weigh. If they are too large for his consumption, or if they are too light, *i.e.*, empty, he kicks them over with his leg. Furthermore, he has four lamps arranged in a square which he can turn on and off, one or more at a time. With these he can communicate his needs. Clearly, with any of his four lamps in two possible states—on or off—he has precisely $2^4=16$ different "words" in his vocabulary. In the present state of affairs, however, these words have as yet no "meaning" whatsoever. This is precisely what we are going to teach him. To this end we have to invent some environmental rules. Let us agree on the following convention (Figure 2): the coordinates of

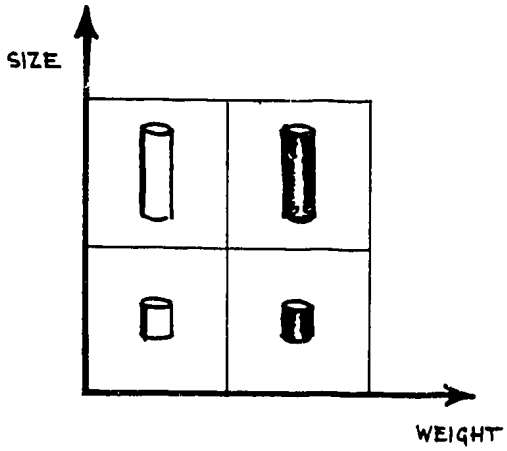

FIGURE 2

his square lamp-box are to represent upwards=size, and horizontal=weight. In other words, if he lights up, say, the upper right lamp, we feed him a tall, filled can, which he will reject, of course. Such a trial will be 100% unsuccessful. If he lights up two lamps simultaneously, say the lower left and right lamps, we will feed him with equal probability either a small empty can or a small filled can. This would make him successful 50% of the time. Clearly, what he has to learn is to light up the lower right lamp only, which causes his environment to feed him the desired small, filled can.

Since our automaton is constructed out of elementary components which follow the points 1 to 4 mentioned earlier, he will, after a series of more or less successful trials, resume the habit of turning on only the lower right lamp when hungry. He knows now the "meaning" of this word; it means "digestable food." At this stage it may become monotonous to feed this critter whenever he turns on his lamp. We do not gain anything by watching this device which has turned into a boring, deterministic system. However, at that instant the automaton presents us with his waste products which happens to be a chocolate bar (Figure 3). This, of course, changes the

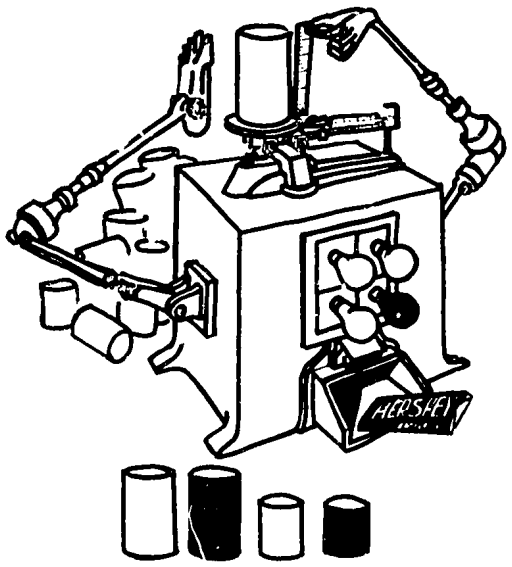

FIGURE 3

situation for us, because it offers certain advantages for us if we are going to form a coalition with this system. Again this may go on for a while—but who wants to live on chocolate forever? Thus we may replace ourselves by an automaton that can't do anything with gasoline, but who thrives on chocolate bars (Figure 4). Furthermore, this new fellow has eyes to recognize light-signals and to distinguish gasoline cans. Wheels permit him to move freely in his environment which consists again of large and small, empty and filled, gasoline cans. He, too, will learn to understand the language of our primary automaton, and he will too, after some time of adaptation, appreciate the advantages of a coalition. Before long you will see these two automata joined together (Figure 5), the one acting as stomach, cracking up the raw-products and transforming them into