A Bet With Peacocke

Ruth Garrett Millikan
University of Connecticut

Professor Peacocke's doubts (chapter 00 above) about the view of mental content offered in "Biosemantics" (chapter 00 above) grow, I believe, from misunderstandings. So I shall try to make the relevant features of my position clearer. Especially, it may help to talk more about the nature of "normal explanations" for cases of novel biological functions, and it may help to say more about the role that "mapping" or "picturing" plays in representing.

It takes many people by surprise that items uniquely associated with individual organisms, individual items quite new under the sun, can have biological functions in a sense defined by evolutionary history. As laymen, we tend to think of that which evolution "designed" in us as being both inborn, and identical in all of us having the same responsible genes. We tend to contrast this with how the environment shapes us and, especially, with what we learn. But that view is not merely an oversimplification, it is faulty at the core. The environment is involved in the determination of every phenotypic trait right from the start, and the environment is highly variable from the start. The environment takes the form, first, of surrounding genes, then of egg or womb, and finally there is the world outside. Genes are selected only if they perform usefully, produce useful phenotypic properties, in a wide variety of environments, first genetic environments, and later external environments.

Sometimes genes are selected because they are capable of performing always the same function despite great variety in their environments. Sometimes, serendipitously, they perform a variety of different useful functions each in a different sort of environment. And sometimes, more systematically, they perform functions that vary in response to some kind of variation in the environment, such that they either utilize or compensate for that variation. In the latter case, and when the environment utilized is outer, we have adaptation of the individual animal to its individual environment. Each of various varieties of learning, such as imitation, conditioning, coming to know by perception, by inference, or by verbal communication, is an example. The expression of open instincts is a second, simpler example. The formation of callouses where the wear is is a third, still simpler example.

One of a pair of identical twins, who plays violin, has callouses on the left hand finger tips, while the other, who plays clarinet, has a callous on the side of the right thumb. These different callous patterns have resulted from exactly the same genetic materials interacting in accordance with exactly the same biological "plan" within different environments. The twins' callouses are not only in different places, they have different functions. Those on the left hand are there to prevent damage to the left finger tips, the one on the thumb is there to prevent damage to the thumb. Surely there is nothing mysterious about this, about the fact that though deriving their functions from the same selective history, these callouses have <u>different</u> biological functions. Because these different callouses have different functions, each must have a different normal explanation for proper performance of its function. The normal explanation for proper performance of the callouses on the finger tips begins with something persistently rubbing on the finger tips but, because the callouses are there, doing no tissue damage. The normal explanation for proper performance of the callous on the thumb begins with something

persistently rubbing on the thumb.

But there is also a way of viewing the matter such that these callouses all have the same function, and the same normal explanation for proper performance. This way is to view the functions of the callouses as relational. Each of the callouses has the same function as each of the others, namely, preventing tissue damage to the site just under it. And the normal explanation for performance of this function is always the same. It begins with something's persisting in rubbing where the callous is, and ends with prevention of damage to the site just under where the callous is. This principle of substituting invariant relational descriptions for variant or novel categorial descriptions of functional items is completely generalizable. Take any variant part, state, aspect or activity of an individual organism that has resulted in accordance with a genetic biological "plan" for producing such variants according to variation in the environment. Every such variant structure can be described relationally. It can be described as an invariant structure, one common to all individuals having (I am putting this too crudely) the same relevant genes. This description defines the structure as embedded in or relative to the individual animal's environment, rather than describing it in isolation or categorially. The structure is described as an invariant structure, exactly the same one as was present also in ancestors of the organism. This univocal relational structure has been systematically selected for over the generations on account of a relational function that it serves. And the way that this relational function is normally effected, the normal explanation for its performance, also has a univocal relational formulation. All novel responses to novel aspects of an individual organism's environment that are truly adaptive responses, that is, that are fitting in accordance with biological plan rather than by an accident of the moment, are responses that could also be described relationally. They could be described as invariant rather than as novel responses. The functions that these novel responses to a novel environment perform when they perform properly are also describable as invariant across the generations. And the manners in which these functions are preformed, the normal explanations for their performances, are likewise invariant when described in relational terms.

Now consider the mapping rules that, according to "Biosemantics," define the content of a mental representation. That is, consider the semantic rules that apply to it. Reference to these rules emerges when one gives an invariant, relational normal explanation for proper performance of the representation and of the consumers of the representation. This relational explanation may have quite a complex derivation. For example, if the representation embodies a particular propositional attitude of a particular human, the derivation will reach back at least to the evolutionary history of those human cognitive mechanisms that are responsible for concept formation, that is, roughly, for human capacities to learn how to learn. Involved will be normal explanations for how these prior learning mechanisms work, which will include as a part (doubly) relational explanations of how their products (the individual concepts formed) work, and ultimately of how particular propositional representations formed using these concepts work, when they work normally.

The truth condition for a mental representation is never, then, merely some novel condition that would be required in the particular case for proper performance. Indeed, the truth condition has nothing whatever to do with the particular case. The truth condition may not be met in the particular case; it may not be physically possible that it be met in the particular case; it may not even be logically possible that it be met in the particular case (compare Chapter 00

note 17). The truth condition is derived, rather, from an invariance in relational function and in the relational explanation for performance of that function applicable over a very long period of evolutionary history.

Any of a great number of combinations of ordinary or freakish circumstances might, in the particular case, suffice to combine with a mental representation to produce some kind of felicitous result. Similarly, worker bees reacting normally to a certain bee dance might succeed in finding nectar because there was nectar, not at the spot indicated, but somewhere (anywhere) on the path between the hive and that spot. Or they might succeed because there was nectar just so far west of the specified spot and a strong east wind was blowing, forcing them off course just enough to find it. Or they might succeed because there was nectar just so far east of the specified spot and a strong west wind was blowing, or because a bee trap was at the specified spot and the owner of the trap was waiting to transport the bees to a good nectar site, or for any of an uncountable number of other reasons. But of course none of these conditions, whether merely possible or actual, has any relevance to what counts as the semantic content of the bee dance. Nor would any such conditions be relevant just because there had in fact been ancestors of the bees who once got to nectar in similar accidental sorts of ways. The bee dance is a representation because there is, we take it, a univocal relational explanation that is invariant over the bulk of cases in which ancestors of our bees found nectar as a result of reacting as they do to bee dances. This explanation begins always in the same way, by pointing to the fact that the relational condition holds that there is nectar at a location projected in accordance with a certain mapping rule from the configuration of the dance, that is, that there is nectar that bears a certain set relation to the dance. This is a set or invariant normal condition for proper performance of all bee dances, or of all worker-bee reactions to bee dances, for a given species of honey bee. Bee dances that lead worker bees to find nectar despite the fact that no nectar bears this set relation to the configuration of the dance do not perform properly in accordance with a normal explanation, but produce their proper effects for merely accidental reasons unrelated to the causes of the evolution of the dances.

Before bringing these observations directly to bear on Peacocke's remarks, there is a second point on the determination of the semantic rules for mental representations that we need also to have in mind. This is a point suppressed in the essay "Biosemantics," where the main task was to distinguish the biosemantic theory from causal/informational accounts, but it lay on the surface in (Millikan 1984) Chapter 6. (It is also implicit in Peacocke's own statement of the biosemantic theory in chapter 00 above.) It concerns normal explanations for production rather than for use or "consumption" of mental representations.

An inner representations has semantic content only by reference to a normal process (or processes) of consumption and normal results or uses, more specifically, by reference to a normal explanation or mechanism for successful use. Equally, it has content only as having been produced by a mechanism that was designed to produce it. That is, the producer was designed to produce representations that map onto the world in accordance with the projection rules to which the normal consumer or normal process of consumption is adapted. The consuming and producing mechanisms must be biologically cooperating mechanisms. Roughly speaking, evolution has designed each such that proper performance of the other is a normal condition for performance of its own functions. The job of the producing mechanism or process is to produce representations that map onto environmental circumstances in accordance with a

certain fixed set of semantic rules. That the representations do indeed map onto circumstances in accordance with these rules is a normal condition for the consuming mechanism's successful use(s) of them.

Now the proper functions of an item are defined as results or activities that a certain (typically very large) set of its ancestors once actually effected. So it could not possibly be a proper function of the representation producers to produce maps that map onto conditions by certain rules unless they have had an effective way of doing this. They must have had some way of systematically producing the required relation between representation and represented. There must be a systematic way in which the makers produce these maps, making them vary as the mapped conditions vary. This does not mean that the producers are always, even usually, able to effect such mappings. But, at a minimum, there must have existed certain kinds of conditions under which they effected such mappings in the past, and thus enabled the representation consumers to perform in a manner that lead to the proliferation of the responsible genes. Nothing is represented by an inner representation without there being a systematic historically realized mechanism or system for getting representations of that type to map onto their representeds by the relevant rule, a system that has worked under at least some conditions. Further, these mechanisms must have produced past effective representations, ones that effectively guided the representation consumers in such a way as to help effect natural selection of the production-consumption system.

One final clarification needs to be made, perhaps, before looking at Peacocke's claims about "reduced content." The fact that a representation maps as it should onto environmental circumstances is, in general, only one of many normal conditions required for proper performance of its various functions. For the bee dance to make its proper contribution to the proliferation of the bees, the weather conditions must be moderate, the air must be pure enough, the sun must not be eclipsed, there must be no bee traps or bee eaters directly between the hive and the indicated nectar spot, a large variety of body parts of the worker bees must be normally constituted, and so forth. But the fact that many bee dances may not be able to perform their proper functions due to absence of various of these other conditions is not relevant to the question concerning the semantic content of the dances. This is because there is no produced and exploited mapping relation between the dances and these conditions. A representation represents only as a part of a system of representation. Where there is no systematic effective and designed way of varying the representation with the conditions to be represented, and no way of exploiting this systematic variance, there is no representing.

Thus there can be no question of interest concerning merely "what explains the proliferation and survival of the belief-producing mechanisms and the organisms containing them when [some particular] p is produced and all is working properly" (Peacocke p. 00 {manuscript p. 177}). First, we are not in any way interested in categorial descriptions of how particular beliefs function. We are interested only in relational descriptions and relational explanations of function. And we are interested in these only as they cover, univocally, very numerous exemplars of concepts and beliefs manufactured by ancestors of the general concept and belief producing systems, ancestors numerous enough and ancient enough for natural selection to have effectively operated over them. Second, we are interested only in those conditions in the environment that the belief producing systems have a way or ways, under certain ideal conditions, of mapping in a principled and effective way. The fact of actual

correspondence of conditions by these mapping rules must figure in a uniform explanation of how very numerous past cases involving consumer uses of such maps happened to aid selection of the organisms harboring them. The form of the normal explanation sought, hence the description of the mapping rules sought, must be general enough to cover the bulk of effective past cases. It cannot possibly be derived by looking just for explanations that tightly cover this or that individual instance of usefulness. For every natural occurrence there exist innumerable correct explanations. The explanation we are seeking is one that covers the relevant past cases univocally, that shows how producer and consumer have historically cooperated with one another under ideal conditions, the one systematically producing certain mappings, the other systematically using them.

"In a nutshell," Peacocke tells us, "the Problem of Reduced Content is this: how is the teleological theorist to block an incorrect assignment of content to beliefs, vis. one which requires for its truth merely the truth of all the logical consequences of p which have a causal impact on the thinker, rather than the stronger condition of the truth of p itself?" (p. 00 {ms.p 177}) The assignment is blocked by the fact that it could not possibly accord with the requirements laid down, namely, that the mapping rules be a single set covering the majority of effective past representations produced by ancestor systems, a set mapping representations onto representeds such that the one varies systematically with the other, such that the representation producers have a systematic way of effecting this mapping, and such that the consuming mechanisms have a systematic way of using the maps that needs conformity to this mapping rule. How could the representation producers possibly manage in a principled way to vary the representations they produced depending upon which consequences were in fact destined at some future date to have an impact on the thinker? And how, if they could do this, would this sort of mapping be of any use to the consumer? We are talking here about a compositional semantics (broadly conceived--see chapter 00 above) for an historically effective representational system, not an accidental coincidence of two facts.

More specifically, consider first Peacocke's worry over beliefs about the past, then later his worries about universal beliefs. Concerning any kind of human beliefs we have the problem that we don't yet know much about the general principles governing normal conceptual development. We don't know, for instance, which kinds of cognitive skills have been selected to vary in individuals following variations in their individual environments, and which skills are, instead, fixed skills. We don't know, for example, whether the ability to have beliefs about the past is learned in accordance with some kind of very general principles governing the acquisition of concepts generally, or whether it merely matures in a roughly invariant way given a suitable internal and external environment. But let us suppose that for beliefs about the past to have semantic contents, there must have been prior beliefs, specifically about the past, that functioned so as to help preserve specialized "beliefs-about-the-past" producing and consuming mechanisms.[1]

---

[1]. In (Millikan 1984) chapter 18 I suggested a different way that the mapping rules for beliefs about the past may be determined—-through the mechanisms of public language taken with the law of contradiction.

Beliefs about the past can be very useful. The most obvious mechanism for this is that they combine with other beliefs to produce beliefs about the present or future, these latter beliefs having immediate utility. A normal explanation for this effectiveness makes reference, probably, to an isomorphism between the patterns by which beliefs about the present and future emerge from beliefs about the past, and patterns by which the corresponding present and future states of affairs emerge from past states of affairs. Peacocke apparently agrees that something like this can happen. I suggest that it not only can but does happen routinely. Most of our beliefs about the past are not about ancient history, after all, but about where we put the glue away yesterday, whether we remembered to pay the telephone bill, and whether the TV said it snowed yesterday on the ski slopes we wish to visit this weekend. Of course when these beliefs find a normal use (one depending on their truth) it is usually because the conditions they map end up having some mediated effect upon the person who holds them.[2] Similarly, for a bee dance to find a normal use (one dependent upon its truth) it is required that the worker bees can get to the nectar mapped, that the bees are not somehow screened off from the nectar. We can conclude that surely many beliefs about the past, like some bee dances, never perform all of their proper functions--for beliefs, those functions beyond, say, participating in inferences--due to the fact that the conditions they map are screened off from having any impact on the believer. Similarly, to rehearse the Millikan litany, many sperm tails don't perform their proper functions, don't succeed in propelling their cargos through the walls of ova, because most sperm don't land anywhere near ova, or are by some other means screened off from ova. All this is true. But none of it has any relevance to questions about the function of sperm tails or about the content of either bee dances or beliefs about the past. It is only the past successes of some bee dances and of some beliefs about the past that are relevant here--certain successes that actually had an effect on the constitution of the relevant gene pools. Surely most of the beliefs about the past that we acquire from history books, for example, are not destined to help preserve the species' general beliefs-about-the-past forming capacities—-though they may of course help preserve the species' memorizing capacities, via unrelated (social) mechanisms.

Turn now to universal beliefs. As before, make the simplifying assumption that if universal beliefs have universal contents, there must be specialized genetically determined mechanisms that have historically made and used beliefs that mapped onto (don't hold me to this formulation) "universal states of affairs".[3] Certainly there are many true universal beliefs for which not all portions of the corresponding universal states of affairs do or even could have an impact upon the believer. For example, as Peacocke points out, if there are stars outside my light

---

[2] It isn't in fact necessary in order that a belief about the past should find itself a normal use (one that requires its being true) that the condition it maps end up having some eventual effect upon the person who holds it. For example, the belief might feed into an inductive inference yielding true belief in a natural law, and this latter belief might find a use, without the first belief having any independent impact on the thinker.

[3] How universal representations map onto the world is discussed in (Millikan 1984) chapter 14.

cone, their contribution to the truth of any universal propositions that I believe about stars can have no impact upon me. This is not to the point, however, so long as mapping rules for universal beliefs could be established through historic cases in which all portions of the relevant domain did make a difference and did help to explain how the belief's truth is needed for normal use. And indeed, such cases are common. The overwhelming majority of our everyday universal beliefs concern very small domains, the exhaustion or lack of exhaustion of which is immediately relevant to our practical concerns. Thus the whole domain matters when I believe that all the cookies I made have been eaten, that all the students in Philosophy 102 showed up for their final examination, that none of the members of our philosophy department lives on campus, that all of the family bicycles are in use. Surely if we do have special mechanisms designed to make and use universal beliefs, these are the sorts of cases that historically put them in business.

Suppose that I make a million pound bet with Peacocke that I will get at least one legitimate reply when I place an add on the front page of the London Times offering his million to the first Londoner over seven feet tall to make himself known to us. If I lose this bet, my loss and Peacocke's gain will very likely be explainable by reference to the truth of his belief that all Londoners are under seven feet tall and the falsity of my belief to the contrary. Of course it's possible that he will win not in this way, but only because all the Londoners seven feet or over neither read nor talk to anyone who reads—-or because they are all locked in jail, or already have too many millions to care, or all get hit by trucks on the way to the post box. But these would not be biologically Normal ways for his belief that all Londoners are under seven feet tall to help him to prosper.

The theory of mental representation proposed in "Biosemantics" (Chapter 00 above) invites a traditional question about the status of psychological explanation. According to the biosemantic theory, both the fact that a belief is a belief at all, and the fact that it has a certain semantic content, rest directly on the history of the belief. Nothing in the belief's current dispositions, hence nothing about its causal potencies, makes it be the belief that it is. How then can the fact that a person has certain beliefs help to causally explain how that person acts? The last sections of "Explanation in Biopsychology" address this question.

Explanation in Biopsychology

VI Biopsychology is a Predictive Science if at all, then only Accidentally

These reflections on the nature of intentional psychology entail that, as a biological science, it does not aspire to be predictive. Biopsychology studies what happens when biological processes proceed normally, but the normal is neither the necessary nor always the statistically average. Prediction and control do of course play an important role under parts of the wide umbrella called psychology --I have mentioned psychological testing, human engineering, psychotherapy, etc.--but prediction and control are not required byproducts of intentional psychology. Indeed, intentional attitude psychology is a rather unlikely candidate to aspire to the detailed prediction of individual human behaviors.

This is true for at least two reasons. The first is diversity among individual constitutions. For psychology to predict individual behaviors, just as a starter, babies would have to be born cognitively and affectively, indeed also physically, alike. But it is abundantly clear that different newborns, inserted into identical environments, would not behave at all alike, unless under the most general and vacuous of descriptions. People are born with predispositions to different cognitive and affective styles, with different cognitive strengths and weaknesses. Non-psychological factors such as body-build, reaction time, energy level and health also play a large role in determining behavior. Further, it is likely that many aspects of our cognitive processes are partly stochastic, hence that which among many possible solutions to a given problem an individual discovers and executes often is not governed by well-defined psychological principles at all. Surely nothing short of complete physical and chemical analysis could in fact predict the detailed behaviors of any individual. The individual is not a replica of its ancestors or of its friends. It is a bundle of heavily redundant unfolding sub-systems adapted each to the others' concrete peculiarities to form co-ordinated larger units, this in accordance with principles of co-ordination and development all of which are, as yet, subject to merest speculation. We are still trying to find out how an individual's muscles and tendons grow the right length to fit the individual's bones, let alone how the various facets of individual cognitive development and function grow into a coherent unit. But there is no reason to suppose that exactly how an individual thinks is any more governed by laws quantifying over individuals than, say, how he walks or plays tennis-- or how he reacts to allergens.

The second reason that intentional psychology cannot be required to predict individual behaviors is that there is no compelling reason to suppose that all or even most of the norms that it describes are usually fulfilled. Most obvious, as has already been noted, is that the environment cannot always be relied upon to do its part in completing the functions of the cognitive systems. Because this is so obvious, it has been equally apparent to all that there could not possibly be any reliable laws of organism distal environment interaction, certainly not for the

case of humans. Hence theorists who take it that psychology's main business is to deal in laws have found it necessary to insist that a scientific intentional psychology would have to be "narrow", that it would have to ignore the environment. But it is also likely that those portions of the cognitive functions that are carried out inside the organism are abundantly vulnerable to failure. The cognitive mechanisms seem to be paradigms of functional redundancy and layered back-up systems, commanding a variety of means to the accomplishment of the same or functionally equivalent projects. If at first you don't succeed, try another way, is a fundamental heuristic for our cognitive functions. Witness, for example, the well-documented variety of forms of compensation employed by those with brain damage. This redundancy strongly suggests the vulnerability of various cognitive techniques taken separately. It follows that there is little reason to suppose that the exact progression of anyone's inner cognitive systems could be predicted on the basis of even the most exact understanding of all types of human cognitive teleo-function, an understanding of all the biological norms involved.

Suppose, for example, that man is indeed a rational animal; that conformity to certain logical principles is a biological norm for human thought-processes. It would not strictly follow that conformity to reason was so much as a common occurrence. Reasoning could be one among other functions of the behavior-controlling systems, one that sometimes worked and was then to the organism's advantage. It could also be one that seldom caused irreparable damage when it failed, due to redundancy and to backing by various cruder behavior-controlling devices such as those found in the lesser animals. Indeed, remembering the way evolution works, it seems that there must at least have been a time when human reason had exactly this tentative status. It is not likely that the ability to reason well or to learn to reason well arrived all at once in a single lucky mutation. And we can raise the question of how well, in fact, the average modern human reasons. Clearly from the fact that drawing rational inferences may be a norm for the human cognitive systems it does not follow that any reliable predictions about inference patterns can be made. Even though man is a rational animal, rational psychology could remain very far from a reliable predictive science.

But a strong contemporary tradition has it that rational psychology must be a predictive science if it is to be a science at all, and that its central job is exactly to predict individual humans' behaviors. It is claimed, further, that our layman's way of thinking about intentional mental states constitutes a "folk psychological theory," the central employment of which is to effect prediction of the behaviors of our fellows, for this is necessary in order to project our own paths through the tangle of other folks' actions. Do I maintain that it is mere illusion that we thus predict the actions of others?

No, it certainly is not an illusion that we do a lot of correct predicting concerning the behaviors of others. Most ordinary forms of social intercourse and social co-operation would be impossible if we could not. But there may be a misunderstanding over the methods that we typically employ for prediction. The tool that we most commonly use, I suggest, is not a theory of the inner mechanisms that lie behind predicted behaviors. The tool is not, for example, belief-desire theory. Most of our predictions are done with a much blunter tool--the method of brute correlation. In many cases there is, of course, some understanding of the outlines of the psychological mechanisms lying behind predicted behaviors, but our predictions do not usually rest on this understanding, either at all, or at least very deeply. They rest mainly on observations of past behavioral regularities for the individual and for the group(s) to which the individual

belongs. Within fairly well-defined limits, people, especially people from the same culture, just do behave uniformly in a theater, on the road, at the grocer's, even when recreating in the park. Most people are more likely than not to meet what others consider to be their business and social obligations, to conform to general expectations concerning what is appropriate or seemly and, very important, to do the things they have said they will do. Beyond this, we project ahead patterns observed in the past for particular individuals. Known personality traits, character traits, and habits serve as our guides. Of course such knowledge merely limits the boundaries of people's likely behavior. It does little or nothing towards actually determining behavior in its variety. But seldom do we make an attempt to predict others' behaviors in much more detail than this. How inept we actually are at predicting behaviors, even of our best friends and family members, when these behaviors are not covered by known regularities, may be illustrated by friends who become separated in a large crowd, say at a fair, each trying in vain to outguess what the other will do in an attempt at reunion.

VII Reasons and Causes

If we are rational, what that means is that rationality is a biological norm for humans, not that rationality is necessitated by special causal laws of human psychology. Compatibly, it is standard nowadays to claim (though on somewhat different grounds) that thoughts categorized in accordance with their semantics are not the sorts of things that could, even in principle, fall under causal laws. On the account of this essay, the semantic category of a thought is determined relative to its biological functions, which depend in turn upon its history, upon its place relative to certain prior events. But having a certain history is not, of course, an attribute that has "causal powers." Hence reasons can not be, as such, causes. More generally, that a thing has a teleo-function is a causally impotent fact about it. Especially, it is never directly because a thing has a certain function that it performs that function or any other function. More nearly the reverse is true. The thing exists and has a certain function because things homologous to it have performed that function (better, had that effect) in the past. Moreover, here the "because" is only partly causal, the other part is constitutive or logical.

But perhaps it will be thought that although things that have functions cannot be supposed to perform these functions either on account of having these functions or in accordance with strict causal laws, still they must perform them in accordance with ceteris paribus laws. Roughly, there have to be conditions under which the functional item would perform its functions since there have to have been conditions under which its ancestors did perform these functions, and the same kind of item in the same kind of conditions would do the same kind of thing again. This ignores defective members of function categories--diseased hearts, injured limbs etc. It also ignores the fact that performance of their functions is, for many items, a relatively rare occurrence. Would we really wish to speak of ceteris paribus laws in cases where ceteris are not paribus most of the time? And it ignores also a third point.

Characteristically, the same function could, at least in principle, be performed by many differently constituted items. But if these items are differently constituted, if they operate in accordance with different principles, then the supporting conditions required for them to effect this function must differ as well. Brain cells performing the division algorithm require oxygen whereas computer chips require electric currents, and so forth. Similarly, the outer world conditions that support the bat's mosquito-locating abilities and those that support his mosquito-catching abilities are different from those that support the same abilities in humans.

(The bat can perform in the dark on silent mosquitos; humans cannot.) The result is that there are no ceteris paribus laws covering all items having a certain function. For ceteris paribus conditions are unspecified conditions that must remain the same from case to case for the law to hold, whereas here the necessary conditions would have, precisely, to vary from case to case. A "law" applying to all such cases could say no more than that the items falling under the law could be made, by adding different circumstances tailored specifically to each case, to perform the function. But surely anything can be made to effect anything if one adds the right intervening media, if one adds enough special enough circumstances. So any such "law" would be empty. There are no causal laws of any kind, then, that directly concern the causal efficacy of reasons as such. The closest we could get would be ceteris paribus laws for human reasons, other ceteris paribus laws for dolphin reasons, still others for Martian reasons, and so forth.

VIII Normalizing Explanations

Our argument suggests that explanation of an agent's behavior by reference to reasons for acting is not best analyzed as explanation by subsumption under causal laws. The question that arises then is what kind of explanation the citing of reasons for acting is, and how it can still be causal-order explanation. Intentional attitude explanations of behaviors proceed, I will argue, by subsumption of behaviors under biological norms rather than laws, and/or by noting departures from these norms and, perhaps, causes of these departures. Following Philip Pettit (1986), to whose views mine run parallel here, I call such explanations "normalizing explanations." The status of explanations of individual behaviors by reference to reasons concerns the relation of normalizing explanations to other forms of causal-order explanation that are, perhaps, better understood.

To explain a phenomenon by subsuming it under norms is to exhibit it as an instance of conformity to or departure from proper operation of some teleological system. A very simple form of normalizing explanation explains the occurrence of a phenomenon by reference, merely, to something whose function it was to produce that phenomenon. For example, the dishes are clean because they have been put through the dishwasher; the washing machine door is locked because the washer is not finished spinning and the door is designed not to unlock until it is finished spinning; the bear is asleep because it is winter and it is (biologically) normal for bears to sleep through the winter (see chapter 00 above).

In order to explain a phenomenon this way it is necessary, of course, to classify it appropriately as the outcome of a teleo-functional process, and this classifying may itself count as a simple form of explanation. What is happening? What is it doing? It's washing dishes, not making soup or just dirtying the water; it's winding a magnetic coil, not storing wire on a spool; it's resting, cooling its motor between cycles, not playing dead, or broken; and so forth.

More complex normalizing explanations tell or implicitly refer to the place an event has in a series or interdependent pattern of functions, or tell where, and perhaps why, malfunction occurred within such a series or pattern. Thus, that cog-wheel's turning in the calculator is its carrying one in a certain addition algorithm; the car went through the light because its brakes failed; the outboard stalls because there's dirt in the carburetor that gets into the needle-valve.

Normalizing explanations often make reference to conditions that must be presupposed for normal operation of a device or system. Thus the outboard won't start because the spark-plugs are wet or because there's no gas it the tank, the scuba-diver passed out because it was too cold or because his tank ran out of oxygen, and so forth.

Finally, the relation between certain conditions of the functional system itself or of the environment and certain states of the system that normally adapt the system's progress to those conditions may be targeted in a normalizing explanation. Thus, the motor is racing because the heavy-load switch is on but the load is not heavy; the washer failed to fill properly because the soap was put into the tub rather than into the dispenser so that the rising suds tripped the water cut-off before the tub was full; the animal's winter-approaching detectors failed because it was kept indoors, which is why it is attired inappropriately or is behaving inappropriately to the season.

IX The Normal Roles of Beliefs and Desires

Notice how natural it would be to say in the last two of these cases that the washer thought it was full when it wasn't and that the animal's system didn't know it was winter. This is because a belief or a bit of knowledge is likewise a teleo-functional item, one whose function is to adapt the containing system so that it can perform its functions under certain conditions, namely, those conditions which the belief is about. Or, being a little more precise, it is the belief-forming mechanisms that produce the adaptations, the adjustments of the organism to the environment, the beliefs. Beliefs themselves are functionally classified, are "individuated", not directly by function but according to the special conditions corresponding to them that must be met in the world if it is to be possible for them to contribute to proper functioning of the larger system in a historically normal way. Somewhat similarly, the water switch's being off will promote the washer's tasks normally only if the condition is fulfilled that the washer is full. And the animal's winter detectors' being off will effect appropriate functioning of the animal in accordance with historically normal reasons only if winter is not yet approaching.

Explicit human beliefs, however, are much more than just biological adapters to certain environmental conditions. They are adapters that perform their tasks in a certain sort of way, namely, through participation in inference processes. A picture that I advocate but will not try to defend here (see chapter 00 above, and Millikan 1984,1986,1991) shows beliefs and desires as working for the organism by modeling (in accordance with very abstract mathematical mapping functions) the environment, modeling the organism's goals, and modeling types of environmental transitions that the organism knows how to bring about. Normal practical thinking then involves tinkering with these models until solutions are found that will effect transitions from the present state of the environment to various desired states. On this picture the teleo-functions of desires (which they may not very often perform), like those of blueprints, are to effect what they model, to get themselves realized. When everything goes according to norm, action guided by the models inside is action conformed to the outside world so as to issue in productive loops through the environment. This happens in accordance with explanations that, made fully general, that is fully spelled out relationally, apply perfectly generally to all successful uses of the (same capacities of the) species' cognitive systems, historical and current. Theoretical inference is then interpreted as a process whereby the internal model of the environment grows or extends itself in accordance with principles that model various logical, geometrical, and causal necessities or regularities or dependencies in the environment.

Be all this as it may, what seems quite certain is that there must exist some sort of systematic teleo-functional organization of the human cognitive systems whereby the making of good practical and theoretical inferences corresponds to normal (but perhaps not average) functioning for beliefs and desires, and whereby it is biologically normal (not average) for desires to be

fulfilled, at least under certain conditions. (Why else the capacity to have desires?). Accordingly, explanations of behaviors by reference to reasons for action are normalizing explanations.

X How Normalizing Explanations Circumscribe Causes

Why it is that normalizing explanations explain, how it is that they fall under a general theory of explanation, is too large a question for this essay. Our question here is only how such explanations connect with simpler kinds of causal explanation. One connecting link is that whatever has a teleo-function has a normal way of operating, a normal way of performing its function. For functional artifacts this may be, in part, the way the designer proposed that the function be performed, for biological devices, it is the way the function has been performed historically. An exhaustive analysis of the way, given its history, that any functional item operates when operating normally, arrives eventually at a description of normal physical structure for such a device and normal physical conditions for its operation, such that physical laws generate performance of this function given this structure and these conditions. By making implicit reference to such causal explanations, normalizing explanations may thus circumscribe quite specific physical explanations without detailing them.

Guided by Cummins (1975, 1983), we notice that the analysis of how a system normally functions may have several parts. First, the larger function or functions of the system may be analyzed into sub-functions that are performed either serially or simultaneously or in some more complicated pattern of interaction. This kind of analysis Cummins calls "functional analysis." Cummins suggests that a functional analysis may generally be represented by a flow-chart, but of course highly parallel processes, especially those that interact to some degree stochastically, must be represented otherwise. Second, the system may be analyzed into sub-systems, which may or may not correspond to discrete physical parts, each of which is responsible for a designated set of sub-functions. This kind of analysis Cummins calls "compositional analysis." Compositional analysis results in a description of the normal (not necessarily actual) constitution of the system by reference to parts described teleo-functionally, that is, normatively rather than dispositionally. (Here I depart from Cummins, who equates functions with dispositions.) Finally, the normal physical constitutions of the elements normally composing the system may be described, along with the surrounding physical conditions required for normal functioning, and it may be shown how these descriptions together account, in accordance with physical law, for cases of normal operation. That is, the system may ultimately be analyzed into a set of physical parts and physical dispositions rather than, merely, functionally categorized parts and normal functions.[1]

---

[1] This does not imply that, given a certain species, there is a classically understood type-type identity relation between, say, normally constituted and normally functioning beliefs and desires about x on the one hand, and certain physiological structures on the other. Certainly if the physical constitutions of human beliefs are typed categorially there is no reason at all to suppose that any such identity holds. If there are bridge laws for humans that map the semantics of thoughts onto physiological structures, surely what these laws map is certain semantic

By reference to the possibility of this kind of physical analysis, explanations of behaviors according to reasons for action may circumscribe physical causes. Compare explaining why a man shakes by saying that he suffers from Brown's syndrome, even though the etiology of Brown's syndrome may not be known. Or compare explaining why a man has brown hair by saying he has genes for brown hair rather than, say, having dyed his hair, though no one knows the constitution of the gene or how it produces brown hair.[2]

That this is not the complete answer to how reasons circumscribe causes becomes evident, however, when we remember that devices falling in the same function category can have widely varying constitutions. For example, we do sometimes explain, say, how John managed to get the can open by noting that he finally found a can-opener, but given the enormous variety among can-openers, the various different principles on which they may work, how could such an explanation possibly do anything towards circumscribing physical causes or types of physical processes lying behind the can's having come open? Similarly, if there really were various other creatures designed quite differently from humans and made of quite different stuffs but who still had beliefs and desires, then explanation of actions by reference to beliefs and desires without mention of the species of creature involved would seem not to circumscribe any particular kind of physical process at all (cf. Block 1990).

But, looking more closely, whether it circumscribes a kind of process depends on how you type your kinds. Behind every normalizing explanation is a device or system with teleo-functions, and an item acquires a teleo-function only by having a very special sort of causal history. For example, if the cat's purr is explained as produced by a purr-box, an organ especially designed, in the smaller cats, to produce purrs, then we know that the purr-box itself has resulted, ultimately, from the operation of prior purr-boxes in ancestor cats which produced purrs, these purrs somehow having survival value, contributing an essential link, at least occasionally, to the historic cat-chain. Thus a salient cause of the purr is a series of prior purrs. Of course when the functions referred to by normalizing explanations are described categorially though they are actually derived from relational functions, no such simple analysis applies. Still, to assign to any phenomenon a place in a functional system is to claim that it has emerged from a very special kind of causal-historical process, a kind that defines functionality. It is to distinguish its particular type of causal origin quite sharply from other etiological patterns.

---

relations among beliefs and desires onto physical relations among these, hence principles of logical interaction onto principles of causal interaction, not categorial meanings onto categorial physiological "shapes."

[2] Compare (Block, 1990).

BIBLIOGRAPHY

[for "Explanation in Biopsychology" only]

Block, N. 1990 Can the Mind Change the World? in G. Boolos, ed., Meaning and Method; Essays in Honor of Hilary Putnam (Cambridge University Press).

Cummins, R. 1975 Functional Analysis Journal of Philosophy 72:

---------1983 The Nature of Psychological Explanation (Bradford Books/MIT Press).

Jackson, F. and Pettit, P. 1988 Functionalism and Broad Content Mind 97 pp. 381-400.

Kim, J. 1989 The Myth of Nonreductive Materialism Proceedings of the American Philosophical Association 63.3:31-47.

Millikan 1984 Language, Thought and Other Biological Categories Bradford Books/MIT Press.

---------1986 Thoughts Without Laws: Cognitive Science with Content Philosophical Review, 95.1:47-80.

---------1991 Compare and Contrast Dretske, Fodor, and Millikan on Teleosemantics, Philosophical Topics

Pettit, P. 1986 Broad Minded Explanation and Psychology, in P. Pettit and J. McDowell, eds. Subject, Thought, and Context (Oxford, Clarendon Press):17-58.

3951 words in Bet With Peacocke
2726 words in sections VII to the end, Explanation in Biopsychology.

****

The thing that was the dime, throttle, brake, belief was the thing that caused the action -- stopping, gumpopping, preying. Davidson or Dretske: but it wasn't because it was a dime, break, throttle, belief that it caused the action. So how is its being a dime, brake, etc relevant more than a simple pointing at the relevant item? It was the action of his that took place at &:03.26 pm that did it; the motion in his cortex at; It wasn't his stepping on the brake that made the car stop. Brakes can be broken!!

Johnny lost in the first round not bcause he is a poor player but because came up against the kid who was junior U.S. champion. that explains it. Whatever it was, whatever programming, technique, or whatever beat all those other kids beat Johnny. The same propensities, structure, that made its ancestors work that way hence that accounted for its existence ansl made it work. But having that structure is not the same thing as or merely an instanciation of being a brake. Being a brake does not = being a disjunct of being a brake's determinates. Being a brake does not supervene on a lower efficatious structure. it is not "coinstantiated" with--(Mcdonalds).

Performance of a proper function-a systematically repeated pattern of causation/explanation.

There also may be a high correlation between stepping on the brake and the car wheels stopping. A high correlation because of a "common cause"? Things that are brakes are usually things that can stop cars. A good predictor even if not the cause. That's just what you would expect given that he stepped on the brake.. *