

Ridgeview Publishing Company

Pushmi-Pullyu Representations

Author(s): Ruth Garrett Millikan

Source: *Philosophical Perspectives*, Vol. 9, AI, Connectionism and Philosophical Psychology (1995), pp. 185-200

Published by: [Ridgeview Publishing Company](#)

Stable URL: <http://www.jstor.org/stable/2214217>

Accessed: 13/11/2013 11:47

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.




Ridgeview Publishing Company is collaborating with JSTOR to digitize, preserve and extend access to *Philosophical Perspectives*.

<http://www.jstor.org>

PUSHMI-PULLYU REPRESENTATIONS

Ruth Garrett Millikan
University of Connecticut and
University of Michigan



I. Introduction

A list of groceries, Professor Anscombe once suggested, might be used as a shopping list, telling what to buy, or it might be used as an inventory list, telling what has been bought (Anscombe 1957). If used as a shopping list, the world is supposed to conform to the representation: if the list does not match what is in the grocery bag, it is what is in the bag that is at fault. But if used as an inventory list, the representation is supposed to conform to the world: if the

list does not match what is in the bag, it is the list that is at fault. The first kind of representation, where the world is supposed to conform to the list, can be called “directive”; it represents or directs what is to be done. The second, where the list is supposed to conform to the world, can be called “descriptive”; it represents or describes what is the case. I wish to propose that there exist representations that face both these ways at once. With apologies to Dr. Doolittle,¹ I call them pushmi-pullyu representations or PPRs.

PPRs have both a descriptive and a directive function, yet they are not equivalent to the mere conjunction of a pure descriptive representation and a pure directive one but are more primitive than either. Purely descriptive and purely directive representations are forms requiring a more sophisticated cognitive apparatus to employ them than is necessary for these primitives. Purely descriptive representations must be combined with directive representations through a process of practical inference in order to be used by the cognitive systems. Purely directive representations must likewise be combined with descriptive ones. The employment of PPRs is a much simpler affair.

Perhaps the most obvious PPRs are simple signals to conspecifics used by various animals—bird songs, rabbit thumps, and bee dances, for example. But PPRs also appear in human language and probably in human thought. Illustrations in language are “No Johnny, we don’t eat peas with our fingers” and “The meeting is adjourned” as said by the chair of the meeting. Human intentions are probably an example of PPRs in thought, serving at once to direct action and to describe one’s future so that one can plan around it. Our inner representations by which we understand the social roles that we play as we play them are probably also PPRs. The natural way that we fall into doing “what one does,” “what women do,” “what teachers do,” and so forth, suggests this. I suspect that these primitive ways of thinking are an essential glue helping to hold human societies together.

I view PPRs from within a general theory of representations developed in other places (Millikan 1984, 1993). I will start by sketching just a little of that theory (though not nearly enough to defend it).

II. The background theory of representation

Brentano took the essence of intentionality to be the capacity of the mind to “intend” the nonexistent. In recent years it has become generally accepted that he was right in this sense: the core of any theory of representation must contain an explanation of how misrepresentation can occur. I have argued that misrepresentation is best understood by embedding the theory of intentionality within a theory of function that allows us to understand, more generally, what malfunction is. For this we use a generalization of the biological notion of function as, putting things very crudely, the survival value of a reproduced type of entity. I call functions of this generalized kind “proper functions” (Millikan 1984, 1993).

Think of the proper function of a type as what it has been doing to account

for its continued reproduction. The possibility of misrepresentation is derived from the possibility that a token may fail to perform a function that has been accounting for continued reproduction of its type. In some cases, failure to function properly may be even more common, statistically, than proper functioning. For example, many biological mechanisms fail much of the time to have those occasional propitious effects that have nonetheless accounted for their proliferation in the species. Thus the eyeblink reflex, exhibited when any object moves too close to the eye, may be triggered uselessly many times for every time it actually prevents foreign matter from entering the eye. In order to proliferate, often a type needs to perform properly only in a some critical but small proportion of cases, a proportion that varies widely depending upon other factors that enter into the mechanisms of evolution by natural selection.

We can apply the notion *proper function* directly to natural languages. Whole sentences are not usually reproduced, but phonemes and words are reproduced, and the syntactic forms in which they are placed are reproduced. They are copied from one generation of speakers to the next, and reproduced by the same speaker on various occasions. If some theory of Universal Grammar is correct, then certain very general grammatical features are also reproduced via the genes. But more immediately, it is clear that speakers of a given language reproduce words patterned into certain concrete syntactic forms, rather than certain other forms, because the effects these patterns have had upon hearers, in some critical proportion of cases, are effects the speakers wish to reproduce. These effects are described with reference to the semantic and conventional pragmatic rules of the language.

The proper effects upon hearers of language forms—the proper reactions of hearers to these forms—are also reproduced. This is not necessarily because hearers directly copy one another's reactions, though this may sometimes happen. These reactions are reproduced because hearers, at least often enough, benefit from reacting properly to the sentences they hear. By moving from sentences into belief or into action in conformity with the rules of the language, often enough hearers gain useful knowledge, or their actions find a reward, hence they reproduce conformity to these rules. Thus each generation of hearers learns to accord in its reactions with the expectations of the previous generation of speakers, while each generation of speakers learns to expect those same reactions from a previous generation of hearers—with sufficient frequency, that is, to keep the language forms in circulation. Similarly, the proper reaction of a bird to the song of a conspecific, though not for most species an imitated response, is also reproduced—in this case, genetically reproduced. In both cases, the proliferation of representations and the proliferation of proper reactions to them are each dependent on and tailored to the other.

The proper function (or functions) of an expression in a public language may contrast with the function that a speaker intends for it on a given occasion. In (Millikan 1984, chapter 3) I have shown how the speaker's intention in use can lend tokens of a language device additional proper functions, as in the case

of Gricean implicature. These additional functions are not functions of the *public* language forms, however. Indeed, the two layers of function, public and private, can sometimes conflict, as in the case of purposeful lying and of certain parasitic uses of language. We will not need to consider this sort of complexity here, but it is well to warn in advance that if there are pushmi-pullyu forms existing in the public language, as I will argue there are, these should not be confused with the well known phenomenon of Gricean implicature. Public pushmi-pullyu forms are double-aspected on the first layer of function, the public language layer. It is not merely the user's intention that produces the pushmi-pullyu effect.

To understand what *inner* representations are (percepts, thoughts), we apply the notion of a proper function not directly to the representations themselves, but to the mechanisms whose function it is to produce and to use inner representations. When functioning properly, inner-representation-producing mechanisms produce representations in response to, and appropriate to, situations in which the individual organism finds itself. In humans, these mechanisms are exceedingly complex, including mechanisms of belief and desire formation, and also mechanisms of concept formation, inference, decision making and action. When the entire system functions properly, the belief-forming mechanisms produce true beliefs and the desire-forming mechanisms produce desires the fulfillment of which would benefit the organism. But it is possible also to sharpen the notion "proper function" in such a way that inner representations themselves are seen to have proper functions, as follows.²

Many biological mechanisms are designed to produce alterations in the organism in response to some aspect of the environment, so as to adapt or "match" the animal to that aspect, hence to serve some further function within that environment. A simple example of this is the mechanism in the skin of a chameleon that rearranges its pigment so that its color will match that of its environment. This then serves the further function of concealing the chameleon from predators. Each *particular* coloring of the chameleon produced is naturally said to have a function too—the same function: concealing the chameleon from predators. Similarly, any state constituting a *stage* in a (proper-) functional process, when the shape of the process and hence of the state is determined by input from and as a function of certain features of the environment, can be viewed as itself having proper functions. The proper functions of this state are to help in the production of various further stages in the process to which this state will give rise if the whole system continues to function properly. In this way, specific inner states—even quite unique inner states, unique because induced by unique organism-environment relations—may have proper functions. In this way, even such representations as someone's desire to climb Mount Everest backward, or someone's belief that persimmons cure mumps, can be considered to have proper functions.

But I have not yet said what any of these proper functions are, what it is that representations, either inner or outer, properly do.

III. Descriptive and directive representations

Elsewhere I have defended a proposal to explain how the content of a representation—its satisfaction condition—is derived. The explanation takes time, and it is different for descriptive and for directive representations. Here I will assume the notion of content, explaining only what I take the difference between descriptives and directives to be.

A representation is *directive* when it has a proper function to guide the mechanisms that use it so that they produce its satisfaction condition. Like a blueprint, it shows what is to be done.

Desires are directive representations. To see how this might be so, it is important to remember that the proper function of an item can be a function that it is unlikely to perform. Perhaps the sad fact is that an overwhelming majority of our desires never become satisfied. Many (e.g., the desire to square the circle) may even be incapable of becoming satisfied, or (e.g., the desire that it rain tomorrow) may be incapable of being satisfied by normal operation of those mechanisms that are designed to help fulfill desires. This has no bearing on the claim that desires are satisfied when things proceed “properly,” that is, when things proceed in the ideal sort of way that accounted for the survival and proliferation of those integrated systems whose job it is to make and use desires. Surely the job of these systems is, first, to produce desires that would benefit one if fulfilled and then, second, given certain additional propitious inner and outer circumstances, to be moved by them to their fulfillment.

Sentences in the imperative mood are, paradigmatically, directive representations. They proliferate in the language community primarily insofar as they (often enough) help to effect the fulfillment of their satisfaction conditions.

Unlike directive representations, what makes a representation *descriptive* is *not* its function. Rather, the descriptive representation’s truth condition is a condition to which it adapts its interpreters or users in the service of *their* proper functions (Millikan 1984 chapter 6, 1993 chapters 4-6). It is a condition that must hold if the interpreter’s (proper) way of being guided by the representation is to effect fulfillment of the interpreter’s functions in accord with design. For example, beliefs are descriptive representations. If my belief that there is an umbrella in the hall closet is to help guide my decision-making and action-guiding apparatus (i.e., the belief’s “interpreter”) such that it serves its function of helping to fulfill my desires, for example my desire to keep off the rain, then there needs to be an umbrella in the hall closet. If it is to help me make a correct inference concerning, for example, whether or not Susan returned my umbrella, one that yields truth not by accident but in accordance with the good design of my cognitive systems, still there needs to be an umbrella in the hall closet—and so forth.

Typical sentences in the declarative syntactic pattern have, perhaps, a descriptive function. Their function is to produce hearer beliefs or, more precisely, to produce *true* hearer beliefs. For it is only when a certain proportion

of what hearers are told is true, hence is interpreted by them into *true* beliefs, that they are encouraged to continue to conform to certain rules of interpretation from English (from the public language) into beliefs. One could not learn to understand a language if too large a proportion of the sentences one heard in it were false. Because its function is to produce *true* beliefs in hearers, the declarative syntactic pattern is descriptive. Roughly speaking, sentences in this pattern affect their interpreters in a proper way only under the conditions that are their truth conditions.

Earlier I remarked that I was not pausing to explain how the satisfaction conditions of representations are determined. But it will be important to grasp this much. On the analysis I have given, the satisfaction condition of a representation is derived relative to its function. The content of the representation turns out to be an abstraction from a fuller affair intrinsically involving an embedding mood or propositional attitude. Put simply, there is no such thing as content without mood or attitude; content is an aspect of attitude. A corollary, as we will soon see, is that it is possible for the very same representation to carry at once two different contents, one relative to each of two different attitudes or moods which simultaneously embed it.

IV. Pushmi-pullyu representations

Consider first a very primitive representation: the food call of a hen to its brood. A proper function of this call is to cause the chicks to come to the place where the food is and so to nourish them. Assume, what is reasonable, that this is the only proper effect that the call has on chicks, the only effect the call has been selected for. Then the call is directive, saying something like “come here now and eat!”. But it is also a *condition* for proper performance of the call that there be food there when the hen calls. So the call is also descriptive, saying something like “here’s food now”. (Note that the descriptive and the directive contents of this representation are different.) Assume further, what is again reasonable, that the effect of the call on the chicks is not filtered through an all-purpose cognitive mechanism that operates by first forming a purely descriptive representation (a belief that there is food over there), then retrieving a relevant directive one (the desire to eat), then performing a practical inference and, finally, acting on the conclusion. Rather, the call connects directly with action. *Its function is to mediate the production of a certain kind of behavior such that it varies as a direct function of a certain variation in the environment, thus directly translating the shape of the environment into the shape of a certain kind of conforming action:* where the hen finds food, there the chick will go. The call is a PP representation.

Other examples of primitive PPRs (probably) are other bird calls, danger signals used by the various species, the various predator calls used by chickens and vervet monkeys, the dance of the honey bee, and so forth. For example, the bee dance tells at once where the nectar is and where to go. Functioning properly

it produces variation in behavior as a direct function of variation in the environment. Actually, there is evidence that the bee has a map in its head of its environment and that the dance induces it, first, to mark the nectar location on this map (Gallistel 1990).³ Still, assuming that the only use the bee ever makes of a mark for nectar on its inner map is flying to the marked position to collect the nectar, then the nectar on the bees' inner map is itself a PPR. And it seems reasonable to count a representation whose only immediate proper function is to produce an inner PPR as itself a PPR.

James J. Gibson did not advocate speculating about inner representations. Yet his notion that in perception we perceive certain affordances (opportunities for action) suggests that perceptual representations are PPRs. Think of perceptual representations simply as states of the organism that vary directly according to certain variations in the *distal* environment. The perceived layout of one's distal environment is, first, a representation of how things out there are arranged—a descriptive representation. It is also a representation of possible ways of moving within that environment: ways of passing through, ways of climbing up, paths to walk on, graspable things, angles from which to grasp them, and so forth. Variations in the layout correspond to variations in possible projects, and in the paths of motion needed to achieve them. The representation of a possibility for action is a directive representation. This is because it actually *serves* a proper function only if and when it is acted upon. There is no reason to represent what *can* be done unless this sometimes effects its *being* done. Compare desires, which serve a function only insofar as they occasionally help to produce their own fulfillment. In the case of perceived affordances, action toward their fulfillment is, of course, directly guided by the percept, variations in the environment, hence in the percept, translating directly into variations in the perceiver's movement.

There are cells in the inferior premotor cortex of monkeys (informally, "monkey see monkey do cells") that fire differentially according to the immediate ends (e.g., grasping small pieces of food with the fingers) of the manual manipulations the monkeys are about to execute, and which also fire when the monkeys see other monkeys perform these same manipulations for the same ends (Rizzolatti et al. 1988). Imitative behaviors in children show up extremely early. One infant was observed in the laboratory to imitate facial expressions (opening the mouth, sticking out the tongue, etc.) at the age of forty minutes (Meltzoff and Moore 1983). We might speculate, on analogy with the monkeys, that these primitive mechanisms of imitation in children employ PPRs, which picture what the other is doing at the same time that they serve to direct what the self is to do. Compatibly, Jeannerod (1994) cites evidence that imagining oneself performing certain movements and actually performing them involve, in part, the same dedicated area of the brain, hence that picturing what one might do and intending to do it may be two sides of numerically the same representing coin. Indeed, one of his suggestions is that imagining an action without at the same time performing it is accomplished by inhibition of normal connections to motor pathways.

It is important to see that PPRs are not merely conjunctions of a descriptive plus a directive representation. PPRs are more primitive than either purely directive or purely descriptive representations. Representations that tell only what the case is have no ultimate utility unless they combine with representations of goals and, of course, representations that tell what to do have no utility unless they can combine with representations of facts. It follows that a capacity to make mediate inferences, at least practical mediate inferences, must already be in place if an animal is to use purely descriptive or purely directive representations. The ability to store away information for which one has no immediate use (pure description), and to represent goals one does not yet know how to act on (pure direction), is surely more advanced than the ability to use simple kinds of PP representations.

V. PP representations in human thought

Organisms often evolve in complexity by modifying less differentiated multipurpose structures into more differentiated dedicated ones. Thus we would expect beliefs (dedicated to facts) and desires (dedicated to goals) to be a later evolutionary achievement than inner PPRs. On the other hand, if there are purposes that could be served as well or better by PPRs than by more differentiated representations, our first hypothesis should be that that is how these purposes still *are* served. I think that there are some such purposes.

One obvious hypothesis is that human intentions are PPRs. If intentions are inner representations, surely they are at least directive ones. They perform their proper functions when they issue in the intended actions. But it is also a common and plausible assumption that a person cannot sincerely intend to do a thing without believing she *will* do it. If one starts in a rather traditional way, assuming that there are only two basic sorts of cognitive representations, purely descriptive ones (beliefs) and purely directive ones (desires), then whether “intending” must involve harboring a descriptive as well as a directive thought may appear to be a matter of “analysis of the concept ‘intention.’” But if intentions are PPRs, then the dual nature of intentions is no conceptual truth but a biological or neurological truth. And there is reason to suppose they might be PPRs.

Suppose that my brain already harbors, for purposes of guiding my action, a representation of what I am definitely going to do. And suppose there is need to take this settled future into account when making further decisions about what else I can compatibly do. It would surely waste space and introduce unnecessary mechanisms for evolution to duplicate the representation I already have. Better just to use it over again as a descriptive representation as well. Notice, however, that this kind of PPR differs from the kinds we have previously discussed in this way. Rather than functioning as do, say, perceptual PPRs, which map variations in the organism’s world directly into (possible) actions, it maps variations in goals directly onto the represented future world. It differs also in that the

contents of the directive and descriptive aspects of the representation are not different but coincide.

A second kind of PPRs that may be fundamental in human thought are primitive representations of social norms and roles. I suggest not that this is the *only* way humans can cognize these norms and roles, but that it may be the primary *functional* way, and that this way of thinking may serve as an original and primary social adhesive. There are good reasons for thinking that humans and other social animals have designed into them mechanisms leading to the coordination of behaviors.⁴ Coordinated behaviors are behaviors that benefit each individual involved in the coordination given that the others also play their assigned roles. Some of the principles governing evolution of such behaviors are now well known. Allan Gibbard suggests that “[s]ystems of normative control in human beings...are adapted to achieve interpersonal coordination (1990, p.64).” His ultimate aim is to cast light on the origins and function of the language of ethics and the thought it communicates, including especially the function of normative discussion in originating what he calls “normative governance.” My project here is less ambitious. My speculations concern only a mechanism for stabilizing and spreading coordinative behaviors that are already in place, and for coordinating expectations.

I have in mind two basic sorts of social coordinations. The first might be called “common norms.” They apply equally to all members of a given society: we drive on the right; we speak at meetings only when duly recognized; we wait in orderly queues; we are quiet at concerts; we honor our contracts; we see to our families first, then to our relatives, then to our friends; and so forth. The second might be called “role norms.” These apply to a person only in so far as he or she is filling a certain role: children obey adults while adults direct children; the chair of the meeting calls it to order and so forth, but does not introduce or speak to motions or vote, while the members do introduce motions, speak to them and vote; pupils raise hands to be called on while teachers speak freely; guests and hosts behave in assigned ways, and so forth.

The norms just mentioned undoubtedly all have coordinating functions. On the other hand, the distinction between common norms and role norms could be applied as well to norms lacking coordinating functions. For example, not eating peas with one’s fingers and not picking one’s nose in public may be noncoordinating common norms, whereas wearing a skirt if you are female and trousers if you are male may always have been a noncoordinating role norm. A mechanism whose biological function is to transmit coordinating norms might well have as a mostly benign side effect the transmission of a good number of non-coordinating norms as well. Thus humans tend to be creatures of convention, exhibiting many patterns of both solo and interactive behavior that are handed on quite blindly, seeming to serve no purpose at all. It may be that our propensity to play games is another side effect of a mentality built to effect coordinations.

The mechanism for stabilizing coordinative patterns of behavior that I

propose is simple. It is the capacity and disposition to understand social norms in a way that is undifferentiated between descriptive and directive. What one does (or what *das Mann* does—remember Heidegger?), what a woman does, what a teacher does, how one behaves when one is married or when one is chair of the meeting, these are grasped via thoughts, PPRs, that simultaneously describe and prescribe. In the primitive mind, these PPRs describe and prescribe what is understood to be The Moral Order, an order taken as totally objective, noninstrumental (absolute) and real, but understood at the same time as stringently prescribed. (In primitive thought, self and others have Sartrean essences with a vengeance!) But it may also be that without the general disposition to think in this way during much of the unreflective parts of our lives, the social fabric would be weakened beyond repair. Yes, I am seriously proposing this as a possible neural mechanism, although supplemented, of course, with more sophisticated mechanisms by which we moderns may also dissect the relevant norms to reveal two faces.

VI. PP representations in human language

If human thought contains PPRs, arguably human languages should contain them as well. Just as they contain forms, on the one hand, whose function is to implant ideas about what is the case, and on the other, to implant ideas about what to do, one would expect them to contain forms whose function is to implant mental PPRs in hearers. For example, if the inner vehicles of our intentions and of our unreflective graspings of social norms are PPRs, it is reasonable that there should be linguistic forms to correspond. There do not seem, however, to be any *dedicated* syntactic arrangements to do these jobs. And indeed, granted PPRs really exist in thought, this lack of a form of expression dedicated entirely to PPRs may do much to explain their near invisibility. PPRs, I believe, are imparted by use of the declarative syntactic pattern. At least this is true for English. That is, the declarative pattern has more than one function. Sometimes it is descriptive, and sometimes it expresses a PP mood.

In the PP mood we say, for example to children, “We don’t eat peas with our fingers” and “Married people only make love with each other.” The job of this mood is to describe and to prescribe, producing at the same time both true expectations and coincident behaviors, the one as a direct function of the other. Notice that the mechanism here is not that of Gricean implicature. Both functions are explicit or literal; the mood is proliferated precisely because it serves both functions at once; both functions are fully conventional.

Strict orders are standardly delivered in the English declarative pattern, which then functions directly: “You will report to the CO at 6 am sharp”; “You will not leave the house today, Johnny, until your room is clean.” This use, I suggest, is more than just directive; it is not just another imperative form.⁵ Its function is to impart an *intention* to a hearer and to impart it directly, without mediation through any decision-making process, for example, without involving

first a desire and a practical inference. This is the PP mood, undifferentiated between directive and descriptive, serving to impart PPRs.

The PP mood may also be used to impart intentions to a group, thus serving a coordinating function. "The meeting will now come to order," when it functions properly, imparts to each member of the group both intentions concerning their own behavior and expectations concerning the behavior of others. The PP mood statement in the university catalog, "Professors will hold office hours every day during registration week," informs both students and faculty, and imparts intentions to faculty.

I would like finally to introduce two more kinds of PPRs that I believe occur in public language but that will require somewhat more lengthy discussions. The first is a variety of explicitly performative sentence. The second are terms expressing "thick concepts."

VII. Performatives

Many performative utterances, I believe, are in the PP mood.⁶ When viewed incorrectly as simple present-tense descriptives, they are puzzling in that they seem to create facts *ex nihilo*, making something to be the case simply by *saying* that it is the case. Suppose, for example, that the chair of the meeting says "The meeting is adjourned," but that nobody pays any attention and three more motions are debated. There does seem to be a sense in which these motions will have been debated not just after the meeting was *called* adjourned but after it *was* adjourned, that is, a sense in which "The meeting is adjourned" is self-guaranteeing. Similarly, once the chair has said "The meeting will now come to order" the meeting has been brought to order even if everyone keeps talking loudly. And if the right person says to the right couple at the right time "I now pronounce you man and wife" then these two are man and wife, even if they don't act that way and even if nobody, including the legal authorities, is disposed to treat them that way. On the other hand, if nobody pays any attention when the chair says "The meeting is adjourned" or, say, "The meeting will now come to order," there is also a tug that says these sentences somehow were *not* true. Let me try to explain why this tension occurs.

Many conventionally or legally molded patterns of activity allow or require the making of, as I will call them, conventional or legal "moves," under designated circumstances, by participants playing designated roles. Examples are performing a marriage ceremony, making a move in a game, appointing someone to a position, making a motion in accordance with parliamentary procedure, legally closing a road, and so forth. These are acts that, once performed, place restrictions on what may or must follow after if things are to accord with convention or law. Moves of this sort may or may not be made with the aid of articulate language forms. For example, in some contexts a bid may be made by saying "I bid," but in others, merely by raising a finger, a vote may be made by raising a hand, and so forth. Most such moves are themselves in essence PPRs.⁷

Their proper function is to channel behaviors that follow so that they take certain forms and not others (directive), and to coordinate expectations accordingly (descriptive). Functioning properly in the usual way, they produce inner PPRs in the various participants which guide them in coordinating their expectations and actions.

The implications of such moves for ensuing behavior are often quite complex, involving complicated mandates or the limiting of options for a variety of affected people playing a variety of roles. For this reason, these implications often could not easily be spelled out in an explicit formula, for example, by saying something simple like “The meeting will now come to order.” Imagine the minister trying to fill in all the blanks in this formula: “You, Jane, and you, John, will now...and the guardians of law will now...and your parents and friends will now...” and so forth, spelling out all that marriage entails. However, there usually are names for conventional or legal moves, or for the situations constituted by their having been performed: “bidding six diamonds,” “checking the queen,” “being married,” “being chairman,” “appointing a chairman,” “voting for a candidate,” “a road being legally closed,” “making a motion,” and so forth. These names are not names for the “shapes” of the moves—not names like “raising one’s finger” or “raising one’s hand.” Rather, they are names or descriptions that classify moves according to their conventional outcomes—by what sorts of things follow after them in accordance with convention or law.

Now there is an obvious and very general meta-convention by which one may use the name of any move in order to make that move. For, assuming that the conventional “shape” that the move has is *merely* conventional, as it generally is, any other shape might be substituted for this shape, so long as everyone understands what move is being made. Saying that I perform the move will generally be enough to perform it, then, granted that the vehicle that usually performs it is arbitrary, and granted that I am the fellow whose move it is to make. For example, even “I move my queen to Q5” may be quite enough to perform that move in chess, especially if circumstances are such as to make an actual board move awkward or impossible (e.g., chess by mail, or while also trying to get dinner).

In English, present-tense declaratives are used in making such moves: “I bid three diamonds,” “I pronounce you man and wife,” “I move that the meeting be adjourned,” “The meeting is adjourned,” “This road is legally closed,” and so forth.⁸ This use of the declarative pattern proliferates because it serves the function of producing the conventional outcomes of the moves named, channeling ensuing activities (a directive function) and coordinating relevant expectations (a descriptive function) at the same time. The proper function of the declarative in each of these cases is thus exactly the same as the function of the move named or described, indeed, the move IS the utterance, in the right circumstances, of the sentence naming it, just as raising one’s hand in the right circumstances IS voting. It is a matter of convention that one can vote by raising one’s hand, and it is a matter of (very sensible) convention that one can make almost any conventional

move by embedding its name in this way in the declarative pattern. These embeddings are called “performative formulas.” Such performative formulas are PPRs. They produce conventional outcomes and they coordinate expectations accordingly, standardly by producing relevant inner PPRs in participants.

But there is another detail. One can, of course, make a conventional move with a performative formula that names that move only granted that one is the right person in the right circumstances to make that move. I cannot, for example, adjourn a meeting by saying “The meeting is adjourned” unless I am chair and we are in the right kind of a meeting and the right time has come. Another function of the performative formula—another thing it does when functioning properly—is to produce true beliefs that the named conventional move is in fact being performed. In this respect the performative is just (is also) an ordinary descriptive. It is true if it is in fact being uttered by an appropriate person in appropriate circumstances, that is, if the move it claims is performed is in fact being performed. Nor does this hang on whether or not it in fact effects its conventional outcome.

Suppose, for example, that I put up a sign on my road saying “This road is legally closed.” If I have reproduced this token of the declarative pattern on the model merely of descriptive sentences I have previously heard rather than copying my use from the cultural pattern of performative uses, then it is not, of course, a sentence in the PP mood at all. It is merely descriptive and, assuming the town has not in fact closed my road, it is plain false. If I have reproduced it instead on the model of PP sentences I have previously heard, perhaps supposing myself to have the legal right to close my own road in this way, then it is in the PP mood. Still, just as if it were a simple descriptive, it is, minimally, false. It is not true that my road is legally closed.

Integrating this now with the previous theme, besides this last sort of truth condition, a performative sentence of this kind also has a directive satisfaction condition and a *second* truth condition as well. It directs that a certain conventional outcome should ensue, and it induces expectations to accord with this. Returning to “the meeting is adjourned,” one of its truth conditions is fulfilled by the fact that it is the chair who says it and at an appropriate time. Whether or not three more motions are debated afterwards, once the chair has said these words, the meeting is in fact adjourned. There is a second truth condition, however, namely, that the proceedings should then actually draw to a close. For only if this happens, will the coordinated expectations which it is the proper job of the performative sentence to induce be true expectations. Last, like any PP mood sentence whose job is to impart intentions, it has the meeting’s actually drawing to a close as a directive satisfaction condition as well.

VIII. Thick concepts

I wish now to speculate about one last example of PPRs in public language. These are sentences that contain words expressing “thick concepts.” Thick

concepts are concepts such as *rude*, *glorious*, *graceful*, that seem to describe a thing and to prescribe an attitude toward it at the same time. Recently there have been absorbing and complex discussions about these concepts that I cannot enter into here. But I would like make a suggestion that might be fruitful to explore.

I hesitate to take any strong position on thick concepts in part because I am not sure, within the framework I have been using, how to understand the nature of those states or dispositions that ordinary speech calls “attitudes.”⁹ It is not likely, for example, that an adequate description of the attitude involved in thinking something rude could be given by reference only to the notions “directive” and “descriptive”. But let us suppose this much: if the function of a term in the language is in part to induce an attitude, then this part of its function is not entirely descriptive. That is, attitudes are other than or more than mere (descriptive) beliefs. If I can make plausible, then, that one function of a term expressing a thick concept is to produce an attitude but that these terms are also descriptive, I shall have shown at least that these terms are close kin to PP representations.

Suppose that certain (perhaps highly disjunctive) configurations of primary qualities tend to produce in us certain attitudes towards their bearers when perceived or contemplated. This might be due to native dispositions, or due to the influence of culture. Secondary qualities are traditionally thought of as powers to produce sensations, but powers to produce attitudes are perhaps similar enough to be called secondary qualities too—“attitudinal secondary qualities”. Secondary qualities are such only relative to a kind of perceiver or, in this case, a kind of reactor. Nonetheless, relative to a certain group or class of reactors, attitudinal secondary qualities are objective properties.

My suggestion is that words expressing thick concepts may, first, *describe* attitudinal secondary qualities relative either to our species as a whole or to the culture of speaker and hearer. That is, declarative sentences using these words attributively will not serve their proper functions in a normal way unless the objects of attribution indeed have certain attitudinal secondary qualities relative to a community encompassing both speaker and hearer. But, second, *this is because their proper function is a directive one*, namely, to produce in the hearer the relevant attitudes. Their function is to cause the hearer to take these attitudes towards these items. That is, these words continue to produce in hearers these attitudes towards designated objects only because the attitudes induced turn out, in a large enough proportion of cases, to be independently “true” by the hearers’ own lights. The objects described are such as actually to *produce*, on direct inspection, or given a more detailed description, *those* attitudes. Words expressing thick concepts thus face two ways at once, describing their subjects and at the same time inducing attitudes towards these subjects. Indeed, perhaps the inner representations induced by these words themselves face two ways at once. Perhaps they are inner PPRs.

Notes

1. And Hugh Lotfing: *The Story of Doctor Dolittle: being the history of his peculiar life at home and astonishing adventures in foreign parts.*
2. Also see, especially, (Millikan 1884) chapter 2, (Millikan 1993) chapter 11 and (Millikan 1994).
3. If the dance says there is nectar at a location where the bee's experience has previously shown there is a large body of water, the bee will be unmoved by the dance (Gallistel 1990).
4. A good informal summary of these reasons is given in (Gibbard 1990).
5. This is a correction of the remarks on army orders in (Millikan 1984) chapter 3.
6. The performative sentences most obviously referred to in the following discussion are those that perform what Strawson (1964) termed "essentially conventional illocutionary acts" or, more accurately I believe, "K-II illocutionary acts" as these are described in (Millikan forthcoming) section VI. The relation of these performatives to those that perform explicit "K-I illocutionary acts" is interesting, but too involved to pursue in a general paper of this sort.
7. I say "most" because some conventional patterns have no functions at all. See (Millikan forthcoming).
8. Some have argued that performatives function by contextual implicature. However, explicit performatives like "I order you to", "I promise that," and so forth, are lacking in some languages, "particularly those with a more developed system of sentence types or those spoken in societies that seem to have less cultural need for formulaic discourse of the kind represented by performative sentences" (Sadock 1988, p.186).
9. Philosophers' "propositional attitudes" of belief and desire are not "attitudes" in ordinary speech at all.

References

- Anscombe, G.E.M. 1957. *Intention*. Ithaca: Cornell University Press.
- Gallistel, C.R. 1990. *The Organization of Learning*. Cambridge: Bradford Books/MIT Press
- Gibbard, A. 1990. *Wise Choices, Apt Feelings* Cambridge: Harvard University Press.
- Jeannerod, M. 1994. "The representing brain: neural correlates of motor intention and imagery". *Behavioral and Brain Sciences* 17:187-245.
- Meltzoff, N. and Moore, K. 1983. "Newborn Infants Imitate Adult Facial Gestures". *Child Development* 54:702-709.
- Millikan, R.G. 1984. *Language Thought and Other biological Categories*. Cambridge:MIT Press.
- Millikan, R.G. 1993. *White Queen Psychology and Other essays for Alice*. Cambridge: MIT Press.
- Millikan, R.G. 1994. "A Bet With Peacocks". In C. Macdonald and G. Macdonald, *Philosophy of Psychology: Debates on Psychological Explanation*. Oxford:Oxford University Press
- Millikan, R.G. (forthcoming). "Proper Function and Convention in Speech Acts". In L.E. Hahn ed., *The Library of Living Philosophers*, Strawson Volume. Open Court

- Rizzolatti, G., Carmada, R., Fogassi, L., Gentilucci, M., Luppino, G., & Matelli, M. 1988. "Functional organization at area 6 in the macaque monkey. II. Area F5 and the control of distal movements." *Experimental Brain Research* 71:491-507.
- Sadock, J.M. 1988. "Speech act distinctions in grammar". *Linguistics: the Cambridge Survey, vol. II*. Cambridge: Cambridge University Press.
- Strawson, P. F. 1964. "Intention and Convention in Speech Acts." *The Philosophical Review* 73(4): 439-60.