

Beliefs and Desires Incorporated

Journal of Philosophy, 91 (8), August 1994, 404-25

Suppose we admit for the sake of argument that "folk" explanations of human behavior--explanations in terms of beliefs and desires--sometimes succeed. They sometimes enable us to understand and predict patterns of motion that otherwise would remain unintelligible and unanticipated. Is the only explanation for such success that folk psychology is a viable proto-scientific theory of human psychology? I shall describe an analysis which yields a negative answer to that question. It was suggested by an observation and an analogy, both of which may initially seem remote from the topic at hand.

I

The observation was simply this: we often talk as if corporate entities such as governments and private corporations have beliefs and desires. Pick up any newspaper and you are likely to find passages such as the following:

The United States has concluded that Serbia is routinely violating the United Nations ban on flights over Bosnia . . . the United States might be prepared to introduce a resolution at the United Nations as early as next week . . . The British Government today gave its strongest public indication that it might be willing to support some degree of enforcement . . . The United States strongly pushed for armed enforcement as part of the initial Security Council resolution . . . But it was forced to retreat because of opposition from both France and Britain, which expressed concern for the safety of their troops taking part in the United Nations peacekeeping operation in Bosnia and Hercegovina.¹

Notice the rich mix of intentional idioms used in this passage. Governments draw conclusions, make public statements, and introduce resolutions. They seem willing to do some things but not other things. They have plans and goals. They indicate their views to one another, and have disagreements. France and Britain are here attributed concern about their troops, which leads them to "express" views opposing those of the United States.

These usages immediately raise a host of questions. We *seem* to attribute beliefs and desires to governments. Do these attributional practices somehow differ from the attribution of such states to people?

¹ *The New York Times*, December 4, 1992, pp. A1, A7.

What does it mean to say that *Britain* is now willing to consider the use of force, or that *France* disagrees? Is it legitimate to talk this way, or must all such talk be reinterpreted?

II

In casting about for answers to these questions, one heuristic is to consult the authorities, and see what they say. Daniel Dennett has over the years carefully articulated a set of conditions which he thinks are necessary and sufficient for the legitimate application of intentional idioms to people or to other "systems." Applying such idioms constitutes what he calls *adopting the intentional stance*. Success legitimates the practice.

I shall argue in this and the next section that every condition that Dennett specifies for adopting the intentional stance can be satisfied as well by corporate entities such as the British Government, General Motors, or the Internal Revenue Service. Dennett sets up a series of hoops; corporate entities can jump (or lumber) through them all. I shall run through the list of conditions sufficient to "adopt" the stance, and show that every one of them is satisfied by our practice of attributing intentional states to governments. Of course this mode of address sidesteps the issue of whether governments--or people for that matter--*really have* beliefs and desires. Some will take my argument to constitute a *reductio* of the "intentional stance." So be it; but I shall be using the argument in pursuit of other ends.

What must one do to "adopt the intentional stance"? What Dennett calls its "flagship expression" is found in his book *The Intentional Stance*.² To adopt the stance, he says

...first you decide to treat the object whose behavior is to be predicted as a rational agent; then you figure out what beliefs that agent ought to have, given its place in the world and its purpose. Then you figure out what desires it ought to have, on the same considerations, and finally you predict that this rational agent will act to further its goals in the light of its beliefs. A little practical reasoning from the chosen set of beliefs and desires will in many--but not all--instances yield a decision about what the agent ought to do; this is what you predict the agent *will* do (*ibid.*, p. 17)

So in order to adopt the intentional stance toward a government, we must decide to treat it as a rational agent, figure out what beliefs and desires it ought to have, given its place in the world and its purposes, and then, using practical reasoning, determine what it ought to do, given those beliefs and desires.

² Daniel Dennett, *The Intentional Stance* (Cambridge: MIT, 1987), p. 3.

In one sense it is trivial to show that we can decide to treat government as rational agents, since we can decide to treat anything as a rational agent. (Dennett for example considers adopting the intentional stance toward a lectern.) We *can* attribute beliefs and desires to governments, since (as our newspaper story shows) we often *do*. The newspaper story also shows that we can predict their behavior using practical reasoning from those beliefs and desires. Pick up any newspaper, or listen carefully to the broadcast news: you are bound to hear governments described as intentional systems. The usage is remarkably frequent. Most of our international news is reported as a high-stakes soap opera. Its players are impulsive, moody, and heavily armed nations.

There is though a less trivial sense in which these questions are significant. Is there some aspect of our practice of attributing intentional states to corporate entities that differs from our attributions of such states to people? In particular, might it differ in such a way as to defeat the claim that governments are intentional systems?

One might first challenge the presupposition that governments can legitimately be treated as rational agents, or indeed as agents at all. Do *governments* do things? We certainly talk as if they do, and we distinguish governmental actions from the actions of private individuals—even when those individuals are the ones carrying out the governmental action. Legally, governments and corporations *are* persons; they own property, enter into contracts, issue statements, carry out actions, and risk liabilities. These properties, contracts, statements, actions, and liabilities do not devolve upon the individuals who at the time were in government employ and acting in their official capacities. It is the United States government that seizes the private property of drug runners; even though particular individuals may carry out that action, *they* do not seize the property, and it does not become theirs but rather becomes the property of the United States government. We distinguish the actions of the agents from the governmental action. It is possible on some subsequent occasion that all the government agents involved pass through qualitatively identical psychological states and perform identical motor activities, yet *not* perform a governmental action. The seizure is not an government action unless (for example) a legal writ has been issued and the agents are standing on ground within their jurisdiction. A perfect forgery of the writ or a late night change in jurisdictional boundaries could render the action null and void.

Can corporate entities have beliefs? Well, as a first step, they certainly make statements. For example, in 1987 the Internal Revenue Service issued an official ruling stating that mortgage interest points charged for refinancing a house could not be deducted entirely in the

year of the loan, but must be pro-rated over the life of the loan. This was the IRS *interpretation* of certain clauses of the Tax Reform Act of 1986. The issue had not been settled in the courts, and the IRS was conveniently stating for us how it planned to interpret the law. We can put it in canonical form: the IRS holds that points for refinancing are not immediately deductible. This attribution of a content to that entity enables us to predict subsequent actions of the agency when returns are examined. Perhaps it helps us to avoid fines and penalties. What makes the attributed content a belief *of* the Internal Revenue Service? We certainly cannot attribute it to the person who wrote the memorandum, since he or she might not personally believe (or even understand) what was written. Perhaps no employee of the IRS believes exactly what is described in the statement—each disagrees with some detail, say—yet they worked it out in committee using the appropriate channels, and the result becomes the official position of the IRS.³ The courts may later *disagree* with the IRS *interpretation* of the laws, and force the IRS to adopt a new official policy. If that happens, it is clear that something has changed—but not necessarily the beliefs held by any particular person.

The intentional states we attribute to the Internal Revenue Service show referential opacity and other features of the beliefs of persons. That agency might hold that the office expenses for my principal place of business are tax deductible (and state this in an official ruling), but deny that my home office is deductible, even though my principal place of business is elsewhere acknowledged to be my home office. One can resort to the courts to oblige the agency to act in accord with the logical consequences of its rulings, or to clarify the semantics of its terms.

The next step in adopting the intentional stance is to decide whether corporate entities are legitimate subjects of desire. Following Dennett, the main grounds for an affirmative answer would be that we do attribute desires to them, and this attribution serves explanatory functions. Consider the account of why France disagrees with the United States. It succeeds by conjoining a belief with a desire: a belief that the use of force might imperil its troops, and a desire to avoid such peril. Corporate desires are less volatile than those of a person, but they do have a finite life span. The attentive reader notices, for example, that Britain's willingness to consider the use of force has (according to the report) palpably changed.

The final step in adopting the intentional stance is to predict the behavior of the system by employing practical reasoning: from

³ See Margaret Gilbert, "Modelling Collective Belief," *Synthese*, 73 (1987): 185-204.

attributed beliefs and desires, infer the action it ought to perform, and predict that it will perform that action. Dennett nowhere requires that the agent actually indulge in the reasoning processes so imagined; but even if that were a requirement, corporate entities would not necessarily thereby be excluded from the ranks of intentional systems. Some organizations draw their conclusions and pick their actions through processes of group discussion and deliberation. If sufficiently unconstrained, those deliberations provide a paradigm case of practical reasoning, in which alternative viewpoints and policies are proposed and evaluated. The participants state their positions, provide reasons for their statements, and then engage in argument and deliberation with one another. "To reason" is in one sense "to talk with another so as to influence actions or opinions"; our exemplars for "reasons," "arguments," and "deliberations" are first and foremost the oral varieties we encounter in groups.

If a group *can* engage in processes of rational deliberation, then one can sometimes predict its result by simulating the reasoning undergone. It seems that governments can sometimes be treated as if they were rational agents.

III

To be an intentional system is, according to Dennett, to be "a system whose behavior is reliably and voluminously predictable via the intentional strategy" (*op. cit.*, p. 15.) The previous section showed that we can (and do) adopt the intentional strategy toward corporate entities. Now we need to show that this strategy is a successful one, yielding many reliable predictions not otherwise available.

What sort of success is required? Dennett describes arranging a rendezvous by telephone. After the conversation, one can predict with fair success that one's interlocutor will subsequently arrive within the designated volume of space-time (*op. cit.*, pp. 26-27). While this prediction would be very difficult to make from the physical or design stance, it is easy from the intentional stance: one takes the conversant's agreement as evidence that he or she *intends* to meet at the rendezvous point; and one knows that generally people do what they intend.

Similar sorts of success attach to adopting the intentional stance toward governments. Consider the movements that become readily predictable once one hears that the European Community is interdicting arms shipments to Bosnia. Here it seems we can predict the movements not just of one person through remote regions of space-time, but of ships and navies--thousands of people, tons of materiel--across scattered seas and tributaries. Folk psychology's successes seem microscopical by comparison.

Dennett requires that the intentional stance yield predictions that were not antecedently available or that cannot be produced using other methods. These predictions must be unavailable from the "design stance" or the "physical" stance: "often the only strategy that is at all practical is the intentional strategy; it gives us predictive power we can get by no other method" (*op. cit.*, p. 23). One way in which the intentional stance yields understanding available "by no other method" is that it allows one to understand

the *patterns* in human behavior that are describable from the intentional stance, and only from that stance, and that support generalizations and predictions (*op. cit.*, p. 25).

As examples Dennett mentions the different molecular patterns that all constitute placing a stock purchase order, and the different ways in which particles are displaced when one drives to a pre-arranged rendezvous. From the physical stance one cannot grasp that all those different motions are instances of placing an order, and it would be extremely difficult to produce an explanation in particle physics for a fact that is easily explicable in intentional terms: that the rendezvous ensues.

Logically analogous arguments apply to corporate activities. Some patterns in corporate behavior are accessible only if one adopts the intentional stance toward those entities. For example, consider the enormous complexity of movements constituting a naval blockade. We have ships, airplanes, mines, missiles, cannon, trucks, troops, and satellites all moving around, emitting signals, receiving data, sometimes shooting at one another. What explains the incidents of gunfire? They are relatively easy to understand when one understands that many of the objects moving around are attached to governments, that some of the governments are in conflict with one another, and that others are aligned with one or another side. We read the official rules of engagement published to govern naval operations, and then can explain why the big boats sometimes shoot at the little boats. Under certain conditions, specified in its rules of engagement, the Navy will fire on anything it believes to have hostile intent. Indeed, this pattern seems accessible *only* if we adopt the intentional stance toward nations. There is no way in particle physics to characterize "neutral shipping," "allied navy," or, for that matter, "Bosnian." These terms also elude all the other natural sciences. For tolerably nimble prediction of the patterns of motion on the high seas, it seems essential to understand that there are governments, that some governments are hostile toward one another while others are friendly or neutral, that certain motions satisfy conditions laid out in the rules of naval engagement, and that the Navy

will fire when those conditions obtain. We must treat governments as entities that can adopt a policy, and the various people and machines in the area as instruments carrying out those policies.

The moody bit players must be recognized before the soap opera gets underway. Seen from the perspective of particle physics, all the drama and excitement of such histories evaporate, leaving just chaos and random death.

If these explanations succeed, and manage to explain behavior not otherwise explicable, then for Dennett no further doubt is possible: governments would *be* intentional systems.

all there is to being a true believer is being a system whose behavior is reliably predictable via the intentional strategy, and hence *all there is* to really and truly believing that *p* (for any proposition *p*) is being an intentional system for which *p* occurs as a belief in the best (most predictive) interpretation (*op. cit.*, p. 29).

The conclusion seems inevitable: if that is all it takes to be a true believer, governments can join the ranks. Beliefs and desires are cited in explanations of corporate behavior with the same degree of success and legitimacy as they are in explanations of human behavior.

IV

We might feel surprise and chagrin at the prospect of treating a corporate entity as a person. But the conclusion can be seen from a different perspective. Recall the ancient idea that a person *is* a kind of corporate entity. The republic is modeled on the soul (not the other way round); Hobbes' Leviathan is an artificial man. That we treat corporations as persons is not so odd if in fact persons *are* corporate entities. Let us step through the looking glass, and examine things from that perspective. The Platonic analogy was repeated by Hume:

...I cannot compare the soul more properly to any thing than to a republic or commonwealth, in which the several members are united by the reciprocal ties of government and subordination, and give rise to other persons, who propagate the same republic in the incessant changes of its parts.⁴

Contemporary functionalism holds that this is no mere analogy: people (or their minds anyway) *are* corporate entities. A mind is a functionally characterized organization, and each part is described by its job or role in the overall system. The description is abstract and the properties it ascribes are purely relational. It says nothing about the stuff constituting the parts. The organization--the structural properties--can

survive a partial or even total change in its office-holders.

Much of the language used to describe mental organization derives from the corporate boardroom. Here is one of Dennett's classic descriptions:

If one can get a team or committee of *relatively* ignorant, narrow-minded, blind homunculi to produce the intelligent behavior of the whole, this is progress. A flow chart is typically the organizational chart of a committee of homunculi (investigators, librarians, accountants, executives); each box specifies a homunculus by prescribing a function *without saying how it is to be accomplished* (one says, in effect: put a little man in there to do the job). If we then look closer at the individual boxes we see that the function of each is accomplished by subdividing it via another flow chart into still smaller, more stupid homunculi.⁵

Lycan's "homuncular functionalism" or "homunctionalism" is even clearer in this respect:

The homuncular functionalist sees a human being or any other sentient creature as a kind of corporate entity--as an integrated system of intercommunicating "departments" that cooperatively go about the business of interpreting the stimuli that impinge on the corporate organism and of producing appropriate behavioral responses. In this model, a psychological description of a human being will consist of a set of flow charts, nested hierarchically. . . . The subsystems will be characterized in terms of their respective corporate responsibilities: "executive unit," "perceptual analyzer," "speech center," "buffer memory," and the like. Now, each of the subsystems will itself be described by a lower level flow chart that breaks *it* down into its own component departments or agencies (sub-subsystems of the containing person), which corporately produce *its* various outputs given inputs, and so on.⁶

Homunctionalism has empirical roots; the claim that a person is a corporate entity is an empirical claim. Indeed, it draws support from on-going research programs in cognitive science. Beneath the disagreements between (say) Fodor and the connectionists, there is at least one claim on which all the parties agree: the mind is a massively parallel system. It consists of layers of hundreds or thousands of "processing units," all working simultaneously and chattering at one another in parallel. Fodor takes them to be "modules," communicating in something like a digital machine language, while connectionists take

⁵ Daniel Dennett, "Toward a Cognitive Theory of Consciousness" in *Brainstorms* (Montgomery, VT: Bradford Books, 1978), pp. 123-4.

⁶ William G. Lycan, "Toward a Homuncular Theory of Believing," *Cognition and Brain Theory*, IV, 2 (1981): 139-59, p. 140. See also Ronald de Sousa, "Rational Homunculi," *The Identities of Persons*, A. Rorty, ed. (Berkeley: University of California Press, 1976), pp. 217-38.

⁴ David Hume, *A Treatise of Human Nature*, L. A. Selby-Bigge, ed. (Oxford: Oxford University Press, 1975), Bk. I, Pt. IV, Sect. VI (p. 261).

them to be "units," which represent and process information by changing their mutual connections; but both parties agree there are *many* of them, working raucously in parallel.⁷ Sample tasks of these processors: pick areas of high contrast, find edges in visual input, calculate retinal disparity, keep the eyes focused while compensating for head movements, pick out phonemes, and so on. Think of the mind, not as one central processor, but as a massively parallel layered network of thousands of them: a *society* of mind.⁸ How are we to map our psychological predicates onto states of such a system?

If a sentient creature is a kind of corporate entity, and it is legitimate to ascribe beliefs and desires to some sentient creatures, then it follows trivially that it is legitimate to ascribe beliefs and desires to some kinds of corporate entities. Not all corporations will have the appropriate organizational structure to sustain such ascriptions, but it follows from homunctionalism that the appropriate organizational structure *can* justify those ascriptions: *no more* is necessary.

This idea is very difficult to believe, or even to understand fully. We do not normally think of people as corporate entities. It is unsettling to suggest that even in solitude one cannot escape committee life. Reporting one's current mental state does not seem to be anything like reporting the current state of General Motors. (It does not help for the homunctionalist to say "Not *that* kind of corporation: *this* kind!") It is difficult to see how to make these ideas more intuitive.

Derek Parfit⁹ provides some assistance, with his analysis of the claim that "a person is like a nation." He argues for what he calls a *Reductionist* view of personal identity, which consists of three claims:

- (1) the fact of a person's identity over time just consists in the holding of certain more particular facts . . .
- (2) these facts can be described without either presupposing the identity of this person, or explicitly claiming that the experiences in this person's life are had by this person, or even explicitly claiming that this person exists. . . .

⁷ See Jerry A. Fodor, *The Modularity of Mind* (Cambridge: MIT, 1983); David Rumelhart and James McClelland, *Parallel Distributed Processing: Explorations in the Microstructure of Cognition* (Cambridge: MIT, 1986).

⁸ See Marvin Minsky, *The Society of Mind* (New York: Simon and Schuster, 1986).

⁹ *Reasons and Persons* (New York: Oxford University Press, 1986), p. 275. See also Peter Unger, "Conscious Beings in a Gradual World," *Midwest Studies in Philosophy*, XII, P. French, T. Uehling, H. Wettstein, eds. (Minneapolis: Minnesota UP, 1988).

- (3) Though persons exist, we could give a *complete* description of reality *without* claiming that persons exist. I call this the view *that a complete description could be impersonal*. (*ibid.*, pp. 210-12).

Parfit points out that most of us are reductionists about nations. The existence of a nation just consists in the holding of certain more particular facts. Those facts can be described without presupposing that nations exist. If an entity is in this way reducible, then one could describe all the particular facts, and give a full account of what there is, without settling the question of whether or not a given nation exists. The identity of the nation is in this way *indeterminate*; the question does not raise a distinction between two distinct states of affairs, but merely between two different ways of describing the same state of affairs. In these respects, Parfit says, "a person is like a nation" (*ibid.*, p. 275). I shall follow Parfit and call the view that a person is a corporate entity *reductionism*. Homunctionalism is a specific variant of reductionism, and it has the three consequences above. Note that we are talking here of *ontological* reduction, not explanatory reduction. It is a thesis about what people are--their constituents, how they are made up--not about our *explanations* of people.

One implication of reductionism about corporate entities is reductionism about their *states*. When we say that the French government feels concern for its troops, we seem to be ascribing a logically simple state to a particular. If reductionism is true, this appearance is misleading. The alleged fact that the French government feels concern for its troops *consists in* a great many more particular facts about people and their relations. It is not a logically simple one-place predicate ascribed to a particular; instead it can be analyzed without remainder in terms of other predicates and other particulars. The "state of the government" turns out to a collection of more particular states (and relations) obtaining among a collection of people.

If the mind is a kind of corporate entity, it would follow that mental states are similar to states of a corporation. I have a student who wants to retain a scholarship to help him graduate. If homunctionalism is true, the fact that he wants to retain the scholarship consists in more particular facts, which can be described without mentioning the desire, any other psychological facts, or even the existence of people. His wanting to retain the scholarship is like the Chinese government wanting to retain favored nation status. The facts are logically analogous.¹⁰ The subject of the desire is the corporate entity as a

¹⁰ Although, as far as I know, Dennett never explicitly endorses the ascription of intentional states to governments or private corporations, he does

whole. The state of desire consists in states and relations of various departments in that corporation; there is no distinct "department of desires" to serve as its subject. Similar reductionist views hold for beliefs and other psychological states. His believing that education will help his job prospects is logically analogous to the United States government believing that low interest rates will help the trade deficit. To say he is worried about his future is similar to saying that Hong Kong is worried about *its* future.

According to reductionism, psychological states are constituted by complicated patterns of relationships holding among sub-personal departments. They *consist in* a great many more particular facts that can be described in an impersonal way.

The reductionist views the mind as a corporate entity: a layered network of interacting processors, each with its own job. One anti-reductionist response to this picture, reiterated several times in the history of discussion of the subject, is to imagine yourself shrinking to the point where you could physically enter this assemblage, as if (with Leibniz) you were entering a mill. Or perhaps you shrink to the point where you could personally join the corporation, take on the job of one of those processors, and watch what happens. Viewed from that perspective, one is expected to find counter-intuitive the claim that the entire network of interacting parts--the vast assemblage--might be experiencing some sensation or thinking some thought. Ned Block's "Great Brain of China" argument--in which Chinese citizens each play the role of a neuron--and Searle's "Chinese room" argument--in which Searle takes on the job performed by a central processor in a computer-

sometimes play with the same analogy. For example, in *Consciousness Explained* (Boston: Little Brown, 1991), he describes the temporal "smear" of consciousness in the following terms:

Every event in your brain has a definite spatio-temporal location, but asking "Exactly when do you become conscious of the stimulus?" assumes that some one of these events is, or amounts to, your becoming conscious of the stimulus. This is like asking "Exactly when did the British Empire become informed of the truce in the War of 1812?" Sometime between December 24, 1814, and mid-January, 1815--that much is definite, but there simply is no fact of the matter if we try to pin it down to a day and an hour. Even if we can give precise times for the various moments at which various officials of the Empire become informed, no one of these moments can be singled out as the time the Empire itself was informed. (*ibid.*, p. 168-9)

The Empire (itself!) is here ascribed intentional states, and the mind likened to it. The same analogy recurs in Daniel Dennett and Marcel Kinsbourne, "Time and the Observer: The Where and When of Consciousness in the Brain," *Behavioral and Brain Sciences*, XV, 2 (1992): 183-248.

-can both be seen as variants of the same thought experiment. But suppose the Chinese *government* can appropriately be said to hold some views, and hope for certain outcomes. It may then not be so counter-intuitive to suggest to the individual Chinese citizens that there is a larger corporate entity, in which they participate, which has intentional states distinct from their own. If the Chinese government can look forward to absorbing Hong Kong, perhaps the Great Brain of China could too.

V

I do not propose to consider all the grounds on which one might agree or disagree with reductionism about people. That question has been discussed at great length in many places, and so in this place I propose to do something different. Let us suppose (for the sake of argument) that reductionism *is* true: people are corporate entities. What would follow of the fate and status of "folk psychology"?

Folk psychology is described as a distillate of the platitudinous generalizations and inference rules underlying our use of ordinary psychological terms in description and explanation. Those generalizations and inference rules characterize relations among hypothetical internal states named by ordinary psychological terms: beliefs, desires, intentions, emotions, and so on. Folk psychology is thought to be a proto-scientific *theory*; it is thought to describe law-like relations among those states, and how they jointly conspire to yield behavior. When we ascribe a particular thought, sensation, emotion, or belief to ourselves or others, we are invoking some part of this theory.

Consider the analogy with governments as intentional systems. Governments not only levy taxes and go to war. They also make official statements, using official spokespersons. The official spokesperson makes statements about what the government believes, hopes, and intends; and he or she describes the official interpretation of the behavior and viewpoint of other governments. Spokespersons use a rich set of intentional idioms in making their official statements. This is the public relations branch of the organization, and so following Dennett¹¹ I shall call the spokesperson *PR*. Just as with the folk psychology of persons, we could collect the generalizations and inference rules upon which official descriptions and explanations rely, and form a kind of folk psychology of *governments*. Would this be a

¹¹ Dennett, *Brainstorms*, p. 156. *PR* was introduced (in 1975) as a "Ron Nessen analogue," and is a homunculus, so the masculine pronoun seems appropriate.

proto-scientific theory of the workings of a government? The difficulty is just that *PR* does not have--and does not need--access to the information that would be required to formulate a theory of how the government works. *PR* communicates with just a few offices within a few subdepartments of the vast bureaucracy of the federal government. He does not see much of how the government works; he just takes orders from some higher ranking homunculi, gives orders to lower ranking ones, and communicates with those few to which he has official access. *PR* is simply in no position to learn how the government as a whole really works. He can later publicize the gossip from a few offices, and while his memoirs might command a hefty advance they certainly do not contain a theory--or even a proto-theory--of how the government works.

When we move from governments to citizens we find a homuncular analogue for *PR*: whatever subagency of the mind is responsible for the formulation of speech contents. If the mind consisted of one Central Processor, it might make sense to suppose that all the operations of mind were self-intimating, accessible, and known by their bearer in some privileged way. The Central Processor must have special access to the goings-on within its registers. But suppose the sub-personal *PR* is not *the* Central Processor, but just one member in a vast society of cooperating agents. His access is finite; his field of action is limited. Much goes on in those other offices of which *PR* is not--and could not become--aware. To suppose that the agency responsible for the formulation of speech contents is privy to the internal operations of those other offices is to lapse again into the view that there is one Central Processor. But that view seems empirically discredited.

On homuncionalist principles a sub-personal *PR* faces the same problems as does our official spokesperson. *PR* has limited contacts with just a few parts of the bureaucracy, restricted sources of information, and his own proprietary modes of operation and communication. How *could* he know how vision works? Why should he care to know how memory is organized, or how to coordinate the thousands of muscle twitches needed to effect a tennis serve? *PR*'s job is public relations: getting along with other people. He does not have the access or the capacity to determine how other parts of the organization do their jobs. In many cases he does not even speak their language.¹²

¹² There is no *a priori* reason to believe that all the processing units making up the mind use systems of representation that can be translated into a natural language, or that all their states have contents readily expressible in English sentences.

If reductionism is true, *PR*'s access to the operations of mind is finite, and much of what goes on must occur beyond his ken. He receives inputs from a subset of offices, but is not even privy to how those inputs were produced. It follows that the sentences forthcoming from *PR* are unlikely to yield a viable explanatory theory of mental operations.¹³ They are at best memoirs from one participant--a partial view, provided by one of the agents involved. The truth of reductionism impugns the explanatory status of folk psychology.

It has often been noted how folk psychology is oddly unconcerned about the questions exercising the experimental discipline: How does vision work? How is memory organized? Why do we dream? What changes inside when we learn things? Folk psychology seems never to address the "how does it work?" kind of question; it spends its energies instead on the "what will he/she do?" kind of question.¹⁴ It does not explain psychological laws; instead it provides commentaries on particular behaviors. Furthermore, the "explanations" seem never falsifiable: one never finds observations posing a serious threat of disconfirming folk generalizations. Instead we inevitably make post hoc modifications in the ascribed antecedents. It is debatable to what extent folk psychology is genuinely *predictive*, as so much of it is only generated after the fact, and merely gives a retrospective gloss on behavior that has already taken place.

I suggest all these facts are easily accommodated by the theory that folk psychology is the product of a sub-agency--*PR*, the one involved with public relations--within a corporate mind, a society of thousands of agents. *PR* produces sentences meant for interaction with conspecifics; folk psychology is a distillate of its memoirs. Some subagency within the bureaucracy does indeed speak for the whole; but its primary purpose is *not* to describe the internal states of the corporate entity.

Our attributions of folk-psychological terms are like official statements by spokespersons in the State Department. Saying "I am

¹³ This claim is not based on an argument from analogy, but is rather drawn as an implication of homuncionalist, or, more broadly, "reductionist" theories of mind. These theories make it hard to see how a subpersonal *PR* could be in any position to formulate a viable explanatory theory of mind. It is unlikely--possible, but unlikely--that study of *PR*'s linguistic habits will reveal much about how the organization works. This argument is inductive. It is bolstered if one can provide an alternative account of *PR*'s activities; see below.

¹⁴ Paul Churchland, *Matter and Consciousness*, revised edition (Cambridge: MIT, 1988), pp. 45-7; Kathleen Wilkes, "Functionalism, Psychology, and the Philosophy of Mind," *Philosophical Topics*, XII, 1 (1981): 147-67; Kathleen Wilkes, *Physicalism* (Atlantic Highlands, NJ: Humanities Press, 1978).

disappointed" is logically analogous to a spokesperson of the U.S. State Department saying "The United States is disappointed." Folk psychology helps predict behavior in the same way that statements by official spokespersons help predict government actions. It will relate to scientific explanations of human behavior in somewhat the way those official statements relate to explanations of government actions.

Official policy statements are not descriptions of inner causes; they are public commitments to a certain line of action. They do not describe internal states of affairs that later cause certain actions; instead they are meant to provide a *rationale* for official actions. A statement of policy serves to justify and rationalize prior actions and, more importantly, signal future ones. If we collect all the platitudes from official spokespersons we would get, not a descriptive theory of the inner workings of the government, but instead a kind of protocol manual for interpreting the meaning of official statements.

Commonsense "explanations" of the actions of individuals share these features. They relate actions to earlier thoughts and feelings of the person so that we can understand the meaning and implications of the action. They thereby serve to assimilate the action to some policy-line that we can readily comprehend and use to predict further consequences. Folk explanations are more like official policy statements than descriptions of causal antecedents. They more often provide an official rationale for a particular behavior than a theoretical explanation for a psychological law. In brief: folk psychology is not *psychology*. It is more akin to diplomacy.

Official policy statements do indeed help one to *predict* the behavior of a government. The statement of an official rationale for current behavior reduces uncertainty about future behaviors; it conveys useful information. When the U.S. announces its intent to blockade the latest miscreant nation, the behavior of thousands of individuals is suddenly predictable. The predictions show the same fecundity as folk predictions: we can, for example, anticipate some of the reactions of neighboring governments, allies, and enemies.

For all this predictive success there is something troubling about treating official statements as *explanations*. It does not seem an adequate explanation of the eventual episodes of gunfire to say that the U.S. government wanted to achieve certain ends. We do not think of the desire as a state of the government, interacting causally with other states, producing the official action. If it were, one would counter the explanation by describing other, countervailing desires of the U.S. government; or perhaps by ascribing it a different set of *beliefs*; or finally by showing some failure in the way the government *reasons*. This would make political science into a kind of depth psychology--one

which details relations among the psychological states of corporate entities. That is clearly not how we go about it when we want an explanation of corporate behavior.

How do we go about it? We leave aside the vocabulary of official statements and descend to the level of interactions among the *people* making up the government. We identify the main factions having power, the key leaders, how they interact with one another, and how they exert control over the governmental apparatus. The explanation ultimately refers to what some key individuals thought, what they were trying to do, and how their actions affected the groups and factions so as to determine the government's behavior. The story gets complicated because of conflict within the government; one needs to explain why this group rather than that group ultimately exerted the control it did over government assets and policy. We are indeed reductionists--this time in the sense of *explanatory* reductionists--about governments. All its properties and behaviors can be explained in terms of the different properties and relations among its constituent parts.

One way to put this is that when we want an explanation of government actions we ignore the official spokespersons, and turn to investigative reporters--the sort who talk to all the key players in the State Department, the Pentagon, the Executive Branch, and the Congress; and who know and can tell us what all those people are thinking and doing. To understand some muddled response to a foreign policy crisis, we need the investigative view. It will not employ the idiom of official statements. No mention will be made of what the United States desires. Instead we want a summary of what the main players are thinking and doing, and of the struggles among them.

Indeed, part of the investigative task is to explain the contents of the official statement. Such statements are explained by referring to the agreements among the individuals who controlled its production. We explain the correlation between the content of the official statement and the subsequent government action by showing how both were controlled by the same group of individuals within the government. This is, after all, what we assume of our own spokespersons. We hope that there is a reliable *causal* link between the formulation of statements to be read by spokespersons in the State Department and the actions of other individuals who carry out U.S. foreign policy. When the link fails, we get a "privatized" foreign policy--which is no longer an action of the government at all.

Perhaps folk psychology can help us understand the rationale for a particular action--and this certainly reduces uncertainty and yields predictions--but, just like an official policy statement, it does not explain *why* the action took place. When we ask why, we want

something more; something based on knowledge of what the key players in the corporation were up to. We want answers to the hard technical questions that *PR* would rather avoid. We want access to the archive of past internal states; the latest official historical narrative or policy statement will not do. Perhaps folk psychology will bear a similar relation to scientific explanations of human behavior as official statements of government policy bear to scientific explanations of government behavior. If predictions based on such statements are subsequently vindicated, we need not assume that the public relations department provided us with true descriptions of the internal state of the corporation. Instead the predictions were vindicated because the same forces within the corporation that produced the statement also eventually produced the policy. The intentional stance can be predictively successful but explanatorily vacuous.¹⁵

VI

This conjunction of predictive success with explanatory failure may seem implausible, at least until one considers the question: *Why* do corporate entities make official statements? How did this mechanism evolve, and what is it for?

I think the answers to these questions are fairly obvious. If governments did not make official statements to one another, their behaviors would be far less predictable; in fact, they would be so *unpredictable* that any kind of trade or interchange between nations would be almost impossible. Think of steaming up the Danube if no governments in that region had issued policy statements, if no rules of engagement had been published, if there were no international conventions (or laws) governing shipping, and so on. Without such official statements one would experience even greater uncertainties than pilots do currently. Intercourse between nations would not be possible.

A key function of official statements is to *make* behavior predictable. They reduce uncertainty about future contingencies of response. Once one knows the Navy's rules of engagement, one has a much better sense of which of one's behaviors will draw fire and which not. Without some such ability to predict the other's behavior, it would be close to impossible to interact. The interchange of goods, services, people, and ideas would be fraught with uncertainties and peril.

¹⁵ This position is not committed to the view that folk psychology is somehow immune to the threat of empirical disconfirmation. On the contrary, official statements can be false; they do presuppose and imply various empirical regularities which may or may not obtain. The point is just that *explanations* of corporate behavior will not employ the language of such statements.

This is true as well at the personal and social level: without some ability to predict the other person's response, social life as we know it would be impossible. Consider the paucity of interaction possible between two people who do not share a language: that would be the *upper* limit on social interaction if we could not make our own "official statements" to one another. The need to predict the other's behavior is ubiquitous. Non-human animals share it. Even if most do not possess a language, they certainly possess the means to communicate: intricate systems of gestures and expressions that have evolved to convey information and to reduce one another's uncertainties about future contingencies of response. Observe the interactions, for example, when a new dog meets the current canine residents of the neighborhood. One sees an intricate pattern of stereotyped gestures and expressions that signal aggressiveness, anger, fear, submission; and a natural sequencing of such signals that manages, mostly peacefully, to negotiate new territories and relative status. Even in those encounters in which bodily harm seems likely, the interchange of increasingly aggressive signals is delicately ritualized, so that one or another can say "uncle" at any point and terminate the proceedings. Our pets manage to predict one another's behavior with remarkable accuracy. The rate of killings among the canine residents of a U.S. city may well be lower than that among the humans.

It is not a new idea to suggest that the expression of emotions in man and animals is a system of stereotyped signs and gestures evolved to reduce uncertainties about future responses and make our behaviors more predictable.¹⁶ This system makes social life possible. Much of folk psychology can be seen as a linguistic means to the same end. It is a more sophisticated system for making our behaviors predictable by others; an extension of "natural psychology" made possible by language.¹⁷ The human ability to *talk* about anger is a cultured variant of the canine's raised fur and bared teeth: we can signal likely future behaviors in a more precise and more discreet fashion, but the underlying function is the same. The folk psychology of anger is a

¹⁶ Charles Darwin, *The Expression of the Emotions in Man and Animals* (Chicago: University of Chicago Press, 1965).

¹⁷ Nicholas Humphrey introduced the term "natural psychology" to describe the systems evolved in social animals to enable them to predict one another's behavior. He argues that animals that are not "natural psychologists" could not cooperate socially or even interact significantly. It is a small step to suggest that folk psychology is just a sophisticated variant of such "natural" psychologies. See Nicholas Humphrey, *Consciousness Regained* (Oxford: Oxford University Press, 1984), and *The Inner Eye* (London: Faber and Faber, 1986).

user's guide to this signal system. It is not a descriptive theory of the makeup of emotions so much as a phrasebook for their ready interpretation.

A reductionist holds that our attributions of ordinary psychological terms are similar to official statements from spokespersons in the State Department. Such statements have empirical consequences and can licence predictions. We have official spokespersons so that governments can better predict one another's actions. The policy-line provides a rationale for behavior. More importantly, it helps one to predict future actions, and so allows social life to proceed with fewer fatalities. Ordinary psychological terms evolved from the similar roots and serve similar functions. Folk psychology is an interpretation manual for such signal systems. It too provides a running commentary, a story line that is of some predictive value. But when we seek explanations of why the system acted as it did, we inevitably endorse explanatory reduction, and seek the inner constituents that jointly conspired to generate the result. *PR* gives us just a partial view--the official gloss--and there is no guarantee that the picture yielded by subsequent investigations will match, or even be recognizable.

VII

One typical response to this entire line of argument merits discussion. Earlier I argued that all of the conditions Dennett specifies as sufficient for ascribing beliefs and desires are satisfied by corporate entities such as the State Department or General Motors. I mentioned that some will take this as a *reductio ad absurdum* of Dennett's view. That is to say, the response will be:

- (1) If the intentional strategy describes sufficient conditions for attributing beliefs and desires, then corporate entities, such as the State Department or General Motors, can legitimately be attributed beliefs and desires.
- (2) But it is ridiculous to attribute beliefs or desires to corporate entities, such as the State Department or General Motors.

-
- (3) Hence the intentional strategy does not describe sufficient conditions for legitimate attribution of beliefs and desires.

I believe that my argument has established the truth of premise (1). The dispute centers on the truth of premise (2). The typical response might be expressed as follows:

Although we talk as if the State Department and General Motors have convictions, plans, and desires, such talk cannot be taken literally. It is a simplifying--and ultimately misleading--analogy, employed by reporters so as to make the news comprehensible to the average fifth grade reader. Such attributions are not literally true.

That is, strictly speaking, all attributions of intentional states to corporate entities are *false*. Indeed, for the *reductio* to work, such claims must be considered to be *quite obviously* false.

This response has some interesting and disturbing implications. Everyday we read sentences in the *New York Times* like the ones in our earlier example:

The United States has concluded that Serbia is routinely violating the United Nations ban on flights over Bosnia . . . The British Government today gave its strongest public indication that it might be willing to support some degree of enforcement . . . both France and Britain . . . expressed concern for the safety of their troops. . . . (*op. cit.*)

According to the *reductio* proponent, all of these sentences are, strictly speaking, false. Governments neither conclude nor fail to conclude anything; they are neither willing nor unwilling to support policies; and they neither express concern nor fail to express concern. Ascription of such intentional properties to a corporate entity is on this view not only false, but ridiculous; e.g., *obviously* false.

The problem ramifies. Much of our discourse about governments, corporations, and other groups is conducted in this idiom. The practice is ubiquitous. It is difficult to imagine how we could proceed without such a convenient shorthand. The suggestion that all of those sentences found in newspapers and broadcast over the airwaves are literally false is apt to be met with incredulity. Do you mean to say that the *New York Times* and *CBS News* routinely disseminate falsehoods? that all of this discourse is systematically deluded?

Consider sentence *F*: "The French government feels concern over the welfare of its troops." What does it mean to say that *F* and sentences like *F* are, *strictly speaking*, false? (or, to put it another way, that sentence *F* is not *literally* true?) Both claims are ambiguous, and two interpretations are possible. One interpretation is that "literally" and "strictly speaking" function as emphatic adverbs, signalling that sentence *F* is *just plain false*. *F* means that some entity (the French government) has some intentional property (a feeling of concern). Since there is no such entity, and governments cannot be subjects of intentional states, sentence *F* is false. The state of affairs it connotes does not obtain.

But a second interpretation is possible. That *F* is not literally true might mean that *F* is false if taken literally--that is, if one takes it to

mean what its surface structure suggests. But, the second interpretation insists, that is not the correct analysis of *F*. Sentence *F* appears to ascribe some intentional state to a particular--the French government--but this appearance is illusory, and the sentence should be analyzed in a different way. On this line, what sentence *F* actually means is something like:

Such and such individuals occupy the following key offices in the French government. They control policy in such and such a manner. These individuals have the following beliefs and desires; and they have interacted in such a way as to produce an official statement which means the same as "The French government feels concern over the welfare of its troops."

Note that this is a reductionist *analysis* of the meaning of sentence *F*. On this line, *F* appears to ascribe a single property to a single bearer; but this appearance is illusory. It should instead be analyzed as describing a great many more particular facts, each of which can be described without presupposing the existence of an entity identified as "the French government." On this interpretation, sentence *F* turns out to be true, and its truth consists in the truth of these many particular claims about particular people and their thoughts and interactions.

A proponent of premise (2) must plunk for the first interpretation and avoid the second. If the *reductio* is to go through, we cannot allow sentence *F* merely to *appear* to be false, or to be false if taken literally (but not otherwise). *F* must *be* false. The position requires that *F* have some clear analysis, which, when applied to governments and other corporate entities, can be shown obviously not to obtain.

This position then puts the ascription of intentional states to people at empirical risk. The view is that *F* means *P*, and *P* is not true of a corporate entity. But now suppose that people are corporate entities. One could imagine discovering that reductionism is the truth about people. It is a logically possible doctrine. If it happens empirically to be the case--as seems likely--it would follow on this line that our ascriptions of intentional states to people are at risk as well. If no corporate entity can "feel concern," then the discovery that people are corporate entities would cast doubt on the ascriptions of such states to people.

Here is a second way to see the connection. Suppose you deny that an appropriately organized corporate entity can sustain intentional states. You subscribe to premise (2). But you also think that people do have intentional states. It follows that you must deny reductionism about people. Our ascription of intentional states to people cannot be viewed as ascription of such states to corporate entities. If reductionism about people turns out to be the truth, you must either retract premise

(2), and allow that appropriately organized corporate entities can sustain intentional states, or you must retract the ascription of intentional states to people.

The final irony is that anti-reductionists here find themselves joining ranks with eliminative materialists. Eliminativists bite the bullet, and deny that *people* sustain intentional states: All such ascriptions are, strictly speaking, false. Just as in the case of governments, that the ascriptions are ubiquitous and useful is no bar to this possibility. It is often claimed that the eliminativist position is somehow self-defeating or incoherent, or that the collection of sentences constituting folk psychology could not be radically false in the way eliminativists suggest. Yet if we turn eliminativist toward our ascription of intentional states to corporate entities, we face arguments of exactly similar form. They are in both cases invalid. One can coherently hew to the eliminativist line concerning intentional states of corporate entities; and the success in that domain demonstrates the feasibility of applying the same line to people. If you want a model for the coherence of eliminative materialism, consider endorsing eliminativism about the intentional states of governments.

I argued that according to the terms of Dennett's analysis it is equally legitimate to ascribe intentional states to governments and other corporate entities as to people. The eliminativist agrees, and says yes, both cases should be assigned the same truth value, namely, the False. The identity theorist--who accepts something like our reductionist analysis of intentional ascriptions--agrees too, and says yes, both cases should be assigned the same truth value, namely, the True. I see no room for a third option.¹⁸

Austen Clark
Department of Philosophy
University of Connecticut
Storrs, CT 06269-2054

¹⁸ For further arguments to this effect see Jaegwon Kim, "The Myth of Nonreductive Materialism," Presidential address, Central Division, *Proceedings and Addresses of the American Philosophical Association*, LXIII, 3 (1989): 31-48.