

Formal Semantics for Causal Constructions

Richmond H. Thomason
Philosophy Department
University of Michigan
Ann Arbor, MI 48109-1003
U.S.A.

rthomaso@umich.edu

October 23, 2010

Abstract

Montague's framework for semantic interpretation has always been less well adapted to the interpretation of words than of syntactic constructions. In the late 1970s, David Dowty addressed this problem, concentrating on the interpretation of tense, aspect, inchoatives, and causatives in an extension of Montague's Intensional Logic. In this paper I will try to revive this project, conceiving it as part of a larger task aiming at the interpretation of derivational morphology. I suggest an alternative to Dowty's approach that, while it does not provide a global interpretation of causality, seems to work well with a wide range of the causal constructions that are important in word formation.

1. The formalization of causal constructions

Among the formal semanticists who work within the language-to-logic tradition that goes back to Richard Montague, and who accept a logic somewhat like Montague's Intensional Logic as the target language, there is a natural division of labor. Those who happen to be linguists are well suited to working out ways to interpret new linguistic constructions in this framework, or better ways to interpret familiar ones. The smaller group who happen to be logicians are better suited to investigating modifications of the logic that might provide improved interpretations of natural language.

For some time, I have been working—very slowly, I'm afraid—on a project of the second sort, having to do with lexical semantics and word formation. Over thirty years ago, David Dowty's ambitious work in [Dowty, 1979] convinced me that formal semantics could not be successfully applied to this area without changes to the logic. (See Section 2, below, for one

reason.) And in fact, lexical semantics has been rather neglected in formal semantics, in comparison with the semantics of sentences.

To provide a theory of the system of lexical operators and to explain logical interactions (for instance, to derive the relationship between the interpretations of ‘bendable’ and of ‘deformable’ from the relationship between the interpretations of ‘bend’ and ‘deform’), it is important to provide a model theory of lexical operators. This would lead naturally, in this case, to a model-theoretic investigation of ability, a project that is philosophical as well as logical, that that also might be helped by ideas from logical Artificial Intelligence.

In this paper, I will be concerned with some lexical processes having to do with causality.

2. Dowty’s modal treatment of causatives

Causality shows up in many morphological processes, including English ‘-ify’, ‘-ize’, and ‘-able’, as well as the resultative (‘wipe clean’) construction. But the causative construction, which allows the verb ‘blacken’, for instance, to be derived from the adjective ‘black’, is the one most overtly related to causality.

Much of [Dowty, 1979] is devoted to the analysis of causative constructions. He begins with a relation P-CAUSE between propositions: P-CAUSE(p, q) is true of propositions p and q iff the truth of p causes the truth of q . A sentence like ‘Alice’s opening the window caused the temperature to drop’ might express an instance of this propositional causal relation. And the relation could be accounted for model-theoretically in various ways.¹

But the derivation of a causative verb from an adjective involves an operator CAUSE₁ taking adjective denotations into transitive verb denotations. In Montague’s Intensional Logic,² this would be a function of type $\langle\langle e, \langle s, t \rangle \rangle, \langle e, \langle e, \langle s, t \rangle \rangle \rangle$ —a function from 1-place propositional functions over individuals to 2-place propositional functions over individuals.

Dowty proposes a definition of CAUSE₁ in terms of CAUSE along the following lines.

$$(2.1) \text{ Cause}_1(P)(y)(x) \leftrightarrow \exists Q^{\langle s, \langle e, t \rangle \rangle} P\text{-Cause}(Q(x), P(y))$$

The proposal is that, for instance, Betty opened the door if and only if there is some property such that the proposition that Betty has this property causes the door to be open.

This definition has the disastrous consequence that, for instance, if anyone opens a certain door then everyone opens it. Let P be the function taking an individual into the proposition that it is open, let b be the door, and suppose that there is some property Q such that the proposition that a has Q causes $P(b)$. Let P denote P , Q denote Q , a denote a , and b denote b . In view of Definition (2.1), we are supposing that CAUSE₁(P)(b)(a) is true. By hypothesis, P-CAUSE($Q(a), P(b)$) is true. By substitution of logical equivalents, P-CAUSE($\wedge[a' = a' \wedge \vee Q(a)], P(b)$) follows. By lambda abstraction, we have P-CAUSE($\lambda x[\wedge[x = x \wedge \vee Q(a)]](a'), P(b)$). So $\exists Q'(P\text{-CAUSE}(Q'(a'), P(b)))$ is true. By Definition (2.1), we have CAUSE₁(P)(a', b).

¹For instance, you could invoke David Lewis’ counterfactual account; see [Lewis, 1977].

²I will use an assignment of types that is different from any of Montague’s, but that is equivalent if we ignore intensional predication.

I don't know of any straightforward way to replace Dowty's definitions with ones that work. Within Intensional Logic, this leaves the alternative of treating CAUSE1 as a primitive. I can't say for certain that this idea couldn't be developed into a reasonably satisfactory theory, but I don't see a way to provide a plausible model-theoretic account of CAUSE1. And a number of related phenomena, such as aspect, some adverbs, and morphological compounds like 'reopen' and 'disentangle', seem to fall into place better in a framework that is equipped with eventualities.

The propositional approach to causatives has other counterintuitive consequences. For instance, the P-CAUSE operator is motivated by intuitions associated with natural language causal constructions that are propositional, or at least close to propositional: constructions involving 'cause' and 'make'. It has often been noted, for instance, that causatives involve a sort of causality that, unlike propositional causality, is somehow immediate.

Apparently, we need another member of the causal family. The leading alternative to propositional causation is event causation, or, more generally, eventuality causation. (The term 'eventuality' includes states, as well as the happenings that are usually called events.) When the death of Charles II of Spain caused the War of the Spanish Succession, this was an example of eventuality causation.

But to have eventuality causation, we will need to introduce eventualities into the logic. This is the first change to Montague's framework that I'll propose. It is a maneuver going back at least to Terry Parsons' study of aspect in [Parsons, 1990], that has been helpful in other settings, and that is well motivated.

3. Adding eventualities to the logic

Eventualities can be incorporated in Intensional Logic without any changes to the logic itself. We merely treat eventualities as individuals (we assign them type e) and add vocabulary and postulates as needed for our theory.³ Even while IL remains as the underlying logic, lexical semantics requires a change in the formalization style and, in particular, the use of postulates.

Montague's use of postulates in formal semantics is very restricted. Postulates are confined to *meaning postulates*, and express semantic constraints on constants of IL that translate lexical items. The quantifier 'every' provides a relatively simple example. If we neglect intensionality, this would be translated by a constant *every'* of type $\langle\langle e, t \rangle, \langle\langle e, t \rangle, t \rangle\rangle$ —the type of a function that inputs a set of individuals and returns a set of sets of individuals. The following meaning postulate will then relate *every'* to the universal quantifier of IL.

$$\forall X^{(e,t)} \forall Y^{(e,t)} [[\mathit{every}'(X)](Y) \leftrightarrow \forall x^e [X(x) \rightarrow Y(x)]]$$

Lexical semantics requires postulates that, rather than merely constraining the meanings of lexical items, provide a theoretical framework for explicating and interrelating these meanings. Postulates may introduce theoretical terms that do not correspond directly to

³I'll follow the usual Davidsonian practice of using e as a variable ranging over events. For the purposes of this paper, this simply means that e is a variable of type e that is used when we have in mind the cases where e takes eventualities as values. I will use different fonts for types and for variables: e is the type of individuals, e is an event variable.

lexical items, but that serve to explicate their meanings. Dowty’s work in [Dowty, 1979] illustrates the phenomenon: to explain causatives and other constructions he develops a theoretical framework involving branching time, temporal intervals, and causal inertia. He doesn’t provide postulates for this framework within Intensional Logic. Although this could in principle be done, it is easier, as Dowty does, to explain the apparatus less formally, in the metalanguage.

Writing in the 1980s, Emmon Bach noted the tendency of such ideas to emerge from the study of lexical semantics; see, for instance, [Bach, 1989]. In Bach’s terms, this sort of investigation calls for a “natural language metaphysics.” But the theories that Bach had in mind also bear a close affinity to attempts in Artificial Intelligence to formalize “common sense knowledge.” This area of Artificial Intelligence may be more relevant to semantics than contemporary metaphysics, because it is more formal. See, for instance, [Davis, 1991, Mueller, 2006].

For my purposes here, I need a theory of eventualities. I won’t attempt a detailed formalization, providing only what is needed for my present purposes. In describing eventualities, and configurations of eventualities, I will need the following elements. The theory could be formulated in Intensional Logic, but I will present it in the metalanguage.

Predicate or Relation	Informal Interpretation
C-CHILD	Relates an eventuality to its causal children.
STATE	Is a state.
TELIC	Is a telic eventuality.
PROCESS	Is a process.
INCEPTION	Relates a telic eventuality to its inception.
BODY	Relates a telic eventuality to its body.
CULMINATION	Relates a telic eventuality to its culmination.

‘C-Child’ stands for “causal child.” I chose a neutral term for this relation because it is not meant to stand for anything in natural language. Instead, it is a theoretical relation, inspired by *causal graphs*. These graphs are used to represent causal relationships between items; for instance, between the variables in a clinical trial. An item *i* is a causal child in the graph of a parent item *j* in case *j* can influence *i* without any intervening causal factors. The idea goes back at least to [Simon, 1952]. Judea Pearl’s use of such networks to model statistical inference, causal relations, and counterfactuals provides a good example of the power of the idea; see [Pearl, 2000]. For a recent extended examination of these methods in a more philosophical context, see [Woodward, 2003].

I will use causal graphs composed out of the C-Child relation to model the causal structure of telic eventualities. In a careful formalization, the relation would be indexed to eventualities. To keep things simple, however, I will not do that explicitly here.

My approach to telicity is inspired by Mark Steedman’s eventuality-based theory of tense and aspect, as described in [Steedman, 1998]. Steedman treats telic eventualities as three-part complexes, consisting of an *inception*, a *body*, and a *culmination*.⁴

The causal structure of a telic eventuality may involve more than just these three parts, because some processes require *maintenance*. Consider the difference between flushing a toilet and driving home from work. The first operation requires no effort on the part of an agent once the inception—operating the lever—has been accomplished. The process that is initiated by the inception is automatic, and normally (and hopefully) will carry itself to completion without assistance from an agent. The second operation requires a great deal of attention, and many actions on the part of the driver. Opening a door could be an operation of either sort, depending on whether there is natural resistance to the opening process—for instance, a spring that will pull the door into a closed position if it isn’t overcome by pushing on the door.

At some points, I will have to consider group activities and plural agents. But I will avoid details having to do with plurality here, and will assume that each component of a telic eventuality, including maintenance activities, has a single individual as its agent.

4. Telic eventualities and causal structure

- (4.1) Consider a door, closed with a latch that is controlled by a doorknob. An agent a is standing next to the closed door. Now a proceeds to open the door in the usual way; let e be the resulting telic eventuality. The inception of e will be an initiating event e_0 in which a grasps the doorknob, turns it, and (say) pulls on the door. The body of e is a process e_1 through which the door swings open; normally, the further advanced this process is, the more open the door is. The culmination of e is a state e_2 in which the door is open.

Telic eventualities are causal structures; they represent simple, recurring causal patterns. The inception causes the body, without any intervening causal factors. The body causes the culminating state, again without intervening factors. Therefore the causal graph will be a three-part causal pipeline.⁵




Figure 4.1

⁴There are alternative accounts of telicity, such as Krifka’s algebraic approach in [Krifka, 1998]. Steedman’s idea is very useful in the project I undertake here, but Krifka’s also has advantages. It isn’t clear, however, that the two approaches are irreconcilable; a three-part structure might well have the algebraic structure suggested by Krifka’s work.

⁵In this and all subsequent graphs, arrows will represent the parent-child causal relation.

Although all telic eventualities have this much causal structure, often there is more than this to the causal story. Suppose, for instance, that the door is equipped with a spring that keeps it in a closed position. Then, throughout the body e_1 of e , the agent will have to exert force on the door to counteract the pressure of the spring. This exertion will be another eventuality, say e_3 , serving as a causal parent of the process e_1 . The resulting diagram will look like this.




Figure 4.2

5. Agency and causatives

A semantic theory of the causative construction should provide an account of how the interpretation of a causative transitive verb is related to the interpretation of the original adjective. With respect to the examples in Section 4, this means showing how *open*'₂ is derived from *open*'₁, where *open*₁ is the interpretation of the adjective 'open', and *open*'₂ the interpretation of the homophonous transitive causative verb 'open'.

5.1. Methodology and the diversity of causality

Common sense provides us with a rather large family of causal notions. The task of sorting these notions out and charting their relationships has puzzled philosophers for a very long time, producing many theories but few definitive solutions. Natural language provides an extensive, but somewhat different family of causal constructions, such as the ones in the following list, including derivational mechanisms as well as explicit causal constructions with various types of arguments.

Causatives:	'Mary opened the door.'
Resultatives:	'Mary wiped the table clean.'
The '-er' construction:	'This is a bottle opener.'
The '-able' construction:	'This bottle is openable.', 'This bottle is unopenable.'
Explicit causality:	'The drop in the temperature made the streets icy.'
Explicit causality:	'The traffic jam made me late for work.'
Enabling:	'Growing up in the South enabled his success as a novelist.'

It would be intellectually satisfying to derive all, or even most of these from a single, basic causal relation. But I don't think it would be a good idea to pin the success of a semantic theory on hope for a unified account. Language, and common sense as well, may provide us with a family of only loosely related causal relations. A more cautious and appropriate

methodology would identify the family members, and provide separate accounts of their semantic interpretations. If unifying connections emerge from these local accounts, so much the better; if not, we will still have made progress.

The two examples in Section 4 provide the material for a local account of causatives. The idea is simple: the causative relation denoted by *open*₂ is the relation between the agent of the telic eventuality e and the theme of the state e_2 that is its culmination. To put it even more simply, causativity is agency, where the agent in question is the agent of a complex telic eventuality.

Although it was motivated by very simple examples, this idea actually seems to hold up well in more complex cases. But it raises other problems that will need to be addressed. (1) How does this idea relate to the notion of ‘direct causation’ or ‘immediate causation’ that is often used to explain causatives? (2) How do we avoid overgenerating causative relations? (3) How is agency in the component eventualities projected to agency for the entire telic eventuality? That is, how are the agents of the component eventualities of a telic eventuality e related to the agent of e ? I’ll discuss each of these in turn. The discussion will reveal that our intuitions about agency in even quite complex causal configurations are well developed and refined. We can use these intuitions to test theories of the causative construction.

5.2. “Immediate” and “inanimate” causality

As far as I can see, the basis for the idea that causative constructions involve a sort of causality that is somehow “immediate” derives from examples like this: Ann turns up the heat, making Bob uncomfortably warm so that he opens the window. Here, Ann caused the window to become open, but Bob, not Ann opened it.

This idea of immediacy can’t be quite right—at best, it is very misleading. To see this, consider a Rube Goldberg⁶ apparatus for opening the window, as complex as you like. An agent, Rube, drops a ball into a slot; the ball rolls along a track, finally settling into the drain hole of a small sink. Water running into the sink overflows, falling into a bucket hanging from a rope. As the bucket becomes heavier, the rope, attached to the door by a system of pulleys, pulls on the door and opens it. In this case, we have no difficulty in thinking that Rube opened the door. And this outcome doesn’t seem to be affected by the length of the causal chain.

Perhaps immediacy amounts to this: the agent of a telic eventuality is the closest agent to the outcome.⁷ When Mary opens the door in the the example of Figure 4.2, she is the agent of both e_0 and e_3 (the inception and the pushing). These are both separated from the culmination e_2 by the body, e_2 , which either has no agent or which, by projection, has Mary as an agent. In either case, Mary is the closest agent to the outcome. And in the Rube Goldberg example, the only component agent is Rube, the agent of the inception.

This idea doesn’t work either. Suppose that Mary sails a sailboat from the shore to an island by hoisting the sail and steering the boat with a tiller. Then Mary certainly sailed the

⁶See <http://www.rubegoldberg.com/>.

⁷Cases in which several agents tie for closeness would be cases of group causative agency.

boat to the island, even though the wind pushed the boat to the island. If we neglect the steering process, Figure 4.2 can serve as a diagram, with e_0 the hoisting, e_2 the movement of the boat towards the island, and e_3 the process consisting of the pushing of the wind on the boat.

Compare this with the case in which two people, Mary and John, cooperate in opening the door: Mary works the latch, and John pushes. Here, it seems more natural to say that John opened the door, or than John and Mary opened the door, than to say that Mary opened the door; this looks like a case of group agency, with one of the two agents getting most of the credit. But, although the causal diagrams are identical, it would be bizarre to say that Mary and the wind sailed the boat to the island.

The difference between these two cases could be explained by saying that the pushing component actually has no agent. When we say ‘The wind pushed the boat’ we are speaking metaphorically, and metaphors are a poor clue to literal semantic content.

This idea doesn’t work either. Consider the following two cases.

(5.1) Mary fells a tree with her chainsaw. The tree falls on a fence, flattening it.

(5.2) The wind blows a tree down. The tree falls on a fence, crushing it.

Both examples have the same causal structure, the one diagrammed in Figure (4.1). But in Example (5.1) we can say either ‘Mary crushed the fence’ or ‘The tree crushed the fence’,⁸ whereas we can describe Example (5.2) by saying ‘The tree crushed the fence’, whereas it is simply false in this case that the wind crushed the fence.

‘The tree crushed the fence’ sounds disingenuous as a description of Example 5.1. I want to say that here, whoever (perhaps Mary) is describing what happened is invoking a truncated telic eventuality that omits her initiating action; this is not a description of the larger telic eventuality that includes Mary’s action.

If we were to neglect inanimate agents in projecting agency, treating these cases as metaphorical, ‘The wind crushed the fence’ should be acceptable, and true, as a metaphorical analogue of ‘Mary flattened the fence’. Moreover, we should be able to treat the tree as a metaphorical agent in Example (5.2), making ‘The wind and the tree crushed the fence’ acceptable. But in fact, this sounds very peculiar.

I draw two conclusions. First, the idea that “immediate” causation involves short causal chains has to be discarded. Second, we have to treat things like the wind in Example (5.2) as genuine, but automatic or involuntary agents.⁹ That makes it necessary to treat voluntary

⁸Although ‘Mary and the tree crushed the fence’ sounds peculiar here.

⁹I replace the familiar ‘animate’/‘inanimate’ distinction with the unfamiliar ‘voluntary’/‘involuntary’ one, because animal and even human agents seem to pattern like the wind or other natural forces when the action is considered to be automatic or involuntary. People and animals can figure, for instance, in Rube Goldberg mechanisms. Say that the ball on the track somehow causes a feather to tickle a sitting person, making him squirm. This activates a squirm sensor under the seat cushion, which in turn opens the door. This causal chain doesn’t seem to be significantly different from the one without the sitting human agent. The term ‘involuntary’ is not as appropriate for what I have in mind as ‘automatic’, since agents who are exercising deliberation and choice in their actions can be considered as involuntary for some purposes. For instance, we can say that Columbus sailed a fleet to the New World, neglecting the roles played by the crews, and even the officers, of the ships. Nevertheless, I will use ‘voluntary’, because of the available of the contrary ‘involuntary’.

and involuntary agents differently in projecting agency. This is the approach I will take below, in Section 5.4.

There is a sense, however, in which using causal structures, represented by causal graphs, and especially telic eventualities, to package causal information does capture a sort of immediate causality. These structures are local: irrelevant circumstances are simply omitted. For instance, Mary may move across the room to the door as part of a larger plan to visit her daughter. That movement prepares the way for opening the door, but is not part of the telic eventuality of opening the door. And causal graphs provide a further sort of locality: an ancestor that is not related to another node through a chain of C-Child relations cannot be a causal influence on that node.

5.3. The problem of overgeneration

Overgeneration is quite a different problem. Suppose, for example, that in Example (4.1), the door is black. Then, it is black in the culminating state. So we should be able to say here that the agent not only opened the door, but blackened it.

Here, what could be a serious problem is solved by a proper use of causal graphs. But, also, the solution seems to commit us to an ontology of states that some ontologically conservative philosophers may not have welcomed.

In particular, we need to distinguish two states here: the state e_2 of the door's being open, and the conjunctive state e_4 of its being open and black. The state e_2 is a C-Child of the process e_1 of opening the door. The state e_4 is not a C-Child of this process, although it would be a C-Child of a process of opening it and painting it black. Alternatively, if we think of states as characterized by feature structures, then the culminating state e_2 of Example 4.1 might be represented as the following feature structure, where d is the door of the example.

$$(5.3) \quad e_2 : \left[\begin{array}{ll} \textit{Aspectual-Type} : & \text{STATE} \\ \textit{Proposition} : & \text{that } d \text{ is open} \end{array} \right]$$

Then we would have to say that the C-Child relation is not monotonic with respect to propositional informativeness. In our example, e_2 is a C-Child of e_1 , but e_4 is not.

$$(5.4) \quad e_4 : \left[\begin{array}{ll} \textit{Aspectual-Type} : & \text{STATE} \\ \textit{Proposition} : & \text{that } d \text{ is open and black} \end{array} \right]$$

But I'm doubtful about using propositions to describe states. It goes well beyond what is necessary for the causative constructions we are attempting to interpret. These apply to simple adjectives, not to their conjunctions, and so provide no incentive to include conjunctive states. So I'll work with a more conservative description of states, which represents just the theme and property. And I'll assume that the property is denoted by an adjective. The state e_2 would then be described by the following feature structure.

$$(5.5) \quad e_4 : \left[\begin{array}{l} \textit{Aspectual-Type} : \text{STATE} \\ \textit{Theme} : \text{d} \\ \textit{Property} : \text{being open} \end{array} \right]$$

5.4. Projecting agency

We will also use feature structures to describe the inception and body of a telic eventuality. Unlike the culminating states, these eventualities can have agents. Our task is to characterize how, given the causal influence diagram for a telic eventuality e and descriptions of the component eventualities, we can assign an agent to e .

Let's begin with the simplest case, like Example (4.1), where the door has no spring; the initial push lets it swing forward. In this case the inception will be a telic eventuality, in which a single agent a unlatches and pushes the door. Representing this of this process would require sequential actions, and I will avoid these details. All that is important for our purposes is that a is the agent of the inception, and that the body has no agent.

The causal influence diagram for this example is given by Figure (4.2), and the relevant features of the inception e_0 and body e_1 will look like this.

$$(5.6) \quad e_0 : \left[\begin{array}{l} \textit{Aspectual-Type} : \text{TELIC} \\ \textit{Agent} : \text{a} \end{array} \right]$$

$$(5.7) \quad e_1 : \left[\begin{array}{l} \textit{Aspectual-Type} : \text{PROCESS} \\ \textit{Agent} : \perp \\ \textit{Theme} : \text{d} \\ \textit{Relation} : \text{opening} \end{array} \right]$$

The progressive form ‘opening’ in Feature Structure (5.7) raises issues that I will discuss later, in Section 6.

The case in which the door is spring-loaded, and a single agent needs to push to counteract the spring, is similar. The body now has an agent—the same agent involved in the inception—and this is the agent of the telic eventuality.

We can generalize these straightforward cases. When no component of a telic eventuality e has an agent, e has no agent. Otherwise, if all components have the same agent a , then a is the agent of the telic eventuality.

As the previous examples indicated, projection is more complicated in the remaining, multiagent case. Here, the crucial question is how to determine the agent of the body, in case the body has many C-parents. I now turn to this problem.

We'll begin with some causal graph relations deriving from the C-Child relation. The *ancestral* C-CHILD* of the C-Child relation is the relation that any node bears to its descendants through a series of C-Child relations. I will spare you the explicit definition of the

ancestral in Intensional Logic, and will state the definition informally. I will do the same with the definition of C-Betweenness.

(5.8) C-CHILD* is the ancestral of C-CHILD.

(5.9) C-BETWEEN relates two eventualities e and e' in a causal graph to any eventuality e'' such that either (i) e'' is a C-Child of e and e' is a C-Child of e'' , or (ii) e'' is a C-Child of e' and e is a C-Child of e'' .

From the examples I have presented (as well as others), a number of generalizations emerge about the attribution of an agent to the body of a telic eventuality. Three important such generalizations can be stated informally, as follows. The idea is that, in the causal graph of a telic eventuality, any agent of a causal ancestor of the body is a candidate for promotion to the agent of the body, but that it is *preempted* from promotion under certain circumstances.¹⁰


(5.10) A voluntary agent is preempted by any voluntary agent between it and the body.

(5.11) An involuntary agent is preempted by any agent between it and the body, and by any causal ancestor of the body with a voluntary agent.

(5.12) The agent of the inception is preempted by any voluntary agent elsewhere in the causal ancestry of the body.

We can illustrate these ideas with the following hypothetical graph, which gives the causal structure, names the agents of causal ancestors of the body, and labels them as voluntary or involuntary. All of the agents a_0 to a_5 are potential agents for the body. But a_0 is preempted by (5.12), because there are voluntary agents elsewhere. a_1 is preempted by (5.11), because of a_2 . And a_4 is preempted for several reasons; for instance, by (5.11), because of a_3 . Finally, a_3 is preempted by (5.11), because of a_2 . The only viable agent here is a_3 , and this will be the agent of the body.

¹⁰These are based on the examples I have presented, but—since, after all, they are generalizations—they cover an unlimited variety of possibilities, many of them too complicated to exemplify easily. Under the circumstances, convincing counterexamples may turn up, and the generalizations are, of course, subject to correction and change.



The informally stated generalizations (5.10)–(5.12) can be reformulated to provide explicit conditions on how the agent of the body of a telic eventuality depends on its causal antecedents. First, we define the set $PE(G)$ of *proximate eventualities* of G , where G is the causal graph of a telic eventuality e .

(5.13) An eventuality e is voluntary if it has an agent, and the agent is voluntary. It is involuntary if it has an agent that is involuntary.

(5.14) Let G be the causal graph of a telic eventuality e . Then $f \in PE(G)$ if and only if:

- (i) f is a C-Ancestor of the body e_1 of e ;
- (ii) if f is voluntary, there is no voluntary eventuality C-Between f and the body; and
- (iii) if f is involuntary then f is a C-Parent of the body, and no C-Ancestor of the body is voluntary.

The set $PE(G)$ can never contain both voluntary and involuntary eventualities. If it contains a voluntary eventuality f , there will be no voluntary eventualities C-Between f and the body. If it contains an involuntary eventuality f , f will be a C-Parent of the body.

All along, I have been assuming that the component eventualities with which we have been working do not have group agents. But if many agents participate in the body of a telic eventuality, they may project a group agent to the body. Bearing this in mind, the following rules of projection will capture generalizations (5.10)–(5.12).

- (5.15) If $PE(G)$ is empty, the body has no agent.
- (5.16) If $PE(G)$ contains a single eventuality, the agent of the body is the agent of this eventuality.
- (5.17) Otherwise, remove the inception from $PE(G)$, create a group that consists of the agents of the remaining eventualities, and project this group as the agent of the body.

Agency can then be projected from the body to the entire telic eventuality: if e is a telic eventuality with body e_1 , the agent of e is the agent of e_1 .

6. The semantic derivation of causatives

Let's recall the original problem. We required a definition of the interpretation of a derived causative in terms of the interpretation of its underlying adjective. For example, given a relation $OPEN_2$ between individuals and states (as in 'The door was open'), the question is how to derive $OPEN_4$, where this is a relation between individuals, individuals, and telic eventualities (as in 'Mary opened the door').¹¹

The detour we have taken through causal structures and the projection of agency may have left it unclear what progress we have made. With eventualities in the picture as arguments of verb and adjective interpretations, the semantic derivation problem is this: to determine what telic eventualities correspond (in the right causal relation) to a given state? In other words, we need to define an appropriate function $LCAUSE$ between telic eventualities and states: $LCAUSE(e, s)$ if and only if e is a telic eventuality culminating in s . This is not trivial, because not every complex with the structure of (4.1) is a telic eventuality culminating in the state e_2 .

Recall the example at the beginning of Section 5.2, in which Ann causes the window to open, but doesn't open it. Let e_1 in (4.1) be the process of turning up the heat (Ann turns the thermostat dial), and e_0 be its inception (Ann begins to turn the thermostat dial). This has the appropriate C-Causal structure, but is not a telic eventuality, because e_1 is not an appropriate body for this particular culmination. To be appropriate, the body must be a process in which the door that is the theme of e_2 becomes open. In other words, the process e_1 must be an inchoative process for the state e_2 .¹²

But that's not all. Let e_0 , as before, be the inception of the process of turning up the heat, while e_1 is now the inchoative of e_2 . e_0 is a cause of e_1 , and will be a C-Parent of e_1 if

¹¹The English sentences are tensed, but we are ignoring tense in the actual analysis, the idea being that tense will more or less take care of itself.

¹²You might also think that the causal structure is wrong in this example, because e_1 is not a C-Parent of e_2 : there are intervening eventualities, such as the process of Bob's opening the window. This, I think, would be a mistake: causal graphs are models of specific causal situations. In constructing a causal graph, the modeler chooses the domain of eventualities that are relevant to the purposes of the model. Some intervening eventualities may well be ignored—for instance, because of considerations having to do with granularity.

we ignore the intervening eventualities. But this is not a telic eventuality, because e_0 is not the inception of e_1 .

With these things in mind, the definition we are seeking of $LCAUSE$ would need to involve five conditions.

$LCAUSE(e, s)$ if and only if:

- (1) s is the culmination of e .
- (2) The body e_1 of e is appropriate for s .
- (3) The inception of e is the inception of the body e_1 .
- (4) The agent of e is appropriate.
- (5) The theme of e is appropriate.

In Section 5.4, we showed how to flesh out (4), providing a characterization of the agent that is appropriate for a telic eventuality, given a causal graph G that provides a causal pedigree of the body. And we have a simple solution to (5): the only appropriate theme of a telic eventuality is the theme of its culminating state.

Here, then, is where the problem stands. In fact, we need a definition of a three-place relation $LCAUSE(e, s, G)$: e is a telic eventuality culminating in s , relative to a causal graph G . We can complete this definition if we can provide a definition of two other relations: (i) a two-place relation $INCHOATIVE$ that holds between a process e_1 and a state s when e_1 is an inchoative process for s , and (ii) a two-place relation $INCEPTION$ that holds between an eventuality e_0 and a process e_1 when e_0 is an inception of e_1 .

As far as I can see at the moment, reasonably satisfactory definitions can be obtained for both of these relations. But that is a topic for another paper. Here, I will only note that Dowty's proposal in [Dowty, 1979] for characterizing the inchoative relation needs to be strengthened. Roughly, the idea is that the inchoative process culminating in s shares the theme of s , ends at the time that s starts, and begins at a time at which the property of s does not hold of the theme of s . The problem here is that, for instance, a tire might be warmed over the same time that it is inflated. (The process of inflation, in fact, warms it.) At the end of the interval, there are two states: a state in which the tire is warm, and a state in which it is inflated. But the inflation process is not the inchoative of the first state, and the warming process is not the inchoative of the second state.

7. Normalcy and causative morphology

In this section, I begin a new, but related topic, and return to the question of whether significant changes are needed in the underlying logic.

If the inception of a telic eventuality occurs—provided that other events that may be required to maintain the resulting process also occur—it initiates changes that normally result in the full unfolding of the telic eventuality, including the culminating state. This is what gives telic eventualities their practical usefulness. We use our knowledge that turning the knob of a door that opens inwards and pulling on it will normally open it when we want to open doors. The semantic account of derived causatives that I sketched above, unlike Dowty's, doesn't appeal explicitly to normalcy, but normalcy is implicitly there in the causal structure of a telic eventuality.

Normalcy appears in [Dowty, 1979] in the form of “inertial worlds”—courses of history in which events unfold as expected. Dowty uses inertial worlds to provide a semantics for progressive aspect. But this is not the only place where normalcy is needed—it is ubiquitous in the semantics of morphological processes. To hammer a nail is to use a hammer on a nail in the normal way that hammers are used. To wipe a table clean is to clean a table by wiping it in the normal way that tables are cleaned by wiping. To towel a glass dry is to use a towel to dry the glass in the normal way that towels are used to make glasses dry. A fastener is a device that is normally used to fasten things together. A door is openable if normally, when the proper (normal?) procedures for opening it are applied to the door, it will then open.

Naturally enough, Dowty treats normalcy as a modal operator: a proposition is normally true on this view if it is true in all inertial worlds. On this approach, the normality that lexical semantics requires is monolithic, and applies to propositions.

However, the examples I listed above suggest that we may be dealing with a multiplicity of normalcies. Some of these examples require normal uses or ways of doing things, and it isn’t entirely clear that these can be defined using a single propositional operator. There may be many different sources of normalcy: the cases I listed have to do with human artifacts and habits, whereas some telic eventualities, like the one expressed by Example (7.1) are natural and have nothing to do with human interests or actions.

(7.1) *The sun melted the ice.*

What is perhaps more important, a single, comprehensive normalcy doesn’t seem to be entirely appropriate in the cases that are central for Dowty. Rather than a global inertia that applies to all telic processes, it seems as if each process has its own local inertia, and that these inertias can clash. Imagine an elderly woman crossing a street at an intersection. At same time, a bus is speeding through the same intersection, on a collision course with the poor woman. We use the progressive form of a telic verb to describe each of these eventualities, but both of them can’t follow their inertial trajectories in this case.

I think that these points—together, perhaps, with a certain amount of suspicion about propositional operators inherited from the failure of Dowty’s CAUSE—provide pretty good motives to look for another approach to normality. And a very promising approach is available now that hadn’t been developed in the late 70s, when Dowty was working out his theory of causative constructions. This is John McCarthy’s circumscriptive approach to nonmonotonic logic. See, for instance, [Lifschitz, 1988].

Typically, a circumscriptive theory of a domain is a first-order theory with the usual sort of vocabulary, together with a set of *abnormality* predicates.¹³ A nonmonotonic consequence relation is obtained by minimizing the extensions of these predicates, but this aspect of circumscription is less important for us than the fact that an *abnormality theory* will be part of the entire theory. The abnormality theory will not only tell us what will hold in cases that are not abnormal, but can also provide information about exceptional cases.

¹³An abnormality predicate, of course, is just the negation of a normality predicate or condition, so the two are interdefinable. As long as I am working with a circumscriptive theory, I’ll adhere to the convention that applies there, and will use abnormality predicates as theoretical primitives.

In connection with the application we envisage, it is significant, and useful, that standard circumscription uses first-order abnormality predicates rather than modal operators. The examples that I gave invite an account in terms of the normality of *eventualities*, which we have already decided to treat as individuals. Another advantage, I believe, is that these predicates are local. A normality about a person crossing the street can coexist with a normality about a bus speeding through an intersection. When they compete, circumscription in general allows the different models in which each normality is satisfied. But it is also possible to exclude some of these models by prioritizing the minimization of abnormalities. We can be noncommittal about what happens when normalities compete, or we can go on to create a theory of what happens in such cases.

Circumscription has another advantage. It is used to formalize nontrivial commonsense reasoning domains.¹⁴ This is strategically important, because the semantics of derived morphology is, as a rule, closely related to everyday commonsense reasoning. Constructions like ‘openable’ and ‘movable’ indicate opportunities for action that are closely related to planning and practical reasoning. I think it will be hard to vindicate even moderately ambitious attempts to formalize lexical semantics by obtaining a few small-scale results, like definitions or meaning postulates that yield plausible consequences. Thoroughgoing connections to a theory of commonsense reasoning would have to emerge from a larger project and will be more difficult to obtain and evaluate. But in the long run, I think this broad-based type of confirmation is what we will have to look for.

8. Conclusion

In this study of causative morphology, the logic served as a framework for a formalization project involving the relevant notions: a rudimentary ontology of eventualities and metaphysics of causality and agency. Most of the effort went into formalizing the relevant aspects of agency and causality; relatively little went into translating natural language constructions. The formalization itself was almost entirely first-order, making little or no use of the type-theoretic apparatus of Intensional Logic, and it became complicated enough so that it was mostly presented informally, using the metalanguage. Diagrams were helpful in conveying many of the important ideas.

These features make for a project that is rather different from “classical” interpretations of natural language, such as the one in [Montague, 1973]. But they are similar to Dowty’s early treatment of derived morphology in [Dowty, 1979]. I think we can tentatively conclude that we can expect methodological differences of this sort when we turn to the semantic interpretation of derived morphology.

Despite these differences, the accounts that emerge of word meaning and syntactic meaning belong to the same general semantic framework, and are perfectly compatible.

The central developments in this paper required no changes to the underlying logic, but they couldn’t have been carried out without a NeoDavidsonian, eventuality-based approach to natural language predication. On the other hand, I suggested that closely related phenomena would require the use of a nonmonotonic logic. My working assumption is that a

¹⁴See, for instance, Lifschitz’ solution to the egg-cracking problem in [Lifschitz, 1998].

nonmonotonic version of Intensional Logic will be adequate for formalizing a broader range of constructions to be found in derivational morphology.

Bibliography

- [Bach, 1989] Emmon Bach. *Informal Lectures on Formal Semantics*. State University of New York Press, Albany, NY, 1989.
- [Davis, 1991] Ernest Davis. *Common Sense Reasoning*. Morgan Kaufmann, San Francisco, 1991.
- [Dowty, 1979] David R. Dowty. *Word Meaning in Montague Grammar*. D. Reidel Publishing Co., Dordrecht, Holland, 1979.
- [Krifka, 1998] Manfred Krifka. The origins of telicity. In Susan D. Rothstein, editor, *Events and Grammar*, pages 197–235. Kluwer Academic Publishers, Dordrecht, 1998.
- [Lewis, 1977] David K. Lewis. Causation. *Journal of Philosophy*, 70(17):566–567, 1977. Reprinted in David K. Lewis, *Philosophical Papers*. vol. 2, Oxford University Press, 1986, pp. 159–213. (Postscripts are added in this reprinting.).
- [Lifschitz, 1988] Vladimir Lifschitz. Circumscriptive theories: A logic-based framework for knowledge representation. *Journal of Philosophical Logic*, 17(3):391–441, 1988.
- [Lifschitz, 1998] Vladimir Lifschitz. Cracking an egg: An exercise in commonsense reasoning. Unpublished manuscript, <http://www.cs.utexas.edu/users/vl/mypapers/egg.ps>, 1998.
- [Montague, 1973] Richard Montague. The proper treatment of quantification in ordinary English. In Jaakko Hintikka, editor, *Approaches to Natural Language: Proceedings of the 1970 Stanford Workshop on Grammar and Semantics*, pages 221–242. D. Reidel Publishing Co., Dordrecht, Holland, 1973. Reprinted in *Formal Philosophy*, by Richard Montague, Yale University Press, New Haven, CT, 1974, pp. 247–270.
- [Mueller, 2006] Erik T. Mueller. *Commonsense Reasoning*. Elsevier, Amsterdam, 2006.
- [Parsons, 1990] Terence Parsons. *Events in the Semantics of English: a Study in Subatomic Semantics*. The MIT Press, Cambridge, Massachusetts, 1990.
- [Pearl, 2000] Judea Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, Cambridge, England, 2000.
- [Simon, 1952] Herbert Simon. On the definition of the causal relation. *The Journal of Philosophy*, 49:517–528, 1952.
- [Steedman, 1998] Mark Steedman. The productions of time. Unpublished manuscript, University of Edinburgh. Available from <http://www.cogsci.ed.ac.uk/~steedman/papers.html>., 1998.

[Woodward, 2003] James Woodward. *Making Things Happen: A Theory of Causal Explanation*. Oxford University Press, Oxford, 2003.