

Event Composition and Event Individuation

Robert Truswell

1 Introduction

This chapter explores a consequence of Davidson's (1967) foundational hypothesis, that events are in some nontrivial way similar to individuals:¹ just as an individual can form part of a larger individual, we expect that an event can form part of a larger event. This would imply that at least some events are composed of multiple smaller events. We call this relation **event composition**.

A related issue is the delimitation, or **individuation**, of events. The semantic structures we describe below imply a very large, richly structured, domain of events, including many events that are not likely to correspond to units manipulated by our cognitive and linguistic capacities. We therefore supplement our exploration of the formal semantics of event composition with a set of statements about cognitive and linguistic well-formedness constraints on events.

Below, we first demonstrate composition relations for individuals and events, and show that such relations have *a priori* plausibility (Section 2). Next, we turn in Section 3 to perceptual and cognitive constraints on event individuation. Finally, Section 4 discusses syntactic aspects of event composition and event individuation.

2 Foundations

In formal semantics, individuals are members of the domain of individuals, D_e , typically a denumerably infinite set partitioned into two subclasses, constants and variables. Perceptually, however, it is not so clear what cognitive units this logical characterization corresponds to. How do we know when we have encountered an individual? How do we recognize the members of that set? Such questions are everywhere when logical and cognitive approaches to semantics meet.

Individuals *qua* logical formatives certainly do not match our intuitive notion of *individual*. For instance, it is likely that London, justice, and the Boston Red Sox (the team) are all individuals in the logical sense, but this seems intuitively absurd: London (see Chomsky 2000) is a city, a strange, amorphous region defined in partly political and partly geographical terms, which also functions as a sort of club with gradient membership (if you were born within the sound of the Bow bells and have

¹I adopt this deliberately vague formulation so as to remain noncommittal about two possibilities: that events simply are individuals, or that the domain of events is disjoint from the domain of individuals, but has a similar structure. We will not attempt to distinguish these two options here.

lived there all your life, you're part of London and that's that; if you live in Islington but only moved there a month ago and travel a lot, then you might feel semi-detached from London; and if you live in the western fringes of Essex, there is probably a whole sociological literature on whether you are part of London). Meanwhile, justice is intangible, an abstract concept that is 'done' or 'served' but, unlike many other things that are done (ballroom dancing, for instance), is somehow not event-like. We have quite clear intuitions about what constitutes justice, but do we really see an individual here? Finally, the Boston Red Sox — that's *nine* individuals (plus substitutes and managers and so on), not one.

We will surely not be able to bridge the gap between the logical definition of 'individual' and our intuitions about individuals, and we should not try — at least, not right away. 'Individual', as a term in our logical vocabulary, is better characterized in terms of its relations to other parts of the logical vocabulary. In standard logics for natural language, individuals are primitive elements from which other categories (such as predicates) are recursively constructed set-theoretically, and how that relates to any intuitions about what counts as an individual is a separate question.

However, there are correspondences between elements of natural language syntax and the model-theoretic translation of those constituents, and these correspondences can help us relate individuals as logical units and as cognitive units. If we have rules of thumb about which natural language constituents denote logical individuals, and intuitions about perceptual correlates of those constituents, then we can infer rules of thumb, imprecise but still useful, about perceptual correlates of the logical notion of *individual*.

Here are two rules of thumb about natural language and logical individuals:

1. Noun phrases canonically denote individuals.
2. Individuals canonically function as arguments to (first-order) predicates.

The qualification 'canonically' is important: there is no way to determine *a priori* the denotational behaviour of elements of natural language, and if there were, the validity of Davidson's hypothesis (along with many other foundational semantic matters) would be settled automatically. There are several well-known exceptions to these heuristics (quantified noun phrases, for example, denote objects of some type other than *e*, usually $\langle\langle e, t \rangle, t \rangle$), but these heuristics show the virtue of intuitively outlandish claims that London, or justice, or the Boston Red Sox, are individuals. First, *London*, *justice*, and *the Boston Red Sox* are noun phrases; secondly, their denotations can all function as arguments to first-order predicates.

- (1)
 - a.
 - (i) London is complex / I resent London.
 - (ii) Jeremy Clarkson is complex / I resent Jeremy Clarkson.
 - b.
 - (i) Justice has been done / I want justice.
 - (ii) Jeremy Clarkson has been done / I want Jeremy Clarkson.
 - c.
 - (i) The Red Sox never make it easy for their supporters / Many people still support the Red Sox.
 - (ii) Jeremy Clarkson never makes it easy for his supporters / Many people still support Jeremy Clarkson.

Now, the crucial point: if London is an individual in this sense, then so is Camden, or the Tube, despite the fact that these are subparts of London. England and Europe are individuals, despite the fact that London is part of these. The same thing with the Red Sox: if the Red Sox is an individual, surely Dustin Pedroia and Major League Baseball are individuals too, in the same sense.

This tells us something basic about D_e : individuals can be part of other individuals. This is possibly not true of the pre-theoretical, perceptually grounded notion of ‘individual’ (if I admit that Jeremy Clarkson is an individual, I would be reluctant to admit that his eyebrows are also individuals), but there you go.

Following Link (1983), I assume a range of mereological, or part-whole, relations among individuals. Link distinguishes between **atomic individuals** and **plural individuals**, broadly mirroring the singular-plural distinction found in many natural languages. *John* and *Mary* denote atomic individuals (say j and m respectively), but the coordinate noun phrase *John and Mary* still denotes an individual, according to the above rules of thumb: aside from the fact that it triggers plural agreement, the distribution of *John and Mary* is very similar to the distribution of *John*. For example, both can function as arguments to predicates like *danced*. If we believe that *danced* denotes a predicate of type $\langle e, t \rangle$, then it would make sense for both *John* and *John and Mary* to be of type e .² Accordingly, we say that *John and Mary* denotes the plural individual $j \oplus m$, and that j and m are **individual parts** of $j \oplus m$.³ An atomic individual is then an individual with no proper individual parts.

Even atomic individuals have parts, though. John has four limbs and 20 digits and 32 teeth and 206 bones and a nose, and so on, but we do not perceive these as individuals independent of John, in the sense in which John remains an individual even when considered as part of $j \oplus m$. We do not look at John and see 264 individuals; we see only j . Nevertheless, there is clearly some mereological relation between John’s nose (call it n) and John: the stuff that constitutes n is part of the stuff that constitutes j . We say that n is a **material part** of j , but not an individual part of j (because j is an atomic individual, and atomic individuals do not have individual parts). If x is an individual, all individual parts of x are also material parts of x , but there can be material parts of x which are not individual parts of x .

Mereological relations among individuals, then, are pervasive. Capturing those relations requires a domain of individuals with a surprisingly rich structure. For example, Link’s analysis entails that multiple individuals can be spatiotemporally coextensive. Facts about natural language seem to require this. As Link discussed, a new ring can be made from old gold. The gold and the ring are coextensive, but are nonetheless different individuals, on the assumption that oldness and newness are mutually exclusive.

If events are similar to individuals, we might expect them to show similar richness of structure. Indeed, they do. In a frequently discussed modification of an example from Davidson (1969), a sphere can rotate quickly while heating up slowly. Assuming that quickness and slowness are mutually exclusive in this case (though not in all cases,

²We ignore the possibility that *John and Mary* is a quantifier, of type $\langle \langle e, t \rangle, t \rangle$.

³Following Link, I use the symbol \oplus for this **individual sum** relation, and $+$ for a **material sum** relation to be introduced presently. For all x_1, x_2 , x_1 and x_2 are material parts of $x_1 + x_2$; x_1 and x_2 are individual parts of $x_1 \oplus x_2$. I also use $x_1 \subseteq x_2$, $x_1 \subset x_2$ for ‘ x_1 is a part (or proper part, respectively) of x_2 ’, and $x_1 \supseteq x_2$, $x_1 \supset x_2$ for the converse.

as shown by Kamp 1975 and Higginbotham 1985), the rotating event and the heating-up event must be distinct, despite being spatiotemporally coextensive.⁴ This chapter explores mereological relations between atomic events and their material parts; see Lohndal's chapter in this volume for discussion of plural events.

We begin by characterizing event composition, again by analogy with individuals. Link's logic guarantees (we omit the details here) that for any atomic or plural individual x , there is some **stuff** (or portion of material) that constitutes x . Moreover, stuff can be subdivided arbitrarily. Finally, portions of stuff are still individuals in their own right. These considerations jointly entail that any individual x can be subdivided into a set of individuals $\{x_1, \dots, x_n\}$, none of which have any material parts in common, which jointly constitute x (the stuff constituting x is the same stuff constituting $x_1 + \dots + x_n$). Let's call the relation between $\{x_1, \dots, x_n\}$ and x **composition**.

Analogues of all of the above can be found in the domain of events (see, in particular, Bach 1986: 5, where the relation 'events:processes :: things:stuff' was proposed). Specifically, the relationship between stuff and atomic individuals mirrors the relationship between stuff and atomic events.⁵ Just as an atomic individual can be composed of a set of portions of stuff, so can an event. That is, a relationship of **event composition** can hold between a set of portions of stuff and an atomic event. We will also talk about **decomposition** of an event into a set of **subevents**. This is to be taken as the converse of event composition.

As with individuals, then, we can construe some events as composed of smaller events, without implying that the composite event is a plurality. (2) is an example.

(2) Michael built a snowman.

Davidson tells us that (2) denotes the proposition that there exists an event temporally located prior to speech time, of snowman-building carried out by Michael. However, building a snowman has its own internal structure: you roll a giant snowball for the body, by pushing a smaller snowball through a patch of snow, roll another for the head, and adorn the result with carrots and scarves and so on. Each of those stages is an event in its own right; jointly, they compose the snowman-building event.

Events can be decomposed recursively. Pushing a giant snowball around is a process which stops when you have a sufficiently giant snowball. This process is composed of a series of iterable smaller events of taking a step and rolling a snowball in front of you. Taking a step involves coordinating a set of muscle movements: internally very complex (if you don't believe me, ask my baby son), even if we, as adults, now often take the complexity for granted. Muscle movements are probably really about things happening to electrons or ions or similar, for all I know.

⁴Davidson (1969) proposed that events are identical iff they have the same causes and effects. It could be true here that the spinning causes the heating but not *vice versa*, further suggesting that the two are nonidentical.

⁵As in the above quote, Bach's term for the event analogue of stuff was *processes*. Link (1997) expanded this use, defining processes as portions of space-time which may be reified as events or as individuals. I maintain Link's use of a single term for portions of material underpinning events or state, but avoid the term *processes*, which is used in other ways in the literature and below. If any reader objects to the conflation of these two types of stuff, the context should make it clear when I am talking about event-stuff, and when I am talking about individual-stuff.

This suggests that the domain of events has a similar internal structure to the domain of individuals: there are discrete atomic events, which can be summed into plural events which we continue to ignore, there are continuous portions of stuff which can be summed and subdivided arbitrarily, and there are two mereological relations, one determining which atomic events are part of which plural events, and one (our focus here) telling us which portions of stuff are part of which events. Finally, we add a relation of event composition between an atomic event e and a set of nonoverlapping events $\{e_1, \dots, e_n\}$, such that the same stuff constitutes e and $e_1 + \dots + e_n$.

A major question for this chapter is how this structure relates to events as perceptual units, manipulated in our reasoning processes, and as described by ‘simple’ natural language predicates.⁶ I will make the simplifying, though perhaps not innocent, assumption that natural language is a good guide to events as cognitive units (see also Zacks & Tversky 2001 and Wolff 2003 for evidence of congruence between events as perceptual and linguistic units). Moreover, I will assume that simple linguistic event descriptions pick out atomic events: a simple VP, for example, denotes a subset of the set of atomic events (Higginbotham 1985). To rephrase the above question, what kind of sets of events can simple event descriptions denote? Or to turn the question on its head, what can we learn about linguistic event descriptions, and perhaps about events in cognition, by examining their denotations in a relatively highly structured domain like the domain of events described above?

As an example, apart from some microscopic modifications, if e is an event of Michael building a snowman, no subpart of e is also an event of Michael building a snowman: if *Michael built a snowman* is a true description of e , it is not also a true description of any $e' \subset e$. In the terms of Krifka (1989), events of building a snowman are **quantized**. In contrast, for a given snowball s , an event e_1 of pushing s , combined with a contiguous event e_2 of pushing s , gives a larger event $e_1 + e_2$, which is also an event of pushing s . In Krifka’s terms, snowball-pushing is **cumulative**, or closed under joins. This is because snowman-building is **bounded**: it ends when the snowman is built. Pushing a snowball is **unbounded**, in contrast: there is no intrinsic point at which you have to stop pushing. This recapitulates some of the distinctions made among aspectual classes by Vendler (1957; see also Mittwoch’s chapter) in terms of denotation in the domain of events: quantized events are telic (accomplishments and achievements), while cumulative events are atelic (activities and states).

We take this to indicate that linguistically and perceptually relevant event types conform to particular ‘shapes’, as characterized by properties of event descriptions such as quantization and cumulativity, and reflected in the fixed set of aspectual classes (see a range of work such as Moens & Steedman 1988, Pustejovsky 1991, and Ramchand 2008 for proposals as to the form and causal origin of those templates). In particular, in this chapter we adopt a common vocabulary whereby an event consists maximally of two components, a temporally extended **process** and an instantaneous **culmination** at which a result state is reached. By including or omitting these two components, we derive four event shapes, as analogues of Vendler’s four aspectual classes.

⁶By a ‘simple’ predicate, I have in mind a noun, adjective, or verb with its arguments, as opposed to a more complex predicate formed by coordinating VPs, negation of events, and so on. These more complex event descriptions are partly covered in Lohndal’s chapter.

1. **Culminated processes** (process + culmination) \approx accomplishments (e.g. *run a mile*)
2. **processes** \approx activities (e.g. *run*)
3. **culminations** \approx achievements (e.g. *hiccup*)
4. \emptyset (neither process nor culmination) \approx states (e.g. *exist*)

Still following Vendler, we adopt diagnostic tests for the presence of a process or culmination. An event with a process is felicitous in the progressive, whereas an event without a process is only felicitous in the progressive if coerced into an iterated reading, for example.

- (3)
- a. John is running a mile.
 - b. John is running.
 - c. John is hiccupping [iterated reading only]
 - d. #John is existing.

Meanwhile, an event with a culmination is infelicitous with *for*-PPs describing the temporal extent of the event, again disregarding possible coercion effects.

- (4)
- a. #John ran a mile for five minutes.
 - b. John ran for five minutes.
 - c. John hiccupped for five minutes [iterated reading only]
 - d. John existed for five minutes.

In Krifka's terms, processes and states are cumulative, while culminated processes and culminations are quantized. With one exception which will not concern us here, quantized event predicates and cumulative event predicates are disjoint. Given that culminated processes and culminations are distinguished by the presence of a culmination, we hypothesize that a major class of quantized event predicates can be distinguished by the presence of a culmination. This is not logically necessary (quantized event predicates could equally well be distinguished by a particular initial state, or a distinctive subevent occurring precisely halfway through the runtime of the event). Rather, it reflects an apparent fact about natural language event descriptions: culminations are special, in that they distinguish a major class of quantized events from cumulative events.

The two subtypes of quantized events, culminations and culminated processes, can be distinguished on the basis of **durativity**. Culminated processes are durative in that they have proper subevents at the same level of granularity. In contrast, culminations (like dying or hiccupping) are construed as instantaneous. Of course, at a microscopic level, culminations do have subparts, but we turn a blind eye to these in our linguistic event descriptions. Such **coarse-graining**, whereby the fine internal structure of a given individual or event is linguistically invisible, appears to be a pervasive part of linguistic descriptions.

Meanwhile, processes and states are distinguished by **dynamicity**: processes typically involve change, and when they don't (events such as resisting or hindering), they

at least involve a dynamic opposition of forces (see Copley's chapter). This is not necessarily the case with states, which describe properties construed as static. In fact, we will not talk any more about states in this chapter, as the vocabulary we develop below cannot obviously be usefully applied to states.

In the following section, we discuss properties of events at different scales, from a ball rolling to an ice age. The four event shapes distinguished above give a unifying organizational principle across events on different scales: events, at any level of granularity, can be partitioned into the same four shapes. In other words, the forms remain the same; the perceptual basis for individuating events according to those forms varies.

3 Constraints on Event Individuation

The mereological relations sketched in the previous section are in principle unlimited in scope: any event, or any individual, can be decomposed into smaller parts. This means that we can represent quite implausible individuals and events, should we want to: decompose e and e' into arbitrary sets $\{e_1, \dots, e_n\}, \{e'_1, \dots, e'_n\}$ of subevents, compose $e_i + e'_j$ into a new event, for arbitrary i, j , and so on. Logically, this ability to represent arbitrary relations among objects is a good thing: notions like plausibility are rarely useful in logics. However, it is natural to complement this logic with a series of statements describing which events are likely to be composed in this way. This is the question of the perceptual and linguistic bases of event individuation.

Almost immediately after Davidson (1967) had argued that events were particulars, described by action sentences, Davidson (1969) raised the question of the individuation of events. Davidson's concern in that paper is rather different from our concern here, though. Davidson is really concerned with identity relations among events, and when statements of the form $\text{te}.P_1(e) = \text{te}.P_2(e)$ are true. Our question is only distantly related, and comes back to the concerns about individuals discussed at the start of this chapter: how do cognitive and linguistic considerations relate to the logical characterization of events sketched above?

We begin by asking what kind of relation can hold between the two components of our event shape, the process and culmination. It is often assumed (following McCawley 1968 and Dowty 1979, among many others) that the process is related to the culmination by a causal relation such as 'directly causes' or 'leads to'. That may be true in many cases: if the falling rock smashes the vase, then there is a process of the rock moving along a particular trajectory, which directly causes the culmination of the vase assuming a very different physical form. Likewise, if the author wrote a novel, then there is a writing process which directly causes the existence of the novel.⁷ However, discussion in Davidson (1969) of the relationship between poisoning, killing, and

⁷Even here, things are more strained, in that there is no miraculous instantaneous event at which the book appears, or becomes fully formed. As with many acts of creation, an author writing a book engages in a process which results in the book increasingly approximating completedness, but the book is only actually finished when the author calls it a day, stops tinkering, and moves on. This is the distinction between culminated processes which are **measured out** by their objects, in that there is a homomorphic mapping between subparts of the event and of the object, and culminated processes where the subparts of the event bear no such direct relation to the subparts of the object. See chapters by Mittwoch and Verkuyl, and references therein.

dying already showed that the directness of direct causation is quite elusive: *A* may kill *B* by pouring poison into his bottle of scotch, but that action did not directly cause *B* to die in any intuitive sense: adding poison to the scotch could be separated from *B*'s death by any amount of spatiotemporal distance, and requires a helping hand from *B* (who must consume some of the scotch if *A* is to successfully kill him). We may agree that *A* killed *B* in such a scenario, but this does not necessarily entail that *A*'s actions (the process) directly caused *B*'s death (the culmination) — see also Fodor (1970). In fact, Copley & Harley (2014) discuss several linguistic structures suggesting that the relationship between process and culmination cannot be one of straightforward direct causation, at least not in the actual world. The occurrence of the process component of a culminated process does not entail the occurrence of the culmination, when on any commonsense definition of direct causation, it should.⁸

The best-known example of this is the so-called **imperfective paradox** (Dowty 1979, among many others).⁹ The progressive form of an activity predicate is taken to entail the perfective form, as in (5).

- (5) a. John was running. →
 b. John has run.

However, the progressive form of an accomplishment predicate usually entails the process component, but not the culmination. That is, (6a) entails (6b) and (6c), but not (6d).

- (6) a. John was painting a still life. →
 b. John was painting. →
 c. John has painted. [More idiomatic: John has done some painting.] ↗
 d. John has painted/will have painted a still life.

The reason for this failure of entailment concerns the semantics of the different aspectual forms. (5a) describes an ongoing event of John running, with the reference time situated within the event time. If some portion of the event time precedes the reference time, then we can conclude that some part of the process of John running has already taken place: John has run. In contrast, painting a still life is a quantized, rather than cumulative, event predicate. If the reference time is situated within the event time, that means that some portion of the process has taken place: John has done some painting. However, the completion of the still life is still in the future, and consequently may not be reached. We therefore cannot conclude that John has painted a still life: (7a) is a contradiction, but (7b) is not.

⁸It may be useful to have a rough definition of direct causation in mind here, for comparison. Following Dowty (1979) and Lewis (1973), causation is often seen as a counterfactual dependency: if *C* causes *E*, in the most accessible worlds like w_0 , if *C* hadn't happened then *E* wouldn't have happened. Such dependencies can be grouped into **causal chains**: e_1 causes e_2 , which causes e_3 ; if e_1 hadn't happened then e_2 wouldn't have happened; if e_2 hadn't happened then e_3 wouldn't have happened. A relation of direct causation holds in a 2-member causal chain, with no intermediate events at the same level of granularity. Although every aspect of this is controversial (see Copley & Wolff 2014 for critical discussion), the yoking together of immediate cause and effect is taken to correspond to the directness of direct causation; it is this which is missing in the examples to follow.

⁹As discussed in Mittwoch's chapter, it is now widely accepted that the imperfective paradox is not actually a paradox, but rather a data point that should shape our theories. The name, however, has stuck.

- (7) a. #John may have been running yesterday, but John has still never (successfully) run.
 b. John may have been painting a still life yesterday, but John has still never (successfully) painted a still life.

All of this is intuitively at odds with a semantic representation in which the painting process directly causes the existence of a still life: (7b) shows that the one can occur without the other, while theories of causation yoke the two subevents together. Dowty, in order to reconcile the imperfective paradox with his analysis of accomplishment predicates as lexicalized instances of direct causation, included a modal component: if John is painting a still life, then the still life may not be completed in w_0 , the actual world, but it is completed in any **inertia worlds**, in which there are no unforeseen interruptions to prevent the painting event running its course.

A second case comes from the now widely documented phenomenon of **non-culminating accomplishments** (Travis 2000, Bar-El et al. 2005). In a range of typologically unrelated languages, the culmination component of an accomplishment predicate is an implicature rather than an entailment, and can be explicitly contradicted. (8a) gives an example from Malagasy, while (8b) gives an example from St'át'imcets (examples taken from Copley & Harley 2014); Mittwoch's chapter contains further examples from Hindi, Mandarin, and Japanese.

- (8) a. Namory ny ankizy ny mpampianatra, nefa tsy nanana
 PAST.AV.meet the children the teachers but NEG PAST.have
 fotoana izy
 time they
 'The teachers gathered the children, but they didn't have time' (Travis 2000:173)
 b. k'ul'-ún'-lhkan ti ts'lá7-a, t'u7 aoy t'u7 kw tsukw-s
 make-TR-1SG.SU DET basket-DET but NEG just DET finish-3POSS
 'I made the basket, but it didn't get finished' (Bar-El et al. 2005:90)

Phenomena like the progressive and non-culminating accomplishments raise doubts about analyses which implicate direct causation in the subevent structure of culminated processes, and I adopted Krifka's terminology above partly because it permits a greater degree of agnosticism about precisely how processes relate to culminations. In fact, I will claim that the nature of the relationship between process and culmination depends on the perceptual nature of the event itself.

Returning to our discussion of snowman-building from Section 2, note that events at different granularities have quite different characters (see also Zacks & Tversky 2001). The smallest events are almost theoretical in nature: we *could* observe them in the lab, in principle, and we know deep down that they exist outside the lab as well, but they are not directly observable in day-to-day life. The smallest events we actually observe are roughly on the scale of muscle movements: we can sometimes see muscles contract as people walk, and we are certainly aware of the alternating pattern of more-or-less bounded movements of the limbs: although it is difficult to decide precisely where one step stops and the next starts, we readily accept that we have seen a string of steps and pushes and so on. At coarser grains still, such as the

rolling of a giant snowball, events are identified by the perceived intentions and goals of agents. For example, Dowty (1979) discusses pauses in events: we recognize an event of Michael rolling a snowball even if he took a 30-second breather in the middle, or left the scene altogether in order to recruit friends who could help him make the snowball even bigger. I take this (unlike Dowty) to be related to a perceived continuity of intention in such cases, even if there is no corresponding continuity of action. Even larger events, such as wars, start and end according to a diplomatic process quite remote from the vast majority of the action on the battlefield. You know when a war (as opposed to other forms of military aggression) is going on because of declarations of war, formal ceasefires, and so on. Likewise, most of the activity in an Apollo mission happens on the ground — the spacecraft moving through space is just the salient tip of the iceberg.

However, there is no perfect correlation between the size of an event, construed as its spatiotemporal extent, and its perceived physical, intentional, or analytical nature. An actor's raised eyebrow might be an exquisitely planned intentional event, but it is still small in scale compared to some physical events like natural disasters.

Rather, these different schemata for describing events reflect different perceptual modes. We perceive physical changes and interactions among physical objects; we attribute intentions to certain animate individuals and recognize physical and other effects as linked to those intentions; and we understand that certain powerful bodies (whether playwrights or presidents) may effectively control the actions of large groups of other individuals in some 'strategic' way (a deliberately vague characterization under which I hope to group everything from the shepherd's control of the sheep through the intermediary of a dog in a sheepdog trial, to the role of the composer and the librettist in an opera, or that of an arch-manipulator using the power of suggestion to get his own way). Finally, an analyst may uncover order in a set of events that was not apparent to any agent involved at the time (emergent phenomena like stock market crashes or the migration out of Africa are likely examples; see also Link 1997 on historical events such as the French revolution). There are surely other types of perceptually and linguistically relevant events (those described by psychological predicates, for instance), but we will concentrate exclusively on the above in this chapter.

Each of these types of perceptual organization comes with its own set of well-formedness constraints. We are more likely to perceive a set of happenings as an atomic event to the extent that they match these constraints. In the next four subsections, we will discuss the articulation of these different types of events and their descriptions, and show that their internal shapes are common to all these event types.

3.1 Physical events

The happenings we perceive as physical events are characterized by a set of common-sense beliefs about the way the world works sometimes grouped together under the heading naïve physics. I will not attempt a coherent statement of a naïve physical model, but give some examples showing relevant aspects of such a model.

Three subtypes of physical events matching the different event shapes from Section 2 are motion and other unbounded physical processes (9); culminations, instantaneous

changes of state (10); and culminated processes (11). A selection of diagnostic tests as described in Section 2 confirms the shape of each event.

- (9) a. The river flowed (for five minutes) / the river is flowing.
b. The flag fluttered.
c. The lava cooled.
- (10) a. The balloon burst (#for five minutes) / #the balloon is bursting.
b. The vase bounced.
- (11) a. The wind blew the ball into the lake (#for five minutes) / The wind is blowing the ball into the lake.
b. The falling tree crushed the car.

The events described above are dynamic and spatiotemporally continuous: a direct interaction between a set of one or more objects associated with tendencies to motion or to rest, as described in Talmy (1988 *et seq.*). For example, any one individual may tend to move in a variety of ways: a river flowing is a fairly stable type of motion, whereas a fluttering flag is less predictable in terms of both orientation and speed of motion. Causal relations like those in (11) then often emerge from local interactions between objects associated with different such tendencies to motion or to rest: the tree has a tendency to fall, the car has a tendency to stasis, and the tree overcomes the resistance from the car.

However, not all examples of physical causation are local like this. A classic example of action at a distance, or nonlocal physical causation, is turning on a light (intentionally or accidentally) by flicking a switch: the switch can be any distance from the light (it is not inconceivable that someone at Ground Control can flick a switch and turn on a light on a space station). The relationship between the switch and the light is otherwise the same as that between the tree and the car, though. This suggests that physical events are not always spatially continuous. ‘Ballistic’ causation may well be the default case of physical causation, in that action at a distance tends to involve a special trigger like a switch, while we understand that any moving object can bump into any other. It does not seem to be the only causal configuration that we perceive, though.¹⁰

Likewise, we understand that when we use a vending machine, the button that we press causes the gummi bears to fall into the tray, even though there are five nerve-racking seconds of indeterminate whirring between the button press and the reward. This is temporal, as opposed to spatial, action at a distance: the button press eventually makes the gummi bears fall into the tray, but we have no idea what, if anything, is going on in between. Although these examples seem more marginal than spatial action at a distance (a switch could turn a light on thousands of miles away, but would we really perceive any relatedness in a switch that made gummi bears appear in hundreds of years’ time?), it seems that physical events are, at least sometimes, temporally as well as spatially discontinuous.

¹⁰Of course, microscopically, such events are spatially continuous: flicking the switch transmits a signal through some medium like a wire, and this causes the effect through a chain of local physical causal relations. The point is that our naïve physics doesn’t see the microscopic intermediate steps, and is happy to associate the more tangible initial cause and final effect directly. See also Rozenblit & Keil (2002).

Finally, if a gust of wind blew a ball halfway down the hill, the ball came to rest, and then a second gust of wind blew the ball to the bottom of the hill, we can report the path from top to bottom with a single event description: *The ball rolled down the hill*. This is another example of coarse-graining: the temporal discontinuity is construed as an insignificant interruption in the event, whereas with the vending machine, the pause is just part of what happens when you push a button.

These exceptions to spatiotemporal continuity suggest that there are actually two different systems lumped together here under the heading of ‘physical events’: spatiotemporally continuous ‘mechanical’ physical motions, and causal relations between cause and effect, which may occur across some spatiotemporal distance, at least given the fairly shallow causal theory that our naïve physics apparently encompasses (see again Rozenblit & Keil 2002 on the ‘illusion of explanatory depth’). Physical events are phenomena which can be understood to this extent purely in terms of such relationships between perceptually accessible individuals.¹¹ In the following section, we will contrast these physical causal relations to causal relations between intentions and goals.

3.2 Intentional events

We construe a subset of individuals (primarily animate individuals) as behaving intentionally:¹² these individuals have goals, and act rationally to reach those goals. I will say that a set of events, construed as an agent’s actions aiming at a goal, jointly compose an **intentional event**, delimited by that goal.

Of course, an agent acting intentionally can also be considered as a purely physical object (animacy entails physicality but not *vice versa*). This drives the ambiguity, already noted in Jackendoff (1972), of (12).

(12) John hit the wall.

On one reading, John is just a lump of flesh, flung against a wall by some other force. On the other reading, John is acting intentionally, propelling his fist into the wall. The former reading describes a purely physical event; the latter is intentional.

Such a distinction is often recognized in the literature, for example in Jackendoff’s (1990) distinction of the ‘action tier’ within his semantic representation. This distinction allows Jackendoff to maintain a largely localist approach to thematic roles, in terms of notions like movement of a theme along a path from source to goal, while acknowledging that this is distinct from questions such as who is acting on who. Jackendoff claims that in the purely physical reading of (12), John is just a theme, while in the intentional reading, John is both a theme and an agent.

¹¹By ‘perceptually accessible’ individuals, I have in mind, first and foremost, visible individuals, but also more abstract individuals which can be perceived but are otherwise quite unlike ordinary individuals, such as the weather.

¹²There are some hints of intentional language in our talk about the weather, for instance *The sun is trying to shine*. The thrust of the present discussion forces me to consider these as nonliteral uses of *try*, rather than intentional in the above sense: we cannot describe a weather forecast by saying that *The sun is planning to shine at 3pm*; nor can we say *The sun is shining* when the sun is presently dispersing cloud cover, which will permit it to shine in the future. No-one attributes such goal-oriented behaviour to the weather, even metaphorically.

Many people claim that the relationship between the intentional event and the physical event in (12) is causal. For example, in the terminology of Ramchand (2008), we might say that in the purely physical reading of (12), *John* is the subject of a process which causes John to come into contact with the wall. In the intentional reading, *John* is also the subject of an initiating event which causes that process. Similar ideas are discussed at length in Pietroski (2000). However, Truswell (2011) argued that such approaches are ultimately unsatisfactory: the relationship between the intentional and the physical event is not merely one of the intentional event causing a physical event which is independently asserted to exist. Rather, the intention defines and delimits the event, and the action realizes the intention.

The evidence for this comes from the progressive test described in Section 2. Following Reichenbach (1947), the progressive locates the reference time within the run-time of the event itself, and so is used to describe ongoing processes or events *en route* to completion.

This means that we can use the descriptive content of VP to tell us what kind of event we are in the middle of. With a purely physical event like the ball rolling down the hill, we have clear intuitions about when (13) can felicitously be uttered.¹³

(13) The ball is rolling down the hill.

(13) can felicitously be uttered from the moment the ball starts moving down the hill, until it reaches the bottom; even (because of coarse-graining) during a sufficiently brief hiatus in the middle. (13) cannot be uttered before the ball starts moving, even if it is clear that the ball is about to roll down the hill (because the wind is picking up, for example); and (13) cannot be uttered when the ball reaches the bottom, even if it carries on moving. The spatial progress from top to bottom delimits the event.

Intentional events can be bigger than this. More specifically, they start earlier. If we see a very round man limbering up at the top of a hill, and we infer that he is preparing to roll down it, we can say something like (14).

(14) Hey, look! The round man is rolling down the hill!

When we say this, the round man is not necessarily moving down the hill at all, but we infer his intention, and also infer that his current actions might rationally be expected to lead to fulfilment of that intention. That is enough for the round man's limbering up to count as part of a rolling-down-the-hill event: the physical rolling down the hill is a proper subpart of the intentional rolling down the hill, and we can use (14) to describe the ongoing intentional event.

Similar effects are reported, from a different perspective, in Wolff (2003). In a series of experiments, Wolff showed that purely physical events were delimited by direct causation, but that intentional events could be more inclusive.¹⁴ One example of this distinction involved a pair of animations. In the first, three marbles were shown.

¹³Do not be misled by **futurate** uses of the progressive like *The ball is rolling down the hill at 3pm next Tuesday*. Such progressives report on established plans, as described by Copley (2008), and so do not tell us about purely physical events.

¹⁴Wolff uses these terms differently to me. I stick to the characterization implied above, while hopefully preserving the essence of Wolff's account.

The first marble rolled into the second, which in turn rolled into the third. In the other animation, the first marble was replaced by a hand, which pushed the second marble into the third. Although the perceived physical relations are essentially identical in the two cases, participants reported seeing two distinct events in the first animation, but only a single event in the second animation. As a linguistic correlate of this difference, participants typically described the chain of causal relations in the first experiment using periphrastic causatives like (15b), but could describe the second animation using lexical causatives like (16a).

- (15) a. #The first marble moved the third marble.
b. The first marble made the third marble move.
- (16) a. The man moved the second marble.
b. The man made the second marble move.

Wolff interprets this result as showing that perceived intention increases the likelihood that a single event will be perceived: participants see that when the hand pushes one marble, the person in the animation intended to move the other marble: moving the first marble **enables** him to move the second. That, according to Wolff, is sufficient to license perception of a single event.

This suggests that intentional events are bipartite: they are **actions** (often processes) which lead to a **goal** (often a culmination). As with physical events, intentional processes and intentional culminations can be found in isolation, or combined in a culminated process. These three possibilities are illustrated in (17)–(19). In each case, the intentional character of the event is not in doubt.

- (17) a. John is working out.
b. John worked out for hours.
- (18) a. John is spitting. [Iterated reading only]
b. John spat for hours. [Iterated reading only]
- (19) a. John is building a snowman.
b. John built a snowman #for five minutes.

In at least the case of the culmination (18), the physical event of spitting is coextensive with the intentional event of spitting; in the other cases, it is certainly not guaranteed that there is a single recognizable physical event that corresponds to the range of activities involved in working out or in building a sandcastle. This has already been discussed for building a snowman in Section 2. Similarly, working out subsumes a range of physically quite distinct activities. There is no clear physical continuity between doing sit-ups and using a rowing machine, for example. It is only the continuity of intention that justifies the grouping of such disparate activities together as a single event.

Intentional events necessarily involve action: a pure intention does not seem to be sufficient to define an intentional event, unless the agent is actually doing something about it. I intend to die happy and fulfilled, but that is not sufficient to license utterances like (20) in my current state (quite healthy, thank you very much).

(20) #Rob is dying happy and fulfilled.

There are further constraints on how actions relate to intentions. For example, we cannot utter (14) if we see the round man at home, eating breakfast before heading to the hillside: even if the round man knows that he is eating a hearty breakfast to prepare himself for the ordeal that lies ahead, we typically would not look at the breakfast and infer a link to a plan to roll down a hill. Likewise, we cannot use (14) if the round man is limbering up, at the top of the hill, intending to BASE jump off the summit, but *we* know that the wind is picking up, and will send him rolling down the hill before he gets a chance to jump. In the first of these cases, the action of eating breakfast is not sufficiently closely related to allow us to read the intention off the action; in the second, we know that the action will lead to a culmination other than the intended one; but we cannot describe that combination of an intention and an unintended culmination with a single verb. This, again, tells us that action and goal form a bipartite structure, analogous to cause and effect in physical events.

A complete theory of physical events would necessarily contain statements of how physical causes relate to their effects, a vast and fraught topic barely touched above (see also Thomason's chapter). Analogously, intentional events raise the question of how we determine the relationship between actions and goals. This latter challenge is equally intractable, and certainly anything like a full theory is beyond us. However, I think we can insist on at least the following three aspects of intentional event descriptions:

1. The agent must believe that there is a relationship of causation or enablement between action and goal.
2. The agent's action must be part of a plan, evident to the speaker, to reach the goal.
3. The plan in question must be minimal, in a sense to which we return below.

The first condition is designed to exclude cases where an observer can see consequences of an agent's action that an agent cannot. For example, let us assume that a common outcome of rolling down hills is broken ribs. A common trait among round men who roll down hills is blissful ignorance of the dangers they face. When the round man is limbering up, we might know that he is preparing to do something that will land him in hospital, but that is not part of the round man's plan. We can still felicitously say (14), but we cannot say (21).

(21) Look! The round man is breaking his ribs!

The second condition is intended as a guarantee of rationality on the part of the agent. Although there is no guarantee of such rationality, rational plans stand in more regular relations to physical events. Intentional events are therefore easier to perceive and describe to the extent that the plans underpinning them are rational: it does not matter if the agent sees a link between his actions and his goals, if the speaker cannot see the link and cannot expect his audience to see the link.

The third condition imposes a minimality requirement on the plans underpinning intentional events. We can think of a plan as roughly analogous to a chain.¹⁵ The intention behind the minimality condition is that elements of that chain must all be related in certain ways, so that there can be no superfluous links in the chain. A rough formulation of the minimality condition is in (22).

(22) **Minimality condition on plans in intentional events**

A minimal plan consists of a series of steps, s_1, \dots, s_n , such that:

- a. for each s_i, s_{i+1} ($1 \leq i < n$), s_i either causes or enables s_{i+1} , and
- b. no well-formed plan can arise from omission of any step s_i ($1 \leq i < n$).

To see the point of the minimality condition, consider that just as intentional events can be larger than physical events, they can also be more temporally discontinuous. We saw that small pauses could be included within physical events, but plans can be put on hold almost indefinitely before they are resumed. The process of a builder building a house involves a lot of basic building activity (putting one brick next to another and so on), but also two different types of non-building activity. On the one hand, there are the preparatory activities, the builder's equivalent of the round man limbering up; on the other hand, there are pauses in the activity of varying lengths where the builder is not engaged in anything directly related to the building of the house. The builder goes home every afternoon, and may disappear for a few days to work on something else altogether. In Canada, a lot of construction work grinds to a halt for weeks or months in the depths of winter. All of these pauses are normal, or even inevitable, but they are of a different status to the preparatory activities. Without the preparatory activities, the house wouldn't get built; without the pauses, the house would still be built. The preparatory activities are part of a minimal plan as characterized above; the pauses are not, and can only be included as part of an intentional house-building event if coarse-graining allows us to ignore them.

Such coarse-graining is partially contextually determined: because Canadian winters are more severe than British winters, Canadians seem to expect long pauses in construction activity over winter, whereas British people do not. In Canada, it is normal to say that someone is building a house next door, even if cold weather has prevented any progress from being made for months. In Britain, if nothing happened for months, people would be more likely to assume that the project had hit the rocks.

Coarse-graining is also somewhat vague. I recently tried to teach myself the programming language Python, which meant that I worked away for a couple of hours each day and didn't think much about it at other times. Sometimes, I would do nothing for a day or two. Despite the fact that the gaps were much longer than the actual activity, I had no problem saying that I was learning Python. Now, though, it has been months since I did anything, and despite my best intentions, it seems clear that I am no longer currently learning Python. It is not clear when I stopped, and I fully intend to

¹⁵Only roughly analogous, because plans can contain multiple independent subplans. Subplans can also act simultaneously as steps towards multiple independent goals. Formally, although plans can be modelled to an extent as a partially ordered set of steps, they certainly do not have to be total orders, as suggested by the metaphor of a chain (see also discussion in Jackendoff 2007). We disregard all of these complexities here, as the points to be made with respect to simple structures are transferrable to the more complex structures.

start again, but right now, it would be disingenuous to claim that I am learning Python.

In other words, intentional events, like physical events, tolerate discontinuities. The discontinuities can be individually longer than discontinuities in typical physical events, and can occupy a greater proportion of the event's runtime. Such discontinuities require us to admit a process of coarse-graining whereby possibly quite large breaks in an activity are subsumed within larger perceived patterns of activity. But not everything can be coarse-grained away like this: until I reopen my textbook or start trying to write little scripts, I'm not learning Python. The minimality condition is intended to ensure that, whatever I am doing right now, it does not count as part of an event of learning Python, because it does not form part of a chain of steps, characterized as in (22), for teaching oneself Python, and is too large a hiatus to ignore.

To summarize, there are several formal similarities between causes and effects in physical events, and actions and goals in intentional events. The differences that do exist between causal and intentional relations can largely be ascribed to differences between our perception of the two types of relation. We perceive physical causes as having almost inevitable, typically proximal effects, but we can see an action as part of a plan to reach a remote goal. As a result, intentional events often have more remote culminations than physical events, the process leading to the culmination can be more internally differentiated in intentional events (because different preparatory steps can be unified by a common intention), and it is more likely that the culmination is never reached.

3.3 Strategic events

Strategic events are similar to intentional events in that the coherence of the stuff constituting such an event is linked to an agent's will. However, the agent in an intentional event is a participant in the event, whereas in a strategic event, the intention may lie with someone who is not an event participant, or may not even be present when the event takes place. This means that strategic events do not necessarily share one of the characteristics of intentional events outlined above; accordingly, the set of strategic events as defined here is larger than the set of intentional events. Although there is little fixed terminology in this area, I will adopt Copley's (2008) term of **director** to describe the non-participant agent in a strategic event.

From a linguistic perspective, we cross a line when we consider directors. As will be discussed in Section 4 (see also Talmy 1988 and Copley & Harley 2014), verbs describe particular relations among event participants. The difference between physical events and intentional events can be seen partly in terms of the nature of those relations. Strategic events are different again, in that the director's intention is not necessarily reflected in direct participation in the event. The following are basic conditions on strategic events:

1. There must be a relationship of causation or enablement between the actions of agentive event participants and the plan or specified instructions of a director.
2. The plan or instructions must be minimal, as above.

Comparing these conditions to those on intentional events, two differences are apparent. First, there is a greater flexibility in argument roles, as the originator of the plan need not be identified with the agent of the event. Secondly, there is a concomitant increase in flexibility concerning how the agent's actions relate to the plan. Although the definitions are not quite parallel, I intend intentional events as a special case of strategic events where the agent is identified with the director. We could also define a complement set of 'strictly strategic events' which are not intentional events (at least one agent is disjoint from the director).

The characterization above could be sharpened. For example, several subcases of strategic events can be distinguished depending on precisely how the agent's actions are related to the director's plan. In one case, both director and agent are aware of the goal, and the agent is acting cooperatively, in accord with the director's plan. This occurs, for example, when a homeowner hires a technician to fix the washing machine: the technician (the agent) intends to fix the washing machine because the homeowner (the director) wants him to, and will take whatever steps he believes will enable him to fix the washing machine. A second case occurs when the director specifies instructions which can be followed multiple times (for example, writing a concerto). The orchestra (the agent) may intend to follow the composer's instructions to the slightest detail, and will do whatever is necessary to enable them to do so, but the composer may well have no idea that the performance is even taking place. Finally, in more Machiavellian examples, the director influences the behaviour of others in accordance with his own goals, without the agent being aware of those goals. The agent has some local goal (to borrow the director's car, say), but in acting towards that goal, inadvertently fulfils the director's plan (removing the evidence from the scene of the crime). I collapse these subcases here, because they all share the common characteristic of delimitation according to a director's plan, and can all be described with reference to that plan.

- (23)
- a. The homeowner got his washing machine fixed.
 - b. The composer had her symphony performed.
 - c. The criminal got rid of the evidence.

Strategic events can be more inclusive than intentional events, because the director's plan can span multiple intentional events, performed by multiple agents. For example, consider a homeowner who wants to sell his house. The homeowner may take the single step of employing an estate agent who will sell his house on his behalf. This is sufficient for the estate agent to assume the intention to sell the house and to act towards that goal (placing adverts, organizing viewings, etc.). In the meantime, the homeowner goes about his daily business and never thinks about selling the house: there is no way of observing the homeowner during these weeks and inferring an intention to sell the house. Months pass, nothing happens, and the homeowner comes to believe that the estate agent is not working hard enough. He tells the estate agent to stand down and employs a second estate agent instead. At this point, the first estate agent stops acting with a goal of selling the house, and the second estate agent starts doing so. The homeowner goes back to not thinking about the house. Weeks pass, and the house is sold. The homeowner says:

(24) I finally sold the house.¹⁶

What is the homeowner's involvement in this process? Mainly, he delegates: he tells other people to align their intentions with his. The estate agents' actions fulfil the homeowner's intention, a phenomenon sometimes known as **secondary agentivity**. The homeowner may well have had no involvement in the actual sale, but it is the homeowner's intention that delimits the culminated process: neither estate agent is involved with the house for the duration of the efforts to sell the house, and other agents, answerable to an estate agent and responsible for smaller tasks such as the preparation of adverts, were presumably involved for even shorter time periods. This is one reason for distinguishing between strategic events and intentional events: intentional events have a single agent, participating in the event; that agent's actions and goals delimit the event. Because directors can coordinate the actions of multiple agents, their goals can delimit larger events.

Such coordination is common among strictly strategic events: wherever there is co-operation and delegation, similar events can be construed from multiple perspectives. For example, various individuals and groups can claim to have fought in a war, but winning the war is a coordinated activity, only attributable to the entire army, or directors coordinating the action of that group of fighters: the head of state or the chief of defence staff. In fact, we see in such examples a kind of hierarchy of direction: directors influence the behaviour of agents, who may in turn be acting as directors constraining the behaviour of subsidiary agents, and so on. The homeowner talks to an estate agent, who asks someone to call the local newspaper and place an advert in the next edition. Such relationships between directors and agents can easily span several levels.

This multiplies quite brazenly the number of events corresponding to a particular portion of stuff. An actor who writes, directs, and performs a solo show simultaneously carries out several strategic, intentional, and physical events. This is quite a necessary feature of a model of event individuation, though: the actor might write quite brilliantly, but perform quite poorly, for example. As with Link's new ring from old gold and Davidson's sphere rotating quickly and heating slowly in Section 2, the writing must be distinguished from the performance, even if both are evidence by the same portion of stuff.

The strategic examples discussed so far all have the shape of culminated processes. As with physical and intentional events, though, strategic events do not have to have that shape. A nonculminating process is described in (25): dogs behave intentionally, but their intentions here are subsumed under those of the speaker, who allows, or causes, the dogs to exercise. As in previous sections, Vendler's diagnostic tests show that walking the dogs is a temporally extended process without an inherent culmination.

(25) a. I'm walking the dogs.

¹⁶Interestingly, the director can be portrayed as the subject of *sell* much more easily than certain, apparently similar cases. When the homeowner hires someone to fix the washing machine, it sounds disingenuous to say *I fixed my washing machine*, and if you order (freshly prepared) takeout, it sounds simply false to claim *I cooked dinner*. I suspect that this is a kind of blocking effect: a speaker is typically quite capable of cooking dinner, so we imagine the speaker is acting as agent. However, few people want to dive into the intricacies of selling one's own house, so the involvement of a specialist can be more or less taken for granted.

- b. I walked the dogs for/in an hour.

As for a strategic culmination, imagine a society in which a suitably powerful person can honour a visiting dignitary by arranging for several soldiers to fire their rifles simultaneously. As described in (26), this is necessarily strategic, as the queen has the role of a director rather than a direct event participant. Vendler's tests diagnose a culmination with no associated process.

- (26) a. #The queen is honouring the visiting dignitary.
b. The queen honoured the visiting dignitary at 1pm/in five minutes/for five minutes.

3.4 Analytical events

We have now seen three variants on a common set of event shapes formed from processes and culminations. A final possibility is that an event may not fit any of the above templates, but may nonetheless be identified by the same formal criteria discussed repeatedly above: individuation as diagnosed by anaphora, coupled with the distinctive event shapes described in Section 2.¹⁷ Any participants in such an event may or may not be aware that they form part of the event (in comparison, at least the agents in intentional and strategic events are normally aware of what type of event they form part of). In fact, there may not be any participants: the last ice age was an event according to the above criteria (it can be referred to anaphorically, and has the shape of a process of expansion of the ice caps, culminating when the ice caps receded beyond a certain threshold), but with no clear participants.

The distinction between climatic events on the scale of an ice age and climatic events on the scale of a rainstorm is a fuzzy one. Many large-scale events such as ice ages are purely physical, and many purely physical events such as rainstorms can only be recognized as events by some observer external to the event itself. However, I maintain the distinction here because rainstorms are accessible to direct sensory perception, whereas ice ages are not: an individual can look around and conclude that it is raining, but may not be able to look around and conclude that he is in the middle of an ice age. The individual may know that he is in the middle of an ice age, but such knowledge is inevitably the product of analytical inquiry, rather than directly related to that one individual's experience.

Events like ice ages are only apparent to analysts, typically divorced from the events themselves.¹⁸ From this perspective, ice ages have something in common with phe-

¹⁷The three non-physical types of event discussed here correspond to three ways in which a rationale clause can relate semantically to its host. In *The pirates sank the ship to collect the insurance*, the pirates are acting agentively, and the rationale clause explains why they did what they did. In *The ship sank to demonstrate the transience of all material things*, we attribute the relation to a director's goal. In *Sunflowers turn to face the sun*, there is no goal (except perhaps metaphorically) on the part of any agent or director. Rather, the explanation is an analyst's causal explanation of the sunflowers' different orientation at different times.

¹⁸Typically, but probably not necessarily divorced. Link (1997) discusses events like the French revolution or the scientific revolution, which appear to have the properties of analytical events on the present typology. However, participants in either event no doubt were aware that something revolutionary was going on, even if they could not have been aware of the nature and extent of the revolution.

nomena like population movements or the behaviour of stock markets: large-scale accumulations of agentive events with emergent properties. Surely no one individual has a desire to contribute to a phenomenon like rural exodus *per se*. Rather, multiple individuals or small groups move independently, in parallel, pursuing smaller-scale goals (jobs, excitement, whatever draws people to cities). The process of, say, rural exodus in France is only apparent to someone who can see the aggregate of those individual histories, just as patterns of change in glaciation are only apparent to someone who sees aggregate data from across the centuries.

We have given characterizations of physical, intentional, and strategic events above. Physical events stand in opposition to intentional events in terms of the basis for event individuation, and strategic events are less rigidly characterized than intentional events, allowing for distinction of a director from an agent. Analytical events, I claim, are even less rigidly characterized: any portion of stuff which fits into one of the event shapes described in Section 2 is an analytical event. As with strategic events, we can then define a set of ‘strictly analytical events’, which do not fall into any other category of events. That is, in principle, analytical events are quite unrestricted; it is up to individuals to make judgements about the set of *actual* analytical events. I take this to be a good thing: quite unruly portions of stuff can insightfully be seen as single events by the right analyst, as with the phenomena supporting the postulation of the Earth’s revolving around the sun, or global warming, as analytical events.¹⁹

Of the examples considered so far, rural exodus and ice ages are processes: people migrate from the country to the city for years, the earth cools for centuries. The migration out of Africa is a culminated process, or at least a set of culminated processes distinguished by their geographical endpoint (for example, people migrated from Africa to Europe in a given amount of time). A clearer example of an analytical culminated process is a typical extinction: although strictly speaking, the extinction of a species is as instantaneous as the death of the last member of that species, the progressive test indicates that we typically (but see below) construe extinction as a culminated process: *the giant panda, and Skolt Sami, are dying out*, we say. Similarly, perhaps the closest nonfiction approximations to analytical culminations are catastrophic global events like mass extinction of dinosaurs: populist presentations often give the impression that dinosaurs were almost instantaneously wiped out globally.²⁰ The mass extinction of dinosaurs, if ascribed these properties, has the form of an analytical culmination: it resists the progressive (*#After the meteor hit the earth, the dinosaurs were dying out*), rejects *for*-PPs (*The dinosaurs died out in/#for decades*), and is too complex to be treated as a purely physical event (*die out* and *become extinct* are collective predicates, which are sortally restricted to species or similar groupings. Although there is clearly an entailment relation between extinction of a species and deaths of the members of that species, a claim of extinction is really a universal claim to be made by an analyst: members of that species used to exist, and now there are no members of that species).²¹

¹⁹The Earth is a sphere spinning quickly and heating up slowly. Davidson would surely approve.

²⁰Real-world culminations are never actually instantaneous, of course: they are simply very quick relative to some contextual standard (compare culmination *hop* to process *climb*, for instance). Even a mass extinction that takes decades or centuries may be construed as a culmination from this perspective. This is another case where temporal coarse-graining is required for a satisfactory empirical account.

²¹It is interesting that *die out* describes both a culminated process in the case of the pandas, and a cul-

Analytical events, then, can have any of the same shapes as other classes of event. They are distinguished by the diminished role of any individual participants, and by the fact that they often occur at timescales which are only apparent *post facto* to an analyst who may not have observed the actual event.

3.5 Summary: Event shape

Across the preceding four subsections, we have identified a set of Vendlerian event shapes: processes, culminated processes, and culminations. Cross-cutting this, there is an essentially argument-structural distinction, although surely not in the classical sense:

1. **Physical events** are reducible to configurations of event participants, associated with different forces. The dynamics of the event are predictable from the orientation and magnitude of the forces associated with particular participants, as well as beliefs about action at a distance as effected by devices like switches.
2. **Intentional events** are grounded in the perceived intentions of the agent of the event: the event is perceived as an action performed by the agent with a given goal, which may be realized as a culmination. If the event is internally homogeneous (processes, or instantaneous changes, but not culminated processes), the event is often perceived as an end in itself. However, for bipartite culminated processes, the process is perceived as part of a plan to reach the goal; as a result, unlike physical events, the goal can be quite remote from the observed process, and may well not be reached.
3. **Strategic events** are like intentional events, except that the intention which defines the event shape is not necessarily the intention of any individual event participant, but rather of a director who may not participate in the event itself.
4. **Analytical events** are like physical events in that their shape is not necessarily defined by the intention of any one individual. They are unlike physical events in that they often cannot be directly perceived, and are instead often identified analytically.

Each of the above event types can have the temporal shape of a process, a culmination, or a culminated process. However, each of the above event types relates differently to the roles of participants in an event. Most importantly, physical and intentional events are individuated based on perceived properties of event participants alone, while strategic and analytical events are individuated otherwise: based on the intentions of a non-participant director in the case of strategic events, and on large-scale inferred patterns which frequently abstract away from individual participants in the case of strategic events. Even if the migration out of Africa must have been composed of individual physical and intentional events, it has only an indirect relationship to the

mination in the case of the dinosaurs, apparently based on world knowledge alone. This could be taken as evidence, as in Mittwoch's chapter, that the distinction between accomplishments and achievements is not as basic as the distinctions between events, processes, and states.

individuals that actually migrated out of Africa. If any one individual, or even any fifty individuals, had not taken part in the migration, the basic analytical fact of the migration would not have changed. Similarly, the migration may well have taken place over a very long period of time: there is no one ‘snapshot’ which would suffice to show that the migration was taking place. The identity of an analytical event is related to the systematicity revealed by generalization and abstraction. The relation of the individual to the analytical event is often indirect in a way which is familiar from any introductory statistical discussion of the relationship between populations and individuals: we see any finite group of participants as instances of the pattern identified by the analyst, rather than constitutive of the pattern.

Cross-cutting this distinction are two general schemata for delimiting events, according to dynamic configurations of individuals (physical or analytical events), or according to the intentions of an agent (intentional and strategic events). In the latter case, a single participant is responsible for the delimitation of the event, whereas this is not the case with physical or analytical events.²² If we hear that *the ball rolled down the hill* or *early humans migrated out of Africa*, the event is delimited not just by the moving theme, but also by the properties of the path-denoting PP. In contrast, with an intentional event description like *Susan carried Jeff into the sea*, we do not care whether Jeff intended to end up in the sea. He may or may not have been a willing participant; we just don’t know. Even verbs which entail things about the intentions of participants other than the agent, like *persuade*, make no commitment as to those intentions prior to the event of persuading. It is an implicature, rather than an entailment, that if *X persuaded Y to Z*, Y comes to intend to Z as a result of the entailment. There is no contradiction in an utterance like (27).

- (27) Jeff persuaded Susan to carry him into the sea, but she didn’t really need persuading: she was intending to do it anyway.

This suggests that our four event types can be recast in terms of two binary distinctions: events are **asymmetric** (delimited according to an individual’s intentions), or **symmetric** otherwise; and they are delimited based on properties of event participants (I will talk below of **directly delimited** events) or otherwise (**indirectly delimited** events). Table 1 summarizes this.

	Asymmetric	Symmetric
Directly delimited	Intentional events	Physical events
Indirectly delimited	Strategic events	Analytical events

Table 1: Binary distinctions among event types

In brief, this section has given a sketch of the degrees of independence of event shape from argument structure: at one extreme (physical events), event shape is entirely determined by the force-dynamic tendencies of event participants; at the other extreme

²²One basic consequence of this latter distinction is that there need not be any participants identified in physical event descriptions like *It rained* or analytical event descriptions like *the last ice age*. We will come back to the significance of this in the following section.

(strategic and analytical events), the relationship, although it undoubtedly exists, is very indirect. We now turn to the linguistic consequences of this independence.

4 Linguistic constraints on event composition

To this point, we have sketched a commonsense theory of events in perception and cognition. The core of the theory has been, on the one hand, a set of general composition relations among events, and on the other, a set of statements about how we tend to individuate events. The discussion has been largely based on event descriptions. The rationale for this is that if something happens, and we have a way of describing it, then that something is an event. Sometimes, the descriptions will be quite short on actual descriptive content (*chaos, that*, etc.); sometimes, as with most examples in this chapter, not.

Less trivially, the linguistic means we used to describe events have shifted as we progressed from physical events to analytical events. At the smaller end of the scale, we described events using clauses like (28a); at the other end, we used noun phrases like (28b).

- (28) a. The ball rolled down the hill.
b. The last ice age.

This is not a coincidence. There are ways of using noun phrases to refer to the event described in (28a), for example in (29a), but there are no obvious verbal equivalents of (28b), at least in English (see (29b)). Although sentences like (29c) are possible, they are hardly verbal event descriptions; rather, they use a verb like *happen* as a means of asserting the existence and temporal location of the event; the event description is still nominal: (29c) can be paraphrased as ‘there is an event *e*, *e* is located prior to speech time, and *e* is the last ice age’.

- (29) a. (i) The ball’s movement/progress/trajectory (down the hill)
(ii) The event we just witnessed
(iii) That
b. (i) *It/there ice-aged (most recently).
(ii) *Ice aged.
(iii) *The age was iced.
c. The last ice age happened/occurred/took place

It seems that we can find a definite noun phrase for any event we perceive (even if, in some cases, the NP in question is just *it* or *that*). However, there are several events which cannot be described by a verb, and it appears that those events tend to be analytical in nature.

We saw earlier that there are argument-structural correlates of the distinctions between different types of events. That suggests a further hypothesis: verbal event descriptions are more restricted than nominal event descriptions because verbal event descriptions, unlike nominal event descriptions like the above (**simple event nominals** in the terminology of Grimshaw 1990), must obey certain constraints on argument re-

alization.²³

Grimshaw shows that, on the one hand, simple event nominals clearly describe events, as shown by their co-occurrence with predicates which are semantically restricted to event-denoting arguments in (30).

- (30) a. The war happened.
b. The race took two hours.

On the other hand, simple event nominals do not make the internal structure of those events linguistically accessible in the same way in which complex event nominals or verbs do:²⁴ simple event nominals do not take *in/for*-PPs (31), and do not take any obligatory arguments (32).

- (31) #The race *in/for* two hours was exciting.
(32) a. John/*There raced ??(against Sam) yesterday.
b. John's/the race (against Sam) took place yesterday.

This means that we can use relatively unconstrained simple event nominals as a comparison class to identify specifically linguistic constraints on other classes of event description: certain events relate participants in a way which does not map well onto a verb's argument structure, given the constraints on the realization of verbal arguments. In such cases, we may nevertheless be able to describe the event using a simple event nominal. If so, we have found an event which is not well-described by a verb precisely because of its internal structure, as described in Section 3.

We focus on near-universally accepted statements concerning argument structure like the following:

1. The event participant hierarchy: syntactically realized event participants obey the ordering INITIATOR > UNDERGOER > RESULTEE, where > represents asymmetric c-command.²⁵
2. A single event participant cannot be described by multiple arguments of V.

These are intended as building blocks in a toy theory of event structure and argument realization. For further ingredients in a fuller theory, see chapters by Lohndal and Levin

²³Much of what I say about verbal argument structure is also true of nominal argument structure in those cases where a noun has arguments (Grimshaw's **complex event nominals**). I am only interested here in the more basic distinction, that nominal event descriptions frequently have no arguments, while verbal event descriptions very rarely do.

²⁴Because it is a fairly subtle task to distinguish different types of nominal event description, I only compare simple event nominals and verbs here. The interested reader can consult Grimshaw (1990), or Moltmann's chapter, for further details.

²⁵Although similar in organization, the event participant hierarchy is seen as distinct from the thematic hierarchies elaborated since Fillmore (1968), and strictly separated from the thematic hierarchy by Grimshaw (1990) in an analysis of the thematically similar *fear* and *frighten* classes of psych-predicates (see Jackendoff 1990 for related ideas). The specific terms used here are borrowed from Ramchand (2008) for concreteness, although there is still some variation among researchers in how many roles are recognized. The general idea is that participants in causally prior subevents are more prominent than participants in caused subevents. For example, an initiator participates in a causing event, which brings about some result involving the resultee, so initiators are more prominent than resultees.

& Rappaport Hovav, as well as an extensive primary literature going back through Hale & Keyser (1993) to work in Generative Semantics such as McCawley (1968). Staggering amounts of research was done in this area in the 1990s and 2000s (see Rosen 1999 for a summary of early results, and work such as Borer 2005 and Ramchand 2008 for more recent proposals). We cannot adequately summarize those results here, and will instead aim to show how constraints on argument realization affect the linguistic description of the different event types outlined above.

The restriction which emerges from these two constraints on verbal argument structure is that verbal event descriptions are usually asymmetric: verbs typically have arguments, and a single argument, the initiator, is more prominent than all the others. Putting alternations such as the passive aside, it is this argument which is realized as the subject (the syntactic argument which asymmetrically c-commands all others in the standard case).

Because of this, for example, there is no verb *schlime* such that (33a) and (33b) are synonymous. This presumably reflects the conceptual difficulty of finding a construal in which the mountain initiates the climbing event.²⁶

- (33) a. John climbed the mountain.
b. The mountain schlimed John.

Essentially the same is true of strategic events: (34a) is an appropriate description of a strategic event, but there is no verb *schefeat* to form (34b).

- (34) a. William the Conqueror defeated the English.
b. The English schefeatd William the Conqueror.

The asymmetric nature of intentional and strategic events feeds the asymmetry in verbal descriptions of such events, as in (33)–(34). In other words, the agent in an intentional event, and the director in a strategic event, are always construed as the initiators of those events, so there is limited scope for variation in mapping of event participants to syntactic positions.²⁷

In contrast, we have so far described physical and analytical events as symmetrical. In particular, a physical event is just a dynamic interaction among a set of participants, while an analytical event is delimited according to the insight of the analyst. This has to be reconciled with the inherent asymmetries among verbal arguments. There are a few ways to do this: some physical and analytical events can be construed as having an

²⁶Pairs of verbs with apparent mirror-image argument structures do exist, with the best known being experiencer alternations such as *John likes pears* and *Pears please John*. In many cases, at least one member of such pairs is stative, in which case notions such as ‘initiator’ are clearly inappropriate. In other cases, such as the *sprayload* alternation (*John loaded the truck with hay* vs. *John loaded the hay onto the truck*), the symmetry reflects two salient ways of construing the event, as bounded by the theme or the goal.

²⁷There is some scope for variation in strategic events, mainly concerning the phenomenon of secondary agentivity discussed above. The secondary agent can be omitted entirely, as in (34) (William the Conqueror didn’t defeat the English single-handedly; rather, he instructed his army to act in a way which led to the defeat of the English). It can also be included with verbs such as *make* or *have* (*William the Conqueror had his men attack the English*, but not *William the Conqueror had his men defeat the English*). A range of subtle consequences follow (for instance, compare *I finally sold the house* with *I finally had the estate agent sell the house* — the former suggests a period of waiting for a buyer; the latter a period of dithering about whether to set the ball rolling).

initiator — if not an agent or a director, a cause, as in (35).

- (35) a. The wind blew the tiles off the roof.
b. Gavriilo Princip [the man who shot Fanz Ferdinand] started World War I.

If there is no obvious initiator, the most prominent argument is typically the theme, or in Talmy's (2000) terms, the figure moving relative to a ground, as in (36).

- (36) a. The tiles fell from the roof.
b. Several small groups of humans migrated out of Africa.²⁸

Finally, restrictions arising from the mapping of event participants to verbal arguments can sometimes be overcome by choosing a verb with a simple argument structure: a 1- or 0-place predicate. The single argument of a 1-place predicate can refer to a group or mass, without differentiating the roles of subparts of that group or mass, as in (37).

- (37) a. (i) The gas dispersed.
(ii) The volunteers spread out across the field.
b. Gondwanaland broke up.

As for 0-place predicates, the best-known are **weather verbs**. There is nothing which forces the events described by weather verbs to be described by 0-place predicates: (38a) and (38b) are equally valid descriptions of the same state of affairs.

- (38) a. It rained.
b. Rain fell from the sky.

In (38b), the rain is construed as a theme, or figure, moving away from the sky, a source. In (38a), there is no such asymmetry, as there are no arguments.

In English, 0-place predicates are quite limited, and very often unavailable to describe particular events, and 1-place predicates can be quite strained. If no such simple argument structure is available, but the complexity of the relations among event participants does not allow for straightforward identification of a single initiator, it will often not be possible to describe an event using a verb and its arguments.

The following is an example. A simple car crash, with two cars, can be described as follows, either verbally with one car identified as the theme/initiator and the other as the goal, as in (39a); verbally with no indication of asymmetry between the roles of the two cars (39b); or nominally (39c).

- (39) a. The red car crashed into the blue car.
b. Two cars collided.
c. The car crash

A more complex car crash, involving 27 cars in various ugly configurations, is more

²⁸A single person, or a single family, cannot migrate. Only largish populations can. But a migration can be made up of multiple small-scale movements of individuals or families, with internal organization visible to anyone other than an analyst. I have no idea if the migration out of Africa actually had such properties; all that matters is that we *could* describe an event with such properties, and it would have to be an analytical event with *several small groups of humans* as its theme.

likely to be described using a nominal: (40ai) is likely to be interpreted as false (it is most naturally interpreted as describing 26 collisions between the red car and a different other car); (40a_{ii}) is better, but still implies a weak reciprocal reading, where 27 cars collided with each other. (40a_{iii}) is perhaps more accurate, but hopelessly circumlocutious — this is a multi-clausal event description, with no single clause accurately portraying the relations among participants. (40b) is looser: if, say, two cars crashed, another went into the back of them, a fourth swerved to avoid them and hit a tree, the distraction caused a fifth to lose control on the far side of the road, and so on, (40b) would be an adequate description, but (40a_{ii}) would not, because the 27 cars did not collide with each other.

- (40) a. (i) #The red car crashed into 26 other cars.
 (ii) 27 cars collided.
 (iii) Several groups of cars collided; 27 cars were involved overall.
 b. A 27-car pile-up.

A second example of the same phenomenon, this time analytical rather than physical, is in (41). Imagine a war involving five countries, A, B, C, D and E, where the following propositions are all true.²⁹

- (41) a. A and B are fighting as allies against C.
 b. A is fighting alone against D.
 c. B is fighting alone against E.

We can describe this situation accurately using a conjunction of the three propositions above, but how else can we describe it? (42a_i) is almost accurate, but oversimplistic, reducing a complex set of interactions to an antagonistic relationship between two ‘teams’. (42a_{ii}) seems unwarranted, implying that A–E all fought each other. The nominal (42b), because it avoids any argument-structural commitments, seems more accurate.

- (42) a. (i) A and B fought C, D, and E.
 (ii) A, B, C, D, and E fought.
 b. The war

In both of these cases, as relations between a large set of participants become more complex, it becomes increasingly hard to shoehorn the event description into a verbal argument structure. As a consequence, it becomes increasingly natural to use an argument-free, nominal event description.

The moral of the story, following Rosen (1999: 8), is that ‘verbs at least in part mean what the syntax allows them to mean.’ Rosen is concerned with linguistic constraints on event descriptions. In fact, those constraints are language-particular. Languages other than English have broader classes of 0-place verbal predicates than the weather verbs described above, whether derived or basic. Perhaps the best known of

²⁹Wikipedia currently lists 114 countries which were implicated in World War II in a variety of ways. Although I imagine that it is possible to see World War II as a two-sided fight between allies and axis, it is clearly also possible to construe it in a way whose complexity far outstrips (41).

these are **impersonal passives** in German and other languages, such as (43).

- (43) Es wurde getanzt
It was danced
'There was dancing.'

A second example, from Serbo-Croatian, uses a reflexive morpheme instead.³⁰

- (44) Ratovalo se godinama.
war.PTCP REFL years.INST
'There was a war for years.'

The Serbo-Croatian example (though not necessarily the German impersonal passive) implies the same kind of complexity, or abundance of activity, which could most felicitously be described by the nominal examples above. Notably, idiomatic English translations of (43) and (44) always resort to a nominal event description. It is the absence in general of verbal 0-place predicates which forces English speakers to use nominal event descriptions in such cases. Verbs, to an extent, mean what the syntax allows them to mean, but the syntax does less to restrict the meaning of simple event nominals. Asymmetries among participants in events with initiators (whether causers, agents, or directors), or amenable to figure-ground articulation, tend to be well-described by verbs; other events, without such an articulation, are often better described by such nominals.

5 Summary

The Davidsonian parallel between individuals and events leads us to expect that events can be individuated at a variety of levels of granularity, just as individuals can. This appears to be true. However, just as with individuals, there are a range of perceptual constraints on event individuation. We identified four different types of events (physical, intentional, strategic, and analytical), corresponding roughly but not precisely to four different granularities, and saw that, despite the distinct individual properties of these different event types, each shares a basic Vendlerian compositional template, consisting maximally of a process leading to a culmination, or nonmaximally of either a process or a culmination in isolation.

A major distinction was drawn between intentional and strategic events, in which there is a single privileged individual whose intention delimits the event, and physical and analytical events, where there need not be any such individual. However, we saw that effects relating to verbal argument realization may impose such an asymmetry on arguments even when there is no such asymmetry inherent to the event: in languages like English, in the vast majority of cases there must be a syntactically most prominent argument corresponding to a semantically most prominent argument. In turn, this entails that many very complex physical and analytical events are most naturally described in English using nominal, rather than verbal, event descriptions: simple event

³⁰Thanks to Berit Gehrke, Dejan Milacic, Ana Werkmann, and Vesela Simeonova for discussion of this and related constructions.

nominals do not need arguments like verbal event descriptions typically do. This means that there are systematic linguistic constraints on a class of event descriptions in English, over and above any perceptual restrictions on the shape of events.

In contrast, other languages have means of circumventing those linguistic constraints, by more productive use of 0-place verbal predicates. We saw two such examples above: the impersonal passive in German and other languages, and a particular reflexive construction found in Serbo-Croatian and elsewhere. In such languages, verbs can be used to describe events which do not readily lend themselves to verbal event descriptions in languages like English.

In sum, we have seen a basic logical relation, of event composition, constrained by perceptual factors relating to the individuation of events, and further constrained by linguistic factors, both universal factors relating to argument structure such as the mapping between event participant roles and syntactic positions, and language-particular factors such as the particular configurations of verbal arguments available in a given language.

References

- Bach, E. (1986). The algebra of events. *Linguistics and Philosophy*, 9, 5–16.
- Bar-El, L., Davis, H., & Matthewson, L. (2005). On non-culminating accomplishments. In L. Bateman & C. Ussery (Eds.), *Proceedings of the Thirty-fifth Annual Meeting of the North East Linguistic Society, Vol.I* (pp. 87–102). Amherst, MA: GLSA.
- Borer, H. (2005). *Structuring Sense. Volume 2: The Normal Course of Events*. Oxford: Oxford University Press.
- Chomsky, N. (2000). *New Horizons in the Study of Language and Mind*. Cambridge: Cambridge University Press.
- Copley, B. (2008). The plan's the thing: Deconstructing futurate meanings. *Linguistic Inquiry*, 39, 261–274.
- Copley, B. & Harley, H. (2014). A force-theoretic framework for event structure. Ms., Université Paris 8 and University of Arizona.
- Copley, B. & Wolff, P. (2014). Theories of causation should inform linguistic theory and vice versa. In B. Copley & F. Martin (Eds.), *Causation in Grammatical Structures*. Oxford: Oxford University Press.
- Davidson, D. (1967). The logical form of action sentences. In N. Rescher (Ed.), *The Logic of Decision and Action* (pp. 81–95). Pittsburgh, PA: University of Pittsburgh Press.
- Davidson, D. (1969). The individuation of events. In N. Rescher (Ed.), *Essays in Honor of Carl G. Hempel* (pp. 216–234). Dordrecht: Reidel.

- Dowty, D. (1979). *Word Meaning and Montague Grammar: The Semantics of Verbs and Times in Generative Semantics and Montague's PTQ*. Dordrecht: Reidel.
- Fillmore, C. (1968). The case for case. In E. Bach & R. Harms (Eds.), *Universals in Linguistic Theory* (pp. 1–90). London: Holt, Rinehart and Winston.
- Fodor, J. A. (1970). Three reasons for not deriving “kill” from “cause to die”. *Linguistic Inquiry*, 1, 429–438.
- Grimshaw, J. (1990). *Argument Structure*. Cambridge, MA: MIT Press.
- Hale, K. & Keyser, S. J. (1993). On argument structure and the lexical expression of syntactic relations. In K. Hale & S. J. Keyser (Eds.), *The View from Building 20: Essays in Honor of Sylvain Bromberger* (pp. 53–109). Cambridge, MA: MIT Press.
- Higginbotham, J. (1985). On semantics. *Linguistic Inquiry*, 16, 547–593.
- Jackendoff, R. (1972). *Semantic Interpretation in Generative Grammar*. Cambridge, MA: MIT Press.
- Jackendoff, R. (1990). *Semantic Structures*. Cambridge, MA: MIT Press.
- Jackendoff, R. (2007). *Language, Consciousness, Culture: Essays on Mental Structure*. Cambridge, MA: MIT Press.
- Kamp, J. A. W. (1975). Two theories about adjectives. In E. Keenan (Ed.), *Formal Semantics of Natural Language* (pp. 123–155). Cambridge: Cambridge University Press.
- Krifka, M. (1989). Nominal reference, temporal constitution, and quantification in event semantics. In R. Bartsch, J. van Bentham, & P. van Emde Boas (Eds.), *Semantics and Contextual Expression* (pp. 75–115). Dordrecht: Foris.
- Lewis, D. (1973). Causation. *Journal of Philosophy*, 70, 556–567.
- Link, G. (1983). The logical analysis of plural and mass terms: A lattice-theoretical approach. In R. Bäuerle, C. Schwarze, & A. von Stechow (Eds.), *Meaning, Use and the Interpretation of Language* (pp. 303–323). Berlin: Walter de Gruyter.
- Link, G. (1997). *Algebraic Semantics in Language and Philosophy*. Stanford, CA: CSLI.
- McCawley, J. (1968). Lexical insertion in a transformational grammar without deep structure. In B. Darden, C.-J. Bailey, & A. Davison (Eds.), *Papers from the Fourth Regional Meeting, Chicago Linguistic Society, April 19-20, 1968* (pp. 71–80). Department of Linguistics, University of Chicago. Reprinted as ch.8 of McCawley (1973).
- McCawley, J. (1973). *Grammar and Meaning: Papers on Syntactic and Semantic Topics*. New York: Academic Press.

- Moens, M. & Steedman, M. (1988). Temporal ontology and temporal reference. *Computational Linguistics*, 14, 15–28.
- Pietroski, P. (2000). *Causing Actions*. Oxford: Oxford University Press.
- Pustejovsky, J. (1991). The syntax of event structure. *Cognition*, 41, 47–81.
- Ramchand, G. (2008). *Verb Meaning and the Lexicon: A First Phase Syntax*. Cambridge: Cambridge University Press.
- Reichenbach, H. (1947). *Elements of Symbolic Logic*. New York: Macmillan.
- Rosen, S. T. (1999). The syntactic representation of linguistic events. *Glott International*, 4, 3–11.
- Rozenblit, L. & Keil, F. (2002). The misunderstood limits of folk science: An illusion of explanatory depth. *Cognitive Science*, 26, 521–562.
- Talmy, L. (1988). Force dynamics in language and cognition. *Cognitive Science*, 12, 49–100.
- Talmy, L. (2000). *Toward a Cognitive Semantics*. Cambridge, MA: MIT Press.
- Travis, L. (2000). Event structure in syntax. In C. Tenny & J. Pustejovsky (Eds.), *Events as Grammatical Objects: The Converging Perspectives of Syntax and Lexical Semantics* (pp. 145–185). Stanford, CA: CSLI.
- Truswell, R. (2011). *Events, Phrases, and Questions*. Oxford: Oxford University Press.
- Vendler, Z. (1957). Verbs and times. *Philosophical Review*, 66, 143–160.
- Wolff, P. (2003). Direct causation in the linguistic coding and individuation of causal events. *Cognition*, 88, 1–48.
- Zacks, J. & Tversky, B. (2001). Event structure in perception and conception. *Psychological Bulletin*, 127, 3–21.