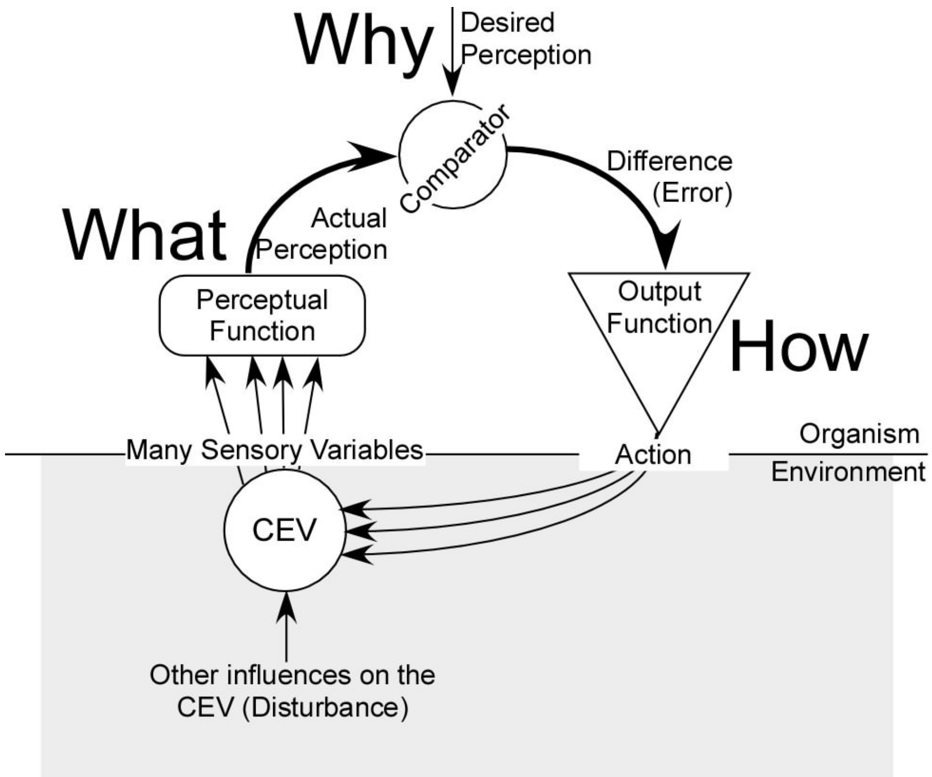

Figure 1.2. The canonical control loop. If the actual perception is not what is wanted, the difference, known as "the error" is processed by the output function to produce action, which through the Environmental Feedback Path influences the CEV, which may be subject to other influences collectively called the Disturbance. The sensory input is processed by the Perceptual Function to produce the actual perception. The form of the Perceptual Function defines the CEV.

2.5 The Behavioural Illusion versus Model fitting

The "Behavioural Illusion" (BI) is easily described, but less easily analyzed. It is the illusion that the "response" that follows a "stimulus" is determined by the processing that occurs inside the organism. This illusion is seductive because obviously if the organism were not there, the response would not occur, so something about the internal structure of the organism that makes it different from a chair or a rock must be producing the "response" to the "stimulus".

But, you may say, does not a spring produce a "response" of lengthening when you pull on it? Yes it does, and experiences like that reinforce the strength of the illusion when the same kind of thing seems to happen when a person is disturbed by a "stimulus". As Shakespear's Shylock says: "If you prick us, do we not bleed?". When the front door bell rings, do I not get up and go to see who is there? — No, that's not the same thing at all. The words are similar, but that's as far as it goes. I may go to the door, and usually do, but on this occasion maybe I believe I know who is there, and do not want them to know I am at home. There's a difference between me and a spring or a rock. I control, and they don't.

The main feature of the BI is that it occurs only when something is controlled. When physical inanimate objects change as a consequence of applied influences, their changes are in principle completely determined by their physical-chemical and structural properties, and can be calculated in advance. But if a person is presented with a "stimulus", they may "respond" in different ways depending on what they want to do and how they perceive the situation.

One of my colleagues told me a story long ago about participating in an experiment when he was an undergraduate. He had been asked to keep a stylus as long as possible on a sensitive area on a rotating disk, so, thinking it was an intelligence test rather than the mechanical skill test intended by the experimenter, he dismounted the disk and held the stylus on the spot. The experimenter, having the "stimulus" of an impossibly perfect result in the test, produced the "response" of expressing anger at my "stupid" friend for not producing the correct "response" to the "stimulus" of the rotating disk. Looking through the lens of Perceptual Control Theory, we can see that each of them was actually controlling perceptions that were not what the other perceived them to be controlling.

Before analyzing similarly complex interactions (the underlying principles of which are examined through much of Part 2 of this book when we deal with "protocols") let's consider the Behavioural Illusion for a single control loop, following Powers (1978). Figure 2.4 shows a simple control loop, but complicates it by introducing unspecified functions into each of the connecting "wires". These functions are monotonic, meaning that for every possible input value there is only one possible output value, and if the input value increases by a small amount, so does the output value. As may become evident, most of these functions have very little effect on what can be seen by an outside observer or experimenter when the perception is well controlled.

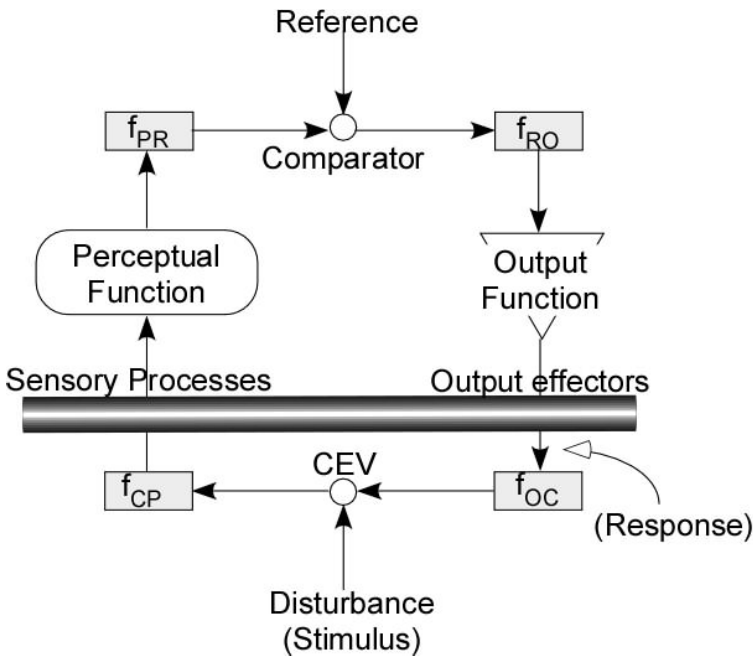

Figure 2.4 A simple control loop, showing possible functions that might occur at different places around the loop. In a typical experiment, the "stimulus" is what a PCT experimenter would call a "disturbance", while the "response" is what the PCT observer would call "output to the environment" (not the output of the control unit's output function).

Following Powers (1978), conceptually but not in detail, we observe that if control were perfect (a physically unrealizable condition), the output of the function f_{OC} would be exactly opposite to the stimulus (change in the disturbance) and the value of the CEV would remain perfectly constant, no matter what the other functions and internal processes of the controller. Everything the experimenter would observe would be a property of f_{OC} and nothing else. The Behavioural Illusion is that what the experimenter observes is a property of the internal organization of the subject, whereas it is actually a property of the environment, f_{OC}.

At the other extreme, if f_{OC} were a switch that had been turned off, so that there was no feedback connection, then the response would be a combined function of f_{CP} and the internal processes of the subject organism. The experimenter would be fully justified in claiming that the relationship between stimulus and response provided evidence to what processes occurred within the organism. But why, then would the subject act at all? After all, *Behaviour is the Control of Perception*, is it not?". The answer is that the action is part of the control of some other perception.

Everyday observation suggests we do act differently depending on the values of perceptions we do not control. If we want to go out, we may take an umbrella if we do not want to get wet, and we perceive the sky to be dark grey, but not if we see the sky a clear blue. Going out is an action in the control of a perception of our location. Taking an umbrella is one way of going out, not taking one is another. The difference allows a person to control for an an imagined perception of future wetness, using an uncontrolled (and uncontrollable) present perception of an aspect of the current environment. Taking or not taking an umbrella when we control for our location to be somewhere outside depends both on our reference value for our perception of wetness and on the uncontrolled perception of the sky. It looks like a simple stimulus-response: see grey and rain (stimulus) → take umbrella (response), but is part of a more complex control loop. After all, on another day the person might like getting wet, but not like being exposed to the sun.

For many experimental cases, we need no such complexity. The simple control loop of Figure 2.4 is nearly enough. The "subject" (who we can call Sean) controls some perception (the output of f_{OC} in the Figure) to some value. The experimenter (who we call Ethel) does something that affects the CEV (though without necessarily being precise about exactly what Sean will observe), and observes something about his action (without necessarily being precise about observing all and only those action components that affect the CEV). Some part of what the Ethel does will influence the CEV[1], some part of what she observes will be Sean's actions to influence the CEV as part of the control loop, and Sean will be only partially successful in controlling the perception. All of these imprecisions affect the experimenter's observations.

Let us ignore Ethel's almost certain imprecisions, and assume that the she is absolutely accurate both in influencing Sean's CEV and in observing what he does in controlling his perception of it. A subject, in any physically possible case, cannot control perfectly. This means that the "response" input to f_{OC} does not produce exactly the opposite of the "stimulus" (change in the disturbance). It is not quite the "right" response, but it is the response produced by the internal processes that include all the other f_{XY} processes that may or may not exist, as well as the Perceptual Function, the Output Function, and the comparator process. Only this failure of control allows any experimenter to infer something about what happens inside the organism in an experiment. The success of control enforces the BI, but control imperfection allows at least some access to the processes in the rest of the loop.

A PCT researcher would propose a model of these processes, generate a software simulation, run the software model, and compare it with what S actually did. In many tracking simulations, the simulation model treats these internal processes except the Output Function as simple pass-through operations, even though it is obvious that something like a Perceptual Function is likely to be extremely complicated, and nowhere in the loop in a real live organism are the f_{XY} processes likely to be trivial. The reason the simple model works is that when control is good, the Behavioural Illusion leaves little scope for distinguishing among different complexities and non-linearities that may actually be there. Because of the BI, the zero'th approximation simple

1. Remember that in Figure 1.2 the CEV is shown as being influenced by many separate paths and as influencing several different perceptions. In trying to influence the CEV, the experimenter is likely to manipulate some or all of these paths, but unlikely to do so in exactly the way they combine to form the CEV. Similarly with the experimenter's view of what the subject senses.

model works so well that it is very difficult to discriminate more precisely the processes that may be operating in the path from CEV to Output.

The "pretty good success" of such simple models is a testament to the power of the Behavioural Illusion when subjects control well. The models cannot have enough resolution to distinguish among different functional forms such as non-linearities (most researchers in perception think, for example, that perceptual magnitudes tend to be something like logarithmically or power-law related to the corresponding physical magnitudes — about which, see below). The simple models can't distinguish such varieties in the internal processing because the fact of control simply inverts the effects of any nonlinearity, as Powers (1978, Expt 3) demonstrated.

If the models simulate the behaviour accurately, does this mean that they correctly represent the underlying processes? Not really, if both the model and the human control well. Any set of internal processes that also control well would agree closely with the human results. More information must be extracted from the experiment. That is best done under conditions where the human controls relatively poorly, or in which feedback of some component of a controlled perception is impossible, as it is in the case of controlling for not getting wet by taking or not taking an umbrella depending on whether the uncontrolled perception of sky colour is grey or blue. In this case, by observing the correlation of "umbrella-ness" with sky colour, the experimenter could determine that the person was able to perceive the sky colour, or something closely associated with it.

In an experiment such as making an on-screen cursor track a moving target in a run of 60 seconds duration, more can be said about the quality of the model than a global estimate of some closeness to the human track such as its RMS error or the correlation of the model track with the human track. The better the model, the more its behaviour should exhibit the same deviation from that of the perfect control system as the human does. Since the only data available to the experimenter is the sequence of sampled values for the track produced by the target, the human, and the model, those sampled values contain all the information about the internal processes of the human that the experiment could ever provide. The question is how to extract it. Other kinds of experiment have parallel difficulties.

In a tracking study there are three tracks, the target, the human, and the model. If after optimizing its parameters the model track is closer to the human track than either is to the target, that is a sign that maybe the model has something going for it. But if the model is further from the human than either is to the target, perhaps the model needs some structural changes.

This criterion seems very simple, and perhaps it is too simple. Nor does it suggest what structural changes might improve the situation. The question is what is being modelled. The simple control model often used does not model neural noise, mouse friction, and such issues. Not even a complicated model could do that, because such problems are independent of the target track, will not be repeated exactly if the same conditions were tested again, and are not what the human was trying to control. The model is concerned only with what the human was trying to control and how the human was doing it.

These independent sources of variation cause the human track to differ from both the model and the target for reasons irrelevant to the accuracy of the model structure. Even if the model with optimized parameter values was exactly representative of a part of what is inside the human, the added noise, if control is good, might easily swamp the small difference between either and the target track. The Behavioural Illusion once again becomes a barrier to sensitive modelling, but less so if control is only moderately good.

However, there may be localized consistent effects in the human track that are not reproduced by the model, such as, perhaps, overshoots near the places where the target direction changes. Such differences, if they are consistent enough to be distinguished from random noise effects, suggest that something is missing or wrong in the model, no matter how closely the model track otherwise matches the human track — perhaps even more particularly if the model track agrees well in other respects with the human track.

The extreme of model fitting comes when the BI is eliminated by eliminating feedback for some perception that affects how another perception is controlled. Returning to the umbrella example, Ethel the experimenter thinks Sean's perception of the likelihood of rain hinges on more than a clear perceptual distinction between blue sky and dark cloud. Let's assume now that Ethel believes that Sean doesn't care one way of the other whether he gets wet. But she asks him please to control for not getting wet, and offers him an umbrella to use when he wants. Sean does control for pleasing Ethel, and to continue perceiving her to be pleased he begins to control for not getting wet. However, he finds the umbrella a bit of a nuisance, so he leaves it home if he perceives the clouds not to presage rain too strongly.

Ethel might be able judge how well Sean could distinguish "rainy" from "dry" cloud patters by observing how often it rained when he went out without an umbrella as compared to how often it stayed dry when an umbrella was taken.

Of course, Ethel might be quite mistaken. On some umbrella-less occasions when it rained, Sean might have been controlling for getting wet. The old PCT mantra "*You can't tell what someone is doing by looking at what they are doing*" really does matter. But although Ethel might be mistaken, that is not the usual case. If Sean always takes the umbrella when it is dark and gloomy and possibly already raining, and never does so when the sky is nice and blue, it's a pretty good bet that he usually does not control for getting wet.

Ethel cannot manipulate the clouds, so she cannot properly do the Test for the Controlled Variable to see whether Sean really is controlling for not getting wet, as she asked him to do. But she can observe that his use of the umbrella correlates with the darkness of the sky. If the sky is darker than some value she can measure, he almost always takes an umbrella, but if it is lighter than some other value, he almost always leaves home without it. Between those two boundaries, he sometimes does and sometimes doesn't, taking it more often near the darker boundary, less near the lighter boundary. If his actions have that characteristic, Ethel has no reason to believe Sean is not doing as she asked. Moreover, she has some evidence about his ability to perceive the cloud state, a perception he is unable to control. Because Sean cannot control his perception of the cloud state's "raininess", Ethel is able to discover how sensitive he is to the raininess of the actual clouds, by plotting the fraction of times it rained when he took or did not take the umbrella as a function of her measures of cloud lightness.

Ethel can perhaps improve her model of Sean's processing of "raininess" if she finds statistical relationships with other environmental variables besides the greyness of the sky she initially imagined him to use. Among cloudy skies, maybe some patterns are actually more likely to produce rain than others. If Sean includes such patterns in his perception of raininess, and Ethel can distinguish them, she can do an analysis to determine whether they make a difference beyond simple greyness in the probability he will take an umbrella.

In this quasi-experiment, the Behavioural Illusion is what allows Ethel to do her analysis. She can't determine the processing that allows Sean to control for not getting wet, but if he controls well for that — meaning he seldom takes an umbrella when it doesn't rain and almost always does when it does rain — then she can discover something about his processing of the "raininess" perception. Her analysis would not make any sense (and she would discover that it made no sense) unless Sean controlled well for not getting wet. Because he has no feedback that would allow him to control the cloud state perception, she can investigate that perceptual process. She could ask in her analysis "Does the patchiness of the clouds make a difference? Does it matter how windy it is? If she obscures his vision of trees that would wave in the wind, does this influence his accuracy?" and various questions along those lines, but only if he controls well enough that a naïve observer would be fooled by the Behavioural Illusion.

With respect to Sean's ability to discriminate "rainy" clouds from "dry" sky, Ethel is in the position of many psychophysicists who use different non-control techniques to study the abilities of people to perceive different things. Just as Ethel asks Sean to take an umbrella if he perceives rain to be probable, they ask the subject to do one thing if they perceive X and another if they perceive Y, but offer the subject no opportunity to influence what is actually is presented. A sound beep may occur in the first or second of two intervals, a luminous patch may be shown with more intensity than its surroundings or less, and so forth. Once shown, it is gone, and the subject can do nothing about it. But the subject does act, because something else is being controlled.

The subject, like Sean, controls for doing what the experimenter asks. The experimenter can tell whether the subject is controlling that perception well by looking at the consistency of the results, though she can't tell whether inconsistency is due to the subject having misunderstood the instruction or to poor control. Maybe the subject is controlling for something similar to what she intended him to control, but different enough to affect her analyses. Maybe there is a "Clever Hans" effect (Pfungst, 1911), and he is perceiving something quite different from what she wanted him to perceive, as Hans, the horse, perceived his trainer's unconscious body language when he seemed to be doing arithmetic. In the days of mechanical relays, subjects sometimes were able to hear differences in the sounds of the relays that corresponded to the "stimuli" they were asked to discriminate. Many such effects were discovered because the data were inconsistent, but nobody can know whether they occurred undiscovered in published papers.

An interesting example of the Behavioural Illusion was described by Marken (2014) following discussions with Powers. A famous psychophysical "law" is Stevens Power Law (e.g. Stevens 1957), which says that the perceived magnitude of a "stimulus" such as a light intensity or the loudness of a sound is a power function of the physical magnitude of the stimulus. Stevens (1966) found also that when another sensory dimension rather than number was used as a match, the result was again a power law with the exponent that would be predicted from the exponents of the individual dimensions when compared to number. Marken notes that the subject is likely to be controlling for the perceived magnitude of the number reported to be equal to the perceived magnitude of the "stimulus". The same presumably would be true when a sensory magnitude is compared to another sensory magnitude. Indeed, it is hard to see any other way the subject could do the task, when you look through PCT glasses.

The issue, then, is how the magnitude of a number or another sensory dimension is perceived. If both the number and the "stimulus" are perceived as logarithmically related to their physical magnitude (as Fechner's Law would suggest), then Marken shows that the output value (number or physical magnitude) will necessarily be a power function of the "stimulus" physical magnitude. In this case, as with Sean's umbrella, the BI allows the experimenter to probe the perception of the uncontrolled variable (the "stimulus") because a higher-level variable that incorporates it (the difference between the number and the "stimulus" magnitudes) is being controlled. The same experiment could be done by specifying the number and asking the subject to adjust the stimulus[2].

2. One should be aware, however, that Garner (1958) long ago demonstrated the unreliability of experimental methods used to generate the power law, at least in the case of perceived loudness. His criticisms do not affect perceived equality judgments, although Garner had previously shown that when people were asked to judge whether sounds were louder or softer than a standard sound, their "equality to half" judgement was very close to the mid-point of the loudness of the

All actual control exists somewhere between the extremes of the unachievable perfect control and no control. Control can be very poor, as in the case of an average citizen trying to control the perceived policies of the government, or very good, as in a tracking study with a slowly varying disturbance. The better the control, the stronger the Behavioural Illusion and the less possible it is to use the relationship between the control actions and the disturbance to say anything about the interior processes of the controller. But on the other hand, the better the control of a higher-level perception that incorporates an uncontrolled perception, the more one can discover about the uncontrolled perception because of the Behavioural Illusion at the higher level.

set of sounds offered for comparison (Garner, 1954). Garner's assessment in 1958 included: "*...it is clear that we are on very dangerous ground in assuming that the loudness scale proposed by Stevens has any real meaning in the experience of normal observers*"