

PART IV
INNER INTENTIONAL SIGNS
CHAPTER THIRTEEN
INNER PUSHMI-PULLYUS

The remainder of this book is about inner intentional signs or "inner representations." The central question that I want to address is how and why, during the evolution of perception and cognition, organisms have acquired inner representations that are more sophisticated than pushmi-pullyu signs. How and why did perception-action cycles, which seem fully to characterize the cognitive character of the simplest animals, slowly give way to or become supplemented with more articulate and differentiated representations such as human beliefs, which are merely descriptive, and human desires, which are merely directive? Only quite speculative answers are possible here, but evidence from contemporary neurology and experimental psychology allow us to discern a few of the turning points pretty clearly, and may help us to tell a good story about some of the rest.

Much the most basic kind of intentional signs are the ones I called "pushmi-pullyu" signs in Chapter Six. I will often call these "P-P signs" or just "P-Ps." P-Ps are signs that are undifferentiated between presenting facts and directing activities appropriate to those facts. They represent facts and give directions or represent goals, both at once. As mentioned in Chapter Eight, we find some P-Ps even in public languages: "No Johnny, we don't eat peas with our fingers!" So far as I know, all intentional signals used between non-human animals are P-Ps. On the one hand rabbit danger thumps mean rabbit danger, but on the other they direct nearby relatives to take cover. Bee dances tell where the nectar is and at the same time tell where the watching bees are to go. The famous leopard, snake, and flying predator calls of the vervet monkeys tell both what kind of predator has been spotted and, simultaneously, direct behaviors appropriate to avoiding that kind of predator. Human smiles and frowns that are not yet overlaid with conscious intentions are simple P-Ps, telling that something potentially rewarding has just been done and to keep doing it or do it again, or telling that something potentially damaging has just been done and to stop doing it or not to do it again. Similarly, the snarls of animals, or their tail lashings or waggings, their mating displays, their signaled invitations to play (the dog's "play bow"), and so forth, are all P-Ps.

P-Ps are also much the most common intentional signs occurring inside organisms. The bottom-most level of inner P-P signs is ubiquitously exemplified, not merely in neural matter, but in the many chemical messengers found in the body tissues and circulatory systems of animals. These are signals that are secreted, sent out, by one part of the animal's body to other parts, usually telling of the condition of the one part, and telling other parts how to respond. These P-Ps are the basic regulators of bodily homeostasis, coordinating the contributions of cells and organs so as to effect routine maintenance of bodily integrity. To suggest that genuine intentionality, genuine aboutness, with the possibility of misrepresentation, actually occurs at this level may at first seem far fetched. But the idea is that there is intentionality here in the sort of way that zero is a number. These are the most humble sorts of limiting cases of intentionality. But by treating such simple signals as intentional signs, just as by treating zero as a number, we will be able to examine their relations to various successors, and see the continuity between them and their more sophisticated relatives.

Simple reflexes, such as the reflex that withdraws the hand from something unexpectedly hot, are mediated by P-Ps. The neural signal that reaches the spinal cord tells what part of the body is exposed to something too hot and directs withdrawal of that part. The very numerous neural mechanisms that work by negative feedback, such as the internal mechanisms that control tropistic behaviors in primitive animals, and portions of the mechanisms that control walking behavior in insects and, indeed, also in mammals, employ P-Ps. Negative feedback is a representation of the discrepancy between the value of a perceived variable and a set target value for that variable, telling what the discrepancy is and directly controlling the strength of the response needed to correct it. The primitive ability to follow a temperature gradient or a light gradient, for example, is mediated through P-Ps, as is the ability to follow a moving target with the eyes.

Moving higher up in the nervous system, the instinctive fear of snakes and heights that is built into many mammals including some humans (all babies instinctively shrink from precipices) are inner P-Ps which, when working as designed, are perceptions on the one hand and directives on the other. Inner states that serve as reinforcers, either positive or negative, are P-Ps. Sweet tastes tell of nutritive value on one hand and direct continued eating or seeking more on the other, although, of course, there may be other systems capable of producing states whose function is to override these P-Ps. Pain tells of damage to body tissues and directs present and future pain-avoidance, hence damage-avoidance, behaviors. That is how sweet tastes and pain work when they work as designed by natural selection. Similarly, perceptions of smiles and frowns, certainly in infants, are P-Ps.

According to J. J. Gibson and contemporary ecological psychologists (Gibson 1969, 1977, 1979; Michaels & Carello 1981; Reed 1982, 1993, 1996) basic perception consists in "picking up" or extracting certain abstract patterns in the ambient energies arriving at the organism's sensory surfaces, which patterns then guide various activities of the organism directly. No inference or calculation is required, but merely sensitivity to certain variants and invariants in the energies impinging on the active organism that, on the one hand, carry information about the relations of significant distal affairs to it and, on the other, directly guide its motions to take account or make use of these distal affairs. Basic perception is thus interpreted as perception of what Gibson called "affordances." Affordances are aspects of the environment that afford the possibility of various activities for the animal, such as walking on, climbing up on, going through or into, chasing or fleeing from (prey or predators) ducking away from (approaching objects), throwing, and so forth.

Gibsonians have generally assumed that if there were such things as inner representations they would have to be things calculated over, vehicles of inference, hence that the perception of affordances does not involve inner representations. But inner processes mediating the perception of and responses to Gibsonian affordances would certainly involve P-P representations, these being far more primitive than the representations Gibsonians reject. Information "picked up" by the organism would have to involve alterations to its inner states which would be intentional signs of environmental affairs, these signs in turn guiding the organism's responses. Further, Gibson's claim that basic perception is perception of affordances is separable from his claim that perception is "direct." Basic perception might involve inference and still be perception of affordances. But in any event, Chapter Nine argued that sensible use of the notion "direct

perception" will include perceptions derived by translation from inner representations of prior external aspects, such as edges, corners, oriented surfaces and so forth. These translations should not be assimilated to inferences. Any response directly guided by a perception of relative size, distance, shape for picking up, angle of incline for climbing, momentum for throwing, and so forth, where dimensions of what is perceived directly guide dimensions of the response, would seem to involve inner P-P representations, whether or not the perceptual representation was derived, as Gibson proposed, as a direct function of invariances in ambient energy inputs. Any such perception could surely be considered perception of an affordance. It should be sufficient that there was a direct mapping between perceived variations in the environment and directly guided variations in behavioral response. At least that is the way I propose to use the phrase "perception of an affordance."

All in all we can conclude, I believe, that P-Ps occur on many different levels within organisms and that they vary greatly in sophistication.

P-Ps may be more articulated or less articulated. Some are merely intentional signals, their only significant variables being time or place. But even the pull of the magnetosome in the bacterium (Chapters Three and Six) is a bit more articulate than that, for it indicates also a direction, which can vary over a solid angle, directing the movement of the bacterium to correspond. The bee dance has variables indicating not only time and place but also direction and distance of nectar, hence direction and distance of the place the bees are to go. Perhaps the neural representation produced in the bee watching the dance is similarly articulated. The image projected onto the retina of the male hoverfly by its eye lens which causes it to dart towards and intercept an approaching female is a P-P (Millikan 1990). It guides the direction of the male's flight according to the direction of the female's flight, determined not as a linear function but as a certain trigonometric function of the retinal pattern she causes (to be exact, 180 degrees away from the target minus $1/10$ the vector angular velocity measured in degrees per second of the target's image across the male's retina (Collett and Land 1878)). There can be no variable articulation of the pushmi face of a P-P sign (the descriptive side), of course, unless this articulation directs coordinate variation on the pullyu face (the directive side). P-P representations are intentional signs; their significant variables are the ones used to guide their interpreting mechanisms.

P-P representations can be very abstract. In contrast, the empiricist tradition in philosophy has quite consistently maintained that what is originally presented to the senses or in perception is concrete. A classical problem for the empiricists was then to explain how mental abstraction was accomplished, so as to separate out representations of kinds and properties from representations of their concrete instances. Given the description of intentional signs used here, this classic problem does not arise. An intentional representation is produced by a system designed to cooperate with an interpreting system in turn designed to use that representation in specific ways. The intentional content of the intentional sign is restricted to what the interpreter can read, that is, can make use of. No matter how rich and nested the natural information carried by an intentional representation is, only the part designed for use by the interpreting or consuming part of the cooperative system is represented intentionally (Chapter Six). Let me quote here a philosopher who, though he does not espouse any form of teleosemantics, makes this particular point very clearly (with a little help from John Locke).

The doctrine that picturelike representations won't do for general or adult or primate concepts involves a conceptual error...Obviously you can't tell how a certain representation functions by confining your attention to the representation alone, or its "resemblances" to things in the world. You must know how the processors that act on it treat it. Thus a pictorial representation can express quite an abstract property, so long as the processors that act on it ignore the right specifications. To take a venerable example, a picture of an equilateral triangle can serve to represent triangles in general so long as the processors that act on it ignore the equality of the sides and angles. Similarly, a picture of a set of twins could represent or express the concept of a pair whose members are identical. (Ned Block 1986 [Stich & Warfield, p. 128])

An interesting result is that even very primitive P-Ps can represent very abstractly. For example, the statolith in the statoreceptor of the fish represents just one very abstract relation, namely, which way is down, hence which way to move to remain right side up. On the other hand, animals that recognize their individual conspecifics by smell also employ simple perceptual means, but in this case the intentional information represented is entirely concrete, for example, it might say "here's Mama." In general, there is no correlation between proximity to sensory input and the abstractness or concreteness of an inner representation.

P-P representations can represent either proximal or distal affairs. For example, even extremely primitive P-Ps, such as the magnetosome which represents the direction of lesser oxygen, can represent quite distal affairs. Perceptual P-Ps represent things whose placement relative to the animal matters to the animal, things that it needs to take account of directly in action. What the pushmi or descriptive face of such a P-P represents is whatever environmental conditions it needs to vary with in order to guide its consumers properly.¹ Sometimes these conditions are absolutely proximal. Whether the skin is rapidly increasing in temperature may make quite a lot of difference to an animal, which is why our heat and cold receptors are designed to perceive just this very proximal affair and not, for example, the objective temperatures of the objects that touch the skin (Akins 1996). Pain and bad tastes also represent affairs absolutely proximal to the animal. But if the energies impinging on the organism, say, the kind and pattern of light or sound impinging on it, makes no particular difference to its well-being when within normal intensity ranges, these proximal patterns of light or sound will not themselves be perceived. Rather, any P-P signs derived from these energy patterns will concern more distal matters about which these patterns carry natural information.

Nor is there, in general, a definite distance at which a given sensory modality, such as sight or hearing, is designed to perceive. Contrary to much of the philosophical tradition, there is no single level of the outer world, such as physical objects versus the mere surfaces of physical objects, or such as the presence of certain physical objects or of events versus mere sounds, of which the eyes or the ears are designed exclusively to produce direct representations. Depending on the animal's needs, various levels of distality of direct perception may be mediated by the same sensory end organs. The affairs naturally signified by retinal patterns, vibrating eardrums, stimulated odor sensors, and so forth, are at various distances and mediated in diverse ways.

Whatever affair a sight or a sound or a scent is a natural sign of, that sight, sound or scent can, in principle, be used to produce an intentional sign of that affair directly, without any intervening intentional signs (Chapters Five, Six and Nine). Simple examples of this may be the

intentional neural signs that mediate between many environmental signs that are "behavior releasers" and the "fixed action patterns" thereby released in many animals (Lorenz and Tinbergen 1939; Tinbergen 1951; McFarland 1981, p. 1990 ff, Gould 1982). The proximal stimulus that guides the mother bird to drop food into the baby bird's open mouth immediately creates a neural pushmi that represents not a patch of red but a hungry baby's mouth, for only if the patch of red is indeed a local sign of a hungry baby's mouth will the behavior thereby released serve its purpose. The pushmi face of the neuronal P-P caused by a certain kind of dark shadow crossing the retina of a male hoverfly that causes him to fly off in a certain direction represents a female hoverfly, not a black moving thing or pattern of light or a moving image on the retina. In each of these cases it is likely that the representation is formed quite directly from retinal stimulations without passing through intermediate stages of representation. On the other hand, a sight or a sound or a scent may produce a direct perception of a distal object by passing through intermediate stages involving translation from prior representations ("direct perception" in the sense defined in Chapter Nine). In this case it is even clearer that there can be direct perception mediated by the same sensory end organ at various levels of distality, hence that pushmi-pullyu representations can represent affordances at various levels as well.

P-Ps can represent affairs that are distal in time as well as distal in space. Probably few if any animals besides humans have developed uses for representations of past affairs, but many have need to represent future affairs. In Chapters Three and Four, local signs of such things as coming rain or winter approaching were discussed. Quite simple animals may need to translate natural signs of such future events into inner P-Ps that stimulate preparatory behaviors. Similarly, no matter how simple, most animals have an obvious need to recognize signs of approaching predators and to translate these into appropriate behaviors. There is nothing the least bit exotic about the production of inner representations of affairs distal in time any more than representations of affairs distal in space. Seeing into the future is exactly like seeing into the distance. The animal for whom a frosty night or the low angle of the sun serves to release winter preparation behaviors is being governed by inner P-Ps whose pushmi faces say that winter is on the way and whose pullyu faces direct what to do about it. For only if winter really is on the way, will the behaviors that result serve the functions for which they were selected, and only if the P-P succeeds in producing these behaviors will it fulfill the function for which it was selected through its normal mechanisms.

Inner P-P representations may or may not require to be joined together with other P-P representations in order to do their work. The P-P neural impulse produced in the frog's optic nerve by a passing fly reports when and at what angle the fly passes and provokes a correspondingly response from the frog's tongue. This impulse forms part of a simple reflex arc which cannot be inhibited, even if the frog is completely sated. It reports a fact and issues an unconditional command. Similarly, during the first few days of its life, a rat pup whose snout comes in contact with a saliva coated nipple grasps the nipple and continues to suck whether or not it is hungry. A few days later, however, this response is inhibited unless the pup is hungry (Hall, Cramer and Blass 1975, 1977). Thinking of this in intentional terms, the pup's system is now sensitive to a new P-P signal indicating a current state of nutritional depletion and directing a response. The response is potentiation of the grasping and sucking reflex. The hunger signal says (roughly: see Chapter Seven), "Nutrients are depleted; if there is a nipple handy, suck on it!"

Similarly, many small animals instinctively take cover if they see a small shadow gliding over the ground, such as would be cast by a flying predator. The shadow produces a P-P that means "predator overhead; if a cover-taking affordance appears, exploit it."

There seems no reason to suppose that affordances irrelevant to current needs are always, or even ever, perceived by most animals. There are, after all, lots of things we humans are capable of perceiving but generally don't perceive unless currently interested. Why would an animal feverishly translate every readable natural sign it encounters immediately into perceptions? Much of the currently perceptible world stays right where it is to explore perceptually later should it then become relevant. There is no need to think of simple animals as perceiving everywhere about them mere possibilities for action. More likely they only perceive what they have motivation, at the moment, to exploit. Indeed, it may turn out that this point is crucial to understanding how represented purposes can be detached from current perception of the means to fulfill them. (This idea will be explored further in Chapter Seventeen.)

To be in a position such that a primary goal, such as having a fly in the stomach, can be achieved by utilizing just one perceived affordance, such as a fly currently passing within reach of the tongue, is a blissful condition. Call such an affordance a "B-affordance." Call a negative condition that threatens immediate disaster an "D-condition." D-conditions are perceived by primitive animals as negative affordances, directing immediate avoidance or escape techniques. Call these "ND-affordances." (Similarly, so called "negative reinforcement" is reinforcement of behavior that has afforded avoidance of or escape from negative consequences.) Organisms such as venus fly traps and sea anemones that do not move about may merely wait for B-affordances to pass by and then seize the moment. Similarly for ND-affordances. More sophisticated organisms make an effort to maneuver themselves into B-affordance conditions. The simplest way may be just to wander about directionlessly hoping to bump into one. This seems to be what clams do, for example. Other animals use more systematic techniques. The newborn baby's response to a touch on the cheek is to turn toward it, thus raising the probability of feeling a nipple on the mouth which will afford immediate nourishment. Very simple animals show various kinds of taxis likely to take them into conditions where B-affordances are more likely to be encountered and D-conditions less likely. The frog recognizes places to approach and sit that are likely to attract flies. Perhaps it also recognizes places to avoid that are likely to attract snakes. One way to view the story of the evolution of perception and cognition is as a story about the acquisition of more and more sophisticated search techniques for maneuvering oneself into B-affordance conditions while staying out of D-conditions. Call an animal all of whose search techniques exploit only chains of perceived affordances so that its behaviors are entirely governed by inner P-P representations a "pushmi-pullyu animal." The central question that I wish to address in these last chapters concerns what the disadvantages of being a purely pushmi-pullyu animal might be and what remedies for these disadvantages may have been supplied during evolution of the higher species.

The cardinal principle involved for any pushmi-pullyu animal in raising the probability of encountering B-affordances is very elementary. Be constructed such that you can perceive affordances that will afford your probable placement in new positions from which you are likely to perceive new affordances that will afford your probable placement in newer positions from which ...and so forth...finally placing you in B-affordance conditions. The trick is that this series

of probabilities should have a product greater than the probability of B-affordances just happening along without any action on your part, the higher the probability the better. Thus the search domain is narrowed and then narrowed again.

Ecological psychologists speak of "perception-action cycles" during which input from the environment is said directly to guide motor output without additional input from the central nervous system. Perception produces action that results in new perception producing further action, and so forth. Robots that work entirely on the perception-action principle have been constructed that perform simple tasks, for example, following walls, avoiding obstacles, picking up soda cans, recharging their batteries at the right times and so forth (Rodney Brooks 1999). The various cycles that govern the behaviors of these robots are arranged so that the activation of some cycles will inhibit the activation of others, an order of importance in activities thus being established.² The activities of insects may be largely or entirely governed in this way, by hierarchies of perception-action chains, or as ethologists call them, chains of "behavior releasers." Thus the digger wasp walks randomly this way and that until it encounters certain signs of a prey; the prey affords stinging hence paralyzing, which affords being dragged to the entrance of the wasp's nest, which affords entering and circling and, if all is well, affords emerging again and dragging the prey within, and so forth. The cycle will be interrupted, of course, if at any point the wasp encounters ND-affordances, signs of danger.

In more complex animals, motivations such as hunger or fear may potentiate perception not just of one but any of numerous alternative affordances, calling on a large set of alternative behaviors to be used contingent on the animal's situation. Extremely complicated long and branching chains of affordances leading to the probability of finding one or another other affordances, leading to the probability of finding one or another...and so forth, may be grasped by some animals, resulting in highly flexible behaviors. And it may be that correctly quantified increases and decreases in potentiations of response dispositions that change the ease with which these can be activated, produced by other relevant stimuli encountered along the way, help account for the tendency of the animal to perceive, from among equally available and relevant affordances, those objectively associated, in the animal's particular circumstances, with higher probabilities of eventual success. The result would be an animal whose behavior was highly flexibly governed by what Gallistel (1980) calls a "lattice-hierarchy."

Gallistel points out that an animal whose behavior is completely governed by a lattice-hierarchy may have acquired that lattice in large part by learning. Suppose that instrumental (operant) conditioning works in the way that classical American behaviorists claimed. The analogy with natural selection is quite strict. Responses to stimuli are strengthened whenever they are followed by "positive reinforcement" or reward and weakened when followed by punishment. Primary reinforcements, those that are not learned, are associated with B-affordances or ND-affordances. Reinforcing a behavioral response to a stimulation thus conditions the animal to perceive an intermediate affordance leading either towards a B-affordance or away from D-conditions. Perceptions of intermediate affordances then become "secondary reinforcers" --say, the introduction into the cage of a bar-for-pressing when a trained rat is hungry, or the lighting of a light that signifies that bar-pressing at this time will now afford food. Thus the animal can be trained to perceive affordances that afforded further affordances in quite a long chain.³ But this instrumental conditioning merely designs another lattice-hierarchy, a

lattice-hierarchy, constructed during ontogeny rather than phylogeny. The result might still be merely a pushmi-pullyu animal.

It is clear then that a pushmi-pullyu animal might be capable of navigating in the space-time-causal order from a great variety of different starting positions relative to its goals so as to reach them with high probability. On the other hand, such an animal might also be subject to failures that strike us as rather ridiculous. Gallistel notes that even in quite flexible animals, available behaviors are by no means always chained so as to apply to relevant situations, even when the increment is very small. As an example, he cites Dilger's work (1960, 1962) describing a hybrid species of lovebirds that, although quite capable of safely carrying the strips of bark used for weaving their nests in their beaks, nearly always carried these strips by tucking them into their tail feathers, losing most of them on the flight back to the nest. The tucking behavior was a leftover from ancestor species that lined their nests with small chips, which more easily stay put in the tail feathers. Although these lovebirds were capable of perceiving and acting on each relevant affordance, these capacities were not chained in an efficient way (Gallistel 1980, 306-8). A more homely example of this is the house cat that washes its ears by rubbing them with its paws and licking them off, and that may pull food scraps out of its dish with its paws to eat on the floor, but doesn't know to use its paw to clean out the yummys at the very bottom of an emptied can of cream of chicken soup. (An occasional cat does figure this out.)⁴ Similarly, a purely pushmi-pullyu animal would lack the ability to recombine various segments of behaviors in its repertoire in new ways so as to achieve new goals. It could achieve new linkages of behavior chains only by reinforcement of accidental connections after the fact, never by inventively looking ahead.

One step toward inventive recombination of behaviors, I will argue, is the articulation of P-P representations that govern behaviors into segments that can be recombined to make novel P-P representations. I will discuss this kind of segmentation in Chapter Fourteen. Also, with the development of more complex articulation in P-P representations, some representations of pure facts become detached, producing "pushmis" that are ready for recombination with various alternative "pullyus." And some representations of facts are not only detached, but stored away for use on other occasions (Chapter Fifteen). The pure pushmi-pullyu animal, on the other hand, represents only facts that it already knows how to use, and represents them only in the context of their use. This means that it never has any extra information lying around, as it were, to employ in situations it has not already been genetically programmed, or trained by conditioning, to deal with. All of its facts are devoted to specific uses. Further, the pure pushmi-pullyu animal doesn't represent, hence doesn't know about, anything it is not currently perceiving. It may have a memory for procedures but not for facts.

A correlate is that the purely pushmi-pullyu animal always represents affairs in its world as bearing certain relations to itself. An animal's action has, of course, to be initiated from the animal's own present location. To serve as an unmediated guide to immediate action, the descriptive face of an inner P-P representation has to represent the relation of the affording situation or object to the perceiving animal. This doesn't require that the animal represent itself explicitly, any more than the bee dance represents nectar, hive and sun explicitly (Chapter Seven), but the self has to be represented at least implicitly. In the simplest cases, the relevant relation may consist merely in the affording situation occurring in roughly the same location and

at the same time as the animal's perception and consequent action. In less simple cases it will include more specific relations to affording objects, such as spatial relation to the animal, size relative to the animal's size, weight relative to the animal's weight or strength, and so forth. We humans, at the opposite extreme, are capable of forming beliefs not only about things and affairs very distant from us, but about things whose spatial and temporal relations to ourselves are completely unknown. Certainly, our descriptive representations do not, in general, represent relations of situations and objects merely to our current selves.

More striking than the failure to recombine behaviors in relevant ways is another sort of failure which was described by Lorenz and Tinbergen (1938). Although the greylag goose apparently reacts intelligently to an egg that has rolled out of the nest by bringing the bill behind the egg and rolling it back into the nest, if the egg slips sideways out of control of the bill, this movement may still be carefully completed, "as if it were a vacuum activity" (Tinbergen 1951). It seems that the goose does not understand the purpose of its own behavior. Dennett (1984) has popularized the example of sphex, the digger wasp, that can be sent into a behavioral loop from which it never emerges by removing its paralyzed prey a few inches away from the door of its nest every time it goes inside to inspect, preparatory to dragging the prey inside. The wasp seems not to understand the purpose of its own activity so as to know when that purpose has been accomplished. I once watched a pair of hamsters repeatedly stumbling over one another as each returned a large cracker to its own corner again and again from the other one's corner just opposite. Neither seemed to notice that its own corner continued to remain empty.

This sort of failure is exactly what we should expect of a purely pushmi-pullyu animal, for such an animal does not represent its goals in a format that enables it to know whether or when it has reached them. The language in which the directive side of a P-P representation is expressed is not the same as the language in which the descriptive side is expressed, so that when acting on the directive produces or fails to produce a corresponding state of affairs, this is not automatically recognizable to the animal. The purely pushmi-pullyu animal does not, as it were, project its goals. Its behaviors are controlled completely from behind by emerging environmental contingencies. It does not represent its goals as purposed future occurrences or states to which actual accomplishments will be compared. Chapter Sixteen will concern the transition from purely pushmi-pullyu animals to animals that represent their goals in the same representational system in which they represent their facts.

FOOTNOTES

1. Granted, that is, that this correspondence is one the representation producer has a means of bringing about through its normal mechanisms of operation (Chapter Five, Chapter Six on fertile female hoverflies).

2. Because there is nothing in these robots that corresponds to a central processor or central nervous system orchestrating the whole by performing inferences, it is commonly assumed that these robots do not employ inner representations at all. But they do employ inner P-Ps in the sense defined here. These robots are not, of course, direct products of natural selection. But they are designed to function as they do by people who are themselves designed by natural selection and by learning processes that natural selection has designed capacities for. The robots have proper functions in the way all designed human artifacts do (Chapter One; Millikan 1984 Chapter Two. For more details, see Millikan 2002).

3. According to Anthony Dickenson, "perhaps the most bizarre demonstration of the power of positive reinforcement comes from a legendary laboratory rat who was prepared to climb to the top of a spiral staircase, 'bow' to the audience, push down and cross a drawbridge, climb a ladder, use a chain to pull a model railway car, pedal the car through a tunnel, climb a flight of stairs, run through a tube, and descend in a lift, all to receive a single pellet of food." (McFarland 1987, p. 472).

4. I don't mean to imply that the domestic cat is a purely pushmi-pullyu animal.