The Science of the Unknowable:
Stafford Beer's Cybernetic Informatics

Andrew Pickering

Abstract

This essay explores the history of Stafford Beer's work in management cybernetics, from his early conception and simulation of an adaptive automatic factory and associated experimentation in biological computing up to his development of the Viable System Model of complex organizations and its implementation in Chile. The essay also briefly pursues Beer into the arenas of politics and spirituality. The aim throughout is to show that all Beer's projects can be understood as specific instantiations and workings out of a cybernetic ontology of unknowability and becoming: a stance that recognizes that the world can always surprise us and that we can never dominate it through knowledge. The thrust of Beer's work was thus to construct information systems that can adapt performatively to environments they cannot fully control.

This essay derives from a larger project exploring the history of cybernetics in Britain in and after World War II.[1] I focus here on Stafford Beer, the founder of the field he called management cybernetics, and his work in informatics. Beer died on 23 August 2002; he was a remarkable and admirable man, and I would like to dedicate this paper to his memory.[2]

Anthony Stafford Beer was born in London in 1926. In 1944, after one year as an undergraduate, he joined the British Army, serving in India and Britain. He left the army in 1949, and between 1949 and 1970 he worked in the steel and publishing industries and ran his own consulting company. From 1970 until his death he worked as an independent management consultant (*Times,* 2002) (Figures 1a and 1b).

I begin with an overview of Beer's general perspective on information science and information systems, intended to bring out the singularity of the cybernetic approach.

From the 1950s on Beer was a remorseless critic of the ways in which computers were being deployed in industry, essentially to replace existing paper systems. He felt that this did nothing to change existing organizational forms and that something more imaginative was required. His argument was that the postwar world was a new kind of world. Specifically the pace of change had increased markedly since the war, and the important thing for organizations was that they should be adaptive—light on their feet and ready to accommodate themselves to the new situations that would arise faster and faster as time went on.[3]

Rendering organizations adaptable, according to Beer, required reorganizing them to make possible specific patterns of information flow and transformation. I will turn to some examples shortly, but first I want to emphasize the gap between the mainstream vision of informatics and the cybernetic one. What we need to think about here is *ontology*—the question of what the world is like.

Mainstream informatics presumes a very familiar ontology. The world is a regular, lawlike place that can be known more or less exhaustively. It is a place that can therefore be controlled and dominated through knowledge. That is the logic behind the creation of bigger databases and faster information systems. Of course, this ontology does recognize the existence of the unknown,

1 See also Andrew Pickering (2002, forthcoming). During 2002–2003 this project was supported by the National Science Foundation under grant no. SES-0094504, and I was based at the Science Studies Unit, Edinburgh University. I thank the unit director, David Bloor, and all its members for their hospitality.

2 Beer provided me with a considerable amount of information before his death. I also thank Beer's partner, Allenna Leonard, and his daughter, Vanilla Beer, for very important assistance and encouragement in my research.

3 See, for example, Beer (1959)—his first book.

Figure 1a. *Stafford Beer in the early 1960s. Source: Harnden & Leonard (1994, p. xii). Reproduced by permission of John Wiley & Sons and Allenna Leonard.*

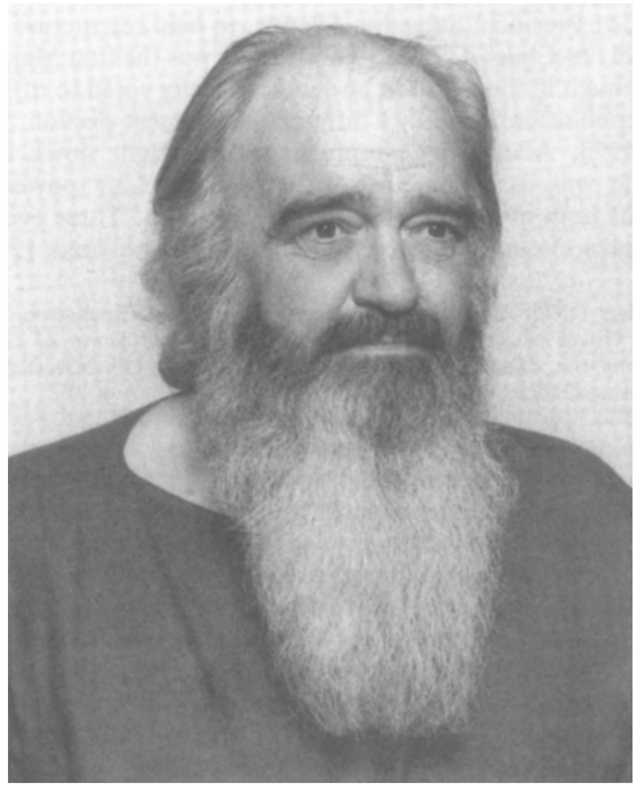

Figure 1b. *Stafford Beer in 1975. Source: Harnden & Leonard (1994, p. 315). Reproduced by permission of John Wiley & Sons and Allenna Leonard.*

but only as something to be conquered, to be drawn into the realm of the known.

Cybernetics turned this picture inside out and exemplified a different and much less familiar ontology. Beer (1959, p. 17) argued that there exists in the world a class of "exceedingly complex systems," including the brain, the firm, and the economy, which are in principle unknowable. However much data we gather on them, we can never know them completely; they can always surprise us. Such systems can never be dominated by knowledge, and instead we have to learn somehow to cope with them. Cybernetics was thus a science of dealing with the unknown—an extremely odd sort of science.[4]

This ontology of the unknowable is the key thing to grasp in thinking about cybernetics, and two corollaries of it are worth mentioning. First, it thematizes time. By definition one has to deal with the unknown in time, as it happens. No amount of information about the past, as stored in conventional information systems, can ever prepare us for genuine unpredictable novelty. As Beer (1972, p. 199) put it: "Look straight ahead down the motorway while you are driving flat out. Most enterprises are directed with the driver's eyes fixed on the rear-view mirror." This contrast between real-time awareness and retrospection is an important angle on the specificity of cybernetics.

The other corollary is that conventional informatics is, as I would say, representational, meaning that it is all about the accumulation of data and knowledge. One might eventually want to draw on that knowledge for action, but that is not the defining feature of an information system. The information system is, as it were, detachable from the action. Cybernetics viewed information systems differently. If one has continually to deal with the unexpected as a practical matter, then the

[4] On the affinity between the cybernetic ontology and that which I arrived at in Pickering (1995), see Pickering (2002). My particular interest in the history of cybernetics is to see how this ontology engaged in a wide variety of real-world projects.

accumulation of representational knowledge seems less relevant. What one wants instead is a performative information system, geared straight into the action, not detachable at all. One would not care exactly what information was flowing through the system and how, as long as its output was an adaptive transformation of the organization to its environment. This contrast between the representationalism of conventional information systems and the performativity of cybernetic ones is very important. In 1962, in one of his more visionary moments, Beer described electronic computers as dinosaurs, looking forward to the day when they would be supplanted by another class of information-processing devices that simply would not have representational intermediate states at all (Beer, 1962b, quote in Harnden & Leonard, 1994, p. 220). When I came across that idea I was amazed. Something beyond the computer? What is this man talking about? Is he mad?

The Homeostat

Beer was not mad. Now we can turn to history, starting with a brief detour through the work of Beer's friend, W. Ross Ashby (1903–1972), the doyen of the English cyberneticians (Figure 2). We need to think especially about a device Ashby built in his spare time in 1948, his famous homeostat (Ashby, 1948, 1952). This was an electromechanical device intended to mimic the biological process of homeostasis—the ability of organisms to maintain such "essential variables" as blood temperature constant in the face of fluctuations in their environment (Figure 3). Without going into details, in Ashby's homeostat the essential variable was the electric current flowing through a movable needle dipping into a trough of water on top of the device, and the machine's environment was constituted by electrical interconnections to other homeostats. The trick in maintaining homeostasis was that when the current within a given homeostat went beyond some preassigned limit, a relay would trip, operating a stepping switch that would change the electrical resistance of the homeostat's inner circuitry, with the sequence of different values for the resistance being determined from a table of random numbers. The homeostat would thus randomly reconfigure itself. If the current were to continue to go beyond its limit, the machine would reconfigure itself again, and so on until homeostasis was achieved. The homeostat was thus, as Ashby called it, an ultrastable device; whatever one did to it, it would eventually find its way back to homeostatic equilibrium. It was a device for staying the same. Another British cybernetician, Grey Walter (1953), sarcastically referred to it as "Machina Sopora."

I need to make three remarks on the homeostat. First, I hope it is clear how it fits in with my earlier remarks on ontology. The homeostat was a device that dealt with the unknown. It did this in real time: it reacted to fluctuations in its environment as they happened. And it did so in a performative rather than a representational fashion: it did not seek to know the world representationally; it simply materially reconfigured itself as the occasion arose. If you understand that, then you understand what was most distinctive about British cybernetics. If orreries—those beautiful early-modern models of the solar system—were the mechanical emblems of the ontology of the knowable, then the homeostat was the emblem of the cybernetic ontology of unknowability.

Second, the homeostat was the centerpiece of Ashby's

Figure 2. *W. Ross Ashby. Source: I thank Ashby's daughters, Jill Ashby, Sally Bannister, and Ruth Pettit, for providing me with this photograph and for permission to reproduce it.*

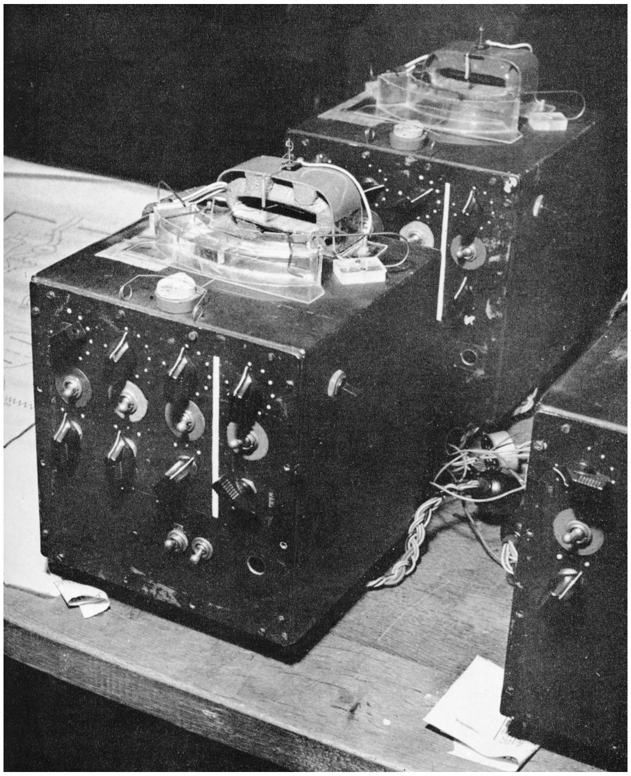

Figure 3. *The homeostat: partial photograph of four interconnected homeostat units. Source: de Latil, P., (1956),* Thinking by machine: A study of cybernetics *[London: Sidgwick and Jackson].*

first book, *Design for a Brain* (1952), and Ashby intended it as a model of the brain inasmuch as it learned to cope with its environment. But it was a performative brain—as distinguished, for example, from the rational representationalist brain that was later exemplified in symbolic artificial intelligence.

Third, we can think of the homeostat as a controller, and much of Ashby's cybernetics focused precisely on questions of control (Ashby, 1956). His key result was the law of requisite variety—Ashby's law, as Beer called it. Variety is a measure of the number of states a system can take up—twenty-five in the case of Ashby's first homeostats, with their different possible internal electrical resistances. The law of requisite variety stated that a system could succeed as a homeostatic controller only if it disposed of as much variety as the environment in which it existed. A homeostat could maintain its ultrastable condition when connected to another homeostat with the same number of internal states, but it might fail against one having twice that number.

The Cybernetic Factory

That is enough to get us back to Stafford Beer. If Ashby practiced cybernetics as a "pure science," then Beer was an applied cybernetician, which is what interests me. I am especially interested in what cybernetics looked like when it was put to work in the world. Ashby's homeostat was at the heart of all Beer's attempts to conceptualize and design adaptive organizations. Now I will discuss some of these attempts as they emerged in the period from the 1950s to the 1970s.

In the 1950s Beer's cybernetics revolved around the contemporary fantasy of the "automatic factory," in which all operations were to be controlled by computers rather than people. Beer likened current visions of the automatic factory to a "spinal dog"—a dog whose nervous system had been severed below the brain (1962b, quote in Harnden & Leonard, 1994, p. 164). Such an animal can apparently continue to live and display bodily reflexes; it is just that it has no brain. To move from the automatic factory to the cybernetic factory thus required adding brain, and this, Beer argued, should be an Ashbean homeostat.

Since there is not enough space to go into much detail here, I will instead discuss some figures from a major paper Beer wrote in 1960 (Beer, 1962b). Figure 4

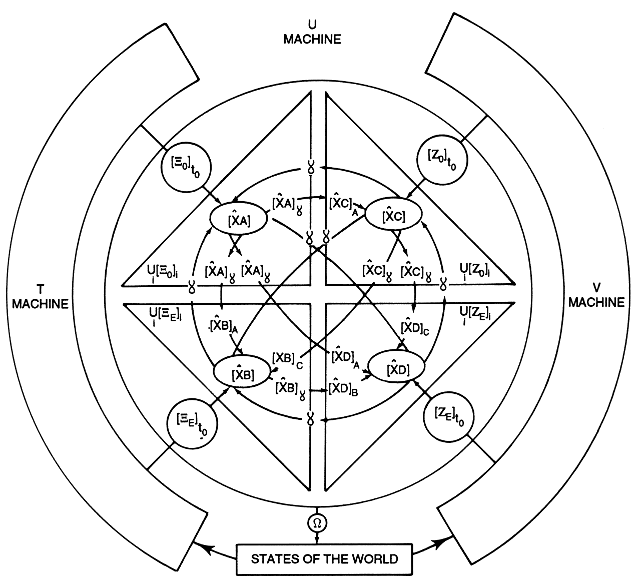

Figure 4. *The cybernetic factory. Source: Harnden & Leonard (1994, p. 192). Reproduced by permission of John Wiley & Sons and Allenna Leonard.*

is a logic diagram of the cybernetic factory. The T- and V-machines are what we would now call neural nets: the T-machine collects data on the state of the factory and its environment and translates them into meaningful form; the V-machine reverses the operation, issuing commands for action in the spaces of buying, production, and selling. Between the T- and V-machines lies the U-machine—the homeostat, or artificial brain—which seeks to find and maintain a balance between the inner and outer conditions of the firm, trying to keep the firm operating in a livable segment of phase-space. Figure 5 is a more suggestive figure, a painting by Beer, labeled "general picture of the whole theory" (the T-, U-, and V-machines are indistinctly labeled in the smaller painting at the lower left).

The cybernetic factory was not pure theory. By 1960 Beer had at least simulated a cybernetic factory at Templeborough Rolling Mills, a subsidiary of his employer, United Steel, and the next figure might help us better understand things. In Figure 6 the lines of circles and squares marked "sensation" to "judgments" correspond to the numerical inputs to the T-machine: for example, "tons bought," "cost of raw material," "cash at bank," and "value of credits." At Templeborough all these data were statistically processed, analyzed, and transformed into twelve variables, six referring to the inner state of the mill and six to its economic environment. Values of these variables were generated at the mill every day—as close to real time as one could get—and each day's numbers were stored as the "generalised gestalt memories" indicated at the lower left and right of Figure 6. Beer claimed to see how all this data collection and processing, including changes in the classification system, could be accomplished automatically, although in fact it was still done clerically in the mill according to protocols devised by operations research scientists: this was one sense in which the mill had become a simulation of a fully cybernetic factory.

The other sense of simulation concerned the U-machine. As indicated in the lower center of the figure, the two gestalt memories of the factory defined two phase-spaces in terms of the relevant parameters, and the job of the U-machine was to strike a homeostatic balance between them. But nothing like a functioning U-machine had yet been devised. The one at Templeborough was still constituted by the decisions of human managers, although now they were precisely positioned

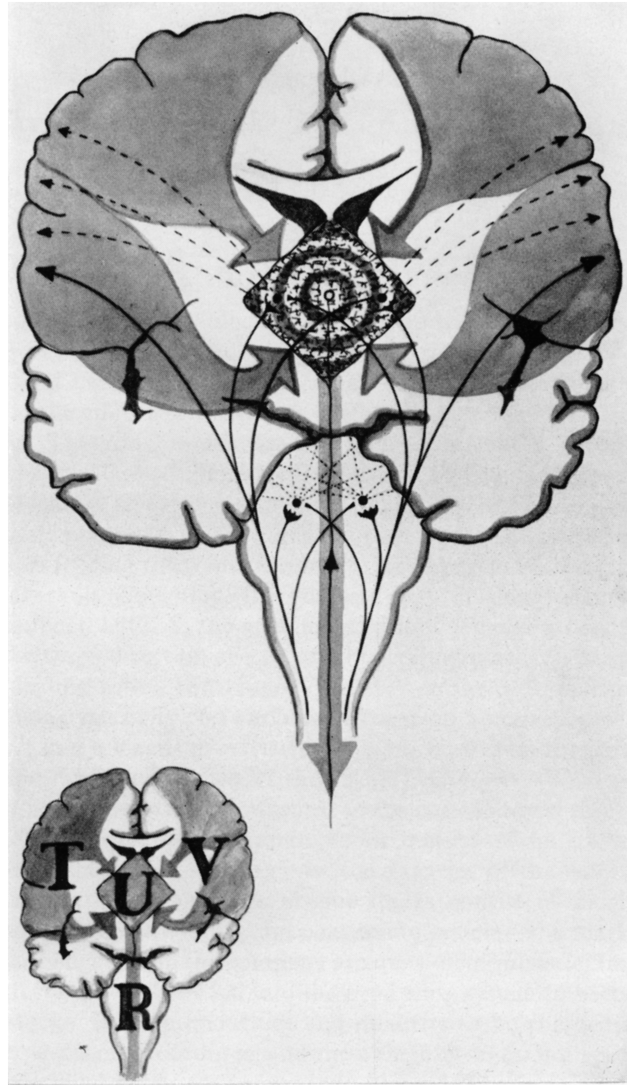

Figure 5. *The factory as brain. Source: Harnden & Leonard (1994, p. 198). Reproduced by permission of John Wiley & Sons and Allenna Leonard.*

in an information space defined by the simulated T- and V-machines.

So by 1960 Beer had constructed a simulation of a cybernetic factory that promised to dispense entirely with human personnel, though humans in fact still filled the gaps for machines not yet in place. Beer could see how to complete the automatic T- and V-machines, but the U-machine remained unspecified. Nevertheless, he wrote: "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)" (1962b, quote in Harnden & Leonard, 1994, p. 212).

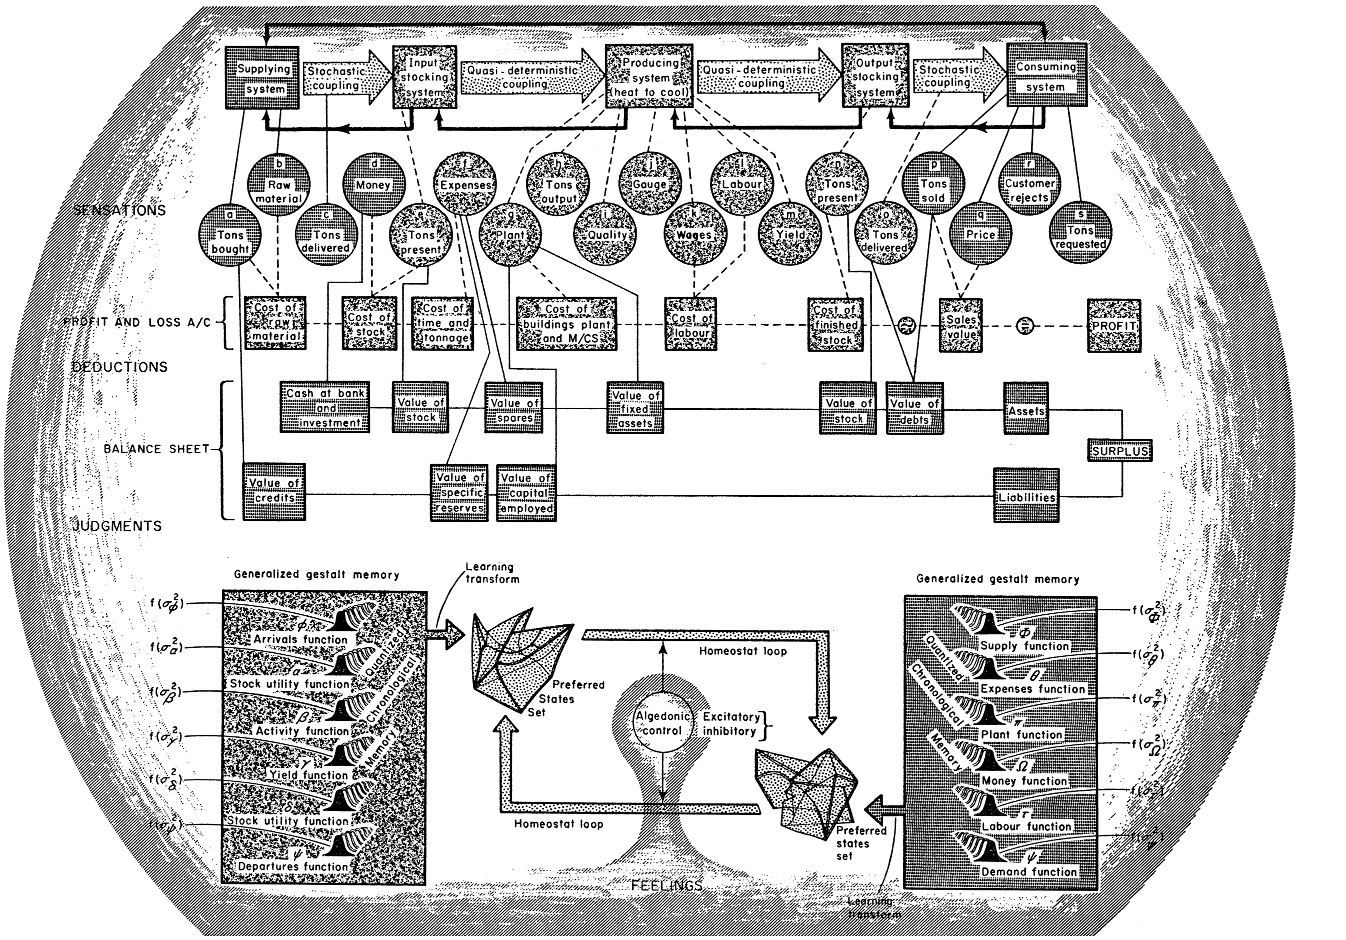

Figure 6. *Simulation of a cybernetic factory. Source: Harnden & Leonard (1994, pp. 200–201). Reproduced by permission of John Wiley & Sons and Allenna Leonard.*

Biological Computing

The vision of the adaptive factory not just running smoothly but also evolving and changing on its own without any human intervention is itself amazing, but Beer's attempts to construct the U-machine homeostat are where the story gets really interesting. The requirements for the U-machine were that first it should be able to internally reconfigure itself, as with Ashby's original homeostat, and that second, in accordance with Ashby's law, it must have high variety in order to have a chance of coping with the complexity of its environment. Nowadays we might think of somehow programming a computer to fulfill this function, but Beer argued that this was not necessarily the way to go. Computers were extremely expensive in the 1950s and 1960s. And besides Beer had come up with a different idea:

> 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. (Beer, 1962b, quote in Harnden & Leonard, 1994, pp. 209, 215)

What is all this about? Ashby had built an electromechanical equivalent of a homeostatic biological system and called it a brain. Beer's idea was to turn Ashby's idea another 180 degrees: he wanted somehow to enroll a naturally occurring homeostatic system as the brain of the cybernetic factory. He had conceived the idea of a nonrepresentational, adaptive, biological computer. This was the machine he hoped would supersede the electronic computer, the referent of his remark about dinosaurs. And during the second half of the 1950s he embarked on "an almost unbounded survey of naturally occurring systems in search of materials for the construction of cybernetic machines" (Beer, 1959, p. 162).

In 1962 he wrote a brief report on the state of the art that makes fairly mind-boggling reading (Beer, 1962a). The list of systems he discussed includes a successful attempt to use positive and negative feedback to train young children to solve simultaneous equations without teaching them the relevant mathematics, and it goes on to discuss an extension of the same tactics to mice. This is, no doubt, the origin of the mouse computer that turns up in both Douglas Adams's *Hitchhiker's Guide to the Galaxy* and Terry Pratchett's *Discworld* series. Beer also reported attempts to induce small organisms—*Daphnia* collected from a local pond—to ingest iron filings so that input and output couplings to them could be achieved via magnetic fields, and he made another attempt to use a population of the protozoan *Euglena* via optical couplings. Beer's last attempt in this series was to use not specific organisms but an entire pond ecosystem as a homeostatic controller. He reported that "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 state of this research at the moment," he said in 1962, "is that I tinker with this tank from time to time in the middle of the night" (1962a, quote in Harnden & Leonard, 1994, p. 31).

In the end this wonderful line of research foundered, not on any point of principle but on Beer's practical failure to achieve a useful coupling to any biological system of sufficiently high variety. However, Beer's imagination in this phase of his work is highly admirable, and it is clear from subsequent developments that the homeostatic system Beer really had in mind was the human spinal cord and brain. He never mentioned this in his work on biological computers, but the image that sticks in my mind is that the brain of the cybernetic factory should really have been an unconscious human body, floating in a vat of nutrients, with electronic readouts tapping its higher and lower reflexes—something vaguely reminiscent of the movie *The Matrix.* This horrible image helps me at least to appreciate the magnitude of the gap between cybernetic information systems and more conventional approaches.

The Viable System Model

Now we can return to something more like normality. Beer's dreams of biological controllers came to an end in the early 1960s, but this provoked a transformation rather than an abandonment of his vision of the cybernetic factory. His 1972 book, *Brain of the Firm,* laid out a new vision of what he called the Viable System Model—VSM for short.[5]

The VSM took up Beer's earlier plan for a cybernetic factory and transformed it along two axes. First, the simulation of the cybernetic factory discussed above became in effect the thing itself. Beer dropped the ambition to dispense entirely with human beings and instead argued that human managers should be positioned within purposefully designed information flows at just those points that would ideally have been occupied by homeostatic ponds or trained mice. Second, Beer extended and elaborated his conception of information flows considerably. The aim of the firm had to be to survive in an environment that was fluctuating and changing. How was this now to be accomplished? The place to look for inspiration, according to Beer, was once again nature. Biological organisms have already mastered the trick of survival and adaptation, and Beer's idea was therefore to read biological organisms as exemplary of the structure of viable systems in general: we should transplant their key features to the structure of the firm. In particular, Beer chose the human nervous system as his model. If his original idea was that the firm needed to contain an artificial brain (made of magnetic *Daphnia* or leeches), the idea of the VSM was that the firm should become a brain, a cyborg brain with human brains lodged within it. (Another weird image, if you think about it too hard.)

The spirit of the VSM is nicely expressed in the juxtaposition of two figures from *Brain of the Firm:* one a schematic of the human body; the other of the firm. Very briefly, Beer argued that one needs to distinguish at least five levels or systems of control in any viable system. In Figure 7a, system 1 consists of four subsidiaries of a larger organization, labeled A, B, C, and D, analogous to arms and legs, the heart, kidneys, and so on (Figure 7b). System 2, the equivalent of the sympathetic nervous system, connects them to one another and to system 3 and seeks to damp out destructive interactions between the subsidiaries. System 3—the pons and

[5] This book was significantly extended in its second edition (Beer, 1981) and eventually formed part of a trilogy with Beer's 1979 and 1985 books. For more on the VSM see R. Espejo and R. Harnden (1989).

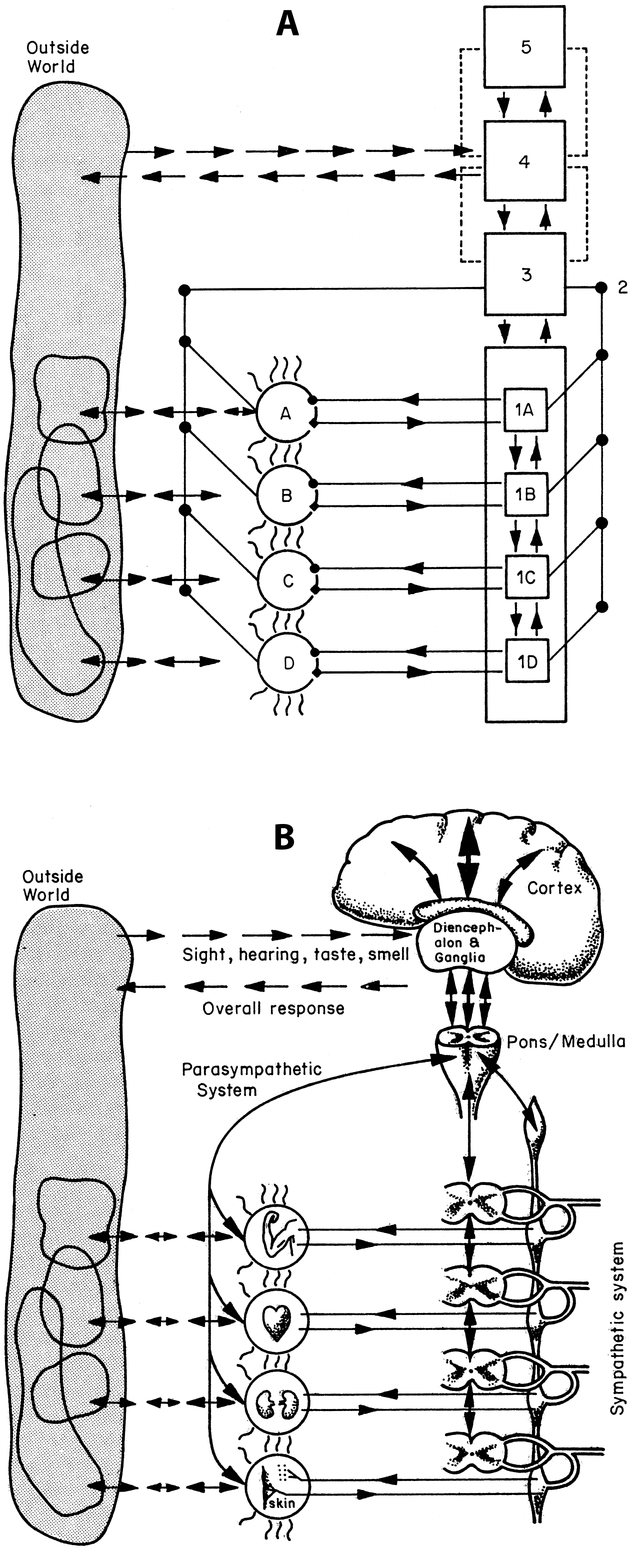

Figure 7. *Control systems in (a) the firm, and (b) the human body. Source: Beer (1981, pp. 130–131). Reproduced by permission of John Wiley & Sons and Allenna Leonard.*

medulla of the VSM—consists of a set of operations research models of production that enables management to react to fluctuations in systems 1 and 2—by reallocating resources, for example. System 4—the base of the brain itself—was envisaged as a decision-making environment for higher management, modeled on the World War II operations room. It would collect and display information from the lower systems and from the outside world, and, very important, it would run a set of computer programs that higher management could consult on the possible future effects of major decisions. At the same time the operations room was intended to function as a club room for senior management—a place to hang out, even when major decisions were not at stake. Finally system 5 was the location of the most senior management whom Beer regarded as the cortex of the firm. Their vision of the firm and its future, whatever it was, was to be negotiated into reality in reciprocally vetoing homeostatic interactions with system 4.

Here I should return to the question of ontology. Despite my earlier emphasis on performance versus representation, it is clear that the VSM did incorporate significant representational elements, especially the computer models running in systems 3 and 4. But one function of the programs running at the system 3 level was statistical filtration—that is, to junk almost all the information that arrived there rather than to store it. And in the VSM the models at levels 3 and 4 were to be continually updated in comparisons between their predictions and the actual performance of the firm and its environment. This updating recognized even in the realm of representation that the world remained an unknowable place; the utility of the models had continually to be found out in real-time experience.

The VSM in Action

The VSM was not a theoretical conceit. Beer's consulting work was based on it, and in the early 1980s he could list among his clients small businesses and large industries, publishers, banks and insurance companies, transportation, education and health organizations, and governments and international agencies (Beer, 1989, pp. 34–35). I will not discuss examples here, but I can note that in much of this work the VSM functioned as a diagnostic tool: comparison with the VSM diagram was a way of singling out organizational problems that needed to be addressed. Only on one major occasion did Beer have a chance to implement the VSM from

Figure 8. *Operations room of Project Cybersyn. Source: Beer, S., (1974), Cybernetics of national development, The Zaheer Foundation Lecture, New Delhi, India, in R. Harnden & A. Leonard (Eds.),* How many grapes went into the wine? Stafford Beer on the art and science of holistic management *(p. 330) [New York: Wiley]. Reproduced by permission of John Wiley & Sons and Allenna Leonard.*

the ground up—when he was invited to help design and implement a control system for the entire economy of Chile, under the newly elected Marxist regime led by Salvador Allende. From 1971 to 1973 Beer threw himself into Project Cybersyn, as it was called (for "cybernetic synergy"); a lot was done in a very short period of time, and I can only summarize what was accomplished.[6]

By requisitioning telex facilities, a real-time communication network called Cybernet was established, linking much of Chile's industrial base to computers in Santiago. A set of programs called Cyberstride was written to process and filter the incoming data at the system 3 level, and another program, CHECO, was written to simulate the overall behavior of the Chilean economy at the system 4 level. The system 4 operations room was also shaping up by 1973, as shown in Figure 8. This cybernetization of the Chilean economy was an extremely ambitious project, which unfortunately never had a chance to go into full operation. On 11 September 1973 General Augusto Pinochet launched a successful coup against the Allende government, and Allende himself was assassinated. Some members of Beer's group fled the country, while others were jailed.

Politics and Spirituality

That is almost as far as I can take the story of Beer's cybernetic informatics in this essay, although there is much more that could be discussed. It would be inappropriate, for example, to close without some mention of Beer's politics, or subpolitics as it should probably be called. It hinged on Ashby's idea that only variety can control variety, an idea that Beer felt Western governments had failed catastrophically to appreciate. The last example of Beer's writing I have seen is dated October 2001. "Last month," he wrote, "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 ways." Low-variety understandings of those planes flying into the World Trade Center and the Pentagon were already feeding back into low-variety responses, aimed this time at "regime-change" in Afghanistan rather than Chile. But "how actually to deal with the crisis provokes further disquiet. Attempts to guard against an infinite number of threats do not have Requisite Variety." This strikes me as a prescient observation, given that while I was drafting this essay (October and November 2002) several hundred people were killed by a bomb in Bali, the Russians killed about a hundred people to end the Chechen siege in Moscow, and a man with a gun picked off about a dozen people at random around Washington, D.C.—all this against preparations for a massive invasion of Iraq.

One should also mention Beer's spirituality—starting as an Anglican, converting to Catholicism for twenty-four years, before ending his life as a self-described Tantric Yogi—and how that fitted in with his cybernetics and the ontology of the unknowable. One might also discuss his painting and poetry. The upshot of such an extended account would be to situate Beer's cybernetic informatics within a very broad assemblage, spanning management, religion, politics, philosophy, engineering, the arts, brain science, and neurophysiology—all as instantiations in different realms of an ontology of the unknowable. One might then even argue that this cybernetic assemblage indicates a path to challenging the hegemony of the modern ontology of the knowable, not just in informatics but far beyond it. But that is a topic that must remain to be explored elsewhere.

6 The second edition of *Brain of the Firm* (Beer, 1981) includes a long history and discussion of the Chilean project.

References

Ashby, W. R. (1948). Design for a brain. *Electronic Engineering, 20,* 379–383.

Ashby, W. R. (1952). *Design for a brain: The origin of adaptive behaviour.* New York: Wiley.

Ashby, W. R. (1956). *An introduction to cybernetics.* New York: Wiley.

Beer, S. (1959). *Cybernetics and management.* London: English Universities Press.

Beer, S. (1962a). A progress note on research into a cybernetic analogue of fabric. *Artorga, Communication 40* (April). Reprinted in R. Harnden & A. Leonard (Eds.). (1994). *How many grapes went into the wine? Stafford Beer on the art and science of holistic management* (pp. 24–32). New York: Wiley.

Beer, S. (1962b). Towards the automatic factory. In H. von Foerster & G. Zopf (Eds.), *Principles of self-organization: Transactions of the University of Illinois Symposium on Self-Organization, Robert Allerton Park, 8 and 9 June, 1961* [sic: actually 1960] (pp. 25–89). New York: Pergamon. Reprinted in R. Harnden & A. Leonard (Eds.). (1994). *How many grapes went into the wine? Stafford Beer on the art and science of holistic management* (pp. 162–225). New York: Wiley.

Beer, S. (1972). *Brain of the firm.* London: Penguin.

Beer, S. (1979). *The heart of the enterprise.* New York: Wiley.

Beer, S. (1981). *Brain of the firm* (2nd ed.). New York: Wiley.

Beer, S. (1985). *Diagnosing the system for organizations.* New York: Wiley.

Beer, S. (1989). The viable system model: Its provenance, development, methodology and pathology. In R. Espejo & R. Harnden (Eds.), *The viable system model: Interpretations and applications of Stafford Beer's VSM* (pp. 11–37). New York: Wiley. Reprinted from Beer. (1984). *Journal of the Operational Research Society, 35,* 7–26.

Beer, S. (2001). What is cybernetics? Acceptance speech for an honorary degree at the University of Valladolid, Mexico, October 2001. Unpublished.

Espejo, R., & Harnden, R. (Eds.). (1989). *The viable system model: Interpretations and applications of Stafford Beer's VSM.* New York: Wiley.

Harnden, R., & Leonard, A. (Eds.). (1994). *How many grapes went into the wine? Stafford Beer on the art and science of holistic management*. New York: Wiley.

Pickering, A. (1995). *The mangle of practice: Time, agency, and science.* Chicago: University of Chicago Press.

Pickering, A. (2002). Cybernetics and the mangle: Ashby, Beer and Pask. *Social Studies of Science, 32,* 413–437.

Pickering, A. (forthcoming). The tortoise against modernity: Grey Walter, the brain, engineering and entertainment. In *Experimental cultures: Configurations between science, art, and technology, 1830–1950* (preprint 213, pp. 109–122). Berlin: Max Planck Institute for the History of Science. To appear in a volume edited by H. Schmidgen & H.-J. Rheinberger.

Times [London]. (2002, 9 September). Obituary: Stafford Beer. Available: www.timesonline.co.uk/article/0,,60-408258,00.html.

Walter, W. G. (1953). *The living brain.* London: Duckworth.