

VISUAL CONTENT: WIDE OR NARROW?

Fred Dretske

For present purposes I shall mean by *content* what Paul Grice meant by non-natural meaning. This is intended to include the propositions specified in the complement clause of propositional attitude attributions. Content, in this case, is what one believes, expects, hopes for and intends. It also includes what representations, of whatever kind, represent to be the case. So, for instance, a thermometer represents the temperature—that it is, say, 42° F—and maps represent the geography of a given area—that there is a river that runs along the road. That there is a river along the road and that the temperature is 42° F are the contents of these representations. This characterization is intended to distinguish content from what Grice called natural meaning (and I have called *information*), what an event or condition means or indicates in a way in which the associated proposition (what is meant or indicated) must be true. Tom's frown cannot mean, not in the natural sense, that Tom is displeased if Tom isn't actually displeased. That he is displeased may be what a frown from Tom *usually* means or what it *would* mean if circumstances were normal, but if Tom isn't displeased, that is not what his frown means. Not in the natural sense. Not this time. It may turn out, though, that by prior agreement and as a signal to friends, that is what a frown from Tom is made to mean. If so, then in deliberately frowning Tom is now "saying" something. He is telling his friends that he is displeased. Tom's frown now means, in a non-natural sense, that he is displeased, and it means this whether or not he is displeased. The frown has content.

I am an externalist about content. I think all content is wide. It is an extrinsic, a relational property. If Tom's frown has content, if it means that he is displeased whether

or not he is displeased, it is because Tom and his associates have, by prior arrangement, given it that meaning. Without these prior arrangements, and the existence of the events and individuals implied by such a prior understanding, Tom's frown might still mean, in a natural way, it might still indicate, that he is displeased, but it couldn't mean this when he isn't displeased. Something meaning P whether or not P is true requires the existence of objects other than the object (event, condition, or state) that means P. In this respect, having content is like being a cousin: it requires the existence of other things.

I am convinced, in fact, that meaning is a special kind of relational property. An object acquires this property, it gains content, in virtue of having a certain kind of history. In this respect it is like being divorced. To be divorced you must have been married. You can, perhaps, create a monster in the laboratory, but you cannot create a divorced monster there. According to this conception of content, content is extrinsic not only to currently *local* conditions but to currently *global* conditions as well: it is not only possible for physically indistinguishable individuals to have different content, it is possible for them to have different content while inhabiting currently indistinguishable worlds.

What I want to do here is say why I think this is so, why I am an externalist of this sort about content. I also want to acknowledge the problems that this form of externalism has and explain why (and how) I prefer to struggle with these problems rather than accept less radical forms of externalism or, even worse, internalism. Historical externalism, as I call my own version of externalism, has its problems, yes, but everybody's got problems, and, frankly, I don't see greener grass in any other yard.

1. Kinds of Visual Content

According to its title, this conference is about content in vision and language. I confine my attention to mental content and, in particular, content in vision. I leave issues about language to others. Even when we restrict our attention to vision, though, there is a source of misunderstanding that should be addressed at the outset. If we distinguish, as I think we must, seeing objects from having beliefs about them, then by visual content one can mean either: (1) the content associated with one's visual experience of an object, the kind of experience required to see an object; or (2) the content of the beliefs one ends up having about these objects. As a result of seeing *x* and *y*, I believe that *x* is longer than *y*, and this belief or judgment can certainly be false, so there is, undeniably, content in our perceptual judgments about the lines we see, but there is also, arguably, content in the visual experience of the lines that is prior to and independent of these judgments. The lines look a certain way. I know they are the same length (I just measured them), but one line nonetheless looks longer than the other, and *looking longer* is a fact about how visual experience, not thought, represents the two lines. If this is so, then perceptual experience has a content of its own that is distinct from the cognitive content of perceptual belief and judgment.

Though I accept something like this account of vision, I don't want to defend it here. I do think it important, though, to keep this distinction in mind when talking about visual content. For some people are convinced that although the content of perceptual belief is broad, the content (if there is any) of the perceptual experiences that cause these beliefs is narrow. A perceptual judgment made by a person looking at Lake Michigan--that there is water in front of him--has (according to externalists) broad content, but the

person's visual experience of the water, according to some of these same externalists, is narrow. Although Earthling and Twin Earthling have different beliefs about the liquid they see, their visual experience of the liquid they see represents this liquid in exactly the same way. Their conscious experiences of it are the same. The intuitions behind these judgments seem to be that sensory or phenomenal affairs supervene on the constitution of an agent although cognitive content, for reasons Putnam and Burge have given us, do not. If you accept this picture of things, you are both an externalist and an internalist about visual content: an internalist about the content of visual experience, an externalist about the content of perceptual judgment. Some visual content is narrow, some broad.

I am, as I said, an externalist about all content. If visual experience represents the objects of vision in a way independent of the way they are represented in thought, I am an externalist about the content of both perceptual experience and perceptual judgment. All content is broad. Others (Gabriel Segal?) may be internalists about both. Whatever you are, we have to be sure we are all talking about the same thing in the rest of this session.

2. Supervenience

When I say I'm an externalist about content, that I think content is broad, I do not mean to deny that it might supervene on something narrow. Content might be broad, it might be a genuine relational property, requiring the existence of other things, and, yet, it might supervene on the intrinsic properties of objects that have this content. Objects that differ in content might always differ in their material constitution. Even if this were so, it would not show that content was narrow. It would only show that differences in content are reflected in an object's non-relational properties.

Age, for example, is a relational, a historical, property of things. You can't be thirty years old without enduring through thirty cycles of whatever periodic processes (presumably external) mark off years (like circuits of the earth around the sun). Yet, at least in the case of trees, age supervenes on the number of rings in the trunk of the tree. Thirty years old, thirty rings. Twenty years, twenty rings. No difference in age without a corresponding difference in the number of rings. Perhaps, for all I know, this is a law of nature. At least it is about as close to a law as we are likely to get in the special sciences. So, at least in the case of trees, the supervenience of age on internal constitution is nomic. Yet, this fact doesn't show that age is a narrow, intrinsic, or non-relational property. No one, I hope, would propose a narrow theory of age—not even for trees.

If the supervenience of A on B is of the right sort, supported, let us say, by natural law, one can determine (find out, come to know, discover) the value of the supervenient property, A, by looking at B, the property on which it supervenes. If you want to know how old the tree is, count the rings. They will tell you how old the tree is. So also with mental content. It may turn out that mental content, the content of both experience and thought, supervenes on the intrinsic properties of agents in such a way as to enable us to tell (someday) what someone is thinking or feeling by looking at their nervous system or, perhaps, just their brain. Even if things turn out that way, it won't mean that meanings are in the head anymore than the corresponding fact about tree rings shows that time is nestled inside the bark of trees. Even if there were a universal law of radiocarbon dating that enabled us to discover the age of *any* object by determining the amount of chemical X it contained, this would not, or should not, tempt one into thinking that time or age is X-ishly intrinsic.

In his book, *A Slim Book About Narrow Content* (MIT Press 2000), Gabe Segal understands this point perfectly well. His argument is that mental content supervenes on the intrinsic microstructure of thinkers. As an argument for narrow content, he realizes that this works only in so far as the supervenience in question is of a special kind, one in which the supervening properties are *explained by* the intrinsic properties on which they supervene. This is the kind of supervenience that the temperature of a gas has to the kinetic energy of the gas's component molecules. This is not the kind of supervenience that a tree's age has to the number of rings in its trunk. You can explain why this gas is hotter than that one by talking about the greater energy of its molecules, but you cannot *explain* why this tree is older than that one by appealing to facts about the number of rings in their trunks. The age of trees is irreducibly relational even though it supervenes, perhaps even lawfully supervenes, on intrinsic facts about trees. If we are going to use facts about supervenience to argue for a theory of narrow content, we must be careful to argue for a very special kind of supervenience. It is this special kind of supervenience I doubt is possible in the case of visual content. Even if content supervenes on something intrinsic, it remains irreducibly relational. You will never explain why events occurring in Clyde's cortex are about football rather than philosophy by describing their electrical-chemical properties. Or so I, an externalist, believe.

3. The Source of Content.

Measuring instruments are representational devices and are, therefore, revealing examples of content-bearing systems. These instruments are supposed to indicate, tell us, how things stand in parts of the world we need or desire to know about. They represent,

sometimes misrepresent, the quantities they have the function of conveying information about. Their power to misrepresent—and, thus, their capacity to harbor content—derives from their informational function, a function they get from us: their designers and users. Without us these objects lack content. They don't say anything, at least nothing that can be false. Like the mercury in the glass tube of an ordinary household thermometer, metal paper clips and flag-poles also expand and contract as temperature varies, thus providing information (natural meaning) about temperature. Unlike a thermometer, though, paper clips and flagpoles (in their usual deployment) do not have the function of providing this information. They have different functions. As a result, defective paper clips and broken flagpoles cannot, while badly constructed thermometers can, misrepresent temperature. That is why thermometers have content while paper clips and flagpoles do not.

I say that measuring instruments are “revealing” examples of representational systems because I take it that what is true of them is true of *all* representational systems: they get their power to represent, their power to say that P whether or not P is true, and thus their capacity to embody content, from their informational function, a function that defines what it is they are supposed to be doing. If an object is supposed to indicate what the pressure (time, distance, speed, etc.) is, if that is what it is there for, then when it fails to perform this function satisfactorily, it misrepresents what the pressure (time, distance, speed, etc.) is. If, on the other hand, there is nothing an object is supposed to be doing, there is nothing it can possibly do that is wrong, mistaken, incorrect, or false. It cannot *misrepresent*. It cannot *say* P when not-P. The source of the norms inherent in the idea of (mis)representation comes from what an object is supposed to be doing and what a

thing is supposed to be doing comes from—where else?—its function or purpose. And function and purpose come from—where else?—history.

Misrepresentation, the sort of things you must be able to do to qualify for content, is merely one of a family of notions that embody this normative flavor and, thus, require of the relevant objects a certain sort of history. To say that a mechanism or bodily organ is defective, for instance, is to say that it is not in the state or condition it is supposed to be in where the condition it is supposed to be in is defined (in the case of bodily organs) by its evolutionary history. And so it is with *injured, healthy, strained, stretched, diseased, flawed, ill, sick, damaged, spoiled, ruined, marred, contaminated, defiled, corrupted, infected, malformed*. Nothing can be any of these things unless it has a certain kind of history, a development that determines how it is supposed to be or behave. In the case of artifacts the standards that define how the thing is supposed to work, come from us, its designers and builders. In the case of bodily organs and mechanisms, the relevant history is independent of our intentions and purposes. It derives from the natural history of the object. That is why a child can be born with a broken leg or faulty vision, but something that materializes randomly out of cosmic dust, even if it happens to be indistinguishable from a broken-legged infant, does not have a broken leg. If you or I deliberately create a thing the way it is as a prototype, in order to illustrate a point, or as a work of abstract art, or if it materializes randomly that way out of cosmic dust, it isn't defective no matter how much it looks and behaves like an object that (with a different history) is defective. The standards implied in describing a thing as defective, malformed, diseased, or damaged come from its history. For this reason, such things (just like divorced people) *have to* have a history. If they don't, they aren't defective, broken, diseased, or malformed. This

is why Swampman, a molecular duplicate of Donald Davidson, could not have a broken leg or defective vision even if Donald Davidson did. This, too, I submit, is why there is nothing in Swampman, as there was in Davidson, capable of misrepresenting anything. To misrepresent you have to be in a state that, given your function, you are not supposed to be in, and, lacking a history, there is nothing in Swampman that has a function. There are things in Swampman that happen to be functioning in the way things in Davidson are functioning, yes, but that doesn't give them that function. As far as I can see, nothing does.

This, in brief, is why I'm an externalist—in particular, a historical externalist--about content. I don't understand where else the power to mean P when P isn't so can come from. If someone will explain to me where, besides history, this power comes from, I'll stop being an externalist.

4. Problems

Externalism has problems and these problems constitute the strongest argument I know of in support of internalism. So let me say, briefly, why two of these arguments, though strong, are not, for me, strong enough.

A. Epiphenomenalism: There is, first, a problem about the causal or explanatory relevance of externally grounded content. Causality seems to be a local affair and, thus, causal explanations are explanations that necessarily advert to the intrinsic properties of the causally involved objects. If two objects are indistinguishable, if they share all their intrinsic properties, their causal impact on the world should be the same. For practical exchange purposes, for example, perfect counterfeit money is as good, as effective, as

real money. Its effect on the rest of the world—bankers, cashiers, and vending machines—is restricted to the intrinsic properties of the paper and since perfect counterfeit has the same intrinsic properties as real money, it has the same effects as real money. Different histories, the histories that make one piece of paper legal tender and the other counterfeit, are completely screened off by the currently observable properties of the paper. If you want to know why the paper had the effect it had, look to its intrinsic, not its historical, properties since anything with the same intrinsic properties, whatever its historical properties happen to be, will have the same effects.

As a result, externalistic theories of content have a problem explaining how what we want, believe, intend and prefer—the content of our internal states--explains what we do. For externalists, content seems as explanatorily relevant to the behavior of animals as is the monetary value of money to the behavior of vending machines. It isn't, as we all know, the monetary value of the money we put in vending machines that explains why these machines behave the way they do. When they give you a coke in response to your insertion of a dollar, it is the intrinsic properties of the paper you put in the machine— its size, shape, weight, markings--that explains the machine's behavior. So too, it seems, it is the chemical and physical properties of our internal states, not extrinsic content, that explains why we move our arms and legs the way we do. If, then, content is an extrinsic (or, even worse, a historical) property of the events occurring in our brains, this content (just like the monetary value of the money in vending machines) is causally impotent. Externalism gets you epiphenomenalism in one quick step.

This is a serious problem. I have grappled with it (in *Explaining Behavior*, 1988) years ago and I won't bore you with the details of my answer. A quick distinction and an

example will have to suffice give you the idea of how I think the relational (in particular, historical) properties of internal states are explanatorily relevant to explaining behavior.

First the distinction that should be familiar to everyone: the difference between a person moving his arm and the movement of his arm. His arm can move without his moving it. The movement of his arm is a result of his moving his arm, but it is not the same as his moving his arm. Moving one's arm is a piece of behavior, something one does, possibly an action (something one does deliberately). The movement of one's arm is not. What, then, is the difference?

This question is important because we expect content—what we believe, intend, and desire—to explain our behavior, why we move our arm, not (at least not necessarily) the result of that behavior—the movement of the arm.

If one identifies behavior, moving your arm, with something inside you making (causing) your arm to move, then explaining behavior, why you moved your arm, is not the same as explaining why your arm moved. Perhaps the causal explanation of why your arm moved must advert to the intrinsic properties of those neurological events that caused your arm to move (thus making neuroscience the relevant authority on why your arm moved), but explaining why these internal events caused your arm to move (which, on the present proposal, is identified with you moving your arm) may require appeal to a quite different explanation—perhaps, even, extrinsic considerations. It is the same with coke machines. If you want to know why a coke came tumbling down the chute when you deposited a \$1 bill, look to the intrinsic properties of the paper you deposited in the machine (and, of course, the present internal mechanical configuration of the machine). If you want to know, on the other hand, why the machine behaved that way when you put a

dollar in it, an extrinsic explanation seems entirely appropriate. Paper with those intrinsic properties caused the machine to release its content because objects with these intrinsic properties have monetary value. It is the fact that (barring extensive counterfeiting) monetary value supervenes on paper of this sort that explains why paper of this sort causes machines to dispense their contents. If it turns out that objects with quite different intrinsic properties acquire the monetary value now had by \$1 bills, then these other objects will soon (after their reconfiguration by manufacturers) cause the machine to dispense its contents. It is the monetary value of the paper that explains, via intentions and purposes of those who make these machines, not why cokes are released, but why internal objects with the intrinsic properties of money cause cokes to be released.

Monetary value, this extrinsic property, becomes causally relevant to machine behavior if we *correctly* identify what we are trying to explain—if, that is, we identify behavior not with the events that result from putting money into the machine, but with the causal process by means of which these events are brought about.

So, too, I submit, in the case of human behavior. Distinguishing human behavior, moving your arm, from bodily movement, the movement of your arm, enables extrinsic facts about your internal states to become explanatorily relevant to why you do what you do. Content, as an externalist understands it, gets its hand on the steering wheel once we correctly identify what sort of thing we are asking content to explain—viz., behavior, a causal process (typically) issuing in bodily movement, not the bodily movement itself.

B. Self-Knowledge.

There is a second problem, an epistemological problem, a problem about how we know what we think and feel if our thinking this and feeling that are extrinsic to the agent

having the thoughts and feelings. You don't look in the mirror, inspect yourself, to find out whether you are divorced or wealthy. If you want to know whether you possess these relational properties you look elsewhere, at the facts that determine your relations to other things. And so, too, should it be with what you think and feel if what you think and feel are relational facts. But we don't look elsewhere to find out what we think and feel. We introspect or, if that isn't quite the right way to put it, we just know these things without *specting* anywhere. We certainly don't discover what we think by poking around in those outside (or historical) facts that externalists tell us determine that (and what) we think and feel. This has suggested to many that the content of our mental states must be intrinsic.

A lot of ink has been spilled over this problem and nothing I can say in the space of a few pages is going to settle matters. I do want to point out, though, that externalists have several options for neutralizing the sting of this problem. The first option is to insist that in the normal course of events extrinsic content will supervene on, it will be reflected in, the intrinsic constitution of an individual in the same way age supervenes on the tree's internal constitution. If a tree wants to know how old it is, if it wants to find out whether it possesses the relational property of being 20 years old, it should look inward, inside its bark, and count its rings. That reveals the determinate value of this relational property. There is nothing in externalism to prevent it from using this model to understand the epistemic value of introspection. We look inward to determine what we think and feel because these extrinsic properties supervene on something inside that reveals, more or less infallibly, what we think and feel. Combined with a reliability theory of knowledge that exempts one from having to know A supervenes on B in order to use B to discover facts about A, this seems like an acceptable answer to this epistemological problem.

There is, though, another option, an option I prefer (since I don't like the simple reliability theory of knowledge that the first answer assumes). Perception tells us what is in the external world: I know you ate a banana, not an apple or an orange, because I saw you eat it. Though perception tells us, in this way, what it is in the external world (in this case, a banana), it doesn't tell us that there is an external world. You could see it was a banana, but you can't see what is implied by this--that it was a physical object, something that enjoys an existence independent of your experience of it. You can see something—viz., that it is a banana--that implies that idealism is false, but you can't see that idealism is false. Likewise, memory is a faculty that tells me what happened in the past—that I had granola, not eggs, for breakfast this morning—but it doesn't tell me what is implied by what I remember: that there is a past. I cannot refute Russell's skeptical conjecture about the past (that the world was created 3 minutes ago) by remembering that I had granola for breakfast (and deducing from this the fact that the past is real) anymore than I can refute external world skepticism by watching you eat a banana. And what is true of perception and memory is, I submit, also true of introspection: it tells you what is in your mind, not that you have a mind. It tells you what you think and feel, the content of your mental states, not that you have mental states with this content. If you want to know that you think that P, you need to find out (at least) two things: (1) that it is P you think, and (2) that you think it. What introspection gives you access to is (1), not (2). Externalists have an answer (supplied, in large part, by Tyler Burge and John Heil) to how we know (1): we think whatever we think we think because the content of the lower order thought—i.e., P--is automatically absorbed into the content of the higher order thought (viz., I am thinking that P) thereby making the higher order thought unerring about the

content of lower order thought. The externalist facts that determine (at the first level) what it is we think are the same externalist factors that determine (at the second level) what it is we think we are thinking. So we don't need access to these factors to know what we think.

But, notice, this neat answer to how we know what we think is an answer to (1), not (2). It tells us how we know what the content of our thoughts are. It does not tell us how, or even whether, we know we have thoughts with this content. It is an answer to the question: "How do you know it is P you are thinking?" not an answer to the question: "How do you know you are thinking it?" To that question externalists need a different answer. But so does everyone else.