5

The Engineering of Purpose: From Water Clocks to Cybernetics

He devoted himself to alchemy, in which he claims to have uncovered miraculous things, and inventions of wonderful furnaces, among them one that will maintain the fire at any degree of heat desired, whether hotter or colder.

—N. C. Fabri de Peiresc (1624; referring to Cornelis Drebbel's thermostatic furnace; translated in Mayr 1970, p. 56)

Let us consider a car following a man along a road with the clear purpose of running him down. What important difference will there be in our analysis of the behavior of the car if it is driven by a human being, or it is guided by the appropriate mechanical sense organs and mechanical controls?

—Arturo Rosenblueth & Norbert Wiener (1950, p. 319)

The Use and Understanding of Feedback Control

Although Bernard and Cannon recognized the self-regulatory nature of the living systems they studied, an explicit, formal understanding of such systems did not develop from this physiological research but rather had to await the attempts of engineers to make purposefully behaving machines using what is now called *feedback control*.

Devices making use of feedback control go back at least as far as the Hellenistic period (Mayr 1970). The first documented device was designed by Ktesibios, a barber and mechanic living in Alexandria during the third century B.C. when that north African city was the scientific and intellectual center of the world (Euclid, Archimedes, and Eratosthenes were just three of Ktesibios's fellow Alexandrians whose names students of astronomy and mathematics will recognize).

Ktesibios's water clock required a steady, unvarying flow of water to measure accurately the steady, unvarying flow of time. But because water flows more quickly from a full container and more slowly when it is less full, Ktesibios had to devise a way to keep the vessel at a constant level while water was flowing from it into the clock mechanism. As he did this in a manner not unlike that of the modern flush toilet to which it is assumed the reader has handy access, I will use this more modern invention instead of the water clock as our first example of a feedback-control device.

The modern flush toilet must have a certain amount of water on hand for each flush to be effective. For this purpose, most residential toilets make use of a holding tank into which water accumulates between flushes. Since too little water in the tank does not allow adequate flushing and too much is wasteful (it will simply flow out through an overflow drain), a mechanism is used to maintain the water at the desired level. This mechanism consists of a float resting on the surface of the water that is connected to a valve. When the water level falls after a flush, the float falls with it and in so doing opens a valve, admitting water into the tank. But as the tank fills and the water level rises, so does the float, eventually closing the valve so that the tank does not overfill.

For the reader who has not already peered inside a flush toilet tank, it is well worth lifting the lid and taking a look. With the tank lid off and the flush lever activated, one can observe in live action the events described: the tank empties, the float falls, the valve turns on, the tank refills, and the valve shuts off. It is also informative to push lightly on the flush lever for a few seconds so that just a portion of the water in the tank escapes into the bowl. This will show that the tank need not be emptied completely before the float valve mechanism acts to refill the tank. If all is operating properly, the float-valve mechanism will not let the water remain very much below the desired level.

What is this desired level? Inside most tanks a line indicates the optimal amount of water for flushing the toilet. If the water level in your tank is above or below this line, it can be changed by adjusting the float's position on the link that connects it to the valve. By changing the distance between the float and the valve, you can control the water level that will be reached before the valve turns itself off.

Notice the phrase I used in the preceding sentence—"the valve turns itself off." Is this actually the case? Isn't it rather that the rising float causes the valve to close? Yes, of course. But what is it that causes the float to rise? Obviously, the water that is filling the tank. And why is the water entering the tank? Because the valve is open. And what will cause the valve to close? The rising water level. So the valve, through a series of events, does in a sense close itself, since the valve's opening eventually causes it to close again.

If it seems that we are going around in a circle here, it is because we are. All feedback-control devices make use of what is called a *feedback loop*, meaning that the effect the device has on its environment is *fed back* to the device. In the case of the toilet tank, the falling of the float causes the valve to open, but the resulting inflow of water causes the float to rise again. So the action of the float is fed back to itself, having the consequence that the float simultaneously *affects* the water level and is *affected by* the water level. And since a low water level results in opening the valve, which raises the water level, this is called a *negative*-feedback system. This contrasts with a *positive*-feedback system, which tends to drive itself to extremes, as when a microphone is placed too close to an amplifier's loudspeaker, resulting in an annoying howl or squeal as sounds are continuously amplified, picked up by the microphone, and reamplified. A positive-feedback toilet tank (if such a useless thing existed) would be one that filled itself when it already had too much water. Since all positive-feedback devices drive themselves to extremes, they cannot be used alone to establish control and so cannot be referred to as feedback-*control* systems (although it is possible to establish certain kinds of control by using a negative-feedback system to control a positive-feedback one).

All feedback control must therefore ultimately rely on negative feedback. We can see now why such a system is called a feedback-control device, since the effect (feedback) of the environment on the device is controlled by the device itself. The operation of the feedback loop should also make it clear that a type of *circular causality* is involved that is quite unlike the one-way, push-pull causality characteristic of physical objects and systems not organized as feedback-control systems.

The usefulness and convenience of the toilet tank feedback-control system becomes more apparent when the system malfunctions. If the valve no longer opens when the water level drops, the human user must then refill the tank manually after each use, taking care not to add too much or too little water. It can be appreciated that the float valve provides a very convenient form of automation that replaces irksome human labor.

Many other feedback-control devices have been designed and used since Ktesibios's water clock, from a Byzantine oil lamp of the third century B.C. that automatically maintains a proper level of oil for burning, to the "fantail" used in eighteenth-century England and Scotland to keep windmills facing the wind. But the device that first attracted worldwide attention and use was the speed governor for steam engines invented in 1788 by Scottish engineer and inventor James Watt (1736–1819).

The invention of the steam engine marked a turning point in human history since it provided a source of mechanical power that for the first time did not depend on the vagaries of wind or water, or the muscles of human or beast. But one problem with the early steam engine was that its speed was sensitive both to the amount of steam pressure generated in the boiler and to the work load placed on the engine. Watt's ingenious solution was to make use of a combination of centrifugal force and gravity acting on a pair of metal balls (called flyweights) spinning on each side of a vertical rotating shaft so that if the speed of the engine increased, the flyweights would spread apart due to centrifugal force. This operated a valve that decreased the flow of steam to the engine so that the slower speed would be restored. If instead the engine's speed decreased, the centrifugal force acting on the flyweights would decrease so that they would be pulled down by gravity, thereby increasing the amount of steam delivered to the engine. In this way, the engine's speed remained constant in spite of fluctuating steam pressure and work loads without requiring a human operator to monitor it and attempt to keep it constant by manually operating a steam valve or changing the amount of heat applied to the boiler. The negative nature of this feedback control is apparent since anything that would tend to decrease the engine's speed would result in an increase in steam delivered to the engine, thereby keeping its speed constant, whereas anything that would tend to increase the speed would result in a decrease in steam delivered to the engine, thereby maintaining its speed.

An early important application of feedback control to electrical systems was achieved by Harold S. Black, an engineer for Bell Laboratories in New Jersey. Black had been wrestling with the problem of designing amplifiers for a transoceanic telephone system. In 1927 he figured out how to use negative feedback to amplify telephone signals by a known amount in undersea cable amplifiers using vacuum tubes that aged and lost amplification year by year and had to be placed on the ocean floor where they were needed to function for perhaps twenty years without maintenance. Black achieved this by building amplifiers with much more amplification than required and then "throwing away" most of it by using negative feedback. The result was an amplifier whose characteristics were almost immune to changes in the vacuum tubes. As a bonus, the fidelity of amplification was greatly increased, changes in available electrical power had practically no effect on the telephone signal, and noise generated in the electronic circuits was markedly reduced relative to the signal (see Bode 1960 for details).

Black's electronic invention used different components from those in the mechanical control systems described above, but the two kinds of systems—the telephone amplifier with negative feedback and the electromechanical negative-feedback control devices—share fundamental similarities, and the same basic laws govern both. In addition, the practice of using schematic diagrams for designing electrical circuits made it clear to Black and other engineers just how feedback-control devices operated: through a feedback loop the system's varying output was used to control its input.[1]

The Birth of Cybernetics

Once the general principles of feedback control were understood, control systems (as engineers refer to them) found widespread use in engineering for automatically controlling processes that were previously not possible or that would otherwise require a constantly attentive human operator. And this brings us back to Walter Cannon, or rather to one of his associates, Mexican physiologist Arturo Rosenblueth.

Rosenblueth, who learned to appreciate the self-regulating nature of living physiological processes through his work with Cannon at Harvard, met and collaborated with MIT mathematician Norbert Wiener and engineer Julian Bigelow. Rosenblueth was knowledgeable about living physiological systems, and Wiener and Bigelow were familiar with new developments in engineering, having developed negative-feedback systems during World War II for aiming antiaircraft guns at enemy airplanes.

They realized that for a machine to behave as a human operator would, it had to be goal directed, and this could be achieved only by designing it as a negative-feedback-control system. This design constraint provided an important clue about the organization and behavior of living organisms. In their influential 1943 paper "Behavior, Purpose, and Teleology," the three men were the first to establish a clear link between animate behavior and that of feedback-control systems designed by engineers. In addition, they maintained that purposeful behavior, whether that of human or machine, did not require the usual impossible teleological assumption of a future cause having a present effect. Instead, purposeful behavior could be explained by present causes having present effects, although now with causation acting in a circular manner.

Pursuing these ideas further, Wiener published a groundbreaking book in 1948, *Cybernetics*, that promised to revolutionize the study of animal and human behavior. In *Cybernetics* (revised in 1961), Wiener continued his application of the principles of feedback control to living organisms and in so doing developed the first formal, mathematical analysis of the types of self-regulatory systems that Bernard and Cannon studied.

But Wiener went beyond physiology. One way of appreciating the breadth of his cybernetic work is to recall Cannon's division of the nervous system into inward-acting (autonomic, involuntary) and outward-acting (somatic, voluntary) systems. Cannon, like Bernard, realized that the function of the autonomic system was to ensure a stable internal environment, maintaining vital conditions such as blood pressure (by varying heart rate and blood vessel constriction and dilation), blood oxygen concentration (by varying respiration), and body temperature (by varying the rate of metabolism and by initiating perspiration or shivering). Cannon, being a physiologist and not a behavioral scientist, was not particularly interested in the function of the somatic or outgoing nervous system, the one that innervates muscles attached to limbs permitting locomotion and other voluntary actions on the external environment. But if the purpose of the autonomous, involuntary nervous system is to control the organism's *internal* environment, why not at least consider the possibility that the purpose of the somatic, voluntary nervous system is to control the organism's *external* environment?

This is essentially what Wiener proposed. Indeed, the word *cybernetics* can be roughly translated from its Greek origin as "steersmanship," referring to the process of steering a ship on a course to a desired destination. Recognition that such behavior was purposeful and was used to control aspects of an organism's external environment (in much the same way as physiological functions controlled aspects of an organism's internal environment) promised a radically new foundation for understanding animal and human behavior. This new perspective is diametrically opposed to the traditional one-way cause-effect view that the environment controls an organism's behavior, either directly through stimulus-response connections or indirectly by initiating intervening cognitive processes between stimulus and response. We will see in the next chapter that this new view has revolutionary implications for behavioral science.