Chapter 4
The Windowing of Attention in Language

1 INTRODUCTION

Using the perspectives and methods of cognitive semantics, this study sets forth the system with which languages can place a portion of a coherent referent situation into the foreground of attention by the explicit mention of that portion, while placing the remainder of that situation into the background of attention by omitting mention of it.[1] Terminologically, the cognitive process at work here is called the *windowing of attention*, the coherent referent situation with respect to which the windowing must take place is an *event frame*, the portions that are foregrounded by inclusion are *windowed*, and the portions that are backgrounded by exclusion are *gapped*. In engaging this subject, the present chapter treats a number of phenomena. It examines five generic types of event frame—a path, a causal chain, a cycle, a participant interaction, and an interrelationship—and it considers the cognitive factors that constitute and bound such event frames. It examines the properties of the windowing process, including its capacity for embedding or for multiple co-occurrence, as well as the functions that this process may serve within the overall organization of cognition. It investigates a number of concomitant cognitive phenomena including the nature of attention, foregrounding and backgrounding, conceptual alternativity, cognitive splicing, goal-schema constancy, causal transparency and the sense of causal immediacy versus distance, conceptual contrast frames, and the systematic relationship of factuality to affect states and explanation types. It speculates on correlations between the windowing structure in language and comparable structuring in perception and motor control, including the ways these are manifested in the experiments of virtual reality. And it observes the commonality of windowing

structure in spoken language and in the sign-language systems spontaneously developed by certain deaf children, a commonality that testifies to the fundamental character of the cognitive structure presented here.

The windowing of attention is just one fragment of the much vaster cognitive system constituting the conceptual structuring of language. In hierarchical terms, the windowing of attention—along with level of attention, center of attention, scope of attention, and network of attention—is part of the larger cognitive structural category in language that can be termed the **distribution of attention**. This category can be considered a **schematic system**. In turn, this system—along with other schematic systems such as configurational structure, location of perspective point, force dynamics, and cognitive state—together constitute the fundamental delineation of conceptual structuring in language.[2]

2 THE NATURE OF ATTENTIONAL WINDOWING

Linguistic forms can direct the distribution of one's attention over a referent scene in a certain type of pattern, the placement of one or more **windows** of greatest attention over the scene, in a process that can be termed the **windowing** of attention. In this process, one or more portions of a referent scene—where each portion has internal continuity but is discontinuous from any other selected portion—will be placed in the foreground of attention while the remainder of the scene is backgrounded. The most fundamental formal linguistic device that mediates this cognitive process is the inclusion in a sentence of explicit material referring to the portion or portions of the total scene to be foregrounded, and the omission of material that would refer to the remainder of the scene intended for backgrounding. This device is the only one to be treated here and the one for which the term "windowing" will be reserved.[3] Although only a certain portion or portions of the referent scene are explicitly specified when thus windowed, it is understood as part of the nature of the windowing process that—given the appropriate context—the addressee will be able to infer the remainder of the scene. Generally, the same referent scene can be windowed in any of several different ways—that is, different patterns of selected windows can be placed over the scene. This latitude is another manifestation of the fundamental linguistic property of **conceptual alternativity** described in chapter I-3, and it will be exemplified in all the categories of windowing treated below.

To introduce some of the terminology employed below, a referent scene that is sequential in nature or that has been sequentialized conceptually can have a window of strongest attention placed over its beginning, middle, or end portion—or, as will be said here, may have **initial**, **medial**, or **final windowing**. On the other hand, such a scene can have a particular portion without a window on it, backgrounded by the lack of sentence constituents referring to it, and accordingly here be said to have **initial**, **medial**, or **final gapping**.

2.1 The Event Frame

To be viable, the concept of windowing requires a basis on which to distinguish between two kinds of material missing from a sentence: a kind whose referent would indeed be understood as belonging to the represented scene, and another kind whose referent would be felt as peripheral or incidental. Serving such a function, something like the following consideration is needed: Arising from whatever causes, whether in part innately universal ones or in part linguistically or culturally specific ones, language users apparently tend to conceive certain elements and their interrelations as belonging together as the central identifying core of a particular event or event type. Other elements, ones that on other grounds might have seemed to share an equally intimate involvement in the event, are instead conceptualized as peripheral or incidental.

A set of conceptual elements and interrelationships that in this way are evoked together or co-evoke each other can be said to lie within or to constitute an **event frame**, while the elements that are conceived of as incidental—whether evoked weakly or not at all—lie outside the event frame. Prominent examples of event frames include the so-conceived entirety of an object's path, that of a causal chain, and that of an interchange of entities (including an exchange of possessions, as in Fillmore's "commercial event"). Typically **not** included within an event frame, however, are, for example, the day of the week on which an event occurred, the geographic locale in which the event occurred, the ambient temperature of the space in which the event occurred, or the state of health of a participant in the event—even though such factors can be fully or even necessarily as much involved in an event as the factors that do get treated as part of the event.

This notion of an event frame is very close to Fillmore's (e.g., 1982) concept of a frame or scene when applied to an event, but there appear to

be several differences of emphasis or of conceptual basis. First, where Fillmore emphasizes mainly the co-presence of certain interrelated conceptual elements, our notion of an event frame is intended to stress as well the exclusion of other conceptual elements from the privileged core. Second, a frame for Fillmore seems to represent a concept or phenomenon that may be specific to a particular language or set of languages and that may be determined only within a particular sociocultural context. Our event frame, however, is generally understood as a more generic category that is quite likely universal across languages, that at least in part corresponds to the structuring in other cognitive systems such as visual perception, and that may well be innately determined. Such a generic status is thus assumed for the event frame types treated below—hence, for the path, the causal-chain, the cycle, the participant-interaction, and the interrelationship event frames. Fillmore's commercial scene, which involves an exchange of possessions, thus might under further investigation come to be seen as constituting only one particular form of a generic type of event frame that consists of an interchange of entities and that is demarcated in accordance with some general factor such as reciprocity or symmetry.

It remains to be determined whether there are relatively general conceptual factors or cognitive principles that govern which clusterings of conceptual material are felt to constitute coherent event frames of particular types. To this end, the analysis below successively posits a number of factors that may contribute to the demarcation of different types of event frames. To preview them, we can at this point indicate the factors that will be proposed. First, in an event frame of motion, the so-conceived entirety of an object's path may be demarcated by periods of stationariness that temporally bound the period of motion, or by "path singularities"—that is, abrupt qualitative shifts in the path direction or in the surrounding medium. It can also be demarcated by a normative scope of perception or by the analysis of a path complex into an embedded structure of one path nested within another. It can further be demarcated by the spatial coincidence of two points of a path when this path is closed or by two bilaterally symmetric elements that represent corresponding points in a reflection about a central axis. Second, in an event frame of agentive causation, the so-conceived entirety of a causal chain may be demarcated by the initiating volitional act of an agent and by the final goal that the agent intends as a result of this act, where this act and goal mark the beginning and the end of the agent's scope of intention. Third, in a cyclic

event frame, the so-conceived entirety of a cycle is generally demarcated by two temporal points that bear the same phase relation to two congruent stretches of occurrence, where these two points are conceptualized as part of a "home" phase. Fourth, in a participant-interaction event frame, the occurrence of two distinct punctual events extrinsic to a certain circumstance that extends through time can mark out a portion of that circumstance and establish that portion conceptually as an event frame. Finally, an interrelationship event frame can be demarcated by the coentailment of its component elements, by the complementary relationship of its component elements where there are only two of these, or by the capacity of its component elements to function as alternative conceptualizations juxtaposed within a single comparison frame.

Given such relatively general factors that help determine portions of conceptual material that will be felt to constitute unitary coherent event frames, is there any still more general cognitive principle that runs in common through these factors or that characterizes the ways in which they function to demarcate the event frames? Such a principle seems to be that the organizing factors function to establish what is conceptualized as a **boundary** around the portion of conceptual material constituting the event frame. This boundary separates that portion from other conceptual material. As might be expected, such a boundary—and, hence, an event frame in general—exhibits various prototype effects such as those described by Rosch (1978) and Lakoff (1987). For example, the boundary might not be a sharp line but a gradient zone, and its particular scope and contour—hence, the particular quantity and portions of material that it encloses—might vary in accordance with the specific context or type of context. Nevertheless, some sense of boundary appears to be present across the relevant cases and to govern certain associated characteristics. First among such characteristics is the definitional one that the material enclosed within the boundary is felt to constitute a unitary coherent conceptual entity distinct from the material outside the boundary. Second, there seems to be some sense of **connectivity** throughout the material enclosed within the boundary and, contrariwise, some sense of **discontinuity** or **disjuncture** across the boundary between the enclosed and the external material. Such conceptualized connectivity and disjuncture might be spatial, temporal, or causal, for example, or might further pertain to information or to perception. Heuristically, thus, it might be spatial, where within the boundary there is access from any one point to any other point without blockage but where the boundary acts as a barrier to

movement from points within to points outside; or temporal, where the material within the boundary extends through a continuous period of time without gaps, but where this material is conceptually excerpted from the surrounding flow of time; or causal, where effects can freely propagate within the boundary but not beyond it. Further, it might be informational, where information or knowledge about particular phenomena held at one point is available at other points within the boundary but not to or from points outside the boundary; or perceptual, where there is perception of all points within the boundary from any point also within it, but not perception of points or from points outside the boundary. Third, the various portions of the material within the boundary are felt to be **co-relevant** to each other, whereas the material outside the boundary is not relevant to that within. This sense of relevance may be able to override the different forms of connectivity—for example, in a commercial scene, bringing together just those participants during periods of action that comprise the exchange of goods and money, excerpted from their spatio-temporal surround.

2.2 Event Frames and Complement Structure

Undoubtedly, something of this sense for what lies inside and what lies outside a conceptual event frame has motivated syntacticians, beyond purely formal evidence, to distinguish between "complements" and "adjuncts," respectively. But the explicit positing of the event frame as a linguistic entity permits an elaboration of complement structure theory that might not otherwise be possible. Current theory recognizes two types of complements to a lexical item that represent its semantic arguments: an obligatory complement, which must accompany the lexical item, and an optional complement that may or may not do so. To these two types of complement we could add a third type, a **blocked complement**, to be adduced where a predicate arguably has an associated argument that cannot be expressed in construction with the particular lexical item.[4] In our terms, such an argument would be felt to be an intrinsic part of a particular conceptually coherent event frame, an argument that might be expressed in construction with some other lexical item that refers to this event frame but one that cannot be expressed in construction with the lexical item in question. These relationships are illustrated in the accompanying diagram. Here, the large rectangle represents a particular event frame. Inside the rectangle, the solid-line square represents an obligatory complement, the dotted-line square an optional complement, and the

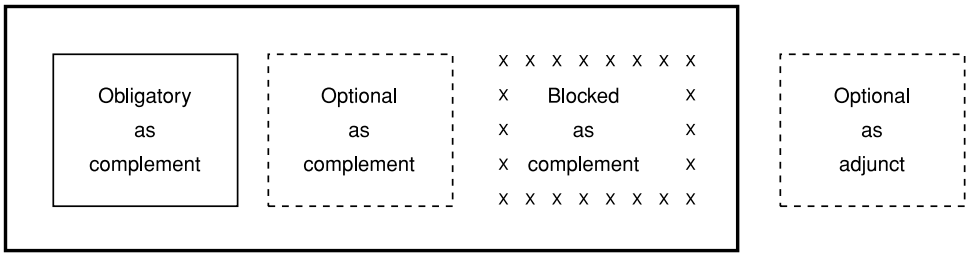

Xed-line square a blocked complement. Outside the rectangle, the dotted-line square represents an element that can be optionally expressed as an adjunct.

All three types of complements, as well as adjuncts, can be illustrated for the verb *spend*, which invokes the Fillmorean commercial scene that includes as arguments a seller, a buyer, goods, and money. Thus, the sentence in (1) shows in italics the verb's obligatory complements, the buyer as subject and the money as object; shows in parentheses an optional complement, the goods as prepositional object; shows in brackets a blocked complement, the seller in an attempted oblique constituent; and shows in braces two optional adjuncts referring to locale and day of week.

(1) *I* spent *$50* (for/on this book) [*from/by/to/for/ . . . the clerk] {at that store} {last Friday}.

To illustrate the potential extent of complement blockage, we can consider what may be posited as the event frame for force dynamics (chapter I-7), which necessarily includes an Antagonist and an Agonist, the two main entities engaged or potentially engaged in an opposing force interaction. The verb *permit* refers to one such force-dynamic event frame and requires complements referring to both of the force entities—the Antagonist as subject and the Agonist as direct object—as seen in (2a). But, in their force-dynamic usage, the English modals regularly block expression of the Antagonist, requiring solely the Agonist as subject. This is seen in (2b) for the modal *may*, which refers to the same type of force-dynamic event frame as *permit*.

(2) a. I permit you to go to the park.
b. You may go to the park (*by/from/ . . . me).

The force-dynamic event frame further serves as the event frame for a more extended form of complement blocking. In one of its constructions, the verb *require*, like *permit*, refers to a particular type of force-dynamic event frame and requires complements referring to both of the force entities—again, the Antagonist as subject and the Agonist as direct object—as seen in (3a). The modal *must*, which refers to the same type of force-dynamic event frame as *require*, can participate in the usual modal construction seen just preceding—blocking expression of the Antagonist and requiring expression of the Agonist as subject—as in (3b). But most modals can participate in a still further construction, characterized in chapter I-7 as involving "Agonist demotion," which exhibits an extreme case of complement blockage. It blocks complements referring both to the Antagonist and to the Agonist—that is, it blocks the entire substantive core of the force-dynamic event frame, as illustrated for *must* in (3c).

(3) a. I require that you let the cookies stay in the jar.
b. You must let the cookies stay in the jar (*by/from . . . me).
c. The cookies must stay in the jar (*by/from/ . . . you, *by/from/ . . . me).

To argue out some of the theoretical issues, we note that one view concerning the complement structure of a lexical item holds that this structure—its requirements, allowances, and exclusions—is exactly consonant with the semantic structure of the lexical item, if that semantic structure is assessed adequately. According to this view, there can be no such thing as a blocked complement, since the semantics of the lexical item could have no component that lacks a corresponding syntactic component. For example, a proponent of this view might argue, on the basis of sentences like *I spent $50 and 100 hours of my time on that ham radio kit*, that the verb *spend* does not really involve the notion of a seller but rather refers to a frame more generic and smaller than a full commercial scene, one that contains an agent expending possessed resources in order to attain a desired goal, so that it is no surprise that the verb's complement structure would exclude reference to a seller. But a closer inspection reveals that when *spend* is used to refer to the outlay of money, as against other kinds of resources, that money must in fact go to a seller engaged with the agent in a standard commercial transaction. For example, in the preceding illustrative sentence, the verb *spend* could not have been used if the $50 had not been given to a seller in exchange for the kit but rather, say, was used as paper ignited to melt solder. Further, the verb *spend*

cannot even be used, say, in **I spent $50 on their ritual mask*, if in exchange for the mask the money was given to native tribes people who wanted it for its aesthetic or curiosity value rather than as part of our standard commercial transaction (as observed by Kean Kaufmann). Thus, when applied to money, the verb *spend* still requires the participation of a knowing seller, even though this participant cannot be expressed by a complement of the verb.[5] With evidence such as this, we would therefore maintain the contrary view that while there is generally much correspondence in language between the system of formal syntactic structure and the system of semantic structure, the two systems nevertheless each have at least in part their own independent structuring patterns and principles. One possibility is that the semantic structure that pertains to event frames derives from, or is simply comprised of, the structure of our conceptual organization, a structure that perhaps is in part innate and universal, while the syntactic complement structure of particular lexical forms in a language can either directly reflect that semantic structure or can partially deviate from it in a kind of frozen grammaticization.

We now examine in sequence several different types of event frame for the forms of attentional windowing that they support.

3 PATH WINDOWING

The first type of event frame considered is that of the so-conceived entirety of a path of motion, here termed a **path event frame**, with respect to which the windowing process can be termed **path windowing**. This windowing process can be treated with respect to three different categories of paths, to be discussed in turn—open paths, closed paths, and fictive paths—all of which can exhibit a cognitive process called cognitive splicing.

3.1 Open Path

An *open path* here will refer to a path that is described by an object physically in motion in the course of a period of time, that is conceptualized as an entire unity thus having a beginning and an end, and whose beginning point and ending point are at different locations in space. To illustrate open-path windowing, the example in (4) pertains to a single particular instantiation of the open-path type but with various patterns of windowing and gapping imposed on it. Thus, (4a) presents the event with maximal windowing over the whole of the conceptually complete

path, while (4b) presents three forms of gapping over one portion of the path and (4c) presents three forms of windowing over one portion of the path.[6] It is understood here that the gapped portions are attentionally backgrounded relative to the foregrounded windowed portions but that, given sufficient context, a hearer would reconstruct each of the partially gapped paths in (4b) and (4c) into the same conceptualization of a complete path.

(4) The crate that was in the aircraft's cargo bay fell—
   a. *With maximal windowing over the whole of the so-conceived entire path*
      —out of the plane through the air into the ocean.
   b. *With gapping over one portion of the path*
      i. Medial gapping = initial + final windowing
         —out of the plane into the ocean.
      ii. Initial gapping = medial + final windowing
         —through the air into the ocean.
      iii. Final gapping = initial + medial windowing
         —out of the airplane through the air.
   c. *With windowing over one portion of the path*
      i. Initial windowing = medial + final gapping
         —out of the airplane.
      ii. Medial windowing = initial + final gapping
         —through the air.
      iii. Final windowing = initial + medial gapping
         —into the ocean.

We can suggest factors that may play a role in the putative cognitive processes by which an open path becomes conceptualized as an event frame—that is, as a unitary event bounded off from surrounding material of space, time, or other qualitative dimensions. One such factor might be the **scope of perception** that one might imagine as being normatively or canonically available at the referent scene. For instance, in generating or in interpreting the sentences of the preceding example, speakers or hearers might imagistically locate a viewpoint for themselves at a canonic position between the aircraft and the ocean whence the crate's path from the plane to the ocean would fall within the available scope of perception and thereby be treated as a unity. Since from such a viewpoint the crate would not be visible either in its prior motion while in the cargo bay nor in its subsequent motion through the water to the ocean floor, such additional

surrounding paths of motion would be excluded from the event frame in the operation of the putative scope-of-perception factor.

Another possible cognitive factor would function to frame together a sequence of phenomena that was assessed as having one qualitative character and separate that off from otherwise adjoining sequences assessed as being qualitatively different. One form of this factor, involving stationary boundary periods, would treat a period of stationariness as qualitatively distinct from a period of motion, so that the attribute of unitary entityhood could be cognitively ascribed to a period of continuous motion that was bounded by two stationary periods. Although perhaps otherwise frequent, this form of the factor would not play a role in the preceding aircraft example since the crate is in fact in motion both before and after the path represented in the sentences.

However, the factor of qualitative difference may have other forms, ones that would apply to the example. One such form might be the treatment of a conceivedly abrupt shift in path direction as marking the distinction between two qualitatively distinct paths and the conceivedly sharp-angled point of the shift as marking the boundary between the two paths. Such a **path singularity** form of the factor could be at work in the aircraft example to mark the beginning point of the crate's fall. Another form of the qualitative factor might address any abrupt shift in the character of the space surrounding a path—for example, change in the ambient medium. This form of the factor could then apply in the example to the passage of the crate's path from air to water, treating that as the end point of the preceding portion of motion.

When they have the requisite character, certain qualitative shifts in a path complex may lead to a conceptual reanalysis of the path into an embedded structure consisting of one smaller distinct path nested within a larger path that can then act as a background reference frame. Thus, though the crate in the aircraft example may be assumed to have objectively traced out a complex path consisting of a horizontal segment followed by a descending parabola, a hearer of the example sentence would probably reconceptualize the motion situation. This reconceptualization would involve a salient straight downward vertical path that is abstracted out as separate from an attentionally = horizontal forward path that preceded the vertical plummet and that the aircraft maintains after dropping the crate. The simpler parts of such a conceptually nested path structure would tend to be demarcated by the so-conceived singularity points located at qualitative shifts.

3.2 Closed Path

The second kind of path, here termed a **closed path**, will refer to the same kind of entity as the open path described in the preceding section with the exception that its beginning point and ending point coincide at the same location in space, so that the path now constitutes a circuit. If this single starting and ending point is treated as lying outside the motional path itself and, hence, outside the event frame, then the initial, medial, and final portions of the event can be additionally identified as being the **departure**, the **away**, and the **return** portions of the path.

The cognitive factors for demarcating an event frame that were adduced in the preceding section might all serve in bounding a closed path as well, with perhaps the factor pertaining to stationary boundary periods as the likeliest to play a role. In the case of a closed path, however, we can perhaps adduce an additional factor, that of spatial coincidence—that is, the fact that two points of the path occupy the same location in space—which permits the conceptualization of the stretch of path looping to and from this location as a unitary entity. This closed-path type will figure below as well in the treatment of cycles with phase windowing.

The example in (5) illustrates this closed-path type. Given the context, the whole event in (5a) can effectively be evoked by any of the alternatives of windowing indicated in (5b): basically, all the possibilities occur except windowing of the departure portion alone. Again, the windowed portions are foregrounded in attention while the gapped portions are backgrounded.

(5) a. [I need the milk.]
(1) Go (2) get it out of the refrigerator (3) (and) bring it here.

b. The whole can be represented by:

i. 2:	Get it out of the refrigerator.	[*medial windowing*]
ii. 3:	Bring it here.	[*final windowing*]
iii. 1+2:	Go get it out of the refrigerator.	[*final gapping*]
iv. 2+3:	Get it out of the refrigerator and bring it here.	[*initial gapping*]
v. 1+3:	Go bring it here.	[*medial gapping*]
vi. 1+2+3:	Go get it out of the refrigerator and bring it here.	[*full windowing*]

3.3 Fictive Path

A spatial configuration that is otherwise understood as static through time can often be alternatively conceptualized so as to be rendered "conceptually sequentialized" and to include a path of "fictive motion" (as characterized in chapters I-1 and I-2). One type of such a **fictive path** is the "trajectory" exhibited by a person's focus of attention shifting over a conceived scene. When the linguistic formulation of a sentence is of the sort that can direct a hearer's attention along such a trajectory, this indication of a fictive path is amenable to the same windowing patterns as is a reference to a path of physical motion.

One English construction that directs one's attentional focus along a spatial path in this way is "X BE across Y from Z." This construction is comparable to the construction "X BE between Y and Z" in that both specify a complex spatial schema that includes two reference points (the Ground objects Y and Z). But the "between" construction calls for a stationary distal perspective point with global scope of attention over the spatial schema as a whole, whereas the "across from" construction specifies a moving proximal perspective point with local scope of attention on elements of the schema taken in sequence. In particular, the construction directs that one's focus of attention describe a path that begins at point Z, that next traverses the extent of Y, and that lastly terminates at point X. This construction thus specifies a fictive equivalent of an open path. The construction is exemplified for two different referent scenes in (6), shown with full windowing, medial gapping, and initial gapping, respectively.

(6) a. *With maximal windowing*
   i. My bike is across the street from the bakery.
   ii. Jane sat across the table from John.
   b. *With medial gapping*
   i. My bike is across from the bakery.
   ii. Jane sat across from John.
   c. *With initial gapping*
   i. My bike is across the street.
   ii. Jane sat across the table.

In the (b) forms, the spatial complex is medial gapped by the omission of the Y component of the construction. Here, the gapped portion is backgrounded and its identity is generally provided by the context or by convention, while the discontinuously windowed portions, the Figure and

"initial Ground," are conceptually abutted against each other (as described further in the next section). In the (c) forms, the spatial complex is initial gapped by omission of the entire "from Z" constituent. In this case, again, the backgrounding of the initial reference point is associated with the assumption that its identity is clear from the context or from convention. To illustrate and elaborate on this idea with a sentence like *The injured cow is across the field*, the implicit initial point is typically (1) a location already in reference (e.g., across from where I had said the tractor had broken down), or (2) the current deictic center (e.g., across from where we are now standing), or possibly (3) a canonical location (say, across from the only gas station on the road).

Again, the cognitive factors for demarcating an event frame that were adduced earlier for an open path of physical motion might all serve in bounding a fictive open path. However, in the case of the *across from* schema and certain other fictive path types, one may perhaps adduce an additional factor of bilateral symmetry, where the two X and Z elements that can be understood as bounding the event frame can in some respect be taken to represent corresponding points in a reflection about a central axis. A factor of this sort seems more evident where the two elements have reversed geometries (e.g., have fronts pointing in opposite directions so as to face each other), as would generally be inferred for the scenes represented by sentences like *Jane sat across from John* or *The couch was located opposite the armchair*. But even in the scene represented by the sentence *My bike is across the street from the bakery*, the bike and the bakery can in some sense be regarded as the bilaterally symmetric "bookends" at either end of a path that lies a bit beyond either side of a geometric strip (the street).

3.4 Conceptual Splicing

With particular regard to the attentional backgrounding that takes place for the medial path portion, consider together all the medial-gapped forms of path windowing above: *The crate fell out of the plane into the ocean, Go bring the milk here, My bike is across from the bakery,* and *Jane sat across from John*. For these and similar cases, the medial portion of the path in some hearers' cognitive representations may reduce to so minimal a state in conscious conceptualization that the discontinuous initial and final phases may seem to run together contiguously, perhaps even seamlessly. This cognitive phenomenon can be termed **conceptual splicing** and may be taken to constitute a particularly significant cogni-

tive process. The next section will present further forms of conceptual splicing and will explore its cognitive ramifications.

4 CAUSAL-CHAIN WINDOWING

What on other grounds and in other cognitive systems can be understood as a "causal continuum" is, instead, in the conceptual organization that seems to underlie much of the linguistic and, no doubt, additional cognitive systems, prototypically conceptualized as a sequence of linked "events," or "subevents"—that is, the equivalent of a so-conceived chunking of the continuum into relatively discrete packets—in which the sense of causality may be associated only with the boundary between each subevent and its linked successor.[7]

A causal chain can constitute another type of sequential event frame, a **causal-chain event frame**, which exhibits windowing of attention in what may be termed **causal-chain windowing**. Analyzed in the way that seems to underlie much linguistic structure and possibly other cognitive structure as well, the type of causal chain understood to be initiated by an intentional agent progresses through the sequence of subevents characterized next and schematized in (7) (see chapter I-8). The cognitive agent first intends that a particular event will occur and that it will result from her action. The agent then generates an act of volition, a subevent that will cause a certain whole-body or body-part motion in the case where the intended outcome is in the physical realm. The resulting bodily motion is a subevent that will then—in the case where it is not itself the final intended outcome—cause a second physical subevent.

To this point, three levels of initiation can be distinguished: the agent's original conceiving of an intention can be regarded as the event that initiates the entire processual complex, with its identifying of a goal and the steps that can lead to it; the volitional act can be regarded as the subevent that initiates the full causal sequence of subevents; and the bodily motion can be regarded as the subevent that initiates the physical portion of this causal sequence.

Resulting from the subevent of bodily motion, there may then ensue an intermediate causally linked chain of subevents. And resulting from the body-motion subevent or from the last in such an intermediate chain of subevents, there may next occur a penultimate subevent, which would thus constitute the immediate cause of the final result. Finally, caused by one of the preceding subevents, there takes place the final resulting

subevent—that is, the goal that the agent originally aimed for as the end of her scope of intention.

(7) Semantic composition of a physical causal chain with an initiatory intentional agent

Agent's scope of intention
[——————————→]
[1] → [2] → [3] → [4] → [5]
Sequence of causally chained subevents

[1]: Agent's act of volition that activates bodily motion
[2]: Bodily motion of the agent (particular body part(s) or whole body) that initiates the physical causal chain
[3]: Intermediate causally chained subevents
[4]: Penultimate subevent = immediate cause of final result
[5]: Final resulting subevent = agent's intended goal within scope of intention

NB: a. [3] may be absent
b. [3] may be absent and [2] may coincide with [4]
c. [3] and [4] may be absent and [2] may coincide with [5]

With regard to factors that might function to cognitively demarcate an event frame of the causal-chain type, certainly in the present kind involving an initiatory agent, the straightforward determiner of such demarcation would be the agent's scope of intention. More specifically, the event frame would consist of the sequence of occurrent or projected causal subevents, beginning with the agent's volitional act and ending with the agent's goal, that is encompassed within the scope of intention assumed for, attributed to, or claimed by the agent.

4.1 Discontinuous Windowing over Agent + Result (+ Immediate Cause)

What is noteworthy about the characteristic or grammaticized structure of constructions that refer to causal chains in most familiar languages is that the entire medial portion of the sequence is gapped, with discontinuous windows solely on the initiatory agent and the finally resulting subevent. For example, a standard English causative construction like *I broke the window* refers to the initiatory agent, "I," and to the final subevent, "the window broke," and indicates that the former intended to, and did, bring about the latter. But there is no indication of what bodily motions the agent undertook to execute the intention—say, my bending down and

moving my hand to grasp a rock on the ground, straightening up and lifting the rock with my hand, swinging my arm while holding the rock in my hand, and releasing the rock from my hand, thus propelling it forward. Nor is there an indication of what intervening causally linked subevents might have occurred—say, the rock's sailing through the air followed by the rock's making contact with the window; nor of what the immediate cause of the final result might have been—say, the rock's forcefully impacting with the window.

Of the material characteristically gapped from the middle of a causal chain, the portion that seems crosslinguistically to have the next-most-ready means for expression is the penultimate subevent of the causal chain—that is, the immediate cause of the final intended result. In English, this penultimate subevent is readily expressed in a *by*-clause, as in the case where the situation in which I intentionally lift, swing, and propel a rock through the air into a window to break it can be expressed by a sentence like *I broke the window by hitting it with a rock*, shown in (8g). This *by*-clause, however, does not accommodate any other subevents in the whole causal chain, from the act of willed bodily motion to the ante-penultimate subevent, as seen in the unacceptability of (8a) through (8e). For many speakers, even a *by*-clause like that in (8f) is not acceptable, and speakers who do accept it do so because they feel that the clause contains within it reference to the penultimate subevent in which the rock actually impacts the window.[8]

(8) *English by-clause reserved for penultimate subevent*
I broke the window
a *by grasping a rock with my hand.
b. *by lifting a rock with my hand.
c. *by swinging a rock with my arm.
d. *by propelling a rock through the air.
e. *by throwing a rock toward it.
f. ?by throwing a rock at it.
g. by hitting it with a rock.

Supporting the next-most-privileged status of the penultimate subevent in a causal chain is the fact that some languages do in fact characteristically or obligatorily identify that event in a causative construction. Thus, in Atsugewi, in most cases a verb root requires a prefix, selected from a set of some two dozen, that specifies the penultimate subevent (see Talmy 1972 as well as chapters II-1 and II-2). For example, consider a situation

in which I used my hands to build a fire with which to destroy a house. To refer to this situation, I can use the verb root *-miq-* '(to cause) an architectural structure to lose its structural integrity' together with the instrumental prefix *mu:-* 'by acting [on the Patient] with heat/fire'. But I cannot use that verb root together with the instrumental prefix *ci-* 'by acting manipulatively [on the Patient] with one's hands'. The reason is that the former prefix refers to the mandated penultimate subevent whereas the latter prefix refers to an earlier subevent.

In a comparable way, in the characteristic English verb + satellite construction in which the satellite expresses the final resulting event and the verb expresses a prior causal subevent, this causal subevent must again be the penultimate one, and nothing earlier (see chapter II-3). Thus, if I have grasped a lever and then used it to pry a lid off a box so as to open the box, I can refer to this causal sequence with *I levered the box open* but not with **I grasped the box open*. Similarly, the previous arson situation—in which I have lit a fire so that a house would catch fire from that and proceed to become consumed in flames to the point of its destruction—can be referred to by the sentence *I burned the house down*, but not by a sentence whose verb expresses any causal subevent prior to the penultimate one, as in **I lit/kindled the house down*.

4.2 Windowing of Causal Chains with Intermediate Cognitive Agents

Following the activities of an initiating Agent, an ensuing causal chain can include additional cognitive entities whose agency is essential in the sequence leading to the final reported result (see chapter I-6). However, to the extent that material referring to such intermediary agents is gapped from a sentence, the intentions, volitional acts, and effects of these agents are attentionally backgrounded, conceptually neglected, and thereby rendered causally "transparent"—that is, subject to the conception of a causal continuity progressing directly through such agents rather than stopping at each agent and being renewed by a fresh act of intention and volition. This effect is seen, for example, in the sentence *I'm going to clean my suit at the dry-cleaning store on the corner*, which omits mention of the cleaners whom the speaker will engage to do the job. Further, the amount of the neglectable intervening material can be enormous, as seen in the referent of a sentence like (9a), which, though mediated by a whole society over decades, can still be conceptualized in terms of a juxtaposition of an individual initiator and a final result.

(9) a. The Pharaoh built a pyramid for himself/*him.
b. The Pharaoh had a pyramid built for himself/him.
c. The Pharaoh had his subjects build a pyramid for *himself/him.

This example further allows us to note that the syntax of the reflexive in English, though usually treated in solely formal terms, nevertheless can be seen to correspond to actualities of conceptualization. In this regard, we can observe that the form in (9a), which windows only the initiator and the final result and distracts little attention onto intermediary factors, requires the reflexive in referring back to the initiator and excludes any use of the nonreflexive for this purpose. However, the (9b) form, whose "*have* + *-EN*" construction adds a window onto the presence of an intermediary agency, though not onto its identity, permits *either* the reflexive *or* the nonreflexive. Further, the (9c) form, with a construction that now also refers explicitly to an identified mediating agency, requires the nonreflexive and excludes use of the reflexive.

In this sequence of forms, we can discern the presence of clines in three different linguistic systems—syntax, semantics, and conceptual structure—and of correlations across these clines. Thus, with respect to syntax, there is a dual cline that involves both a successively lengthening verb complex and a shift along an obligatory-optional axis. In particular, proceeding through (9) above, the cline progresses from a simplex "V" (*build*) with a requirement for the reflexive in (9a); through the form "*have -EN* + V" with the allowance of either the reflexive or the nonreflexive in (9b); to the complex "*have* + NP + V" with a requirement for the nonreflexive in (9c).

In correlation with this syntactic cline, there is a cline in referential semantics—that consisting of the specification of the intermediary agency—which ranges from null specification in (9a), through indication of the presence of such agency without specification of its identity in (9b), to specification of both its presence and its identity in (9c).

And, in correlation with these syntactic and semantic clines, there is an attentional-conceptual cline with dual aspects. In this cline's progression from (9a) to (9c), there is an increase in the strength of attention directed to the presence of the intermediary agents (as distinguished from the mention and identification of them that was treated in the preceding cline). Further, there is a qualitative shift in the conceptualization of the relationship between the initiator and the final outcome that ranges from

a sense of direct causal immediacy in (9a)—another case of the cognitive splicing effect—to a sense of causal distance in (9c).

The lexicosemantic "logic" here is presumably that a reflexive form suggests a more direct connection between two references to a single entity, thus according better with the conceptual immediacy of the initiator-outcome relationship in (9a), whereas a nonreflexive form suggests a more distant connection between two references to a single entity, thus according better with the conceptually greater causal distance between the initiator and the final outcome in (9c).

The middle form, (9b), is the most telling for a demonstration that the role of semantics is here more determinative than that of syntax. For while there may be solid syntactic arguments for the necessity of the reflexive in (9a) and for the nonreflexive in (9c), there is no immediately obvious non–ad hoc syntactic justification for open use of either the reflexive or the nonreflexive in the (9b) form. But the semantic-conceptual account involving a gradient in the cognitive salience of the intermediate causal factors does accord neatly with the overt linguistic behavior.[9]

4.3 Cognitive Underpinnings of Causal Windowing and Gapping

Again, what is cognitively noteworthy in the characteristic medial gapping of causal sequences is the great degree to which the middle portion is reduced in one's field of attention, and sometimes seemingly eliminated from it, in the cognitive process of conceptual splicing noted earlier. With its patterns of causal windowing and gapping, language structure here appears to reflect a cognitive structuring in which a sentient agent's intention for the occurrence of a particular state or event and its actual occurrence are characteristically conceptualized together as a kind of melded unity in the foreground of attention, with little or no attention directed to the intervening mediating stages. This conceptual arrangement would seem to match a presumed kind of experience recurrent from earliest age on in which an intention and its realization, both in awareness, feel seamlessly linked. This experience includes little or no awareness of mediating actions and events—ones that, if considered, might be taken for granted as automatic bodily movements and expectable physical occurrences.[10]

One may speculate that biological evolution has resulted in this form of cognitive structuring of attention for its selective advantages, namely, that it constitutes a functionally relevant type of invariant or constancy in

cognition while allowing for other forms of necessary plasticity. The constancy here is the goal of achieving a correspondence between an intention to effectuate some particular circumstance and seeing to it that that circumstance in fact becomes realized through whatever activities prove necessary. Where cognitive organization must remain plastic is in the determination and marshaling of such necessary activities, since the conditions attendant on realizing some purpose can vary greatly.

There are two main categories of such variation. First, the physical and functional constitution of any individual organism can change, whether by ontogenetic development or by environmental impact, including injury. Second, the characteristics of an organism's surroundings, both physical and social, can change during its lifetime or can vary in accordance with where the organism is born. The overall function of the cognitive processes here posited to be in operation would thus be to maintain a goal schema as constant and to execute it through variously appropriate means across constitutional and environmental variety and change.

To illustrate these notions, consider as a candidate for a commonplace cognitive invariant the intention to move forward while avoiding obstacles. With respect to constitutional change through ontogeny, as a human individual develops from an infant into an adult, she will replace crawling on all fours by bipedal walking to effectuate this forward-motion intention, thus ontogenetically changing the means marshaled while maintaining the goal schema intact. As for constitutional change due to external impact, if that individual were to suffer the loss of a leg, the baby crawling on three limbs or the adult walking with crutches would now execute a new movement pattern while still realizing the same goal of forward motion with avoidance of obstacles. To exemplify environmental variety, if the adult learns to drive a car, he replaces the use of alternating leg movements for that of a slight pressure of the right foot on a pedal to effectuate the same goal of forward motion, and he replaces judging lateral clearance for the span of his shoulders by assessing instead the clearance for his car's fenders in maintaining the same goal of avoiding obstacles.

In fulfilling the function of maintaining goal-schema constancy, the degree of plasticity of execution can clearly be enormous, as evidenced, for example, by a human's ability to learn to move forward across a range of implementations as disparate as crawling on all fours, limping along on crutches, driving a car, swimming underwater, or propelling herself in the

microgravity of a space shuttle. In all such cases, the primary attentional window can remain constant, encompassing only the intended goal and its realization, and, once proficiency has been achieved, little or no attention may be directed to the particular physical means and movement patterns engaged in to execute this goal. A cognitive concomitant of this attentional restriction can be an experience in the individual of the maintenance or continuation of a single "sense of body" across all the variation of physical means for executing the goal—or, from a dynamic perspective, an experience in the individual of the projection of his baseline sense of body into the divergent new means employed to execute the goal. This phenomenon is evident, for example, in the way that a driver can invest his car with the experiential property of being an extension of his body or even of constituting his body, or in the way that the operator of a remote robotic device (such as a mechanical arm) often has the experience of being present at the distal location in what has come to be termed "telepresence."

Note further that the implementational range of the disregarded intermediate causal phenomena can encompass the role not only of the body and mechanical extensions of the initiating agent but also of the voluntary cognitive and physical contributions of other mediating sentient agents. Linguistic evidence of this expanded plasticity was given in the preceding section. Comparably, for the preceding conceptual case of intended forward motion, an individual who has, say, taken a bus part of the way in getting to town can experience his going into town in terms of his intention to do so and its realization, with little or no attention directed to his reliance on a bus driver to transport him in the bus over a portion of the path. Evidently, our cognitive system of executional plasticity can include the utilization of the actions of other agents so that these, too, subserve our cognitive constancy system for intentions and their realization.

Given the familiar examples of plasticity in motor execution and in bodily identification noted in the preceding, little surprise should be caused by the recent successes of computer-based "virtual reality" in placing an individual in circumstances unusual for perception and motor control. Virtual reality simply makes extended use of plasticities long since selected for and everywhere evident. If anything, virtual reality systems at present are still shy of incorporating certain commonplace capabilities of our everyday executional plasticity—for example, our inclusion of the actions of other agents as being within our control in addition to our control over our bodies and their direct extensions. Where the tech-

niques of virtual reality can in fact prove most instructive is, complementarily, in ascertaining the constraints on and limitations of human plasticity. For example, would it be feasible for a subject in a virtual reality system to adapt to conditions where the more slowly she moves, the faster the scene shifts, and vice versa, or where the softer the pressure she exerts, the more forcefully the objects in the scene behave, and vice versa? Could a subject learn to bodily identify with an octopus figure depicted in the virtual scene, integratedly controlling each of the eight limbs with eight different kinds of motion of her own body?

The perspectives and evidence arrayed above argue for the selective advantage in the evolution of a cognitive system of intentional constancy—which maintains certain abstract schemas of intention and its realization—beside a cognitive system of executional plasticity. In the same way that cognitive linguistics has proposed other close correspondences between linguistic structure and the structure of nonlinguistic cognitive systems, the thesis proposed here is, specifically, that the portion of an agentive causal chain characteristically windowed in linguistic structure corresponds to the cognitive system of intentional + realizational constancy, while the characteristically gapped material corresponds to the cognitive system of executional plasticity.

5 PHASE WINDOWING

A further type of event frame consists of an event that iterates in a cycle—what will be termed a **cycle event frame** here. A sentence referring to such an event can direct the positioning of a window of strongest attention over a particular phase of that iterating cycle—a cognitive process that is termed **phase windowing** here. The overall event comprised of an iterating cycle is sequential but may have no clear beginning, middle, or end portions in reference. However, each component cycle when abstracted out can be thought to have the usual initial, medial, and final portions of a sequential event. Further, though, now that this sequence repeats and can be interpreted as additionally having a rest state between iterations, it can be considered to have an **initial**, **medial**, and **final phase** as well as a **base phase** that occurs after the final phase and before the initial phase. In the specific case where the overall event is a motion event and one component cycle constitutes in particular a closed path of the type treated in section 3.2, then the earlier distinctively labeled portions of a closed path now become its "departure phase," "away phase," and

"return phase," while the base phase can now be distinctively labeled as its **home phase** and be understood as constituting the state of locatedness at the spatially coincident point of the closed loop.

The conceptual event frame associated with a cyclic occurrence may have a hierarchical structure, unlike previous cases. Rather than comprising a larger frame that directly spans the overall event, it would seem instead to consist of a frame around just one cycle's worth, but with the sense that successive iterations of this cycle are superimposed on each other within the single smaller frame. Thus, with respect to the factors that can cognitively define an event frame, one can here posit a further factor, that of the part-for-part congruence of one segment of occurrence with another—a property of direct mappability between segments. This factor can function to cognitively delimit a portion of occurrence that can constitute such a segment and that can thus be conceptualized as a unit event. Here, any two temporal points that bear the same phase relation to two such congruent stretches of occurrence can be taken to constitute the boundaries of one cycle's worth. And, in particular, points of this sort that occur within what can be conceptualized as the "basic" or "home" portions of occurrence have a privileged status for constituting the boundaries of a cycle.

To illustrate cycles with an iterated closed path, the sentences in (10) can all—given a sufficiently constrained context—be taken to pertain to the same cyclic event frame in which the home phase consists of a pen lying on a table, the departure phase consists of the pen falling off the table onto the floor, the away phase consists of the pen lying on the floor, and the return phase consists of my picking the pen up from the floor and placing it back on the table. Exhibiting alternative options for attentional windowing, however, the sentence in (10a) windows greatest attention on the departure phase of this cycle (or, more precisely, on just the earlier portion of the departure phase, comprising the pen's falling down off of the table but not down onto the floor), leaving the remainder of the cycle in the background of attention. The sentence in (10b) windows only the return phase (or, more precisely, only the later portion of the return phase, comprising my lifting the pen up onto the table but not up off the floor). The sentence in (10c) places discontinuous windows over the departure and return phases while leaving the remainder of the cycle in the background—as schematized in the accompanying diagram. Thus, here as before, the language affords the speaker alternatives of attentional windowing on essentially the same event frame with the addressee feasibly

able to infer the different gapped portions for each alternative so as to reconstruct back to the same single event frame. Further, the sentence in (10c) can be taken to induce cognitive splicing in the hearer by conceptually running together the departure and return phases, with the extreme backgrounding or loss now not only of the medial phase but also of the base phase (i.e., of the static home and away phases).

(10) a. *With departure-phase windowing*
The pen kept falling off the table.
b. *With return-phase windowing*
I kept putting the pen back on the table.
c. *With departure-phase plus return-phase windowing*
The pen kept falling off the table and I kept putting it back.

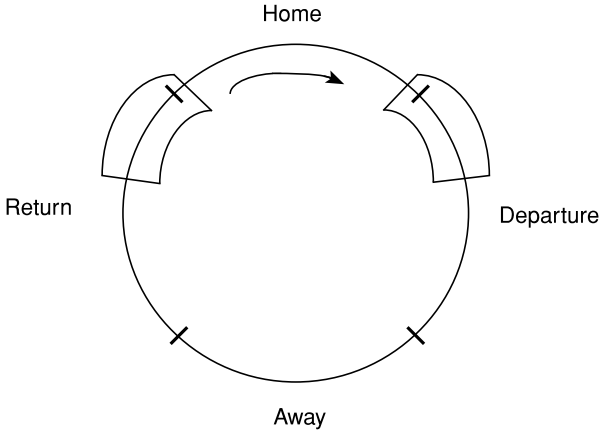

In this chapter's examples, including the preceding example, alternatives of windowing constitute different attentional patterns, but these patterns are placed over what can otherwise be the same single referent. However, the cycle event frame can also support referentially nonequivalent phase windowings. This can arise where a particular phase window is established by some reported external coincident event, rather than by the speaker's predilection. To illustrate, the main cyclic event could be an iterated closed path undertaken by a Mr. Smith with respect to his office: being in the office (home phase), leaving it for another location (departure phase), being at that other location (away phase), and going from that location back to his office (return phase). And the external coincident event could be my repeated telephoning of Mr. Smith always during the same particular phase of his path cycle. The three sentences in (11) express

such a coincidence for three different phases of this cycle; the diagram schematizes the coincidence pattern for sentence (11c).[11] The phase windowings selected out in this way are clearly part of three referentially distinct situations.

(11) Whenever I phoned,
  a. Smith was always just about to step out of his office.
  b. Smith was always just stepping out of his office.
  c. Smith had always just stepped out of his office.

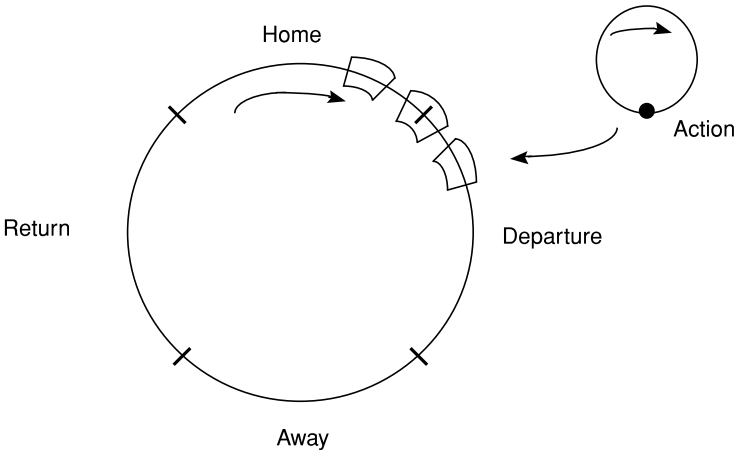

6 PARTICIPANT-INTERACTION WINDOWING

Consider a complex situation that consists of two parts: (1) a primary circumstance, and (2) some participant(s) interacting with that circumstance on (at least) two different occasions. A "participant" here can be a participant either of the expressed referent event or of the current speech event. A participant's interaction with the circumstance can be direct, as in observing or considering the circumstance, or indirect, as in asking another participant about the circumstance. In referring to the whole of such a situational complex, some languages have provision for the alternative placement of a window of heightened attention on one or the other of these two interactions. In particular, linguistic devices direct an addressee to adopt one of the two participant interaction times as the point at which to locate his temporal perspective point, and to place around the interaction there an attentional window that could include such elements

of the interaction as the activity, the surrounding scene, or the cognitive content of the participant.

The whole situational complex can be understood to constitute a new type of event frame, the **participant-interaction event frame**, which permits alternatives of **participant-interaction windowing**. This type of event frame shares a characteristic with the preceding types (and in this respect all these types differ from the type treated next), namely, that it constitutes a sequence of phenomena differing through time and that, accordingly, the alternative windows of attention differ with respect to their temporal placement.

Considering again the kinds of factors that can demarcate the boundaries of an event frame, the present type of event frame may exhibit a further such factor: the occurrence of two distinct punctual events extrinsic to a certain circumstance that extends through time can mark out a portion of that circumstance and establish that portion conceptually as an event frame. Here, in particular, a portion of the primary circumstance is marked out by two participant interactions with it.

For a first illustration, the two short segments of discourse in (12) can be interpreted as referring to a single situational complex that is of the sort just outlined.

(12) a. John met a woman at the party last week. Her name was Linda.
b. John met a woman at the party last week. Her name is Linda.

Here, the primary circumstance is the temporally unbounded state of a certain woman's having the name Linda. It can be argued for (12a) and perhaps even more strongly for (12b) that each of these segments of discourse equally evokes the same concept of a pair of participant interactions with this circumstance.

The first interaction, an indirect one, is that of John with the primary circumstance of a woman's being named Linda, namely, his encounter at the party last week with the woman bearing that name. He may have asked and/or been told her name, or the discourse may be providing that information without his having learned it. The second interaction, a direct one, is my—that is, the speaker's—consideration of the woman's name at the present moment of speaking. In the second sentences of (12a) and (12b) referring to the woman's having a name, the use of the past tense in (12a) and of the present tense in (12b) then signals the differential placement of an attentional window over one or the other of these interactions.

This past tense in (12a) might seem peculiar, since the past tense is largely associable with an event that has terminated before the present moment, whereas in fact the state of the woman's bearing her name is continuous. The explanation for the use of the past tense here, we would argue, is that, despite the overt syntax, it does not apply to the main referent of the sentence—that is, to the woman's being named Linda—but rather to the time of the first participant interaction: to John's encounter with the woman. The window of attention placed around that temporal point would then include aspects of the interaction, some of them inferred or imagined, such as John's interchange with the woman or the surrounding party scene. On the other hand, the present tense of (12b) signals the adoption of the temporal perspective of the second participant interaction—that is, the present moment—and directs the placement there of an attentional window that includes something of the interactional context, such as my contemplation of the woman's name-bearing state either in its current relevance or in its ongoing unbounded character.

Accordingly, we have here in the participant-interaction case a type of windowing rather comparable to those in the preceding sections where each of two formulations evokes the entirety of a particular event frame while explicitly indicating only certain subportions of that event frame and thus establishing a selective window of attention on it.

Although it was just argued that each of the discourse sequences in (12) at least implicitly evoked a pair of participant interactions with the primary circumstance, nothing in the sequences explicitly specified the duality of interaction. But in (13), the word *again* unmistakably indicates that there were at least two interactions in the situation.

(13) a. What was your name again, please?
b. What is your name again, please?

In the situational complex here, the primary circumstance is the unboundedly continuous state of your having a particular name. The two interactions with this circumstance are, at an earlier moment, your or someone's saying your name in my presence, perhaps with my having asked you for it, and, at the present moment, my asking you for your name.

The initial interaction, while taking place, would have been in the present tense, consisting, for example, of my asking you *What is your name?*, or of your saying *I'm Susan*, or of someone's saying *This is Susan*. But my subsequent question to you would need to be something like one

of the forms in (13). This is because, by the requirements of English for reflecting the pragmatic circumstances here, if I have forgotten or did not catch your name the first time and assume that you and I remember my prior presence at the name's utterance, my subsequent asking must include a marker specifically acknowledging the repetition. Such a marker could be the word *again* or the English "echo question" intonation pattern. Thus, both sentences of (13) explicitly indicate that the present question is the second of two interactions with the same intent on my part to learn your name.

But this second-question formulation with *again* permits the use of either the past or the present tense. And, as before, the past form may at first seem paradoxical in its usage with a temporally unbounded referent (your having a name). However, the explanation of this behavior, as posited previously, is that the choice of tense in the main verb does not pertain to the overt referent of the clause but rather to my two interactions with that referent—that is, the earlier or the later instance of my hearing or asking about the referent. In particular, the past tense of (13a) selects the time of my initial interaction as the point at which one is to locate one's temporal perspective so as to place a window of heightened attention over that interaction, while the present tense of (13b) requires the performance of these same cognitive processes for my later interaction, the one occurring at the present moment of speaking.

Certain observations can serve to reinforce and refine our proposal that attentional windows are placed over participant interactions with the primary circumstance. First, the view that the past and present tenses in the preceding examples direct the placement of windows only over the two participant interactions that we have cited is buttressed by our clear English-speaker intuition that they could not refer to any other bounded temporal periods. Thus, the past in (12a) could not refer to a moment between the time of last week's party and the present moment—say, to a moment three days ago—nor to some time before the party. Comparably, the past in (13a) could not refer to a point between the last time I heard your name and the present moment, nor between any previous occasions of my hearing your name if there were more than one of these.

Second, some might note that the overt tense that appears in the example sentences is expressed as part of a reference to the primary circumstance, rather than as part of some explicit reference to the participant interaction that we have posited. Accordingly, some might prefer to see an alternative analysis in terms of the primary circumstance alone. Such

an analysis might hold that a delimited portion of the unbounded primary circumstance is conceptually marked out for consideration in isolation and that only this portion is placed in a tense relation with the present moment. However, this approach is easily faulted. The primary circumstance does not have to be continuous and unbounded, as it has been in the previous examples, but can also be a punctual event that occurs only once. Yet here, too, the account in this section will still hold. Thus, each of the sentences in (14) equally reflects two interactions I have had with you over your knowledge of a plane schedule. And they would seem to differentially window respectively the earlier interaction and the present interaction. But there is now no possibility of interpreting the tense as applying to some marked out subportion of the primary referent, since this is now the punctual and upcoming plane departure.

(14) a. When was her plane going to leave again tomorrow?
b. When is her plane going to leave again tomorrow?

Further, if there really were a tense-located referent of the sentence that indeed consisted of a temporally delimited subportion of an otherwise unbounded circumstance, then that referent portion should be compatible with an overt constituent that explicitly refers to the delimited time period in question. But such additional constituents, on the contrary, render the sentence unacceptable, as seen in (15).

(15) a. John met a woman at the party last week. Her name was Linda *while he was there. / *when he asked her for it. / *when she told him.
b. What was your name again *when I asked you for it before? / *when you told me it before?

The unacceptability of these sentences further indicates that the use of a past or present tense in the example sentences of this section cannot be accounted for simply as some automatic syntactic reflex involving, say, some sequence-of-tense rule that is triggered by some other time-specific constituent but, rather, must genuinely reflect a semantic option.

Note that some participant-interaction-type sentences can support an alternative ''evidentiary'' reading that does permit temporally specific adjuncts. Thus, the segments of discourse in (16a) and (16b) would probably first be read in accordance with the participant-interaction analysis of earlier examples—with the unbounded iterative activity of a geyser's spouting replacing the unbounded static state of bearing a name. In this

reading, the two segments would refer to approximately the same situational complex and differ only as to their pattern of attentional windowing.

(16) a. I was in Yellowstone Park last year. Old Faithful spouted regularly.
b. I was in Yellowstone Park last year. Old Faithful spouts regularly.
c. I was in Yellowstone Park last year. Old Faithful spouted regularly (—at least) while I was there.

But, in addition to such participant-interaction readings, these two sentences can have evidentiary readings that now have meanings substantially different from each other. Under such readings, in (16a), I report only what I witnessed during my visit and suggest no inferences about activity outside that scope. But in (16b), I use what I witnessed as evidence to confirm the general notion that there is continuous unbounded activity. And, in (16c), a temporally delimiting constituent of the type seen above to be unacceptable appears compatibly with the past tense form here. Such evidentiary readings do not instantiate participant-interaction event frames or windowing and, accordingly, their tenses apply directly to the overtly expressed referent in the usual way.

In all the preceding examples of this section, the primary circumstance being referred to is unchanging through the progression of time. However, reference to a circumstance that does change with time can constitute a further case that looks like participant-interaction forms, but really is not. In such a case, the choice of tense applies directly to the primary referent rather than to a participant interaction, and a temporally specific adjunct is permissible. This can be seen in (17a), where the changing primary circumstance is the time of day.

(17) a. The time was 10:53 when I asked for it.
b. *The woman's name was Linda when I asked for it.

Accordingly, with the *again*-question frame used earlier, if I am now asking you for the time of day for the second time—where the first time I asked was sufficiently earlier to render the answer you then gave pragmatically useless—I cannot felicitously use the past tense, as in (18a), but must rather use the present, as in (18b).

(18) a. #What time was it again, please?
b. What time is it again, please?
c. What time was it again when I asked you before?

On the other hand, I can felicitously wish to know the answer you delivered on the earlier occasion—for example, where I was recording in a notebook the ongoing results of an experiment. In this case, the past tense would be acceptable and could now be used with an overt constituent explicitly referring to the past moment in question, as seen in (18c).

Note that the *again* in the sentences of (18) still pertains to the speaker's dual interaction with the primary circumstance, namely, to the fact that I have now twice heard or asked you for the time. But—unless the time of my first interaction is pragmatically recent enough, in which case the tenses in (18a) and (18b) can revert to their participant-interaction windowing usage—the tense can now no longer be used to window one of these interactions since its use is preempted for pertaining to a particular subportion of the changing circumstance.

Finally, consider again the original examples of participant-interaction windowing. Here as elsewhere, although the choice of window placement does not affect the principal situational complex being referred to, it does have further semantic consequences. Thus, the tense used can suggest the relevance that the primary circumstance has to current concerns, with the past suggesting lack of relevance and the present suggesting the presence of relevance. For example, the past tense in (12a) can suggest that John's association with the woman last week at the party ended there, while the present tense in (12b) can suggest that their association has continued to the present and is of current relevance.

7 INTERRELATIONSHIP WINDOWING

A frequent type of language-relevant cognitive entity is a conceptual complex that contains or is comprised of parts not autonomous in themselves but intrinsically relative with respect to each other, where the presence of one such part necessarily entails the presence of the other parts. A conceptual complex of this sort is here called an **interrelational complex** and can constitute a further type of event frame, the **interrelationship event frame**. Such an internally self-entailing complex could logically be considered a single-unit entity, but our conceptual and attentional systems are so organized as to be able to conceptualize the whole as if portioned out into quasi-independent elements to which heightened attention can be differentially directed. With respect to its linguistic expression, such a complex can be conceptually partitioned—in a way that may be universal—into parts expressed by syntactically distinct constituents. Frequently, a

language will permit alternatives of windowing over one or another part of such a complex, while mention of the remaining parts is omitted—although their presence is still understood. Such alternatives of **interrelationship windowing** allow the selection of a locus of strongest attention within a complex or the adoption of a particular perspective over the complex while—given the appropriate context—still conveying the whole of the complex.

Note that the earlier types of windowing do not seem to fit this notion of an intradependent interrelationship. For example, for path windowing, a later path segment is not entailed by an earlier one but is rather represented as being additionally present. By contrast, in interrelational complexes, the relevant components co-define each other. Accordingly, once again considering the factors that can function to demarcate an event frame, the boundaries of an interrelationship event frame can apparently be determined by a new factor, that of co-entailment. Apart from these differences, however, what is common to both the earlier types and the present type is that each event-frame type supports alternatives of the placement of attentional windows over it, and the gapped portions are largely recoverable by the hearer—whether by inferences involving entailment or by inferences involving familiarity with other event frame–determining factors at work in a particular context.

We examine here two kinds of interrelationship event frames, one based around Figure and Ground roles, and the other around factual and counterfactual conditions.

7.1 Figure-Ground Interrelationship

As they are characterized in chapter I-5 for their function in language, the Figure and the Ground in a spatial scene are relative concepts necessarily characterized with respect to each other. The Figure is a moving or conceptually movable entity within the scene whose site, path, or orientation is conceived of as a variable of which the particular value is the relevant issue and that is characterized with respect to the Ground. The Ground is a stationary reference entity within the scene with respect to which the Figure's site, path, or orientation is characterized. As described in chapter II-1, the Figure and Ground are components of an event of Motion (covering both motion and location) that includes two further components, as in the semantic structure in (19).

(19) [Figure + Fact-of-Motion + Path + Ground]

This Motion event well exemplifies the kind of conceptual entity that is intrinsically irreducible—that is, of which no part can exist without the rest—but that in general is conceptually and linguistically partitioned into components that can be treated differentially as to attentional distribution. This conceptual entity, then, constitutes a particular type of interrelationship event frame, the **Motion event frame**, and it can support a particular type of attentional alternativity, **Figure-Ground windowing**.

To illustrate this type of windowing, consider a scene in which paint is peeling off a wall, where the paint would be understood to function as the Figure relative to the wall as Ground. For mention of both the Figure and the Ground within a single sentence, English often has available two counterpart constructions (analyzed in detail in chapter 10 of Talmy 1972), one in which the Figure appears as the subject and the Ground in an oblique phrase, as in (20a), and another in which these grammatical relations are reversed, as in (20b).

(20) a. The paint is peeling from the wall.
b. ?The wall is peeling of its paint.

If there were a need to gap reference to the Figure or the Ground, the constituent referring to it would have to be omitted. Since English does not generally permit the omission of a subject NP but can often omit an oblique constituent, as here, (21) shows two further counterpart constructions based on the preceding pair but with the oblique constituents missing.

(21) a. The paint is peeling.
b. The wall is peeling.

Given the appropriate context, then, (21a) refers to the original scene but with windowing of the Figure (plus the activity) and gapping of the Ground, whereas (21b) windows the Ground (plus the activity) while gapping the Figure.[12] Thus, with such alternative constructions, one can refer to basically the same interrelational spatial complex of codependent Figure/Ground elements and selectively window one or the other of those elements.[13]

7.2 Factual-Counterfactual Interrelationship

A linguistic construction can have the semantic property of presenting the referent of its overtly expressed material as being the case or, alternatively, as not being the case. In traditional terminology, these are, respectively, factual and counterfactual constructions.

Further, a language can have a pair of constructions, one of them factual and the other counterfactual, such that if their overtly expressed materials are positive-negative counterparts of each other, then both constructions make the same overall statement. Given the availability of a particular doublet of such paired constructions, a speaker can make the same overall statement in choosing either one of the constructions, but the speaker would also thereby select whether to direct greater attention to something that was the case or to something that was not the case. Since each member of such a pair of counterpart factual-counterfactual construction types entails the other, their referent types together can be considered to constitute a certain kind of interrelationship event frame, a **factuality event frame**, and the directing of heightened attention to one or the other of these referent types can be called **factuality windowing**.

A factuality event frame exhibits a still further property. Under selective attentional windowing, it can support not only the exclusive consideration of one chosen alternative by itself, but also the placement of the two alternative conceptualizations within a single frame of consideration, so that, although main attention is on only one of the alternatives, the other alternative is still present in a backgrounded way to act as a foil for comparison. An event frame that in this way evokes larger-frame juxtapositions of alternative conceptualizations, can be further said to constitute a **comparison frame**. The characteristic of constituting a comparison frame can then function as one further factor for demarcating an event frame, and the factuality event frame seems to derive some of its characterizability as an event frame from this factor. Certain constructions and lexical forms in a language tend to evoke comparison frames, and the following do so for the occurrence versus the nonoccurrence of some referent.

First, a syntactically negative clause (e.g., *I didn't go to John's party last night*) overtly names something that did not take place but tends to evoke consideration of the corresponding unrealized positive event—and in this respect it differs from a simple positive clause, which tends not to evoke consideration of its negative counterpart. Second, even a syntactically positive main clause when it is adjoined by a *because*-clause (e.g., *I went to the movies last night because they were playing my favorite film*) tends to evoke its unrealized counterpart (a failure to go to the movies) since the inclusion of a reason or cause that has given rise to some realized phenomenon suggests that, in the absence of that cause, the phenomenon would not have occurred. Third, a nonsimple positive clause that also

includes a constituent placing the referent event at some point along a scale of certainty or realizedness (e.g., *Sue may have gone to John's party last night, | Perhaps Sue is at John's party now, | I just barely got to the movies last night*) brings into consideration the existence of such a scale and thereby evokes the consideration of points nearer the opposite pole of the scale. Fourth, an interrogative form, even of an otherwise simple positive clause (e.g., *Did Sue go to John's party last night?*), has as its main semantic point the issue of the occurrence or nonoccurrence of the situation it refers to and, of course, naturally contrasts the occurrence status of its overtly expressed material against the opposite occurrence status. And fifth—the topic of this section—a grammatically counterfactual construction (e.g., *I would have gone to John's party last night if I had had the time*) overtly names a counterfactual event that did not take place (*I . . . have gone to the party*), but it also evokes its factual complement, what actually took place (my staying away from the party). These five types would all seem to be "space builders" in Fauconnier's (1997) terms.

Among sentence types, perhaps mainly it is a simple positive factual declarative clause (e.g., *I went to the movies last night*) that raises in consciousness only the named event without the backgrounded accompaniment of its unrealized alternative. Although it may be the case that a positive statement of this kind is generally made only if its referent is taken to be news to the hearer, unanticipated relative to some baseline of expectation, it seems that such a statement is not usually experienced as an assertion averred contrastively against the potential of its nonoccurrence. Apparently at work here is a cognitive asymmetry that accords to the positive and to the factual the status of having primacy and of being basic, so that the negative and the counterfactual are on the contrary conceptualized as secondary and nonbasic, perhaps as somehow derived from the basic by some cognitive process of reversal.

In addition to construction types like those above, certain lexical items seem to incorporate within their lexicalization a scope encompassing both realization and nonrealization. Thus, the verb *miss*, as in *I missed the target*, seems not to simply refer directly to a projectile's passing to one side of a target, but rather to evoke a two-stage bipartite conceptualization consisting first of the projectile's hitting the target and then the denial of such an occurrence, with a conceptual shifting of the projectile's path off to one side. Comparably, the verb *regret*, as in *I regret that I lent him money*, though referring directly to an actually occurrent event, nevertheless conjures up the wished-for nonoccurrence of that event. Similarly,

the use of the verb *succeed*, as in *I succeeded in opening the window*, shares in common with its nonuse, as in *I opened the window*, a reference to an actually occurring event. But its use differs from its nonuse in that (among other effects) it sets this occurrence of the event within a comparison frame for a contrast with the possibility of the event's nonoccurrence.

As indicated earlier, given that a construction can evoke within a single comparison frame both the factual and counterfactual alternatives of a situation, the issue of windowing enters where the same situation can be referred to by either of two constructions, where one construction names the factual form of the situation while evoking its counterfactual alternative, and where the other construction does the opposite. Why might languages afford ready syntactic means for focusing on what has not occurred? In explanation, one can adduce for the systems of discourse or narrative such factors as the motivation to achieve a heightened effect by specifying a goal that was vainly sought (in the case where the nonoccurrent was preferable to the occurrent), or by specifying a danger that was avoided (in the case where the nonoccurrent was less desirable than the occurrent).

Notationally in the specific analyses that follow, the symbol A, as a mnemonic for "Actual," will represent any particular factual alternative, while the symbol ∼A will represent the corresponding counterfactual. For any particular example, in addition, a P may be used to indicate a clause whose overt syntactic form is positive, while not-P would indicate a syntactically negative clause. Thus, the sentence *I didn't go to the party* can here be represented symbolically as A (not-P) to suggest a paraphrase like "What actually happened is that it was not the case that I went to the party." In truth-value terms, A and ∼A entail each other with the sign of their proposition reversed—that is, A (P) is equivalent to ∼A (not-P), and A (not-P) is equivalent to ∼A (P)—but in terms of conceptual organization, it is necessary to discriminate an A/∼A factual-counterfactual parameter separately from a P/not-P syntactically positive-negative parameter.

The symbols A/∼A are chosen over the symbols T/F of truth-conditional semantics for several reasons. First, the truth-conditional symbols are used in an objectivist system of reference, whereas the orientation here is of a conception-based system of reference, whose theoretical distinctness can be better kept in attention by the use of distinct symbols.[14] Second, it is clearer to show explicitly the counterpart relationship between a matched factual-counterfactual pair with the use of a reversal-type operator like "∼" than with the use of two separate symbols like T

and F, which obscures the fact and the nature of their interrelationship. Third, the assignment of the simpler symbolic representation, A, to the factual and of the more complex and derived representation, ∼A, to the counterfactual corresponds to the cognitive asymmetry that accords basic status to the factual and nonbasic, possibly derived, status to the counterfactual.

7.2.1 Affective States Associated with Factuality States Our first specific demonstration of a factuality interrelationship is of the linguistic representation of the counterpart affective states that are experienced with respect to a pair of factual and counterfactual complements. We first consider the case where the counterfactual circumstance is held to be more desirable than the actual circumstance. Here the affective pattern consists of two emotional states: 'regret' over what factually happened and a 'wish' for what counterfactually did not happen. These two states are understood to refer to the same single situation, as represented in (22), and to differ essentially only as to their placement of attention.

(22) ∼A more desirable than A—associated affective states:
regret over A "=" wish for ∼A

That is, as we typically understand them, each of these emotions conjures up the full comparison frame of the factual-counterfactual interrelationship, but focuses attention on only one of the alternative factuality states while evoking the other as a background comparand. In the terms used above, each of these states windows attention on one alternative of the interrelational complex.

English constructions that represent these two affective states and their attentional windowings are shown in (23a) and (23b), respectively, here exemplifying a case where the factual circumstance is an absence of activity ("I didn't go to the party").

(23) a. *Windowing A—i.e., what did take place*
I regret that I didn't go to the party. / I regret not having gone to the party.
It's too bad I didn't go to the party.
b. *Windowing ∼A—i.e., what did not take place*
I wish I had gone to the party.
If only I had gone to the party. / Would that I had gone to the party.
I should have gone to the party.

We next consider the inverse condition where the counterfactual is held to be *less* desirable than the actual. With the desirability thus reversed, the associated emotions—again ones whose character must depend on bringing both factuality alternatives into a single frame of comparison—would seem to be, on the one hand, pleasure over the actual realization of what has occurred considered against the possibility of its not having occurred, and on the other hand hypothetically contemplated displeasure over what did not occur considered against the knowledge of what has in fact occurred, as indicated in (24). In English, at least, it is evident that there are fewer constructions and lexicalizations that represent this arrangement of factors than in the case where the nonoccurrent alternative was the preferable one. Some of the most serviceable forms that do occur for this poorly represented pattern are given in (25)—here again illustrating a case where the factual circumstance is an absence of activity (''I didn't go to the lecture'').

(24) A more desirable than ∼A—associated affective states:
Pleasure at realizing A as against ∼A
''='' hypothetically contemplated displeasure with ∼A as against A

(25) a. *Windowing A—i.e., what did take place*
It's a good thing that I didn't go to the lecture.
I am (sure) glad that I didn't go to the lecture.
b. *Windowing ∼A—i.e., what did not take place*
It would have been too bad if I had gone to the lecture.
I would/could have gone to the lecture to my misfortune.

The differential favoring of the former case (the counterfactual as preferable) over the latter case (the factual as preferable) is evidenced, first, by the greater availability of open-class lexical forms that directly lexicalize the favored affectual patterns. For example, here English has the fully specific lexical forms *regret* and *wish* for the first case as against nothing but the partially serviceable *sure glad* or the too general *glad* for the second case. In addition, the favored case exhibits a greater representation by closed-class forms, which, as chapter I-1 argues, collectively represent the fundamental conceptual structuring system of language. Thus, many languages express the 'wish' notion by subjunctive-like morphemes or by unique constructions like the English *would that* and *if only* or by specific modal forms comparable to English *should*. And the 'regret' notion has at least some closed-class representation—for instance, in Yiddish by the

particle form *nebekh*. This particle can be glossed as 'poor me/you/him/...' (and is hence comparable to English *alas* except for being fully syntactically integrated within the sentence) as in *Ikh bin nebekh nisht gegangen oyf der siimkhe*, "I alas didn't go to the party." But closed-class representation for the unfavored patterns—the 'sure glad that' and 'would have been too bad if' notions—is not immediately apparent.

This difference in closed-class representation can be highlighted by noting that the favored pattern can be represented (as it was in (23b)) by a basic member of the modal system, *should*, whose meaning can be approximately characterized as 'would to one's betterment, benefit, and pleasure' (see chapter I-7). However, the unfavored pattern has no counterpart modal with the meaning 'could to one's worsening, detriment, and displeasure', which could have fit into a sentence in (25b), as if to express something like "*I would-to-my-misfortune [= Modal] have gone to the lecture."

This observation of more and less favored affective patterns suggests a program of investigation. In sequence this program would involve (1) isolating the factors that, occurring together in patterns, appear to underlie affective and cognitive states with obvious lexical or constructional representation; (2) recombining those factors so as to generate a full array of potential patterns; (3) searching various languages for lexical or constructional representation of all such generated patterns; and (4) seeking explanations for the apparent distribution of well and poorly represented patterns.

7.2.2 Explanation Types Associated with Factuality States Our second specific demonstration is in the general semantic domain of **explanations**—in which one circumstance "A′" is proposed to account for another circumstance "A"—where we observe that complementary explanation types can be associated with the two complementary factuality states. The basic equivalence of explanation types across the factual-counterfactual distinction can be formulated as in (26).

(26) A because A′
"=" ∼A-[conditional] if ∼A′

This generic formulation can be considered to encompass distinct subtypes of explanation on the basis of additional parameters, such as whether A or ∼A is held to be the preferable circumstance and whether there is an Agent either in A or in A′ who is deemed to be responsible for

or in control of the specified event by dint of his intentions and actions. However, at least in English (though other languages must be checked), the explanation constructions generally do not overtly mark any such subtypes—unlike the affect constructions, which explicitly distinguish different affective states. Accordingly, the different explanation types proposed next generally correspond to constellations of solely inferable factors. However, given the ascription of a particular explanation type to a presented construction, there will be a specific counterpart explanation type to be ascribed to the construction whose factuality is complementary to that of the first construction.

As before, we begin further analysis with the case where what has not occurred, ∼A, is held to be more desirable than what has occurred, A. Consider in addition the case where a particular cognitive agent is deemed to be responsible for A but not for A′. Here, then, an actual circumstance A′ that is outside a particular agent's control and that is offered to account for another actual but undesired circumstance for which the agent is responsible, A, can be construed to constitute an **excuse** for A. Complementarily, explicit reference to the nonoccurrent but desired circumstance ∼A can be construed as **reassurance** (or **bravado** for a first-person report) about the agent's capacity to realize ∼A in the potential case where cause A′ remains nonoccurrent as ∼A′. These relationships are symbolized and illustrated in (27).

(27) ∼A more desirable than A—associated explanation types where:
An Agent is responsible for A but is not in control of A′.
a. *Structure of the explanation types*
Excuse for A: A because A′
"=" reassurance (bravado) as to ∼A: ∼A-[conditional] if ∼A′
b. *Example with* A
I didn't catch the frisbee, A′: the car was in the way
factual—excuse
I didn't catch the frisbee because the car was in the way.
A (not-P) because A′ (P′)
Counterfactual—reassurance|bravado
I would have caught the frisbee if the car hadn't been in the way.
∼A-[conditional] (P) if ∼A′ (not-P′)

Proceeding now to the case where the occurrent A is held to be more desirable than the nonoccurrent ∼A, we further consider the case in

which a specified agent is in control of A′ but not of A (an unspecified agent is in control of A). Here, A can be understood as the compensation or **reward** that follows from the agent's execution of A′. Correspondingly, any potential nonexecution of A′ by the agent would be understood to result in the nonoccurrence of A, hence to constitute the **threat** of noncompensation—relationships symbolized and exemplified in (28).

(28) A more desirable than ∼A—associated explanation types where: A specific agent is in control of A′ but not of A (which another agent controls)
   a. *Structure of the explanation types*
      A as a reward: because A′
      "=" ∼A as a threat: if ∼A′
   b. *Example with* A
      He got a raise, A′: he worked hard
      Factual—reward
      He got a raise because he worked hard.
      A (P) because A′ (P′)
      Counterfactual—threat
      He wouldn't have gotten a raise if he hadn't worked hard.
      ∼A-[conditional] (not-P) if ∼A′ (not-P′)

The explanation types that are complementary with respect to factuality states also bear specific relations to each other with respect to force dynamics (see chapter I-7)—that is, the semantic component of language that pertains to the interactions of opposing forces such as an object's intrinsic tendency toward motion or rest, another object's opposition to this tendency, resistance to such opposition, the overcoming of resistance, and the impingement, disimpingement, or nonimpingement of blockage. Employing the terminology of chapter I-7, we can note that, for all the explanation types, the A circumstance functions as the Agonist—that is, the force-bearing entity of focal attention—while the A′ circumstance functions as the Antagonist, or the opposing force-bearing entity. We can also see that the Agonist has an intrinsic tendency toward rest—in this case, toward nonoccurrence—and that the Antagonist is the stronger of the two circumstances. In the factual explanation types, like the excuse and reward types, the Antagonist circumstance impinges on the Agonist circumstance and thus overcomes its tendency toward rest—that is, it forces it into occurrence. On the other hand, the counterfactual explanation types, such as the reassurance and threat types, depict a potential

world in which the Antagonist circumstance does not impinge on the Agonist circumstance, which is thus free to manifest its intrinsic tendency toward rest—that is, toward nonoccurrence.

Although abbreviated, the analysis in this section serves to demonstrate the existence of an integrated system that interrelates four semantic-syntactic domains that might otherwise have been thought to be independent: the windowing of attention, factuality states, affective-cognitive states, and force dynamics. It further shows that factuality and counterfactuality are complementary states within a single conceptual interrelationship and that languages afford devices for placing a window of primary attention over either of the two states.

8 MULTIPLE AND NESTED WINDOWING

Although the windowing process has so far been treated separately for each type of event frame, in fact multiple instances of windowing can occur at the same time, each with respect to several concurrent event frames. In some cases, one instance of windowing would have to be understood as nested within another, whereas in other cases, two instances of windowing would have either an indeterminate hierarchical relationship or an equipollent status. The sentences in (29) exhibit a successively greater number of instances of windowing.

(29) a. The ball rolled off the lawn back onto the court.
b. The ball rolled back onto the court.
c. The ball rolled back.
d. I rolled the ball back.
e. I kept rolling the ball back.
f. If I hadn't kept rolling the ball back, there would have been no game.

The initial sentence (29a) here exhibits a simple path event frame, complete perhaps except for a medial gapping. Sentence (29b) refers to the same path event frame but now with initial and medial gapping, hence windowing only the final portion of the path. Sentence (29c), treating the path event frame as an interrelationship event frame—in particular, as an event of motion with a Figure and a Ground—retains the Figure (the ball) within its windowing but gaps the last remaining indication of the Ground (the court). Sentence (29d) now adds an agent-initiated causal chain to the previous already-gapped motion event, thus representing a

causal-chain event frame, and, as is typical for English, windows only the agent and the final resulting subevent, while gapping specification of all the intervening causal actions. Sentence (29e) puts the previously gapped referent into an iterated cycle, thus representing a cycle event frame, but windows solely that referent as the return phase of the cycle, while gapping mention of the home, departure, and away phases of the cycle. Finally, sentence (29f) places the windowing complex to this point within a comparison frame, in particular, within a factuality event frame that windows consideration of the counterfactual while gapping consideration of the factual.

To regard one entire windowing complex that results from the concurrent or nested application of several distinct windowing processes, consider sentence (29e) as an example. Of the entire event that it refers to, this sentence windows the presence of a path but has gapped virtually the entirety of its particulars except for an indication that it is a return path (*back*); it windows the presence of a motion event and, within that, the Figure (*the ball*), but has gapped the Ground; it windows the presence of an agent-initiated event frame, but within this it windows only the agent (*I*) and the final resulting subevent (rolled the ball back) while gapping mention of all the intervening actions such as my volitionally bending down, grasping the ball, and propelling the ball into motion; and it windows the presence of an iterated cycle (*kept*), and within this the return phase, but it gaps the remainder of the cycle, including the ball's use within the court, its path from the court to the lawn, and its resting on the lawn. It is thus evident that a sentence can allude to quite an extensive referential complex while gapping an enormous amount of conceptual material from this complex.

9 SOME EVIDENCE FOR THE FUNDAMENTAL CHARACTER OF THE WINDOWING PROCESS

A range of types and alternative patterns of windowing are exhibited by the communicative signing systems that deaf children generate spontaneously and autonomously in certain circumstances. As studied by Susan Goldin-Meadow, such children have hearing parents who aim without success to communicate aurally and who employ gestural indications no more extensively or elaborately than most hearing parents use with their hearing children. To express themselves to their parents, such deaf children develop their own signing systems, ones whose structure and com-

ponents are largely not based on any external exemplars. Accordingly, one may interpret the characteristics of such systems as reflecting fundamental properties of cognition and of conceptual organization, where perhaps these properties are themselves innately determined. Thus, the fact that windowing figures prominently in such spontaneous signing systems argues for the conclusion that the cognitive processes of attentional windowing and gapping are sufficiently fundamental that they are not specific to spoken languages but appear at least through the whole cognitive domain of natural communication systems.

To illustrate, we can describe the alternatives of path-windowing patterns and of causal-chain-windowing patterns that were exhibited by a deaf child, David, observed between the ages of two years ten months and four years ten months (Goldin-Meadow 1979; Goldin-Meadow and Mylander 1990; Goldin-Meadow, personal communication). Consider first the circumstance where David would want another person to move a particular object from where it was located to a new location. One way he indicated this idea was first to point to the particular object by extending an index finger at the object and then retracting the finger a bit, and next, with the hand reoriented, to point in the same way to the new location. The initial pointing was aimed directly at the object, whether this was resting at some inanimate location or was already in the grasp of the other person. The subsequent pointing was aimed directly at the new location if this was an inanimate site, with the whole gesture perhaps adequately translated with the English verb *put* as in *Put that there*. But if a person—whether a third person or David himself—was to be the recipient or new possessor of the object, the subsequent pointing gesture was aimed at the person's chest, not hands. The whole gesture is now perhaps well translated with the English verb *give* as in *Give that to him/me*.

It is not clear whether for David the conceptualization underlying the initial pointing was of the object alone or of the object at its initial spatial location. It is further unclear whether subsequent pointing at a person's chest was conceptualized solely as marking that person as a recipient or also as a spatial location. Nevertheless, the fact that the overall gesture does indicate initial and subsequent regions of the surrounding space that approximate and are temporally iconic with the beginning and ending points of a desired motion and the fact that the gesture does not indicate any intermediate regions of the space suggests that the gesture is much like a spoken-language indication of a path with initial and final windowing and with medial gapping.

Another way David would represent a desired object transfer was to begin as before by pointing at the object but then to indicate the path through the surrounding space that the object should follow. He would trace out this path usually again with his index finger, now reextended, or, on occasion, with a new hand shape that represented how the other person might hold the particularly contoured object while moving it (e.g., a fist shape for holding a long thin object such as a spoon). He might then finish by pointing at the desired posttransfer location, or simply stop after a sufficient execution of the path-tracing gesture where the continued trajectory and terminus of this path could be inferred. Accordingly, to continue the comparisons, this gestural complex without the final pointing would seem to correspond to spoken-language forms of initial plus medial path windowing with final gapping, while the gestural complex that included the final pointing would seem to correspond to a full-path windowing.

David employed a still further type of gesture to express a desired object transfer, one exhibiting yet another path windowing pattern. For example, to indicate to the experimenter that she should go put her coat in the closet, David, without any initial point at the coat, began his gesture with a flat hand held palm downward (a hand shape used to signal carrying an object so as to place it) moving in a line toward the closet, and finished by pointing at the closet. We can now interpret this further gestural type as exhibiting medial plus final windowing with initial gapping. Thus, David demonstrated a process of selection among alternative patterns of windowing over a path event frame.

David's gestural communication also exhibited what may be interpreted as alternative patterns of causal-chain windowing. Consider, for example, the two ways in which David would represent his using drumsticks to beat his toy drum. He could clench his hands as if each were holding a drumstick and alternately swivel his hands as if swinging the drumsticks repeatedly down onto and up off of a drumhead. Alternatively, he could extend the index finger of each hand as if these were the drumsticks themselves and alternately swivel his hands as if his fingertips—the ends of the "drumsticks"—were hitting the drumhead.

It seems likely that David formed both these gestural complexes out of the one framework of a single conceptual structure, an event frame of the causal-chain type. This causal chain would have consisted of a precursor subevent [0], comprised of an intentional Agent's exercise of volition on

his body; the resulting bodily movement [1], which is the initial subevent of the physical part of the causal chain that here consists of the hands clenching and alternately swiveling; the resulting medial subevent [2], consisting of the drumsticks alternately swinging; and the resulting final subevent [3], consisting of the tips of the drumsticks alternately hitting the drumhead at the bottom of their arc.

By a narrow windowing interpretation, where only the overtly visible gesture is taken to be within the window, David's first gestural complex windows only the initial subevent [1] of the causal chain—that is, the subevent in which the hands clench and swivel. (Or, if one takes the first gestural complex to include the whole of David's person as well as his hands, it windows together the precursor subevent [0]—that is, the agent exercising volition—along with the initial subevent of the causal chain.) By the narrow interpretation, the second gestural complex would then window the medial subevent [2] of the causal chain—that is, the subevent in which the drumsticks swing.

A wider windowing interpretation would include in a window the overt gesture plus its most directly suggested concomitant. Under this interpretation, the first gestural complex windows both the initial subevent [1] of the clenching swiveling hands, which it shows overtly, plus the directly suggested medial subevent [2] of swinging drumsticks—that is, it windows the initial plus medial portion of the causal chain. Comparably under a wide interpretation, the second gestural complex windows both the medial subevent [2] of drumstick swinging, which it shows overtly, plus the directly suggested final subevent [3] of drumstick tips hitting the drumhead—that is, it windows the medial plus final portions of the causal chain.

Under either the narrow or the wide interpretation, it is strongly to be inferred that David was windowing only portions of a full causal event frame while intending to communicate the whole of the event frame, and was thus spontaneously exhibiting the cognitive windowing process in the causal domain much as in the spatial path domain before.

Such spontaneous and autonomously generated manifestations of a windowing process acting on implicit event frames, occurring in a gestural system in a way that seems fully parallel with the same phenomena earlier demonstrated for spoken language, strongly suggest that these attentional phenomena are a fundamental part of conceptual structuring in the human cognitive system for communication and perhaps also in much of human cognition in general.

10 LINGUISTIC WINDOWING AND THE COGNITIVE SYSTEM OF ATTENTION

We can now briefly consider the functions that the linguistic windowing process serves with respect to the overall organization of cognition, looking in particular at the functions served by windowing, by gapping, and by the alternatives of patterning that these can enter into.

Since the fundamental characteristic of windowing is the selective distribution of attention with respect to a conceptual complex, we must first consider more closely the nature of attention before we can determine the cognitive functions of the windowing process. Our view is that the faculty of attention is the operation of a particular cognitive system. This attentional system is able to establish active connections with aspects of other cognitive systems. The attentional system appears to have extreme flexibility as to what it is able to link up with in this way (perhaps as much flexibility as any cognitive system has), and it seems able to shift these linkups with great rapidity.

In a linkup of this sort, the attentional system lends its own processing properties to the usual functioning of the other system. These properties may be quantitative as well as qualitative and executive in character.

Thus, quantitatively, the posited attentional system may include an especially fine-grained and finely differentiated set of neural connections that allow it to function in the following ways: It enhances the processing of the other linked-up system. It differentiates factors in the other system in a more fine-structural fashion. It processes concurrently a greater number of factors present in the other system than that system itself can process. And it lowers the threshold above which certain kinds of activation in the other system can lead to further neural consequences (i.e., as a form of increased "alertness," it permits or enhances a response to weaker signals).

In addition, the attentional system may have certain special processing capabilities that allow it to function qualitatively and executively in the following ways: It selects certain factors within the other linked-up system for special processing. It compares and contrasts various factors in the other system with each other. It detects incompatibilities across such factors and brings them into an encounter for potential resolution. It brings in processing from still other cognitive systems to form a larger field of integrated processing. And, in the execution of this last function, it modulates or brings about interactions between such other cognitive systems

whose forms of processing might otherwise have little or no compatibility with each other.

It is possible that different proportions of the attentional system can be engaged in a linkup with another cognitive system in a process that gives rise to attentional gradience. The operations of the other cognitive system would thus be able to occur, over a range: more in the foreground or more in the background of attention. It is further assumed that the attentional system is able to link up at any given moment with only limited portions of other systems, so that its distinctive processing capabilities are in effect a limited cognitive resource.

We can now apply these observations to windowing in language. The establishment of a linguistic window over certain portions of a conceptual complex correlates with the linkup of the attentional system with the corresponding aspects of the cognitive system processing that conceptual complex. On the positive side, one function served by this establishment of windows of attention over certain portions of a conceptual complex is that the enhanced processing capabilities of the attentional system can thereby be associated with only those conceptual areas currently assessed as the most relevant or important relative to larger concerns and goals.

In a complementary fashion, the gapping of certain portions of a conceptual complex permits certain conceptual areas that are assessed as less relevant, more redundant, or more obvious (i.e., capable of being filled in by the hearer) to continue on unenhanced at their usual background level of processing. In addition, gapping allows the limited resource of the enhancement system to be reserved for the more important areas. These two properties of gapping thus subserve the function of the efficiency of communication of conceptual material.

The phenomenon of alternativity in linguistic windowing would clearly arise from the flexibility characteristic of the attentional system. If the attentional system were rigidly connected with the system processing a conceptual complex, one could attend only to certain portions of that complex, never to other portions. The function served by this alternativity is that approximately the same conceptual complex can be differentially adapted to different patterns of concerns that occur within different contexts.

11 CONCLUSION

The present chapter has examined a fundamental form of conceptual and attentional organization as this is evidenced primarily in language, though

its more general cognitive counterparts have also been addressed. We have seen that human cognition appears to systematically segment the occurrence of phenomena into certain types of unitary coherent conceptual packets, here termed *event frames*, where each type of event frame includes certain kinds of conceptual material but not other kinds. We posited a number of conceptual factors that help determine which phenomena are in this way packeted together into an event frame. A common cognitive principle was posited as running through these different factors: we conceptualize an event frame as demarcated by a boundary, one that encloses a region of coherence, co-relevance, and connectivity. The different types of event frame are understood to constitute generic conceptual categories that are probably universal across languages, possibly innate, and apparently in correspondence with conceptual structures present in cognitive systems outside that of language.

This chapter has treated several types of event frames: a path, a causal chain, a cycle, a participant interaction, and an interrelationship. This last type of event frame includes both the Figure-Ground interrelationship and the factual-counterfactual interrelationship, and in the latter we demonstrated a systematic relationship that affect states and explanation types bear to factuality.

Our cognition has the further capacity to select particular portions out of an event frame and to direct greatest attention to those portions while placing the remainder of the event frame in the background of attention. This cognitive process has here been termed the *windowing of attention* when it is realized in language by the inclusion of explicit linguistic material for the portions to be foregrounded (*windowed* portions) and the exclusion of any explicit material for those portions to be backgrounded (*gapped* portions). As part of a general cognitive capacity here termed *conceptual alternativity*, we are further able to perform the selective windowing process in different patterns for the same event frame. Several event frames are able to co-occur or to be embedded one within another, each with its own windowing pattern, so as to form a rather extensive referential complex with a corresponding complex of composite windowing.

For any event frame, those portions that are selected for placement in the foreground of attention may be experienced as forming a seamless continuous unity in a cognitive process here termed *cognitive splicing*. This process may well constitute one of the major psychological constancies, though one perhaps little recognized. Such a constancy could

have evolved for the selective advantage of (among other things) maintaining a single goal schema, consisting of a particular intention plus its realization, invariant across a wide range of executional variation.

Finally, we observed the strong parallels between windowing in spoken language and what seems to be a fully comparable process in the spontaneously developed signing systems of certain deaf children. Here, as well as in the parallels between linguistic windowing and perception or motor control, and in several further respects, the linguistic structures examined in this chapter can be seen as reflecting general and fundamental forms of cognitive organization.

Notes

1. This chapter is a moderately revised version of Talmy 1996a.

For their advice and assistance, my thanks go to Kean Kaufmann, Ruth Shields, Robert Van Valin, and David Wilkins.

2. Chapter I-1 outlines this framework and discusses another portion of the attentional system, the "level of synthesis."

3. This factor, the presence versus the absence of overt language material, is only one linguistic device for the setting of attentional salience. Other devices, to be treated in subsequent work, include the following: hierarchy among grammatical categories, hierarchy of grammatical relations, positioning at certain sentence locations instead of other locations, head versus nonhead constituency within a construction, degree of morphological autonomy, solo expression versus joint conflation, phonological length, and degree of stress. While most of these other devices can place attention along a gradient, windowing is taken to set attentional salience at two discrete levels: relatively foregrounded or backgrounded.

4. Some precedent for the notion of a blocked complement is present in Jackendoff's (1990) "constant argument," which can be expressed in an optional complement when specific but which, in effect, is blocked in standard speech when generic. An example of a constant argument is the argument pertaining to money in connection with the verb *buy*. Thus, one can say *I bought the book for $50*, but not **I bought the book for money*.

5. This argument is further strengthened by the fact that it has been conducted over examples containing the more liberal of the two prepositions—*on* instead of *for*—that *spend* permits with its goods-specifying complement, for instance, above with *that ham radio kit*. With *on*, not only can the expenditure of nonmonetary resources be mentioned, but the money itself could have been used either to buy the goods mentioned or to purchase other things—for example, paint, tools, insurance, expert advice—for use in the maintenance of the goods. But the use of the alternative preposition, *for*, permits reference only to money used in exchange for the goods and precludes reference to the expenditure of other resources: *I spent $50 (*and 100 hours of my time) for that ham radio kit*.

6. These alternative patterns of path windowing are presented here as if they might simply be a matter of the speaker's choice. But, of course, various discourse and pragmatic factors play a role. Though such factors are not further addressed here, it can be noted that narrative style can affect the selection. Thus, in the presentational style of the oral literature of the Clackamas Chinook (Jacobs 1958), there is a strong tendency to provide full windowing for all the open paths referred to. Examples are the following excerpts from Jacobs' literal translations.

(i) a. They left him, they went on, they came to the third mountain.
b. When it was dark, then they went, they went along, they got to there.

7. In less prototypical conceptualizations, the causality can encompass not only direct causation but also allowance or enablement, and it can occur not only at the boundary marking the end of one subevent and the start of another subevent (onset causation) but also throughout the duration of a single subevent (extended causation) (see chapter I-8).

8. Although this formulation in terms of a requirement for penultimacy may lie in the right direction, refinements and emendations are clearly needed. For example, although the sentence *?*I broke the window by throwing a rock* seems rather marginal, its close kin, *I broke a window by throwing rocks* seems relatively acceptable. In search of an explanation, we can note that, in general, a contributing factor in acceptability may be the issue of granularity or chunking—for example, the amount of the causal continuum that is conceptually framed together for consideration as a penultimate event. Thus, in the more acceptable sentence here, the window did not break as a result of my aiming some particular rock at it. Rather it broke as a chance consequence of my hurling rocks in various directions, so that the relevant chunk size of the penultimate event may be felt to extend from the act of throwing to the chance impact of one of the missiles with a window—a larger subevent that perhaps metonymically can be referred to as "throwing rocks." Further sentences pose additional challenges—for example, why it is fine to say *He killed himself by jumping out the window* instead of *He killed himself by throwing himself onto the pavement*—and it is not clear if the factor of granularity alone can resolve them.

9. Kuno (1987) has extensively investigated the conceived degree of immediacy or distance between two references to the same agent.

10. One indicator of the degree of backgrounding of the medial causal material is the fact that even linguistic analyses of agentive expressions failed to explicitly note the necessary presence of a bodily act by the agent until this was pointed out in Wierzbicka 1975 and in Talmy 1976b (here, chapter 8).

11. Actually, these sentences exhibit an additional factor beyond windowing, "direction of viewing." The window in (11a) is located in the latter portion of the home phase but includes a prospective viewing ahead to the initial point of departure, while the window in (11c) is located in the earlier portion of the departure phase but includes a retrospective viewing back to the initial point of departure.

12. This analysis shows a point neglected in previous work (e.g., Keenan and Comrie 1977), which posited the advancement or demotion of a term along a hier-

archy of grammatical relations. That work emphasized advancement as a process for increasing the prominence of a referent, but spoke little of demotion as a process for getting a referent into an oblique constituent that could then be deleted in order to background that referent.

13. Although the inclusion of the oblique Figural phrase in (20b) is awkward for that particular example, other examples exhibit all four of the construction types treated in (20) and (21), e.g., the forms in (i) and (ii) below. The Figure or Ground roles of the noun phrases are indicated symbolically here.

(i) a. The gasoline [F] slowly drained from the fuel tank [G].
    b. The fuel tank [G] slowly drained of gasoline [F].
(ii) a. The gasoline [F] slowly drained.
    b. The fuel tank [G] slowly drained.

14. Truth-value semantics and logic assume or proceed as if assuming the view that there is a direct relation between a linguistic expression and what is held to be its counterpart (its "referent") in the world. Cognitive linguistics, on the other hand, maintains that the relation between a linguistic expression and something in the world cannot be direct but must, in effect, "pass through" the mind of the language user. In particular, the relevant primary relationship is between the linguistic expression and the mind of the language user, who must first cognize the expression. Thus, a linguistic expression must first evoke a particular conceptual content in the language user's mind, being considered there by the imaginal cognitive system. This content can then be further related to other conceptual contents in the same mind, including concepts about the world.