

Chapter 1

The Problem of Intentionality

When a person suspects or supposes, realizes, regrets, infers or imagines, doubts or discovers that something is true, he is taking an attitude or changing his attitude toward a proposition. To understand such propositional acts and attitudes, in terms of which the cluster of activities that constitute inquiry must be explained, we need to be clear about what kind of thing propositions are. Are they essentially linguistic things, or abstract objects, or mental constructs? Do they have constituents and, if so, what kinds of constituents? What are the identity conditions for propositions? That is, under what conditions do two sentences express, or two persons believe, the same proposition? What kinds of relations hold between propositions, between propositions and the sentences that express them, between propositions and the subjects of propositional attitudes? I will begin my exploration of the structure of inquiry with an exposition and defense of an account of propositions which is intended to answer some of these questions. The definition I will defend is short and simple, but it has struck some philosophers as an obscure and unhelpful metaphor rather than a real definition, and it may seem to involve problematic commitments and to have some unpalatable consequences. My defense will try to make clearer the content of the definition, as I understand it, to free it from some of the commitments and consequences that have been attributed to it, and to argue that others of its supposed commitments and consequences are more acceptable than it might seem.

The definition I will defend has its origin in formal semantics, and many have assumed that what interest it has lies in its application to technical problems in semantics. But I think a more important virtue of the definition, and of the framework that it presupposes, is the contribution they make to the philosophical understanding of the foundations of semantics, to our understanding of the notions of mental and linguistic representation. My main task in this chapter will be to tie the definition of proposition that I will defend to an independently plausible general conception of the nature of intentional mental states.

I will begin with a brief statement of the definition and of some of the problems that must be overcome if it is to be acceptable. Then I will sketch, in impressionistic terms, two contrasting pictures of mental and linguistic activity, one of which motivates the definition of propositional content that I want to defend. I will argue that only this picture of mental and linguistic activity provides us with a plausible way of solving the problem of intentionality—the problem of explaining the nature of intentional or representational mental states. I will outline a strategy for solving this problem, and then, in chapter 2, discuss critically some strategies for solving the problem in the spirit of the alternative picture of mental representation.

The analysis I will defend defines propositions in terms of possible worlds, and so one of the burdens of my defense will be to explain and justify this familiar but controversial notion. In chapter 3 I will try to explain what I think possible worlds are and why I think the kind of commitment to their existence which the definition makes is a reasonable one.

*A proposition is a function from possible worlds into truth-values.*¹ That is all there is to the definition, but to make it clear we need to say something about the terms with which propositions are defined. Discussion of the most problematic of these—the notion of a possible world—will be deferred to chapter 3, but I will make a brief remark about truth-values and functions.

There are just two truth-values—true and false. What are they: mysterious Fregean objects, properties, relations of correspondence and noncorrespondence? The answer is that it does not matter what they are; there is nothing essential to them except that there are exactly two of them. We could formulate the definition of proposition in a way that did not mention truth-values at all without changing its essential character: a proposition may be thought of as a rule for selecting a subset from a set of possible worlds. The role of the values *true* and *false* is simply to distinguish the possible worlds that are members of the selected subset from those that are not. But is there not more to truth than this? Should not an adequate theory of truth include some explanation of why curious people seek it, honest speakers aim at it, and good arguments preserve it? Shouldn't it help us to understand and solve metaphysical problems, such as disputes between realists and idealists? Somewhere in a theory of propositions and propositional attitudes such explanations must be given, but according to the account to be developed here, these questions are concerned less with truth itself than with belief, assertion, and argument, and with the relation between the actual world and other possible worlds.

A function may be thought of intuitively as a rule for determining

a value relative to any member of a specified domain of arguments. But the identity conditions for functions are purely extensional: if functions f and g are defined for the same arguments, and have the same values for each argument, then they are the same function. So a proposition is fully determined relative to a domain of possible worlds by the subset of that domain for which the proposition takes the value *true*.

Let me mention two distinctive features of this definition—features that are the source of the strengths of the account, as well as of its weaknesses. First, according to this conception, propositions lack structure of the kind that reflects the semantic structure of the sentences that express them. The account thus contrasts with accounts of content given, for example, by Frege, Russell, and Wittgenstein, according to which propositions are compounded out of individuals, concepts, properties, or senses. Second, according to this conception, propositions are defined independently of language and linguistic behavior. The definition thus contrasts with any account which tries to explain content in terms of the uses of linguistic expressions. The two most serious objections to the proposed definition derive from these two features.

It is a consequence of the first feature that propositions are individuated much less finely by this account than by contrasting accounts. The analysis implies that propositions are identical if they are necessarily equivalent—true together and false together in all possible circumstances. But, it may be objected, these are not plausible identity conditions for the objects of mental states such as belief since it is obvious that one may believe a proposition while disbelieving one that is necessarily equivalent to it. It may, for example, be a nontrivial mathematical problem to see that two expressions are necessarily equivalent, and where a person has not yet solved such a problem, his attitude toward the content of one may be different from his attitude toward the content of the other. The implausibility of the identity conditions imposed on propositions by the definition is particularly striking in the case of necessarily true propositions. Since all necessary truths are necessarily equivalent to each other, it follows that there is only one necessary truth. Since all mathematical truths are necessary, this means that there is only one true thing that can be said in mathematics, although it can be said in many different ways.

The second distinctive feature of the definition—its language independence—opens it to a second objection: the charge of ontological extravagance. It may be argued that if the account is committed to the existence of possible worlds, then it is false, since there obviously are no such things. But if the theory treats possible worlds as only a convenient fiction, then whatever its heuristic value in helping to see the

consequences of certain assumptions, it has no real explanatory power. At the very least one can reasonably insist that the elusive concept of a possible world be given more specific intuitive content before it is asked to bear the burden put on it by the proposed definition of proposition.

Both of these problems are serious ones, and I have no final and decisive solutions to them. But I will try to make a case that they are not insurmountable obstacles to the success of the account of propositional content that I will be defending.

Here is one impressionistic picture of the nature of human activities which involve mental representation—call it the pragmatic picture. Rational creatures are essentially agents. Representational mental states should be understood primarily in terms of the role that they play in the characterization and explanation of action. What is essential to rational action is that the agent be confronted, or conceive of himself as confronted, with a range of alternative possible outcomes of some alternative possible actions. The agent has attitudes, pro and con, toward the different possible outcomes, and beliefs about the contribution which the alternative actions would make to determining the outcome. One explains why an agent tends to act in the way he does in terms of such beliefs and attitudes. And, according to this picture, our conceptions of belief and of attitudes pro and con are conceptions of states which explain why a rational agent does what he does. Some representational mental states—for example, idle wishes, passive hopes, and theoretical beliefs—may be connected only very indirectly with action, but all must be explained, according to the pragmatic picture, in terms of their connections with the explanation of rational action.

Linguistic action, according to this picture, has no special status. Speech is just one kind of action which is to be explained and evaluated according to the same pattern. Linguistic action may be a particularly rich source of evidence about the speaker's attitudes, but it has no special conceptual connection with them.

This picture suggests that the primary objects of attitude are not propositions but the alternative possible outcomes of agents' actions, or more generally, alternative possible states of the world. When a person wants a proposition to be true, it is because he has a positive attitude toward certain concrete realizations of that proposition. Propositions, the picture suggests, are simply ways of distinguishing between the elements of the relevant range of alternative possibilities—ways that are useful for characterizing and expressing an agent's attitudes toward those possibilities. To *understand* a proposition—to know the content of a statement or a thought—is to have the capacity to divide the relevant alternatives in the right way. To *entertain* a proposition is

to focus one's attention on certain possibilities, contrasting them with others. To *distinguish* two propositions is to conceive of a possible situation in which one is true and the other false.

Here is a contrasting impressionistic conception of mental representation—call it the linguistic picture. Rational creatures are essentially speakers. Unspoken thought is something like inner speech—"saying in one's heart."² Representational mental states represent the world because of their resemblance to, or relation with, the most basic kind of representations: linguistic expressions.

Those attracted to this picture are inclined to say things like this: Assertion is "not the expression of an interior act of judgment; judging, rather, is the interiorization of the external act of assertion."³ "The thought that *p* is an episode which might also be referred to as the mental assertion that *p*."⁴ "A creature cannot have thoughts unless it is the interpreter of the speech of another."⁵ "Representational characteristics of mental states derive from representational characteristics of sentences of the language of thought."⁶

It is not essential to the linguistic picture that every thinking creature be capable of outward speech or that every one of our thoughts be expressible in our public language. All that is essential is that thought be explained by analogy with speech. Every thinking creature, according to this picture, does something like talk to itself in the language of thought, even if it lacks the capacity to translate its utterances into any public language. Our inexpressible thoughts are inner utterances for which no adequate translation into a public language is available to us.

Proponents of the linguistic picture are inclined to be skeptical of any notion of an abstract object of thought and speech which can be identified across languages. Particular episodes of speech and thought have sentences or sentence-analogues as their objects, and there may be various more or less indeterminate translation relations that hold between them. But translation, according to this picture, is not an attempt to preserve some independently understood property like meaning or content. Rather, translation is the notion in terms of which the vague preanalytic notions of meaning and content should be explained—or better, with which those vague notions should be replaced.

I want to emphasize that in characterizing the linguistic picture I am describing a cluster of metaphors and analogies which guide the construction of theses and theories and not attempting to describe even the broad outlines of a unified approach to the explanation of mental representation. As will become clearer in the next chapter, the development of the linguistic picture leads in two quite different directions which emphasize different analogies between speech and thought. One

hypothesizes a language of thought, which may be different from any language used for communication; the other argues for the dependence of thought on the social activities of speech. There are sharp conflicts between these different developments of the analogy.

My main task in this chapter is to bring the two pictures into sharper focus by formulating and discussing one issue that divides them. The issue concerns conflicting strategies for solving the problem of intentionality. Let me first say what the problem is and then describe the alternative strategies for solving it.

The problem of intentionality is a problem about the nature of representation. Some things in the world—for example, pictures, names, maps, utterances, certain mental states—*represent*, or *stand for*, or are *about* other things—for example, people, towns, states of affairs. Some philosophers have suggested that the capacity to represent, and to confer representational properties, is a distinctive and essential capacity of thinking things. Persons can represent because they have minds; inanimate objects can represent only because of the way people use and regard them.

For various familiar reasons, intentional or representational relations seem unlike the relations holding between things and events in the natural world: causal interactions, spatiotemporal relations, various notions of similarity and difference. One can, it seems, picture, describe, or think about such things as gods and golden mountains even if they do not exist. And one can picture, describe, or think about a triangle or a sunset without there being any particular triangle or sunset that is pictured, described, or thought about. Some philosophers have used these distinctive features of intentional relations to argue that they are irreducible to natural relations. From this conclusion it is argued that mental phenomena cannot be a species of natural phenomena. Any account of thinking things as natural objects in the material world, these philosophers argue, is bound to leave something out. The challenge presented to the philosopher who wants to regard human beings and mental phenomena as part of the natural order is to explain intentional relations in naturalistic terms.

The linguistic and pragmatic pictures each suggest strategies for giving a naturalistic explanation of representation—both mental and linguistic representation—but the two strategies differ in what kind of representation they take to be more fundamental. The pragmatic picture suggests that we explain the intentionality of language in terms of the intentionality of mental states, while the linguistic strategy suggests that we explain the intentionality of mental states in terms of, or by analogy with, the intentionality of linguistic expressions.⁷ If we opt for the former course, then to avoid circularity we need an explanation in

terms that make no reference to language of the representational character of mental states. If we opt for the latter course, then we need an explanation in terms that make no reference to mental states of the representational character of linguistic expressions. I will argue that the problem can be solved in the first of these ways and that it can be solved only in this way—that attempts to solve the problem of intentionality in the alternative way are doomed to failure. In trying to make this case, I will discuss a number of arguments from a very illuminating paper by Hartry Field, "Mental Representation."⁸ I find this paper useful for my purposes because I agree so thoroughly with the terms in which it poses the problem of explaining mental representation and because I disagree so thoroughly with the solution it offers. Field and I agree in taking the problem of intentionality to be the central philosophical problem that an account of mental representation must solve. We also agree, I believe, that the problem can be solved and that only one of the two options listed above describes a way of solving it. We disagree only about which of the two ways it is.

This informal characterization of Field's project may be misleading, so I will add a qualification. There is a sense in which Field might agree that the intentionality of thought is prior to the intentionality of language: it might be, according to Field as I understand him, that the representational properties of expressions of a *public* language are derived from, and explained in terms of, the representational properties of beliefs and intentions. But Field would insist that the representational properties of beliefs and intentions must be explained in terms of the semantic properties of a language or language-like system: mentalese, or the language of thought. Field's essential thesis is that only by assuming that mental states have something like a linguistic structure can we explain how they can represent the world. His strategy, like Jerry Fodor's, is to explain propositional attitudes as nonintentional relations to sentences of the mental language and then to explain the intentional properties of propositional attitudes in terms of the semantic properties of the sentences of that language.

Sometimes the problem of intentionality is posed in terms of the notion of *content*. Some objects, or states of objects, have content, where to have content is to be related to a proposition. Hartry Field, in the paper cited, suggested that the problem of intentionality (at least for those who take mental states such as belief at face value as relations between persons and the contents of the states) is the problem of giving "a materialistically adequate account of a relation between a person and a proposition."⁹ This is a problem, I think, and it is closely tied to the problem of intentionality, but I want to suggest that it is not quite the same problem. There may be relations between persons and

propositions which are not intentional relations—which do not themselves involve representation—and which are unproblematic from the point of view of the problem of explaining representation. Such non-intentional relations between persons and propositions are important, since one strategy for solving the problem is to reduce intentional relations (such as belief and attitudes pro and con) to nonintentional relations between persons and propositions.

How can a person be related to an abstract object such as a proposition? Let me ask, first, a parallel but easier question: how can a person be related to an abstract object such as a number?¹⁰ Note that while some relations between persons and numbers, such as those that are instantiated when I think about or name the number 42, are intentional, others are not. I take it that height in inches, weight in pounds, age in years, are all nonintentional relations between persons, or other physical objects, and numbers, and that while the existence of such relations may call for some explanation, the problem they pose is not a problem about representation. Nevertheless, I want to look briefly at the question, how can physical objects be related to numbers, in the hope that it will throw a little light, by analogy, on the question, how can persons be related to propositions, and ultimately on the question, how can persons stand in intentional relations to propositions.

The analogy between numbers and propositions, and the examples of physical quantities such as height and weight, are useful, I think, for suggesting different ways that one might understand how people could be related to propositions. Some philosophers seem to assume that we must respond in one of the following three ways to the apparent fact that propositional attitudes relate people to propositions: (1) One might deny that propositional attitudes correspond to such relations, taking the fusion or orthographic accident line. "The fusion story is the proposal that sentences like 'John believes it's raining' ought really to be spelled 'John believes-it's-raining'; that the logical form of such sentences acknowledges a referring expression ('John') and a one-place predicate with no internal structure ('believes-it's-raining'). 'John believes it's raining' is thus an atomic sentence, similar *au fond* to 'John is purple.'"¹¹ (2) One might hypothesize some kind of mysterious non-natural connection between persons and abstract objects, for example, that people have "a special intellectual capacity (*theoria*) wherewith one peers at abstract objects."¹² (3) One might suppose that people are related to propositions in virtue of being related to a sentence token, or mental analogue of a sentence token, which expresses the proposition. The first response is obviously unsatisfactory, and the second is incompatible with a naturalistic account of human beings. So we are left with the third response, in which the psychological relation between

a person and an abstract object is factored into two relations: first, a psychological relation between a person and a less problematic entity (a sentence token) which will have a physical form; second, a semantic relation between the sentence token and an abstract proposition.

The analogy between propositions and numbers suggests that there may be further alternative strategies for explaining how a person can be related to a proposition. No one would be tempted to tell the fusion story about occurrences of the numerical expression 'two hundred' in the statement 'George weighs two hundred pounds,' nor would anyone be attracted by the hypothesis that George's weight is a mysterious nonnatural relation between George and the number. Are we forced, then, to say that George must weigh two hundred pounds in virtue of containing within him something that counts as a token of a numeral denoting the number?

What is it about such physical properties as having a certain height or weight that makes it correct to represent them as relations between the thing to which the property is ascribed and a number? The reason we can understand such properties—physical quantities—in this way is that they belong to families of properties which have a structure in common with the real numbers. Because the family of properties which are *weights* of physical objects has this structure, we can (given a unit, fixed by a standard object) use a number to pick a particular one of the properties out of the family. That, I think, is all there is to the fact that weights and other physical quantities are, or can be understood as, relations between physical objects and numbers. There is, of course, much more to be said about physical quantities, for example, about what it is for properties to belong to a family and for such a family to have a certain structure. The theory of measurement provides rigorous and detailed answers to such questions. But if one were inclined to think that there is some mystery, in general, about how physical objects can stand in relation to abstract objects, the informal explanation I have given ought to be enough to dispel the mystery.

Some might be inclined to argue that the kind of explanation measurement theory gives of physical quantities shows that they are not really relational properties at all—at least not properties that relate physical objects to numbers. They are instead either intrinsic properties of the physical objects or properties defined in terms of the relations holding between the object and other physical objects. There is a grain of truth, according to this line of argument, in the fusion story since, ontologically if not semantically, physical quantities may be nonrelational properties. On some explications of the rough intuitive distinction between intrinsic and relational properties, this may be right, but the fact remains, as even a nominalist would agree, that the numerical

expression that goes in for x in a statement of the form 'George weighs x pounds' is a semantically significant constituent. It is important to keep the extremely counterintuitive *semantic* fusion thesis distinct from a much vaguer, but perhaps much more plausible, *ontological* thesis about what kinds of properties are expressed by certain semantically complex expressions. The thesis that propositional attitudes should be analyzed as relations is compelling only if it is understood as a denial of the semantic fusion thesis.

Fodor, in his defense of the thesis that propositional attitudes should be analyzed as relations, does not clearly separate the ontological from the semantic issue. Discussing an example of Dennett's, Fodor says that the surface grammar of such apparently relational expressions as 'Mary's voice' is ontologically misleading—that it does not really express a relation between Mary and an object, her voice.¹³ But I am sure that Fodor would agree that the semantic structure of 'Mary had a little voice' exactly parallels that of (one reading of) 'Mary had a little lamb,' even if voices have a more tenuous place in our ontology than lambs and even if the references to voices are more easily paraphrased away than references to lambs. Whatever a proper metaphysics might say about voices, the semantic fusion thesis is as implausible about voices as it is about weights, beliefs, or lambs.

In a postscript to a reprinting of "Mental Representation," Field discusses the analogy between propositional relations and the relations between physical objects and numbers expressed by physical magnitude terms.¹⁴ He suggests that a solution to the problem of intentionality which parallels the explanation of physical magnitudes provided by measurement theory would support his thesis that mental representation requires a system of internal linguistic representation. A parallel solution, he argues, would "have to postulate a system of entities inside the believer which was related via a structure preserving mapping to the system of propositions. The 'structure' that such a mapping would have to preserve would be the kind of structure important to propositions: viz. logical structure. So the system of entities inside the believer would have to have logical structure, and this I think means that the system of entities inside the believer can be viewed as a system of sentences—an internal system of representations."¹⁵

I agree that the relations essential to a system of propositions are logical relations, if this is understood to mean relations of entailment, compatibility, logical independence, and so forth, but I am not convinced that entities standing in these relations must be sentences, or sentence-like. According to the possible worlds analysis, propositions have logical structure but nothing corresponding to linguistic structure. Field argues that we need a more fine-grained conception of proposition, one that

would support the internal linguistic representation hypothesis, but the analogy with numbers does not, by itself, support this point. I think, however, that the analogy does point the way to a strategy for answering the question of what conception of proposition is required for an adequate account of propositional attitudes. Just as the empirical relations that fix the reference of physical magnitude terms determines which features of numbers are physically significant and which are not, so the empirical relations which a functional theory uses to explain propositional attitude concepts will determine which features of abstract propositions are significant, and so what conception of proposition is appropriate.

Before looking at propositional attitudes and mental representation, I want to consider some simpler relations between persons and propositions, relations that I think can be understood by analogy with physical magnitudes. The analogy suggests that to define a relation between a person or a physical object and a proposition is to define a class of properties with a structure that makes it possible to pick one of the properties out of the class by specifying a proposition. I will give three examples of relations that I think can be understood in this way. They are all artificial and over-simple examples, but I hope they will make the point that it is at least possible to define relations between persons and propositions that are unproblematic from the point of view of the problem of intentionality and do not require or support a hypothesis of internal linguistic representation.

First example: Consider a concept of *need* defined as follows: an organism *needs it to be the case that P* (at a certain time) if and only if the organism would survive (beyond that time) only if *P*. The concept might be generalized: it is not essential that the thing to which needs of this kind are ascribed be an organism; anything for which a notion of survival, or an analogue of the notion of survival, could be specified will be something which might be within the domain of this relation. The organism or object need not represent the propositions that define its needs, and no one would be tempted to ascribe any mysterious nonnatural states to things in need, in this sense. But this simple relation does have some of the properties that philosophers have found puzzling about intentional relations. One can need food (strictly, given my artificial definition, need it to be the case that one eats food) without there being any particular food that one needs. And an unfortunate organism may need something (may need it to be the case that it has something) which does not exist. I take the fact that the simple need relation has such properties to show, not that this relation is intentional, but rather that these properties are not restricted to intentional relations

and are not, in themselves, problematic from a naturalistic point of view.

Second example: Suppose we have an organism, physical object, or system of physical objects whose behavior is explainable on the assumption that certain of its states are equilibrium states. When the object or system is in equilibrium, it tends to remain there, while when it is in disequilibrium it tends to change in ways that bring it into equilibrium. The equilibrium state might be an internal state of the object or system, or it might be a relation between it and its environment. The object might move toward equilibrium either by undergoing an internal change or by causing its environment to change in the relevant way. Such objects and systems are, of course, familiar. Feedback mechanisms such as thermostats are examples, but so are simpler systems, about which it would be farfetched to think in teleological terms, such as a closed volume of gas in which the kinetic energy of the different parts tends to equalize. We might define a general relation between objects and propositions in order to talk about such a system: call it *tendency-to-bring-about*. It will be true that *x tends-to-bring-about* that *P* if and only if *P* is a logical or causal consequence of *x* being in its equilibrium state. Again, the relation has some of the allegedly problematic properties of intentional relations, but there is no mystery about how ordinary physical objects and systems can be related to propositions in this way.

My third and last example of a naturalistic propositional relation, suggested by an analysis in a paper by Dennis Stampe, brings us closer to a relation that looks like a relation of representation.¹⁶ Consider an object which has intrinsic states that tend, under normal or optimal conditions, to correlate with its environment in some systematic way, and where the object tends to be in the state it is in *because* the environment is the way it is. For example, the length of a column of mercury in a thermometer tends to vary systematically with the temperature of the surrounding air; the pattern of light and dark on the ground on a sunny day may correlate with the shape of the tree that is between it and the sun; a pattern of radio waves emitted by a transmitter may tend to correlate with the pattern of sounds made in the vicinity of a microphone; the number of rings on a cross section of a tree trunk may correlate with the age of the tree.

To characterize precisely the kind of situation we are considering, one would have to specify three things: first, the relevant set of alternative states of the object doing the representing; second, a one-one function taking these states into the corresponding states of the world; third, the normal or optimal conditions, or what Stampe called *fidelity conditions*. If *a* is a variable ranging over the relevant states of the

object and f is the function, then the relation that must hold between the object and its environment is as follows: for any a , if fidelity conditions obtain, then the object is in state a if and only if the environment is in state $f(a)$. Furthermore, the explanation for the correlation must be that the world's being in state $f(a)$ tends, under the relevant conditions, to cause the object to be in state a .

We could all think of examples of objects of this kind which are very reliable indicators of their environments, but we can also imagine cases for which conditions are often not normal or optimal because of distorting influences of one kind or another. In such cases, even though it is true that under the fidelity conditions, the object will be in state a if and only if its environment is in state $f(a)$, it might be often true in fact that the correlation failed to hold: that is, that the object was in state a while the environment was in some state different from $f(a)$.

Given such an object or system, reliable or not, we might define a relation, which I will call *indication*, between the object and propositions as follows: the object *indicates* that P if and only if, for some a in the relevant set of alternative states of the object, first the object is in state a , and second, the proposition that the environment is in state $f(a)$ entails that P . So, for example, if a tree trunk has 78 rings, then it indicates that the tree is 78 years old, and also that the tree is between 70 and 80 years old. This is true even if, because of a disease that infected the tree, or because of unusual climatic conditions, the number of rings fails to reflect the tree's real age.

I have not defined a general relation of indication; I have just given a schema for defining such a relation for particular kinds of objects against a particular theoretical background of the kind sketched. To make such a relation precise, one would have to spell out the fidelity conditions and characterize the relevant alternative states. But there clearly are objects and systems, both natural and constructed, which have the right properties, and this presents no problem from a physicalistic or naturalistic point of view.

Relations of indication of this kind, unlike the two earlier kinds of relations between objects and propositions that I discussed, do seem to me to be relations that it is reasonable to call representational. Stampe, in fact, argued in the paper cited that a notion like indication can provide a general naturalistic account of the concept of representation. It is certainly appropriate, and common, to use the notion of *information* to describe cases in which a notion of indication could be defined. It seems natural to say that when an object indicates that P , it contains or conveys the information that P . So we might expect notions something like indication to play a role in a naturalistic explanation of paradigm intentional mental states such as belief and attitudes pro and con.

Before going on to look at intentional mental states, let me make a few general comments about the kind of propositional states that I have been discussing. First, note that while I have suggested that my examples—at least the first two—are not intentional in the sense of representational, they are obviously all intentional in a different sense of the word, the sense sometimes called “intensional with an *s*.” Sentences of the form *x needs it to be the case that P*, for example, will create opaque contexts within the sentence that replaces *P*. All of my example relations were defined in terms of causal connection and counterfactual dependence— notions which are intentional in this related but different sense. Philosophers have, of course, found this broader kind of intentionality problematic as well, but the problems it presents are different. While some philosophers might reject, or try to explain away, causal connection and counterfactual conditionals, none would argue that intentionality, in this sense, marks a boundary between the mental and the physical. Some philosophers might argue that concepts of causal connection and counterfactual dependence have their origins in the structure we impose on the world, and not in the world itself, but even if this is so, it is the physical world to which we apply concepts which are intentional in this sense.

Second, note that the examples I discussed include both forward-looking and backward-looking propositional relations. That is, in one case a state of an object was defined in terms of what it tended to cause, while in another case a state of an object was defined in terms of what tended to cause it. (The other example—the first—is not really either of these: in this case a state of an object was defined in terms of what would tend to cause the object to be in some different fixed state.) I will argue that we need to use both forward-looking and backward-looking propositional states in order to explain intentional mental states.

Finally, let me summarize the general point that I am trying to make with the examples of propositional states of physical objects. The claim which the examples are intended to illustrate and support is that there can be relations between objects (or persons) and propositions which are unproblematic from a naturalistic point of view. If the claim is right, then we have a rebuttal to a general argument of Hartry Field’s against the possibility of using a functional analysis of mental states to solve the problem of intentionality. As I understand it, Field’s argument in broad outline is this: all that a functionalist theory of mental states does is to reduce claims about *mental* relations between persons and propositions to claims about the existence of unspecified *physical* relations between persons and propositions. It does nothing to show how it is possible for there to be physical relations between persons and

propositions. But this question—how is it possible for a person to be physically related to a proposition—is the problem of intentionality. For this reason, Field concludes, “functionalism does not either solve or dissolve [the problem of intentionality].”¹⁷ My examples are an attempt to respond to this argument by pointing to unproblematic physical relations between objects or persons and propositions.¹⁸

Of course Field is right that functionalism itself—the thesis that mental states are states of an organism or mechanism which are individuated by their function in the working of the organism or mechanism—does not, by itself, solve our problem. We need to say what the relevant functional states are, and to explain them in terms of unproblematic propositional relations of the kind I have been discussing. I will conclude this chapter by looking at a familiar functionalist strategy for explaining intentional mental states, at some problems with it, some ways of responding to the problems, and some consequences of the strategy concerning the structure of propositions.

The strategy I have in mind is the one suggested by the pragmatic picture of mental acts and attitudes. 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 a 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.¹⁹

Could this, or some more sophisticated variant of it, be an adequate explanation of the nature of belief and desire—one that points the way to a general explanation of intentional mental states? The kind of propositional relations used to explain belief and desire are like one of the examples that I claimed was an unproblematic propositional relation: the tendency-to-bring-about relation. But even if we accept that such relations as that are unproblematic, there remain some difficulties with this strategy for solving the problem of intentionality. First (a problem I will just mention), the account is obviously a gross oversimplification. We need to be sure that the qualifications and distinctions needed to make it more realistic will not smuggle in unexplained notions that are as problematic as what we are trying to explain. Second, and more serious I think, is a problem that has often been noted. The dependence of the explanation of each of the two intentional notions on the other, which is such a striking feature of the pragmatic strategy, presents a threat of vicious circularity. Is this theory simply a shell game that hides the problem of intentionality under belief while it explains desire, and under desire while it explains belief? I think there is a problem of circularity, but it will be important to look carefully at just why there

is a problem, and at just what it is. I will argue that the circularity shows, not that the pragmatic analysis is mistaken, but only that it needs to be supplemented in order to provide a solution to our problem.

One reason—not, I think, a good one—for worrying about vicious circularity in the pragmatic analysis is a verificationist worry. Because of the mutual dependence of the analyses of belief and desire, it is clear that distinct and apparently incompatible hypotheses about the beliefs and desires of an agent might both be compatible with the same behavioral data—even with all possible behavioral data. One can vary one's hypothesis about an agent's beliefs without varying the predicted behavior so long as one makes compensating changes in one's hypothesis about the agent's desires. Therefore, it might be argued, there can be no fact of the matter about which of such alternative hypotheses is correct, and so no fact of the matter about what an agent believes or wants.

If the pragmatic analysis had a behavioristic or verificationist motivation—if it were an attempt to reduce unobserved inner states to patterns of behavior—then I think this argument would show that the analysis had failed. But this is not the problem which the analysis is attempting to solve. Belief and desire are problematic, not because they are inner states which are not directly observable, but because they are intentional; the analysis is an attempt to explain the intentional in terms of the nonintentional, not an attempt to explain the unobservable in terms of the observable. So we can, without undercutting the job that the analysis is attempting to do, understand dispositional properties such as belief and desire as real causal properties of persons, and not simply as patterns of actual and possible behavior.²⁰

Imagine a machine whose inner states are inaccessible to us and which moves in certain complicated ways. On the basis of its behavior, we formulate a hypothesis that its movements result from two separate but interacting mechanisms inside. Someone else formulates a rival hypothesis with different interacting mechanisms which produce exactly the same behavior under all possible circumstances. Obviously we could not tell, simply by observing behavior, which hypothesis was correct. Perhaps we could not tell at all. But there would still be a fact of the matter about how the machine worked. There could be a fact of the matter, even if the mechanisms hypothesized were defined in abstract functional terms—in terms of their role in producing the behavior.

The pragmatic analysis, with dispositional states understood as real causal properties, treats a rational agent as something analogous to such a machine. In ascribing beliefs and desires to a person, we not only make conditional predictions about how the person will behave; we also commit ourselves to claims about the kind of mechanisms

which explain why a person behaves the way he does. The mutual dependence of belief, desire, and action is a reflection of the fact that the hypothesized explanation says more than the conditional predictions which it entails. This fact makes it harder than it might be to verify claims about beliefs and desires, but in itself it presents no conceptual problem to the analysis as an explanation of intentionality.

But there is, I think, a conceptual problem: a fatal relativity in the propositional relations defined by the pragmatic analysis which shows that this analysis cannot, by itself, solve the problem of intentionality. Let me use an example to bring out the problem.

Mary is angry at Fred, her neighbor. She wants him to suffer, and believes that he will suffer if she plays her cello badly at three o'clock in the morning. So she does play her cello badly at three o'clock in the morning. That, at least, is one hypothesis for explaining why Mary did what she did. Here is another: Mary wants *Albert* to suffer, and believes that *Albert* will suffer if she plays her cello at three in the morning. *That* is why she did what she did. Now one might find the second hypothesis less plausible—even perverse—since Mary has no reason to want Albert to suffer; she has never met or heard of him. And she has no reason to believe that playing her cello badly will cause him to suffer, since he lives 3,000 miles away. Suppose the defender of the perverse hypothesis, when pressed about the implausibilities in his explanation, elaborates his hypothesis by saying that Mary believes Albert, rather than Fred, to be her neighbor, believes that Albert, rather than Fred, insulted her, believes that Albert's name is "Fred." In fact, all the attitudes that a sensible observer would say Mary takes toward Fred, the defender of the perverse hypothesis says that Mary takes toward Albert. The *only* difference between the two proposals, let us suppose, is that in the perverse hypothesis, Albert is everywhere substituted for Fred. The two hypotheses will, of course, predict exactly the same behavior, but there is also a stronger equivalence between them. Not only do belief and desire interact to produce the same actions, according to the two hypotheses, but also there is an exact correspondence between the beliefs hypothesized and the desires hypothesized by the two competing accounts. So not only are the two accounts equivalent with respect to the behavioral phenomena, they are also equivalent with respect to the mechanisms they postulate to explain the phenomena. The shift from Fred to Albert looks, from the point of view of the pragmatic analysis, like an innocent shift in the conventional units used to describe Mary's attitudes and relate them to each other, and not a shift in the claims made about the attitudes themselves.

The same point will hold for any such substitution, not only of

individuals for individuals, but of properties for properties, or whole propositions for whole propositions. All that is required is that certain internal structure be preserved. And we need not substitute *persons* for persons: we might take Mary's attitudes toward Fred to be attitudes toward a mathematical point, or a class of events in her brain, without affecting the substance of the explanation of her behavior.²¹ The *content* of belief and desire cancels out on the pragmatic analysis. Even if that analysis does give us an account of the structure of explanations of rational action, by itself it gives us no account at all of how beliefs and desires can represent the world.

It is obvious from our example what got left out. Our mental states represent what they represent not only because of the behavior they tend to cause, but also because of the events and states that tend to cause them. The reason that Mary's emotions and beliefs are directed at Fred, and not at Albert, is that Fred, and not Albert, caused Mary to be angry, and to have the beliefs that Mary has which the sensible hypothesis says are beliefs about Fred. That's not quite right. Fred might be the innocent victim of a misunderstanding, or of a malicious practical joke played by Albert. But it is clear, as has been emphasized by most recent work in the theory of reference, that the fact that Mary's attitudes are attitudes toward Fred is to be explained, somehow, in terms of Fred's causal role in producing Mary's attitudes. The total relativity of content that is a feature of the pragmatic analysis is the result of ignoring this essential element of mental representation.²²

The pragmatic analysis tries to explain mental representation entirely in terms of forward-looking propositional relations such as the tendency-to-bring-about relation. The relativity of content that results forces us to recognize that belief is a backward-looking propositional state. What I want to suggest is that belief is a version of the propositional relation I called *indication*. We believe that *P* just because we are in a state that, under optimal conditions, we are in only if *P*, and under optimal conditions, we are in that state because *P*, or because of something that entails *P*. But a causal account of belief—an account that treats belief as a kind of indication—cannot, I think, *replace* the pragmatic analysis, it can only supplement it. For an account of belief must explain, not only how belief can represent the world, but also what distinguishes *belief* from other kinds of representational states. Consider the retinal images that form on the eye. That is a kind of indication; the state of the retina can be described in terms of a relation between a person and a proposition. And if a bald head is shiny enough to reflect some features of its environment, then the states of that head might be described in terms of a kind of indication—in terms of a relation between the person owning the head and a proposition. But no one would be

tempted to call such states belief states. The reason is, I think, that they are not connected in the appropriate way with tendencies to produce action.

If belief is a dispositional state of the kind postulated by the pragmatic analysis, and also a kind of indication, then we have a fixed point with which to break into the circle that is responsible for the relativity of content. Beliefs have determinate content because of their presumed causal connections with the world. Beliefs are *beliefs* rather than some other representational state, because of their connection, through desire, with action. Desires have determinate content because of their dual connection with belief and action. Both the forward-looking and the backward-looking aspects of these mental states are essential to the explanation of how they can represent the world.

Might such an account be right? It is, I think, intuitively clear that however often we may fail to act according to our beliefs, there is a presumption that we do. Where people don't do what is appropriate, given their beliefs, we expect there to be some explanation for this; we may appeal, for example, to incapacity, absentmindedness, or self-deception. It is also intuitively clear, I think, that there is a presumption that people's beliefs will correlate with, and be caused by, their environments. Where beliefs are false, or only accidentally true, we also expect some explanation for the deviation from the norm: either an abnormality in the environment, as in optical illusions or other kinds of misleading evidence, or an abnormality in the internal belief-forming mechanisms, as in wishful thinking or misremembering. These intuitions suggest that we do have the conceptions of normal or optimal conditions which make it possible to understand belief as a kind of indication, and belief and desire together as a kind of tendency-to-bring-about, and so might make it possible to explain the intentionality of such states in naturalistic terms.

I have tried to describe a strategy for analyzing intentional mental states in a way that is motivated by the pragmatic picture, and that solves the problem of explaining how mental states can represent the world. I recognize that what I have offered is only a strategy—only the bare outlines of an account of intentional mental states. But there is enough in this outline of a strategy to allow us to draw some philosophical conclusions about the notions of mental representation and propositional content which are implicit in any way of carrying it out.

First, let me make some remarks about the causal dimension of the causal-pragmatic strategy and causal theories of reference. The strategy suggests that if reference is a causal relation, it is because belief and intention are causal relations, and because reference is to be explained in terms of the intentions and beliefs of speakers. The argument from

the relativity of content was an attempt to show that, at least within the framework of the pragmatic picture, belief and desire *must* be considered partly in causal terms. If this is right, then this framework implies that the claim that some kind of causal theory of reference is correct is not a thesis about how language happens to work, but a thesis about how it has to work if reference is to be a device for representing the world.

The hypothesis that belief should be explained in terms of indication does not directly imply any particular causal account of reference. The indication account of belief explains representation in terms of what *would* cause the believer to be in a certain state under certain possibly counterfactual conditions rather than in terms of what does in fact cause the believer to be in that state. In any case, one would need an explanation of reference in terms of belief and intention in order to make the connection. But the indication account of belief shares with causal theories of reference the rejection of the idea, which I think is implicit in earlier accounts of reference, that representation is an internal matter: that one's words and thoughts represent in virtue of the intrinsic properties of speakers and thinkers. Both kinds of accounts argue instead that what we mean, and what we believe, is in part constituted by facts about the environment in which our thoughts and linguistic acts occur, the facts that help to explain why we have the thoughts and say the things we do.

The theses that names refer in virtue of causal connections with things, and that the meanings of certain common nouns "ain't in the head," were initially defended with thought experiments about particular examples rather than with general arguments about the nature of representation, and I think this has led some observers to draw the wrong conclusions from them. I want to comment on two assumptions that have been made about causal theories of reference: that a causal theory of reference requires an atomistic account of representation and that a causal theory of reference applies only to a limited range of expressions or concepts. Both, I think, are mistaken.

Donald Davidson contrasts

... two approaches to theory of meaning, the building-block method, which starts with the simple and builds up, and the holistic method, which starts with the complex (sentences, at any rate) and abstracts out the parts. . . .

The two approaches are, I think, naturally associated with two views of proper names. With the building-block approach goes the causal theory of proper names. . . .

The other view holds that interpreting the sentences (and hence,

by abstraction, the names) used by a speaker depends solely on the present dispositions of the speaker (or a community of speakers) and so the causal history of names is strictly irrelevant.²³

I think that both of Davidson's contrasts are useful, but I object to the way they are associated. There is no reason why a theory of meaning, or an account of the source of intentionality, cannot be both holistic, explaining representation on the level of propositions, or perhaps even larger representational units, and also causal, explaining representation in terms of causal relations between the agent or a community of agents and the world. An explanation of representation in terms of indication is an example of a kind of account which combines these two approaches.

The most striking of the examples used to defend causal theories of reference have involved proper names and common nouns denoting natural kinds, and some philosophers have assumed that only such expressions should be explained in causal terms. Nathan Salmon, for example, contrasts *directly referential* terms with those that are *descriptive*. He suggests that Putnam's thesis that "meanings ain't in the head" applies only to the former kind of term. For the latter, "whose senses *do* consist solely of general properties, grasping the sense of the term is a wholly internal psychological state."²⁴ Purely general terms, Salmon assumes, are descriptive in this sense, and so one can represent them in virtue of being in a purely internal psychological state. Salmon takes Putnam's twin earth argument to show that natural kind terms such as "water" and "tiger" are not purely general, but he seems to assume that other terms, perhaps qualitative predicates such as "yellow," are immune to such arguments. This seems to me a mistake. If representation is essentially a causal relation, then no predicate, and no mental state, can represent in virtue of the intrinsic psychological properties of the person who is using the predicate, or who is in the mental state. Purely general properties may still be properties of things in the world, and representing such properties requires interaction with such things.

Second, I want to comment on the consequences that the strategy I have outlined has for the relation between thought and language. The strategy points toward an explanation of attitude and content which makes no essential reference to language. If the strategy can be carried out, then we will have a foundation for the kind of semantic theory that explains the meaning and content of linguistic expressions in terms of the intentions and conventions of language users, and which explains what it is to understand a language in terms of the capacities of speakers and hearers to use the language to serve their needs and desires. But

although linguistic representation is not essential to the kind of account the strategy points to, it is essential that there be some form of internal representation in any creature that is correctly said to have beliefs and desires. I emphasized, in discussing the threat of circularity in the pragmatic analysis, that this account, as I understand it, is not behavioristic in the sense that it identifies mental states with patterns of behavior. In attributing beliefs and desires, we are attributing certain kinds of internal causal properties which have a structure that tends to reflect the world in ways that make it appropriate to call them representations. These representations could conceivably take the form of sentences of a language of thought written in the belief center of the brain, but they also could take the form of pictures, maps, charts or graphs, or (most plausibly) a diversity of redundant forms, none of which are very much like any of the forms which our public representations take.

The pragmatic picture and the possible worlds definition of proposition does not then deny that beliefs are internally represented. But it remains neutral on the form that those representations must take. It should be emphasized that the possible worlds conception of proposition does not present an account of the form in which propositions are represented which is a rival to a linguistic account. Accepting this account of propositions does not, for example, commit one to a psychological hypothesis that our minds represent a space of possible worlds point by point, with individual representations of individual worlds. The aim of the definition is to give an account of the structure of what is represented while leaving open questions about the means by which this is accomplished.

One could stretch the concept of language to include all the possible forms of representation, and in this way reconcile the kind of account I am promoting with the thesis that internal *linguistic* representation is essential to mental representation. Some of Gilbert Harman's remarks, in his book *Thought*, seem to be suggesting this kind of reconciliation. On the one hand, Harman accepts a functionalist account of intentional mental states similar to that suggested by the pragmatic picture. "Mental states and processes," he argues, "... are constituted by their function or role in the relevant program. To understand desire, belief and reasoning is to understand how desires, beliefs and instances of reasoning function in a human psychology."²⁵ But on the other hand, he argues that "we can simply take mental states to be instances or 'tokens' of appropriate sentences of a language of thought."²⁶ Harman's motivation for taking the functional states that are states of belief and desire as tokens of sentences is that such states, in order to play their functional role, must have a structure which parallels the structure of sentences.

In particular, the contents of attitudes must be things which can stand in logical relations such as entailment and incompatibility, which can be negated and conjoined and disjoined with each other, which can be *about* things, which can be singular or general. Are not sentences things that have just the right properties, and so can they not provide a useful model for the contents of propositional attitudes? We do, after all, use sentences to express the contents of mental states, and as Harman remarks, "this connection between states and sentences is no accident."²⁷ One might argue that to talk of a language of thought is to do no more than to talk of a means of representation appropriate to explanations of rational behavior in accordance with the pragmatic picture.

Even if linguistic representation is construed broadly so that any reasonably complex system of internal representation constitutes a language of thought, there remains an important difference between the pragmatic and linguistic pictures. According to the kind of account I have outlined, the *form* in which beliefs and desires are represented is not essential to their content. Two different agents might have the same beliefs even if the forms in which the beliefs are represented are radically different. The conceptual separation between form and content is, I think, the central feature which distinguishes the conception of thought implicit in the pragmatic picture from the one implicit in the linguistic picture.

The aim of the possible worlds definition of proposition is to assign to the *contents* of representations just the structure that is motivated by the pragmatic account of the functional role of representations. Propositions, defined this way, *are* like sentences in some ways: for example, they stand in entailment relations, can be related as contradictories or contraries. But they do not have constituents which correspond to the semantically simple constituents of sentences, and do not have an analogue of grammatical structure. If our internal representations have such structure, this is, the pragmatic account implies, not a matter of *what* is represented but of *how* it is represented.

It is essential to rational activities such as deliberation and investigation that the participants represent alternative possibilities, and it is essential to the role of beliefs and desires in the explanation of action that the contents of those attitudes distinguish between the alternative possibilities. The particular ways in which alternative possibilities are represented, or the particular means by which distinctions between them are made, are not essential to such activities and explanations, even if it is essential that the possibilities be represented, and the distinctions be made, in some way or other.

The conceptual separation between form and content is reflected in the identity conditions for propositions which the causal-pragmatic

analysis implies. Whatever propositions are taken to be, and however we make precise the propositional relations of indication and tendency-to-bring-about in terms of which the analysis explains belief and desire, it is clear from the general schemas for the definitions of those relations that the following will be true: if the relation holds between an individual and a proposition x , and if x is necessarily equivalent to proposition y , then the relation holds between the individual and y . This implies that the thesis that necessarily equivalent propositions are identical—the main substantive consequence of the possible worlds analysis of propositional content—is a thesis that is tied to, and motivated by, the causal-pragmatic explanation of intentionality. This does not show that the identity conditions are right, or solve the problem of equivalence that these identity conditions create, but it does show that they, and the possible worlds analysis of proposition that goes with them, have a deeper philosophical motivation than has sometimes been supposed. If this definition had been proposed simply as a technical apparatus meant to systematize brute intuitions about the structure and identity conditions for objects of belief, then the examples of necessary truths and other nontrivial equivalences would show that the definition had missed the mark. The proper response would be to replace the technical apparatus with one that could make finer discriminations between the contents of attitudes and expressions. But since we have an argument to show that the identity conditions are right, as well as examples that seem to show that they are wrong, the proper response is not so clear.

The problem is, I think, that the alleged counterexamples are not just counterexamples to a particular analysis, but cases which are problematic in themselves. We lack a satisfactory understanding, from any point of view, of what it is to believe that P while disbelieving that Q , where the ' P ' and the ' Q ' stand for necessarily equivalent expressions. One can understand easily enough what it is to *assent* to a statement while dissenting from an equivalent one, but belief cannot be explained in terms of assent, among other reasons because one can assent to a statement without understanding it correctly. In order for a person's assent to a statement to show that he believes what it says (and not just that he believes that whatever it says is true), he must know what it says. What is unclear is how to explain *knowing what a statement says* in a way that does not have the consequence that a person knows what both of two necessarily equivalent statements say only if he knows that they are equivalent.

Could we escape the problem of equivalence by individuating propositions, not by genuine possibilities, but by epistemic possibilities—what the agent takes to be possible? This would avoid imposing implausible identity conditions on propositions, but unfortunately, it would

also introduce intentional notions into the explanation, compromising the strategy for solving the problem of intentionality. If belief and desire are to be explained in terms of naturalistic relations such as indication and tendency-to-bring-about, then the possibilities used to individuate propositions must be the ones that are relevant to these relations, and these clearly must be genuine, and not merely epistemic possibilities.

The problem of equivalence is part of a broader problem concerning deductive reasoning. The problem is to explain how it is possible for the conclusion of a deductive argument to contain any information not already contained in the premises and, as a special case of this, how it is possible for a necessary truth to contain any information at all. An answer to this question is needed to explain how drawing deductive inferences can be a way of increasing one's knowledge, and how knowledge of necessary truths can be knowledge at all. The problem does not arise from any easily identifiable philosophical dogma which might be given up to avoid it. It is true that it was empiricists who explicitly drew the conclusion that necessary truths and deductive inferences were empty of content, forced to this conclusion, apparently, by the doctrine that all knowledge has its source in sense experience, together with the belief that mathematical truths are not confirmed or refuted by sense experience. But the conclusion really derives not from any substantive assumption about the source of knowledge, but from the abstract concept of content or information. The difficulty is, I think, that any way of conceiving of necessary truths as having content is at the same time a way of conceiving of them as contingent—as one way things could have been among others. This is, I think, because we do think of content and information in terms of alternative possibilities. Whether the source of my information is my senses, authority, or a faculty of intellectual intuition with access to a Platonic realm of abstract entities, its deliverances are not news unless they might have been different.

The problem of deduction may ultimately be a reason for rejecting our intuitive notion of content, and the possible worlds definition of proposition which expresses it, but before making this move, we should have an alternative conception, and an alternative strategy for solving the problem of intentionality. In the next chapter I will consider some alternative ways of explaining the notion of mental representation. Then in chapters 4 and 5 I will return to the problem of equivalence, and to the question of how one can have inconsistent beliefs or fail to believe the consequences of one's beliefs. I shall try to show that there is at least the possibility of reconciling the apparent facts of deductive ignorance and inquiry with the pragmatic picture, and with the conception of content which I am defending.

