Organizations as True Believers

Deborah Tollefsen

In everyday discourse and in legal and moral contexts we often attribute beliefs, intentions, and desires to organizations like corporations.[1] Pick up any newspaper and you are likely to find passages like the following:

> Microsoft believes that its status as an accused monopolist entitles it to use the judicial process to delve into its competitors' most sensitive commercial information. (*USA Today*, May 9, 1999)

> The lawsuit filed by Washington Attorney General Christine Gregoire accuses the tobacco industry of conspiring to violate antitrust and consumer-protection laws, restraining and intentionally suppressing health research, manipulating cigarette nicotine levels without disclosure and threatening retaliation against companies that marketed safer cigarettes. (Associated Press, June 9, 1999)

> U.S. steel companies, alarmed by a flood of imports, are poised to file lawsuits Wednesday charging their foreign rivals with dumping metal in the USA at cutthroat prices. . . . In response to U.S. criticism, Japanese steel companies already have promised to cut back on their exports to the USA. (Rich Miller, *USA Today*, June 9, 1999)

> Ethiopia maintains its objectives are to destroy the Eritrean army, secure border territory it claims as its own, then withdraw. . . . Ethiopia has overwhelmed and humiliated its neighbor with surprising ease. (Associated Press, May 23, 2000)

The intentional idiom is applied lavishly in these passages. Corporations make promises, threaten competitors, conspire to break the law, are alarmed, and believe propositions. Governments communicate their objections and are subject to psychological states like humiliation. These attributions are made in the context of explaining and predicting the actions of organizations.[2] For instance, based on the fact that Japan has promised to cut back on steel exports we can predict with some accuracy that there will be a decrease in Japanese steel exports.

Faced with the fact that we do ascribe intentional states to organizations, there are several alternatives open to us:

1. We may take the hard line and claim that our attributions of intentional states to organizations are merely metaphors and although useful are, strictly speaking, false.[3]

JOURNAL of SOCIAL PHILOSOPHY, Vol. 33 No. 3, Fall 2002, 395–410.

2. Second, we may say that while such ascriptions are neither mistaken nor metaphorical, what we really mean when we say that organization *G* believes that *p* is that all or some of the members of *G* have the belief that *p* or they have certain other intentional states.
3. We may decide that corporations really do have intentional states.

In section I of this paper I argue that options 1 and 2 are implausible. In sections II and III, I defend option 3 by arguing that, given a certain plausible account of intentionality, organizations *really do* have intentional states. In section IV, I consider and reply to some criticisms of the thesis I defend.

I

The view that our ascriptions of intentional states to organizations are mere metaphors is a common view. But this way of making sense of our attributions can be ruled out on the following grounds. Our practice of attributing responsibility to organizations (consider, for instance, current tobacco lawsuits) seems to presuppose that organizations *literally* have intentional states. For we could not hold them legally and morally responsible for an act unless they *intended* to commit the act.[4] Since we do not hold organizations metaphorically responsible (much to the dismay of tobacco companies), the attributions on which our ascriptions of responsibility rest should be, at least initially, considered nonmetaphorical. Further, although false ascriptions could be explanatorily powerful (just as false theories are sometimes explanatorily powerful), explanatory power is prima facie evidence that our ascriptions are not simply false. We might also note that if option 1 is correct and our ascriptions of intentional states to organizations are simply false, then we, the media, social scientists, lawyers, political scientists, etc., are continually disseminating lies. This seems like an odd result and again, prima facie, evidence that our ascriptions are not mere metaphors. Given these considerations, option 1 seems implausible.

Option 2 proposes that our ascriptions be analyzed in terms of the intentional states of individuals in the organization. Raimo Tuomela (1993, 1995) and Margaret Gilbert (1988, 1996) have each developed an account of group intentionality in this manner. Tuomela argues that in order for an organization to believe that *p* all the operative members (those responsible for decision making and policy) must accept *p* as the view of the group (Tuomela 1995, 314–16). According to Gilbert, in order for a group to believe that *p* the members must form a plural subject of belief. Members form a plural subject of belief when they form a joint commitment to believing as a body that *p*. The formation of a joint commitment occurs when each member expresses his readiness to be so committed with others under conditions of common knowledge (Gilbert 1996, 352–55).

Though there is much about these accounts that is illuminating, providing a conceptual analysis of our collective belief statements is not a promising approach. The failure of conceptual analysis in other areas of philosophy should justify such skepticism. But more importantly, these analyses overlook the fact that our attributions of intentional states to organizations are often,

if not always, made in ignorance of the intentional states of the members (even the operative members). The point is analogous to one made by Daniel Dennett concerning our practice of ascribing intentional states to individuals. Dennett writes, "we all use folk psychology knowing next to nothing about what actually happens inside people's skull's . . . our capacity to use folk psychology is quite unaffected by ignorance about brain processes—or even by large-scale misinformation about brain processes" (Dennett 1987, 48). Analogously, our attributions of intentionality to organizations are made with only a vague idea of the inner processes of the organization and often without any information about the intentional states of the members. Yet our explanations of the actions of organizations in terms of their beliefs, intentions, and desires are successful. This suggests that meaning analysis is not the way to approach group intentionality, for it does not adequately explain why our attributions of intentionality to organizations work so well to predict and explain the actions of organizations in the absence of information about the intentional states of their members.

Given the implausibility of options 1 and 2 we are left with option 3: organizations *really do* have intentional states. Such a position is likely to meet with staunch resistance, especially from those who think beliefs and intentions are states of the brain or mind. Organizations don't have brains or minds, so how could they have intentional states? But the possibility of group intentionality is not so easily dismissed. Whether something has a mind depends on whether it has intentional states and not the other way around.[5] Further, most philosophers of mind now reject any crude type-identity theory, so the fact that organizations lack grey matter is not evidence that they lack intentional states. Indeed, some might argue that the favored account of mind, functionalism, supports the view that organizations have intentional states. According to contemporary functionalism, the mind is a corporate entity. It is characterized in terms of its functional organization. Each part is defined in terms of the role it plays in the overall system. These roles can be multiply realized, just as the roles in an organization can be filled by multiple individuals. Given this way of characterizing the mind, one could argue that organizations not only have intentional states but are collective minds. But I will not argue in this way.[6] Instead, I suggest that we look to interpretationism for an account of group intentionality. This will provide us with an explanation of why our practice of ascribing intentional states to organizations is so successful and will allow us to be realists (of a sort) about collective intentionality.

II

Interpretationism is the view that if an agent is interpretable, the agent is an *intentional* agent. This view has been developed and defended by, most notably, Donald Davidson and Daniel Dennett.[7] It is an approach to intentionality that starts not with metaphysical speculations about the nature of the mental, but with our practice of attributing intentional states. It asks, what are the constitutive features of our practice that guarantee its explanatory power? That is, what assumptions do we need to make about an agent in

order to interpret her behavior successfully? Approaching intentionality from this third-person perspective allows us to avoid having to speculate about what beliefs *are* or the conditions under which one can be said to believe a certain proposition. If interpretation is successful, then the assumptions we make about the agent are justified. Revealing the constitutive features of our practice, then, reveals the nature of the mental.

When we attribute intentional states to agents we do so in order to explain and predict not simply an action or an utterance, but the agent. The subject of intentional state attribution is, then, the whole system or person. When we say of Susan that she believes it will rain and desires to stay dry, this explains why *Susan* carries an umbrella to work. I want to know why *she* acted in this manner. So the primary interest is in explaining Susan and not simply explaining her bodily movements or what happened in her brain. The point is made more forcefully by Jennifer Hornsby: "An action explanation is not a reply to a question about why some event occurred, and, in revealing what an agent thought and what she wanted, it does not introduce any singular term for the cause. . . . Since its focus is how things were with her, it is no wonder that no 'purely causal' statement can be extracted from the explanation. The objective is to see a causally complex whole—a person—in a certain, intelligible light" (Hornsby 1997, 93). Prediction and explanation of an agent's behavior involves making sense of them as unified subjects.

When we provide an explanation of an agent's behavior, we provide reasons for her behavior. Not just any reasons will do. The reasons must justify the behavior. As Dennett points out, "explanation of actions citing beliefs and desires normally not only describes the provenance of these actions, but at the same time defends them as reasonable under the circumstances" (Dennett 1987, 48). We are not looking simply for reasons why someone would behave in a particular way (any number of reasons may justify the behavior) but the reasons *for which* she acted the way that she did.[8] Again, we are interested in *her* reasons, not simply reasons per se. This, in turn, requires that when we attribute beliefs, intentions, and desires to an agent in order to provide her with reasons for her behavior these must be ones that can be seen as intelligible from the agent's perspective as well as our own. As William Taschek notes, "We want it to make sense to us that she acted as she did for the reasons she did. Reflection on these cases suggests that an action (including an utterance) will be intelligible to us—in the relevant sense—only if it makes sense *to us* that it made sense *to the agent* that she did what she did" (Taschek 1988, 3). This reflexive constraint guides us in our interpretative endeavors.

In order to make sense of others and hence make their behavior (both linguistic and nonlinguistic) intelligible to us in the appropriate way, we must assume that the agent shares our norms of rationality, that is, that the agent is rational. If we did not make this assumption the agent's behavior would remain unintelligible to us and interpretation would be undermined. We can understand the thoughts and words of another only if we suppose that his beliefs are incorporated in a pattern that is similar to the pattern of our own beliefs, which is to say, a pattern essentially governed by the norms of consistency and truth.[9] Therefore, the constraints we recognize as applying to our own beliefs, constraints that we could loosely call epistemological, are to be

used as constraints on possible interpretations; do not act contrary to your best judgment, draw inductive inferences on the basis of all available, relevant evidence, believe only things you take to be true, and don't believe inconsistencies.

If our explanation of the agent in terms of beliefs and desires makes her behavior intelligible to us only if it can be seen by us as intelligible from the agent's perspective, then our practice of interpretation involves the positing of an alternative point of view, a point of view from which the intelligibility of one's own behavior can be assessed. I am introducing here what many have called a *first-person* point of view. For Locke, a first-person point of view was a phenomenological point of view—a unified consciousness. Contrary to Locke, we need not associate this point of view with consciousness. Nor need we associate it with a soul or brain. A rational point of view is not a point located somewhere inside someone's head or soul. It is a perspective, and one that can be adopted by other agents. When we make sense of others we often project ourselves into their rational point of view in order to be able to better provide reasons that are intelligible in the way the reflexive constraint requires.

A rational point of view is, then, a conceptual point rather than a metaphysical point. It is a perspective from which one can assess one's cognitive life. If we did not assume that the agent had such a perspective the rationality constraint and the associated reflexive constraint would have no force. Agents could not be assumed to be governed by our norms of rationality if we did not also assume that they had a rational point of view from which they could assess their beliefs and intentions for consistency, truth, and intelligibility.

In *The Bounds of Agency* (1997) Carol Rovane suggests a similar account of the notion of a rational point of view. The rational point of view, according to Rovane, need not be tied to a particular soul or animal or a particular phenomenological viewpoint. A rational point of view is a view from which a person deliberates. According to Rovane, we engage in the following sorts of activities from the rational point of view: accept the implications of one's attitudes, resolve contradictions and conflicts, rank preferences, assess opportunities for action, determine means for arriving at ends, and consider the consequences of certain actions and attitudes. The rational point of view is essentially a normative notion. It is something in which contradictions and conflicts *ought* to be resolved, in which preferences *ought* to be weighed, etc. More generally, it is something from which all-things-considered judgments ought to be reached and implemented.[10]

When we attribute intentional states to others we do so on the assumption that they are agents with a rational point of view—a point of view governed by the same sorts of normative constraints we conceive ourselves as governed by and from which certain cognitive activities originate. It is from this rational point of view that one can assess the truth and consistency of beliefs, resolve contradictions between other intentional states, determine means to ends, etc. Thus, it is from this point of view that reasons can be seen as intelligible. Without such a point of view presupposed it would be impossible to make sense of others in the way that we do.

Once the assumption of rationality (with its associated assumption of a rational point of view) is in place we attribute to the agent intentional states that a rational agent *ought* to have. We begin by attributing what beliefs the agent *ought* to have given its environment and function in the world. Then we figure out what desires it *ought* to have given the same considerations. These attributions are not made in isolation but holistically. The assumption of rationality involves the assumption that the agent shares with us a dense pattern of belief and thought, and the contents of thought are, in our everyday practice, attributed on the basis of these other beliefs and their content.[11] Once we have attributed to the agent the beliefs and desires a rational agent *ought* to have, we can predict what a rational agent with those beliefs and desires ought to do or appeal to these beliefs and intentions in order to explain the agent's behavior.[12]

The fact that our practice of interpreting others is so successful is evidence that our assumption of rationality is justified. Our interpretative practice works because agents are rational to some extent and designed by evolution to be so.[13] If this is so then we need not look for deeper metaphysical facts. The nature of intentionality is revealed in our practice of making sense of others. If the assumption of rationality is justified then we are dealing with an agent who has a pattern of belief and thought that is similar to our own and that is governed by the same norms of intelligibility. If our best efforts to make sense of an individual fail, then there is no reason to believe that we are dealing with a rational agent and hence, there is no reason to believe we are dealing with an intentional agent. According to interpretationism, then, "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" (Dennett 1981, 68).

III

This brief account of interpretationism will have to suffice. In what follows, I want to suggest that the interpretationist framework can be extended to organizations.[14] Now, according to some interpretationists[15] even thermostats are believers of a rudimentary sort.[16] My attempt to extend interpretationism to organizations may seem, then, to be rather trivial. But it is not at all obvious how this extension is possible. Interpretationism claims that a constitutive feature of our practice of making sense of individuals is the assumption of rationality and this, as we have seen, involves the positing of a rational point of view. But when we assume an organization is rational aren't we simply assuming that all or some of its members are rational? And given that organizations are made up of individuals with their own rational points of view, what sense is there to be made of the rational point of view of an organization? Organizational theory, in particular rational system theories, can shed some light on these questions.[17]

Rational system theorists view organizations as instruments designed to attain specified goals. This approach emphasizes that the behavior of individual agents within an organizational context must be governed by specific limits. Within an organization the number of alternatives for action is great

and the information needed to evaluate consequences too vast for an individual to handle on his own. Even if every agent in the organization is rational this does not guarantee that they all will act rationally together to attain the goals of the corporation.[18] The distribution of labor within an organization reduces the responsibility of individuals, and the formalization of rules, authority structures, and organizational roles constrains choices and makes available to the agent the information necessary to perform in her role. By providing integrated subgoals, stable expectations, required information, necessary facilities, routine performance programs and in general a set of constraints within which required decisions can be made, organizations direct individual participants. It is the formalization of rule and role and the specificity of goals and subgoals that creates an environment in which agents can act individually and jointly to achieve organizational goals.

The structure of the organization provides a way of synthesizing the disparate perspectives of individuals into a unified perspective from which goals and subgoals can be set and achieved. This rational point of view is expressed in organizational policy, in the structure of the organization, and in the history of its decision making. Recall that a rational point of view is a perspective from which deliberation takes place. When individuals deliberate in an organizational setting, they adopt the rational point of view of the organization. It is from the point of view of the organization, rather than their own personal point of view, that deliberation ought to take place. In fact, one explanation of the failure of individuals to act jointly and the failure of group decision making is that individuals are deliberating from their own rational points of view rather than adopting and sharing the perspective of the organization.

Organizations establish hierarchies of goals. Means for achieving general goals become themselves subgoals. Each level is considered an end relative to the levels below it and as a means relative to the levels above it. "Through the hierarchical structure of ends, behavior attains integration and consistency for each member of the set of behavior alternatives is then weighed in terms of a comprehensive scale of values—the ultimate end" (Simon 1976, 63). An organization's hierarchy can be viewed as a congealed set of means-ends chains promoting consistency within the organization. Rationality, then, resides in the structure of the organization itself and not merely in the individual participants. It is to be found in "the rules that assure participants will behave in ways calculated to achieve desired objectives, in control arrangements that evaluate performances and detect deviance, in reward systems" (Scott 1998, 54).

Given the rational systems perspective, we can now clarify what it means to assume that an organization is rational. When we assume that an organization is rational, we assume that it is structured in such a way as to achieve its goals efficiently. The norms of intelligibility that govern our own cognitive enterprises and our interpretations of others are applicable at the organizational level. The rules and policies of an organization present a unified perspective from which decisions can be made, information assessed, inferences drawn, consequences considered, means to ends determined, inconsistencies avoided, and joint actions performed. Achieving goals efficiently

involves engaging in the same sorts of activities rational individuals engage in, except in an organization these activities are often carried out jointly.

Once the assumption of rationality is in place we attribute beliefs, intentions, and desires to organizations in the same way we do to individuals. We ascribe to the organization those beliefs that an organization ought to have if it wants to achieve its goals efficiently given its function in the environment. We attribute desires in the same manner. And such attributions are made holistically as well. They presuppose other beliefs, goals, and intentions. These attributions are then the basis for explanation and prediction of the actions of organizations.

Do we actually apply this method in our explanation of organizational behavior? That is, do we adopt this interpretative stance when we try to make sense of organizations? I think reflection on our everyday practice reveals that we do. Consider the following example: Suppose we want to explain the incidents of gunfire during a naval blockade. We read the official rules of engagement published to govern Navy operations. These rules embody the Navy's rational point of view. We can then explain why the Navy sometimes fires at other ships. Under certain conditions, specified in its rules of engagement, the Navy will fire at anyone it believes to have hostile intent.[19] Or consider the predictions we make about car companies. What will Ford Motor Company do in response to the rise in gas prices? If Ford is rational then it will act so as to maximize its profits. Because Ford wants to avoid losing money on its line of large vehicles and believes that individuals are less likely to buy large vehicles during a time at which gas prices are high, we can predict that Ford will discount its large vehicles.

I submit that predictions and explanations like this are commonplace and very successful. Our practice of interpreting the actions of organizations is just an extension of our practice of making sense of individuals, and it is governed by the same constitutive rules. Our attempt to make sense of the actions of organizations would fail unless we assumed that the organization itself is rational. This involves assuming that the organization has a rational point of view from which members engage in the same sorts of cognitive activities individuals engage in and that the organization is governed by the same norms of rationality.

The success of our practice of explaining and predicting organizational acts is a result of the fact that organizations are designed to be rational—just as the success of our practice of making sense of individuals is a result of the fact that individuals are designed, for the most part, to approximate ideal rationality. If our assumption of rationality is justified in the case of organizations then this is all we need to know in order to count organizations as intentional agents. If all there is to being a believer is to be a system whose behavior is interpretable from the intentional stance then organizations are true believers.

IV

A full defense of the thesis I develop in this paper would require, among other things, a lengthy defense of interpretationism. I do not have space here

to provide such a lengthy defense. I will, however, consider and respond to a few obvious objections that one might raise to the view I have developed.

Throughout this paper I have claimed that our explanations and predictions of organizational behavior in terms of the intentional states of the organization are successful. Critics may charge, however, that their success is exaggerated and that predictions and explanations in terms of the individual actions and intentional states of organizational members have more success. The intentional stance may be easier and more practical, but according to this objection, it is not explanatorily powerful. The best explanation and the most accurate predictions occur at the level of individual psychology. Call this the objection from individualism. Explanation of social phenomena, claims the individualist, requires reduction.

An analog of this objection has been raised against interpretationism at the level of individuals. Belief-desire explanations may be useful to some extent, but true explanation comes at the physical level. The physical stance provides us with a more complete and powerful explanation. Dennett argues quite convincingly, however, that the intentional stance yields predictions and explanations that are available from no other stance. "The patterns of human behavior describable from the intentional stance and only from that stance support generalizations and predictions" (Dennett 1981, 64). As evidence Dennett provides a thought experiment. Imagine that a race of superintelligent beings comes to earth. These beings are LaPlacean superphysicists. They are able to discern the movements of subatomic particles and, based on this, to predict with accuracy future events. Now imagine that one of these aliens engages in a prediction contest with a human well-versed in folk psychology (any human will do). The two contestants observe the following scenario. A woman receives a phone call. She answers and says: "Hello? Hi dear! Oh, Tim's coming with you? O.K. Be sure to stop and get a bottle of wine so we can offer him some at dinner." Their task is to predict what will happen next. Here is Dennett's account of the results of the contest:

> On the basis of this observation, our Earthling predicts that a large metallic vehicle with rubber tires will come to a stop in the drive within one hour, disgorging two human beings, one of whom will be holding a paper bag containing an alcoholic fluid. The prediction is a bit risky, perhaps, but a good bet on all counts. The Martian makes the same prediction, but has to avail himself of much more information about an extraordinary number of interactions of which, so far as he can tell, the Earthling is entirely ignorant. . . . The Earthling's performance would look like magic! How did the Earthling know that the human being who got out of the car and got the bottle in the shop would get back in? (Dennett 1981, 65)

The human's predictive success is a result of the fact that certain patterns of behavior are discernible only from the intentional stance. Logically analogous arguments apply to organizations. The pattern of a naval blockade is constituted by the movements of ships. This pattern is accessible only from the intentional stance. There is no way in particle physics to characterize

"neutral shipping" or "allied navy." These are patterns of *social* behavior discernible only if we take the intentional stance toward organizations.

Our individualist may claim, however, that I have misunderstood his objection. Although particle physics won't successfully yield predictions and explanations (because it misses certain social patterns), these social patterns are nothing more than the patterns of individual human behavior. The individualist claims that we must provide explanations and predictions not in terms of physics, but in terms of individual psychology.

In his article "Beliefs and Desires Incorporated" (1994) Austen Clark argues that the official policy statements of organizations are essential to our practice of predicting organizational behavior but that such official statements should not be seen as legitimate *explanations* of the actions of organizations:

> 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. . . . How do we go about it [explanation]? We leave aside the vocabulary of official statements and descend to the level of interactions among people making up the government. (Clark 1994, 419)

Underlying this position is the view that belief-desire explanation is a type of causal explanation. This view is advocated by some interpretationists as well, notably Davidson. According to Davidson, reasons would not have the explanatory force they have unless they were also causes (Davidson 1963). If belief-desire explanation is causal explanation then it doesn't seem right to explain the actions of organizations in terms of desires and beliefs. How could these collective beliefs and desires be causally efficacious in the production of organizational acts? They do not seem to be physical states of the organization, and only physical states (or events) can engage in causal relations. Organizational beliefs and desires are epiphenomenal and hence causally inert, according to this objection.

Again, there is an analog of this objection that applies at the level of individual psychology. If one thinks beliefs and desires are not reducible to brain states (because, for instance, their content is determined by their relation to the environment, and hence intentional states do not supervene on brain states) then beliefs and desires appear to be epiphenomenal and are not causally responsible for the production of other states and behavior. The need to salvage the causal explanatoriness of belief and desire has led many down a reductionist path. Beliefs and desires must be brain states in order for belief-desire explanation to be true causal explanation.[20]

Although I do not have the space here to give this objection and its analog extended discussion, here is the beginning of a reply to both. Reductionism is motivated by a view of causality that is too stringent. This point has been argued quite forcefully by Lynne Baker in *Explaining Attitudes* (1995). Baker points out that the conception of causal explanation on which reductionism is based is too restrictive. It is restrictive in the sense that it rules out certain

types of explanations as causally explanatory. This is prima facie unacceptable because these explanations comprise a great collection of successful everyday explanations as well as explanations in economics, politics, and psychology. Such explanations make no reference to intentional states but, nonetheless, presuppose intentional states. Baker uses the following as an example: "Al's application for a gun permit was turned down. The causal explanation is that Al is a convicted felon. If he had not been a convicted felon, he would have received the gun permit" (Baker 1995, 154). Baker argues that if we adopt the criterion of causal explanatoriness present in the views of, for instance, Jaegwon Kim (1993), this perfectly good causal explanation must be rejected. This suggests that one ought to reject the accounts of causality underlying the reductionist project.[21] If we adopt a less stringent notion of causality, perhaps a counterfactual analysis of causation, we can salvage the causal explanatoriness of beliefs and desires at the level of individuals and organizations.

A second objection that might be raised to the view I have developed involves challenging my claim to have defended option 3, the view that organizations really do have beliefs. Interpretationism, according to some critics, is just a form of instrumentalism, and if all I am claiming is that from a certain perspective one can view organizations as having intentional states then organizations don't *really* have intentional states, we just treat them as if they do. Therefore, I haven't defended option 3 but option 1—our ascriptions of intentional states to organizations are useful fictions but, strictly speaking, false. A brief look at Dennett's discussion of realism will, I think, provide some response to this concern.

Interpretationism is a "mild and intermediate sort of realism" (Dennett 1991, 29). The patterns of human behavior discernible from the intentional stance are objectively out there to be interpreted. In "Real Patterns" (1991) Dennett develops this point more fully. He appeals to Gregory Chaitin's definition of mathematical randomness. "A series (of dots or numbers of whatever) is random if and only if the information required to describe (transmit) the series accurately is incompressible: nothing shorter than the verbatim bit map will preserve the series. Then a series is not random—has a pattern—if and only if there some more efficient way of describing it" (Dennett 1991, 32). A pattern, then, is real, if there is a description of the data that is more efficient than the bit map. There may, of course, be several different ways to efficiently describe the data, but that fact does not mean that the pattern is not really out there to be described.

Dennett applies these considerations to the arena of intentional state attribution. The patterns of human behavior are real patterns because there is a description of human behavior that is more efficient than a description that cites micro physical properties. The more efficient description is the one offered at the level of folk psychology. It is only from the intentional stance that we can discern these real patterns. There may be several ways of interpreting an agent, each equally successful, but this fact does not mean that these patterns are any less real.

We can extend Dennett's point to organizations. The interpretative stance we take toward organizations is able to discern *real* patterns of social behav-

ior, patterns that are missed if one attempts to explain the social world by appealing only to individual intentional states.[22] Interpretationism, then, allows us to be "mild and intermediate" realists about the intentional states of organizations.[23] We need not posit a group mind or consciousness in order for our ascriptions of intentionality to organizations to have legitimacy, nor do we have to think of these ascriptions as mere metaphors. Organizations and individuals really have beliefs, but their ontological status is more akin to centers of gravity than to tables and chairs.

One final objection needs to be addressed. A typical response to my line of argument is that if interpretationism yields the result that organizations can be true believers then this is all the more reason why interpretationism (or my version of it) should be rejected as a plausible account of intentionality. My thesis constitutes a reductio of interpretationism. The objection can be formulated in the following way:[24]

1. If interpretationism is correct then organizations really have intentional states.
2. It is ridiculous to think that organizations really have intentional states.
3. Hence, interpretationism must be false.

The argument loses its force, I think, when we uncover some of the presuppositions motivating premise 2. The idea that organizations could have beliefs and desires is often rejected because organizations do not have minds or brains. But we have already noted that this does not support the view that they cannot have intentional states. Whether something has a mind is dependent on its having intentional states, not the other way around, and the failure of type-identity theories suggests that having a brain is not necessary for intentionality.

But the fear of group minds is not the only thing motivating premise 2. Traditionally, the notion of belief and intentionality has been thought to be linked to notions like a first-person perspective, consciousness, and self-awareness, and these concepts just don't seem applicable to organizations. Organizations do not have a consciousness that is separate and distinct from the consciousness of individuals, and it would be difficult to say in what an organization's self-awareness consists. It is the absence of these things from groups rather than simply the absence of a mind or brain that motivates skepticism with respect to group beliefs. Groups, according to some, do not seem to have the features necessary for intentionality. Hence, they cannot be true believers.

We have seen, however, that groups do share one important feature with individuals. They can, under certain conditions, constitute a rational point of view. This is not a first-person perspective but a *plural* perspective.[25] Such a perspective must be presupposed if interpretation is to be successful. But what about consciousness and self-awareness? The fact that organizations lack consciousness and self-awareness need not lead to skepticism about collective belief. There is another conclusion that can be drawn. Consciousness and self-awareness are plausibly thought of as properties of individual subjects. Since groups do not have these properties the conclusion one should

draw is that groups do not have the features necessary for *individual* subjecthood. This, of course, is not a strange conclusion to draw. Groups are *not* individual subjects, they are groups! In Gilbert's words, they are *plural* subjects. Note that this does not rule out the possibility that groups have the features necessary for intentional agency. If the differences lie in the nature of the subjects rather than in the nature of intentionality, there is reason to think that groups can be intentional subjects, just subjects of a different kind. The fact that my version of interpretationism acknowledges organizations as true believers should not, then, be seen as a reductio. On the contrary, it should be seen as evidence of its strength as a theory of intentionality. It provides a univocal concept of intentionality that is applicable to different kinds of subjects.[26]

Notes

[1] I define organizations as collectives oriented to the pursuit of relatively specific goals that exhibit a high degree of formalization. It is the high goal specificity and relatively high formalization that distinguish organizations from other types of collectives such as primary groups, families, communities, and social movements. I follow rational system theorists in this definition. For a discussion of rational system theories and various other approaches to organizations, see Scott 1998.

[2] When I refer to the action of an organization I am not presupposing that the organization is something over and above its members that acts independently of its members. I adhere to David Copp's (1979) account of collective action. The actions of collectives are secondary actions constituted by the primary actions of individuals.

[3] I mean to include here all figurative uses of speech, not just metaphor. This would include the view that our ascriptions of intentional states to groups is an instrumental use of folk psychology.

[4] There are some immediate objections one might raise here. First, one might claim that we do not hold corporations *morally* responsible. Reflection on the way we talk about corporate acts seems to suggest otherwise, however. We certainly praise and blame corporations in our everyday moral discourse. One might also object that corporations are not held legally responsible in the same way that persons are. Legal cases involving corporations are civil cases rather than criminal cases. In civil cases there is no need to prove intent. Further, not the corporation itself, but the individuals in the organization responsible for organizational acts, are held legally responsible. Both these claims are false. Organizations are now frequently brought up on criminal charges, and such charges need not be made against any individual in the organization. For an interesting discussion of the law and criminal charges against corporations see ⟨www.usdoj.gov/04foia/readingrooms/6161999.htm⟩. This is a memorandum from the Deputy Attorney General on bringing criminal charges against corporations.

[5] A similar point is made by David Vellman in "How to Share an Intention" (1997).

[6] The reason I do not adopt this line of argument is that I am interested in saying something about our practice of ascribing intentional states to organizations. Since functionalism does not make much of the interpretative side of belief ascription, I find interpretationism to be a more applicable framework.

[7] The version of interpretationism I present here preserves elements of both Davidson and Dennett but differs from both accounts in various ways that should be apparent to those familiar with their work.

[8] Davidson (1963) makes this distinction.

[9] This does not mean that in interpreting others we must assume that they believe everything that we believe or that we must always agree. Disagreements are inevitable. But as Davidson points out disagreements are only possible against the backdrop of agreement.

[10] The notion of a rational point of view as a point of view from which certain cognitive activities proceed provides us with a much richer conception of rationality than can be found in either Dennett's or Davidson's discussion of interpretationism. The norms of rationality require not just that our beliefs be consistent and mostly true, but that we engage in certain activities that promote consistency and truth. Rationality, then, is not a state one reaches (i.e., the state of having all consistent beliefs) but an activity. This is what makes my account of interpretationism novel. Rovane does not develop an account of interpretationism and does not link her conception of a rational point of view with intentionality. She is concerned with defining the notion of personhood and providing an account of personal identity. From her account of personhood she argues for the *possibility* of group persons and a plurality of persons in one body. In some ways my thesis is both weaker and stronger than Rovane's. It is weaker because I do not think organizations are persons or could be persons. It is stronger because I argue that, because organizations are interpretable, our assumption of rationality is justified, and hence organizations are rational agents with rational points of view. I argue not just for the *possibility* of group belief but the *actuality* of group belief.

[11] Davidson and Dennett adhere to what Kirk Ludwig (1993) calls a meaning or content molecularism. This is the view that to have an attitude with a certain content or speak a language in which a term has meaning, one must have some attitudes or expressions with meanings from a range of related contents or meaning.

[12] How does one determine what desires and beliefs an agent ought to have? Following Dennett, we can adopt the following guidelines: attribute as beliefs all of the truths relevant to the system's interests that the system's experience to date has revealed. An implication of the intentional stance is that believers believe mostly truths. If we are to attribute the beliefs that a system *ought* to have then we will not attribute to the system many false beliefs. False beliefs will not promote a system's interests, and so they are not beliefs a system *ought to have*. This is not to say that we cannot attribute false beliefs, but such attributions are made against the backdrop of mostly true beliefs. If attribution of falsehoods is required there must be a story to tell that justifies such an attribution. Attribution of desires involves attribution of the most basic desires: survival, absences of pain, procreation, play, etc. And the attribution of absurd or harmful desires requires a special story, just as the attribution of false beliefs requires a special story.

[13] The view that humans are designed by evolution to be rational is Dennett's view in "True Believers" (1981).

[14] One may immediately object to an extension of interpretationism to organizations on the following grounds: in order to interpret organizations we will need to assume that organizations are rational. That is, we will have to assume that they have a network of beliefs similar to our own and a rational point of view from which cognitive activities originate. But this is the very thing for which I argue! Interpretationism appears to beg the question with respect to collective intentionality. I think there are several things we can say in response to this objection. First, when we descend to the level of individuals one can raise the same objection to interpretationism. Interpretationism seems to presuppose that an agent is interpretable in order to interpret it. This is an accurate description of the process, but it need not be problematically circular. Because of the interdependence of meaning and thought there is no way to approach intentionality from outside of the intentional realm. Further, we can think of the interpretative process as a pseudoscientific method. Scientists often assume that a theory is true in order to see what predictions it will yield. If the theory yields successful predictions adoption of the theory is justified. This simplifies the scientific method a bit, but we can see interpretationism as engaging in the same sort of enterprise. See Dennett 1987 (50) for a discussion of the unproblematic circularity involved in the intentional stance.

[15] Oddly, Dennett himself never considers organizations or groups as possible intentional systems. Since for Davidson neither animals nor children are intentional agents he would probably not accept organizations as intentional agents. The only person I am aware of who has suggested extending interpretationism to groups is Austen Clark (1994). Clark extends Dennett's version of interpretationism to organizations and argues that one is either forced to accept that belief-desire explanation is not legitimate

explanation (because it works on organizations just as well) or to accept eliminativism. Though I have profited from reading Clark's article I disagree with his view. I also think more needs to be said about the way in which interpretationism can be legitimately applied to organizations. I do so in this paper. In *Collective and Corporate Responsibility* (1984) Peter French argues that organizations are believers. He does this by adopting Davidson's account of events and descriptions rather than his account of intentionality. Just as the same mental event can be described from the intentional stance and from the physical stance, French suggests that the same social event can be described as the action of individuals or the action of the organization. Although I find this approach interesting, I think it fails to say what it is about organizations that licenses such a redescription. French argues that it is the internal decision structure that licenses the redescription. But this doesn't go far enough. I would argue that the internal decision structure makes it possible to see the organization as having a rational point of view and, hence, as a rational agent. Once the assumption of rationality is in place we can then interpret the organization. If our interpretations are successful our assumption of rationality is justified.

[16] See Dennett 1981, 68.

[17] There has been, in the past fifty years, a tremendous amount of literature produced on organizations and organizational theory. See Scott 1998 for an overview of organizational theory.

[18] The failure of individuals to act rationally in groups, even though outside of groups they act rationally, is best documented in the literature on mobs. See Kendrick 1987 for a discussion of some of the literature in social psychology on the psychology of mobs.

[19] This example is a version given by Clark (1994).

[20] See Davidson's "Mental Events" (1970). Also, see Kim 1993 for a view motivated by the same sorts of causal worries.

[21] I am aware that there are those who would claim that this is not a "perfectly good" causal explanation. There are complicated issues here that I do not have the space to develop. Let me just say that even if one finds Baker's position untenable there are other ways of avoiding the causality problem. One solution at the level of individual beliefs and desires is to say that although beliefs and desires are not causally efficacious they are causally relevant. I am currently exploring the possibility of extending Philip Pettit's (1993) account of the causal relevance of second-order properties to argue for the causal relevance of collective beliefs.

[22] It will be helpful here to extend Dennett's thought experiment. Imagine that two people are engaging in a prediction contest. Their task is to predict what Ford Motor Company will do in response to the enormous increase in gas prices. One of the participants is an individualist. He is certain that social phenomena can be explained and predicted by appealing to individual intentional states. The other is a collectivist, and she believes that predicting the behavior of an organization involves viewing the organization as a rational agent. In order to predict what Ford will do the individualist will have to find out who the operative members of the organization are, how each member voted and why he voted that way, and whether he is telling the truth about his intentional states. The collectivist, on the other hand, knowing that individuals are likely to stop buying large vehicles during a time at which gas prices are high and knowing that Ford sells a great deal of these vehicles and wants to continue to maximize its profits, will predict that Ford will discount these vehicles in the near future. The prediction is a bit risky perhaps but a pretty good one nonetheless. The individualist will, perhaps, make the same prediction but will have expended a considerable amount of time and energy. I say "perhaps" because the individualist may miss the fact that when individuals act in their organizational roles they often act differently than they would outside of the organizational context. Because the individualist is working at the level of individual psychology he may not have room in his theory for notions like organizational context or role. These concepts are concepts of the social scientist, not the individual psychologist. Compared to the individualist the collectivist performance will look like magic. How did she know Ford would lower its prices on large vehicles without even talking to the president of the company?

[23] This might be seen as an additional benefit of adopting the interpretationist framework over, say, a functionalist account of group intentionality.
[24] Clark (1994) considers a version of this argument.
[25] I think this is what Gilbert really has in mind when she talks of plural subjects.
[26] This conception of intentionality applies to animal subjects as well. Note, however, that the requirements of rationality as I develop them (rational point of view, rationality as activity) may exclude certain types of subjects that Dennett's version of interpretationism includes in the class of intentional systems.

References

Baker, Lynne. (1995) *Explaining Attitudes*. Cambridge: Cambridge University Press.
Clark, A. (1994) "Beliefs and Desires Incorporated." *Journal of Philosophy* 91:404–25.
Copp, D. (1979) "Collective Actions and Secondary Actions." *American Philosophical Quarterly* 16, no. 3:177–183.
Davidson, D. (1963) "Actions, Reasons and Causes." *Journal of Philosophy*, 60:685–700.
Davidson, D. (1970) "Mental Events." In *Experience and Theory*, ed. L. Foster and J. Swanson. London: Duckworth.
Davidson, D. (1975) "Thought and Talk." In *Mind and Language*, ed. S. Guttenplan, 7–23. Oxford: Clarendon Press.
Davidson, D. (1984) *Inquiries into Truth and Interpretation*. Oxford: Clarendon Press.
Dennett, D. (1981) "True Believers: The Intentional Strategy and Why It Works." In *Scientific Explanation*, ed. A. Heath, 53–75. Oxford: Oxford University Press.
Dennett, D. (1987) "Three Kinds of Intentional Psychology." In *The Intentional Stance*, 44–68. Cambridge: MIT Press.
Dennett, D. (1991) "Real Patterns." *Journal of Philosophy* 88:27–51.
French, Peter. (1984) *Collective and Corporate Responsibility.* New York: Cambridge University Press.
Gilbert, M. (1988) *On Social Facts.* New York: Routledge.
Gilbert, M. (1996) *Living Together.* Lanham, Md.: Rowman & Littlefield.
Hornsby, J. (1997) Simple Mindedness: *A Defense of Naive Naturalism in the Philosophy of Mind*. Cambridge, Harvard University Press.
Kendrick, Clyde. (1987) *Group Processes: Review of Personality and Social Psychology.* London: Sage Publications.
Kim, Jaegwon. (1993) *Supervenience and the Mind*. Cambridge: Cambridge University Press.
Ludwig, K. (1993) "Is Content Holism Incoherent." In *Holism: A Consumer Update*, ed. J. Fodor, 45–67. Amsterdam: Rodopi.
McDougal, William. (1920). *Group Mind.* New York: G.P. Putnam's Sons.
Pettit, Philip. (1993) *The Common Mind.* Oxford: Oxford University Press.
Pettit, Philip, and Frank Jackson. (1988) "Functionlism and Broad Content." *Mind* 97:381–400.
Pettit, Philip, and Frank Jackson. (1990) "Program Explanation: A General Perspective." *Analysis* 50:107–17.
Rovane, Carol. (1997) *The Bounds of Agency.* Princeton: Princeton University Press.
Scott, R.W. (1998) *Organizations: Rational, Natural and Open Systems*. Upper Saddle River, N.J.: Prentice Hall.
Simon, H. (1976) *Administrative Behavior: A Study of Decision Making Processes in Administrative Organization.* New York: Free Press.
Taschek, W. (1988) "Making Sense of Others: Donald Davidson on Interpretation." Paper prepared for the Fellows Seminar Institute for the Humanities, University of Michigan.
Tuomela, R. (1993) "Corporate Intention and Corporate Action." *Analyse und Kritik* 15:11–21.
Tuomela, R. (1995) *The Importance of Us.* Stanford: Stanford University Press.
Vellman, D. (1997) "How to Share an Intention." *Philosophy and Phenomenological Research* 57 (March).