

To see that the relevant moments in a durational adverb are contextually determined, note that (83) is not odd in the same way as (84):

- (83) John has been working in San Diego for the last five years.
He usually spends his weekends at the beach.
- (84) ?John has been serving his prison sentence for the last five years.
He usually spends his weekends at the beach.

Because of our knowledge of facts about the real world, we know that the relevant moments included in *the last five years* in (83) do not include weekends, vacations, etc., whereas the relevant moments covered by the same quantifier in (84) are much more inclusive. I doubt that anyone would claim that the time adverb itself has a different logical structure in (83) and (84). I realize that "relevant psychological moment" may sound like a vague notion at this point, but it seems that we must either adopt it for the time being or else stop using the universal quantifier to represent durational adverbs. Note that in the analysis presented above the actual number of moments in an interval is not important; as long as there are at least two, then (75) is contradictory.

A second objection to the analysis would be that there are potential counter-examples to it in the form of sentences like (85):

- (85) John found his son's tricycle in the driveway for six weeks.

(85) appears to be well-formed, despite the fact that it contains an achievement verb, a durative time adverbial, and no indefinite plural or mass noun.

Part of the solution to this problem is that (85), on its acceptable reading, is understood to be elliptical, in that a second time adverbial of some kind is implicit:

- (85') John found his son's tricycle in the driveway

}	every day
	once a week
	frequently
	etc.

for six weeks.

That is, the different occasions of "finding" are separated by intervals. (The same observation should perhaps be made about (77), but this is only part of the difference.) I am not sure what the best way of handling this matter is.

A second difference between (85) and (75) is that *discover* in (75) is more likely to mean "come to know the existence of" whereas *find* in (85) is more likely to mean "come to know that NP is at x place at y time." Coming to

know the existence of something is a once-and-for-all event (barring memory lapse), whereas an object that reappears in unexpected places presents ever-new "facts" to be discovered. The nouns *buried treasure* and *tricycle* were thus not chosen at random; a buried treasure, once discovered, is not likely to surprise one a second time by reappearing unexpectedly, but a tricycle is just the kind of object that would. The main claim, that one does not discover the same fact more than once, still seems valid, and I think that this way of treating the anomaly of (74) vis-à-vis (77) is viable in spite of the problem of vagueness in durative adverbs.

2.3.4. Carlson's Treatment of 'Bare Plurals'

Though the treatment of indefinite plurals and achievements just given (which come from Dowty 1972) seems adequate as far as it goes, it leaves one important question unanswered: if it is correct to analyze an indefinite plural like *fleas* as involving an existential quantifier (i.e. as equivalent to *a flea* or *some fleas*), then just why must this quantifier have narrower scope than the durative adverbial? From what has been said, it might be supposed that the contradictoriness of the wide scope reading with an achievement verb is all that prevents this second reading from being apparent, but this is not so. Examples like (86) and (87) (from Carlson, 1977, p. 27) have stative and activity verbs respectively, yet the (b) sentences only appear to have readings in which this putative existential quantifier has narrower scope than the adverbial, while the (a) sentences with an explicit quantifier *a* or *some* clearly have a reading with the existential quantifier taking wider scope (as well as perhaps a less obvious reading with the quantifier taking narrower scope):

- (86) a. { A cat has } been here since the Vikings landed.
 b. Cats have been here since the Vikings landed.
- (87) a. { A tyrant } ruled Wallachia for 250 years.
 b. Tyrants ruled Wallachia for 250 years.

This is only the beginning of a long story, however. Carlson (1977) examines a number of quantifier-like constructions (negation, other NP quantifiers, durative and frequentative adverbs, aspectual verbs like *continue*, anaphoric constructions) that might be expected to bring out a scope ambiguity with

these indefinite plurals, and in every case the only possible reading is one in which the “existential quantifier” underlying the indefinite plural appears to have narrower scope than the other quantifier or operator.

A further peculiar fact is that indefinite plurals (or what Carlson calls *bare plurals* following Chomsky) elsewhere seem to be interpreted as having a kind of universal, or *generic* quantifier, yet it is hard to find a single sentence (at least in certain tenses) in which the bare plural is truly ambiguous between an existential reading (as in (86b), (87b), and earlier examples) and a generic reading. The sentences of (88), for example, have to be taken as referring to smokers, cats, or elephants in general, not just a particular group of smokers, cats or elephants:

- (88) a. Smokers are rude.
b. Cats meow.
c. Elephants are quite easily trained.

Note that in *Tyrants ruled Wallachia for 250 years*, some particular tyrants or other are clearly referred to, not tyrants in general. The same comment applies to the examples of indefinite plurals in the previous section.

Even more striking are cases observed by Carlson (1973; 1977) in which an anaphoric pronoun and its bare plural antecedent differ in whether a generic or “existential” interpretation is given. Consider the examples in (89) and (90):

- (89) a. May hates *racoons* because *they* stole her sweet corn.
b. *Racoons* stole May’s sweet corn, so now she hates *them* with a passion.
- (90) a. I didn’t think that *goats*_i actually liked *tin cans*_j until I saw *them*_i eating *them*_j.
b. Before I actually saw *goats*_i eating *tin cans*_j, I didn’t think that *they*_i liked *them*_j.

(Anaphorically related pronouns and bare plurals are italicized.) In each case, a bare plural and a pronoun can only be understood in one way, but an existential can be the antecedent of a generic pronoun (as in (89b) and (90b)) or a generic can be the antecedent of an existential pronoun (as in (89a) and (90a)). This failure of pronominalization to heed the difference between existential and universal quantification is unheard of elsewhere.

All of these unusual syntactic facts really point unambiguously to one conclusion, but it is so bizarre that it almost escapes notice. This is that

first, there is no scope ambiguity with indefinite plurals (“existential” bare plurals) simply because there is no quantifier involved in these noun phrases at all¹⁰ and second, the apparent difference between “generic” and “existential” interpretations is due to the meanings of the verbs they interact with, *not* the meanings of the noun phrases themselves. Inspection of the sentences in (89) and (90) shows, for example, that the bare plural subjects and objects of *steal* and *eat* always have an existential interpretation while those of *like* and *hate* always have generic interpretations. (There is a general class of exceptions to this statement involving generics in subject position; cf. Carlson (1977, pp. 247ff.)) Thus the pronouns and their antecedents in (89) and (90) should be said to have the same meaning in a real sense, though the interaction of the meanings of the verbs with these must somehow obscure this fact.

Though it might at first appear impossible to come up with an explicit semantic analysis of bare plurals that satisfies these syntactic desiderata, Carlson is able to do so by means of some fundamental ontological innovations. Bare plurals, we are told, are the proper names of *kinds*. There are as many kinds as there are bare plurals, so we must note that Carlson’s kinds are not just the *natural kinds* of Kripke (1972) and Putnam (1975) but include “unnatural” kinds as well, such as pillows, coffee mugs, and pipe wrenches. A kind cannot, however, be identified with the set of individuals that “make up” the kind or even with the property they all share, but for various reasons must be taken as a basic entity (member of D_e in the UG model) in its own right. A relation R is then introduced in the semantic apparatus that specifies what things *realize*, or “make up”, a kind; if a is a thing and b is a kind, then $R(a, b)$ asserts that the thing a realizes the kind b , as for example a particular cat realizes the kind cats.

It is important for Carlson to provide a somewhat parallel ontological treatment of individuals themselves. He distinguishes between *individuals* and what he calls the *stages* of individuals – these might be thought of as “temporal slices” of individuals, their manifestations in space and at individual times. An individual is that “whatever-it-is” that ties stages together and makes them a single unit. This ontology is similar to views of individuals suggested at times by Kaplan (1973), Gabbay and Moravcsik (1973), and Montague (1973), but not quite identical with any of these. It is crucial that the same R relation that relates kinds to their members also relate individuals to their stages. (This may seem curious at first, but is justified by the consequences that result.) If c is a stage and d is an individual then $R(c, d)$ asserts at any time that the stage c realizes the individual d at that time. It is also to

be noted that R is transitive, so if stage c realizes individual d at a time and individual d realizes kind b at that same time, then stage c also realizes kind b at that time.

Now it turns out, according to Carlson, that some verbs and adjectives that apparently predicate things of individuals and kinds actually amount to predications about *stages* that realize those individuals or kinds at the current time, while other verbs and adjectives really do predicate things of the individuals (or kinds) themselves. Thus while the relation *loves* is true of individuals x and y at a time just in case the individual x stands in the *love*-relation to the individual y , the relation *eats* is true of individuals x and y just in case there exists some stage x' that realizes x at that time, some stage y' that realizes y at that time, and the stages x' and y' stand in some relation defined on stages, which we may call the *eat'*-relation. For example, *Goats like tin cans* would have the representation (91), but *Goats were eating tin cans* would have roughly the representation (92), ignoring tense. Here, g is the constant denoting the kind goats, t the kind tin cans.¹¹

(90) $like(g, t)$

(91) $(\forall x)(\forall y)[R(x, g) \wedge R(y, t) \wedge eat'(x, y)]$

(This is somewhat similar to the way Montague “decomposes” a relation between an individual and a property of properties to an extensional relation between individuals, but the semantic entities involved are here quite different.) In the case of ordinary individuals (i.e., individuals that are not kinds), there is at most one stage that realizes the individual at any given time, so the difference between predicates applying to individuals themselves and predicates applying to their stages is likely to go unnoticed. (However, there may be a few observable syntactic consequences of this difference, some of which will become relevant later in this work.) But with kind names (bare plurals and a few other expressions such as *that kind of animal*), predicates that apply to realizations give rise to an “existential” interpretation, since there will be more than one realization of the kind, and the predicate merely asserts that at least *some* realization has the relevant property. If the predicate applies to kinds themselves rather than stages, then the generic interpretation arises because nothing is being predicated of any realization of the kind, i.e. of any ordinary individuals or stages of them.

I do not have the space here to go into the numerous technical details of Carlson’s proposals that are required to make it complete (for example, a three-sorted logic) nor the impressive evidence Carlson amasses for his

proposals and against the obvious alternatives to it. Because the proposal may initially sound somewhat implausible, I encourage the reader to refer to Carlson (1977; 1977a) for these details and arguments.

What is important for the present discussion of achievement verbs and durative adverbials is that Carlson’s analysis attributes an existential quantifier binding the variable over realizations to the meaning of the verb, not to the meaning of the indefinite plural noun phrase. The indefinite plural noun phrase itself is a proper name wherever it occurs and so it obviously cannot have scope wider than an adverbial quantifier or any other quantifier in the sentence. (Carlson formulated his solution in the PTQ theory, so the “decomposition” of predicates into realization-predicates is accomplished through the translation procedure; when we compare translational decomposition with classical GS decomposition, we will look in detail at ways of insuring in each theory that this existential quantifier necessarily comes from *within* the logical structure of the verb and has narrowest scope. For now, I am assuming that his solution can be accommodated in the generative semantics theory under discussion.) Incorporating Carlson’s analysis into the BECOME analysis of achievement sentences like *John discovered fleas on his dog for six weeks* would result in a logical structure roughly represented by (93), where f denotes the kind fleas:

(93)

$(\wedge t: t \in \text{six weeks})AT(t, \text{BECOME}[John\ knows\ that\ (\forall x[R(x, f) \wedge x\ is\ on\ his\ dog]])]$

The BECOME analysis here, as before, explains why we understand this example to mean that John did not discover the same fleas over and over (i.e., the same realizations of the kind fleas). If Carlson’s analysis is correct, then it is possible to retain the insight from Dowty (1972) that the acceptability of examples like this is to be explained with the BECOME analysis in terms of an existential quantifier with narrow scope but to add the independently motivated account of the narrow scope quantifier that was lacking earlier. Mass terms turn out to have all the same distributional properties as Carlson discovered for bare plurals, and though he does not provide a detailed analysis of mass terms, these parallels suggest that a similar treatment ought to be possible if Carlson’s proposal is correct (cf. Carlson, 1977, pp. 462ff., for some suggestions).

2.3.5. Degree-Achievements

There are some cases of verbs which would seem to be achievements on some semantic and syntactic grounds but which nevertheless allow durational adverbs (even without indefinite plurals or mass terms):

- (94) The soup cooled for ten minutes.
- (95) The ship sank for an hour (before going under completely).
- (96) John aged forty years during that experience.

These seem to express a change of state like other achievements: *cool* is definitely an inchoative meaning "come to be cool", *sink* here means "come to be not afloat", and *age* is "come to be old." Yet there is no contradiction in (94)-(96), no implication that the same change of state took place over and over.

Upon inspection, it turns out that the class of inchoatives that can occur with durative adverbials are just those which have been called *degree words* by linguists (Sapir, 1949; Bolinger, 1972) and *vague predicates* by philosophers (Lewis, 1970; Kamp, 1975). These involve properties such as *big*, *wide*, *good*, *tall*, etc. of which we cannot definitely say once and for all how to determine what their extension is, but can only say so relative to some agreed-upon standard of comparison or some particular context of use. The most typical vague predicates seem to be adjectives, specifically, those that form the comparative without semantic anomaly. As we have *this is cooler than that*, we also have (94) with the adjective *cool*, but as it is strange to say *Mary is more pregnant than Sue* (on a normal interpretation of *pregnant*), it is strange in the same way to say *Mary got pregnant for a month*.

Recent proposals for a model-theoretic treatment of vague predicates (Lewis, 1970; Ginet, 1973; Kamp, 1975) have all been based in one way or another on an appeal to multiple ways of resolving the vagueness of these predicates by assigning a definite extension to them, i.e. different ways of drawing the "boundary" between cool and non-cool things, big and non-big things, etc. Kamp (1975; pp. 136-137) explains it in this way:

At the present stage of its development – indeed, at any stage – language is vague. The kind of vagueness which interests us here is connected with predicates. The vagueness of a predicate may be resolved by fiat – i.e. by deciding which of the objects which as yet are neither definitely inside nor definitely outside its extension are to be in and which are to be out. However, it may be that not every such decision is acceptable. For there may already be semantical principles which, though they do not determine of any

one of a certain group of objects whether it belongs to the extension or not, nevertheless demand that if a certain member of the group is put into the extension, a certain other member must be put into the extension as well. Take for example the adjective *intelligent*. Our present criteria tell us of certain people that they definitely are intelligent, of certain other people that they definitely are not, but there will be a large third category of people about whom they do not tell us either way. Now suppose that we make our standard more specific, e.g., by stipulating that to have an I.Q. over a certain minimum is a necessary and sufficient criterion for being intelligent. Further, suppose that of two persons u_1 and u_2 of the third category u_1 has a higher I.Q. than u_2 . Then, whatever we decide this minimum to be, our decision will put u_1 in the extension if it puts u_2 into it. Finally, let us assume for the sake of argument that any way of making the concept of intelligence precise that is compatible with what we already understand that concept to be is equivalent to the adoption of a certain minimum I.Q. Then there will be no completions in the partial model that reflect the present state of affairs and in which u_2 is put into the extension of the predicate but u_1 is not.

This approach leads directly to a way of deriving comparative adjectives from positive (non-comparative) adjectives (rather than deriving the positive form from the comparative, as earlier semantic treatments of comparatives have suggested). That is, *x is taller than y* will in effect count as true if and only if, for all "acceptable" ways of resolving the vagueness of *tall* by separating the tall from the non-tall, if y counts as tall then x counts as tall also by that method, but not vice versa.

Kamp's proposal is the most detailed that I have seen. He adopts an analysis based on Van Fraassen's supervaluations (Van Fraassen, 1969); that is, the basic interpretation is a partial model that leaves certain predicates undefined (neither true nor false) for certain individuals. Associated with the partial model are a set of (acceptable) completions of that model which fill in the "gaps" in the partial model in various ways – i.e. they assign a truth value for the undefined arguments in the partial model in different ways but otherwise agree with the partial model. In addition to providing a way of treating comparatives, Kamp can also assign a numerical degree of truth (between true and false) to vague sentences like *John is tall* by means of a probability function defined over the acceptable completions. (This method has considerable advantages over attempts to assign degrees of truth to vague sentences by means of multi-valued logics.) The partial model, its completions, and the probability function together form a *vague model* for a language. In a further development, contextual disambiguation of vague sentences is represented by a function from context to models which are less vague than the basic model. Though *John is tall* might be undefined for the basic model, it might well come out true or false (or have greater or lesser degrees of intermediate truth) for certain contexts.

Given such apparatus, an intuitively satisfactory solution to the problem of degree-achievements with durative adverbs begins to emerge. A sentence like *The soup cooled for ten minutes* should be analyzed as saying that for each time t within an interval of ten minutes duration, there is some resolution of the vagueness of the predicate *cool* by which *the soup is cool* is true at t but not true at $t - 1$. Conditions on the acceptable resolutions of the predicate *cool* will in effect require that a different, higher threshold of coolness (i.e. a *lower* temperature for the threshold) be chosen for each successive time in the interval; otherwise the soup could not simultaneously count as cool with respect to one time and resolution of vagueness and also count as not cool for the next time and its resolution of vagueness. This seems to accord well with intuitions about how we understand the sentence, and also avoids having to derive *The soup cooled* from the morphologically unmotivated BECOME[*the soup is cooler*] rather than simply BECOME[*the soup is cool*].

What is necessary for this analysis to work is that the way of resolving vagueness must be capable of being chosen differently for each time t within the interval represented by the durational adverb. Just how this is best done is not yet clear to me. If we do not mention resolutions of vagueness at all in the recursive clauses of the semantic truth definition but merely let a complex sentence be true in a context if the context gives a resolution of vagueness for the elementary predicates that makes the whole sentence true, then different resolutions cannot be used for each time covered by a durational adverb. If on the other hand each recursive semantic definition counts a sentence as true if merely some resolution of vagueness makes it true under appropriate conditions, then *The soup cooled for an hour* could be vacuously true though the soup's temperature remained unchanged; there would still be some resolution at such time t which treated the soup as cool and some resolution (a different one) which made it not cool at $t - 1$, so BECOME[*The soup is cool*] would be true at each time in the hour. What we must apparently do is this: a sentence BECOME ϕ should be true at t if and only if there is some resolution of vague predicates that makes ϕ true at t but false at $t - 1$; then $(\wedge x: x \in \text{an hour})\phi$ must be true if and only if for all times t' within the interval *an hour* there is some resolution of vague predicates that makes ϕ true at t' . This effects the right restriction on the "scope" of quantification over times and resolutions. Whether some way will come to light of avoiding this explicit appeal to resolutions in the recursive clauses I do not know, and so I will leave the matter at this point.

2.3.6. Accomplishments and CAUSE

The verb *kill*, which appeared as an example in McCawley's influential article on word meaning (McCawley, 1968), is an accomplishment verb. If one examines the large literature on "causatives" in GS, the class of verbs there referred to as causatives seems to be co-extensive with the class of accomplishments, though aspectual syntactic tests like those in 2.2.3. have not been used to define the class. This convergence is not surprising when one recalls that Kenny considered all accomplishments to be describable as "bringing it about that p " for some proposition p . (This use of *causative* contrasts with the way it is used in traditional linguistics, according to which it refers only to verbs derived by a causative affix, an affix whose meaning is paraphrasable as "cause to", "cause to be" as English *ize* in *randomize*. When it is necessary to distinguish among syntactic and morphological varieties of causatives, generative semanticists generally distinguish among *lexical causatives*, such as *kill*, *derived causatives*, such as *randomize*, and *periphrastic causatives* – phrases containing a general causative verb plus a separate complement verb, such as *make him leave*, *cause him to leave*.)

In fact, I suggest that in the aspect calculus we construe all accomplishments as having the logical structure $[\phi \text{ CAUSE } \psi]$, where ϕ and ψ are sentences. These embedded sentences ϕ and ψ may have various forms, the most common being the case where ϕ is a BECOME sentence or contains an activity predicate, and ψ is a BECOME sentence. For example, an accomplishment sentence like *John killed Bill* would have a logical structure with roughly the form of (97), and that of *John painted a picture* would have roughly the form of (98):

(97) $[[\text{John does something}] \text{ CAUSE } [\text{BECOME} \neg [\text{Bill is alive}]]]$

(98) $[[\text{John paints}] \text{ CAUSE } [\text{BECOME} [\text{a picture exists}]]]$

This analysis differs from McCawley's original version in that CAUSE is here treated as a kind of two-place sentential connective, rather than as a relation between individuals and propositions.

This so-called "bisentential analysis" of CAUSE did not originate with Dowty (1972) but had been suggested in various contexts (Vendler, 1967a; Geis, 1970; Fillmore, 1971; J. McCawley, 1971;¹² Lee, 1971; N. A. McCawley, 1973; Rogers, 1972; Givón, 1972). I will not attempt to survey thoroughly the reasons for choosing one or the other analysis in a generative semantics theory, but merely cite a few advantages of the "bisentential" analysis and

refer the reader to the above literature, Wojcik (1974; 1976) and Shibatani (1976) for further details.

An obvious motivation for CAUSE as a "subject-complement verb" in generative semantics is Ryle's observation (Ryle, 1949, p. 150) that accomplishments are semantically bipartite in a way that activities are not, that "some state of affairs obtains over and above that which consists in the performance . . . of the subservient activity." Vendler (1967, p. 154) and Geis (1973, p. 211) make essentially the same observation in pointing out that accomplishment sentences like (99) are elliptical; one can conclude (100) and (101) from (99):

- (99) John dissolved the Alka Seltzer.
 (100) John dissolved the Alka Seltzer by doing something.
 (101) John's doing something dissolved the Alka Seltzer.

Geis suggests that (101) is the underlying structure of (100), (100) being derived by a transformation of Agent Creation, a transformation that breaks up the subject complement into an agentive subject and a post-posed *by*-phrase. This transformation may derive some plausibility from the fact that its operation is quite similar to that of the well-motivated Raising (to Subject) transformation, the rule that derives (102a) from (102b) (compare with (101) and (100)):

- (102a) John would be unlikely to win the contest.
 (102b) John's winning the contest would be unlikely.

For what we may call general causatives like *kill*, *open* and *make* (in the sense of *create*) the sentential subject analysis might seem unmotivated, since the meaning of these verbs does not seem to specify anything about the kind of activity that is used to bring about the result, but only the result itself. One can kill a person or animal by any number of activities or procedures; one may open a door by pushing, kicking, striking it, by throwing something at it, by setting off an electronic device or maybe even by saying a magic word, and the ways of making a picture are likewise varied. However, many monomorphemic accomplishments do specify this associated activity in more or less detail. In the class of homicidal verbs (always popular as linguistic examples) are examples like *electrocute*, *strangle*, *poison*, *drown*, *hang*, etc. which give a specific method of bringing about a death (as well as examples like *assassinate* and *execute* which specify a particular motive

though not a means¹³), and one can not only make a picture, but can also paint, draw, sketch, etch, carve, or stencil a picture, these activities indirectly giving indications of the kind of picture that results. Thus we want to suppose that the embedded subject sentence of CAUSE in the underlying structure of general causatives like *kill* or *make* contains a quite general activity or event verb, while other accomplishments have a more specific predicate in this place. (Even *act* is not general enough for the causal event of *kill*, since its subject can be an inanimate (so-called "instrumental") subject, as in *The falling tree killed John*; perhaps *do something* is sufficiently general.)

An even more notable motivation for bisentential CAUSE is a kind of accomplishment construction called *factitive* in traditional grammar and *instrumental* in generative semantics (Green, 1970; 1972; McCawley, 1971):

- (103) Jesse shot him dead.
 She painted the house red.
 She hammered the metal flat.
 He swept the floor clean.
 (104) He drank himself silly.

(The term *instrumental* is really inappropriate since the construction clearly includes examples like *She slammed the door shut*, *He shook her awake*, *She pulled it free* in which no "instrument" is involved.) Here, an activity (or accomplishment) verb combines with an adjective and an object noun phrase to give an accomplishment in which the verb describes the causal activity (or accomplishment) and the adjective gives the result state that the direct object comes to be in as a consequence. Given the sentential subject analysis of CAUSE, examples in (103) would have the kind of structure represented in (105):

- (105) [[*He sweeps the floor*] CAUSE [BECOME[*the floor is clean*]]]

An interesting feature of the construction is that though the object of the causal clause is usually identical with the subject of the result-clause (cf. (105)), this need not necessarily be the case. In (104) the understood object of the simple verb *drink* is not the person denoted by *himself*, though *himself* clearly functions as object of the "whole phrase" *drink silly*, in the sense that his becoming silly was brought about by his drinking (something). Constructions semantically similar to (103) exist in which a predicate nominal or prepositional phrase replaces the adjective, such as *elect John chairman*, *cook the steak to a crisp*, as a parallel to (104) is *read oneself to sleep*.

(Sentences like (103), (104) and these last examples will be treated explicitly in 4.7 below.)

Another class of sentences that may motivate a bisentential analysis of CAUSE is a subset of the verb-particle constructions (cf. Fraser, 1965; 1974), those in which the particle expresses a location that the direct object comes to be in as a result of an activity identified by the basic verb, such as *put the book away*. Within the lexical restrictions of English it is often possible to hold the activity constant and vary the result state as in (106), or to hold the result constant and vary the activity as in (107):

- (106) throw NP away
 throw NP down
 throw NP aside
 throw NP in
 throw NP up

- (107) put NP away (aside, etc.)
 throw NP away
 send NP away
 drive NP away
 call NP away


The point of these paradigms is to suggest that at least a restricted subset of the verb-particle constructions should not be treated as single lexical units consisting of verb and particle together, but that they are to some real extent compositional accomplishment constructions of activity verb and particle that expresses a result state.

The alternative in GS to deriving *by*-phrases as just proposed is to treat CAUSE as a relation between an individual and a proposition as McCawley originally did and then derive *by*-phrases from yet a different abstract operator. Such an analysis is proposed by McCawley (1971), according to which all of (108)-(110), if not even more sentences, are derived from the structure (111):

(108) He made the metal flat by hammering it.


(109) He flattened the metal by hammering it.

(110) He hammered the metal flat.



Here, BY is treated as a sentential connective (or in strict GS terms, a “two-place predicate”, the NP arguments themselves dominating two S-nodes) in accord with the prevailing GS view that adverbials are derived from sentential operators (“predicates of higher sentences”). The derivation of (108) and (109) from (111) is fairly straightforward, but the derivation of (110) from the same source is somewhat more dubious. McCawley “conjectures” (1971: 31) that after predicate raising has attached BECOME and FLAT to CAUSE in S₁ and *the metal* from S₄ has become the derived direct object of this complex verb (via subject raising?¹⁴), (i) Equi-NP Deletion deletes the subject of HAMMER in S₂ on the basis of its identity with *he* in S₁ (the same deletion would take place in the derivation of (108) and (109)), then (ii) “a highly suspect transformation deletes the object of HAMMER under God knows what identity condition with S₁”, and (iii) predicate raising combines CAUSE-BECOME-FLAT with BY, and then finally (iv) “a transformation hereby christened means-incorporation” combines this derived verb with the remaining verb HAMMER in S₂.

Georgia Green (1972, p. 97) finds that the derivation of (110) works out to her satisfaction if the underlying structure is not (111) but (112):



She claims that the derivation of (110) from this structure can be accomplished using only the three rules Subject Raising, Equi-NP Deletion and Subject Formation (a rule that Chomsky-adjoints a subject NP to the left of its verb) – plus lexicalization rules of course – though her derivation in fact involves no less than fourteen applications of transformations in this group and the assumption that transformations apply to their own outputs on the same cycle.

The apparent syntactic simplicity of the derivation I proposed at first might seem to give it an advantage over these two, but given the complexity of accepted GS derivations at that time, this complexity would not likely be taken as a very serious argument. (Needless to say, the proposal of GS derivations of this complexity has given rise in some quarters to the suspicion that potentially any form of surface structure must be derivable from any form of underlying structure whatsoever in a GS grammar, this suspicion then leading to despair over the possibility of ever actually testing whether a GS grammar could generate all and only the well-formed sentences of English or some fragment of English. This is a suspicion I am not unsympathetic with.) The source of all this complexity is of course the unquestioned GS assumption that (110) must have the same underlying syntactic structure as (108) and (109), despite its superficial dissimilarity. If one gave up this assumption, then it would seem much more natural syntactically to derive (108) and (109) from a structure like McCawley's and Green's and to derive (110) from a structure like (105).

Another possible reason for preferring a sentential connective CAUSE over McCawley's CAUSE plus BY is that the intuitive interpretation of BY (ϕ, ψ) seems quite similar to that of [ϕ CAUSE ψ], except that the order of arguments is reversed.¹⁵ If BY could be eliminated in favor of CAUSE, a kind of economy could be achieved that is much desired in the GS methodology. A more pragmatic reason for preferring CAUSE as a sentential connective in the present context is that the model-theoretic interpretation of [ϕ CAUSE ψ] I want to consider requires that it be a sentential connective (or else that we in effect define McCawley's CAUSE in terms of this sentential connective).

Of the many problems that arise in attempting to analyze accomplishments from an underlying structure containing CAUSE, one deserves discussion here (others will be attended to later). It was noticed at the very first discussion of this kind of analysis that sentences with derived causatives may not be exactly paraphrasable by sentences with the English verb *cause*, though this is sometimes hard to judge. Hall (1965, p. 28) notes that "one

argument that probably does not convince anyone who does not already agree is that causing a window to break and breaking a window simply do not mean the same thing," adding examples where she finds a derived causative ungrammatical but the periphrastic causative paraphrase acceptable:

- (113) a. A change in molecular structure caused the window to break.
b. *A change in molecular structure broke the window.
- (114) a. The low air pressure caused the water to boil.
b. *The low air pressure boiled the water.
- (115) a. The angle at which the door was mounted caused it to open whenever it wasn't latched.
b. *The angle at which the door was mounted opened it whenever it wasn't latched.

("Ungrammatical" may be too strong a term for (113b), (114b) and (115b) according to some people – I find them merely a little odd – but there is clearly some kind of difference between the (a) and (b) examples which has to be accounted for.) But as Hall immediately points out, this difference is not automatically evidence against the analysis of causative *break*, etc. in terms of CAUSE. The operator CAUSE is an abstract element and need not be considered identical in meaning with the English "surface verb" *cause*; this surface verb might contain other abstract predicates besides CAUSE in its underlying structure, or it might differ from CAUSE in its presuppositions. This possibility, however, presents the GS theory with a methodological dilemma that potentially all structuralist decomposition analyses are subject to: just how do we decide whether a given decomposition analysis in terms of completely abstract elements adequately represents the meaning of the analyzed word or not, given that the test of a decomposition analysis is not just whether a putative English paraphrase containing the "decomposing" words of the analysis is really synonymous with the analyzed word or not? If we say *kill* is CAUSE BECOME NOT ALIVE but have no independent way of deciding exactly what the meaning of these abstract elements is (once we admit that comparing them to *cause*, *become*, *not* and *alive* is no adequate test), then the analysis is in danger of approaching complete vacuity. Even if we were to accept the structuralist's doctrine (which I don't) that we only need to isolate the primitive semantic contrasts of a language, not further analyze these, we still face the problem of knowing whether the theoretical construct CAUSE used to analyze one kind of word is really representing the same meaning as it does when it is used in analyzing another kind of word.

In traditional linguistic analysis, the keen semantic intuitions of the linguist are the only test of whether the significance attached to an abstract element is really constant wherever that element is used, but such judgments are very tricky, especially when each analysis contains more than one abstract element, so that it may be difficult to know just what "part" of the meaning of a real word is being attributed to each abstract element.

In the case of the semantics of causation, further research has magnified rather than diminished the importance of the problem Hall observed. It is now widely assumed that there are at least two kinds of causation evidenced systematically in natural languages, direct (or *manipulative*) causation and indirect (or *directive*) causation (Shibatani, 1976, pp. 31-39) and some writers suggest even more distinctions (Talmy, 1976). Manipulative causation is said to necessarily involve the physical manipulation of the object affected by the agent, while directive causation does not; perhaps the clearest example of the distinction in English is *John stood the child up* (manipulative) vs. *John made the child to stand up* (directive). Shibatani claims that not only in English but in other languages (Korean, Japanese) as well, manipulative causation tends to be expressed by lexical causatives and directive causation by periphrastic causatives (though the generalization is not absolute). But granted that the distinction is well-motivated, the question of how best to analyze the distinction in GS remains open. Should we postulate two causative operators, CAUSE_m and CAUSE_d? Should we assume directive causation is expressed by a primitive causation operator and that manipulative causation is produced by combining this with an adverbial element meaning "by direct manipulation"? Or do we take manipulative causation as basic and posit an adverbial meaning "by indirect means"? Or are there a "general" causation operator and two kinds of specializing adverbials? Are any of these solutions equivalent to any others? (The distinction between kinds of purposeful and non-purposeful causation may possibly be captured by a DO operator introduced below, but this will not help with the kind of difference observed in (113)-(115) above, where no animate subjects are involved.)

The only sure remedy I can see for this problem is to attempt to assign an explicit model-theoretic interpretation to every such abstract element postulated. (Alternatively, one could provide a system of deductive rules which make the entailments derivable with such elements precise, but the model-theoretic method also makes entailments precise and defines meaning in terms of non-linguistic objects as well.) Only in this way will the entailments of a decomposition analysis be really clear, and only in this way can we be sure the same abstract element is used to the same semantic effect in different

analyses. The only case where we can satisfactorily make an exception to this rule is the one where one of the elements of a decomposition analysis *can* explicitly be equated with the meaning of an independent English word – for example, in McCawley's decomposition of *kill* it would seem acceptable to take the meaning of the abstract element ALIVE to be that of the adjective *alive*. For when we do this, we can still test the entailments of an analysis precisely in terms of other English sentences, even though these contain non-logical constants that are not given a standard interpretation. For example, if our analysis of *John killed Harry* gives (by virtue of the explicit analysis of CAUSE, BECOME and negation) the formal entailment that *Harry is not alive* is true under just the right conditions, then this serves as an adequate test of the analysis of *kill* even though we leave the stative predicate *alive* unanalyzed.

Accordingly, in the section that follows I will take CAUSE to be a logical operator (rather than as representing the meaning of English *cause* exactly) and attempt to give a model-theoretic interpretation for [ϕ CAUSE ψ]. As this is an ambitious undertaking which remains in the preliminary stages, I therefore do not feel the need to apologize for ignoring the apparent distinctions among the various kinds of direct and indirect causation mentioned in the literature, since I regard what I am doing as a necessary preliminary to exploring these distinctions coherently. In Chapter 6 I will present one way of dealing with unsystematic divergences of derived causatives from their predicted meanings, and I think it could still turn out that no more should or can be said about manipulative (as opposed to directive) causation than this. (Also, see McCawley (1978) for arguments that at least some of the above distinctions in kinds of causation, if not all of them, can be accounted for in terms of conversational implicature.)

2.3.7. CAUSE and Lewis' Analysis of Causation

In the long philosophical literature on causation, an intuitive connection between causal statements and counterfactual statements has frequently been observed. For example, G. H. von Wright (1963; 1968) observed that to assert that an agent has brought about an event (as in (116)), the speaker must believe that three kinds of facts obtain, in this case those in (117):

- (116) John opened the door.
- (117) a. The door was not open just before John acted.
b. The door was open just after John acted.