that the accessing, reading, and even monitoring can take place without consciousness. If those operations are needed to fix belief, then belief can be an output of unconscious mental performances.

Although I do not think that the view of believing as consciousness is right, it does point to something important. There are circumstances (sleep, incapacitation, lack of attention, etc.) when the central mind does very little, if anything at all. Does the system believe anything under those circumstances? Even if some memory representations are explicitly stored as beliefs, a possibility we shall discuss later on, I do not think we can say that it does. Storing beliefs is not yet believing. The reason should be obvious by now. Without an appropriate context of incremental cognition nothing is being believed, for no incremental information is being fixed without cogitation, and in the circumstances we are talking about there is no cogitation.

Let us consolidate what has been said so far by weakening competing views.

4. WHAT BELIEVING IS NOT

Philosophers have said or implied many things about belief. Three, however, stand out as perhaps the most popular and resilient dogmas about belief. One is the dispositional, another the memory, and the third the registrational, or modular, view of belief.

Disposition?

An excellent characterization of the dispositional notion of belief was given by Ryle. He writes that believing 'is a propensity not only to make certain theoretical moves but also to make certain executive and imaginative moves, as well as to have certain feelings'.[7] In a sense, I could not agree more: believing must presuppose such a disposition if my incremental account is to have any plausibility at all. But the dispositional view does not want a relation of presupposition between belief and the required disposition; it wants some sort of identity. This is why the dispositional view is both too weak and blatantly incomplete. It is too weak because it accommodates too much and therefore says very little which is specific about belief. The dispositional view, in Ryle's very formulation, also fits other

[7] G. Ryle, *The Concept of Mind* (New York, 1949), p. 135.

mental attitudes, such as desiring and planning. This is why it lacks the conceptual specificity needed to isolate and explicate belief.

The weakness is not only conceptual, it is also explanatory. Consider the standard counterfactual analysis of belief proposed by the dispositional view: '*S* believes that *p*' amounts to 'If it were the case that *q*, *S* would (do, think, imagine, choose, feel, etc.) that *r*'. For example, 'Babette believes that Pusha wants an ice-cream' amounts, on this analysis, to 'If Babette had money, she would buy Pusha an ice-cream'. If this analysis is meant to be reductive or eliminative of belief (and other mental attitudes), as it typically is, then it is not going to work. The point is familiar, but I want to enlist it here in the service of my position. The fact is that the presumed disposition to act (think, imagine, feel, or whatever) which characterizes belief would not be activated unless Babette's other relevant attitudes, other beliefs too, were brought into the picture. Thus, for instance, the analysans does not explain anything unless we further assume that, in the context, Babette also *wants* to please Pusha (after the Bruckner-related punishment) and *believes* that buying her an ice-cream would achieve the desired result. Further beliefs and desires may be needed to fix these beliefs, and so on. So, even if belief were a sort of disposition, it would not be activated, and would not become cognitively and behaviourally efficacious, without an aggregation of other interacting attitudes being activated at the same time. But, as argued earlier, this is already a characterization of what is going on as a performance. Saying, by way of analysis, that the other attitudes are also dispositions is not going to help because, even if they were, it is again their occurrent aggregation and interaction, not merely their being dispositions, which have been shown to do the explanatory work. In other words, we have to go from disposition to performance to specify what information a representation encodes and what role it plays, hence, whether it is a belief and what belief it is.

This brings us to the incompleteness of the dispositional view. The view leaves a huge gap between the existence of a mere disposition (which says something like, 'Form a belief whenever the conditions are appropriate', meaning, 'Do *x* whenever *y*, other things being equal') and the fixation of the specific mental content which is going to do the real work of a belief. How is the gap between disposition and content going to be bridged? The dispositionalist faces a dilemma. Each horn of the dilemma forces him to

go beyond talk of mere dispositions. The question we are asking is, how does a disposition encode a mental content? The encoding can be either explicit or implicit. If it is explicit, we are really talking of mental representation. We may be talking of syntactic and intentional dispositions to form representations but, as we saw, this is hardly talk of belief. Or, more likely, we may be talking of memory representations, given that the point of the dispositional talk is to emphasize stability and latency of belief. We shall see shortly what this implies. If, as is more likely, the encoding meant by the dispositionalist is implicit or procedural (not a representation, rather a way of handling representations), then he may be talking of primitive or architectural constraints on information processing. But such constraints, it turns out, can plausibly be attributed only to the Representer's modules. In that case, the encoding of belief is modular, not central. Modular belief, I argue, is not a sensible notion. Modular representations are certainly necessary for belief but they cannot be beliefs. Procedures to handle representations can also be acquired and stored. But such procedures only characterize the attitudinal part of a belief, not its content.

To summarize, then: if belief is analysed as a central disposition or procedure, then it is too weak a notion, for it fits many other attitudes, and is explanatorily incomplete, for it does not show how the disposition or procedure encodes the content which does the cognitive and behavioural work expected of a belief. As a disposition, belief will do better if located either in the Representer or in Memory. But this new choice is dilemmatic, as neither component displays genuine doxastic virtues. Let us see why.

Memory?

Memory belief is a tricky topic, so I will try to work in successive approximations. Let me begin with a blunt implication. If belief is an attitude towards incremental information (i.e. towards something specifiable in terms of some uncertainty, some issue, competing alternatives, some background knowledge, and so on), then a belief is stored as such only to the extent to which incremental information is stored as such. But it is very unlikely that memory stores incremental information, therefore it is very unlikely that memory stores beliefs. The main reason memory is unlikely to store incremental information is that, as we saw, the latter is manufactured in a context where a current issue meets available knowledge

and resources of handling it. True, most of the latter are stored in memory but they are not necessarily, not even likely to be, stored in the very form or configuration needed to specify the incremental information involved. We store representations and computational resources out of which incremental information can be manufactured in a context, but it is not obvious at all that we store those representations and computational resources explicitly, in arrangements and configurations which *already* specify particular forms of information. There are practical and principled limitations facing such explicit storage. Practically, it is quite uneconomical for memory to store a huge variety of explicit configurations of representations specifying so many forms of incremental information when it is much easier and possible to generate those forms of information when necessary. The principled limitation is this. Incremental information is defined relative to, among other things, an uncertainty to be reduced, some competing alternative, and a pool of relevant data. It is very hard to make sense of the notion that memory constantly stores representations of specific values of these parameters. We may store representations of particular issues or problems and contemplated solutions, in case we have not made up our minds or for some other such deliberate reason, but this is typically temporary. Once we have reduced or eliminated some current uncertainty in a context of cogitation, we move on to other issues. We may, and probably do, retain general procedures of dealing with similar classes of issues (cogitational routines, if you like), but that is still different from retaining and storing beliefs. The procedures, unlike the beliefs, are ways of handling, not representing, information.

So we have a conflict between two powerful intuitions. One is that memory stores a lot of beliefs. The other is that an essential aspect of belief content is information, most likely in something like the incremental format considered in this paper. But it looks as if our current notion of memory cannot accommodate both intuitions. One intuition has to go. Since I find the connection between belief and information quite holy and untouchable, I am prepared to let the notion of memory belief go. So my first blunt conclusion is: there are no memory beliefs.

Before I weaken this conclusion, I want to strengthen it a bit further. For memory to store a belief it must store a configuration of representations encoding some information. My earlier argument

was meant to undermine the plausibility of the view that memory can store such information by storing the required *configuration* of representations. Now I want to undermine the plausibility of the *antecedent* presupposition that in general memory stores, explicitly, the very *representations* which constitute particular beliefs. This brings us to the question of tacit versus explicit beliefs.

Do I believe that 465 is larger than 2? Now that I ask myself and quickly go through the incremental moves required (or transfer moves made in similar contexts), yes, I do. Did I believe it two minutes ago, before I asked myself this question? How? Explicitly? Not likely. I do not think that the specific representation '465 is larger than 2' (in this or any other form) was ever copied by my memory because it was never sent there. This is because I have never thought of it before. And even if I once formed it (say, one dull afternoon at school), I trust that an efficiency device in my memory quickly disposed of it, as it probably does with most of our very specific representations. But did (do) I believe that 465 is larger than 2 implicitly or tacitly? In what sense? As a potential but never represented implication of some basic truths of arithmetic which I also happen to believe? This will not do. First of all, it is not wise to extend the notion of belief to tacit, unrepresented logical implications, for the simple reason that a content which is not represented in some form is a content which does not contain information, and belief *is* an attitude to information. Second, even if we decide to extend belief, tacitly, to logical implications, as many urge, we should note that the decision is motivated on syntactic and semantic grounds. The notion of logical implication *is* a creature of syntactic and semantic considerations. I have nevertheless argued that such considerations underdetermine the attribution of belief because they are incapable of capturing the information a belief encodes.

So how should we think of my current belief that 465 is larger than 2? My suggestion is: as an occurrently manufactured piece of information whose representation was generated by some appropriate process. We have a recognized ability to generate representations for various purposes, including those of cogitation and action. Why shouldn't we assume that we also generate beliefs out of basic representations under information-sensitive rules and constraints. There may well be procedural constraints (in the form of routines, habits, schemes, and so on) which often guide quite rigidly

the manufacture of many of our beliefs about time, space, arithmetic, logical relations, objects, cats, and so on. But this is different from saying that particular, explicit representations are stored in the very form in which they surface as beliefs. It may well be that what we *mean* by 'long-standing' or 'dispositional' beliefs are not explicit, specific-content-encoding representations but, rather, enduring and fairly standardized procedures of generating the representational configurations needed for belief. It certainly looks more rational for memory to work in this way.

The moral of this part of the argument can be formulated as follows: to have a memory belief that *p* is to store the representational resources and have the procedural dispositions to manufacture the belief that *p*, occurrently, in some context of cogitation, not necessarily a conscious one. The content 'that *p*' is of course type-individuated according to several type of parameters, incremental parameters included.

And yet, we may want to say, even though incremental information is occurrently manufactured in the appropriate contexts of cognition, and even though the representations encoding that information are generated rather than explicitly stored in the very form of that encoding, we still seem to store and have long-standing and explicit beliefs. How should we account for them? Here is my line of speculation.

We obviously store much more than we believe. We store representations originating in perception, and in imagining, desiring, or intending various states of affairs. Unless, for some reason, the original cognitive context (perception, imagination, intention, etc.) is itself copied, indexed, or filed in some form, the resulting representations, if stored at all, are likely to be stored without special distinction. In other words, those representations are not intrinsically beliefs or desires or intentions; they are just inert, informationless inscriptions, patterns on the memory disc, as it were. If, however, a representation is stored, since copied or indexed or filed, *as a belief* (or intention, etc.), this must mean that it was initially manufactured as a belief (intention, etc.) in an occurrent context of cogitation. The copy or index is a record of the attitude generated in that context. I am saying, in other words, that a representation is stored as belief only when (*a*) it is the output of some earlier cogitation and (*b*) a copy of the attitude formed in the cogitation context or an index signifying 'believe what follows' or

some explicit filing in a doxastic dossier or address, accompanies the representation in question.

To illustrate: if I have always believed that $2 + 2 = 4$, it is because the representation '$2 + 2 = 4$' was copied or indexed or filed with a record of the original attitude. This package was later reinforced on many, mostly philosophical, occasions and has become fully routinized. As a result, I now have a vicious disposition to yell '*Yes!*' whenever asked whether I believe that $2 + 2 = 4$. (I wish Descartes' demon would one day do something about this.) Such storage is possible. In this sense we have a number of explicit memory beliefs.

But even this conclusion need not be conceded too quickly or at least not without significant modulation. Even if something like 'BELIEVE (that *p*)' is explicitly stored, I want to say that what is stored thus is still a sort of virtual belief. A belief is virtual if it is stored explicitly, as attitude-and-content, but *without* an appropriate context of cogitation. When the latter is supplied, given the record of the attitude, the stored content goes more or less automatically into a belief position, without much cogitation. But the context can still make a difference, even for virtual beliefs of this sort. This goes to show that, when stored, those beliefs are just that, virtual, that is, less than actual. The occurrent context of incremental cognition can, for example, make a difference if the competing alternatives change their strength. If the alternative of God's manipulating my mental representations becomes serious, my virtual belief that $2 + 2 = 4$ may require re-evaluation. Silly sceptical games apart, the point I am making is that virtual beliefs must still struggle to become beliefs in an actual context of cognition.

While conceding the presence of a good number of virtual beliefs, I still find it more plausible to think that memory stores the resources out of which beliefs are manufactured than to think that it stores already manufactured beliefs. This line of speculation is not that outlandish. There is good introspective as well as experimental evidence to suggest that, relative to their originals, memory copies decay fairly fast and become simplified, schematic, shallow. It is common observation that very soon after we read or hear a sentence we typically forget its actual wording, grammar, intonation, theme, and informational context, yet retain for a longer time, or are able to recover quickly, its meaning. Quite often, it takes grammatical features such as passive/active constructions, cleft constructions

and others, intonation, prior text, and surrounding context to shape the incremental information that is being conveyed in a sentence. The fact that these information shapers are not stored is indication that the information they shape is not stored either. We know this from experience, and there is also ample psycholinguistic literature which shows it.[8] But then, if the specific incremental information is missing, so must be the attitudes that are designed to be sensitive to it.

The point just made suggests an interesting parallel with what Fodor says about the outputs of the Representer's modules,[9] (about which more in a moment). Fodor calls those outputs 'shallow' in the sense that only their syntactic and (in my terminology) intentional features are computed. It takes the central mind to read more into them, for example, speech-act force or irony or (the most important to my mind) information. It is worth speculating that, special indexing or filing apart, there may be some sort of reverse process through which memory strips the incoming representations of their cogitative ambience and richness, as it were, and stores them as counterparts of shallow outputs, for further use by the central mind. The picture then is the following: when a particular context of cognitive performance provides the occasion, the Centre receives shallow outputs from the Representer and also reaches into Memory for some more. Given other constraints, the Centre 'reads' both sorts of outputs for the incremental information it needs to handle the issue raised by the occasion. In so doing the Centre 'attitudinizes' the representations conveying that information and makes them into beliefs, among other things.

Modular beliefs?

Does the Representer, in its various modules, form any beliefs? Are there (such creatures as) modular or registrational beliefs? A familiar, and occasionally useful, strategy, in the search for an answer, is to formulate a number of plausible constraints on the notion of belief (constraints which in fact encapsulate the reasons for having this notion in the first place) and then asking which components of the mind can possibly satisfy the constraints and thus exhibit believing. In some form or another we have already used

[8] See E. and H. Clark, *Psychology and Language* (New York, 1977). I discuss some of the evidence in Bogdan, 'Mind, Content and Information'.

[9] Fodor, *The Modularity of Mind*, pp.

this strategy. At this point in the argument, if we look at a few selected constraints on believing, we must conclude that modular belief is a very implausible notion.

As mentioned earlier, arguments have been put forward (particularly by Stich and Fodor) to the effect that belief cannot emerge at the modular level. The basic idea is that, on the one, doxastic, hand, we expect beliefs to be accessible, consciously or not, to various cognitive and behavioural processes; to be inferentially integrated; to interact with other mental attitudes and thus be cognitively penetrable; to be abstracted from (or not necessarily carrying information about) their causal origin, hence not to be bound to a particular sensory domain; and not to be hard-wired. On the other, modular, hand, the Representer's formation and handling of modular representations, prior to the shallow outputs, display precisely the opposite properties. These representations are inaccessible to other processes in the system; are inferentially unintegrated and cognitively impenetrable; are tied to their causal origin, hence bound to a specific sensory domain; and are hard-wired. That should suffice, I suppose, to exclude belief from the modular theatre of cognition. If a modular representation is a mere registration of some input in the form of a syntactic and intentional structure, then such a registration cannot be a belief. This is not because, as many philosophers would put it, such a registration fails to perform a role or engage an intention, or because it is too lawful to accommodate the idea of false belief, but rather because, as treated at the modular level, such a registration fails to display the sort of information we want to associate with belief. The other failures are symptoms of this 'informationlessness'.

But if you still have reservations, consider quickly these further constraints. You surely want beliefs to be possibly false. Not likely, if they are modular. The modules do not make many mistakes and are indeed fairly lawful at the appropriate (not necessarily physical) level of description. Even illusions are typically excellent and lawful modular accomplishments. When you know how the visual module works and the physics of the environment, the stick perceived to be bent in water is what you would expect from a reliable visual module. If the stick were perceived to be straight, you would take the module either to be faulty in a strange way or else manipulated by the central mind, hence no longer a module. You may also want beliefs to be changeable, revisable. Not if they are modular.

Modules do not really understand this sort of cognitive weakness and hesitation, unless you drug them. The fact that cats are animals which seem to hesitate in doorways goes to show that cats are not (entirely) modules. Pusha certainly is not.

Consider, finally, this important constraint on (what can be called) the maximal specificity of a belief. The rough idea is that, other things being equal, we want a belief to have (be) the most specific information which a set of representations can syntactically and intentionally encode and which makes a functional difference to the organism's cognition and behaviour. Suppose we have two syntactic forms P and Q, of whatever complexity, which are intentionally and semantically equivalent. If Q carries CINF and P does not, then Q is to be treated as a belief, but not P. The implication is that if P is a shallow output of the Representer (or one culled from Memory, for that matter) and Q is its central, CINF-added reading, then Q, but not P, is the belief being fixed in that context, in spite of the assumed syntactic, intentional, and semantic equivalence of P and Q. Since, in general, shallow modular outputs do not by themselves convey incremental information, whereas their central readings do—which is why the latter are more specific than the former—beliefs must be centrally debriefed, as opposed to modularly formed, representations.

In canvassing several serious possibilities concerning the mental realization of believing, that is, components of the mind which could functionally execute believing, we have also constructed an argument from elimination. If belief is neither a mere underlying disposition of some sort, nor originally a memory representation, let alone a modular representation, then (given the alternatives) it must be a central representation assigned an informational structure of an incremental sort. Since, moreover, the central mind is an office where input and shallow representations, from either the Representer or Memory, are read and debriefed for specific information and then directed to do some work, belief must be manufactured occurrently, as an output of these very undertakings by the central mind.

Before concluding with some reflections on this story of belief and some of its consequences, let me speculate just a little on one psychological aspect of the story. Neuroscientists have been aware for quite a while that the brain processes information on many channels at once or in parallel. Needless to say, this makes better

and faster processing than serial processing. More recently, cognitive psychologists and workers in artificial intelligence have paid more attention to cognitive models which employ parallel processing. Philosophers have been slower in seeing the relevance of this development. It may not mean very much if one is interested only in, say, the purely semantic or formal aspects of cognition, as most philosophers are. Yet it may mean a lot to one who is also interested in central processes, such as thinking or inference, or, in particular, in the way such processes manufacture information. The assumption of serial processing *may* force the theorist to view each complex representation, such as a sentence or formula, as a self-contained unit fully valued syntactically, informationally, and functionally, hence as a belief. Most philosophers, I guess, tend to take this line. My story of belief, on the other hand, also allows the assumption of parallel processing when it comes to the manufacture of incremental information. The idea is that the CINF-manufacturing programme may access and read several representations at once in order to extract the needed information. None of these contributing representations by itself, not even the principal input representation (typically encoding the new information), is a belief *prior to* the parallel and aggregate manufacture of incremental information. The didactic upshot of this line of thought is that if we become accustomed to the idea of parallel processing underlying central mental processes, then we may also become accustomed to the idea that it takes several syntactic and intentional forms to fix the information to which a belief is sensitive. This in turn may have the salutary effect of containing the temptation to posit beliefs left and right, whenever a representation (*qua* syntactic and intentional form) is assembled, stored, recalled or operated on. Believing takes more than that. But I would urge this very same conclusion even when we start from the premiss of serial processing.

5. SUMMING UP

Where does all this leave the notion of belief? Well, belief is still an attitude to some mental representations having some formal structure, intentional content, and semantic co-ordinates. It is just that the representations focused on by the belief attitude now constitute a configuration which, when read by an appropriate programme, specifies some incremental information which in turn engages the

organism's cognition and behaviour. The argument has been that the belief attitude tracks this sort of information, not just its syntactic and intentional encoding—or rather tracks the last two only *in so far as* they carry some incremental information. So, on my account, one still believes the good old 'that *p*' when one believes something. It is just that the type-identity of what the *p* in question stands for is no longer a mere syntactic and intentional object which is semantically valued, as the standard notion has it; the type-identity of *p* (i.e. of what is believed) is now shaped by an additional dimension, the informational one.

There is one disturbing challenge with which the informational dimension confronts the philosophical notion of belief. The reader must have detected the challenge from the very beginning, but only now are we in a position to state and defend it clearly. Let me formulate the challenge by way of an objection. 'Haven't you given us all along' (you may want to say) 'an account of belief formation or fixation and not really an account of belief as such, that is, an account of what it is to *have* a belief? Surely, most beliefs are formed or manufactured inductively and possibly incrementally. But their genesis must be different from their intrinsic nature as well as from their continued existence and operation in the mind. You give a genetic account of coming to believe, an account with which many would agree. But you give no account of what it is to believe, that is, to have beliefs and deploy them.'

Quite so, I answer. Mine is an account of belief fixation. But, I claim, this is *all* there is to believing, namely, its occurrent fixation. In other words, I claim that you, the objector, beg the whole question. You assume what you have to prove, and what I have tried to disprove, namely, the existence of a principled, explanatory distinction between forming a belief and having a belief. When you have a belief, you in fact form or manufacture its content, even though the procedures used may be standardized and already in place. Beliefs come into existence when certain conditions are met, even though most elements (basic representations as well as procedures) out of which they are manufactured are typically stored but only in a pre-doxastic state, as it were. Do we store thoughts as such? Not likely. We manufacture them when we need them. But beliefs *are*, for the most part, thoughts with functional obligations. If you think of beliefs which are not thoughts, you are probably thinking of procedures and behavioural routines, or of modular

representations (i.e. raw registrations), or, finally, of memory representations. But none of these, we saw, are beliefs.

There is, however, a stronger answer to your objection. It generalizes an earlier argument about memory beliefs. Suppose you accept that beliefs are fixed incrementally in the way I suggest they are, given that beliefs represent information. How is it, then, that what is fixed in this way in an occurrent, incremental context of cognition is *different in type* from what is stored in memory or imprinted as a disposition? The type difference, as I have argued, is due to the absence of specific information in both 'dispositional' and 'memory' beliefs. This clearly means that what I say is fixed is type-different from what you say is stored or operating as a disposition. Which one is belief? Once again, the choice is forced upon us by the conflict of two very strong intuitions: on the one hand, that belief is information, on the other, that beliefs exist in some enduring form (stored, dispositionally, etc.). Since what is information for an organism is contextual and ephemeral, the two intuitions cannot coexist. At least not in the form in which they are currently explicated and theorized about. For me, again, the belief–information combination is the winning combination.

The third answer to your objection comes in the form of an exercise in metaphilosophical diagnosis. Its moralistic point is this: let me explain to you the temptations of calling *belief* something which is less than belief. To get to that answer, we need first a synoptic view of what we have been doing here.

The picture of belief that emerges from this paper can be conveniently summarized in Figure 3. This is no neat way of analysing belief but then belief is no neat thing to analyse. The key notions, those of incremental information and belief, have been given no explicit definition, for I have none to give. The strategy, rather, has been to let a number of constraints implicitly characterize them. Thus understood, the notions of belief and CINF are a feat of abstraction. But so is any notion. What we do in setting up such notions is to build equivalence classes around factors we take to make an essential difference to the type-identity of the phenomenon we study, in our case, mental content and belief. The five columns in the middle of Figure 3 indicate basic types of factors thought to make a difference to what sort of thing a belief content is. It all depends on how one casts the equivalence relation, both within and across the columns. Putnam, for example, has argued

SAME			DIFFERENT	
Representation			Information	Role
Parameters External	Internal			
Semantic	Syntactic	Intentional	Incremental	Functional
causal source input reference	symbols forms computations	meaning concepts sterotypes	theme issue given information	inference intention behaviour
truth conditions	etc.	recognition schemes	new information projection	etc.
World	Representer & Memory		Cogitator (Centre)	
		The Mind		
W	[S +	m	+ CINF] ↦	r
	Mental content			Role
From			To	
		BELIEF		

Fig. 3

that syntactic and intentional parameters are not enough to characterize mental contents and beliefs, in particular, by showing that semantic parameters *do make a difference.*[10] So more types of parameters have had to be included. In the same spirit, my argument pushes further and suggests a still wider, content-defining equivalence class which also incorporates informational parameters of the incremental sort. The strategy was to show that incremental parameters do make a difference to our type-individuation of belief content. Which is what we have done.

The result is not going to be very intuitive. The reason, very roughly, is this. Belief is different things to different people in different contexts. One takes from the notion of belief what one needs in a particular context. That is to say that one concentrates on one or more dimensions at the expense of others. We can visualize this process of selection of dimensions as one in which one slides the notion of mental content across the five central columns in Figure 3, thus setting up appropriate equivalence classes. In this way the columns selected are going to delineate the aspects of mental content to be considered in a context, hence, what is believed in that

[10] See, e.g. H. Putnam, 'The Meaning of Meaning', in K. Gunderson (ed.), *Minnesota Studies in the Philosophy of Science*, Vol. 7 (University of Minn. Press, 1975), reprinted in Putnam, *Philosophical Papers*, Vol. 2 (Cambridge UP, 1975).

context. For example, if we are interested in the truth or actual aboutness of a belief, we are likely to assume the belief's informational and functional dimensions and focus only on its syntactic, intentional, and semantic dimensions. We then treat belief as a syntactic truth carrier. But this does not mean that a belief is (type-identical with) a syntactic truth carrier. It only means that in the context in question we assume that the other dimensions (functional, informational, etc.) have done their prior work and shaped the belief in question, after which we can consider only the output of that shaping and examine it only in the light of its semantic, syntactic, and intentional aspects. In the same abstractive spirit we may, for purposes of psychological explanation, treat belief as a mental form with causal impact, i.e., treat it syntactically, intentionally, and functionally, while assuming the belief's semantic and informational dimensions. But that, again, does not make belief (type-identical with) a mental form with causal role. And so on.

Suppose we now ask, 'When are two beliefs, yours and mine, the same (meaning type-identical)?' Sameness, we know, is a relative notion. We have to specify the aspect under which it obtains. In the case of belief, these aspects correspond to the dimensions just discussed. You and I may have the same belief in the semantic dimension, either because both beliefs are true or, more concretely, because they share the same truth conditions. Or we may share the same belief in the intentional dimension because we apply the same concepts and therefore assign the same meaning to our belief. This relative (type-) sameness may be secured even when the beliefs are different relative to the remaining dimensions, if we decide to ignore those differences for some reason. The linguistic practice is such that we still attribute, and explain in terms of, sameness of belief even though that sameness covers only one or two dimensions, not others. Strictly speaking, this is less than the whole notion of belief, but we do not care. The practice of selecting dimensions of interest and then parading them as full beliefs is so deeply ingrained that we find it counter-intuitive to be reminded that, strictly speaking, belief is a plurality of dimensions. Thus, you will almost surely find the following argument counter-intuitive: 'Strictly speaking, the notion of belief must contain the informational dimension. But then, two or more people rarely share the same type of belief since very rarely are they going to manufacture the same sort of

incremental information, given that very rarely are they going to share the same data base, alternatives, and evaluation metrics, among other things.' The argument seems counter-intuitive because in most ordinary contexts (which is where intuitions are shaped) we choose to ignore informational differences if other dimensions matter more. So, the moral of my story is, not only do we manufacture belief because we occurrently manufacture an essential dimension of belief content, namely, incremental information. We also manufacture the very notion of belief we find convenient and appropriate in a certain context by choosing the dimensions which we think fit that context best.

So far so good. We approach many other mental and wordly phenomena this way. The trouble starts when philosophers, for all sorts of reasons, allow these normal abstractive moves to freeze into programmatic prejudices which explicate belief. This is when belief gets conceptually assimilated to a particular dimension or particular sets of dimensions which happen to fit a certain conceptual paradigm. If introspectible phenomenalism is the fashionable paradigm, as it was not so long ago, then belief is (as a type, no more than) a conscious experience, a feeling, a vivid impression, in brief, a mental representation defined by our access to it. Then behaviourism takes over and, lo and behold, belief is (as a type, nothing but) a disposition to behave, a functional connection between representational input and behavioural output. Central-state materialism becomes fashionable and belief ends up being (as a type, just) a representation-encoding brain state. Now, in the age of language and its philosophy, belief is maltreated accordingly. Since the leading obsessions in philosophy of language are, how does language connect with reality? what and how do words and expressions mean?, and the like, belief itself gets squeezed into its syntactic, intentional, and semantic dimensions, with some acknowledgement of the functional dimension. In ordinary contexts, dimensions neglected are dimensions set aside for tactical reasons. In philosophical contexts, dimensions neglected are dimensions ideologically ignored.

Since many things will change in our understanding of belief, the prudent metaphilosophical moral to be derived from our story is: expect more, not fewer dimensions of belief. If, as I argue, belief is manufactured by the central mind, then the expectation of its multi-dimensionality is not unwarranted. The central mind is not a tidy,

well regulated, and well behaved module, so forget about Occam's razor.[11]

[11] I want to thank my friends at Tulane and the University of New Orleans, and particularly Harvey Green and Norton Nelkin, for asking tough questions, making many observations and shaking their heads in disbelief.