

Synthese 70 (1987) 205-227.

RADU J. BOGDAN

MIND, CONTENT AND INFORMATION

1. INTRODUCTION

What is it that one thinks or believes when one thinks or believes something? A mental formula? A sentence in some natural language? Its truth conditions? Or perhaps an abstract proposition? The current story of content is fairly ecumenical. It says that a number of aspects, some mental, other semantic, go into our understanding of content. Yet the current story is incomplete. It leaves out a very important aspect of content, one which I *call incremental information*. It is information in a specific format, information as a limited or local increment, structured by a number of underlying parameters. It is in the form of such increments that information drives cognition and behavior. This is why, perhaps of all aspects of content, it is incremental information which matters most when we want to understand cognitive attitudes and performances. This in turn must have an impact on our philosophical notions of content, propositional attitudes, inference, justification and knowledge.

Yet the impact has not materialized so far. Whence my puzzle. In different versions the story of incremental information has been around for a while. Linguists, psychologists and AI people have done a good deal of work on it and so have, from more specialized angles, theorists interested in the logic and semantics of questions and epistemic verbs, in inductive inference and the pragmatics of explanation and in epistemology.¹ In philosophy of mind there is no discernible interest in incremental information. Although this is a time when the philosophical concern with content is quite intense, there is no clear perception of the connection between content and incremental information. Why? My speculation focuses on three reasons. One is that incremental information is simply not regarded as an *aspect* of content. Another is that,

when perceived as an aspect of content, incremental information is regarded as an *epiphenomenal* aspect, that is, one which is reducible to and explained by other, more fundamental aspects of content. The third reason is that, even when seen as, a genuine and irreducible aspect of content, incremental information is not regarded as a mental aspect of content, hence not as something that drives cognition and behavior.

One must therefore invalidate these reasons if one wants, as I do, to put the notion of incremental information to serious work in philosophy of mind and metapsychology. To invalidate the reasons just mentioned amounts to showing that incremental information is an intrinsic, irreducible and mental aspect of content. Showing this is the major task of my paper. To motivate the enterprise and put it in a larger perspective I will also address some objections and, in conclusion, anticipate some implications.

2. ASPECTS OF CONTENT

To say that one thinks or believes or in general cognizes something in some particular modality is to say that one has a specific attitude toward a content. Philosophers used to call such attitudes *propositional*. Strictly speaking, this is question-begging since it assumes without argument either that a proposition is content or else that, among the various aspects of content, it is the propositional aspect that cognizers have attitudes toward. Neither alternative should be conceded. My argument will show both alternatives to be false. So I would rather talk of cognitive or *content attitudes* when we talk of belief, thought and the rest. But I am not going to worry about content attitudes. The focus here is on content. I am, however, going to assume that whenever incremental information is involved a content which it inhabits is subject to one attitude or another. I am going to show that the attitude is sensitive to incremental information, not just to a proposition or sentence or mental form.

A content must be described. When we say that Big Lem believes that it is raining, the clause *it is raining* functions as a content description. I will occasionally call it a *surface description of content*. It describes what Big

Lem believes. But a content description is not yet a content specification. What Big Lem believes was just described, in English, as *it is raining*. Yet it does not necessarily follow that what Big Lem believes is the very English description of the content. Big Lem may not know English. Or he may not believe what he does under the very concepts presupposed by the description. And there are other differences between content and its surface description, as we will soon find out. To assume, therefore, that a content description is a content specification is to beg quite early in the game the very issue we want to address. Since, then, a content description is not necessarily a content specification, the question is: How do we get from a content description to a content specification? As anticipated, content has many aspects or parameters, or so we assume nowadays. So when we are asking: when does a content description deliver content?, what we are really asking is: which parameters of the content description would fix content? And which parameter in particular would fix the incremental information it carries? This is the *individuation* question. It is the first but not the only question we want to ask about content and incremental information. We not only want to identify them, but also to understand them in terms of some further, deeper features and regularities. We want, in other words, a principled account of content and its aspects. This is the *characterization* question. The main task of this paper is to address the individuation question. But if its answer to this question is right, it points the direction in which we should look for an answer to the characterization question. The direction is that of the cognitive mind.

One caveat before we go. It is not my aim here to either define or explain content in cognition. I only want to show that the aspects we tend to think of most, when we think of content, are not capable of fixing a *further* aspect, incremental information. I am not, however, saying that with incremental information we have a complete story of content, let alone that incremental information is content. I am only saying that without incremental information we have no story of content to tell.

How do we individuate content? On the one hand, we must rely on an intuitive notion of content, one which views content as what-is-believed or thought or, in general, cognized. This intuitive notion helps us decode and interpret the surface description of a content, given other things we know

about the circumstances. On the other hand, we propose various philosophical explications of content. To a large extent, the worth of the philosophical explication consists in the degree to which it matches and rationalizes the intuitive notion, given a larger philosophical picture where content has a place. Let us see how this works.

To individuate something is to establish its identity conditions and thus be able to tell when something is, and is not, the same as something else. Since, as assumed, we start from a (surface) content description, the question of individuation becomes, "which parameters of content description individuate content?" Given that an individuation question rests on the notion of sameness, which is an equivalence notion, we can work with the notion of content equivalence. Equivalence in general is aspect-relative. Things are always equivalent under some aspect or another. So must be contents. The parameters of content description are precisely the aspects that interest us. So we can say, of two content descriptions, that they are or are not equivalent under a certain aspect, such as truth value or meaning. The content descriptions satisfying a thusly relativized equivalence relation will be said to form an equivalence class. If content descriptions are equivalent under one or more aspects deemed both necessary and sufficient to individuate content, we say of those descriptions that they have the same content. Then, presumably, we know what content is, not its nature, for that is a characterization problem, just its (type) identity.

Let us see now how various standard parameters meet our intuitive expectations about content. We will begin from the semantic end and work our way toward the mental end. A content-descriptive sentence must have a *truth value*, that is, it must be either true or false. But we do not expect truth value to individuate content, for truth-valuationally all true content descriptions would have the same content, which is absurd.² So we need a finer-grained parameter which would group true content descriptions according to what makes them true. This is the job of the notion of *truth condition*. The notion specifies what in the world (things, properties) relates to what parts in the content sentence such that the latter is true of the former. A sentence describing my belief that Pusha is a nice fat cat is true if and only if what 'Pusha' refers to is what 'nice fat cat' refers to. Yet truth conditions cannot fix

content either. *Pusha is a nice fat cat* and *The only four- legged animal residing illegally in Escondido Village is a nice fat cat* happen to share the same truth conditions but do not describe the same content. One may believe the former but not the latter for a variety of reasons, all having to do (to put it generally) with one's mental states. Thus one may be unaware that Pusha lives in Escondido Village or one may lack the concept of illegality. Further parameters are needed. So far, we notice, the head of the cognizer has been black-boxed. The policy of content individuation was publicly transparent and nonintentional because it was insensitive to the ways the cognizer represents the facts. No longer. We are now stepping inside the head. Content individuation becomes opaque and intentional.

There is a lot of stuff inside the head that may affect content individuation. Should content identity require *computational* identity, that is, identity of content-encoding formulae in some inner language of mental representation? Not necessarily. The same formula may have different interpretations, hence different contents. Thus, closer to the surface of the natural language, a sentence like, *The nuts are all over*, can describe various contents. So can a bunch of symbols in the inner mental code. It is unlikely, since it is uneconomical, that mental symbols have intrinsic and unique values. Nor is formal difference automatically reflected in content difference. *Pusha ate the nut* may share the same content but not the same form with *It was the nut which was eaten by Pusha*. For all we know, similar discrepancies may also operate at the level of the inner code. So mental form and content go their own separate ways. Nor does *conceptual* sameness fare better in fixing content. An indexical sentence, as we all know first hand, may involve the same concepts and yet be descriptive of as many different contents as there are persons thinking that sentence. Take your favorite candidates contemplating the sentence *I am the greatest philosopher north of New Orleans*: same English sentence, same concepts, same inner formula in the mental syntax, perhaps same emotion in the respective breasts, yet as many contents as vanity allows.

This last example drives home the point that no single standard parameter, outside or inside the head, can fully individuate content. A joining of forces seems necessary. This would be a hybrid view of content, that is, the view

that it takes several parameters, both semantic and mental, to fix content. The hybrid view was initially defended by Putnam in his analysis of meaning. On that analysis, it takes syntactic categories, concepts and stereotypes, on the mental side, and indexical extension, on the semantic side, to fix meaning. This sort of analysis has been extended to content in general. Some theories also add a social dimension to the determination of content. Although no longer elegant and simple, the hybrid analysis appears to many philosophers as powerful enough to individuate cognitive content.³ I disagree. On the mental side of content, incremental information is still elusive.

3. INCREMENTAL INFORMATION

First Approximation

To see why, consider the following story. Babeau and Babette are told in so many words that Pusha emptied the bottle of wine. They register the same input, ultimately caused and thus made true by the same event, Pusha-emptying-the-bottle-of-wine, compute the same input sentence the same way, from phonology to grammar to concepts to meaning, yet then end up behaving differently. Babeau is relieved and starts dancing; Babette is rather sad and goes for a walk. Both had the same desire: to drink that bottle of wine at dinner. So the beliefs formed by the input must have been different. Indeed, prior to the input, Babeau was worried that Pipus, their intellectual and slightly paranoiac cat, might have done it. Babeau was worried because he has trained Pipus to avoid wine bottles. Babette, on the other hand, has hoped that the fat, relaxed and fun loving Pusha has emptied something else, not their last bottle of wine. So Babeau and Babette have extracted different information from the same input. Extracting information in a case like this requires thought. Information acted upon is belief. So both thought and belief must differ in Babeau and Babette although the causal input and many psychological parameters (syntactic form, concepts, etc.) are the same. Belief and thought are content attitudes. Their difference in Babeau and Babette is one of content. It is incremental information that makes that difference. But then incremental information must be part of content.

Let us see how the informational difference works. Babeau and Babette had distinct informational expectations, that is, distinct zones of uncertainty because their background knowledge was not the same. Babeau knew that

(1) Someone emptied the wine bottle

and even suspected one of their overimaginative cats but did not know which one did it. Babette, on the other hand, knew that

(2) Pusha emptied something

but did not know what. When the input sentence to the effect that

(3) Pusha emptied the wine bottle

becomes available, it fills different informational slots because it bridges, as it were, different informational gaps in the knowledge that Babette and Babeau have about the situation. For Babeau the gap to be bridged is from *someone to Pusha*, for Babette from *something to wine bottle*. These are precisely the incremental differences in information between (3) and (1), and (3) and (2), respectively. These differences help individuate the incremental information that (3) adds to the prior information Babeau and Babette already had.

If (3) were the literal and full specification of what Babeau and Babette believe, then there would be no content difference between

(4) Babeau believes that Pusha emptied the wine bottle

and

- (5) Babette believes that Pusha emptied the wine bottle.

Yet the information they have access to, upon learning that (3), is different and this leads to different behaviors. Since the notion of belief is half information, half behavior-from-information, we have to conclude that (4) and (5) cannot describe the same beliefs. This is because, in an incremental sense, the content sentence (3) does not individuate the same information for Babeau and Babette. Yet in (4) and (5) the content clause (3) has the same truth value and truth conditions, the same grammatical structure and logical form, the same concepts and possibly stereotypes, hence the same hybrid meaning. (4) and (5) cannot be construed as describing type-identical beliefs because of a content-bound difference in incremental information. It follows that the standard parameters surveyed so far cannot fix incremental information. This is why a content description like (3) which is associated with such parameters fails to provide a complete specification of what one believes or thinks or in general cognizes.

Individuation

What does it take to individuate incremental information? The essential structure is that of given and new information. For Babeau the fact that someone emptied the wine bottle is given (or background) information, relative to what he learns later. The new information for him is that that someone is Pusha. The given information is what is held fixed and unchallenged in a particular cycle of incrementation or information expansion. The structure of the given information may contain one uncertainty slot or more to be filled with new information, although practically any slot, even when not perceived as one of uncertainty, can be expanded informationally.

We can construe the Uncertainty slot as a variable ranging over several alternative candidates for the new information compatible with what the agent already knows. We can think of these candidates as forming a class. Let us call it the *alternatives class*. For Babeau the alternatives are potential answers (compatible with everything else he knows about the situation at hand) to the question, Wholwhat emptied the wine bottle? For Babette, whose given information is that Pusha emptied something, the burning question is, What did Pusha empty? Again, as far as Babette knows, there are a number of alternative targets for Pusha's emptying routines.

It is important to see that in other circumstances the same input sentence (3) could have had different incremental contents. Following custom, let us capitalize the new information in the content clause and assume that the rest is part of the given information. If, in the new circumstances, Babette knew in advance that Pusha has done something, then she could have acquired from (3) the new information that

(6) Pusha EMPTIED something

or that

(7) Pusha did something to the WINE BOTTLE

or finally that

(8) Pusha EMPTIED THE WINE BOTTLE.

We can even imagine Babette just having the (given) information that something happened in the house. The alternatives class in this case is understandably large and so is, correspondingly, Babette's uncertainty. Now it is the entire

(9) PUSHA EMPTIED THE WINE BOTTLE

that, relative to the given information that something happened, constitutes the new information. It could have been THE ROOF FELL or PIPUS ATE A

REIDEL BOOK. The difference between (3) and (9) then is simply this. (3) is an informationally insensitive and hence inadequate specification of content; that is, one which does not incorporate incremental information into content, whereas (9) is an informationally sensitive and hence adequate specification of content. It does not matter how the difference in informational sensitivity is represented, as long as it is represented. It is no good to say that both (3) and (9) are the same sentential description of the same facts or that they express the same proposition or have the same meaning or that they are identical under any other standard parameter of content. This is true but irrelevant. For, as earlier argued, all those parameters that go into explicating (3) fail to capture incremental information, although there is no question that they capture the fact involved, the proposition expressed, the concepts applied, the meaning supplied and so forth. The capitalization of (9) is just a graphic way of indicating that increment-individuating parameters (given/new information and others) are at work. Intonation in a verbal context, the use of a cleft construction or other tricks can do the same job. (3) and (9), as descriptions of content, are identical under many aspects but not under incremental information.

There is perhaps a more helpful way to represent the given and new information and therefore incremental cognition and the content attitudes associated with it. We can think of incremental cognition as progression from given to new information. The starting point, or base, which is the given information, is assumed and not subject to current incrementation. We can represent this by indicating that the given information is outside the scope of the current incrementation. This will come out as

(1 0) Of [given information] incrementation to the effect that [x = new information].

To take one instance of our story, we can represent Babeau's original incrementation as

(1 1) Of [Some x emptied the wine bottle] Babeau's incrementation is that [x = Pusha].

Babette's on the other hand is representable as

- (12) Of [Pusha emptied some y] Babette's incrementation is that [y = wine bottle].

Also, to give just another illustration, if Babette's given information was that something happened, then her incrementation to (9) would appear as

- (13) Of [some event p] Babette's incrementation is that [p Pusha emptied the wine bottle].

A schema like (10) and its instances are obviously artifacts of analysis. There is nothing in ordinary language in which to naturally paraphrase (11) or (12) or (13), nor is there anything in our common sense psychology which approximates (10). This should not surprise us. There is no notion in ordinary language which just captures the computation of meaning as such or pattern recognition or grammatical analysis of a sentence at the deep structure level, nor is there anything in common sense psychology to enable us to conceptualize such phenomena. Incrementation is in the same boat, one more aspect contributory to the overall fabric of content. Yet there is a further reason why the incremental schema described by (10) does not surface *as such* in ordinary attributions of content. Incrementation leads to updating one's information about a particular topic. This is to say that the new information resulting from incrementation is *integrated into* the structure provided by the given information. Recall that the latter contains an uncertainty slot in need of updating. The new information, once acquired, will occupy the slot in question within the earlier structure. As a result, we can say that the final description specifies not the increment as such but rather (what we may call) the *terminal* information. The latter is the outcome of updating locally one's given information with some new information. We can represent this as

- (14) Of [given information that a-b-x] after incrementation to the effect that [x = c = the new information] the terminal information is [a-b-c].

If we assume that incrementation is made possible by a process of inference and also that the terminal information is encoded in the overall content of some attitude such as perception or beliefs, we can represent (14) as

(15) Of [given information] upon inferring that [new information] one believes that [terminal information].

In our story, in Babeau's case, this will amount to

(16) Of [some x emptied the wine bottle] upon inferring that [$x = \text{Pusha}$] Babeau believes that [Pusha emptied the wine bottle].

(16) is a description-under-analysis of what it is informationally for Babeau to believe that Pusha emptied the wine bottle. Unlike the surface description (4), (16) reveals the underlying informational expansion involved and thus the incrementation to which Babeau's initial belief is subjected.

What else is needed to individuate incremental information? Obviously there is a *theme* to one's cognition at a given time, something which defines the area of interest or attention. A theme marks the outer boundaries of potential incrementation. A theme keeps track of antecedent incrementations and thus provides continuity and relevance. A visual scene is a theme, and so is a story, a problem to be solved or a plan to be executed. To fix the format of incrementation we also have to specify the *categorial articulation* of the information involved. For example, we can categorize a situation in terms of object-property or agent-action-object or in terms of other structures. Incrementation can then take place along some dimension of the structure. Thus Babette's incrementation to the effect that it was the wine bottle that Pusha emptied was along the *object* dimension in the categorial structure agent-action-object. Other individuation parameters may be cited but it is not my present purpose to go into too many details. The intention is first to convey some idea of what it takes to individuate information in cognition and second, in so doing, to strengthen the argument that incremental information is an aspect of content which cannot be reduced to the standard parameters mentioned earlier.

Let us now summarize. Consider again the case of Babette coming to believe that

(17) Pusha EMPTIED the wine bottle.

The analysis of incrementation sketched so far says that an informationally sensitive description like (17) reveals the incremental aspect only if it is relativized to the following parameters:

- (A)i. theme: what is attended to (event in the house)
- ii. given information: what is held fixed (Pusha did something to the wine bottle)
- iii. uncertainty: identity of action (what did Pusha do to the bottle?)
- iv. projected: a relevant alternatives class (emptied, played with, broke, etc.)
- v. inferred: value for the new information (emptied)
- vi. categorial articulation: agent-ACTION-object
- vii. integrated: new information (emptied) into given information (Pusha did something to the bottle) → terminal information (Pusha emptied the wine bottle).

It takes at least a specification of (i) through (vii) in (A) to individuate the incremental information brought about by (17) and thus fix Babette's specific belief. Take another example. Do I believe that $2 + 3 = 5$ is true? Yes, I do. If we stop here, the content individuation of my belief is incomplete. I do believe that $2 + 3 = 5$ is true, but as opposed to what? What are the relevant alternatives? What is at stake? What is the context, the theme? The nature of mathematical truth? A particular mathematical truth? Skepticism? What is the uncertainty? What do I or should I hold fixed as given information? Notice that all these increment-individuative questions do not bear on the very truth of $2 + 3 = 5$ or on my believing it *to be true*. Truth is not yet information and belief is information-hungry. Literally construed, believing-something-to-be-true is not yet fully believing if it takes only truth values or truth conditions to fix it. To fully believe, what I believe must also be thematic, categorially articulated, an alternative to other possibilities, in short, informative. The philosophical slogan that to believe is to believe it to be true

is nonsensical if understood to imply that the truth aspect is sufficient to characterize what is believed.

4. OBJECTIONS

There are several types of objections to my story of incremental information. One is to firmly separate information from content and claim that information is an effect or by-product of using content, not a part of content. The shortest answer is that the objection begs the question, for it assumes instead of arguing that information is not part of content, of what one cognizes. Another answer is that if content does not encode and carry information, then by definition it is not *cognitive* content since cognition is understood as information processing. But then the objection is irrelevant since our overall explicandum here is cognitive content. A third answer is that an informationless notion of content cannot be the notion of content we associate with content attitudes such as belief or thought. There is no informationless thought or belief.

There is another type of objection which recognizes that content is information but maintains that the notion of information involved can be analyzed in terms of meaning or proposition or truth conditions. This is the semantic notion of information, a notion which has received wide attention in the philosophical literature for the last four decades. This is a legitimate and useful notion for specific inquiries. So is the statistical notion of information. However, neither notion is subtle enough to capture what constitutes information for a specific cognitive system, although what each of the former analyzes may be necessary for what the latter analyzes. Thus, for example, the statistical notion may capture aspects of information as input or message, whereas the semantic notion may capture the logical discriminations that a certain form or sentence is capable of. We should think of incremental information as mapping the statistical and semantic aspects onto incremental parameters like those discussed earlier in the previous section.

Nevertheless, the current consensus in philosophy seems to be that content either is informationless or else is informational only in a semantic way. One wonders why. I think that a bit of history can help. The source of the consensus can be traced back to Frege and Husserl, among others, and the

specific questions about content they were interested in. *Their* major question was: What is it across ages, languages, concrete cognitive circumstances and other such pragmatic inconveniences that, say, Pierre and Peter understand when one registers that *Le sicritaire general du Parti a accepti la position parce que son pire l'a refusie* and the other that *The general secretary of the Party accepted the position because his father refused it*? Their answer was: A Noema or Gedanke, that is, some abstract propositional object. Such an object by itself is cognitively inert, a meaning or some sort of structure without information. But remember the focus of the question: what-is-being-understood. My point is, what-is-being-understood is not what-is-being-believed or in general cognized, although understanding may be necessary for belief. Since what is understood is not a full cognitive content, understanding itself cannot be a genuine content attitude. Understanding is perhaps computing a meaning or figuring out truth conditions or something along these lines. Believing is more than that. But we rarely if ever understand without at the same time believing, inferring or imagining. This is to say that we rarely if ever compute a meaning without involving it in a larger cognitive endeavor. Furthermore, Frege, Husserl and countless other philosophers assume without argument that the object of understanding is also the object of thought or belief. This assumption need not be granted and in fact is mistaken, as I have shown earlier.

The moral is this. If your "philosophical intuitions" have been trained and conditioned by theories of "content" in the Frege-RussellHusserl tradition, that is, if there is a theory-motivated sense in which you accept that we, as cognizers, literally think abstract propositions or believe meanings or remember truth conditions, depending on which is your favorite aspect of content legislated into Content Itself, then you will fail to see why incremental information must be part of content. But then, my argument is, you also fail to see why content encodes information at all and therefore why content is cognitive. But then again you fail to see why content is subject to a cognitive attitude or another, since those attitudes are information-bound. But then, finally, if neither information nor attitudes constrain your notion of content, it is not content that your notion is about.

Another semantic challenge is posed by a view which construes our explicandum as what I *call emphatic reference*. Suppose we say that

(18) Socrates' *drinking hemlock* at dusk caused his death.

On one view of emphatic reference, that of Peter Achinstein,⁴ the emphasis in (18) refers to that aspect of the event which caused the death. A shift in emphasis, say to *at dusk*, picks out another causally efficacious aspect and thus changes the meaning and the truth value of the resulting statement. Indeed as a causal statement,

(19) Socrates' drinking hemlock *at dusk* caused his death

is false. Achinstein does not explain how the emphasis picks out the right causal aspect. The choice, I think, is this: either emphatic reference is primitive, somehow self-contained and independent of cognitive processes OR ELSE it is parasitic on the latter. I will consider and reject the first alternative in a moment. So emphatic reference must be parasitic. I maintain that it relies on informational incrementation. My point is that (18) and (19) are implicit causal claims whose explicit informational structures are

(18a) Socrates' drinking hemlock at dusk caused his death and it is drinking hemlock which caused his death

and

(19a) Socrates' drinking hemlock at dusk caused his death and it is his doing so at dusk which caused his death.

In both (18a) and (19a) it is the second conjunct which provides the new and specific information concerning causal responsibility. In one case the information is right, in the other it is not.

The other alternative, as I said, is to construe emphasis as genuinely and irreducibly referential. This, it appears, is what Fred Dretske does.⁵ His view is very roughly the following.

(20) PUSHA emptied the wine bottle

and

(21) Pusha emptied the WINE BOTTLE

are two different embodiments or allomorphs of the proposition that Pusha emptied the wine bottle. Allomorphs have "worldly counterparts", namely, "event allomorphs" which propositional allomorphs like (20) and (21) refer to. This becomes obvious, Dretske thinks, if we look at the gerundive nominalizations of (20) and (21), namely

(20a) PUSHA's emptying the wine bottle

(21a) Pusha's emptying the WINE BOTTLE

and ask what they refer to. They cannot refer to different events because (20) and (21) describe the same event, yet they refer to something objective out there because the causal efficacy of what (20a) refers to is different from the causal efficacy of what (21a) refers to. The former, we may recall, caused Babeau's joy, the latter Babette's sadness. Allomorphic events like those designated by (20a) and (21a) are, on Dretske's view, causal aspects of the event described as *Pusha emptied the wine bottle*. Events cause in virtue of their allomorphs.

How is allomorphic reference itself to be explained? Dretske does not appear to have an answer but we may be able to fabricate one by piecing together elements of his story of allomorphic reference. Allomorphic events have causal potency. Moreover, it is when we focus on their causal aspects that we can identify allomorphic events as targets of emphatic reference. Finally, Dretske notes, most allomorphically sensitive contexts are epistemic and they are thus sensitive because they are causal. In other words, the reason why, say, a perceive-that context is allomorphic is that there is a causal relation between an allomorphic event out there and the perception that registers it. The same applies to remembering or believing.

Now here is my reconstructive fabrication. The overwhelming majority of contexts which Dretske counts as allomorphically sensitive, because causal, is

epistemic. Reference itself is not among them. This poses a simple choice. Allomorphic reference either is or is not part of a larger cognitive process such as perception or memory. If it is, then it is very likely that its emphatic, allomorph-tracking function is parasitic on the informational accomplishments of the process it inhabits. If, on the other hand, allomorphic reference is not cognitive, then it must rely on an independent mechanism. What can it be? Typically, philosophers think of three mechanisms of reference: description, immediate, demonstrative acquaintance and causation. The first relies on information and the second cannot handle the complexity and abstractedness of the contents we are considering here. So only causation is left. But there are problems with a purely causal theory of emphatic reference. One is that in cognition causation takes the form of an information flow. Dretske would be the first to agree with this. Another serious problem is that a causal theory cannot handle propositional structures without relying on information, and emphasis can be propositional. Suppose an entire sentence is emphatic. What that sentence represents cannot be characterized in purely causal terms. A fact out there does not simply cause one to refer to it with some sentence. A fact or event may cause an organism to first register some information which, in turn, after many processes, may invite a sentential representation. But once we conceptualize the issue this way we are no longer talking of mere, precognitive causation. We are now talking of information processing. It is the latter, I conclude, that emphatic reference is parasitic on.

There is finally the pragmatic view.⁶ To a large extent, this view is legitimized by the informationless notions of content that I have attributed earlier to the Frege-Husserl tradition. Indeed, if content is informationless, then its informationality must come from outside. The first move in the pragmatic game is to locate incremental information in linguistic communication. The next move consists of arguing that incremental information must be a pragmatic creature of contexts, use and communication because incrementation in communication cannot be specified unless we fix all sorts of contextual, functional and socio-communicational parameters. This is, to begin with, a superficial argument. Even if it is true that communication produces and handles incremental information, the question we want to ask is: why and how does it do it? If we think hard enough, it seems to me, we will find that the answers to both questions must eventually involve cognitive

design, hence mental reasons. But the argument is also fallacious. It is a fallacy (of the behaviorist or verificationist sort) to argue from the manner of specification to the nature of what is being specified. To specify temperature we need some measurement conventions yet we do not want to say that those conventions characterize the very nature of temperature as a physical phenomenon. Likewise in the domain of the mental: it does not follow from the fact that some mental phenomena need pragmatic specification that the phenomena themselves are less mental or in fact pragmatic. On the contrary, one can make a good case that the need for pragmatic specification is good evidence that we deal with highly complex mental phenomena. After all, we do not need much pragmatics to figure out the amoeba's cognition! So much, critically, for the second pragmatic move. But the first move need not be granted, to begin with. It is simply not true that incremental information is exclusively a matter of discourse. It is generally a matter of mind. This is the issue to which we now turn.

5. MIND AND INCREMENTATION

The argument so far has been that incremental information is an irreducible aspect of content. It remains now to show that it is a mental aspect of content. Showing this is important here for two reasons, one particular, the other more general. The particular reason is this. It is only as mental that information is incremental *relative to* the parameters identified two sections ago. For any information-carrying message and any receiving system endowed with memory, it is true that the message adds some information to what the system already knows. In other words, for any system with memory, some input information is incremental. But not in the sense pursued here. It takes specific mental capabilities (for computation, language, concepts, belief interaction, etc.) to map some input information into the incremental output satisfying conditions (i)-(viii) in (A) above. The general reason is that it is only as mental that incremental information drives an agent's cognition and behavior and thus figures in our understanding of her specific cognitive attitudes and performances. Generically, it is only as mental that incremental information can be part of our explanation of mind, intentionality and content attitudes, among others. So much, then, for why the mentalization

of incremental information matters. The question now is how to construe this position.

There is a trivial sense in which incremental information is mental but this sense is not the one we need. Trivially, incremental information is mental because incrementation itself is executed by mental means and also because the outcome of incrementation, the information itself, is encoded in mental states. Nobody would deny this fact. The problem is that this fact does not show that information is incremental *for mental reasons*. It only shows that, *whatever* the reasons for its being incremental, information is encoded in some mental form or another. This is perfectly compatible with the reasons for incrementation being merely semantic or communicational or generally pragmatic. In any of these cases, the information would be mental because it is cognitively handled, and incremental because this is the format in which the information is used. There is, however, another, nontrivial sense in which incremental information is mental. This is a sense in which the reason the mind handles information in increments is internal to the mind itself and has to do with the way the mind is designed and the way it works. I propose to argue for this nontrivial sense from two perspectives, one principled, the other empirical. The principled perspective focuses on the design of the mind. The empirical perspective focuses on relevant psychological and common sense evidence concerning the ways the mind operates with information.

There is a principled reason why the mind must have been designed to deal with information incrementally. It has to do with its finitude and in particular with how much information a cognitive process can handle at any given time. It turns out that it is not that much. We know for example from George Miller that there is a limit (7+2) to how many unrelated units short-term memory can service. Other such capacity limitations have been found in attention, recognition, judgment tasks and so on. The actual human limits don't matter, which is why the problem is one of principle. One can imagine, perhaps even devise, a short-term memory capable of holding 7000 ± 35 units. It is still much less than the (input) information assaulting an organism. It is the limitation itself (whatever its magnitude) that requires a design solution. So how can we make the best of such limitations? What would be the optimal format for handling information, given the limitations? If we think of the format

problem as a design problem, then one rational solution would be to choose fewer units to be held in memory but to pack more information into them. For example, if the system can process a natural language, it would be counterproductive to ask it to just store simple symbols or letters as basic units of information. Words are more efficient and groups of words and entire sentences will do even better. Another good design policy is to bring the units under some meaningful unity and make them part of a story, of a theme. Relevance and continuity help memory; irrelevance and disconnectedness don't. Miller calls the organization of information into relevant and flexible units chunking. To enable the system to do such chunking we have to endow it with adequate mental capabilities.

Notice now that a nonmentalist argument would make chunking a matter of learning and context and thus trivialize (in the first sense earlier described) the mental aspect of informational incrementation. On the other hand, the design argument, which is mentalist, claims that the organism is programmed to look for chunks given some further constraints, just as it is programmed to find relevant connections between chunks. The context only provides the matter and opportunity for chunking the input for the increment-slots. But the context does not dictate the chunking itself, as a cognitive policy.

I turn now to the psychological evidence. The strongest evidence is the very pervasiveness of incrementation, its presence in various cognitive modalities and processes. Consider vision. One simple reason why vision is incremental is that it interacts with memory. Features already registered and identified are held in memory as given information while the occurrent intake has room for the new information. Another reason for incrementality is that vision may be operating relative to a set of frames or schemata or some such sort of "structural expectations" which organize the input along chunk-like dimensions and thus allow incrementation. Marvin Minsky seems to believe that a frame organization of visual perception is a matter of optimal design and there appears to be experimental evidence to this effect. Fred Dretske is the first philosopher to seriously examine the incrementality of perception. In his view, a visual achievement is incremental. Thus, in a given context, Babeau could see that Pusha is emptying the wine bottle without having to see, in the same context, that it is Pusha who does it. He may have identified Pusha in an

earlier context either visually or by other means. Once the Pusha-information is fixed, Babeau's visual achievement in the current context is just that the action in question is that of emptying a wine bottle. This is the new information visually acquired. If we consider only the phenomenal quality of vision, which is what we typically do, we will find this incremental account fairly odd. After all, when Babeau sees Pusha emptying the wine bottle he sees Pusha among other things. This is the phenomenal but not the structural story of Babeau's visual achievement. Dretske does not make the phenomenal/structural distinction but I think it should be made. Quite roughly, the phenomenal aspect of what is seen can be explained in terms of what the transducers register and deliver to the higher processes, an image. The structural aspect, on the other hand, is a matter of format, of slots within frames within schemes, or whatever account would structurally fit the individuation parameters identified earlier in (A). Incremental information is a matter of structure, not of phenomenal input as such. A third internal, cognitive reason why the output of vision is incremental is that it interacts with other types of mental processes, not only memory, but also beliefs and thoughts, for example. If these processes interact informationally, as they do, it is quite normal that they will do it along incremental lines.

This brings us to the next story, that of memory. If Babeau can limit his visual increment to the emptying of the wine bottle, this is because the earlier information on Pusha was already stored in memory. But now, visually, Babeau adds an increment to the stored information, so memory itself is enriched incrementally, and also incremental is the very belief formed as a result. There is then the phenomenon of memory interference which again indicates incremental structures at work. Elizabeth Loftus, for example, reports an experiment in which subjects are shown the film of a traffic accident. They are then asked somewhat misleading questions geared to the given information, that is, to the information the subjects are assuming and not paying much attention to. Later on, when asked to report on what they saw in the film, the subjects include the misleading information. Loftus concludes: "This seems to happen because the information in the questions, whether true or false, can become integrated into the person's recollection of the event (which was visual), thereby supplementing that memory".⁹ For language-perception-memory interference there is interesting evidence showing that

increment-serving syntactic structures such as active-passive constructions or cleft-constructions can influence the visual information people extract from the scenes they see, relative to what they already know and expect.¹⁰

Finally, consider inference. A good number of psychologists take the basic computations underlying perception or language processing to be inferential and to consist of an incremental hypothesis formation and testing.¹¹ Whether in this fine-grained form or in a more global form, as thinking, inference displays an obvious incrementality. It operates relative to some prior beliefs or a body of evidence (given information), a zone of uncertainty (the conclusion slot), and an alternatives class (candidates for the conclusion), among other parameters. This, after all, should not surprise us: inference is the engine of information expansion.

6. LOOKING AHEAD

This concludes my argument about incremental information being an intrinsic, irreducible and mental aspect of content in cognition. What follows from this? What larger picture does it open up? I see the following pattern of implications emerging.

We rationalize and explain cognition and behavior in terms of mental aspects of content. We also measure and assess the cognitive virtues and accomplishments of mental contents. This is to say that the notion of mental content is central to both intentional and psychological explanations, as it is to epistemological evaluation. If, as I have argued, incremental information itself is central to mental content, it follows that these forms of explanation and evaluation must be retooled if they are to capture and account for incremental information. The retooling must begin, I think, with the very notions of content attitudes, for these are precisely the notions with which we conceptualize, explain, rationalize and evaluate cognition and behavior. This assumes that we can show that some content attitudes such as belief and thought can and should be construed not only as sensitive to, but in effect as driven by, the incremental component of mental content. This is what I intend to show elsewhere.¹² It is, then, through these revised notions of content attitudes that

one would expect the story of incremental information to have its impact in psychology, philosophy of mind and epistemology.

NOTES

1 See, for example, N. Smith and D. Wilson: 1979, *Modern Linguistics*, Indiana University Press, for a linguistic survey; E. and H. Clark: 1977, *Psychology and Language*, Harcourt Brace Jovanovich, for a psycholinguistic survey; J. Hintikka: 1976 *The Semantics of Questions*, North-Holland, Amsterdam, for the semantics of questions and epistemic verbs; I. Levi: 1967, *Gambling With Truth*, Knopf, for inductive logic; A. Garfinkel: 1981, *Forms of Explanation* Yale University Press, for the pragmatics of explanation; and F. Dretske: 1969, *Seeing and Knowing*, Chicago University Press, Chicago, for epistemology. As my text and notes will often acknowledge, Fred Dretske has put informational incrementation on the philosophical map, mainly in epistemology. Yet his work on this notion has received relatively little attention on the part of philosophers and then only in linguistic contexts. I speculate on the reasons a few sentences later.

2 Could we imagine truth value ever individuating content? Apparently, some people can when they say seriously and literally things like "science is after truth" or "all I want to know is the truth". Imagine someone in court taking literally the injunction to tell the truth, the *whole* truth and *nothing but* the truth! What we want in cognition are specific representations, already constrained in various ways, which are also true. For the sake of the argument, however, the only instance I can think of where truth value might conceivably fix content is that of God's knowledge. One can test one's intuitions about this by going through the successive parameters to be introduced shortly and checking whether they are needed to fix the content of God's knowledge. A serious conceptual and theological problem may arise as parameters such as truth conditions, concepts, computations, and the rest, can be construed as bringing unneeded specificity and therefore limitations to God's knowledge. My hunch is that those ulterior parameters are not really needed, which would mean that, for God, truth value delivers content. But then God knows without cognizing, that is, without handling information.

3 For a representative sample of the hybrid view of content, see A. Woodfield (ed.): 1982, *Thought and Object*, Oxford University Press.

4 P. Achinstein: 1979, 'The Causal Relation', *Midwest Studies in Philosophy*, IV, 1979; also, 'Causation, Transparency and Emphasis', *Canadian Journal of Philosophy* V(I), 1975, 1-23.

5 F. Dretske: 1977, 'Referring to Events', *Midwest Studies in Philosophy*, II.

6 See, R. Stockwell: 1977, *Foundations of Syntactic Theory*, Prentice-Hall, Englewood Cliffs; also F. Dretske: 1972, 'Contrastive Statements', *Philosophical Review*, 81, 411-437. We could have also considered a further view, call it the *epistemic view*, which holds that the incremental structure is brought about by the operation of epistemic verbs. Fred Dretske has proposed this view a number of years ago in his 'Epistemic Operators', *Journal of Philosophy*, LX-VU, 1970, 1007-1023. I find the epistemic view plausible yet geared only to symptoms, for it does not answer the deeper question, Why do epistemic verbs create incremental structures in the content clauses? The epistemic view only notices that they do.

7 See D. A. Norman: 1976, *Memory and Attention*, 2nd ed., Wiley, where Miller's paper 1956 is partly reprinted; also R. N. Haber: 1976, 'Information Processing', in E. C. Carterette et al. (eds.), *Handbook of Perception*, 1, Addison-Wesley, Reading, Mass.

8 Minsky: 1977, 'Frame System Theory', reprinted in P. N. Johnson-Laird and P. C. Wason (eds.), *Thinking*, Cambridge University Press. Also, U.

Neisser: 1976, *Cognition and Reality*, Freeman, San Francisco. Dretske's work is *Seeing and Knowing*.

9 E. Loftus: 1980, *Memory*, Addison-Wesley, Reading, Mass. pp. 46-47.

10 PA. Hornby: 1974, 'Surface Structure and Presupposition', *Journal of Verbal Learning and Verbal Behavior* 13, 530-538; also Clark and Clark, op. cit., chap. 3.

11 See, for example, Jerry Fodor: 1975, *The Language of Thought*, Harvard University Press, Cambridge, Mass.

12 'The Manufacture of Belief', in Radu J. Bogdan (ed.), *Belief*, Oxford University Press, 1986.