a difference.

Turning to the critics, it is significant that they seemed unable ever quite to get the VSM into focus. Beer's overall cybernetic aim, to bolster the adaptability of organizations, was never, as far as I can make out, mentioned by the critics; neither was the key cybernetic idea of adaptive coupling between levels. Instead, the critics focused on a cybernetically denatured version of the VSM, a version from which the distinctively cybernetic elements had been removed, turning it into a nightmare of command and control. The critics mapped the VSM onto a distinctively modern space in which it did not belong, and they found it wanting there. This inability to contemplate the thing in itself I take to be further evidence that ontology makes a difference.⁴⁰

Having said that, I have also recognized that the critics' concerns about the VSM were not empty. It does seem clear that systems like that envisaged in Project Cybersyn could be readily stripped down in practice and turned into rather effective systems of command, control, and surveillance, the very opposite of what both Beer and the critics aimed at. But as I have said before, the object of this book is not to resurrect any specific cybernetic project, including Cybersyn. It is to exhibit and examine a whole range of such projects—as a demonstration of their possibility and their difference from more

conventional projects in cognate areas, and as models for the future. A future cybernetic politics that followed Beer's lead into subpolitics might well want to bear in mind the democratic fragility of the VSM—while contemplating algedonic meters as, shall we say, a desperate but entertaining attempt to open up a politically deadening status quo.

Pinochet's coup in Chile was not the end of Beer's involvement in politics at the governmental level, especially in Central and South America. He went on to consult for the governments of Mexico, Venezuela, and Uruguay, as well as, in other directions from the United States, Canada, India, and Israel (Beer 1990a, 318–21), and “bits and pieces of the holistic approach have been adopted in various other countries, but by definition they lack cohesion” (Beer 2004 [2001], 861).⁴¹ I will not pursue that line of development further here; instead, I want to explore Beer's cybernetic politics from another angle.

The Politics of Interacting Systems

LAST MONTH [SEPTEMBER 2001], THE TRAGIC EVENTS IN NEW YORK, CYBERNETICALLY INTERPRETED, LOOK QUITE DIFFERENT FROM THE INTERPRETATION SUPPLIED BY WORLD LEADERS—AND THEREFORE THE STRATEGIES NOW PURSUED ARE QUITE MISTAKEN IN CYBERNETIC EYES. . . . ATTEMPTS TO GUARD AGAINST AN INFINITE NUMBER OF INEXPLICIT THREATS DO NOT HAVE REQUISITE VARIETY.

STAFFORD BEER, “WHAT IS CYBERNETICS?” (2004 [2001], 861–62)

So far we have focused on the internal politics of the VSM—on social arrangements *within* a viable organization. Here, the organization's environment was conceptualized in rather amorphous terms, simply as that to which the organization needed to adapt. As we saw in the previous chapter, in the 1950s Ross Ashby was led to think more specifically about environments that themselves contained adaptive systems and thus about interacting populations of adaptive systems, including the possibility of war between them. Beer's experiences in Chile and of the subversion of the Allende regime by outside states, especially the United States, led him to reflect along similar lines from the 1970s onward. These reflections on the interrelations of distinct systems, usually conceived as nation-states, themselves warrant a short review.

Beer's basic understanding of international relations followed directly

from his cybernetic ontology along lines already indicated. Nation-states are obvious examples of exceedingly complex systems, always in flux and never fully knowable. Their interaction should thus take the usual form of reciprocal vetoing or mutual accommodation, exploring, respecting, and taking account of the revealed variety of the other. Beer found little evidence for such symmetric interaction in the contemporary world, and thus, much of his analysis focused on what happens when it is absent. At the other pole from homeostat-like explorations lies the attempt to dominate and control the other, and Beer's argument was that this must fail. According to Ashby's law, only variety (on one side) can control variety (on the other). Any attempt simply to pin down and fix the other—to make it conform to some given political design—is therefore doomed to make things worse. The imposition of fixed structures simply squeezes variety into other channels and manifestations which, more or less by definition, themselves subvert any imposed order.

Beer's general analysis of macropolitics was thus, throughout his career, a pessimistic one: conventional politics is bereft of cybernetic insight and thus continually exacerbates crises at all levels. This rhetoric of crisis is a resounding refrain from his earliest writings to his last. In Beer's first book, the crisis is one of the West in general (the only instance of Cold War rhetoric that I have found in his writing) and of British industry in particular (Beer 1959, ix): "The signs are frankly bad. . . . The index of industrial production has not moved up for four years. We desperately need some radical new advance, something qualitatively different from all our other efforts, something which exploits the maturity and experience of our culture. A candidate is the science of control. Cybernetic research could be driven ahead for little enough expenditure compared with rocketry, for example. And if we do not do it, someone else will." In his later and more political writings, the crisis was often said to be one of the environment and of the conditions of life in the third world, as well as the more usual sense of political crisis: a socialist government in Chile as a crisis for the Americans and British being a prime example.⁴²

When I first encountered this language of crisis in Beer's writing, I tended to ignore it. It seemed self-serving and dated. On the one hand, the rhetorical function of "crisis" was so obviously to motivate a need for cybernetics. On the other, we all used to talk like that in the 1960s, but, in fact, the world has not come to an end since then. As it happens, though, while I have been writing about Beer, his stories have started to seem very relevant and, indeed, prescient. Everything that has happened since those planes flew into the World Trade Center and the Pentagon speaks of an American attempt (abetted by




Figure 6.16. The cybernetics of crisis. Source: S. Beer, *Brain of the Firm*, 2nd ed. (New York: Wiley, 1981), 354, fig. 48.S.

the British) at command and control on a global scale, seeking to freeze the world, to stop its displaying any variety at all—running from endless “security” checks and imprisonment without trial to the invasions of Afghanistan and Iraq. And the aftermath of the American invasion of Iraq—what we have been taught to call “the insurgency,” the killing, destruction, mayhem, and torture in the name of “democracy”—speaks vividly of the negative consequences of seeking to repress variety.

Little more can be said here—this book is not a treatise on recent world history—but I do want to note that Beer’s “cybernetics of crisis” included an analysis of how crises like the present one can arise. Again, Beer’s focus was on transformative flows of information. Figure 6.16 is his basic diagram for considering such processes: the hatched area denotes a crisis affecting three different interest groups, which might be nation-states, A, B, and C. The details are less important than Beer’s general analysis of the information flow from the crisis region into A (“sensory input”) and the return action of A on the crisis (“motor output”). What Beer emphasized was that such information flows necessarily impoverish variety, and that in a systematic way. His argument was that *representations* of crises are inevitably filtered through low-variety conceptual models, models through which governments interpret crises to themselves and the media interpret them to the public. These models then feed into a low variety of potential actions which return to intensify the

variety of the crisis along axes that are unrepresentable in the models, and so on around the loop.

Let me close this section with three comments. First, we can note that this last discussion of the role of models in the production of crises is of a piece with Beer's general suspicion of articulated knowledge and representation. Models might be useful in performance, as in the VSM, but they can also interpose themselves between us and the world of performances, blocking relevant variety (hence the significance of the inarticulacy of Beer's algedonic meters, for example). Second, Beer died before the invasion of Iraq; the above thoughts on that are mine, not his. But, again, I am struck now not by any self-serving quality of his rhetoric, but by the prescience of his analysis. The highly simplified story of information flows and variety reduction that I just rehearsed illuminates how global politics could have collapsed so quickly into one-bit discriminations (Beer 1993a, 33) between "us" and "them," the goodies and the baddies; how it could have been that a majority of the American population could believe there was some connection between Al Qaeda and Iraq prior to the invasion and in the existence of what we were taught to call "weapons of mass destruction"; how it is that the American public and, perhaps, their government could have expected the invaders to be greeted with flowers and kisses rather than car bombs; and (turning back to the question of controlling variety) why mayhem should have been expected instead. Of course, third, one does not have to be Stafford Beer or a cybernetician to be critical of the war on terror, a "war" in which, "allies are expected to go into battle against an abstract noun, and to assault any nation unwilling to mobilize in such folly" (S. Beer 2001, 862–63). What interests me, though, is the generality of Beer's cybernetic analysis. We all know how to generate simplistic stories of heroes and villains, and much of the political talk of the early twenty-first century takes that form. Take your pick of the goodies and baddies—Saddam Hussein and Osama bin Laden or George W. Bush and the neocons. Such reversible stories will no doubt always be with us. Beer's analysis, instead, did not focus on the particulars of any one crisis. He actually began the most extended exposition of his analysis by mentioning the British abdication crisis of 1936, arguments over Indian independence from Britain in 1946, and the Suez crisis of 1956 (Beer 1981, 352–53). His analysis did not hinge on the question of whether George W. Bush was evil or stupid; his argument was that something was and is wrong at the higher level of large-scale systems and their modes of interaction that persistently produces and intensifies rather than resolves global crises. I take the novelty of this style of analysis to be another example of the ways in which ontology makes a difference.

Team Syntegrity

HOW SHALL WE EVER CONCEIVE

HOWEVER EXPRESS

A NEW IDEA

IF WE ARE BOUND BY THE CATEGORIZATION

THAT DELIVERED OUR PROBLEM TO US

IN THE FIRST PLACE

?

STAFFORD BEER, *BEYOND DISPUTE* (1994B, 8)

From the time of Project Cybersyn onward, the VSM was the centerpiece of Beer's management consultancy. In parallel to the VSM, however, he also developed a rather different approach to organizations that he called "team syntegrity." This grew from the 1950s onward, "flared into considerable activity 20 years ago, and occupied me throughout 1990 in a series of five major experiments" (Beer 1994b, 4). In the 1990s also, the conduct of "syntegrations" became partly a commercial business for Beer and his friends, associates, and followers.⁴³ Beer only published one book on syntegrity, *Beyond Dispute* (1994b), as distinct from three on the VSM, but he and his collaborators developed and reflected upon syntegration in considerable detail.⁴⁴ I am not going to attempt to do justice to that here. My aim is to sketch out the basic form of the approach, to connect it to the cybernetic ontology, and, continuing the above discussion, to examine it as a form of micro-sub-politics.⁴⁵

Put very crudely, the substance of team syntegrity was (and is) an evolving format or protocol for holding a meeting, a rather elaborate meeting called a "syntegration," and we can explore this format in stages. First, there are the connected questions of what the meeting is about and who should come to it. On the latter, Beer offered no prescriptions. The idea was that syntegration was a process focused on some topic of interest to its participants. His model for thinking about this was a group of friends who met regularly in a bar and found themselves returning to some topic, perhaps current politics, but an early example in the development of the technique involved members of the British Operational Research Society seeking to redesign the society's constitution in 1970, and the first experiment in 1990 involved a group of friends and friends of friends thinking about world governance (Beer 1994b, 9, 35). The participants were, then, characterized by their common concern and interest in whatever the syntegration was about. Beer called such a group an

“infoset,” and, for reasons that will become clear, the basic form of an infoset would comprise thirty people.⁴⁶

But just how should the topic of such a meeting be defined? This was a matter of pressing concern for Beer, a concern that ran along much the same lines as his critique of opinion polls mentioned earlier. The usual way of structuring such a meeting would be to distribute in advance an agenda listing specific topics for discussion and action. Beer’s point was that such an agenda prefigures its outcome within lines that can already be foreseen, and “anything truly novel has two minutes as Any Other Business” (Beer 1994b, 9). His idea, therefore, was that the first element of a syntegegration should itself be the construction by the infoset in real time of a set of relatively specific topics for discussion. In the mature form of syntegegration this entailed a fairly complicated protocol extending over some hours, but, in essence, the procedure was this: Knowing the general topic of the meeting—world governance, say—each participant was asked to write down at least one brief statement of importance (SI) relevant to the topic, aiming to encourage original discussion of some aspect of the overall focus of concern. These statements would then be publically displayed to all of the participants, who would wander around, discussing whichever SIs interested them with others, elaborating them, criticizing them, or whatever (all this, and what follows, with the aid of experienced “facilitators”). Finally, after a prescribed length of time, the participants would vote for the developed SIs they considered of most importance, and the top twelve SIs would be chosen as the focus for the remainder of the meeting (27). In this way, something like a specific agenda would be constructed, *not* as given in advance but as emergent itself in the process of the meeting.

Given a set of thirty participants and twelve SIs, what happens next? In a short but complicated process, participants are each assigned to a pair of SIs, respecting, as much as possible, their preferences. Then the process of syntegegration proper begins, and things get complicated to explain. How do you organize the discussion of twelve topics by thirty people? A completely unstructured agora-like situation is imaginable, but experience dictates that it would get nowhere. One might try to structure the meeting by, say, ranking individuals or topics in terms of priority, but this would return to Beer’s critique of agendas, one step down the line. Inspired by Buckminster Fuller’s geodesic domes (Beer 1994b, 12–14), the solution that Beer arrived at was to structure discussions in the form of a geometric figure, the icosahedron (fig. 6.17).⁴⁷

An icosahedron has thirty edges and twelve vertices, and hence the appearance of these numbers above. Each of the twelve topics is assigned to




Figure 6.17. The syntegration icosahedron. Source: S. Beer, *Beyond Dispute: The Invention of Team Syntegrity* (New York: Wiley, 1994), 338, fig. S6.2.

one of the twelve vertices of an imaginary icosahedron; each participant is imagined to be placed on an edge and engages in discussions of the two topics assigned to the vertices at the end of his or her edge. In turn, this implies that each participant is a member of two discussion groups of five people, each associated with the five edges that meet at any vertex (plus some additional niceties, including the participation of “critics” from disconnected edges, which I will not go into). These groups then meet repeatedly (three or more times) over a period of days for discussions that take off from SIs at their vertex, adding to, refining, and elaborating these statements in the course of their interactions. These discussions cannot all take place at once—one cannot be a member of two groups discussing two topics simultaneously—so participants alternate in time between their topics. And, according to Beer, the effect of this is that discussions *reverberate* around the icosahedron. On the first occasion, the discussion of any topic has a *sui generis* quality defined by the interaction of the five people concerned. But by the second iteration, their positions have each been inflected by different discussions at the other end of

their edges, themselves inflected by discussions at one further remove. And by the third iteration these inflections are traveling around the geometrically closed figure, and there is the possibility that an earlier contribution returns “to hit its progenitors in the back of the neck” (Beer 1994b, 13). This is what Beer meant by reverberation: ideas travel around the icosahedron in all directions, being transformed and becoming progressively less the property of any individual and more that of the infoset as a whole. At the end of the process, each vertex has arrived at a final statement of importance (FSI), and these FSIs are the collective product of the syntegration (Beer 1994b, 32–33).

Thus, in outline, the form of the syntegration process, and to put a little flesh on what the products of such a process can look like, we can look briefly at a syntegration held in Toronto in late 1990 (Beer 1994b, chap. 6). “The group who came together were recruited mainly by word of mouth. . . . Thus the Infoset was assembled in an unusually arbitrary way: we may call it such a unity only because of its members all being drawn to the heading on the poster: ‘What Kind of Future do You Want?’” (87). The first three of the SIs constructed at the start of the syntegration were: ‘God is a verb not a noun,’ ‘Each child spontaneously desires to develop responsibilities commensurate with its abilities,’ and ‘Censorship is a personal issue.’” In Beer’s précis, the first three of the FSIs, the products of the syntegration, were (97–98)

1. *Local Empowerment*: the need to push decision making downwards, especially in the case of abolishing nuclear war.
2. *Law and Government*: the move from ownership to stewardship, control to guardianship, competition to cooperation, winners and losers to winners alone.
3. *How to Make World Peace*: sovereign individuals acknowledge and accept the responsibility of a (human) world social contract, towards environmental protection, security, and evolution of the planet.

What can we say about this example? First, it shows that syntegration can be a genuinely dynamic and open-ended process: the SIs and FSIs were in no sense contained in the original topic; they evidently emerged in the syntegration itself. But what about the statements of importance themselves? They hardly come as singular revelations, at least to scholars interested in such matters, but, as Beer put it, “it could not be claimed that the FSIs . . . embodied major new discoveries, although they may have done for some present. . . . [But] they are hardly banal” (97).

This and similar experiences in other syntegrations led Beer to remark that “amongst many others I have often claimed that in planning it is the process

not the product that counts” (Beer 1994b, 97), and *Beyond Dispute* documents in various ways the fact the participants in syntegegrations generally found them enjoyable and productive. The phrase “consciousness-raising” comes to mind, and we will see below that such phrases had a very literal meaning for Beer—his idea was that a genuine group consciousness could arise from the reverberations of syntegegration. Let me close this section, however, with some general reflections on the syntegegrity approach to decision making, with the critique of Beer’s VSM in mind—“a topic of which” Beer declared himself in 1990 to be “heartily sick” (Beer 1990b, 124).

Like the VSM, syntegegrity can be described as a form of subpolitics, this time at a microscale of small groups. Like the VSM, syntegegrity had at its heart a diagram, though now a geometric figure rather than a neurophysiological chart. Again like the VSM, syntegegrity, Beer argued, staged an inherently democratic organization, an arrangement of people in which concrete, substantive, political programs could be democratically worked out—indeed, he often referred to syntegegration as “complete,” idealized,” and “perfect democracy” (Beer 1994b, 12; 1990b, 122). And, unlike the VSM, in this case it is hard to dispute Beer’s description. Beer’s critics were right that the VSM could easily be converted to a system of surveillance, command, and control, but it is hard to contrive such fears about syntegegrity. By construction, there are no privileged positions in the syntegegration icosahedron, and there is no evident way any individual could control the syntegegration process (short of wrecking it beyond recognition).

Once more, too, we can see how ontology and subpolitics are bound up together in syntegegrity. As exceedingly complex systems, the participants cannot know in advance what topics will emerge from the syntegegration process, and this emergence is orchestrated as a process of multihomeostat-like reciprocal vetoing and creative mutual accommodation between participants and statements of importance. Of course, there is some prestructuring entailed in the assembly of an infoset around a broad topic and in the geometric arrangement of persons and topics, but here we can note two points. First, the syntegegration process was even more fully open ended than that of the VSM. If a set of formal if revisable mathematical models were intrinsic to the latter, no such formalisms intervened in syntegegration: topics, statements, and goals were all open-endedly revisable in discussion as they reverberated around the icosahedron. Second, the icosahedral structure did undeniably constitute an infringement on individual freedom: individuals could only contribute to the discussion of topics to which they had been assigned. In this sense, and as usual, syntegegrity staged a hybrid ontology, partially thematizing and acting out an

ontology of becoming, but within a fixed framework. Beer would no doubt have remarked, as he did of the VSM, that any form of organization exacts its price, and that the price here was worth paying for the symmetric openness to becoming that it made possible. One can also note that syntegegration was a finite and limited process; participants were not locked into it, in the way that they might be within a business or a nation. So, in the next syntegegration participants could take other positions within the diagram, and, of course, the entire general topic could shift.



Throughout this book we have been concerned with the socio-ontological mismatch between cybernetics and modern institutions, with the amateurism of Beer's work on biological computing as our latest example. In the earlier chapters we also ran into examples of a constructive response to the mismatch: Kingsley Hall, for example, as providing a model for a new social basis for cybernetic forms of life, the germ of a parallel social universe as Alexander Trocchi envisaged it. Beer, too, contributed to this constructive project. As we saw, he played a key role in the establishment of the Department of Cybernetics at Brunel University—a partly successful attempt to implant a sustainable cybernetic presence in the established academic order. From another angle, the VSM can be seen as an attempt to reconfigure the world of organizations along cybernetic lines, to make that world an explicitly and self-consciously cybernetic place. And we can understand the team syntegegrity approach to decision making similarly—not now as the construction of enduring institutions, but as making available a finite and ephemeral social form lasting for just a few days, that could be mobilized ad hoc by groups at any scale for any purpose, from reorganizing the British OR society up to world governance.⁴⁸ One does not have to subscribe to the details of the VSM or team syntegegrity; the point here is that Beer's work can further enrich our imaginations with concrete examples of what Trocchi's parallel universe might look like, and that those forms would indeed be importantly different in specific ways from the hegemonic forms of our present social, political, and subpolitical arrangements. Again, ontology makes a difference, here in the domain of subpolitics.

Cybernetics and Spirituality

IN INDIA THERE ARE MANDALAS—PICTURES CONVEYING SACRED INSIGHTS NOT EXPRESSED IN WORDS. OUR MODERN CHIPS MAY NOT BE SACRAMENTALS, BUT

THEY USE NO FORM OF WORDS. COME NOW (SOMEONE MIGHT PROTEST), WE KNOW WHAT THE CHIP DOES, THE FUNCTIONS IT PERFORMS. SO (IT SHOULD BE REPLIED) DID THE YOGIS OF INDIA, THE LAMAS OF TIBET, ALSO UNDERSTAND THEIR OWN MANDALAS.

HANS BLOHM, STAFFORD BEER, AND DAVID SUZUKI,
PEBBLES TO COMPUTERS (1986, 37)

And now for something completely different. Well, not completely. The previous chapters have looked at some of the connections between cybernetics and Eastern, nonmodern, forms of spirituality, and we can continue the examination here. Beer rigorously excluded all references to spiritual concerns from his writings on management cybernetics, and one can certainly take the latter seriously without committing oneself to the former—many of Beer’s associates and followers do just that. But of our cyberneticians it was Beer who lived the fullest and most committed spiritual life, and I want now to explore the relations between his spirituality and his cybernetics, beginning with an outline of his spiritual career.

Beer was born into a High Church family and, according to his brother, before the family moved to Wales to escape the bombing of World War II,

we all attended the Church of St John the Evangelist, Shirley, where our Father and Stafford were Servers in the choir—indeed both were members of the Guild of Servers and wore their medals. . . . Stafford always sat sideways in his choir stall with one side of his glasses over his ear and the other in his mouth and frowned. The glasses, I believe, had plain glass in them as he wanted to look older than he was. At some moments when the vicar said something (I assume outrageous to Stafford) he took the glasses off and turned to glower at the pulpit. I felt very proud of him. . . . To me they were happy times and prepared us both to take the spiritual dimension of our lives seriously, wherever it took us from that traditional Anglo-Catholic Church in the thirties.⁴⁹

The spiritual dimension of Stafford’s life took him in two directions. Sometime after his military service in India, he converted to Catholicism (1965, 301), but he later “gave up Christianity and discovered Christ,” and toward the end of his life he described himself as a Buddhist, a tantric yogi. According to Allenna Leonard, he had been fascinated with Eastern philosophy since he was a small child. In his year at University College London he wanted to study Eastern philosophy, but the subject was not taught: “My dear boy, go to

SOAS”—the School of Oriental and African Studies. Instead, as we have seen, he went to India with the British Army in 1944, returning in 1947 “as thin as a rake, a very different person. . . . He was almost totally absorbed in Indian mysticism, had read endless books and had seen death, etc, I recall he told me there was no such thing as pain; it was in the mind and mind over matter and so on. To prove his point he allowed people to press lighted cigarettes onto the inside of his wrist to burn a hole while he felt nothing.”⁵⁰ So, we have these two sides to Beer’s life: the scientific (cybernetics) and the spiritual (Catholicism, Eastern mysticism, and strange performances). There is, of course, nothing especially unusual about that. Many physicists, for example, are deeply religious. But in respect of modern sciences like physics, the scientific and the spiritual are usually held apart, existing, as one might say, in different compartments of life, practiced in different places at different times, in the laboratory during the week and in church on Sunday. Bruno Latour (1993) speaks of the “crossed-out God” of modernity—the Christian God as both almighty and absent from the world of science and human affairs. As usual, cybernetics was not like that. Beer’s cybernetics and spirituality were entangled in many ways, and that is what I want to explore here, focusing first on Beer’s overall perspective on nature and then on the more esoteric aspects of his spiritual understandings and practices. The earliest of Beer’s spiritual writings was an essay published in 1965, “Cybernetics and the Knowledge of God,” and this provides a convenient entrée for both topics.

Hylozoism

First, Beer’s perspective on nature. “Cybernetics and the Knowledge of God” begins not with nature itself but with a discussion of the finitude of the human mind. “Each of us has about ten thousand million neurons to work with. It is a lot, but it is *the* lot. . . . This means that there is a strict mathematical limit to our capacity to compute cerebrally—and therefore to our understanding. For make no mistake: understanding is mediated by the machinery in the skull” (Beer 1965, 294). As a corollary, beyond our cerebral limits there must exist in the world things which we cannot know.⁵¹ Here we recognize the cybernetic ontology of unknowability—Beer was writing for a readership of nonspecialists; otherwise, he could simply have said that the cosmos was an exceedingly complex system, as he had defined the term in *Cybernetics and Management* in 1959. There is, though, a difference in the way in which Beer develops this thought in this essay. One can think of the economic environment of a firm as being exceedingly complex in a mundane fashion: we can readily

comprehend many aspects of the economy; it is just impossible to hold all of them and their interrelations in consciousness at once. In the religious context, in contrast, Beer reaches for a more absolute sense of unknowability, invoking repeatedly “an irreducible mystery: that there is anything” (Beer 1965, 298). And this is where God comes in: “Here is another definition [of God], which I would add to the scholastic list of superlative attributes: *God is what explains the mystery*” (299). This is an odd kind of explanation, since Beer could not offer any independent definition of the *explanans*. One mystery, God, is simply defined here as that which explains another, existence. In ordinary language, at least, there is no “gap” between the two terms, so I am inclined to read Beer as saying here that matter and spirit are one, or that they are two aspects of an underlying unity. This is part of what I want to get at in describing Beer’s appreciation of nature as hylozoist—the understanding that nature is infused, one might say, by spirit.

At any rate, we can see here that the ontology of unknowability was a straightforward point of linkage, almost of identity, between Beer’s worldly cybernetics and his spirituality: the correlated mysteries of existence and of God are simply the mystery of exceedingly complex mundane systems taken to the *N*th degree, where *N* is infinite. And along with this ontological resonance, we can find an epistemological one. I have remarked several times on Beer’s cybernetic suspicion of articulated knowledge and models, as a not necessarily reliable detour away from performance, and he expressed this suspicion, again to the *N*th degree, in relation to the spiritual (Beer 1965, 294–95, 298):

To people reared in the good liberal tradition, man is in principle infinitely wise; he pursues knowledge to its ultimate. . . . To the cybernetician, man is part of a control system. His input is grossly inadequate to the task of perceiving the universe. . . . There is no question of “ultimate” understanding. . . . It is part of the cultural tradition that man’s language expresses his thoughts. To the cybernetician, language is a limiting code in which everything *has* to be expressed—more’s the pity, for the code is not nearly rich enough to cope. . . . Will you tell me that science is going to deal with this mystery [of existence] in due course? I reply that it cannot. The scientific reference frame is incompetent to provide an existence theorem for existence. The layman may believe that science will one day “explain everything away”; the scientist himself ought to know better.

Epistemologically as well as ontologically, then, Beer’s cybernetics crossed over smoothly into a spiritually charged hylozoism. And we can follow the



Figure 6.18. The Gatineau River, Quebec. Source: Blohm, Beer, and Suzuki 1986, 51. (Photo: Hans-Ludwig Blohm. © Hans-Ludwig Blohm, Canada.)

crossover further by jumping ahead twenty years, to a book published in 1986, *Pebbles to Computers: The Thread*, which combines photographs by Hans Blohm with text by Stafford Beer and an introduction by David Suzuki. It is a coffee-table book with lots of color pictures and traces out a *longue durée* history of computing, running from simple counting (“pebbles”) to digital electronic computers. The history is not, however, told in a linear fashion leading up to the present, but as a topologically complex “thread”—drawn by Beer as a thick red line twisting around photographs and text and linking one page to the next—embracing, for example, Stonehenge as an astronomical computer and Peruvian *quipus*, beautiful knotted threads, as calculational devices. Here Beer develops his ontological vision further. Under the heading “Nature Calculates,” he comments on a photograph of the Gatineau River (fig. 6.18) that catches the endless complexity of the water’s surface (Blohm, Beer, and Suzuki 1986, 54): “This exquisite photograph of water in movement . . . has a very subtle message for us. It is that nature’s computers *are* that which they compute. If one were to take intricate details of wind and tide and so on, and use them . . . as ‘input’ to some computer simulating water—what computer would one use, and how express the ‘output’? Water itself: that answers both

those questions.” And then he goes on to reproduce one of his own poems, written in 1964, “Computers, the Irish Sea,” which reads (Blohm, Beer, and Suzuki 1986, 52; reproduced from Beer 1977):

That green computer sea
with all its molecular logic
to the system’s square inch,
a bigger brain than mine,
writes out foamy equations from the bow
across the bland blackboard water.

Accounting for variables
which navigators cannot even list,
a bigger sum than theirs,
getting the answer continuously right
without fail and without anguish
integrals white on green.

Cursively writes recursively computes
that green computer sea
on a scale so shocking
that all the people sit dumbfounded
throwing indigestible peel at seagulls
not uttering an equation between them.

All this liquid diophantine stuff
of order umpteen million
is its own analogue. Take a turn
around the deck and understand
the mystery by which what happens
writes out its explanation as it goes.

In effect, this poem is another reexpression of the cybernetic ontology of unknowability, where the unknowability is conceived to reside in the sheer excess of nature over our representational abilities. The water knows what it is doing and does it faultlessly and effortlessly in real time, a performance we could never emulate representationally. Nature does “a bigger sum than theirs”—exceeding our capacities in way that we can only wonder at, “shocked” and “dumbfounded.”⁵² But Beer then adds a further point (Blohm, Beer, and Suzuki 1986, 54): “The uneasy feeling that [this poem] may have caused derives, perhaps, from insecurity as to who is supposed to be in charge. Science (surely?) ‘knows the score.’ Science does the measuring after all. . . .

But if art is said to imitate nature, so does science. . . . Who will realize when the bathroom cistern has been filled—someone with a ruler and a button to press, or the ballcock that floats up to switch the water off? *Nature* is (let it be clear that) *nature* is in charge.” There is a clear echo here of Beer’s work with biological computers (which, as mentioned earlier, also figure in *Pebbles*): not only can we not hope to equal nature representationally, but we do not need to—nature itself performs, acts, is in charge. This idea of nature as active as well infused with spirit is the definition of hylozoism, which is why I describe Beer’s ontology as hylozoist. We could even think of Beer’s distinctive approach to biological computing as a form of hylozoist, or *spiritual*, engineering. Aside from the reference to spirit, we can also continue to recognize in this emphasis on the endless performativity of matter the basic ontology of British cybernetics in general.⁵³ And we can make further connections by looking at Beer’s thoughts on mind. In *Pebbles*, he refers to the Buddhist Diamond Sutra: “Think a thought, it says, ‘unsupported by sights, sounds, smells, tastes, touchables, or any objects of the mind.’ Can you do that?” (Blohm, Beer, and Suzuki 1986, 67). The implicit answer is no. Sensations, feelings, cognition—all emerge from, as part of, the unrepresentable excess of nature, they do not contain or dominate it. And under the heading “The Knower and the Known Are One” Beer’s text comes to an end with a quotation from *hsin ming* by Sengstan, the third Zen patriarch (d. 606) (105):

Things are objects because of the mind;
The mind is such because of things.
Understand the relativity of these two
and the basic reality: the unity of emptiness.
In this emptiness the two are indistinguishable
and each contains in itself the whole world.

I cannot give a fully cybernetic gloss of this quotation; the notion of “emptiness” presently eludes me. But one can go quite a way in grasping the Zen patriarch’s sentiment by thinking about Ashby’s multihomeostat setups—one homeostat standing for the brain or mind, the others for its world—or perhaps even better, of the configuration of DAMS in which a subset of its elements could be designated the mind and the others that to which the mind adapts. In the dynamic interplay of mind and world thus instantiated, “objects” and “mind” do reciprocally condition each other.⁵⁴

I can sum this section up by saying that there are two perspectives one might adopt on the relation between cybernetics and Beer’s spiritual stance

as discussed so far. If one balks at any reference to the spiritual, then one can see Beer's hylozoism as an *extension* of cybernetics proper, adding something to the secular part that we have focused on elsewhere. Then we could say: *This* is how one might extend cybernetics into the realm of the spiritual if one wanted to; *this* is the kind of direction in which it might lead. On the other hand, if one were prepared to recognize that there is a brute mystery of existence, and if one were willing to associate that mystery with the spiritual realm, itself defined by that association, then one could say that Beer's hylozoism just is cybernetics—cybernetics taken more seriously than we have taken it before. Beer's spirituality can thus be seen as either continuous with or identical to his worldly cybernetics—a situation very different from the discontinuity between the modern sciences and the crossed-out God of modernity. Once more we see that ontology makes a difference, now in the spiritual realm—the cybernetic ontology aligning itself with Eastern spirituality rather than orthodox Christianity and, at the same time, eroding the boundary between science and spirit.⁵⁵

Tantrism

YOGA MEANS UNION, WHETHER OF SELF AND COSMOS, MAN AND WOMAN, THE DIFFERENT CHAMBERS OF THE MIND. . . . IN THE LIMIT, THEREFORE, OF THE A AND THE NOT-A.

STAFFORD BEER, "I SAID, YOU ARE GODS" (1994 [1980], 385)

The second spiritual topic I need to discuss has to do with esoteric knowledge and practice, and here we can also begin with Beer's 1965 essay "Cybernetics and the Knowledge of God." One might think that having named existence as the ultimate mystery and having defined God as its explanation, Beer would have reduced himself to silence. Instead, the essay opens up a discursive space by thinking along the same lines as Beer did in his management cybernetics. In the latter he insisted that the factory's economic environment was itself ultimately unknowable, but he also insisted that articulated models of the economy were useful, as long as they were treated as revisable in practice and not as fixed and definitive representations of their objects, and the essay follows much the same logic in the spiritual realm.

Just as the factory adapts to its economic environment in a performative fashion without ever fully grasping it, so it might be that, while our finite brains can never rationally grasp the essence of God, nevertheless, the spiri-

tual bears upon us and leaves marks upon the “human condition” (Beer 1965, 294). Beer gives the example of suffering. “The child of loving parents is suddenly seized by them, bound and gagged and locked in a dark cellar. What is the child to make of that? It must be evident to him that (i) his parents have turned against him; but (ii) they have done so without any cause, and therefore (iii) the world is a place where things can happen without causes.” In fact, in this story, “what has actually happened is that the home has suddenly been raided by secret police, seeking children as hostages. There was no time to explain; there was too much risk to the child to permit him any freedom” (296). Like the parents in this story, then, Beer had the idea that God moves in a mysterious way *which has effects on us*, though, as child analogues, we cannot grasp God’s plan. The marks of God’s agency are evident in history.

That means that we *can* accumulate knowledge, though never adequate, of God, just as factory managers learn about their environments. And that, in turn, implies, according to Beer in 1965, that there are two authorities we should consult in the realm of the spiritual. One is the Catholic Church—the “admonitory church” (Beer 1965, 300)—as the repository of our accumulated wisdom in brushing up against and adapting to the spiritual. But since Beer later renounced Catholicism, his second source of authority bears emphasis. It is “the total drift of human knowledge. Though compounded of the work of individual brains . . . the totality of human insight can conceivably be greater than the insight of one brain. For people use their brains in markedly different, and perhaps complementary ways.” In cybernetic terms, many brains have more variety than one and thus are better able to latch onto the systems with which they interact. And the reference to “complementary ways” here asserts that there is even more variety if we pay attention to the historical drift of knowledge over a range of spiritual traditions rather than within a single one (301): “Anthropologist friends point out so many alien cultures produce so many similar ideas about God, about the Trinity, about the Incarnation. They expect me to be astonished. They mean that I ought to realise there is something phoney about my specifically Christian beliefs. I am astonished, but for opposite reasons. I am cybernetically impressed . . . by Augustine’s precept: ‘securus judicat orbis terrarum’—the world at large judges rightly.” Beer perhaps verges on heresy in his willingness to find spiritual truths across the range of the world’s religions, but he saves himself, if he does, by seizing in this essay on just those truths that the church itself espoused: God, the Trinity, the incarnation of God in Christ. Later, when he had left the church, he seized on other ones, as we will see. For the moment, let me repeat that here Beer had developed a cybernetic rhetoric for smuggling all sorts of positive

spiritual knowledge past the ontology of unknowability, and it is worth noting that one example of this figures prominently in the “Knowledge of God” essay (Beer 1965, 297): “In fact, we—that is men—have a whole reference frame, called religion, which distinguishes between orders of creation precisely in terms of their communication capacity. The catalogue starts with inanimate things, works up through the amoeba and jellyfish to the primates, runs through monkeys to men—and then goes gaily on: angels, archangels, virtues, powers, principalities, dominations, thrones, seraphim, cherubim.” So here, in the writings of the world’s greatest management cybernetician, then director of one of the world’s first OR consulting groups, we find the medieval Great Chain of Being, running continuously from rocks and stones to angels and God. There is, of course, no integral connection between this and cybernetics, but, at the same time, it is hard not to read it back into the development of Beer’s cybernetics. The recursive structure of the VSM, as discussed so far, is nothing but the Great Chain of Being, sawn off before the angels appear—and, as we shall shortly see, Beer subsequently insisted on recontinuing the series, though in non-Christian terms.



As I said, these maneuvers in “Knowledge of God” open the way for a positive but revisable domain of spiritual knowledge, and we can learn more of where Beer came to stand in this domain from a book that he wrote that was never published, “Chronicles of Wizard Prang” (Beer 1989b).⁵⁶ Wizard Prang is the central character in the twenty chapters of the book and clearly stands for Beer himself: he lives in a simple cottage in Wales, has a long beard, wears simple clothes, eats simple food, describes himself as “among other things . . . a cybernetician” (133) and continually sips white wine mixed with water, “a trick he had learned from the ancient Greeks” (12). The thrust of the book is resolutely spiritual and specifically “tantric” (103). Its substance concerns Prang’s doings and conversations, the latter offering both cybernetic exegesis of spiritual topics and spiritually informed discussions of topics that Beer also addresses in his secular writings: the failings of an education system that functions to reproduce the world’s problems (chap. 2); the sad state of modern economics (chap. 15); the need to beware of becoming trapped within representational systems, including tantric ones (chap. 15).⁵⁷ We are entitled, then, to read the book as a presentation of the spiritual system that Beer lived by and taught when he was in Wales, albeit a fictionalized one that remains veiled in certain respects. And with the proviso that I am out of my depth here—I am no expert on the esoteric doctrines and practices to follow—I




Figure 6.19. Beer meditating. (Photo: Hans-Ludwig Blohm. © Hans-Ludwig Blohm, Canada.)

want to explore some of the resonances and connections between Beer’s tantrism and his cybernetics.

Tantrism is a hard concept to pin down. In his book *Stafford Beer: A Personal Memoir*, David Whittaker notes that “the word ‘tantra’ comes from the Sanskrit root *tan* meaning ‘to extend, to expand.’ It is a highly ritualistic philosophy of psycho-physical exercises, with a strong emphasis on visualization, including concentration on the yogic art of mandalas and yantras. The aim is a transmutation of consciousness where the ‘boundary’ or sense of separation of the self from the universe at large dissolves” (Whittaker 2003, 13).⁵⁸ And we can begin to bring this description down to earth by noting that meditation was a key spiritual practice for Beer.

Here, then, we can make contact with the discussion from earlier chapters—of meditation as a technology of the nonmodern self, aimed at exploring regions of the self as an exceedingly complex system and achieving “altered states of consciousness” (Beer 1989b, 41).⁵⁹ Like the residents in the Archway communities, but in a different register, Beer integrated this technology into his life. Beyond that we can note that, as Whittaker’s definition of tantrism suggests, Beer’s style of meditation involved visual images. He both meditated upon images—mandalas, otherwise known as yantras (fig. 6.20)—and



Figure 6.20. A yantra. Source: Beer 1989b, chap. 14, 105.

engaged in visualization exercises in meditation. In the tantric tradition this is recognized as a way of accessing a subtle realm of body, energy, and spirit—experiencing the body, for example, as a sequence of chakras ascending from the base of the spine to the top of the head, and eventually aiming at union with the cosmos—“yoga means union,” as Beer was wont to put it.⁶⁰

Three points are worth noting here. First, we again find a notion of decentering the self here, relative both to the chakras as lower centers of consciousness and to the higher cosmos. As discussed before, we can understand this sort of decentering on the model of interacting homeostats, though the details, of course, are not integrally cybernetic but derive specifically from the tantric tradition. Second, we should recognize that yantras are, in a certain sense, symbolic and representational. Interestingly, however, Beer has Perny, his apprentice, stress their *performative* rather than symbolic quality when they first appear in Beer’s text. Perny remarks on the yantra of figure 6.20 that “I was taught to use this one as a symbol on which to meditate.” Another disciple replies, “It’s sort of turning cartwheels.” “I know what you mean,” Perny responds. “This way of communicating, which doesn’t use words, seems to work through its physiological effects” (Beer 1989b, 106). We thus return to a performative epistemology, now in the realm of meditation—the symbol as integral to performance, rather than a representation having importance in its own right.

Third, and staying with the theme of performance, we could recall from chapter 3 that Grey Walter offered cybernetic explanations not just for the altered states achieved by Eastern yogis, but also for their strange bodily performances, suspending their metabolism and so on. We have not seen much of these strange performances since, but now we can go back to them. Wizard Prang himself displays displays unusual powers, though typically small ones which are not thematized but are dotted around the stories that make up the book. At one point Prang makes his end of a seesaw ascend and then descend just by intending it: “Making oneself light and making oneself heavy are two of the eight occult powers”; Prang can see the chakras and auras of others and detect their malfunctioning; Perny “change[s] the direction of a swirl [in a stream] by identifying with it rather than by exerting power”; the logs in the fireplace ignite themselves; spilled wine evaporates instantly on hitting the tiles; Prang sends a blessing flying after two of his disciples, “with the result that Toby [slips] and [falls] over with the force of it.” More impressively, Perny remarks that “you give me telepathic news and I’ve seen you do telekinetic acts,” and at one point Prang levitates, though even this is described in a humorous and self-deprecating fashion: “The wizard’s recumbent form slowly and horizontally rose to the level of where his midriff would be if he were standing up. He stayed in that position for ten seconds, then slowly rotated. His feet described an arc through the air which set them down precisely, smoothly onto the floor. ‘My God,’ breathed Silica, ‘What are you doing?’ . . . ‘Demonstrating my profession of wizardry, of course.’ ‘Do you often do things like that?’ ‘Hardly ever. It’s pretty silly, isn’t it?’”⁶¹ I paid little attention to these incidents in Beer’s text until I discovered that the accrual of nonstandard powers is a recognized feature of spiritual progress by the yogi, and that there is a word for these powers: *siddhis*.⁶² Beer’s practice was securely within the tantric tradition in this respect, too.

In these various ways, then, Beer’s spiritual knowledge and practice resonated with the cybernetic ontology of exceedingly complex performative systems, though, as I said, the detailed articulation of the ontology here derived not from cybernetics but from the accumulated wisdom of the tantric tradition. Having observed this, we can now look at more specific connections that Beer made between his spirituality and his cybernetics.



Beer’s worldly cybernetics as I described it earlier is not as worldly as it might seem. This is made apparent in *Beyond Dispute*. For the first ten chapters, 177 pages, this book is entirely secular. It covers the basic ideas and form of team




Figure 6.21. The enneagram. Source: S. Beer, *Beyond Dispute: The Invention of Team Syntegrity* (New York: Wiley, 1994), 202, fig. 12.4.

syntegrity, describes various experiments in syntegration and refinements of the protocol, and elaborates many of Beer's ideas that are by now familiar, up to and including his thoughts on how syntegration might play a part in an emergent "world governance." In chapters 11–14, 74 pages in all, the book takes on a different form. As if Beer had done his duty to the worldly aspects of the project in the opening chapters, now he conjures up its spiritual aspects and the esoteric knowledge that informs it. I cannot rehearse the entire content of these latter chapters, but we can examine some of its key aspects.

Chapter 12 is entitled "The Dynamics of Icosahedral Space" and focuses on closed paths around the basic syntegration icosahedron, paths that lead from one vertex to the next and eventually return to their starting points. Beer's interest in such paths derived from the idea mentioned above, that in syntegration, discussions reverberate around the icosahedron, becoming the common property of the infoset. In chapter 12, this discussion quickly condenses onto the geometric figure known as an *enneagram* (fig. 6.21), which comprises a re-entrant six-pointed form superimposed on a triangle. Beer offers an elaborate spiritual pedigree for this figure. He remarks that he first heard about it in the 1960s in conversations with the English mystic John Bennett, who had in turn

been influenced by the work of Peter D. Ouspensky and George Ivanovich Gurdjieff; that there is also a distinctively Catholic commentary on the properties of the enneagram; and that traces of it can be found in the Vedas, the most ancient Sanskrit scriptures, as well as in Sufism (Beer 1994b, 202–4). Beer also mentions that while working on Project Cybersyn in Chile in the early 1970s he had been given his own personal mandala by a Buddhist monk, that the mandala included an enneagram, and that after that he had used this figure in his meditational practices (205).⁶³ Once more we can recognize the line of thought Beer set out in “Cybernetics and the Knowledge of God.” The enneagram appears in many traditions of mystical thought; it can therefore be assumed to be part the common wisdom of mankind, distilled from varied experience of incomprehensible realms; but its significance is performative, as an aid to meditation, rather than purely representational.

So what? Beer recorded that in the synteintegration experiments of the early 1990s he had acquired a new colleague in Toronto, Joe Truss, who had once founded a business based on an enneagrammatic model, and that Truss had then succeeded in finding reentrant enneagrammatic trajectories within the synteintegration icosahedron.⁶⁴ Truss and Beer were both exceptionally impressed by the fact that these trajectories were three-dimensional, rather than lying in a single plane as in figure 6.21 (Beer 1994b, 206): “Joe came to my house late at night to show me his discovery, and he was very excited. Well, all such moments are exciting. But I was unprepared that he should say, ‘Do you see what this means? The icosahedron is the actual *origin* of the enneagram, and the ancients knew it. Could it not be possible that the plane figure was coded esoteric knowledge?’ Obviously (now!) it could.” From this point on, if not before, synteintegration took on for Beer an intense spiritual as well as practical significance, especially as far as its reverberations along closed pathways were concerned.⁶⁵ Here, then, we have an example of the sort of very specific and even, one could say, technical continuities that Beer constructed between his worldly cybernetics and his spiritual life, with the enneagram as a pivot between the everyday geometry of the icosahedron and meditative practice. This immediate continuity between the secular and the spiritual contrasts interestingly, as usual, with the separation of these two realms that characterizes modernity. It points to the unusual “earthy” and hylozoist quality of cybernetic spirituality, as a spirituality that does not recognize any sharp separation between the sacred and the profane.

I mentioned earlier the appearance of the Great Chain of Being in Beer’s “Knowledge of God” essay, and that this reappeared in a truncated version in his published discussions of the viable system model. We might doubt,




Figure 6.22. "Theory of recursive consciousness." Source: S. Beer, *Beyond Dispute: The Invention of Team Syntegrity* (New York: Wiley, 1994), 253, fig. 14.2.

however, that this truncation had much significance for Beer personally; it seems highly likely that the VSM's recursive structure was always spiritually charged in Beer's imagination—that again in this respect the mundane and the spiritual were continuous for Beer. And in *Beyond Dispute*, the truncation was explicitly undone, as shown in figure 6.22, displaying nine levels of recursion, running from individual neurons in the brain, through individual consciousness ("cerebrum"), up to Gaia and the cosmos. Chapter 14 of *Beyond Dispute* is a fascinating cybernetic-political-mystical commentary on these different levels, and I can follow a few threads as illustrations of Beer's thought and practice.

One point to note is that while the labeling of levels in figure 6.22 is secular, at least until one comes to "Gaia" and "cosmos," Beer's discussion of them is not. It is distinctly hybrid, in two senses. On the one hand, Beer accepts current biological knowledge of the nervous system, as he did in developing the VSM, while, at the same time, conceptualizing it as a cybernetic adaptive system; on the other hand, he synthesizes such biological knowledge with esoteric, mystical, characteristically Eastern accounts of the subtle body accessible to the adept. The connection to the latter goes via a cybernetic analysis of *consciousness* as the peculiar property of reentrant structures.⁶⁶ The human brain would be the paradigmatic example of such a structure (containing an astronomical number of reentrant neuronal paths), but Beer's argument was that

any reentrant structure might have its own form of consciousness. “Plexus,” at the third level of recursion in figure 6.22, refers to various nervous plexuses, concatenations of nerve nets, that can be physiologically identified within the body. Beer regards these as homeostat-like controllers of physiological functions. At the same time, he imputes to them their own form of “infosettic consciousness,” remarks that it “is not implausible to identify the six ‘spiritual centres’ which the yogi calls chakras with plexus activity in the body,” and finally asserts that “knowledge of the existence of infosettic consciousness within the other five chakras [besides the chakra associated with brain] is possible to the initiate, as I attest from yogic experience myself” (Beer 1994b, 247). With this move Beer deeply intertwines his management cybernetics and his spirituality, at once linking the subtle yogic body with physiological cybernetic structures and endowing the recursive structure of the VSM with spiritual significance, this double move finding its empirical warrant in Beer’s meditational practice.

Beer then moves up the levels of consciousness. At the fourth level of figure 6.22 we find “cerebrum,” the level of individual consciousness, which Beer identifies in the subtle body with Ajna, the sixth chakra. Then follow four levels having to do with social groupings at increasing scales of aggregation. “Neighbourhood” refers to the small groups of individuals that come together in syntegration, and “it must be evident that the theory of recursive consciousness puts group mind forward as the fifth embedment of consciousness, simply because the neighbourhood info-set displays the usual tokens of consciousness” (Beer 1994b, 248). From this angle, too, therefore, syntegration had a more than mundane significance for Beer. Not simply an apparatus for free and democratic discussion, syntegration produces a novel emergent phenomenon, a group mind which can be distinguished from the individual minds that enter into it, and which continues the spiritually charged sequence of levels of consciousness that runs upward from the neuron to the cerebrum and through the yogic chakras. At higher levels of social aggregation, up to the level of the planet, Beer also glimpses the possibilities for collective consciousness but is of the opinion that practical arrangements for the achievement of such states “work hardly at all . . . there is no cerebrum [and] we are stuck with the woeful inadequacy of the United Nations. The Romans did better than that” (249).

The ninth level of figure 6.22, “cosmos,” completes the ascent. Here the mundane world of the individual and the social is left entirely behind in yogic union with “cosmic consciousness,” accessible via the seventh chakra, Sahasra, the thousand-petaled lotus. And Beer concludes (255):

For the yogi, the identification of all the embedments, and particularly his/her own selfhood embodied at the fourth embedment, with the cosmos conceived as universal consciousness, is expressed by the mantra Tat Tvam Asi: “That You Are.” These are the last three words of a quotation from one of the Ancient Vedic scriptures, the Chhandogya Upanishad, expressing the cosmic identification:

That subtle essence
which is the Self of this entire world,
That is the Real,
That is the Self,
That You Are.

— — — — —

Rather than trying to sum up this section, it might be useful to come at the topic from a different angle. Beer, of course, identified the spiritual aspect of his life with the tantric tradition, but it strikes me that his fusion of cybernetics and spirituality also places him in a somewhat more specific lineage which, as far as I know, has no accepted name. In modernity, matter and spirit are assigned to separate realms, though their relations can be contested, as recent arguments about intelligent design show: should we give credit for the biological world to God the Creator, as indicated in the Christian Bible, or to the workings of evolution on base matter, as described by modern biology? What interests me about Beer’s work is that it refused to fall on either side of this dichotomy—we have seen that his science, specifically his cybernetics, and his understanding of the spiritual were continuous with one another—flowed into, structured, and informed each other in all sorts of ways. This is what I meant by referring to the earthy quality of his spirituality: his tantrism and his mundane cybernetics were one. Once more one could remark that ontology makes a difference, here in the realm of the spiritual.⁶⁷ But my concluding point is that Beer was not alone in this nondualist space.

I cannot trace out anything like an adequate history of the lineage of the scientific-spiritual space in which I want to situate Beer, and I know of no scholarly treatments, but, in my own thinking at least, all roads lead to William James—in this instance to his *Varieties of Religious Experience*, as an empirical but nonsceptical inquiry into spiritual phenomena. James’s discussion of the “anaesthetic revelation”—transcendental experience brought on by drugs and alcohol—is a canonical exploration of technologies of the nonmodern self, Aldous Huxley *avant la lettre*. Huxley himself lurked in the margins of chapters 3 and 5 above, pursuing a biochemical understanding

NAME OF CHAKRA	POSITION ON SURFACE	APPROXIMATE POSITION OF SPINAL CHAKRA	SYMPATHETIC PLEXUS	CHIEF SUBSIDIARY PLEXUSES
Root	Base of spine	4th Sacral	Coccygeal	...
Spleen	Over the spleen	1st Lumbar	Splenic	...
Navel	Over the navel	8th Thoracic	Celiac or Solar	Hepatic, pyloric, gastric, mesenteric, etc.
Heart	Over the heart	8th Cervical	Cardiac	Pulmonary, coronary, etc.
Throat	At the throat	3rd Cervical	Pharyngeal	...
Brow	On the brow	1st Cervical	Carotid	Cavernous, and cephalic ganglia generally

Figure 6.23. Chakras and plexuses. Reproduced from *The Chakras*, by C. W. Leadbeater (Leadbeater 1990, 41, table 2). (With permission from The Theosophical Publishing House, Adyar, Chennai-600 020, India. © The Theosophical Publishing House, Adyar, Chennai-600 020, India. www.tw-adyar.org.)

of spiritual experience without intending any reduction of the spiritual to base matter, with John Smythies and Humphrey Osmond's research on brain chemistry lurking behind him. Gregory Bateson, Alan Watts, and R. D. Laing likewise ran worldly psychiatry and Buddhism constructively together, with a cybernetic ontology as the common ground. One thinks, too, of the Society for Psychical Research (with which both James and Smythies were associated) as a site for systematic integration of science and the spiritual (chap. 3, n. 62).

From another angle—if it is another angle—a canonical reference on the chakras that loomed so large in Beer's thought and practice is C. W. Leadbeater's *The Chakras*, continuously in print, according to the back cover of my copy (Leadbeater 1990), since 1927, and published by the Theosophical Society. Much of Beer's esoteric writing echoes Leadbeater's text, including, for example, the association between the chakras and nerve plexuses just discussed (fig. 6.23).⁶⁸ Theosophy, too, then, helps define the scientific-spiritual space of Beer's cybernetics. And coming up to the present, one also thinks of

certain strands of New Age philosophy and practice, already mentioned in chapter 5, as somehow running together science and spirituality, mind, body, and spirit.⁶⁹

We should thus see Beer's cybernetic tantrism as an event within a broader scientific-spiritual history, and I can close with two comments on this. First, to place Beer in this lineage is not to efface his achievement. On the one hand, Beer went much further than anyone else in tying cybernetics—our topic—into the realm of the spirit. On the other hand, from the spiritual side, Beer went much further than anyone else in developing the social aspects of this nonmodern assemblage. Esoteric writings seldom go beyond the realm of the individual, whereas the VSM and team synteegrity were directed at the creation of new social structures and the rearrangement of existing ones in line with cybernetic and, we can now add, tantric sensitivities. Second, placing Beer in relation to this lineage returns us to questions of institutionalization and marginality. The entire lineage could be described as sociologically occult—hidden and suspect. Even now, when New Age has become big business, it remains walled off from established thought and practice. Despite—or, perhaps better, because of—its elision of mind, body, and spirit distinctions, New Age remains invisible in contemporary debates on the relation between science and religion. Like Gysin's Dream Machines, New Age spirituality and Beer's spirituality fail to find a place within modern schemata of classification. And, to change direction again, perhaps we should regret this. The early twenty-first century seems like a time when we should welcome a form of life that fuses science and spirituality rather than setting them at each other's throats. Again, this exploration of the history of cybernetics offers us a sketch of another future, importantly different from the ones that are more readily imagined.

Brian Eno and New Music

[BEER'S WORK] SO FUNDAMENTALLY CHANGED THE WAY THAT I THOUGHT ABOUT MUSIC THAT IT'S VERY DIFFICULT TO TRANSLATE INTO INDIVIDUAL THINGS, IT JUST CHANGED THE WHOLE WAY I WORK. . . . STAFFORD FOR ME WAS THE DOORWAY INTO A WHOLE WAY OF THINKING.

BRIAN ENO, QUOTED IN DAVID WHITTAKER, *STAFFORD BEER* (2003, 57, 63)

We touched on relations between cybernetics and the arts in chapters 3 and 4 as well as briefly here in connection with biological computing, and I want

to end this chapter by returning to this topic. The focus now is on music, and the intersection between Beer's cybernetics and the work of the composer and performer Brian Eno. If Beer himself is not widely known, many people have heard of Eno and, with any luck, know his music (which is, like any music, impossible to convey in words, though what follows should help to characterize it). Eno's first claim to fame was as a member of Roxy Music, the greatest rock band to emerge in the early 1970s. Subsequently, he left Roxy and went on to develop his own distinctive form of "ambient" and "generative" music (as well as important collaborations with David Bowie, Talking Heads, and U2), with *Music for Airports* (1978) as an early canonical example.⁷⁰ The content of this music is what I need to get at, but first I want to establish Eno's biographical connection to cybernetics.

In an interview with David Whittaker (Whittaker 2003, 53–63), Eno recalled that he first became interested in cybernetics as an art student in Ipswich between 1964 and 1966. The principal of the art school was Roy Ascott, Britain's leading cybernetic artist, who will reappear in the next chapter, and the emphasis at Ipswich was on "process not product. . . . Artists should concentrate on the way they were doing things, not just the little picture that came out at the end. . . . The process was the interesting part of the work" (53). Eno was drawn further into cybernetics in 1974 when his mother-in-law lent him a copy of Beer's *Brain of the Firm*, which she had borrowed from Swiss Cottage Library in London. Eno was "very, very impressed by it" and in 1975 wrote an essay in which he quoted extensively from *Brain*. He sent a copy to Beer, who came to visit him in Maida Vale (Whittaker 2003, 55–56).⁷¹ In 1977 Beer invited Eno for an overnight visit to the cottage in Wales, where Eno recalled that dinner was boiled potatoes and the following conversation took place (55):

[Beer] said "I carry a torch, a torch that was handed to me along a chain from Ross Ashby, it was handed to him from . . . Warren McCulloch." He was telling me the story of the lineage of this idea . . . and said "I want to hand it to you. I know it's a responsibility and you don't have to accept, I just want you to think about it." It was a strange moment for me, it was a sort of religious initiation . . . and I didn't feel comfortable about it somehow. I said "Well, I'm flattered . . . but I don't see how I can accept it without deciding to give up the work I do now and I would have to think very hard about that." We left it saying the offer is there, but it was very strange, we never referred to it again, I wasn't in touch with him much after that. I'm sure it was meant with the best of intentions and so on but it was slightly weird.

Now we can turn to the substantive connection between Eno's music and Beer's cybernetics: what did Brian get from Stafford? Speaking of this connection and apparently paraphrasing from *Brain of the Firm*, in his interview with David Whittaker Eno said that "the phrase that probably crystallised it [Eno's cybernetic approach to music] . . . says 'instead of specifying it in full detail; you specify it only somewhat, you then ride on the dynamics of the system in the direction you want to go.' That really became my idea of working method" (57).⁷² And the easiest way to grasp this idea of riding the dynamics of the system is, in the present context, ontologically. Beer's ontology of exceedingly complex systems conjures up a lively world, continually capable of generating novel performances. Eno, so to speak, picked up the other end of the stick and focused on building musical worlds that would themselves exhibit unpredictable, emergent becomings. And we can get at the substance of this by following a genealogy of this approach that Eno laid out in a 1996 talk titled "Generative Music" (Eno 1996b). This begins with a piece called *In C* by Terry Riley, first performed at the San Francisco Tape Music Center in 1964 (Eno 1996b, 2–3):

It's a very famous piece of music. It consists of 52 bars of music written in the key of C. And the instructions to the musicians are "proceed through those bars at any speed you choose." So you can begin on bar one, play that for as many times as you want, 20 or 30 times, then move to bar 2, if you don't like that much just play it once, go on to bar three. The important thing is each musician moves through it at his or her own speed. The effect of that of course is to create a very complicated work of quite unpredictable combinations. If this is performed with a lot of musicians you get a very dense and fascinating web of sound as a result. It's actually a beautiful piece.

Here we find key elements of Eno's own work. The composer sets some initial conditions for musical performance but leaves the details to be filled in by the dynamics of the performing system—in this case a group of musicians deciding on the spot which bars to play how often and thus how the overall sound will evolve in time. Eno's second example is a different realization of the same idea: Steve Reich's *It's Gonna Rain*, first performed in 1965, also at the Tape Music Center.⁷³ In this piece a loop of a preacher saying "It's gonna rain" is played on two tape recorders simultaneously, producing strange aural effects as the playbacks slowly fall out of phase: "Quite soon you start hearing very exotic details of the recording itself. For instance you are aware after several minutes that there are thousands of trumpets in there. . . . You also

become aware that there are birds” (Eno 1996b, 3). Again in this piece, the composer specifies the initial conditions for a performance—the selection of the taped phrase, the use of two recorders, and then “rides the dynamics of the system”—in this case the imperfection of the recorders that leads them to drift out of synchronization, rather than the idiosyncratic choices of human musicians—to produce the actual work.

Eno then moves on to one of his own early post-*Brain* pieces composed in this way, from *Music for Airports*. This consists of just three notes, each repeating at a different interval from the others—something like $23 \frac{1}{2}$, $25 \frac{7}{8}$, and $29 \frac{15}{16}$ seconds, according to Eno. The point once more is that the composer defines the initial conditions, leaving the piece to unfold itself in time, as the notes juxtapose themselves in endless combinations (Eno 1996b, 4).

In his talk, Eno then makes a detour through fields like cellular automata and computer graphics, discussing the endlessly variable becomings of the Game of Life (a simple two-dimensional cellular automaton, developed by John Conway: Poundstone 1985), and simple screen savers that continually transform images arising from a simple “seed.” In each case, unpredictable and complex patterns are generated by simple algorithms or transformation rules, which connects back to Eno’s then-current work on a musical generative system—a computer with a sound card. Eno had contrived this system so as to improvise probabilistically within a set of rules, around 150 of them, which determined parameters such as the instruments and scales to be employed, harmonies that might occur, and steps in pitch between consecutive notes.⁷⁴ As usual, one should listen to a sample of the music produced by this system, but at least Eno (1996b, 7) found that it was “very satisfying,” and again we can see how it exemplifies the idea of riding the dynamics of what has by now become a sophisticated algorithmic system.

Thus the basic form and a sketchy history of Brian Eno’s ambient and generative music, and I want to round off this chapter with some commentary and a little amplification. First, back to Roxy Music. Eno does not include his time with Roxy in any of his genealogies, and one might assume a discontinuity between his Roxy phase and his later work, but the story is more interesting than that. Eno played (if that is the word) the electronic synthesizer for Roxy Music and, as Pinch and Trocco (2002) make clear in their history of the synthesizer, it was not like any other instrument, especially in the early “analog days.” In the synthesizer, electronic waveforms are processed via various different modules, and the outputs of these can be fed back to control other modules with unforeseeable effects. As Eno wrote of the EMS synthesizer, for example, “The thing that makes this a great machine is that . . . you can go

from the oscillator to the filter, and then use the filter output to control the same oscillator again. . . . You get a kind of squiging effect. It feeds back on itself in interesting ways, because you can make some very complicated circles through the synthesiser” (quoted in Pinch and Trocco 2002, 294). Here we find the familiar cybernetic notion of a feedback loop, not, however, as that which enables control of some variable (as in a thermostat), but as that which makes a system’s behavior impenetrable to the user.⁷⁵ We can think about such systems further in the next chapter, but for now the point to note is that analog synthesizers were thus inescapably objects of exploration by their users, who had to *find out* what configuration would produce a desirable musical effect.⁷⁶ “The resulting music was an exchange . . . between person and machine, both contributing to the final results. This may be why analog synthesists can readily recount feelings of love for their synthesisers” (Pinch and Trocco 2002, 177). In this sense, then, one can see a continuity between Eno’s work with Roxy and his later work: even with Roxy Music, Eno was riding the dynamics of a generative system—the synthesizer—which he could not fully control. What he learned from Beer was to make this cybernetic insight explicit and the center of his future musical development.

Second, I want to emphasize that with Eno’s interest in cellular automata and complex systems we are back in the territory already covered at the end of chapter 4, on Ashby’s cybernetics, with systems that stage open-ended becomings rather than adaptation per se. Indeed, when Eno remarks of *It’s Gonna Rain* that “you are getting a huge amount of material and experience from a very, very simple starting point” (Eno 1996b, 3) he is singing the anthem of Stephen Wolfram’s “New Kind of Science.” In this sense, it would seem more appropriate to associate Eno with a line of cybernetic filiation going back to Ashby than with Beer—though historically he found inspiration in *Brain of the Firm* rather than *Design for a Brain*. We could also recall in this connection that no algorithmic system, in mathematics or in generative music, ever becomes in a fully open-ended fashion: each step in the evolution of such systems is rigidly chained to the one before. Nevertheless, as both Eno and Wolfram have stressed in their own ways, the evolution of these systems can be unpredictable even to one who knows the rules: one just has to set the system in motion and see what it does. For all practical purposes, then, such systems can thematize for us and stage an ontology of becoming, which is what Eno’s notion of riding the system’s dynamics implies.⁷⁷

Third, we should note that Eno’s ambient music sounds very different from the music we are used to in the West—rock, classical, whatever. In terms simply of content or substance it is clear, for instance, that three notes repeating

with different delays are never going to generate the richness, cadences, and wild climaxes of Roxy Music. Whatever aesthetic appeal ambient music might have—“this accommodates many levels of listening attention and is as ignorable as it is interesting” (Whittaker 2003, 47)—has to be referred to its own specific properties, not to its place in any conventional canon.⁷⁸ And further, such music has a quality of constant novelty and unrepeatability lacking in more traditional music. *In C* varies each time it is performed, according to the musicians who perform it and their changing preferences; *It’s Gonna Rain* depends in its specifics on the parameters of the tape players used, which themselves vary in time; the computerized system Eno described above is probabilistic, so any given performance soon differs from all others even if the generative parameters remain unchanged.⁷⁹ Perhaps the easiest way to put the point I am after is simply to note that Eno’s work, like Alvin Lucier’s biofeedback performances (chap. 3), raises the question, Is it music? This, I take it, again, is evidence that ontology makes a difference, now in the field of music. I should add that, evidently, Eno has not been alone in the musical exploitation of partially autonomous dynamic systems, and it is not the case that all of his colleagues were as decisively affected by reading the *Brain of the Firm* as he was. My argument is that all of the works in this tradition, cybernetically inspired and otherwise, can be understood as ontological theater and help us to see where a cybernetic ontology might lead us when staged as music.⁸⁰

Fourth, these remarks lead us, as they did with Beer himself, into questions of power and control. Usually, the composer of a piece of music exercises absolute power over the score, deciding what notes are to be played in what sequence, and thus exercises a great deal of power over musical performers, who have some leeway in interpreting the piece, and who, in turn, have absolute power over the audience as passive consumers. In contrast, “with this generative music . . . am I the composer? Are you if you buy the system the composer? Is Jim Coles and his brother who wrote the software the composer? Who actually composes music like this? Can you describe it as composition exactly when you don’t know what it’s going to be?” (Eno 1996b, 8). These rhetorical questions point to a leveling of the field of musical production and consumption. No doubt Eno retains a certain primacy in his work; I could not generate music half as appealing as his. On the other hand, the responsibility for such compositions is shared to a considerable extent with elements beyond the artist’s control—the material technology of performance (idiosyncratic human performers or tape players, complex probabilistic computer programs)—and with the audience, as in the case of computer-generated music in which the user picks the rules. As in the case of Beer’s social geometries,

a corollary of the ontology of becoming in music is again, then, a democratization, a lessening of centralized control, a sharing of responsibility, among producers, consumers, and machines.⁸¹

My fifth and final point is this. It is ironic that Eno came to cybernetics via Beer; he should have read Pask. The musical insights Eno squeezed out of Beer's writings on management are explicit in Pask's writings on aesthetics. As we can see in the next chapter, if Pask had handed him the torch of cybernetics, Eno would not need to have equivocated. Pask, however, was interested in more visual arts, the theater and architecture, so let me end this chapter by emphasizing that we have now added a distinctive approach to music to our list of instances of the cybernetic ontology in action.