

Protothought had no logical names

James R. Hurford

1. Introduction

The evolutionary history of any complex system, such as human cognition or the human language capacity, necessarily starts with something simpler. This is not to deny that evolution can sometimes simplify. But the dominant evolutionary trend is from simple to complex, and the original protosystems were undoubtedly simple. It is accepted as a working hypothesis that a precursor of the modern human capacity for complex syntactic language was a capacity for protolanguage (Bickerton 1990), a kind of communication system with no syntax. In protolanguage, although words may have been uttered in short sequences, there were no rules defining the wellformedness of strings, and therefore words in protolanguage could not be said to belong to separate syntactic classes, such as Noun or Verb.

A view commonly encountered about the emergence of true syntactic language is that during the protolanguage period there existed a somewhat complex system of mental representation, capable of representing the structure of events with their participant agents and patients. In this view, the move to syntactic language was in large part an externalisation of this pre-existing system of representation: "The mechanism [of syntax] was there all the time, but it was not being used for language" (Bickerton 1998: 350). Let us call such a system of mental representation "protothought". It is often assumed that the structure of this protothought was something like Predicate Calculus (but possibly without quantifiers).

2. Predicate logic, semantics, and psychology

Assuming something like a modern Predicate Calculus form of mental representation for our prelinguistic ancestors actually attributes to them at least a five-way distinction between types of atomic mental entities, as the following simple formulae illustrate. The first formula: $CAME(john)$. The translation of this would be 'John came'. The second formula: $\exists x[TALL(x) \& MAN(x) \& CAME(x)]$. The translation of this would be 'A tall man came'.

In such representations, one has to distinguish the following types of atomic elements, each of which functions in the representations and relates to external meanings in a distinctive way:

- predicates (CAME, TALL)
- individual constants (john)
- individual variables (x)
- connectives (&)
- quantifiers (\exists)

The proposals which derive the basics of modern syntactic structure from such prelinguistic mental representations are seldom fully explicit on points of detail, but typically seem to assume that individual constants are the evolutionary sources of Proper Nouns and perhaps also of Common Nouns, whereas the predicate constants are the sources of Verbs. One account which is completely explicit on this point is found in the computer models of Simon Kirby (Kirby 2000a and 2000b). Here it is clear that individual constants are the evolutionary source of what emerge as Proper Nouns, and predicate constants the source of Verbs.

Such an apparatus for prelinguistic mental representation is implausibly rich. I will argue that a reduced apparatus is more plausible. Specifically, in my less complex format for mental representations, there are only predicates, individual variables (but no individual constants), just one connective (namely conjunction: &), and just one quantifier (the existential “ \exists ” implicit in all representations). The implication of this proposal is that formulae such as the first given above, CAME(john), were not available as prelinguistic mental representations, specifically because they contain individual constants. Why do I believe that individual constants cannot be justified as elements of mental representation pre-existing language? A brief look at the place of predicate logic in modern scholarship will help to set the scene.

The originators of modern logical calculi did not intend them to be taken as psychological entities. Frege, for example, was averse to “psychologism”, believing that individuals’ private concepts are beyond investigation. The central tradition stemming from Frege interprets logical formulae by a direct relationship with a presumed objectively given world, thus bypassing any psychological stage which might be located in a user’s mind or brain. In this view, meaning is a relation between (logical) language and the world; no mind intervenes or mediates in this: “According to Montague, the syntax, semantics and pragmatics of natural languages are branches of mathematics,

not of psychology" (Thomason 1974: 2). This approach to meaning has been characterised (and criticised) by George Lakoff as "objectivist semantics" (Lakoff 1988: 119-154).

The designers of Predicate Calculus assumed, surely correctly, that the world contains individual objects like the sun, the moon, particular people, dogs, cats, tables, and chairs. These individual objects are the denotata of the individual constants in logical formulae, such as *john*. Since, in this tradition, no minds mediate the relation between language and the world, the function which assigns denotations to individual constants is merely stipulated; it is typically explained by the author of a text on logic with enough ordinary-language prose to make the idea of an inherent (or stipulated) link between the logical name and the object it denotes clear enough to the reader. Often, the author conjures up some fictitious little world with fictional characters in it, and can be confident that the reader will get the idea. Textbooks such as Cann (1993), Guttenplan (1997), and McCawley (1981) as well as more advanced accounts of meaning in logic assume that giving a real semantics to sentences involves mapping them to an assumed objective world, without involving any psychological entities such as concepts, thoughts, or ideas.

When linguists began to build generative models of languages including their semantics, it seemed appropriate to assume that the classical and simple languages of logic, such as Predicate Calculus, could be taken in their entirety and used as mental representations of meaning, much as standard IPA phonetic transcriptions can be used to represent pronunciation. Generative grammar, as indeed most modern schools of linguistics, aims to model linguistic competence, that is, psychological representations of language in speakers' heads. But it cannot be safely assumed that a formalism, such as predicate logic, which was not designed to suit psychologistic purposes, will be appropriate for such purposes. A semanticist who accords psychological reality to semantic representations like logical formulae, faces not only the task of stating what in the world the components of these representations map onto, but also the task of accounting for how these mappings could come to be known by the language user.

Theorists of language evolution who assume something like Predicate Calculus representations preceded language in the minds of early hominids also face the problem of accounting for the mappings between the mental entities (such as predicates and individual constants) and parts of the world. An account of how mental predicates (LITTLE, GREEN, MAN) come to be mapped onto properties of the world is, I claim, relatively unproblematic,

whereas an account of how mental individual constants (john, mary) come to be mapped onto particular individuals is impossibly problematic.

3. Perception and cognition of objects

It is crucial to distinguish between perception and cognition. Mental representations of events and their participants are involved differently in these two psychological processes. Let us start with perception or, more precisely, “perception-as”.

3.1. Perception

3.1.1. Unfocussed perception

Some experiences come uninvited. A creature may notice a particular smell, or feel cold, or hear thunder. In prelinguistic hominids, these experiences would undoubtedly have caused different reactions and can be regarded as separate mental predicates whose denotations are the classes of stimuli giving rise to them. In hearing thunder or noticing a smell, the perceiver perceives only the smell or the thunder; no other object or person is perceived. Such perceptions can be represented schematically as isolated zero-place predicates: SMELL₃₁, COLD, and THUNDER. These could be paraphrased in modern English as ‘There is a certain smell (around here)’, ‘It’s cold’, and ‘Ah, thunder’. Note the lack of any contentful grammatical subjects in these paraphrases.

If the creature has no language, there is no question of these categories being connected with any words. I assume that nonlinguistic creatures have representations which involve mental categories and that they are able to judge whether or not perceived objects belong in these categories. Such mental categories are partitions of an innately given structured space of possible categories, acquired through experience. Activation of different categories, whether by immediate experience or by rumination, will stimulate different patterns of behaviour; this is what enables nonlinguistic creatures to behave systematically in relation to the world. Our imagined creature somehow carries categories like SMELL₃₁, COLD, and THUNDER in its mind. Their existence in its mind has been distally triggered by its previous experiences. In the act of perception, these categories are reactivated.

3.1.2. Attending to perceived objects

We now bring particular objects into the discussion. For concreteness, let's say a creature looks, with some attentive deliberation, at a fresh green sycamore leaf. In this case, the focus of attention is the leaf itself, which is an object in the world. In attending to it and judging it to be a green leaf, the creature applies certain mental categories to it, which we can relatively safely represent as the mental predicates GREEN and LEAF. Let us assume that this much is possible, the creature's species having evolved to be able to recognise greenness and the class of leaves.

Just as with thunder and coldness, the categories GREEN and LEAF are somehow carried in the creature's mind, having been planted there by its previous contacts with the world. In the act of perceiving a green leaf, these mental categories are brought to bear on an object of attention in the external world. It is natural, for illustrative purposes, to represent the principal elements in such an act of perception by a formula such as GREEN(x) & LEAF(x) where x denotes the real-world object attended to, and the predicates correspond to the mental categories which the creature judges to apply to this object. (To the best of my knowledge, the first use of such formulae, suggested as prelinguistic representations of meaning, is in Batali 2000.) Such a formula, then, has a mixed, partly external objective, partly mental, semantics. This is appropriate because an episode of "perception-as" is an encounter between mental categories and real-world objects. The x in the formula is essentially deictic (or "indexical"). It is not the constant name of any particular object. It only stands, in this convenient expository representation of an act of perception, for whatever object happens to be the object of perception. If on another occasion the creature sees a different green leaf, the same representation as before, "GREEN(x) & LEAF(x)", would be appropriate, as this is another act of perceiving an object as a green leaf. If I may call this x a variable, it is attention that associates the variable with a real perceived object in the world.

The same mental predicate can be either a zero-place predicate or a one-place predicate, depending on whether the stimulus comes from an object to which the perceiver is attending. Thus if the creature deliberately sniffed the leaf, its perception might also include SMELL₃₁(x).

There is a fuzzy borderline between attention to a specific object and awareness of ambient impressions. Attention involves selection of one source of stimulation over others which have an equally good chance of being selected or attended to. A creature might decide to have a good look at a hole

in a tree (as opposed to the tree's bark or an upper branch). A hole in a tree is something of a size that one can reasonably pay selective attention to. Compare this with a large cave, especially one in which the creature is sitting. As the cave is all around it, it is not clear whether to count a perceptual judgement about the cave as an unfocussed perception or the perception of an object attended to. That is, it is not clear whether this should be represented as $\text{DARK}(x) \ \& \ \text{CAVE}(x) \ \& \ \text{SMELLY}(x)$ or as $\text{DARK} \ \& \ \text{CAVE} \ \& \ \text{SMELLY}$. (To keep things simple, I allow the connective "&" to conjoin more than two conjuncts. Because I will not be dealing with disjunction, this does not at all affect the argument.) It will pose no problem to the account proposed here if we allow that mental representations cannot be fully determined, or theorised about, in all cases. Either the creature represents its perceptions as simply ambient (the zero-place case) or more specifically about the place it is sitting in as opposed to some other place it has chosen not to attend to (the one-place case).

3.2. Cognition: representing previously perceived objects

The important aspect of cognition, for our purposes, is the memory or store of experiences which are no longer current. Consider how the nonlinguistic creature, while not actually perceiving a green leaf, stores in memory the fact that it did once see a green leaf. When objects are no longer attended to, the association of an external object with an attention variable is lost, since it is the very act of attention which brings about this association. In the transition from perception to cognition (that is, the process of putting into memory), then, the variable loses its former (extensional) meaning. The creature can remember that it once had a " $\text{GREEN}(x) \ \& \ \text{LEAF}(x)$ " experience, but there is now nothing pointing, as attention once did, to the object itself. The closest that Predicate Calculus can get to this is formulae with existential quantifiers, like $\exists x[\text{GREEN}(x) \ \& \ \text{LEAF}(x)]$ and $\exists[\text{DARK}(x) \ \& \ \text{DAMP}(x) \ \& \ \text{SMELLY}(x) \ \& \ \text{CAVE}(x)]$. The second formula here might be paraphrased as '(I remember) there was a dark damp smelly cave'. Memory of complex experiences (but not involving more than one object) is captured by conjunctions of such atomic proposition-predicates. Each conjunction or bundle of features is a separate entry in the mental database, a memory of a separate experienced event.

Implicit in this account, and the central plank of the whole argument, is the theme that in prelinguistic cognition there are no components of mental

representation with the logical status of individual constants, that is, elements, that, as the philosophers say, allow their possessors to re-identify particulars with absolute reliability.

3.2.1. There are no particulars in animal cognition

The individuals most important to us are our near kin, mothers, and children. The following quotation demonstrates the *prima facie* attraction of the impression that animals distinguish such individuals, but simultaneously gives the game away:

The speed with which recognition of individual parents can be acquired is illustrated by the “His Master’s Voice” experiments performed by Stevenson et al. (1970) on young terns: these responded immediately to tape-recordings of their own parents (by cheeping a greeting, and walking towards the loudspeaker) but ignored other tern calls, even those recorded from other adult members of their own colony (Walker 1983: 215).

Obviously, the tern chicks in the experiment were not recognising their individual parents. They were being fooled into treating a loudspeaker as a parent tern. For the tern chick, anything which behaved sufficiently like its parent was recognised as its parent even if it wasn’t. The tern chicks were responding to very finely grained properties of the auditory signal and apparently neglecting even the most obvious of visual properties discernible in the situation. In tern life, there usually aren’t human experimenters playing tricks with loudspeakers, and so terns have evolved to discriminate between auditory cues just to the extent that they can identify their own parents with a high degree of reliability. Even terns presumably sometimes get it wrong: “[A]nimals respond in mechanical robot-like fashion to key stimuli. They can usually be ‘tricked’ into responding to crude dummies that resemble the true, natural stimulus situation only partially, or in superficial respects” (Krebs and Dawkins 1984: 384).

A creature can appear to recognise a particular individual reliably if it gets sensory impressions which are finely grained enough to distinguish this individual from others with which it is likely to come into contact. To manage successfully in “recognising its mother”, a creature need not have a mental distinction between individual and property. All that is needed is the ability to distinguish just the set of properties which in practice differentiates its mother from other candidates.

English dictionaries often distinguish between two senses of the word “recognise” as below:

1. If you recognise someone or something, you know who someone is or what someone is when you see them or hear them, because you have seen them or heard them before. EG *She didn't recognise me at first...*
2. If you recognise something ... you are able to identify it when you see it, for example because you have learned about it or had some previous experience of it. EG *They are specially trained to recognise the symptoms of radiation sickness... There was a field planted with what he recognised as maize (Collins Cobuild English Dictionary).*

This lexicographic distinction rests on the ontological distinction between individuals or particulars (I use the terms interchangeably) and universals. The central claim here is that prelinguistic creatures could not have had mental representations containing elements corresponding to particular individuals; they could only recognise things in sense 2 above, but not in sense 1.

3.2.2. Individual particulars, universal properties

Prelinguistic creatures could not have been able to re-identify particulars with absolute reliability because, in fact, neither can we. Given what we know about the world, it is, in the limit, impossible. This is not to deny that much of the time we are quite successful at re-identifying particulars, especially where it matters. But situations where re-identification is impossible are also common. For instance, every morning I pick up a bottle of milk from my doorstep. If I have a backlog of undrunk milk from previous days, I add today's bottle to the two or three that are already in the fridge in a particular spatial sequence so that I know which bottle is which. If someone disturbs my milk filing system, then I no longer know which is which because I cannot re-identify particulars (my dairy does not datestamp its products). Where it really matters to us to keep individuals apart, we either mark them somehow (as with my milk filing system) or we develop more refined perceptual categories so that we can remember the difference between individuals in terms of their slightly different properties. This is presumably why humans are so proficient at face recognition. But it is both impossible and unnecessary to refine one's perceptual categories to the point where every individual object is distinguishable from all others. One paperclip is

much the same as another, as are sheets of paper, grains of sand, first-year students, and so on.

The argument above has concentrated on the difficulty with logical individual constants as components of mental representations. The impression may have been given that logical predicates are not problematic in the same way. Philosophers have emphasised to me that, in the limit, predicates are also subject to similar problems. Putnam's Twin Earth argument (Putnam 1975: 223-247) establishes that there could be objectively (that is, from a God's-eye view) two different sorts of stuff that we label water. Our concept of water could fit both kinds of stuff, and the stuff that is chemically XYZ could, to all intents and purposes, serve exactly the same range of purposes as the stuff that is chemically H₂O to the degree that we wouldn't know the difference: "Extension is not determined by psychological state" (Putnam 1975: 222).

There is, however, a great practical difference between an ability to re-identify particulars (such as the precise lion that was stalking me yesterday) and an ability to re-identify universals (such as lions in general), and hence a great practical difference between individual constants and predicates.

Invariant factors in the environment in which humans have evolved have given rise to our perceptual space and the discriminability of regions within it. The species has evolved to see just the colours it can see, hear just the sounds it can hear, smell just the smells it can smell, and so forth. Despite some individual variation, the same inputs from the world trigger broadly similar perceptions in everyone, across a wide range of possible inputs. More importantly, there are relatively few cases in which objectively disjoint stimuli trigger the same perception (objectively, that is, to the best of our current scientific knowledge). Putnam mentions such a case: "Although the Chinese do not recognise a difference, the term 'jade' applies to two minerals: jadeite and nephrite. Chemically, there is a marked difference. Jadeite is a combination of sodium and aluminium. Nephrite is made of calcium, magnesium, and iron. These two quite different microstructures produce the same unique textural qualities!" (Putnam 1975: 241). Examples such as jade pose no problem for a theory that prelinguistic creatures could have had mental predicates corresponding to various classes of stimuli. Their mental predicates would have had a perceptual and functional basis. If it looks like a duck, sounds like a duck, smells like a duck, feels like a duck, and tastes like a duck, it's a duck.

Properties are procedurally prior to individuals. That is, we use an object's properties to try to determine which individual it is. We humans are often obsessed by this question of which individual is presenting itself to us. We can never be certain, but we get as close as we can by finer and finer

subclassifications of the properties we are observing. In life's ongoing calculation of "what is happening", the fundamental empirical premises are the properties which the world presents to us. The conceptual identities of individuals appear only in our conclusions.

Central to the concept of an individual are the related ideas of uniqueness and spatio-temporal continuity. These are not part of the idea of a property. Many things are red, in different places, at the same time. But an individual cannot be in two places at once. Consequently, the evidence needed to assign the same property (a mental predicate) on different occasions is weaker than the evidence needed to conclude that something is the very same individual. To assign a property, in many cases all we need is the evidence of our senses: "No doubt about it, this is green". But without keeping a constant eye on an object or insulating it somehow, we can never be sure that it hasn't been replaced by a twin: a different individual.

3.3. *Language hints at knowable identity*

Modern human languages contain devices which are interpreted as if their identification of particulars posed no problem. Proper names are the prime example. But in fact our ordinary usage with proper names is more rough and ready than the usage of logicians with individual constants. We all know many people called *Jim*, and usually the context of the speech situation clarifies which *Jim* we are talking about. Thus *Jim* is in fact only a shorthand for whichever of the people who have been regularly labelled *Jim* in the experience of the interlocutors is most salient in this context. Clearly, *Jim* in English does not pick out a single individual, though doubtless there are some rare names which do happen to pick out single individuals. But we would not wish to distinguish logically or grammatically between different proper names on the basis of their relative rarity of application. Markers of definiteness, like English *the*, are also basically shorthand for whichever of the things concerned is most salient in this context.

Individual objects are less enduring over time than properties. The objects to which proper names are most commonly attached, namely people, only live for a few generations. But almost all of the properties which we see in an individual person can be seen, in different combinations, in other people, long before and long after the brief life span of the particular individual. *Ars longa, vita brevis* might be adapted to the case of properties and particulars, which is related to why, in philosophical discussions, properties

are often called universals. The relative impermanence of particulars, as opposed to properties, undermines the utility of constant expressions for particulars in the continuing language of a community. Place names are the most useful designators of enduring individual objects in modern languages, but our early ancestors would have had much less use for them. In any case, there is evidence that in modern languages all place names in fact derive historically from descriptions, that is, from expressions involving property terms, like *the dark pool* or *the new castle*. Probably, all person names are ultimately so derived as well.

Linguistic devices suggest that humans are capable of mentally representing unique individuals using the mental equivalent of logical individual constants. We might even grant, for the sake of this argument, that modern humans do have this capacity. But it is only language that pushes us toward this conclusion. Without an insight into our language, outside observers of modern human behaviour could not tell that we are able to mentally represent particular individuals. We, in our turn, when we consider the minds of nonlinguistic creatures, must adopt the simplest hypothesis compatible with the observed facts. There is no need to hypothesise mental individual constants in animal minds. When animals respond systematically to impressively finely graded stimuli – even to the extent of reliably recognising what human observers believe to be unique individuals – there is no need to postulate a special category of mental representation like individual constants. Finely graded predicates or conjunctions of less finely graded predicates will do the job.

4. Social life without individual constants

We are considering the emergence of language with modern characteristics from pre-existing cognitive capacities in prelinguistic creatures. I have maintained that the cognitive representations available to prelinguistic hominids for remembering the events of their lives are most unlikely to have contained the equivalent of logical individual constants. A more plausible form of prelinguistic mental representation would have had only predicates, individual variables, a conjunctive connective, and an implicit existential quantifier. The baby australopithecine or habiline would have had many stored memories of episodes involving its mother, many containing a form something like this: $\text{SMELL}_{31}(x) \ \& \ \text{SQUAWK}(x) \ \& \ \text{WARM}(x) \ \& \ \text{CUDDLY}(x)$. (Here and henceforth, the existential quantifier is omitted from such re-

presentations. It is implicit. The variables, however, cannot be eliminated because they are needed for representing memories of events involving several different objects.)

Our immediate prelinguistic ancestors certainly had complex social lives, as do modern chimpanzees. In modern human terms, they remembered who were their friends and enemies, who was whose child, who had recently fought with whom, and who had groomed whom. But the *whos* in this last sentence tend to impose our modern logical analysis on the animals' behaviour. It is possible, with finely graded predicates, to represent as accurately as is necessary any of the information that australopithecines or habilines would have stored about members of their social group. That this is generally true is shown by a transformation that one can apply to any first-order logical formula containing individual constants, replacing the constants by corresponding (suitably fine-grained) predicates. For example, CAME(john) can be replaced by CAME(x) & JOHN(x) where JOHN is a predicate which is only satisfied by one individual. Of course, expressing it in terms of a predicate being "only satisfied by one individual" is again to express things from a modern, language-informed perspective. To a prelinguistic hominid, as to a tern chick, it was only important that the mental predicate be specific enough to be reliable in all the circumstances in which it was likely to be used. Two-place predications involving different individuals can be handled in this way, but with distinct variables. Thus *Kanzi is Panbanisha's child* could be represented as KANZI(x) & PANBANISHA(y) & CHILD(x,y). Clearly, representations without individual constants are sufficient to account for the Machiavellian intelligence of apes.

It is assumed here that our prelinguistic ancestors were capable of forming associations between different concepts, allowing inferences between thoughts involving different predicates. Thus SMELL₁₄(x) → EDIBLE(x) might have been a mental rule useful in foraging. Similarly, "portmanteau" predicates could have been created, abbreviating longer conjunctions of more basic predicates. Thus the baby australopithecine we mentioned earlier could have had a rule such as [SMELL₃₁(x) & SQUAWK₇(x) & GAIT₁₂(x)] → MOTHER(x).

5. So what?

How does this argument relate to other issues in language evolution? The origin of syntax is a central problem. At the heart of syntactic structure is the

Subject-Predicate dichotomy. So any claim about the origin of the Subject-Predicate dichotomy is central to theorising about the origins of modern human language.

There are two tendencies in the search for explanations of features of language, whether these are evolutionary or otherwise. One tendency is to look in the direction of possibly innate, not necessarily communicatively functional, characteristics of the minds which gave rise to language. Bickerton, in particular, has stressed the representational function of language, as opposed to its communicative function. Accounts of the rise of Subject-Predicate structure from the pre-existing-representation point of view assume a pre-existing dichotomy between two types of concepts, a predicate type, and a subject (or argument) type. Individual constants are the principal instances of concepts of the subject/argument type. In this view, pre-existing propositional representations of Subject-Predicate form get rather straightforwardly translated into linguistic clauses with corresponding structure.

Such proposals push the question of the origin of Subject-Predicate structure back into prelinguistic representations, and that is considered far enough to absolve the proposers from further explanation. But as I have shown, it is unnecessary to attribute such pre-existing mental structure to prelinguistic creatures. They could have got along fine without it, and so they probably didn't have it.

The opposing tendency in linguistic explanation is to appeal to communicative function. Communicative function cannot be a consideration relevant to prelinguistic (or precommunicative) mental structure. Pressure from communicative function can only arise once mental representations are externalised in public, conventionalised symbols. From this perspective, an alternative explanation for the origin of Subject-Predicate structure is communicatively functional Topic-Comment structure. It must be emphasised that Topic and Comment are not categories of classical (modern) logical systems. Topic and Comment are categories rooted in pragmatics, in the shared knowledge of speakers and hearers, and thus are not, and never have been, the concern of truth-conditionally oriented logicians.

The argument presented in this paper has been largely negative. I have argued that our prelinguistic ancestors could not have represented the world in terms of formulae involving individual constants, where these individual constants were destined to become the evolutionary antecedents of modern referring expressions by a process of externalisation into public language.

This negative argument is a prelude to a positive argument that will describe a mechanism whereby essentially communicative pressures gave rise to the emergence of modern Subject-Predicate structure, based on prior representations in which the only constants were predicates. In brief, some predicates became selected as prototypically serving a Topic function, while others became selected as prototypically serving a Comment function. And this happened concurrently with the emergence from protolanguage of longer, more coherent strings of words. But that is another story.

Acknowledgements

This work was supported by grant (Number R000237551) from the UK Economic and Social Research Council and by a Fellowship at the Collegium Budapest Institute for Advanced Study. I have benefited from helpful comments and advice from John Batali, Manfred Bierwisch, Peter Carruthers, Terry Deacon, Daniel Dor, Dan Nettle, Bruce MacLennan, Luc Steels, and Paul Taylor, but don't blame them for any faults in the paper.