

JOSEPH THOMAS TOLLIVER

KNOWLEDGE WITHOUT TRUTH

(Received in revised form 20 November, 1987)

1. INTRODUCTION

Knowledge implies truth. No one can know that Howard the Duck is the President of the United States, because no one can know that p when p is false. People come to have knowledge by coming to stand in an appropriate relation to a true proposition; this relation is usually called "belief," "assent," or "acceptance." Belief relates a person to a proposition because belief is both a state of a person and a state with propositional content, i.e., intentionality. Thus, people assert, claim, or communicate knowledge by uttering sentences that have as their meaning propositions that are contents of their beliefs. The appropriate answer to the question, "What does S know?" is a list of true propositions that S believes. We have here the propositional paradigm for knowledge, i.e., knowledge is best represented as a set of true propositions a person believes.

Recently the central assumption of the paradigm has come under attack. Churchland (1979), Field (1984) and (1986), and Stich (1983)¹ have argued that there are good reasons to abandon the notion that people have propositional attitudes. The arguments purport to show that ascriptions of semantic properties, such as truth or reference, to states of persons is unnecessary or down-right counterproductive to the enterprise of empirical psychology. They suggest that people do not have belief states with truth conditions or desire states with satisfaction conditions. The arguments are of course many and various, but they take two basic forms: (1) arguments that psychological explanation can proceed without use of content notions, and (2) arguments that psychological theories ought not to make use of content notions. The first set of arguments are based on the claim that semantic properties are explanatorily redundant because whatever can be explained by means

of them can be explained without them.² The second is based on the claim that content attribution is always in some measure idiosyncratic. Again, this objection is pressed in a number of different ways, two prominent ones being that: (i) the alleged holism of meaning renders generalizations over meanings unlikely to hold across persons³ and (ii) that the underlying rationality assumptions of content ascriptions limits the applicability of psychological explanations to creatures with inferential competencies similar to ours.⁴ Ignoring for the moment the contribution of a justification condition on knowledge, we seem to lose our principal means for distinguishing knowers from nonknowers, viz., that knowers have a relation to the world, i.e., satisfaction of the truth conditions of their internal states, that nonknowers lack. Suppose for a moment that we take these arguments at face value, that we take them as establishing that we have no propositional attitudes. Accepting this has a couple of immediate consequences: (i) belief does not relate a cognizer to a proposition, and (ii) knowledge does not relate a person to a true proposition. This requires rejection of the propositional paradigm for knowledge, since we no longer have reason to think that we can represent a person's knowledge by means of the propositions he knows to be true. We therefore seem to lose the ability to distinguish knowers from nonknowers.

Is there any alternative to reliance on truth conditions for distinguishing knowers from nonknowers? I think there is, because I think that it can be shown that representation is not necessary for knowledge. Knowledge surely requires some form of epistemic correspondence relation, i.e., a relation that determines what one has knowledge of when one has knowledge, between states of a person and the world. Surely representation relations such as reference and truth condition can be used to model such relations, but they are not the only means at our disposal for representing the relation between beliefs states of a person and some state of affairs in the person's environment when that person has epistemic access to that state of affairs. I will suggest that even if the internal states of persons have no semantic or representational properties, they still bear epistemically important relations to the world. Let us say that some internal states of cognitive systems have *objects*. I intend to show that the cognitive object relation can be given a purely causal nonintentional analysis and that it can serve as a form of epistemic correspondence relation.

The account of the cognitive object relation will be teleological and control theoretic. In brief, a cognitive state c of a person S has an event or state of affairs o as its object just in case o has a certain kind of influence on the internal control function of c under conditions of proper functioning in S . Considered control theoretically, objects, events, and states of affairs in the world are not objects or contents of belief and desire because they are objects or contents of *de re* or *de dicto* propositional attitudes, but rather because they are states that modulate the control function of belief and desire states. The account of knowledge will be information theoretic. Again in brief, a person has knowledge when he is informed with respect to those states of affairs that are the objects of his beliefs.

2. RELIABILITY AND TRUTH

Rejection of the propositional paradigm for knowledge requires rejection of the notion that we can represent a person's knowledge by means of the propositions he knows to be true. What we need instead is an account of knowledge other than one based on representational states linking the knowing mind to the world by means of their semantic properties, i.e., we need a nonrepresentational account of the epistemic correspondence relation. I intend to build such an account around a reliability theory of knowledge. Reliability theories tend to portray the knowing mind as a reliable measuring or indicating instrument. D. M. Armstrong (1973) compares a knower to a properly working thermometer.⁵ A thermometer reading that fails to correspond to the actual temperature is like a false belief, one which does correspond is like a true belief. A correct reading on an improperly working thermometer is like true belief that is not knowledge. A correct reading of T degrees on a properly working thermometer ensures that the local temperature is indeed T degrees. This is analogous to a person who has knowledge. This analogy suggests to me the possibility of an account of knowledge in terms of a more direct relation between the knowing mind and the world than one mediated by the semantic properties of belief. But as the above summary of Armstrong's view suggests, most reliability theories of knowledge depend on the supposition that beliefs have truth conditions.

The concept of truth usually plays two important roles in a reliability account of knowledge. It provides the basis for an account of epistemic reliability, i.e., that kind of reliability necessary for knowledge. Epistemic reliability is a property of a belief or of a belief forming process. An epistemically reliable belief is one that guarantees that its truth condition obtains, or increases the probability that it obtains.⁶ An epistemically reliable belief forming process is one that produces more true beliefs than false beliefs.⁷ The other link to truth is related to the first. The truth condition of a belief determines what features of the world one must be epistemically reliable with respect to in order to have knowledge, i.e., it determines the epistemic correspondence relation. If S believes that there is a cold beer in his refrigerator, then, roughly, to know that there is he must be epistemically reliable with respect to the state of affairs that would make his belief true.⁸

Reliability is not a property of the proposition that is the content of the belief, but of the belief state relative to that proposition. It is a second-order property of the belief state, i.e., a property the state has in virtue of properties of the first-order states that instantiate the belief. These second-order properties are variously understood as causal, nomological, or conditionship properties of the first-order states. Since reliability is a property of the belief state, one is led to wonder what theoretical work the reference to a proposition is doing here and whether some other construct might do the same work. Clearly, the proposition is functioning as a marker here, as a way of picking out a particular set of second-order properties. Which set? The set that constitutes the reliability of a particular belief with respect to a state of the world. Which state of the world? The state of affairs that is the truth condition of the proposition.

Now it is just this theoretical resource that the eliminativist arguments aim to deny us. These arguments purport to show that belief states bear no theoretically important relations to propositions. So the decision to take these arguments at face value implies that we must find some other means of indexing the relevant second-order properties. We need some other way of connecting a belief state to a particular state of the world and of defining its reliability properties without relying on any semantic properties of the state.

3. THE PRAGMATIC ANALYSIS OF PROPOSITIONAL ATTITUDES

It is a commonplace observation these days that belief and desire are correlative concepts. They are used to ascribe mental states whose contents are interspecificable relative to the role they play in the production and control of rational action. S's belief that Dobbin will win the next race is the belief that it is in part because of the role it plays, along with S's desires, in producing S's betting activity. This commonplace underlies the so-called pragmatic account of the propositional attitudes. This account has it that the contents of mental states are determined by the structure of the causal role they play in the production and control of rational action. Here is Robert Stalnaker's characterization of the approach.

Belief and desire, the strategy suggests, are correlative dispositional states of a potentially rational agent. To desire that P is to be disposed to act in ways that would tend to bring it about that P in world in which one's beliefs, whatever they are, were true. To believe that P is to be disposed to act in ways that would tend to satisfy one's desires, whatever they are, in a world in which P (together with one's other beliefs) were true.⁹

Stalnaker speaks of a person as being the subject of an attitude of belief or desire towards some state of affairs. We can of course extend the pragmatic strategy to include assigning contents to particular states of a person. Let \mathcal{A} be the set of types of actions that S can perform. Let B be the set of belief states possible for S, and D be S's possible desire states. Let o be the name of some particular type of state of affairs. Then o is part of the content of a particular b in B if, and only if there is a member A of \mathcal{A} and there are members d_1, \dots, d_m of D and other members b_1, \dots, b_n of B such that it would be rational to perform A given d_1, \dots, d_m and b_1, \dots, b_n only if b says that o obtains. A particular d in D has o as part of its content if, and only if all of the states of affairs that are included in the contents of S's beliefs obtained and d caused S to perform some action A in \mathcal{A} as part of the execution of a rational plan of action, then S's Aing would tend to cause o to obtain.

In addition to suggesting that the content of mental states is determined by their role in processes that explain behavior, the pragmatic approach also suggests that any state of affairs that can be a content of

belief can be a content of desire and *vice versa*. For any state of affairs that can be an object of desire its being an object of belief requires only that there be beliefs and desires such that, for some type of action, it would be rational to perform it only if some belief state of the person said that the state of affairs in question obtained. Surely some such beliefs, desires, and action type exists for each possible object of desire. Betting situations would seem to be constructible for each such case that would yield the required results. Beliefs about the past seem to be an exception to the converse proposition. We obviously have beliefs about the past, but it is not obvious that we can desire the past, at least not in the sense of desire spelled out above.¹⁰ The problem is that the pragmatic account says that an object of desire is something that, under certain circumstance, our actions would tend to bring about. Given the nomological impossibility of changing the past it would seem that our actions cannot bring the past about — it already exists! This is a problem only if the person's beliefs include the belief that the state of affairs in question is past. A person surely can desire a state of affairs that is past, but still awaited. The classic example is the child awaiting an injection, desiring that it be over, only to be told that it has been over with for some time. So, surely, we can have past events as object desire. It should also be clear that we can have the pastness of an event as an object of desire, for it is the pastness of the injection which the child desires. So the belief and action types relative to which past events are objects of desire on the above approach are simply ones that do not include the belief that the state of affairs is past.

If every content of belief is a potential content of desire, and *vice versa*, then for every belief that has a particular object there is a corresponding desire that has the same object. So the problem of specifying the objects of belief reduces to the problem of saying when a given belief state has the same object as a given desire state and saying what that object is. Stalnaker's version of the pragmatic analysis will not be of use to us because it employs the very semantic notions that we are seeking to avoid, for it is based on the role various internal states play in the explanation of behavior via various practical reasoning frames. The rationality assumptions involved in this mode of explanations means that content notions cannot be eliminated from them. Control theory provides us with one way of explaining the functional role of

various internal states in the production of behavior without the need for attributing semantic properties to those states. Control theory explains behavior by showing how it is goal directed. The goal directedness of behavior is understood by means of the role of feedback in the modulation of responses. Feedback is information about the effects of one's actions on the world. A goal state is one the detection of which removes its internal index for its role in controlling behavior. The system has obtained what it "wants".

My version of the pragmatic account characterizes desire states as internal indices of goals. They play a role in the construction and activation of plans of action. Plans of action are complex internal structures of instructions that actually direct the behavior of the organism. Ideally a plan of action remains active during its execution until the motive of that plan is realized. Desires, by means of their place in the hierarchy of preferences, their relations to beliefs, and their relations to plans of action, exercise influence on the flow of control within the system. The flow of control from one part of an action plan to another and from one plan of action to another is, in part, determined by the goals of the system, indirectly, via their internal index. Beliefs are states that indicate the satisfaction of desire, and as such remove desire states from a role in the flow of control of behavior. A belief state should have this effect when it says that the system has what it wants. We must assume the system is operating properly, for we want to be able to say that x is the goal of S's behavior even if, due to malfunction, the system stabilizes in circumstances that do not include x .

These relationships between belief, desire, and the control of behavior are most easily illustrated in the case of seeking and finding behavior. You see someone making their way about a room, looking under piles of books, on shelves, in end table drawers. You do not need to know anything more about the person for you to reach the confident conclusion that he is looking for something. He is engaged in search behavior. From the places he looks you can even get information about what the object of his search is — a letter perhaps, but certainly not a rake. Whatever it is, he has not found it yet. Suddenly he ceases looking in the room and makes directly for the den, to a desk where he opens a book and removes a scrap of paper. It is a shopping list. He turns and

proceeds to leave the house. His search behavior has ended. He obviously has a desire state that had that list as its object. When he opened the book his perceptual state produced a belief state that ended the influence of that desire state on his present behavior. Assuming that all is working properly, that this is the kind of belief state that should remove that desire from the flow of control, we should interpret these states, control theoretically, as being about the same thing. Control theoretically, nothing more could be relevant. The pragmatic approach implies that nothing which makes no difference to the explanation of behavior of a person S makes a difference to the content of S's mental states. This is so whether the mode of explanation is folk psychological or control theoretic. A control theoretic version of the pragmatic account of the objects of the attitudes will form the basis of my account of the epistemic correspondence relation.

4. BELIEF/DESIRE CORRESPONDENCE

Let us say then that a desire state is *satisfied* when, under conditions of proper functioning, a belief state cancels its influence on the flow of control within a cognitive system. Let us call the process wherein a belief cancels the control influence of a satisfied desire "matching". The belief state that does so is a *correspondent* of that desire state. The correspondent of the desire indicates the satisfaction of desire. It indicates that included in what is going on around the system now is the object of desire. This is what the system wants.

There is a problem here in that there can be other reasons for a system to end the control influence of a desire state. Consider the following case that illustrates this.¹¹ Sally wants a *Dos Equis*. She has been thinking about having one all day as her office got hotter. But when she got home, all she found in the refrigerator was a Miller Lite. Sally is very thirsty, and a cold *Dos Equis* would be just perfect, but she decides that it is not worth going out to get one and drinks the Miller Lite. Sally satisfies. Having drunk the Miller, Sally's beer-interest is satisfied for the day. She would not accept a *Dos Equis* now if one were offered to her. Here a belief state, under what we can suppose are conditions of proper functioning, defeats the control influence of a desire state, i.e. matches it. The belief state indicated that Sally had a

Miller Lite. If that belief is the correspondent of Sally's desire, then what Sally desired was a Miller Lite. But we wanted to say that Sally desired a *Dos Equis*. The above characterization of belief/desire correspondence does not allow us to distinguish the object of desire from that which would satisfy the desire. In addition a system might acquire a new desire with higher priority in the system's hierarchy of desires and be matched for that reason, or the system might decide that a current plan of action simply cannot be fulfilled. The belief indicating this should remove some desires from the flow of control, but not because the belief indicates the satisfaction of desire. Under conditions of proper functioning a belief state matching a desire might mean any of these things. Is there a condition under which it means only that the desire is fulfilled?

To specify such a condition we must look at the dynamics of control within a cognitive system under various, perhaps counterfactual, conditions. In particular we must look at what the system does in ideal circumstances. This will tell us what the goals of the system are relative to its current plans of action. I will assume that if when executing a well conceived plan of action, i.e., one which ought to lead to successful action most of the time, under conditions of proper functioning (more on this later), the system defeats the control influence of a particular desire state, it tends to be because that desire is fulfilled. We can incorporate this assumption into a condition of belief/desire correspondence. A belief state *b* and a desire state *d* have the same object for a person *S* if, and only if under conditions of proper functioning, if *d* were currently influencing the execution of a well conceived plan of action, *b* would tend to cancel *d*'s influence on the flow of control within *S*. I speak of tendencies here because the best laid plans can go awry through no fault of the person. But the pragmatic approach presumes that, by an large, well laid plans tend to guide successful action. Thus the fact that, in ideal circumstances, Sally's desire is matched by a belief that signals a Miller Lite does not indicate that the two states have the same object, for this is not what should tend to happen. By an large Sally's plan of action should result in her enjoying a cold *Dos Equis*.

At this point we would like to have a completely general theory of good planning. We are unlikely to ever have such tool. Goodness in

planning is no doubt as diverse in its nature as the goals that might be pursued. Certain high-level generalizations are possible: take care and use imagination in the construction of alternatives; review relevant information; provide for failures of intended results; and keep things as simple as possible. But there probably is no single set of principles of planning to guide preparations for nuclear war, a search for part-time employment, and the conduct of a public opinion poll.

4.1. *Proper Functioning*

Time for a few words about proper functioning. Suppose we view persons, insofar as they are knowers, as complex cognitive systems. There are obviously many aspects to proper functioning for cognitive systems, but only one aspect is important to epistemic correspondence relations — the system's evaluation of the significance of the environmental input with respect to its present control function. Of crucial importance will be what states of affairs serve as a basis for the beliefs of the system. To ease our way into the kind of analysis I intend to offer let us first consider what this kind of proper functioning would come to for a system whose internal states do possess propositional content. I will show how an adequate account for this case can be adapted to deal with the no-content case.

Let me start with the following characterization of proper functioning for a cognitive system. Let me say that a cognitive system is functioning properly if the contents of its belief states are conditioned in an appropriate way, i.e., by the information content of the states of affairs that constitute their basis.

- (PF) A cognitive system *S* is functioning properly, iff, for every proposition *P*, *S* believes that *P* only if every state of affairs that could serve as a basis for what *S* believes contains the information that *P*.

(PF) expresses a requirement that the belief contents of a properly working cognitive system be informationally based. But this requirement is surely too strong. A cognitive system would have to be infallible to meet it.

What we want to require is that the system make the best use it can of the information it has available. In this context this amounts to two requirements: (i) the state of affairs that functions as the basis for the system's belief, call it *e*, does not so function because the system does not possess complete information about the state of affairs that is the content of the belief, call it *c*, i.e., even if the system were maximally informed about *c* it would still include *e* among its bases for its present belief; (ii) the system correctly evaluates the relevance of the information content of its basis to the truth of its present belief. By "the proper evaluation of significance" I mean that a cognitive system should take information that has positive (probabilistic) relevance to the truth of its belief (i.e., its belief is more likely to be true, given this information, than it would be otherwise) to have positive relevance (i.e., its estimate of the probability that its belief is true should be higher, given this information than it would be otherwise) and similarly for information with negative or zero relevance. This observation suggests that we eliminate those states of affairs which can serve as a basis for a belief that *P* in part because either they do not contain enough information about *P*, or because the cognitive system mistakes the relevance of the information it receives. I will call states of affairs whose functioning as a basis for the beliefs of a cognitive system is not explained by the deficit or mistakes about the relevance of information that the basis contains about the states of affairs that are the contents of the system's of beliefs *informationally transparent bases for belief*.

(PF) can be revised to take account of this new restriction.

(PF)' A cognitive system *S* is functioning properly, iff, for every proposition *P*, *S* believes that *P* only if every state of affairs that could function as an informationally transparent basis for what *S* believes contains the information that *P*.

(PF)' is an adequate characterization proper functioning for a cognitive system whose internal states have a propositional content determined by the system's causal/informational interaction with the world. The key suggestion is that a properly functioning cognitive system should have an informationally transparent basis for its beliefs. This is just the

notion needed to understand “. . . under conditions of proper functioning . . .” in the characterization of belief/desire correspondence.

(BDC) A belief state *b* of a cognitive system S corresponds in its object to a desire state *d* of S, iff, if *d* were currently influencing the execution of a well conceived plan of action directed towards bringing about the goal state of which *d* is an index and there were an informationally transparent basis for *b*, *b* would tend to cancel *d*'s influence on the flow of control within S.

I of course need to reconstrue informationally transparent basing so as not to presuppose the attribution of semantic properties to internal states. This will involve eliminating all reference in our characterization of informationally transparent basing to information *content*, restricting reference to measures and comparisons of the *amount* of information contained in events and states of affairs. I will also need to replace the relevance notion employed. These notions concern the relevance of the information content of events to the truth of beliefs. These will be replaced by notions of the relevance of informative events to the control function of internal states. Also, evaluations of the relevance of events to the truth value of beliefs will be replaced with evaluations of the relevance of events to the control function of internal states.

5. INFORMATIONALLY TRANSPARENT BASING

5.1. *Information*

Information, in the ordinary understanding of the notion, is a semantic and epistemic notion — to receive information is to learn that some proposition is true.¹² But the standard technical understanding of information is that to receive information is to have uncertainty about the occurrence of events reduced. This is the so-called “reduced uncertainty” concept of information. If we measure the uncertainty concerning an event by means of its probability, the information contained in an event is inversely proportional to its probability. Thus, on the reduced uncertainty view information is conveyed from one event to another in virtue of nomological/causal relations among changes in

the probabilities of events. Thus the newspaper contains information about the starting times of local movies because the nomological/causal relations between the listing in the paper and the screening of the films is such that the probability of the films starting at a particular time is higher given the listings than they would be otherwise. So understood to talk about the information that one event contains about another is merely to talk about the degree of lawful correlation between changes in the probability of these events. Especially, it is not to attribute any semantic content to these events. Another way of putting this is that so understood information laden states have no intentionality, at least not in virtue of containing information. This is the notion of information that I will presuppose in my characterization of informationally transparent basing.

5.2. *Relevance*

Control theoretically, the relevance of environmental input for a cognitive system consists in the effect — enhancing, inhibiting, or indifferent — that receiving that input has on the chances for the successful control of the system's behavior. Successful control is that which results in the system getting what it desires. Some responses to environmental input have positive relevance to control function because they increase the probability of successful action. A feedback system uses input about the results of its last action to modulate the production of control impulses. Let us call these control impulses “volitions”. Volitions are the last kind of states in a given perception/action cycle that are subject to control by instructions or modulation by feedback. The elements of an action that flow from a volition are “purely ballistic” in the sense that no further guidance or correction of the action is possible.¹³ Volitions have a relevance to the success of a plan of action equal to the objective probability of the success of that plan given the occurrence of the volition, minus the probability of success given that the volition does not occur. So an event has positive relevance to the control function of some internal state of a cognitive system just in case the occurrence of that event increases the probability that the state in question will be a cause of the occurrence of a volition that has positive relevance to a plan of the system. So we can replace notions of positive (or negative)

evidential relevance with notions of positive (or negative) control theoretic relevance in our understanding of informationally transparent basing.

Since positive relevance consists in progress towards a goal condition, a system's taking some input to have positive relevance consists in the input causing an increase in the probability of matching. I argued above that a belief state's removing a desire state from any influence on the flow of control means, under certain circumstances, that the desire has been satisfied. Changes in the internal structure of a cognitive system that increase the probability of that happening indicate that the system takes itself to be making progress toward satisfying a desire. To correctly evaluate the relevance of an input then is for the system to take it to have the relevance that it in fact has.

5.3. *Basing*

Belief basing is normally understood as a relation between a content individuated belief state of a person and a proposition, thus the canonical expression of this relation is, "Person S bases his belief that P on proposition Q". I will construe the basing relation as a relation between a noncontent individuated belief state of a cognitive system and a state of affairs. There are two conditions to this basing relation: (i) the state of affairs that is the basis of the belief in question is a cause of the belief, either efficient or sustaining, and (ii) the cognitive system takes the state of affairs to have positive relevance to the control function of the belief state.¹⁴

5.4. *Informationally Transparent Basing*

We are now in a position to have a reconstructed notion of informationally transparent basing suitable for serving as our standard of normal functioning for cognitive systems

- (ITB) A state of affairs e is an informationally transparent basis for the control function of a belief state b of a cognitive system S, iff,
- (i) e is a basis for b ,

- (ii) no increase in the amount of information in *e* or in the amount of information about *e* contained in *b* would result in *e* not being a basis for *b*,
- (iii) *S* correctly evaluates the relevance of *e* for the control function of *b*.

With informationally transparent basing so understood, a properly functioning cognitive system will be one all of whose belief states have an informationally transparent basis.

6. THE OBJECTS OF THE ATTITUDES

I said earlier that the matching of a desire, under certain circumstances, indicates the satisfaction of desire. It indicates that included in what is going on around the person now is the object of desire. This is what the system wants. But since this includes everything that obtains now, the matching does not indicate what in particular is the object of desire. The mode of indication is the system's causal/informational interaction with the world. This course-grained extensional relation does not enable any fine-grained individuation of particular objects of desire, for there are possible worlds where the internal states of a cognitive system are causally/informationally related to every state of affairs at that world.¹⁵

But the conditions on belief/desire correspondence require that there be a tendency for a particular belief state to match a particular desire state under informationally transparent conditions. This tendency will exist in virtue of certain commonalities among the circumstances that occasion the satisfaction of desire. These commonalities are guaranteed by the fact that the person in question is seeking its goal condition, using a well-conceived plan of action, and basing its beliefs on feedback from the world in an informationally transparent manner. Under these conditions the tendency of a belief to match a given desire means that the circumstances in which this tendency is manifested have something in common — the object of desire. Let us then say that an informationally transparent basis (a certain type of state of affairs) that is, or could be, a cause of a belief state's matching a desire state in circumstances that satisfy the conditions on belief/desire correspon-

dence is compatible with the belief state in question. Let us say that a possible world is compatible with a belief state if, and only if, it is compossible with a state of affairs that is compatible with the belief state. Given these stipulations we can define the objects of belief as follows.

- (OB) A state of affairs e is an object of belief state b of person S , iff, every possible world w that is compatible with b includes e as a part.

Likewise, a given state of affairs is an object of a particular desire state just in case it is included in every possible world compatible with that desire state. The set of possible worlds compatible with a particular desire state will, of course, be identical to the set of possible worlds compatible with its corresponding belief state.

7. PERCEPTUAL KNOWLEDGE

An analysis of perceptual knowledge is naturally suggested by this account of objects of belief. A state of affairs can only be an object of belief relative the capacities of the control system of S to "get it right." That is, it can be an object of b only if it can be an informationally transparent basis for b (although (OB) does not require that an object of belief be an informationally transparent basis for the belief in question). Of course even when e is an informationally transparent basis for b , b might be relatively uninformed about e , for all that is required by (ITB) is that increasing the information content of e , or the amount of information about e in b would not cause the system to abandon e as a basis for b . So, a natural additional condition for objects of belief to be objects of knowledge is that b be maximally informed about its object. Let us suppose that the event e consists in some object o having some property F , then perceptual knowledge of this state of affairs can be defined as follows.

- (PK) S has perceptual knowledge of the state of affairs of o 's being F , iff,
 (i) o 's being F is an object of some belief state b of S , and

- (ii) *o*'s being F is an informationally transparent basis for *b*,
and
- (iii) *b* contains the same amount of information about *o*'s
being F as is contained in *o*'s being F.

(PK) characterizes perceptual knowledge as a relation between a person and a state of affairs. Whatever plausibility (PK) has vindicates the suggestion that the belief object relation, as characterized in (OB), can serve as an epistemic correspondence relation, as a relation that determines what in the world the knower knows. Since the (PK) does not require the attribution of any semantic properties to subjects of knowledge attributions, it meets our requirements for an account of knowledge compatible with the eliminativist's position. Notice that S can have knowledge of *o*'s being F even though there is no guise, description, or mode of presentation by means of which S knows this. This because of the directness of the epistemic correspondence relation between S and the state of affairs that is the object of his knowledge. (PK) therefore characterizes a completely transparent epistemic relation. Unfortunately, it must be left to another occasion to discuss the phenomena that tend to drive epistemologists away from this kind of view.

A less strict standard of perceptual knowledge can be obtained from (PK) by replacing the requirement of maximal information with a parametric requirement that fixes the level of informativeness necessary for knowledge by context.

- (PK)^{*k*} S has perceptual knowledge of the state of affairs of *o*'s being F, iff,
- (i) [same as (PK)], and
 - (ii) [same as (PK)], and
 - (iii) *b* contains *k* amount of information about *o*'s being F.

The value of *k* can be fixed by a number of features of the case being evaluated. For example, one feature that sometimes influences our attributions of knowledge is the gravity of the consequences should the object of S's belief not obtain. If all prospects for happiness for S and all his loved ones depends on the existence of the belief object, then we

would use (PK) as our standard of knowledge, i.e., we would want S's epistemic status to come as close as possible to guaranteeing that the F object exists before we agreed that S knows of *o*'s F-ness.

In addition to (PK) and (PK)^k we might adopt the following view of perceptual knowledge.

- (PK)^c S has *c* amount of knowledge of the state of affairs of *o*'s being F, iff,
- (i) [same as (PK)], and
 - (ii) [same as (PK)], and
 - (iii) *b* contains *c* amount of information about *o*'s being F.

Here *c* does not, as *k* does, function as a context sensitive marker of the minimum informativeness necessary for knowledge, rather, it is a continuously valued (perhaps) measure of S's knowledge. The more informed S is about an event the more he knows about it. (PK)^c implies that, in a sense, knowledge is much more plentiful than (PK) or (PK)^k would suggest. S has knowledge whenever he has some information about some state of affairs that is an object of belief. Since some capacity for information pickup is necessary to have any objects of belief in the first place, every individual that has beliefs about the world likely has some knowledge.

I find support within my concept of knowledge for all of the above analyses of perceptual knowledge. I have intuitions that accord with each of them and can use notions associated with any one of them to generate counterexamples to the other two. So appeal to counterexamples hardly seems to be a useful means for choosing among them. A more salient consideration seems to be that we might want some way of factoring out the contribution that epistemic status makes to successful action.¹⁶ Many things obviously bear on S's chances of getting what he wants: the resources available in S's environment, the nature of the functional capacities S has to exploit the environmental resources (i.e., their efficiency, number, stability, ease of use, etc.), and S's ability to detect the presence and absence of resources and recognize their relevance to satisfaction of his desires. We might want a distinct representation of the contribution epistemic status makes to the chances of successful or adaptive interaction with an environment. Relative to this theoretical purpose, some analyses of perceptual knowledge might show themselves to be clearly superior to others.

If we are concerned to understand S's chances of success in a single case or some short run of cases, (PK) or (PK)^k seems to provide a more useful explanatory construct. For example, if we want to find out whether S's knowledge of the age of his mating partners partially explains his reproductive success, we probably would want some yes/no determination of his epistemic status with respect to each female he mates with. If, on the other hand, we are not just concerned with one individual, or one sex of a species, but, rather, all of S's conspecifics, an indication of epistemic status that failed as reliable predictor of individual success might still be useful as a predictor of species success. The suggestion here is that choices among analyses of epistemic concepts should be influenced not so much by preanalytic intuitions about the content of concepts, for there appear to be many incompatible concepts of knowledge conditioning our judgments, but on some clear understanding of what theoretical work one wants an analysis of epistemic notions to do in the first place.

8. IS THIS REALLY KNOWLEDGE WITHOUT TRUTH?

It has appeared to some, especially those who are sympathetic to causal theories of reference, that despite the fact that no explicit reference is made to truth, the above account of knowledge is tacitly committed to truth in a couple of different ways. First, the account implies that S has knowledge with respect to *o*'s being F only if *o* is F. This, plus the disquotation scheme for truth, implies that S has knowledge with respect to *o*'s being F only if "*o* is F" is true. Second, the only reason the account of the objects of belief looks plausible is that the states of affairs picked out by (OB) as objects of belief are reasonably regarded as denoted by those beliefs. Likewise, (PK) is plausible as an analysis of perceptual knowledge only because it specifies conditions under which it is reasonable to say that S's belief state has *o*'s being F as its *de re* propositional content. What we seem to have here is just a control theoretic account of reference and content, the very notions our eliminativist's arguments suggested were unnecessary for psychological explanation. It seems that the control theoretic account shows one way in which the semantic properties of internal states reduce to their causal/informational properties, not that knowledge does not involve states that have such semantic properties.

The first objection points out that the proposed analysis is compatible with a truth requirement for knowledge. This is certainly right, at least insofar as truth is understood disquotationally, i.e., as a mere formal device of semantic assent.¹⁷ The soundness of the eliminativist position poses no threat to the legitimacy of relating the knower to a true belief in this sense, for it only amounts to a convenient abbreviation for a collection statements such as, “if S has knowledge with respect to snow’s being white, then snow is white.” No semantical facts are implied by such a statement. The belief state *b* need not refer to any individual nor express any proposition for *b* to be disquotationally true. Is there an adequate *correspondence* notion of truth that will imply, along with (PK), that if S has knowledge, then there is some proposition P such that *b* expresses P and *b* is true? Clearly so if (PK) itself contains an adequate causal/informational account of reference.

Two replies are appropriate. First, that analysis is not intended to be one that is incompatible with attribution of content to internal states. It is merely intended to show that this is not necessary for the development of an epistemic theory. If, contrary to the expectations of the eliminativists, a successful account of the semantics of internal states is forthcoming, then we will know that there is no need to do what this essay argues can be done — get by without truth. But secondly, while control theory might provide an adequate general framework for the explanation of successful action, there is no reason to assume that it must do so by some single, comprehensive, and systematic correlation of internal states with external events. This form of correlation would be necessary to render plausible the claim that the belief object relation could function as a reconstruction of the reference relation. There might be a number distinct domain-specific systems of control theoretic explanation of behavior. If so, this account of knowledge will not imply the adequacy of any correspondence notion of truth.¹⁸

9. CONCLUSION

The appearance that eliminativist positions in the philosophy of psychology are incompatible with knowledge attribution seriously is a mistaken one. There is at least one form of knowledge attribution, employing a reliability conception of knowledge, that does not depend

on the attribution of truth or satisfaction conditions to internal states of knowers. I have suggested at times above that we should understand this as showing us that epistemic correspondence relations might be established in some other way than via the semantic properties of internal states. Control theory seems to provide the means for one alternative construction of this relation. I would hazard the guess that there are more.¹⁹

NOTES

¹ Paul Churchland (1979), *Scientific Realism and the Plasticity of Mind*, Cambridge: Cambridge University Press; Harry Field (1984), 'Thought Without Content,' presented at MIT-Sloan Foundation Conference, May, 1984 and Field (1986), 'The Deflationary Conception of Truth', in G. MacDonald and C. Wright (eds.) (1986), *Fact, Science, and Morality: Essays on A. J. Ayer's, Language, Truth, and Logic*, Oxford: Blackwell, pp. 55–115; and Stephen Stich (1983), *From Folk Psychology to Cognitive Science*, Cambridge: The MIT Press/Bradford Books.

² Stich (1983), pp. 170–183 and Field (1984) and Field (1986) pp. 83–85.

³ Harry Field (1978), 'Mental Representation', *Erkenntnis* 13, pp. 9–61 and Stich (1983), pp. 54–56.

⁴ Donald Davidson (1970), 'Mental Events' in Davidson (1980) *Essays on Actions and Events*, Oxford: Oxford University Press/Clarendon Press, pp. 207–227 and Davidson (1974) 'Psychology as Philosophy' in Davidson (1980), pp. 229–244. For extended critique of both forms of argument see, Jerry Fodor (1987), *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*, Cambridge: The MIT Press/Bradford Books, Chapter 3 and William Lycan (1982), 'Psychological Laws' in Biro and Shahan (eds.) (1982), *Mind, Brain, and Function*, Norman: University of Oklahoma Press, pp. 9–38.

⁵ D. M. Armstrong (1973), *Belief, Truth, and Knowledge*, Cambridge: Cambridge University Press, p. 166.

⁶ Cf. Armstrong (1973), Fred Dretske (1981), *Knowledge and the Flow of Information*, Cambridge: The MIT Press/Bradford Books, and Marshall Swain (1981) *Reasons and Knowledge*, Ithaca: Cornell University Press.

⁷ Cf. Alvin Goldman (1986), *Epistemology and Cognition*, Cambridge: Harvard University Press.

⁸ I say "roughly" because there are number of ways to relate S's epistemic reliability to the truth conditions of S's beliefs. One might require that S's beliefs about the contents of his refrigerator be always or generally true, or that S's beliefs about beer in his refrigerator be reliable, or that his belief about the location of a particular can of beer be reliable, or that his memory beliefs be reliable. These alternatives result from modulating the scope of the reliability required for knowledge. Must the belief that p be reliable with respect to its truth condition, or merely be a member of a class of beliefs generally reliable with respect to their truth conditions? For an extended discussion of this and related problems see Richard Feldman (1985), 'Reliability and Justification', *The Monist* 68, pp. 159–174.

⁹ Stalnaker (1984), p. 15.

¹⁰ This problem was pointed out to me by Lynne Rudder Baker in correspondence.

¹¹ This a case of Ted Morris's. It was offered in comments on an earlier version of the present paper presented at the Pacific Division meeting of the American Philosophical Association, March, 1987.

¹² Cf. Dretske (1981), Chapter 3.

¹³ Cf. Fred Adams (forthcoming), 'Feedback About Feedback: Reply to Ehring', *The Southern Journal of Philosophy*.

¹⁴ This not strictly speaking an adequate account of the basing relation. The problem is that a basis need not, in general, be a cause of the belief that it is a basis for. For a discussion of this problem and the adjustments needed in an adequate account of a basing to deal with it and other problems see, Tolliver (1982), 'Basing Beliefs On Reasons', *Grazer Philosophische Studien* 15, pp. 149–161.

¹⁵ Of course if some version of the causal theory of reference is true, or even approximately true, then the mechanism of that theory could be taken over and recharacterized as a causal theory of cognitive objects. I am ignoring this possibility in the context of this discussion. I am also ignoring Dretske's informational analysis of *de re* content. Both approaches seem to presuppose the adequacy of some theory of psychological content and thus run counter to my intention here to accept the eliminativist's argument at face value.

¹⁶ Thomas R. Alley (1985), in 'Organism-Environment Mutuality, Epistemics, and The Concept of an Ecological Niche,' *Synthese* 65, 1985, pp. 411–444, argues that representations of epistemic relations and capacities is essential to an adequate representation of an organism and its ecological niche.

¹⁷ For some illuminating recent discussions of disquotational truth see, Michael Williams (1986), 'Do We (Epistemologists) Need a Theory of Truth?', *Philosophical Topics* 14, Spring 1986, pp. 223–242 and Harry Field (1986).

¹⁸ For more on the role of correspondence notions of truth in the explanation of behavior see Field (1986) and Tolliver (forthcoming), 'Beliefs Out of Control', in Silvers (ed.) *Representation: Readings in the Philosophy of Mental Representation*, Dordrecht, D. Reidel.

¹⁹ Thanks are extended to Kent Bach, James Bogen, John Biro, John Heil, Keith Lehrer, Ted Morris, Kim Sterelny, and Frederick Suppe for comments on a earlier drafts of this paper. A version of this paper was produced at and presented to an NEH Summer Institute on Theory of Knowledge conducted by Alvin Goldman and Keith Lehrer in 1986. Other versions were read at the Pacific Division Meetings of the American Philosophical Association, March, 1987 and at the Thirteenth Annual Meeting of the Society for Philosophy and Psychology, June, 1987. Work on this paper was also supported by a Summer Research Grant from the University of Maryland, College Park and a Ford Foundation Fellowship for Minorities during the academic year 1987–88.

REFERENCES

Adams, Fred: forthcoming, 'Feedback About Feedback: Reply to Ehring', *The Southern Journal of Philosophy*.

Alley, Thomas R.: 1985, 'Organism-Environment Mutuality, Epistemics, and The Concept of an Ecological Niche,' *Synthese* 65, pp. 411–444.

Armstrong, D. M.: 1973, *Belief, Truth, and Knowledge* (Cambridge, Cambridge University Press), p. 166.

Biro, John and Shahan, Robert, eds.: 1982, *Mind, Brain, and Function* (Norman, University of Oklahoma Press).

Churchland, Paul: 1979, *Scientific Realism and the Plasticity of Mind* (Cambridge, Cambridge University Press).

- Davidson, Donald: 1970, 'Mental Events' in Davidson (1980), pp. 207—227.
- Davidson, Donald: 1974, 'Psychology as Philosophy' in Davidson (1980), pp. 229—244.
- Davidson, Donald: 1980, *Essays on Actions and Events* (Oxford, Oxford University Press/Clarendon Press).
- Dennett, D. C.: 1982, 'Beyond Belief,' in Andrew Woodfield ed., pp. 1—98.
- Dretske, Fred: 1981, *Knowledge and the Flow of Information* (Cambridge, The MIT Press/Bradford Books).
- Feldman, Richard: 1985, 'Reliability and Justification', *The Monist* 68, 2, pp. 159—174.
- Field, Hartry: 1978, 'Mental Representation', *Erkenntnis* 13, pp. 9—61.
- Field, Hartry: 1984, 'Thought Without Content,' presented at MIT-Sloan Foundation Conference, May, 1984.
- Field, Hartry: 1986, 'The Deflationary Conception of Truth', in MacDonald and Wright eds., pp. 55—115.
- Fodor, Jerry: 1987, *Psychosemantics: The Problem of Meaning in the Philosophy of Mind* (Cambridge, The MIT Press/Bradford Books).
- Goldman, Alvin: 1986, *Epistemology and Cognition* (Cambridge: Harvard University Press).
- Lycan, William: 1982, 'Psychological Laws' in Biro and Shahan eds.
- MacDonald, G. and Wright, C. eds.: 1986, *Fact, Science, and Morality: Essays on A. J. Ayer's Language, Truth, and Logic* (Oxford, Blackwell).
- Stalnaker, Robert: 1984, *Inquiry* (Cambridge, The MIT Press/Bradford Books).
- Stich, Stephen: 1983, *From Folk Psychology to Cognitive Science* (Cambridge, The MIT Press/Bradford Books).
- Swain, Marshall: 1981, *Reasons and Knowledge* (Ithaca, Cornell University Press).
- Tolliver, Joseph: 1982, 'Basing Beliefs On Reasons', *Grazer Philosophische Studien* 15, pp. 149—161.
- Tolliver, Joseph: forthcoming, 'Beliefs Out of Control', in Silvers ed. *Representation: Readings in the Philosophy of Mental Representation* (Dordrecht, D. Reidel).
- Williams, Michael: 1986, 'Do We (Epistemologists) Need a Theory of Truth?', *Philosophical Topics* 14, pp. 223—242.
- Woodfield, Andrew ed.: 1982, *Thought and Object* (Oxford, Oxford University Press/Clarendon Press).

Department of Philosophy
University of Arizona
Tucson, AZ 85721
USA