

Perceptual Control Theory: A model for understanding the mechanism and phenomenon of control

Carey, Goldstein, Marken, Nevin, Powers, Taylor:
20100920 plus content from WTP20100916 (BN20100914, WTP20100915, TC20100915)

The phenomenon of control is integral to psychology. Even a cursory glance through databases and journals reveals a staggering number of references to control. Terms such as *perceived control*, *locus of control*, *cognitive control*, *subjective control*, and *vicarious control* speak directly to the phenomenon. If we include implicit references to the phenomenon, such as *self-determination*, *agency*, *learned helplessness*, and *emotional dysregulation*, the number of references grows exponentially.

Recognition that control has an important place in the process of living, therefore, is ubiquitous. While the phenomenon of control is often front and center in research programs, however, the same cannot be said for the mechanisms of control. It seems odd to us that investigations into the phenomenon of control have for the most part proceeded in the absence of a clear understanding of the mechanisms that produce that phenomenon. The aim of this paper is to demonstrate a testable model of how control occurs, and to lay out the properties that are essential for control in any system, living or artificial. This model elucidates many other important concepts as well, such as cooperation and conflict, training, learning, and motivation. For simple kinds of behavior, the model has been developed in the form of a working simulation that tests the underlying theory in a rigorous way and, perhaps more importantly, makes the phenomena of control far more understandable than words alone can accomplish. Conclusions that are supported by testing the performance of a generative simulation of actual behavior justify a degree of confidence that cannot be accorded conceptual or statistical models.

Why has the recognition of control been divorced from an understanding of how control works? The reason is not hard to find.

Simple control mechanisms that have been in use for many centuries operate by clear and obvious principles. Anyone can understand how a float-valve regulator can keep the water level in a reservoir constant. As water flows out of the reservoir, the float (connected to operate a valve) descends just enough to increase the inflow of water to balance the outflow while the water level falls only by the small amount needed to open the valve by the right amount. If the outflow varies, the float and the water level rise or fall by the very small amounts needed to make the inflow remain equal to the outflow. Ktesibios of Alexandria worked out this clever negative feedback control system in about 230 B.C.¹ That is how the then-standard water clock managed to keep reasonably accurate time.

¹Landels (1978).

However, it wasn't until 70 or 80 years ago that engineers came to understand systems like this mathematically. Some scientists, those who would soon found the new disciplines called cybernetics and engineering psychology, quickly saw their relevance to the understanding of life processes. Many followed who recognized the parallels between artificial control systems and living systems, but they had difficulty believing what they saw. This is because those principles cannot be fully grasped without relinquishing the old ideas of lineal cause and effect, input and output, stimulus and response, independent variable and dependent variable—concepts of causality that are at the very foundations of most of the life sciences.

Scientists kept reinserting these conventional notions into their descriptions of control phenomena, perhaps without realizing it, simply because the old concepts had become intuitive and were taken for granted in all lines of thought. Instead of seeing the control loop as a complete entity, the new cyberneticists usually discussed it as a sequence of events, as if each part of the system waited for its turn, responding to a message from the previous part by sending a message to the next part. First the water flows out and the float drops. Then the valve opens. Then the water flow increases. Then the float rises and stops the inflow. Of course even Ktesibios two millennia ago could see that his regulator didn't work that way. But it is one thing to see a continuous process, and quite another to *express* it in the mathematics of continua. And even with a quantitative understanding, the appropriate verbal descriptions are actually quite difficult to achieve because words can only refer to one part at a time.

The sequential description makes it easy to keep thinking in the old terms of simple causality, each event causing the next. In reality the working of a control system is like the turning of a wheel; all parts carry out their functions at the same time. As the outflow increases, the valve opens more and the inflow increases while the float is descending slightly. But trying to think consistently in this way was difficult, and is still difficult for those new to these ideas. It is all too easy to fall off the tightrope and think, "First this, then that." But, as we will show, control systems do not work that way, and neither do living systems, not at any level from DNA to the cerebral cortex.

There is another problem. Living systems are and always have been purposive. They have goals and, if the systems are complex enough, intentions and desires and hopes and aspirations. The founders of biology, scientific psychology, neurology, and physiology had ruled these subjective notions out of the scope of objective science, but they are essential for describing in ordinary language how control really works.

As will be seen, the outward behavior of living systems is one phase of an orderly network of closed causal loops, of which the inmost phase is an assertion of intended results. We will try here to use ordinary language as much as possible, while keeping in mind the underlying principles to be discussed. For the mathematics we will for the most part leave the heavy lifting to demonstration programs that you will be invited to exercise. (Source code is available for inspection if you wish to probe more deeply.)² The first part of the exposition begins by introducing some concepts that go beyond the mere fact of control, so as to address questions that inevitably arise but which need a bit more discussion before we can delve into how all this works.

The what, why and how of control

²For the set of demonstrations published with Powers (2008), sometimes referred to as the 'LCS3 series'.

Consider an ordinary, everyday situation. I have my finger on a button beside a door. If I pause to think about what I am doing, it seems simple enough: “I’m ringing the doorbell”. But is that why I am here? Am I not trying to get someone to open the door? I am visiting Aunt Mary. That is why my finger is on the button. If you were a stranger passing by, you wouldn’t know this. You might guess that I am trying to add to the expected vote total for my preferred candidate in an upcoming election, or that I am making some money by delivering pamphlets. Whatever motivation you imagined, you would probably not propose that my objective was simply “to press my finger on a button”.³

This illustration may seem to belabor the obvious—that in everyday life, people have purposes, that the purpose is what the person is ‘doing’, and that their actions are the means for doing it. Equally obvious is that one usually cannot determine a person’s purpose by observing their actions. What is perhaps less obvious is that there are three aspects to this, a *what*, a *why*, and a *how*.

In the doorbell illustration, the first *what* that an observer might guess is “he is ringing the doorbell”. Although its *why* is obscure to an observer (but not to the doorbell ringer), its *how* is clearly “by pressing the doorbell button”. However, even this *how* has its own *what-why-how* pattern. *What* is “seeing and feeling my finger pushing the button”, *why* is “to make the bell ring” and *how* is “by moving my hand and arm to the appropriate place”.

If our observer looks at the possible *why* of pressing the doorbell button, another pattern of *what-why-how* emerges, for which making the doorbell ring is the *how*. *What* might be “to get someone to open the door”, or put better “to see someone open the door”. And so it goes. Every “what someone is doing” is part of a *what-why-how* structure. In every case, *why* is because some state of the world is not as the person would like it to be, and *how* is a means of making the world a little or a lot closer to what the person would wish.

All of this sounds self-evident, albeit anecdotal and not very scientific. But it can be scientific. The *what-why-how* complex describes *control*. This does not mean forceful dominance of people or the environment, it is a technical term of art that means bringing some particular condition toward a desired state and maintaining it there, like controlling the water level in a reservoir. That is the engineering definition of control, and the thesis of this paper is that control is what living organisms do. Indeed, it is what you are doing, on many levels and in many ways concurrently, as you read this paper.

This informal account of what-why-how points to a fundamental model of behavior that has been under development since the early 1950s. First published in 1960 as simply ‘feedback theory’, it was named Perceptual Control Theory (PCT) during the 1980s by members of the interdisciplinary, international group of researchers and practitioners who have engaged with it. This paper is a summary of the PCT paradigm as it is presently understood in this community of research and praxis, including methodology, results, and applications. It will become evident that PCT is largely concerned with finding explanations for *how* living systems must be organized to behave as they do, while properly acknowledging prior observations of *what* they do. The *why* of behavior, as noted above, is found in *what* they do at a higher level of organization.

Behavior as the means of control

³See Vallacher & Wegner (1985); Kozak, Marsh, & Wegner (2006); Marsh, Kozak, Wegner, et al. (2010). Links to these papers are at <http://www.wjh.harvard.edu/~wegner/actid.htm>.

The basic thesis of PCT is not difficult to describe. The behavior of organisms—their observed activity—is not the final product of prior causes. Rather, it is understood to be a variety of means to ends. From an observer’s standpoint, the ends are manifested in the way the behavior brings aspects of the local environment to particular states, static or dynamic, and maintains them in those states against disturbances. From the organism’s standpoint, the ends are certain experiences (or cessations thereof) that are intended or preferred. A moment’s thought will suggest to you how the subjective purposes of the organism may become objectively known by careful investigation of resistance to disturbances. PCT is the science of purposive behavior.

Consider another example. I hear music. In the language of PCT, I have a reference value for how loud I like this kind of music. At the same time, I perceive the current loudness of the music. I compare the loudness I perceive with its reference value, and if there is a difference, I do something that changes the physical environment to alter the loudness I perceive. Maybe I put in earplugs, maybe I move to another room closer to or further from the source, maybe I ask someone to turn the volume knob or turn it myself to make the music I hear louder or softer.

At the same time, other things might influence the loudness of the music I perceive. Maybe someone closes the door of the room where the music is playing, or turns the volume knob. I continuously perceive the loudness of the music, and at any time that it differs from my reference value, which may change from moment to moment, I may behave in such a way as to bring the loudness that I perceive nearer to the reference value that I currently have for it. As William James would exclaim “Again, the fixed end, the varying means!”⁴ In PCT the varying means are explained as actions opposing the effect of various disturbances on some variable that a person is controlling—here the perception of loudness.

We are describing here a feedback loop, as illustrated by the diagram of the canonical PCT control loop in Figure 1.

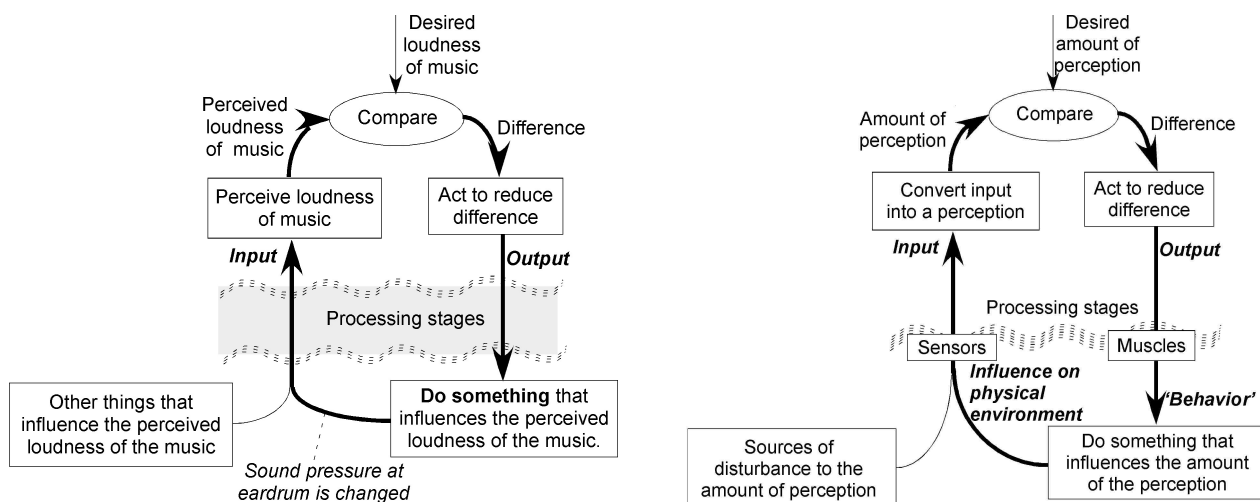


Figure 1. The canonical PCT control loop. (Left) controlling the loudness of music (Right) the generic loop. The key point is that what is controlled is the value of some perception, by means of the behavior that influences the physical environment: ends (purposes) achieved by variable means.

⁴James (1890/1950:4).

Not only is it a feedback loop, but the feedback is *negative*. In popular parlance ‘negative feedback’ is equivalent to criticism, whereas ‘positive feedback’ suggests encouragement.⁵ In the original engineering meaning, however, positive feedback increases the difference between the received amount of perception and the reference (desired) amount, the opposite of what is needed for control, while negative feedback decreases the difference. The float mechanism in Ktesibios’s regulator allowed never too much or too little water. A negative feedback control system can be designed so it can reduce the difference or error until it is at the limits of measurement (if it is worth the energy needed to generate corrections for every microscopic error). Only negative feedback results in full error-correcting, disturbance-resisting control.

Obviously, negative feedback control is not observed for every assemblage found in nature. A living system must have a particular kind of internal organization in order to be capable of controlling. A central concern of PCT from the outset has been to deduce the necessary properties of that internal organization by creating and testing generative working models of the actual behavior of individual organisms. Because behavior both results from changes and is the means to create and stabilize specific conditions of the organism and its environment, causality in this kind of system is circular. What appear at first to be ordinary physical consequences of motor activities are recognized to be perceived states of the world which the controlling organism actively seeks and defends against disturbances.

In the generic control loop of Figure 1, we can discern the features of the two main concepts of behavior that preceded PCT. Following the path from ‘Desired amount of perception’ through difference, output, and behavior, we have the same organization proposed by early neurologists⁶ and accepted by many modern neuroscientists and cognitive psychologists. According to this view, a high-order plan or goal is converted, step by step, into the simple or patterned behavior needed to achieve it. The other concept is seen by following the path from sensors through intervening processes (omitting the comparison and reference signal) to output and the muscles. There we have the organization known as a stimulus-response or input-output system, in which the environment causes behavior through causal paths connecting input to output. Both of these classical ideas omit the feedback path through the environment, although variations on the basic themes have been offered to take the effects of feedback—incorrectly—into account. Both concepts solidified into schools of thought before engineers discovered the right way to analyze systems having this circular kind of causal organization (H.S. Black, 1934).

Toward a new science of life

The biological, psychological, and social sciences have commonly studied organisms as simply one more possible arrangement of matter and energy, subject to the same laws of physics and chemistry as any other arrangement. PCT satisfies this requirement—control systems do not require any violation of the laws of physics and chemistry—but PCT recognizes additional laws that are emergent from the negative feedback control properties of suitably organized physical and chemical systems. The recognition of these properties enables a systematic accounting for the behavior of organisms, individually and in groups; without that recognition, observations of behavior can only be treated statistically or with unspecific generalizations.

⁵This usage stems from a misunderstanding when the term was taken up by the human potential movement, and has spread thence to fields such as counseling, education, and management.

⁶For example, Sherrington (1906), and many others.

To live is to control. As has long been known, life is negentropic.⁷ Organisms exploit the orderliness in the world around them as a means of increasing their own orderliness and stability. This effect is a consequence of control. Control theory explains how an organism can impose order on its local world at the expense of order elsewhere. Control does not confer totally arbitrary intervention in the processes of the environment, but it often seems to do so, in that the organism and its world both behave quite differently when an organism is in control. A car left to steer itself would soon run off the road or collide with another car if only momentum, gravity, wind, and potholes affected it. But add a driver to the car, and it—along with a huge number of other cars with drivers—stays on the road in its proper lane for hundreds of miles, traveling to a destination with great reliability. This is a highly improbable outcome when a controlling agent is not present. With control by the driver added to the picture, the same outcome becomes highly probable.

A shift in perspective comes with the recognition of the phenomenon of control. Without that recognition, behavior can be considered only from the point of view of an external observer, who as a scientist has little choice but to try to explain the activities of behavior as a mechanistic outcome of external forces acting on an otherwise inert 'preparation'. This has for centuries been a point of contention between scientists and laymen. Of course no one would deny that there are human experiences that cannot be investigated from outside the individual who is having them. 'Beauty', for example, is proverbially "in the eye of the beholder". Yet a requirement of objectivity seems to place such experiences out of the reach of science. This is not a problem for physics and chemistry, which need not take any account of 'points of view',⁸ since material objects do not 'initiate' any actions. But it is more than just a matter of 'physics envy'. The material of subjective experience has seemed intractable. When scientists have tried to introspect, they have failed to come up with predictable, reproducible observations on which all can agree. A science of psychology has seemed unattainable, except by converting it to something that looks as much as possible like physics and chemistry.

Control theory changes all that. A negative feedback loop is a strange organization, in that even without any external inputs, even when it appears to be doing nothing, it is functioning and behaving. If it is controlling something like the speed of a steam engine as the loads vary (or in an organism, the speed and manner of walking uphill and downhill in gusty winds), it is continually active, needing only an amorphous energy input without any external stimuli from outside the loop to direct the variations in its behavior. It resists disturbances of many kinds—even variations in properties of its own actuators—without being told to do so and without ever having encountered them before. And it acts on its environment to maintain a physical variable such as speed, position, orientation, repetition rate, and distance between objects—the list is almost endless—exactly as if it had a goal and as if it used its actions to produce and sustain that goal state in the external world. All of these capabilities, before the advent of control-system engineering, would have been (and were) dismissed as mystical fantasies by those who were trying to be faithful to scientific principles as they understood them.

Before control was recognized as a phenomenon, it seemed obvious and necessary that behavioral activities (responses) are controlled—in the sense of 'determined'—by perceptual inputs (stimuli). These are all that can be observed on the outside of the behaving organism. Once we recognize negative feedback control, it becomes obvious that it is the perceptual input that is controlled by means of behavioral activities. The relevant point of view becomes that of the organism, not that of an external observer. We cannot account for the *how* and *why* of the organism's activity until we have determined the *what* of it. The *what* is, as may be expected by this point, *controlling*, understood in a

⁷Schrödinger (1944), Brillouin (1953).

⁸Relativistic and quantum developments in physics have somewhat undermined that idea—but not in the life sciences.

sense unknown to the life sciences before the 1930s.

PCT proposes a new answer to the question of what it is that distinguishes a natural arrangement of matter and energy that is alive from one that is not. An organism constructed of the kinds of control systems described in PCT can have purposes of its own—that is, it can spontaneously select as goals future states of the world around it and alter its own behavior to achieve and maintain such goal-states. It can automatically, without external guidance or instruction, adjust its actions to oppose the effects of random and otherwise unpredictable disturbances (if they are not too sudden or too powerful for it) quickly and accurately enough to prevent their having any important effects. In larger assemblies of systems (more complex organisms), it can control hierarchically; that is, it can adjust one set of goals as a means of achieving other, higher-order goals. It can control many different variables in parallel at the same level of the perceptual hierarchy⁹, and by those means control multiple variables of a higher order at the same time. Certain kinds of control can operate within the hierarchy to make it learn and adapt, altering aspects of its organization so it changes the way it controls variables that matter to it less, in order to control variables that matter to it more.

By this shift in perspective, PCT reconciles the objective approach of science and engineering with subjective experience. It provides a clearly mechanistic model of behavior that can be implemented and studied as a computer simulation, and which also explains how human beings can have goals, intentions, preferences, desires, and other experiences that have sometimes been thought to be illusory or errors of interpretation.

The question naturally arises: if PCT has been building into a coherent model for 50 years or so, with an increasingly vigorous and growing research community gathered around it, why doesn't everyone know about it? Perhaps the most important reasons are found in an unfortunate development that occurred almost as soon as control engineers had elucidated the phenomenon of negative feedback control.

A discovery abandoned

Devices employing negative feedback control were first documented in the time of Ktesibios¹⁰, but it was only in the 1930s that the principles were first formalized by engineers. This was the basis of the wartime automation revolution of the 1940s. Norbert Wiener, a mathematician at MIT, working with his friend and MIT colleague, the neurophysiologist Arturo Rosenblueth, and a young engineer named Julian Bigelow (later a pioneer in computer science), developed automatic rangefinders, electro-mechanical negative feedback control devices that aimed anti-aircraft guns where aircraft overhead would be when the projectiles arrived.¹¹ The new field of cybernetics continued to draw interdisciplinary talents to the systematic study of the behavior and organization of this new genre of 'servomechanisms'. A cybernetic revolution in the life sciences began to gather momentum in the late 1940s and early 1950s, spurred and inspired by annual, sometimes biannual cross-disciplinary meetings, called the Macy Conferences, organized primarily by Warren McCulloch and the Macy Foundation from 1946 through 1953.

⁹The description of the perceptual hierarchy, well developed in the study of sensory physiology and psychophysics in the 19th century, is epitomized in Hayek (1952). Control in a hierarchy of perceptions is an important part of PCT, see Powers (1973, 1998, 2005). For the development of the hierarchy in children, see van de Rijt & Plooi (2008-2010) and references cited there.

¹⁰Wiener (1948 [1961]), especially the introduction to the first edition; Mayr (1970); Bennet (1979).

¹¹de Rosnay (2000).

But the revolution came to a halt, essentially dead, in a decade. Negative feedback control was abandoned as the best model of purposive living systems by its main original proponent, the prominent cyberneticist W. Ross Ashby, four years after it was adopted. His first book, *Design for a Brain*,¹² was about negative feedback control, and was instrumental in launching the investigations that led to PCT. In his second book on the subject, *An Introduction to Cybernetics*,¹³ Ashby pointed out how a perfect controller might be designed and, perhaps inadvertently, planted the idea that there was something inferior about negative feedback control because there always had to be some error to drive the output. He failed to mention that this error could be as small as a thousandth or even a millionth (in a real device) of the desired magnitude of the result. In place of the negative feedback model, Ashby and others offered a different idea. A perfect controller, they proposed, would analyze the environment, determine what actions would be needed to produce desired results, and then issue the commands necessary to make the muscles generate those actions. This represented a return to an idea of brain operation originally offered by Sherrington in 1906,¹⁴ in which the cerebral cortex formulated general commands that were then elaborated, level by level, into the detailed commands reaching the 'final common path' to produce organized behavior. The notion of a perfect controller is still very much with us. 'Modern Control Theory' appears to be an elaboration of it.

From that time onward, negative feedback control has been regarded by many as old-fashioned. In 1961, Alfred Chapanis, then president of the Society for Engineering Psychologists, wrote "The servo-model, for example, about which so much was written only a decade or two ago, now appears to be headed toward its proper position as a greatly oversimplified inadequate description of certain restricted aspects of man's behavior."¹⁵ This was published the year after the first paper leading to PCT appeared (after seven years of development). Writing in *Purposive Systems*, the 1968 proceedings of the first annual symposium of the American Society for Cybernetics, Ralph Gerard, a founder of the Society for Neuroscience and a contributor to the Macy Conferences on cybernetics, said "I have always regarded a drop of water sliding down a slightly inclined plane as showing all the manifestations of purposeful behavior." The reference was to the erratic path followed by the drop as gravity pulled on it, and to the movement toward the lowest point. The drop of water inevitably reached the lowest point not because it had that as a goal, but because of simple laws of physics. Its path was erratic not because it was resisting variable disturbances, but because irregularities in the surface pushed it this way and that. Gerard had clearly not learned anything about control systems.

In the 1952 book, Ashby wrote extensively about negative feedback control. In the 1956 book, he argued persuasively (but incorrectly) that the more complex design based on analyze-compute-execute processes would be more accurate since it did not need to allow any errors to occur. Ashby probably did not intend to abandon negative feedback control, but others picked up this idea and claimed that the open-loop design would operate faster than the negative feedback control system, eliminating delays. It could even, it has been proposed, anticipate disturbances and, by predicting their occurrence, generate actions to oppose them at the same instant they occurred. Since evolution would naturally have shaped organisms to operate in the best possible way, the followers of Ashby assumed that this model should also be used to explain the behavior of organisms.

Unfortunately for this assumption, real organisms seldom behave as perfect control systems. It is, in

¹Ashby (1952).

¹Ashby (1956).

¹Sherrington (1906); Burke (2007); Ashby (1952, 1956).

¹Chapanis (1961:126).

fact, easy to design artificial control systems that control much better than people do, but that amounts to making a model of the behavior of a perfect robot in a simplified environment, not of an imperfect human in a real one. To make a model that behaves as much as possible the way a real person does—in, for example, a tracking task—it is necessary to resurrect the negative feedback control model. Ashby had the right idea when he explained the importance of negative feedback control in his first book. PCT would have been accepted long ago, at least in cybernetics, if he had not written the second one. His second book did not explicitly say so, but was interpreted as saying that negative feedback control was a bad second choice when control really mattered.

What kind of theory is PCT?

The simplest form of theory in psychology, and the most prevalent, is a set of statements describing what will be observed under certain experimental conditions, such as “Mothers hold their babies on the left side.” The only test of such a theory that is possible is to observe whether this is how mothers really behave. It has no generality, no necessary connection to any other observation. Either most mothers do this, or they don't. A more complex example is Piaget's stage theory¹⁶ of the cognitive development of children, that is, a theory of changes in their cognitive abilities and processes. Piaget saw actions as the basis of early development, and mental operations coming in later. Stage theory is supported if in fact it is observed that cognitive development is first based on actions and later on mental operations. It must be tested, but the only test possible is to find whether this pattern is repeated, or fails to repeat, when observed again. There is no hypothesis to test concerning how the child must be internally organized in order to show these patterns.

PCT advances hypotheses, many of them testable even now, about the structure and organization of the nervous system, including the brain, at the level of functions if not anatomy: the *how* of behavior. If Piaget's ‘theory’ (proposed observation) is correct, we would expect eventually to use PCT to try to explain how the child manages to do these things.¹⁷ From the PCT point of view, Piaget has offered some data that, if valid, require theoretical explanation; he did not offer a theory of the same sort as PCT. His theory describes a phenomenon, but does not explain how it works. PCT neither supports nor denies theories of the type offered by Piaget: It accepts them to the extent that they are valid, and then, if all goes well, offers an explanation for why and how the phenomenon in question is generated—and suggests other similar phenomena that might prove to be observable.

When we look in the behavioral sciences for a theory of the same type as PCT, probably the oldest one to be found, dating back to Descartes, is the proposition that stimuli acting on sensory nerves are the cause of motor behavior. Like PCT, this theory is an attempt to explain all behavior, not any particular one. The other main theory of this type, deriving from ideas such as Ashby's analyze-compute-execute hypothesis and Sherrington's map of the brain, is that such stimuli cause centers in the cerebral cortex to generate plans of action, which are then executed by lower systems to produce desired ends. It is with theories at this level of coverage and intent that PCT must be compared, theories of the kind that aim to explain how all behavior is produced, rather than attempting to describe or predict what specific behaviors will be observed under specific circumstances.

¹Piaget (1937/1954).

¹See van de Rijt & Plooi (2008-2010) and references cited there.

To recapitulate the basic principles of PCT: Behavior is not a linear result of prior causes, it is the variable means of achieving goals that the behaving organism specifies within itself.¹⁸ The activities of behavior are only one among many of the causes that affect some aspect of the environment about which the organism has some preferences; or more exactly, affect the organism's perception of that aspect. The difference between that perception (a neural firing rate) and the preference, an internally specified reference value for that same perception (another neural firing rate) is the cause of the behavioral activities (by way of the propagation of the resulting neural signal downward through the control hierarchy). By this control loop of circular causation, the organism does whatever works to maintain that difference at or near zero; behavioral actions vary precisely as needed to achieve consistent aims.

This answers philosophical objections that fortunately are not much heard since the fall of logical positivism into disfavor. Criticisms were made that any notion of goals or purposes must be disallowed in a science of behavior, because it would require the future to affect the present, or effects to become their own causes, or infallible predictions to be made. Those objections may be outdated, but they were responsible for a general rejection of the idea of purposive behavior at the time when important psychological theories were just starting to form. The consequences are still with us. The efficacy of PCT models demonstrates, however, that all that is needed to account for purposive behavior is continuous and concurrent perception, comparison, and action.

Quantitative and qualitative theories: variables and categories

Following Ashby's perhaps unintended lead, the conventional ideas of control most widely accepted today propose that an organism achieves goals in steps, by first analyzing the environment, then calculating the actions and trajectories of action needed to bring the goal-state about, and finally by executing the actions. Despite the doubts raised here, the evidence for this model seems clear: the actions required to achieve a goal-state are indeed produced with the normal result of successful goal-attainment.

The evidence, however, is far richer and more informative when we measure the variability of behavior rather than counting instances of categories of behavior. Closer inspection shows that the actions within a category are not as regular and repeatable as they seem at first, and that in fact repeated goal-seeking actions have regular effects precisely because they are *not* repeated exactly. Sometimes exactly the opposite action must be used to repeat a regular effect. The reason is that those regular effects are influenced by more than just the organism's actions; there are also independently varying influences in the environment, including past and present states of the organism itself. Results can be repeated only by varying the actions so that they precisely counter those unpredictable disturbances and changes in environmental conditions which are also simultaneously influencing the result. It is not just that many different actions *can* produce the same result, a qualitative observation that Skinner proposed in his definition of the 'operant'; different actions *must* be employed, and just the right different actions each time. As we will see in demonstrations later in the paper, actions must vary quantitatively in exactly the right way if the same result is to recur.

The only reason that behavior (the observed activity) seems to repeat is that human observers tend correctly (in PCT terms) to describe the actual movements of limbs and the forces they create by

¹⁸Ultimately, the evolutionary history of the species enters into the meaning of "within itself."

referring to what they accomplish rather than what they are. A driver “making a left turn” seems to be generating a stereotyped behavioral pattern that is qualitatively the same each time it is executed, as if it were a simple repetition of what has been done before. This has been taken to imply that repeating the behavior means that the nervous system must be issuing the same commands to the muscles each time, with the muscles then having the same effects on the environment each time. The actual behavior, consisting of the operations of turning the car’s steering wheel and applying force to the pedals, is skipped over. But the impression of a repeating behavior is dissipated as soon as an engineer’s or a physicist’s eye is brought to the scene. The car never approaches the intersection of roads along exactly the same line or at the same speed as the last time; the tires distort, bounce, and slide by different amounts each time they encounter smooth or rough spots on a road that may be dry or slippery; crosswinds require more or less effort to be applied to the steering wheel to achieve the same turning path; the speed of the car influences the turning radius, as does the number of passengers in the car. Yet somehow every time there is a left turn, the steering wheel turns in just the manner required during that particular turn for the car to move in very nearly the same stereotyped fashion from the lane it is in to its proper place in the crossing lane.

It is the result that is stereotyped, not the muscle action that produces it or the neural commands that operate the muscles. Watching a race-car negotiating the turns at the ends of the straightaways, we as amateurs see that the driver actually turns the front wheels *the wrong way* if the speed is high enough. Conversely, repeating precisely the same neural output signals or actions each time would not produce the same consistent result. The detailed steering-wheel movements that got the race-car driver around the curve on the slow first lap would send his car into the wall if repeated exactly on this lap at full speed. That fact will prove below to have deadly consequences for the theory of reinforcement.

After sufficient quantitative observation of behavior, it becomes clear that it is not an organism’s neural outputs or motor actions that repeat, but the consequences of those outputs and actions. The outputs and actions themselves vary exactly as required to keep the consequences the same. The small disturbances revealed by close inspection—as well as large ones that are easy to see, if not to notice—have multiple independent causes that arise from different environmental sources on different occasions, at unpredictable times, in unpredictable directions, and to unpredictable degrees. Yet what we observe is exactly the kind of variation in behavior that is needed, given all the other visible and invisible influences acting at the same time, to make the critical consequences repeat.

By conventional ways of thinking this is impossible. But control systems do not operate in a simple input-output way. They can control consequences because they continually monitor the state of the consequences, and when that state differs, moment by moment and instant by instant, from what is expected or intended, the difference is used as the basis for altering the action in exactly the way that will keep the difference as small as possible. That is how the needed variations in behavior are produced, and why they do not need to be calculated in advance. Nothing prevents the organism from calculating actions in advance if it has the higher levels of organization needed to do this, but trying to anticipate disturbances and calculate what actions are going to be needed is very difficult and slow, and not likely to work very well in a world that is even slightly unpredictable or subject to disturbance. Negative feedback control is by far the simplest, fastest, and most accurate kind of control possible for real systems in the real world.

Multidimensional and multidimensional control

Any single control process can be modeled in isolation, as in the initial diagrams in this paper, but a model of the behavior of organisms must represent many control processes acting at once. In PCT, multidimensional control is modeled not as if complex signals or vectors were under control, but in the style called by Oliver Selfridge 'pandemonium',¹⁹ in which many one-dimensional controllers are acting at the same time. Because each controller senses just one dimension of variation, complex control requires many one-dimensional controllers to be working in parallel. While this seems wasteful of neural resources, with considerable duplication of function, the resulting models are in fact computationally simple, and the bottom line is that they reproduce real behavior accurately, the *sine qua non* of model-based analysis.

The Russian physiologist Nicolas Bernstein²⁰ anticipated cybernetics in many ways, and in the 1950s came to the same conclusion that was being developed at the same time in the ancestral "feedback theory" that became PCT: behavior has to be multiordinal—organized hierarchically, in layers. The same simple problem led to this conclusion both in PCT and in Bernstein's work: if the spinal reflexes act to stabilize limbs against disturbances, they will prevent centers higher in the brain from using those limbs to carry out behavior. Any disturbance will cause a reflexive reaction against the disturbance. Since the brain obviously does use the spinal systems in producing behavior, there must be a principle that allows the higher systems to operate by incorporating the reflexes, not just by overcoming them or turning them off. This principle can be extended to higher feedback loops, each higher loop presenting the same problem to subsystems above it.

Bernstein never completely settled this problem. He was on the right track, but he lacked expert knowledge of the engineering principles of negative feedback control which inform PCT. The secret lies in the reference signal, the (variable) "Goldilocks standard" against which perceptual signals are judged as being too small, too large, or just right (though on a continuous scale). To use a reflex-type control system as means of control, all that the higher systems have to do is vary the reference signal.

This casts new light on Sherrington's concept of a 'final common path',²¹ which he took to consist of signals carrying commands telling the muscles how much to contract. In a control-system model of the reflexes, the muscles are operated not by reference signals or command signals, but by error signals. The signals from spinal motor cells carried by alpha efferent axons to muscles result from two inputs to the motor neuron: an excitatory input descending from higher centers, and an inhibitory input coming from sensors in the tendons measuring the force applied by the muscle. The net signal leaving the motor neuron represents the excess of excitation over inhibition, and the feedback loop at this level makes the sensed tension in the tendon (due to the force exerted by the muscle) match the constant or changing reference signal received from above. Thus the brain (or a system higher in the spinal cord) sends the motor neuron a signal saying, in effect, "Make the tension signal match *this* signal." The feedback loop alters the output to the muscle, in just a few milliseconds, until the match is achieved. The reference signal is not a command to produce a certain amount of action; it is a request for a certain amount of perceived force or tension. As the reference signal varies, so does the perceived—and actual—muscle tension.

This establishes a principle of hierarchical control that seems to apply equally well at many levels of organization. Higher systems act to control their own perceptual inputs by telling lower systems to produce a specific amount of the variable they are specialized to sense and control, not what action

¹See http://en.wikipedia.org/wiki/Pandemonium_Architecture.

²Bernstein (1967).

²Sherrington (1906).

they should perform in order to produce the requested perception. How much to sense, not what to do. The lower systems autonomously act on their environments to make their own perceptual inputs match the specified reference condition of the moment.

Conflict and cooperation

The concurrent control of input variables by different controllers can result in conflict. An everyday example of conflict within the hierarchy occurs when a parent wants to warn someone of a hazard, but does not want to wake the baby. Warning the person calls for using a loud voice; letting the baby sleep calls for quiet. Two control loops are controlling the same perceived loudness of one's voice, trying to produce very different values of the same variable at the same time. The person may resolve the internal conflict in this case by gesturing to get the other person's attention and then whispering.

The two controllers may be in different people. One person approaches an open doorway at the same time as someone else coming the opposite way. One may stand aside and wait, or both may turn sideways to slip past each other. Most conflicts are routinely resolved just this easily. While a conflict remains unresolved, neither controller can achieve its goal; both are impaired, and one or both may effectively be removed from functioning. Psychological difficulties with this basis are addressed by the Method of Levels (MOL), which will be described in a later section.

When conflicts between control systems inside one organism or in different organisms are not resolved, the result can be a serious loss of function. Each system tries to make the same physical variable match a different reference condition. If the difference is moderately large, both systems will experience large control errors, because one variable can't have two different values at the same time. As a consequence at least one of the control systems will produce as much output action as possible, limited only by strength and endurance, while the stronger system will have a reduced range of available output variations.

Both conflict and cooperation have the same formal description in the PCT model: two or more controllers are controlling their perceptions of one common variable in their environment. In the case of conflict, the control actions of each are a disturbance to control by the other. In the case of cooperation, these conflicts are resolved as they arise, usually by each party taking responsibility for one part of the task that is independent of other parts. If two cooks interfere with each other, let one of them make the salad dressing while the other seasons the broth. When very skillful control of exactly the same variable is involved on both sides of a conflict, even small differences in the goals can cause large degrees of opposing efforts. This explains why cooperation, even when highly valued, can be difficult to put into practice.

The resolution of conflicts requires changes in some part of control systems that create behavior—in the given environmental situation, the perceptions or the actions (as physical forces or reference signals for lower systems) must become different. In PCT such changes are described as a general process of reorganization.

Changes of organization

The final facet of PCT is concerned with adaptation and ontogeny. Both are accomplished by changing

the structure of the control hierarchy, either to change the way it behaves (adaptation) or to develop capacities for control that it did not previously have at a more primitive stage (ontogeny). In accord with the general principles of PCT, this process of changing the structure of control systems is seen itself as a control process, in which variables of basic importance, referred to by Ashby as 'essential variables', are maintained near reference states by altering the organization of the organism. The resulting theory of *reorganization* incorporates one of Ashby's most important and still viable ideas, that of 'ultrastability' achieved not through systematic control or direct changes in behavior, but through random variations of the properties of a system.

The worst case for a learning theory is that in which the organism needs to learn some control process that has never been learned before, which cannot be extrapolated or generalized from past experience, and which cannot be worked out logically, either because it is unique or because the organism has not developed (or may never develop) the requisite capacities for using logic. This problem led Ashby to a control-system theory (adopted into PCT) of the most basic kind of learning, the kind that has to precede the learning of any systematic methods of learning. It is this sort of worst-case learning that is usually meant when the term *reorganization* arises in PCT discussions. The first version of this theory was not given much weight early in the development of PCT, but under the tutelage of a small bacterium it has developed into a testable model of basic learning.

E. coli reorganization

B. F. Skinner explained the acquisition of the first successful behavior in conditioning experiments by saying that organisms spontaneously "emit" random variations of behavior. PCT adopts that idea but in a different form: the basic theory of reorganization is that the organization of system parameters (and hence behavior) varies randomly at a rate that depends on the amount of *intrinsic error*. A seldom-mentioned condition for successful operant conditioning is that animals are deprived of whatever is to be used as a reinforcer; conditioning with food reinforcers, for example, does not work on satiated animals. Deprivation is not just an 'establishing condition' as Skinner called it. It causes errors in control systems that are organized to produce inputs of food, water, warmth, light, and so on and keep them at levels in the range necessary for comfort and even life. The theory of reorganization that is part of PCT proposes that when such basic control errors are large enough or protracted enough, they bring reorganizing processes into action.

Intrinsic error means a difference between the state of some essential variable, such as blood sugar, and a genetically-determined reference condition. This difference is proposed to result in random changes of organization. The kind of learning involved is fundamental, the kind that occurs when there is no systematic method available for higher levels of control to pursue, and when there is no prior experience to guide changes. This has long been known as 'trial and error' learning. Because the changes are unstructured, they are not constrained by anything but the existing organization, so the possibility of finding solutions to new control problems is maximized.

Clearly, if the random changes of organization produce new behavior patterns that eliminate the deficit in blood sugar, the intrinsic error driving those changes will be eliminated and the changes will stop. That will leave the latest result of reorganization in effect, and behavior will show the new patterns from then on, just as if something—a 'reinforcer'—had somehow told the organism that the new pattern was a good one for maintaining the right blood sugar levels.

This concept has been part of PCT since the first published paper in 1960, but it seemed at first too

inefficient. Doubts about random reorganization are reasonable; at first blush it doesn't seem very likely to work. Not until 1980 was it taken seriously. In that year, Daniel Koshland published a small book on bacterial chemotaxis which contained a principle that vastly increases the effectiveness of random reorganization. This principle enabled successful simulations of multidimensional systems reorganized in accord with it.

The bacterium *E. coli* cannot steer, but it can make its way up and down chemical gradients very effectively. It does so by swimming in a (more or less) straight line and occasionally 'tumbling', changing direction in a way that Koshland reported was actually random. The explanation of the gradient-climbing is found in the observation (verified by perfusion experiments and biochemical analysis) that *E. coli* senses the time rate of change of concentration of chemical substances, change that is normally induced by its swimming in the gradient. If the rate of change of an attractant is positive, *E. coli* continues in a straight line. The attractant is diffusing radially from a source in the fluid medium, so the straight-line path of the bacterium may be visualized as a tangent line across concentric circles around a point, gradually reaching a closest approach to the source. As the path then starts to draw away from the source, the time rate of change of concentration goes negative, and *E. coli* immediately tumbles.

Since the tumbles change the direction of swimming at random, the result is just as likely to be worse as it is to be better. If the rate of change is still negative, however, another tumble ensues immediately, and tumbles keep repeating until the rate of change is once again positive. The bacterium does not swim far—a few body lengths—before tumbling again, so it does not travel much between successive tumbles. The result is that it travels much farther and faster up the gradient than down it. For repellents, meaning substances that *E. coli* avoids, the relationships are reversed. According to Koshland, *E. coli* can behave in this way in relationship to more than 20 different substances—simultaneously, apparently.

To translate this principle into terms of reorganizing a control system in a simulation, the spatial dimensions in which *E. coli* moves become parameters of multiple control systems. Swimming in a straight line becomes, in a simulation, adding small increments again and again to each parameter being varied, the direction of travel in parameter space being determined by the relative positive or negative amount of change per iteration in each dimension. A tumble corresponds to altering randomly the proportions in which different parameters are changing. To make sure the process does not overshoot its purpose, the amount added to each parameter is reduced as the control errors decrease.

In comparison simulations, the *E. coli* principle has proven to be over 50 times more efficient than a method based on random point-mutations of parameters, the proposal in the initial 1960 paper. This is because it makes use of information about the changing size of control errors. A slow progressive parameter change that continually reduces control errors simply continues as long as improvement continues. Only when the control error worsens does a 'tumble' take place, and then tumbles occur rapidly until the errors are declining again. The 50-fold gain in efficiency is seen when only two parameters are varying; the larger the number of parameters being reorganized, the greater is the gain in efficiency. It is possible that this principle will provide the final rebuttal to arguments that natural selection with random variability of individuals in a population is unlikely to account for the facts of evolution. If evolution is actually carried out at the level of the genome by an organism-generated

process of *E. coli*-type reorganization, it may easily prove to be as efficient as necessary. (The idea that organisms generate their own evolution is not entirely new.)²²

We will now examine how reorganization theory relates to other theories with claims to an account of learning, training, and kindred topics. Two of the main theories are Pavlovian or classical conditioning theory and Skinnerian operant conditioning theory. Both of these widely accepted theories deal not primarily with learning in the sense of reorganization, but with carrying out already-existing control processes in a changed environment. Reorganization—change in the internal organization of the behaving system—occurs only at special stages of these processes, and much of the behavior that looks like learning or adaptation can be produced by control systems without any change of internal organization.

Classical (Pavlovian) conditioning

Pavlovian conditioning begins, we propose, with an existing control process, either learned or inherited (a “reflex”). We will consider the example of thermoregulation. An animal subject to cold air blowing on its body will eventually start to shiver: seemingly, the stimulus of cold being sensed by skin receptors is causing the response of shivering. But bodily temperature control is most directly concerned with maintaining the core temperature or even the brain temperature close to a reference level. When loss of heat to a cold environment disturbs the core temperature, sensors in the brain report the sensed temperature as a perceptual signal, and the perceptual signal is compared with a reference signal (or equivalent) that defines the current reference temperature. Cooling thus disturbs the sensed core temperature and produces an error signal, which activates the shivering, which creates heat that opposes the effect of heat loss on the temperature: a simple negative feedback control loop. This would happen even if the skin temperature sensors were not involved.

If core temperature is a controlled variable, then heat-loss has its primary effect through disturbing core temperature, and the sensory effects occurring at the skin are not necessarily involved at first. The shivering that results from cooling is an action that generates heat, and this heat counteracts the temperature drop in the core. However, there is a time lag between onset of cooling of the core and the buildup of shivering, and during that time the core temperature can drop significantly. This would be classed as a change in an essential variable, per Ashby, and that will result in the start of random changes in organization of existing control systems. Those changes of organization will persist until they have the effect of reducing or eliminating the error in this basic inherited control system. Changes that do not accomplish that result simply allow reorganization to continue so further changes occur. What sort of change would correct the average temperature error and thereby stop the changes?

If shivering started earlier after the onset of heat-loss, the core temperature error would be counteracted sooner and perhaps might not occur at all. This is where the skin temperature, treated now as a stimulus that will become a CS, a conditioned stimulus, comes in. If the shivering could be started as soon as the skin temperature dropped significantly, it would start adding heat to the volumes where core temperature is sensed before the cooling can have much or any effect. This would happen if the input function of the core temperature control system were to be reorganized.

²²At the time of writing, [title] Yu Li of Manchester University, UK, has created a new algorithm that also uses strictly random changes, but has at least 5 times the efficiency of the method displayed by *E. coli*. The behavior it produces seems more like the paths taken by winged insects, such as mosquitoes following a CO₂ plume toward an animal.

Reorganization changes the weights given to signals at random, raising and lowering them at a rate that depends on the amount of “intrinsic error,” as we call an error in one of these basic control systems. If an input function is being reorganized, sooner or later many different input signals will become connected and then disconnected again as reorganization continues: the synaptic strengths will rise, then fall again; neural fibers might grow or atrophy. Let us suppose that one of the input signals that becomes connected to the functions that sense core temperature is the signal from the skin sensors that report a drop in skin temperature: “cold receptors”. The controlled variable will then become not just the core temperature, but the core temperature minus the “cold” signal²³. With this arrangement, if either core or skin temperature drops, the perceptual signal in the control system will decrease. If it falls enough below the reference temperature, shivering will start.

The result will be what is observed: A drop in skin temperature will cause shivering to start before the core temperature actually begins to fall, or before it has fallen enough to cause reorganization to start. And more to the point, the same process can happen if there is any sensory signal that regularly and repeatedly precedes the core cooling and that is capable of being connected with the correct sign to the input of this basic control system. The anticipatory signal could be a tone sounded before the experimenter causes cold air to blow on the organism, provided that the tone repeatedly and reliably stays in the correct temporal relationship to the core cooling.

Reorganization is not logical; no reasoning is involved. The tone ends up causing shivering because when that connection is made at some point during the random shuffling of connections, the shivering starts before the core temperature falls, and the reorganizing process stops.²⁴ If that did not happen, that connection would not have formed or persisted after forming in the first place. If it already existed because of prior training, removing the relationship between the tone and the cold blast would allow the cold to affect the core temperature, so reorganization would continue. That connection would be reorganized away. If the tone had been causing shivering, it would cease to cause it: that process is called extinction.

This model has not yet been simulated to see what other phenomena of classical conditioning it could reproduce, or to reveal any hidden flaws in the above proposal. That test will no doubt be done in the near future.

Operant (Skinnerian) conditioning theory

We can now see classical conditioning as a process of learning new perceptions to control. The same reorganizing process that creates the phenomena of classical conditioning can also explain operant conditioning. The main difference is that reorganization now works primarily on the output side of the control system rather than the input side.

In operant conditioning, the situation is more complicated because behavior is ‘instrumental’: it affects one of the variables that is also considered a cause of the behavior. If that sounds like the kind of closed loop on which PCT is based, it should. In all the basic forms of operant conditioning such as a fixed-ratio experiment, there is first a deprivation of something like food that is important to the organism. Again, we begin with an error signal in some basic and presumably inherited control system.

²³The “cold” signal will have to make an inhibitory connection for reorganization to stop.

²⁴If shivering starts too soon, core temperature will be increased above the reference level. If the increase is large enough, it could start reorganization and reduce the anticipatory shivering. Eventually, just the right amount of shivering will occur.

Subjected to this ‘establishing condition’, the animal ‘emits’ whatever behaviors have already been learned that can counteract the deprivation. It is very likely that the animal already has learned or inherited some strategy for finding food, such as executing a search pattern that bring all parts of the local environment within range of food detectors employing vision, or more likely smell. During this search, the animal’s actions by chance actuate something in the experimental apparatus (depressing a lever, typically) to cause a bit of food to be dropped into a feeding cup, where the animal finds it and eats it. Typically, the behavior then changes so that the animal begins exploring in the area where the food was obtained. This increases the probability that the lever will be contacted again and more food will be delivered. So far, no learning, no change of organization has occurred. But now, gradually, the random-looking explorations begin to focus on the lever, and after some time the animal is pressing the lever in a way that looks organized and purposeful so as to provide for itself as much food as it needs (if the experimenter has been generous).

Two different processes appear to be working here. The first one is simply the search for food and the narrowing of the search for more to the area where food was found. This is most probably an organized behavior that all rats learn, or it may be innate behavior due to a control structure that they are born with. In the second process, the rat’s accidental and then purposeful use of the lever to obtain food, it is the progressive refinement of the behavior pattern that makes it instrumental—reliable and organized to produce a specific effect in the given environment. Only the second process requires any change of internal organization. Together, these two processes take place in what we may call the learning phase of a conditioning experiment. That phase is followed by a maintenance phase when the animal routinely uses the new technique to alleviate its hunger. The reorganizations in this kind of conditioning are primarily on the output side, where errors give rise to changes in reference signals being sent to this or that lower-order control system that produces already-organized behaviors.

We have noted how in operant-conditioning experiments it is necessary to use ‘establishing conditions’; if food is to be the reinforcer, the organism must first be deprived of food long enough to reduce its weight to 80 or 85 percent of its free-feeding weight. Of course, depriving the organism of water does not increase the ‘reinforcing’ properties of food—and anyone would say “of course”—but the theory of reinforcement can’t explain why. Researchers in this area do—informally—understand that deprivation of something essential, like food, produces a ‘drive’ to obtain things that correct the deprivation. Removal of the deprivation reduces or removes the reinforcing value of the so-called reinforcer. When we realize that the food input is being controlled,²⁵ the reason for the deprivation effects becomes self-evident. When successful control is finally learned, that ends the deprivation and thus ends the random variations. The new organization of control systems can now maintain food intake close to whatever its reference level is. Thirst has its own intrinsic reference, and depriving the organism of water results eventually in its controlling the water intake, not the food intake, to bring it back to its reference level.

Behaviorist principles do not allow hypotheses about what is happening invisibly inside an organism, but PCT offers a model based on scientific principles that has a place for things that an organism wants, intends, wishes for, and so on. A control-system model of operant conditioning proposes a perceptual signal and a reference signal.²⁶ The perceptual signal represents the current state of some perceivable variable (in this case related to the nutritional state of the body); the reference signal or equivalent

²⁵There is a *why* for this, of course: control of food input is the means of controlling a perception of hunger. Hunger is understood to be intrinsic, having no superordinate *why*.

²⁶A model also requires detailed data about the behavior of individual specimens. In many types of research these data are routinely discarded as statistical aggregate data are assembled.

specifies the wanted amount of that perception. A likely kind of learned or innate behavior driven by the difference or error is to use existing behavioral control systems to explore every inch of the cage in a way that a layman might interpret as a search for food.

Reinforcement is said to increase the probability of the behavior that produced it. This has a descriptive basis in observations during the learning phase of an experiment. Observation of what happens in the operant cage shows, however, that it is the convergence of exploratory activities below, near, and above the lever that increases the probability of producing the reinforcement. The PCT alternative to reinforcement, up to this point, is simply to say that this is normal control behavior. Hunger could be a perceptual signal in some basic, perhaps inherited, control system. The reference signal would specify a level of zero for this perceptual signal. The error then causes already-learned organized exploratory behavior with the possible result of reducing the hunger toward its reference level. When the error is reduced, the tendency to go on exploring is decreased; when error is reduced enough, the exploring ceases.²⁷

This model leads us to expect, for different reasons, essentially the same series of events that the theory of reinforcement suggests, so for the initial learning, either theory accounts for the described facts. Simply having a plausible alternative to the theory of reinforcement, however, is useful in itself. It shows that reinforcement *is* a theory, not simply a description of a fact.²⁸

When we pass beyond the initial learning phase to later stages of performance, however, we encounter anomalies that the theory of reinforcement is unable to explain. We would expect the same relationship to continue to hold for further changes in reinforcement. If the rate of reinforcement increases, the behavior rate should also increase, and conversely if the rate of reinforcement decreases.²⁹ But that that is not what happens.

During the second, or maintenance phase, assuming that the animal gets all its food from lever-pressing, it comes to press the lever often enough to generate and consume the food at the average rate which suffices to keep the body weight very constant for long periods of time. Furthermore, if food is consistently added, even by tube feeding, to what the behavior is producing, the behavior rate does not increase, but instead decreases just enough to keep the total food intake constant and the weight constant. An animal offered high-calorie, tasty food in unlimited quantities (increasing the rate of reinforcement) will not, despite the implications of reinforcement theory, eat faster and faster until it explodes. Conversely, consistently removing food (lowering the rate of reinforcement) does not lead to slower lever pressing, it leads to faster lever pressing, so that again the original rate of food intake is maintained. This was long ago confirmed in experiments with obesity. Rats that are obese because of certain hypothalamic lesions maintain a higher body weight, but they, too, defend it against disturbances in either direction, keeping their total food intake constant. Experiments with rats obtaining all their food by lever pressing (Collier, 19xx) showed that these animals maintain an approximately constant calorie intake when the apparatus varies the required ratio of presses to pellet deliveries over a range from 20:1 to 1000:1, even if the reward size is varied, too. So the direct evidence shows that the effect of changes in reinforcement is actually the opposite of what reinforcement theory predicts.

²⁷Of course special hypotheses can be made about properties of the perceptual function to give small changes in food intake an exaggerated effect, but such details are best left for experiments to settle.

²⁸It also provides a physically and physiologically plausible alternative to the claim that a bit of food can change a probability, which of course is not a variable of the kind that can be changed by doing something to it. Behaviorism may not be concerned with *how* or *why*, but PCT is.

²⁹For simplicity, we will now assume a fixed-ratio schedule in our examples.

For the comparison with PCT, we must review this process in somewhat more formal terms. In operant conditioning, an animal learns to press a lever to get food. Each bit of food is said to increase the probability of pressing the lever again and obtaining more food. When the variables are measured as rates, this theory is expressed by saying that the rate of pressing is increased by each reinforcement. Clearly, this can't continue indefinitely, and it doesn't. The behavior rate comes to some steady value and the reinforcement rate also comes to a steady value.

Various explanations are offered for this limit on pressing and reinforcement rate such as satiation and fatigue, but it can be summed up simply by saying that the behavior rate stops increasing when R , the reinforcement rate, comes to an observed steady-state rate R' .

$$R = R'$$

In operant-conditioning language, the reinforcement rate, when equal to R' , is said to be 'maintaining' the behavior rate (the behavior rate is also clearly maintaining the reinforcement rate). Now, if P is the rate of pressing, dP/dt is the acceleration of pressing, the rate at which the rate of pressing is increasing. The above summary, then, is expressed by the equation

$$dP/dt = g^* (R' - R)$$

When R is zero, the rate of pressing increases at its maximum rate, $g^* R'$. While R is less than R' but rising toward it, there is a diminishing increase in rate of pressing. When $R = R'$, the rate stops increasing and the reinforcement rate is constant. The constant g is adjusted to make the equation fit the observed data.

The reinforcement rate will also be constant at a value of $R = P/N$, where N is the number of presses required per reinforcement (fixed-ratio schedule). When g has the right value, it will found that the above equation will be satisfied at the same time that the equation $R = P/N$ is satisfied.

These relationships can be put into correspondence with the canonical control-system diagram, as shown in Figure 2. The input quantity is the reinforcement rate R . The pressing rate P is the output quantity and the feedback function is a multiplier of $1/N$. (This is not labeled as such in Figure 2.) The magnitude of the hypothetical perception of the input quantity r can be numerically equated to R by suitable choice of units in which to measure signals inside the controller. The reference signal r' is then numerically equal to the observed asymptotic value of R , which is R' . The error signal is equal to $r' - r$, and the output function is a time-integrator with a sensitivity factor of g (also not labeled as such). The factor g determines how rapidly the steady state will be reached but does not influence the magnitude of the steady state values of P and R .

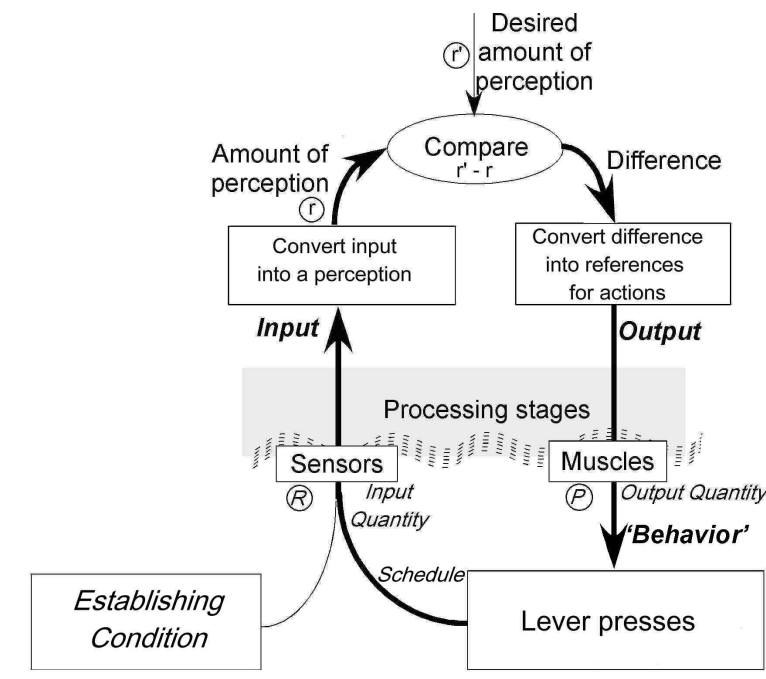


Figure 2. The canonical PCT control loop, as in Figure 1, showing the correlation with variables identified in reinforcement theory. The reinforcement rate R is converted to a perceptual signal r , a neural firing rate within the organism. The reference signal r' corresponds in the same way to R' , the observed asymptote of the reinforcement rate (not shown). The difference, which in control theory is called the error signal, is converted to reference signals for actions which, through the environment, influence the input quantity—in this case, any actions that repeatedly depress the lever. The input quantity R is a measurement corresponding to the controlled perception r . The only disturbance that is permitted by the experimental apparatus is the enabling condition, which is maintained by means of the structural constraints of the apparatus under the given reinforcement schedule.

As we have seen, the PCT model explains this behavior as control of perceived hunger by means of varying food intake, which affects some important variable (such as blood sugar) that is directly dependent on food intake. It is simply another demonstration of negative feedback control.

In recent years, feedback phenomena have claimed more and more attention as researchers discover closed loops of causation as isolated phenomena at every level of organization in living systems. PCT shows how these observations fit into a systemic whole, but to grasp this we must begin with simple cases. The following example is taken from “LCS3”, a suite of demonstrations (Powers, 2008), more of which will be discussed later.

Simulations and models

In the following demonstration of a negative feedback control model, a person uses a mouse to make a cursor track a moving target for one minute. Data are sampled 60 times per second. The data for a single experimental run are shown in the upper plot of Figure 3 representing a one-minute experimental run. The red trace shows the target movements; the green trace shows the mouse and cursor movements. The black trace shows the difference between target and cursor—the tracking error.

There is a consistent small time delay, hard to see in Figure 3, between target movements and cursor movements (upper plot). The delay is not removed by anticipatory mouse movements as Ashby and others after him claimed would happen. In the upper part of Figure 3 the results of fitting a negative feedback control model to the data are summarized; the best-fit delay in the model's response is 7/60ths of a second, which is 7 frames of the computer display running at 60 Hz, or 116 milliseconds. That is how far behind the target movements the participant is moving the cursor, on average. The delay does not get smaller with practice.

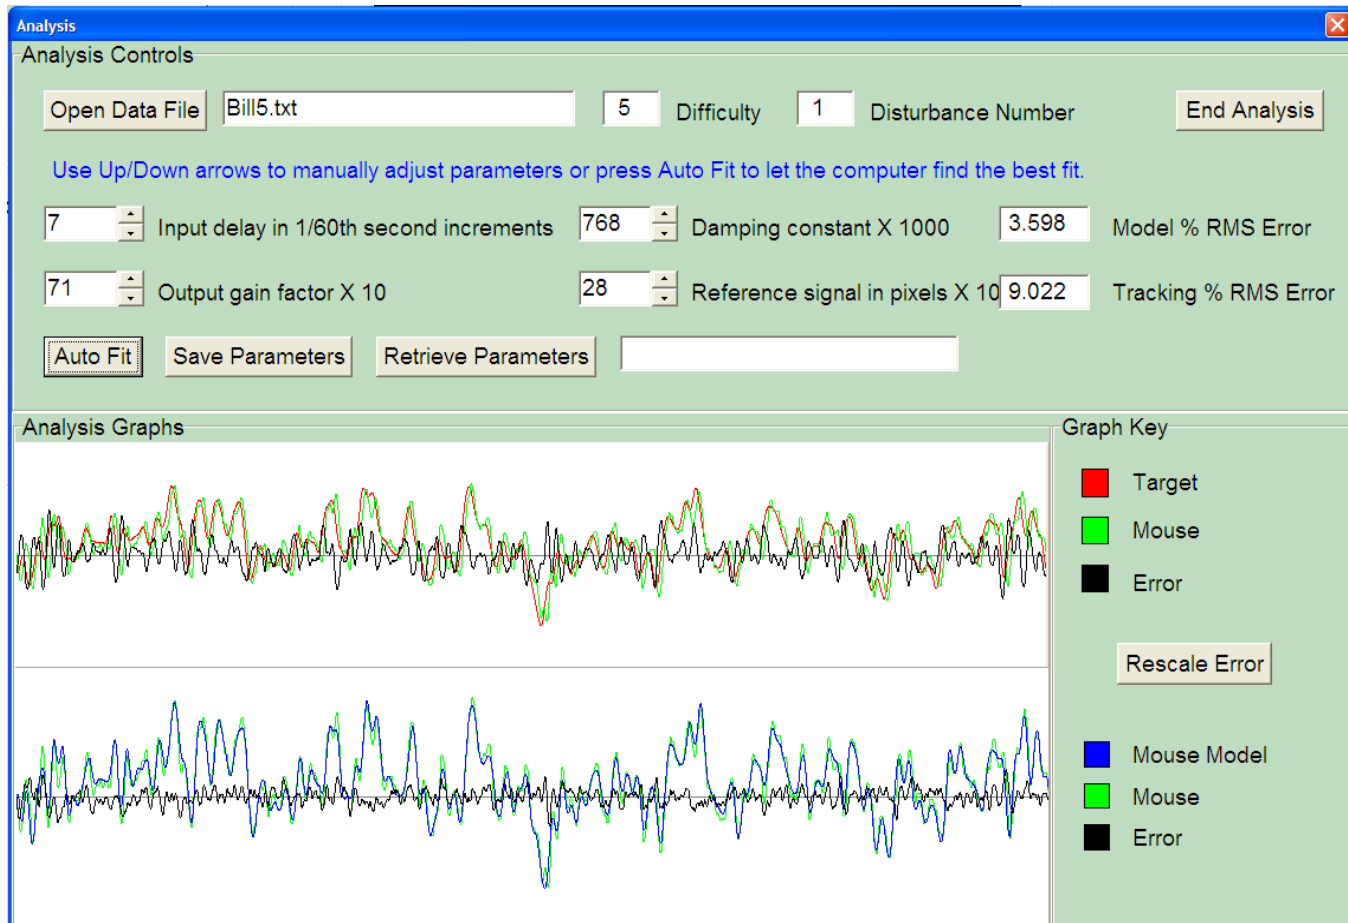


Figure 3. Analysis of human tracking run and fit of negative feedback control model to the data.

The lower plot compares real human performance with the performance of a computer implementation of a negative feedback system controlling the proximity of the cursor to the target, resisting unpredictable random disturbances to the target position just as the human subject did. The same target movements are used for model and human runs. The model's simulated mouse movements (blue) are compared with the person's real mouse movements (green). They are very nearly identical, with the same delay relative to the target movements. The mean difference between model and real behavior is 3.6% of the range of target movement. In this run, the target movements are rapid enough (maximum difficulty) that the tracking error is 9% of the target range, which to the subject seems like very poor tracking. The model fits the person's performance well within the tracking error, showing that the model is making similar mistakes. This same model will control even better with the delay set to zero,

but it will work too well. With no delay and all the remaining parameters optimized, the mismatch with the real behavior rises from 3.6% to 6.0%. The delay is real.

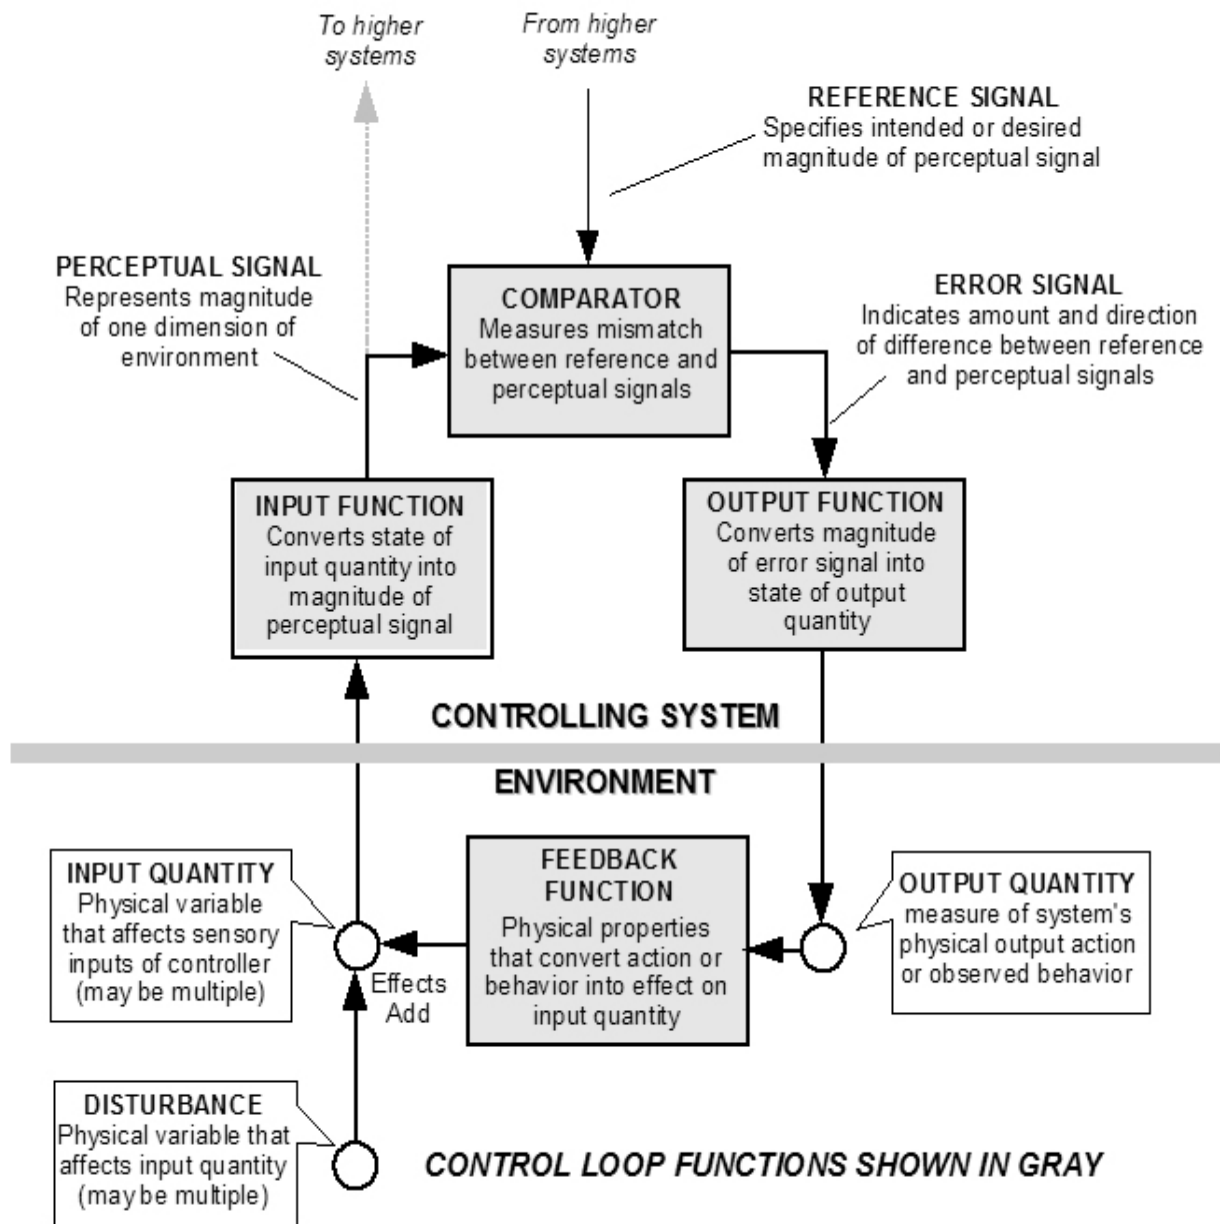
Soon after control theory started to appear in the psychological literature, various commentators noted that all real systems contain time delays. It was thought, apparently, that with any time delay at all, a negative feedback control system would have to become unstable. Error-correcting actions would start too late to prevent disturbances from having immediate effects, and would persist after disturbances disappeared, generating self-disturbance; and the time delay would convert negative error-opposing feedback into positive error-amplifying feedback, with the likely result that the whole system would oscillate violently. "Feedback is too slow" became a mantra for those resisting the new propositions. That is the result of making qualitative judgments about what is actually a quantitative problem.

While time delays can result in pathological behavior, all that is needed to correct it is to make the output driven by the error signal proportional to the time-integral of the error rather than to the error itself. This is equivalent to making the rate of change of output proportional to error. The constant of proportionality is adjusted so that during the time-delay that exists, the feedback effects from the output cannot change by more than the size of the perturbations caused by the disturbance. We will show this in greater depth in the next section.

This adjustment is sufficient to stabilize the system given any fixed or maximum time delay in its response. Even more important, as we have just seen above and will further demonstrate presently, a working model of a control system incorporating this principle can reproduce experimental behavior of a human participant, including delays, with an accuracy of three to four percent of the range of variation of observed disturbances and responses, equivalent to a 25 to 30 sigma fit of model to data. There can be no practical possibility that this model fits the observations by chance, since $p < 1E-12$ or much less.

So while it is true that the success and stability of a control process depends on a number of static and dynamic aspects of the system and its environment, and that general treatments of the stability and accuracy of control systems can become very complex, nevertheless in applying control theory to organisms there is a shortcut to a solution: the living system's performance is observed to be stable and accurate, so a biological answer to the problem of stability, even if unknown, clearly exists. Given that observed performance is stable, and that we know of one way of stabilizing a model that accurately reproduces real behavior, we are assured that PCT gives a correct general picture of how control works without requiring that the exact method of achieving stability be known.

The very simple model that provided the illustration above exemplifies a method of analysis that originated in the 'operations research' of World War II and the field of engineering psychology that grew up right after the war. Like PCT itself, it is basically a simple idea; but also like PCT, the power that it proves to have as an aid to understanding far exceeds what its simplicity seems capable of providing. It is important to understand both the simplicity and the power of a model constructed in this way, so we will take some time to study that here. Figure 4 shows a generic model of a single control system, one system among many at one level among many: the building-block of the hierarchy of control systems that constitutes PCT.



The model assumes that inside the participant there is a perceptual signal, some kind of neural signal the magnitude of which literally and quantitatively represents the magnitude of the input quantity. In the tracking task, the input quantity is the vertical distance between the target position **T** and the cursor position **C**, and the random variation of the target position acts as a disturbance of that input quantity. This suggests that the perceptual signal **p** quantitatively represents the cursor position **C** minus the target position **T**, as expressed in the equation $p = C - T$.

Actually, of course, as we have noted, there is a delay involved in going from the perception of target

and cursor to the signal representing the distance between them. If the delay is τ seconds, the working perceptual signal at time t represents the target-to-cursor distance at a prior time, $t - \tau$, so the corrected equation as used in the model is

$$1. p(t) = C(t-\tau) - T(t-\tau)$$

The basic negative feedback control system receives a reference signal (r) from elsewhere within the organism which specifies the currently intended or desired magnitude of the perceptual signal. The “comparator” emits an error signal e indicating the magnitude and sign of the difference between r and p (the time index is omitted but understood):

$$2. e = r - p$$

Experiment has shown that in the best model for the output function, the mouse velocity is proportional to the error signal. A positive error (perception less than reference) causes an upward velocity of the cursor that is proportional to the error by a gain factor G (that is, $V_{\text{cursor}} = G * e$).

The next position of the cursor C_{new} is the current position C_{old} plus the velocity V_{cursor} times the duration dt of one iteration of the program. Making the substitution for V_{cursor} yields a third equation:

$$3. C_{\text{new}} = C_{\text{old}} + G * e * dt$$

That is the totality of the simplest version of the model: a set of three simple equations or program steps which, evaluated over and over with the same pattern of target positions that the human participant experienced, duplicates the participant's actions in the tracking task above within 4.0% of their peak-to-peak range, in great detail. The model whose performance is illustrated in Figure 4 adds one more term to equation (3), a damping factor d , and that is what reduces the discrepancy between the model and the human participant to 3.6% with a maximum-difficulty disturbance, a small but consistent improvement. With this damping factor, the third equation (as it actually is implemented in the demonstration program) is

$$3'. C_{\text{new}} = C_{\text{old}} + [(G * e) - (d * C_{\text{old}})] * dt$$

It is remarkable that these simple equations do so well in simulating real behavior, considering that we are ignoring possible nonlinearities such as the Weber-Fechner law, potential noise in the system, continuously varying angles at the joints, and many other possible causes of poor performance of a simple linear model. In this light, examine the lower plot of Figure 3 again, showing the mouse/cursor positions of the real person and the model. The black trace representing the difference between model and person consists mainly of small high-frequency oscillations that are too fast for this system to suppress. Within the bandwidth of good control, the errors must be far smaller than the 3.6% to 4.0% of the range of target movement that is measured. There must be something fundamentally right about this hypothetical model.

A set of demonstration programs

Of course, tracking experiments involve only a very narrow range of behavior. They are a legacy from the engineering psychologists and physiologists of the 1950s, some of whom had worked in the war years on problems of aiming guns to track enemy targets. However, it must be acknowledged that we are in a position analogous to that of Galileo with his pendulums and inclined planes. As they are demonstrated and accepted, the principles of PCT can be applied to any behavior at all, but the most reliable experiments are still simple ones that can be implemented on a computer. Even so, many of the computer demonstrations of control processes that have been made publicly available by PCT researchers involve other kinds of behavior.

The most recent demonstrations are in a collection referred to as the 'LCS3 set', named for and explained in *Living Control Systems III* (Powers 2008). The program set, on a CD in the book, can also be downloaded from <http://www.livingcontrolsystems.com/lcs3.html> and runs on a Windows computer. The reader is advised to explore the set, because actually running the demonstrations is unquestionably the most effective way to learn what PCT is about (just as the best way to understand a float-valve regulator in a water clock is to see one in operation).

A second set of demos, also downloadable, are in a tutorial program called (for historical reasons) Demo3. The following refers to the demonstrations shown in various steps of this tutorial.

'Responding' to an invisible stimulus

The first three demonstrations in the "Demo 3" set explain how the mouse affects the cursor on the screen and the way numbers are used to determine positions. The first control task, step 4, is a tracking task: "compensatory" tracking in which the goal is to hold a cursor aligned with a stationary target and stabilize it against an invisible disturbance. After the 30-second experimental run is finished, a graph of the results appears. Figure 5 shows the result of one run. It differs from what you see when you exercise the program because the disturbance and the subject's resistance to it both differ from one run to another.


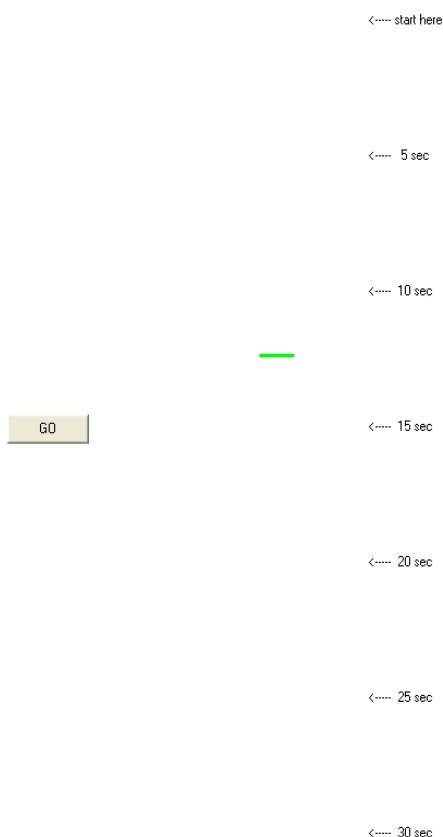


Figure 5: Compensatory tracking. Black line shows mouse movements, green line shows cursor movements. Target position is horizontal red line. The purple trace shows an invisible disturbance that varied during the run: mouse position relative to the centerline is equal and opposite to the disturbance at all times down to a moderate level of detail.

The main point of this demonstration is the way the participant moves the mouse so as to cancel the effects of an invisible disturbance (purple line) which, without these efforts, would move the cursor up and down so as to coincide with the disturbance trace. The green line shows the resultant cursor position during the 30 seconds of the experimental run (from left to right). There is no stimulus on the screen that corresponds to the purple disturbance plot, and clearly the green cursor line would be of no use in indicating the magnitude or direction of the disturbance. Thus there is no basis for claiming either that the mouse movements were a response to the cause of the perturbations of the cursor, or that the participant's brain was observing the disturbance and planning the actions needed to keep the cursor near the target. The information required to carry out either of those modes of action is simply not available in this demonstration. This is emphasized by the fact that one's performance improves over repeated exercise of these demonstrations, even though a new disturbance pattern is generated each time any step of the demonstrations is re-run. *Learning takes place, but there is no pattern of behavior to learn: what is learned is control.*³⁰

Hierarchical control through reference signals

In demonstration 5, the participant is told to make the cursor descend from the top to the bottom of a range marked off in seconds, so that it passes each mark on schedule. An unseen disturbance is still being applied to the cursor, so the participant must move the mouse so as to resist the effects of the disturbance and keep the cursor descending at a uniform rate. Figure 6 shows the appearance of the screen at about the 12-second mark. The participant is counting off the seconds, trying to make the green cursor move down so it passes each arrow at the time marked beside the arrow.



³⁰The control-system functions change for signals change. This is also how new neural nets, a process that largely involve

the relationships between neural actions of as-yet-unorganized which the neonate begins.

Figure 6: Demonstration 5 at about the 12-second mark. The green cursor begins at the start line and is moved up and down by unseen disturbances as the subject, resisting these disturbances, attempts to move it smoothly downward so as to reach each successive mark at the indicated time.

The graph for one run of demonstration 5 is shown in Figure 7.




Figure 7: One run of demonstration 5. The mouse cursor (green) approximates a steady progression from top to bottom (y axis) over the time elapsed (x axis). Effects of mouse movements (black) add to effects of the disturbance (purple) to accomplish this descending movement. No stimulus corresponds

to the disturbance, and no planning of mouse movements is possible.

The black trace showing the mouse movements executed by the subject's hand does not resemble either a mirror image of the disturbance pattern in purple or the pattern followed by the cursor in green. The result when the mouse movements are added to the disturbance is the requested slow movement of the cursor from high to low, as shown by the diagonal green line. But again, it is clear that this result cannot be accounted for in terms of responses to any visible stimulus, nor could the mouse movements have been planned in advance and then executed.

This demonstration shows what is meant by control through varying reference signals. The steady downward velocity of the green cursor bar is, according to PCT, the controlled variable for some system fairly high in the perceptual hierarchy, having to do with control of perceived rates of change of position. This system generates a slowly-varying reference signal³¹ for a lower system concerned with maintaining the cursor in some particular position against disturbances. This lower system is just as in the previous demonstration, except that now the 'particular position' where the lower system is maintaining the cursor is being changed through time by the higher system. In both cases, the lower system acts to make the cursor position match the reference position at all times (as well as it can) by resisting disturbances. The difference is that in the previous demonstration the reference is stationary, but in this demonstration a higher-level system is changing the reference signal in the direction from positive toward negative, so that the lower system creates the requested perception of a slowly descending cursor—by, of course, using a still a lower level of organization to move the mouse up and down in whatever way works to make the cursor on the screen actually descend.

That the cursor is under positional control at all times is shown by the way it resists a disturbance that is trying to push the cursor up and down, away from its steady descent. The mouse position varies oppositely to the disturbance, not only canceling it as in the first demonstration, but also adding enough additional variation to maintain the steady downward velocity. This happens automatically at the level of position control. The higher control system concerned with downward velocity does not have to do much to resist residual effects of disturbances. Most of the resistance has been accomplished at the lower level.

In addition to illustrating principles of control, these demonstrations are one way of testing the theory through attempts to prove hypotheses wrong—attempts to falsify them. The remaining demonstrations illustrate and test additional principles of PCT.

Challenging PCT with experiments and simulations

The first book-length treatment of what is now known as PCT (Powers 1973, 2002) was finished before the advent of inexpensive desktop computers and the exponential growth of computing speed and memory storage. Some 12 years later, the first interactive computer demonstrations of the principles of PCT began to take shape, in time for the first meeting of the Control Systems Group in 1985.³² At this meeting, a tracking experiment was shown in which a subject used a joystick to make a cursor on the computer screen track a moving target, the controlled variable being the separation of cursor from target and the reference condition (defined by instructions) being zero separation. Demonstration 4, the first one that we discussed above from the latest book in this series, recapitulates that demonstration.

³¹The output function of this control system could be classified as a simple form of 'central pattern generator'.

³²This 1985 tracking experiment is included in DEMO 1, part of the series of DOS programs listed under Resources.

This was also the first instance of a computer simulation of a PCT-type control system designed as a model of the person doing the tracking task. The parameters of the simulated control system were an integration sensitivity and a constant reference signal which were adjusted to make the performance of the model match the real person's performance with as little difference as possible. The RMS difference between modeled joystick movements and the real movements could be reduced to less than 10 per cent of the range of movement of the target. More recent versions have reduced the RMS error of fit to less than 4 per cent.

The most important aspect of this early simulation was that in a moving-target version it could be used with either a single smoothed-random disturbance moving the target, or with a second uncorrelated disturbance added that made the cursor movements differ randomly and by large amounts from the joystick movements. With the second disturbance acting, the subject would move the joystick in a way that corresponded neither to the target movements nor to the second disturbance, but was exactly the movement needed to minimize the tracking error.

This demonstration illustrated the important point that the behavior observed in a control situation generates a regular result without itself being regular. This is the main feature of PCT that distinguishes it from the calculate-and-execute models of control behavior: it is not possible for the organism to calculate in advance the joystick movements that will be required, because the disturbances are being generated from random numbers during the experimental run, and are unknown in advance. A calculate-and-execute model necessarily fails in the presence of unpredictable disturbances. This is only noticeable if working models are made and tested.

Among the 'LCS3' demonstration programs introduced so far, of particular interest is the demonstration called "Square circle". In this demonstration, a white dot is used by the participant moving a mouse to trace around the sides of a red square. At the end of one complete tracing, the path of the mouse is revealed: it is a circle. In a variant mode, the revealed path is a triangle—a bit more difficult to execute, but even more unexpected by the participant. The point is to show that what a person *experiences* as his or her own behavior is actually the perception that they are controlling, not their actions, which are often markedly different from it.

Other demonstrations are distributed with (Powers 2008) or can be downloaded from the Internet³³. In one of these, the participant's task is to keep a small green circle aligned inside a slightly larger red circle in one corner of the screen. A white tracing shows the actual path of the mouse, which at the end of the run is seen to spell out in script the word "hello". This is caused by a patterned disturbance of the green circle which (invisibly) traces out "hello" upside down and backward. Even though the observed behavior of the participant (the movement of the mouse) is essentially unrelated to the control task being accomplished—the behavior as the person experiences it—nonetheless, when we overlay the disturbance on the mouse movements they are very highly and negatively correlated (estimated to be in the -0.99s).³⁴

The first demonstration of the LCS3 set may be the most philosophically interesting. A red ball is shown drifting left and right while it rolls vertically and changes shape from short and wide to tall and thin. Each aspect is affected by a smooth disturbance, the three disturbances being nearly uncorrelated.

³³See Resources

³⁴A similar result is seen with Richard Marken's simulation of a bimanual control experiment described in Mechsner et al. (2001). The reader may exercise the demonstration at <http://www.mindreadings.com/Coordination.html>

The mouse affects all three variable aspects of the ball—shape, orientation, and position—at the same time and by the same amount. The participant's task is to pick one of those aspects and keep it constant: shape as *round*, position as *centered*, or orientation as *level*. That this can be done at all is of considerable interest, but of equal interest is the fact that the computer can determine reliably which single aspect is being controlled and which two aspects are varying as side-effects. The computer deduces which effect of the action was intentional and which others were accidental. It is very seldom wrong even if control is poor.

Intention, in PCT, refers not to behavioral acts but to the consequences of those acts. The intended consequence of controlling the orientation of the red ball is to keep its axis pointing toward the viewer. Because each aspect of the ball is being influenced by a different pattern of disturbances, the same actions that stabilize orientation can't simultaneously stabilize position or shape; in fact they increase the variance of those two variables because the actions aren't systematically opposed to the relevant disturbances. The result is a rather puzzling combination of correlations: the actions that stabilize orientation correlate almost perfectly (-0.99) with the disturbance that tends to alter orientation, yet those actions and those disturbances show only a low correlation, close to zero, with the orientation that is being controlled. The mouse movements always correlate much better with the aspects that are *not* being controlled.³⁵

To return to a subject at the beginning of this paper, a general-purpose demonstration called “LiveBlock” shows a basic control system as a ‘live block diagram’. Here we have a control system with an adjustable transport lag, time constant, gain factor (output amplification), and environmental feedback factor, plus an adjustable reference signal and disturbance. The model runs continuously in the background so the effects of changing system parameters and independent variables can be seen as they occur. The method of stabilizing a system with time lags in it is illustrated, as are many other basic properties of a negative feedback control system. It is hoped that this demonstration can finally counteract many of the false ideas offered over the past 60 years about the limitations of negative feedback control as a model of behavior.

The discovery of the principles underlying a phenomenon for which there is a widespread naive (and incorrect) explanation is always of particular interest to many readers. Such is the discovery of what baseball players are really doing when they catch fly balls, as reported by McBeath, Shaffer, & Kaiser (1995), and the explication of their naive projection of perceptions as reported by Shaffer & McBeath (2005). The simulation at <http://www.mindreadings.com/ControlDemo/CatchXY.html> demonstrates this.

The methodology of PCT research

Near the beginning of this paper, we made note of an inherent difficulty of the experimental investigation of living things. An organism controls its own perception of some aspect of its environment, but that privileged point of view from inside the observed organism is unfortunately not available to scientific observers of other organisms. As observers of a different organism we do not have access to that perception, we only have our own perceptions from our own points of view, external to the organism. For that reason it has been crucially important to devise tests for determining which aspects of its perceived environment the organism is controlling.

³⁵Trying to explain this fact is a good test of one's understanding of PCT.

The principal datum in PCT methodology is the controlled variable. All of the demonstrations that we have reviewed have clearly displayed three variables: the controlled variable (e.g. distance between the mouse cursor and the target), the disturbance (producing movements of the mouse cursor independent of the user's movements of the mouse), and the relevant behavioral actions (indicated by the changing mouse position). Obviously, the disturbance can't be identified until we know just what the controlled variable is and how mouse movements affect it. If the user, unbeknownst to us, is ignoring the moving target and instead trying to draw a large circle with the mouse cursor, the measured and graphed results will not make sense to us. There will be no relationship between what is expected to be controlled (the position of the cursor relative to the target) and what the subject is actually controlling (following the outline of an imagined large circle). Even in this simplified, artificial, two-dimensional laboratory environment it is difficult to see what is actually under control; we would expect more naturalistic settings to present still more difficulties. Yet the technique for determining what perceptual variable is being controlled is essentially the same everywhere. The requirements are few. We must be able to make intelligent estimates of which aspects of the environment the organism can perceive and influence with its activities, and we must be able to also influence those aspects of the environment.³⁶

The fundamental step of PCT research, the Test for controlled variables, is the *slow and gentle* application of disturbing influences to the state of a variable that the researcher surmises is already under control by the observed organism. If the organism changes its action and thereby prevents the disturbing influence from having the expected effect on that variable, that is strong evidence that the experimental action disturbed a controlled variable. It may take a number of variations of the disturbance to isolate just which aspect of the environmental situation is under control. This strategy gives a definition of the controlled variable in terms of the observer's way of perceiving the organism and its environment.

In order to build working generative models of behavior, like the simulations we have been exercising, there is one further requirement. We must be able to measure these influences affecting the state of the environmental variable that we have decided to test. Until a simulation produces very nearly the same numbers as were produced by measurement, it needs refinement; and when it does, we have a strong basis for the claim that the simulation models essential aspects of the unseen internal structure of the organism whose behavior we measured.

Our understanding of the inner workings of a hierarchical perceptual control system reorients our thinking in a number of fields.

Relation of PCT to conventional psychology

PCT terms, as seen in Figure 4 above, originate in the field of engineering and are based on the concept of a negative feedback control system. The following table of correspondences suggests how PCT could integrate the field of Psychology.

Table 1. Correspondence of key terms in PCT to concepts in conventional psychological theories. These are not equivalences, but rather indications that the correspondent terms have similar referents, albeit from different points of view and with different implications for theory and practice.

³⁶Insofar as each level of perception is the environment for the next higher level in PCT, the test for the controlled variable may involve "environmental" variables that are available only to the person being studied (as in Robertson et al. 1999).

Conventional Psychology	PCT
Stimuli	Perceptual input; when measured, input quantities
Sensations, Perceptions, Concepts	Perceptual signals at different levels of the perceptual hierarchy
Motives	Reference signals
Drives	Error signals
Responses; Behavior	Observable means of control; when measured, output quantities
Learning and Development	Reorganization; Changing existing control systems; Altering connections from higher to lower control systems in novel situations.
Cognitive map	The subset of perceptions from a person's perceptual hierarchy, which is involved when a person is controlling many variables at once

Some of these terms correspond, but with essential and radical differences of meaning, as for example *stimuli* and *responses*. Some PCT terms, such as *error signals*, have no correspondence in conventional psychology. However, error signals relate indirectly to terms used in conventional psychology such as stress, distress, and emotions. An account of emotions, an obviously important topic, is beyond the scope of this paper.³⁷

The PCT way of doing research differs from conventional psychological approaches in the criteria for what is a good result. The correlation between actual behavior and model behavior in the PCT work done with pursuit tracking is in the .90 and higher range. This shows that it is possible to obtain much higher correlations than is typically obtained in conventional psychology research when one has an understanding of the mechanisms involved within the person as described by a correct functional model. The 'revolutionary' aspect of PCT is the message that one should continue research until the correlation between theory and data moves into the .90 range and higher so that one can make predictions for an individual, not just 'on the average'. By striving for this goal, psychology will be on much firmer ground when applications are attempted.

The PCT way of describing a person is useful in various applications of PCT, including psychotherapy.

PCT applied to psychotherapy

The field of diagnosing and treating psychological problems affords an excellent example of how PCT can provide a unifying framework to an otherwise fragmented area of research and practice. It is widely recognized, for example, that current classificatory systems of psychological disorders such as the

³⁷Briefly, our present view is that an emotion perception is constructed from sensed physiological conditions—the sensations that we call feelings—in combination with other perceptions that are those being controlled (or are associated with them). An error signal in a control system has two coordinated effects. Firstly, it provides reference signals for the actions which reduce the error signal. Secondly, it activates the endocrine and autonomic nervous systems to prepare the body for that activity. Thus, feelings are always potentially present when we are controlling, because there is always some body state present when a person controls. See Powers (2005), Chapter 17.

DSM IV-TR (American Psychiatric Association, 2000) do not easily map onto the lived experience of psychological distress. Despite the invocation of concepts such as ‘comorbidity’,³⁸ there is a growing awareness that this system of classification is unsatisfactory in important ways. In recent years there has been great interest in processes that are said to be ‘transdiagnostic’ across DSM categories as an explication of underlying pathways by which diverse symptoms become manifested (Harvey, Watkins, Mansell, & Shafraan, 2004). A PCT account of such processes explains why this approach has merit and why categorizing symptoms is problematic.

We have already outlined in detail how behavioral output varies in order to control perceptual input. Consequently, internal problems of control do not give rise to recognizable, standardized symptoms. That is why there is so much variation within current classificatory categories of symptoms. This also accounts for the lack of clear differentiation between categories. In the real world, where unpredictable disturbances to control occur, behavior must vary as a person repeats attempts to solve a given problem. The appearance of symptom patterns is analogous to the appearance of constellations in the night sky. They are arbitrary groupings in the eye of the beholder that reflect no underlying order or structure. Categories of behavior—that is, of variable control system outputs—cannot reveal the order or structure of internal malfunctions.

There is a developing acknowledgement that it is the causes of distress associated with particular symptoms rather than the distress or the symptoms themselves that need to be understood. Kazdin (1999) argues that functional impairment, rather than symptoms, is the main reason people seek psychotherapy. Large-scale population surveys, for example, have demonstrated that many people experience psychotic symptoms without requiring treatment (Bentall, 2009, p. 107). Auditory hallucinations are one type of symptom that can be accompanied with or without the debilitating experience known as psychosis. Who has never had cause to complain of being unable to stop a tune from replaying itself in imagination?³⁹ Auditory hallucinations can even be helpful. A singer who can’t mentally ‘hear’ a note before singing it will not sing very well. However, when a person is distressed by auditory hallucinations, the hallucinations are problematic. In PCT, the problems which are discussed in therapy are the ones which are distressful to the person; other people may be distressed by the person’s symptoms, but this is considered to be more of a social problem than a psychological problem, to be settled by social processes like law, mediation, conflict resolution, or negotiation, not by therapy. PCT provides no way to decide which party in a social conflict needs therapy.

Each person’s problems are understood to be unique to that person. There is no justification in PCT for applying what worked for someone else to this person, just because that other person had similar symptoms or DSM IV-TR diagnoses. The same symptomatic behavior can result from an entirely different set of internal conflicts; similar conflicts can lead to entirely different sets of symptoms.

Psychotherapy has focused, understandably, on pathology. PCT contributes a useful perspective in understanding psychological disorders by first providing a model of satisfactory psychological functioning. PCT portrays dysfunction in terms of disruption of successful control. Conflict between control systems, as we noted earlier, is a problem because it effectively removes the control abilities of both systems. Conflict is usually transitory. It is when conflict is unresolved and becomes chronic that the symptoms recognized as psychological disorder become apparent. Distress is the experience that

³⁸A relatively empty term meaning simply that one or more disorders or diseases are diagnosed in addition to a primary one.

³⁹Discussion of how imagination is modeled in PCT is beyond the scope of this paper. See Powers (1989:277).

results from the person's inability to control important experiences.⁴⁰ It clearly can't be "treated" as if being distressed is itself the problem. Restore the ability to control and the distress will disappear.

Conflict, as it is conceptualized in PCT, occurs between two control systems at the same level. These control systems, however, are located within a hierarchical network of control systems so their conflicted arrangement will influence and be influenced by lower and higher level systems. This account of psychological distress may explain why no reliable biological explanations of mental illness have ever been discovered. From a PCT perspective, control systems that are in conflict are not dysfunctional or broken. In fact, it is quite the reverse. The better the control systems would control in isolation, the more intense the conflict will be when they oppose each other. Some kinds of mental illness, perhaps most, may be a result not of broken brains but of well functioning control systems locked in chronic conflict.

It is the hierarchy that provides a clue as to where treatments should focus to help conflicts resolve. Systems at one level receive their references from the next higher level. When two control systems are conflicted, it is the signals being sent to each from the next highest level that need to be altered. The senders, not the receivers, need reorganization. Otherwise competent control systems are being misused by higher systems.

Most therapies assume that attention to problems facilitates change, and people in general tend to be most aware of painful or dramatic consequences of conflict. This is seldom helpful in itself. Attention is drawn to the symptoms rather than the causes of loss of control, symptoms such as apathy, confusion, fear, or despair. Often a person will try to strengthen the 'good' side of a conflict, which usually just makes conflict more extreme because the other side resists the effort to change and starts to look good for other reasons.

PCT suggests that resolving conflicts requires the learning process of reorganization that was explained previously. Initially, reorganization can do this by modifying components of higher-level control systems that send conflicting reference signals to others at a lower level. That change may well result in other reorganizations being needed at higher and lower levels. The therapeutic approach that is based on the principles of PCT is called the Method of Levels (MOL). Its principle is to redirect attention 'up a level' to the control systems responsible for generating the conflict, away from a preoccupation with the symptoms and the immediate efforts on both sides of the conflict. It appears that reorganization and awareness are linked in such a way that the systems in awareness become the particular focus of the reorganizing processes. Reorganization is an automatic response to intrinsic error; it can't be controlled voluntarily. But awareness can be redirected, and apparently the focus of reorganization follows awareness to a useful extent.

Consider wanting to stop smoking to avoid lung disease and at the same time wanting to continue it to relieve withdrawal symptoms caused by smoking; or wanting to leave a partner to avoid abuse and, at the same time, wanting to stay with the partner for the sake of love. Ultimately, attempts to modify the actions of conflicted systems, or to give preference to one goal by will power, cannot permanently alter the conflict. The conflict is resolved only when awareness is shifted to the level above the conflicted systems so that reorganization can be directed to the systems creating the conflict—the systems that are establishing these conflicting goals. Hence the name, the Method of Levels.

⁴⁰Other problems that can also arise, such as feeling overwhelmed by environmental forces, require more extended discussion than is possible within the scope of this paper.

For the person in therapy, MOL is an experience of describing in detail a current area of distress to a therapist who understands PCT. The therapist's approach is a blend of questioning about subjective experiences and selectively drawing a client's attention to seemingly tangential or peripheral subjects, usually on the basis of comments the client makes—comments that the therapist familiar with PCT recognizes as possibly reflecting involvement of a higher-level system. In this way, clients show the therapist what path to follow, and when the therapist helps them focus in the right place, their own reorganizing capabilities generate new perceptions and goals that may resolve the conflict, or uncover the deficiencies that cause trouble.

MOL has been used over a number of years by different clinicians in a variety of clinical settings. Evaluations have been conducted of the way in which MOL is experienced by routine clients in routine clinical contexts.⁴¹

It is telling that despite the demonstrated effectiveness of various psychotherapies there is still no generally accepted account of how these effects are achieved, and in fact, it has been shown that psychotherapies based on quite different models of disorder can have similar effects (e.g., Wampold, 2001). There has been an increasing call to move away from developing new techniques and strategies based on diagnosis and instead to focus on underlying common principles and mechanisms. PCT provides a common underlying process (conflict) and a common change mechanism (reorganization) that might be particularly significant for understanding this peculiar situation. The paradigm of perceptual control may provide the means to make sense of these otherwise puzzling results.

Recent research in neurobiology has indicated that psychotherapy can have effects in the brain that are similar to the effects that pharmacology achieves.⁴² Again, this result would come as no surprise from a PCT perspective. The hierarchy of PCT is a hypothesized neuronal architecture for the full range of control. At every level, brain chemistry and brain function are aspects of the same processes.

It is certainly the case that, at this stage, PCT raises more questions for research in this field than it answers. Do conflicts at different levels of the hierarchy, for example, result in different types of pathology? Does the rate of reorganization affect the experience of internal conflict? What influences the mobility of awareness such that some conflicts are resolved satisfactorily while others become chronic? The possibilities for new research, as usual with new ideas, proliferate.

While some of the propositions about the application of PCT principles to psychotherapy remain speculative, there is also indirect but strong evidence for this approach. Problems of control (such as behavioral control, impulse control, emotional control, and thought control) are widely recognized as important in psychological functioning. Many approaches to psychotherapy use conflict formulations to explain psychological distress. Many approaches also discuss the importance of awareness in resolving problems as well as recognition of the need to consider problems from higher levels of thinking (such as important life values or belief systems). Finally, there is a growing body of literature that recognizes that the change involved in the resolution of psychological distress is not a linear or predictable process.

In fact, full-time MOL practitioners, almost all of whom came from other schools of thought, agree that MOL is probably an explanation of why other therapies succeed when they are successful, and when they fail, why that happens. In the unanimous view of these practitioners, methods that succeed

⁴¹For details about MOL and its use in clinical practice, see Carey (2006).

⁴²For example, Schwartz, (1996) on OCT symptoms.

are consistent with MOL, and those that fail are not. Many therapists have independently developed methods that come close to MOL, simply by weeding out what doesn't work. For some, such as Rogerians, a switch to pure MOL would involve relatively minor changes. For others, of course, such a switch would call for so many deviations from customary practice that an orderly transition would be essentially impossible.

Exploring psychological disorders and their treatment from the perspective of perceptual control provides a new direction for psychotherapy researchers and practitioners. There is a growing possibility that it will enable a clearer understanding of the nature of psychological distress that is developed from a model of function rather than dysfunction. It may also promote the distillation of the important components of psychotherapy such that therapists can be clearer about their roles and treatments can become more efficient. Moreover it can, and already does, provide a guide regarding the purpose of psychotherapy. PCT, then, will have an impact on long standing debates such as the dodo bird hypothesis.⁴³ A unifying focus such as the one provided by PCT will allow a more consistent and coherent approach to emerge that will go a long way towards preventing the debilitating impact of psychological distress that is currently on the increase in many countries.

There may also be other implications of this approach that cannot easily be predicted at this stage. Perhaps the stigmatizing nature of mental illness will change with a more accurate explanation of these problems that is inherently psychological (yet firmly grounded in neurobiology) and intuitively optimistic and hopeful. The nature of the delivery of psychological treatments might also change as researchers and clinicians become more familiar with the reorganizing capabilities of individual systems. Perhaps we will learn to use both psychotherapy and pharmacotherapy more judiciously. While the outcomes may not be entirely obvious there seems to be sufficient justification at this stage to step into the paradigm of control and to build our knowledge of the mechanisms of psychotherapy from the foundations of these functional and rigorously tested models.

Directions in PCT research

Despite being more than 50 years old, and having the principles of negative feedback control and many known facts of neurology and physiology behind it, much of PCT is still provisional and hypothetical. The main lines of research that PCT theoreticians are concerned with have to do not with applications but with testing the core concepts of the theory. The demonstrations cited earlier were constructed as deliberate challenges to the basic concept of control, in which the assumptions of the theory were made explicit as properties of a simulated system. In some cases they were used to make quantitative predictions of behavior for experimental test. The fact that it is possible to give confident descriptions of what will happen when any randomly-selected person participates in the interactive demonstrations shows the progress that has been made at this basic level.

It is now quite clear that simple kinds of behavior are well described in great detail and under many variations in experimental conditions by a negative feedback control model. But the theory contains hypotheses about much more than a single control system operating in a simple experiment. Eleven

⁴³This is the proposal that all therapies are equally effective, depending upon the practitioner, and "all must win prizes" (alluding to a pronouncement by the Dodo Bird in *Alice in Wonderland*). See http://en.wikipedia.org/wiki/Dodo_bird_verdict.

levels of control have been proposed⁴⁴. Do they really exist? If they exist, have they been correctly defined? Are there too many or not enough levels? It is proposed that reference signals are derived from remembered perceptions, a possibility raised by the way people can remember goals and act to match current perceptions to perceptions that have been experienced some time in the past. How long do remembered reference signals remain accurate? Are there differences at different levels of perception? Does imagining a perception interfere with present-time control? Does degree of awareness or attention affect the quality of control, and if so, in what way? Does reorganization really occur—can we detect it by measuring control parameters at frequent intervals? Will the changes prove to be random? Will they become smaller as an asymptote is approached? Does reorganization really follow awareness? Is it really driven by intrinsic errors, such as error in normal biochemical and homeostatic control systems?

The ‘pandemonium’ form of the model, in which control systems control only scalar variables in one dimension, is probably wrong in several regards, the main one being a lack of interaction between different perceptions. But there is a way of transforming this model into one with more realistic functions that correspond to sensory and motor nuclei or larger areas in the brain, and studies of interactions will give clues about what the groupings really should be.

Another large area for research concerns the way higher-level systems use lower-level systems. In PCT, the only link from higher to lower systems is the reference signal, but there are reasons to believe that higher systems can vary some parameters of lower-level control systems, for example the sensitivity of an output function to error signals (reducing the sensitivity to zero is equivalent to turning the control system off—but there are other ways to do that as well). It is also possible that higher systems can act to alter lower-level perceptual functions, changing the very nature of a controlled perception so the same perceptual signal represents a different aspect of the external world.

The inability of neural signals to change sign (a negative rate of firing is impossible) requires that bidirectional control be accomplished by pairs of control systems, each acting in only one direction. This implies that paired systems employing agonist-antagonist muscles to accomplish bidirectional motor control might have measurably different properties in the two directions. And more than two directions of control in one dimension are possible; muscles set at varying angles can contribute to X and Y control when they are shared by several control systems.⁴⁵

Adaptation is handled in PCT by the basically random reorganizing process. However, control systems exist that provide more systematic ways of gaining control when it is difficult (that is, solving problems), and these have been explored to some extent. In theory, one result of random reorganization is to create control systems that are so much more effective than the random process that errors in a given area of experience no longer get large enough to activate reorganization. True or false?

If older theories had been subject to the kind of testing that PCT has undergone even at this early stage, most of them would have been discarded long ago. The flaws are simple and obvious (as they were for PCT in the beginning). When the premises of a theory are taken as unquestionable facts without even being tested (as in calling every action a “response”), flaws are perpetuated and go undetected, and the quality of science suffers. Much as we have tried to avoid that error in PCT, our success is no doubt partial at best. We have tried to be skeptical about what seem our best ideas, to demand demonstration

⁴⁴Briefly: intensities, sensations, configurations, transitions, events, relationships, categories, sequences, programs, principles, and system concepts.

⁴⁵Powers (1979, Part 3; p. 10 of the reprint as `byte_aug_1979.pdf`).

of every claim. Whether that has worked remains to be judged.

The approach that has led through many years to this paper has entailed going through all the assumptions on which PCT is based and, one by one, testing to see if they can be supported or ruled out by experimental data. That task is far from finished. The current form of the model is the result of half a century of challenging the theory, with many failures and subsequent improvements. At first, just one person was doing this work, and in terms of actual computer modeling, never more, until recently, than half a dozen. One person or a small handful cannot develop or explore a new idea in all the necessary ways. Real progress and the building of a real science require the vast resources of a whole discipline, with thousands upon thousands of independent minds, each reorganizing in unpredictable ways, looking for difficulties with and improvements of the theory. PCT as it stands today is no more than a pilot study, a definition of a problem and a possible kind of solution. It is difficult to imagine what will happen when the full power of a scientific community is turned to developing it in all the directions that are possible. Will it ignominiously disappear? Or will it turn into the direction for all the life sciences before the 21st Century is out?

Resources

Computer simulations

The most recent set of simulation and demonstration programs designed for and included with (Powers 2008) can be downloaded at <http://www.livingcontrolsystems.com/lcs3.html>.

Earlier DOS and Windows programs by Powers can be downloaded at http://www.livingcontrolsystems.com/demos/tutor_pct.html

Programs by R. Marken at <http://www.mindreadings.com/demos.htm> can be run in a web browser.

Reference websites

Introductions and discussions of Perceptual Control Theory can be found at several web sites. Some of the major reference sites are: <http://www.livingcontrolsystems.com>, <http://www.pctweb.org/>, <http://www.mindreadings.com/>, <http://www.perceptualcontroltheory.org/>

References

American Psychiatric Association. 2000. *Diagnostic and statistical manual of mental disorders DSM-IV-TR*, Fourth Edition (Text Revision). Arlington, VA: APA.

Ashby, W. Ross. 1952. *Design for a brain*. London: Chapman & Hall. [Repr. 1960, New York: Wiley. Online: <http://www.archive.org/details/designforbrainor00ashb>.]

Ashby, W. Ross. 1956. *An introduction to cybernetics*. London: Chapman & Hall. [Online: <http://pcp.lanl.gov/ASHBBOOK.html> & <http://pespmc1.vub.ac.be/books/introcyb.pdf>.]

- Bennet, S. 1979. *A history of control engineering 1800-1930*. New York: Peter Peregrinus Ltd.
- Bentall, R. P. 2009. *Doctoring the mind: Is our current treatment of mental illness really any good?* New York: New York University Press.
- Bernstein, Nicolas. 1967. *Coordination and regulation of movements*. New York: Pergamon Press.
- Brillouin, Leon: 1953. "Negentropy principle of information", *J. of Applied Physics*. 24.9:1152-1163.
- Burke, Robert E. 2007. "Sir Charles Sherrington's *The integrative action of the nervous system*: a centenary appreciation". *Brain* 130.4:887-894. doi: 10.1093/brain/awm022
- Carey, Timothy A. 2006. *The Method Of Levels: How To Do Psychotherapy Without Getting In The Way*. Hayward, CA: Living Control Systems Publishing.
- Chapanis, Alphonse. 1961. "Men machines, and models". *American Psychologist* 16:113-131.
- de Rosnay, Joël. 2000. History of cybernetics and system science. Principia Cybernetica Web: <http://pespmc1.vub.ac.be/cybshist.html>.
- Harvey, A. G., E. Watkins, W. Mansell, & R. Shafran. 2004. *Cognitive behavioural processes across psychological disorders: A transdiagnostic approach to research and treatment*. Oxford: Oxford University Press.
- James, William. 1890/1950. *Principles of Psychology*. New York: Henry Holt/Dover Publications.
- Kazdin, A. E. 1999. The meanings and measurement of clinical significance. *Journal of Consulting and Clinical Psychology*, 67(3), 332-339.
- Kozak, M., Marsh, A. A., & D. M. Wegner 2006. "What do I think you're doing? Action identification and mind attribution". *Journal of Personality and Social Psychology*, 90.543-555.
- Landels, J.G. 1978. *Engineering in the Ancient World*. Berkeley: University of California Press; London: Chatto & Windus.
- Marsh, A. A., M. N. Kozak, D. M. Wegner, M. E. Reid, H. H. Yu, & R. J. R. Blair. 2010. "The neural substrates of action identification". *Social Cognitive and Affective Neuroscience* 2010; doi: 10.1093/scan/nsq004.
- Mayr, Otto. 1970. *The origins of feedback control*. Cambridge, MA: MIT Press.
- Mechsner, F., K. Kerzel, G. Knoblich, & W. and Prinz. 2001. "Perceptual basis of bimanual coordination". *Nature* 414:69-73.

McBeath, M. K., Shaffer, D. M., & Kaiser, M. K. 1995. "How baseball outfielders determine where to run to catch fly balls". *Science*, 268:69–573.

Piaget, J. 1937/1954. *La construction du réel chez l'enfant / The construction of reality in the child*. New York: Basic Books.

Powers, William T. 1973. *Behavior: the control of perception*. 1st ed. Hawthorne, NY: Aldine de Gruyter.

Powers, William T. 1979. *The nature of robots*. Part 1, "Defining behavior". *BYTE*, 4.6:132-144 ; Part 2, "Simulated control systems", *BYTE* 4.7:134-152 ; Part 3, "A closer look at human behavior", *BYTE*. 4.8:94-116 ; Part 4, "Looking for controlled variables", *BYTE*. 4.9:96-112. [Online at http://www.livingcontrolsyste.ms.com/intro_papers/ as four PDF files.]

Powers, William T. 1989. *Living control systems*. New Caanan, CT: Benchmark Publications. [Selected papers 1960-1988.]

Powers, William T. 1998. *Making sense of behavior*. New Caanan, CT: Benchmark Publications.

Powers, William T. 2005. *Behavior: the control of perception*. 2nd expanded ed. New Canaan: Benchmark.

Powers, William T. 2008. *Living control systems III: The fact of control*. New Canaan, CT: Benchmark Publications. [Includes a CD with a set of demonstration programs.]

Robert A. Rescorla 1988. "Pavlovian conditioning; it's not what you think it is". *American Psychologist* 43.3:151-160.

Robertson, R. J., D.M. Goldstein, M. Marmel, & M. Musgrave. 1999. "Testing the self as a control system: Theoretical and methodological issues." *International Journal of Human-Computer Studies* 50:571-580.

Schrödinger, Erwin. 1944. *What is life: The physical aspect of the living cell*. Cambridge: Cambridge University Press.

Shaffer, Dennis M., & Michael K. McBeath. 2005. "Naive beliefs in baseball: Systematic distortion in perceived time of Apex for Fly Balls". *Journal of Experimental Psychology: Learning, Memory, and Cognition* 31.6:1492–1501.

Sherrington, Charles S. 1906. *The integrative action of the nervous system*. New York: Scribners. [Repr. 1947 New Haven, CT: Yale University Press.]

Schwartz, J.M. (1996). *Brain lock*. New York, NY: HarperCollins Publishers.

Vallacher, R. R. & D. M. Wegner. 1985. *A theory of action identification*. Hillsdale, NJ: Lawrence Erlbaum Associates.

van de Rijt, Hetty, & Frans Plooi. 2008-2010. *The Wonder Weeks. How to stimulate your baby's mental development and help him turn his 10 predictable, great, fussy phases into magical leaps forward*. Arnhem, The Netherlands: Kiddy World Promotions B.V. [First published in 1992 as *Oei, ik groei!* by Zomer & Keuning Boeken B.V., Ede and Antwerp. Updated with interactive participation of researcher-parents.]

Wampold, B. E. (2001). *The great psychotherapy debate: Models, methods, and findings*. Mahwah, NJ: Lawrence Erlbaum Associates.

Wiener, Norbert: 1948 [1961]. *Cybernetics or control and communication in the animal and the machine*. Cambridge, MA: MIT Press. [Repr. with new Introduction 1961.]